Adverse Selection and Competitive Market Making: Empirical Evidence from a Limit Order Market

Patrik Sandås
University of Pennsylvania

This article presents a new methodology for testing economic restrictions on the price schedules offered in a limit order book that are based on (i) break-even conditions for marginal limit orders and (ii) rational updating conditions for order book revisions over time. Using order flow data from the Stockholm Stock Exchange, I find strong evidence of insufficient depth in the limit order books relative to the theoretical predictions. An extended model, which allows the model parameters to depend on market conditions, captures some of the systematic variation in the observed order book depth.

During the last 15 years many securities markets have introduced electronic limit order book systems.[1] This development has spurred a growing interest in the theoretical and empirical properties of such markets.[2] In order to account for the popularity of this trading institution, we are interested in understanding traders' order placement decisions and their contribution to liquidity and price formation. This article takes a step toward that goal by showing how the restrictions that the order books must satisfy in a relatively frictionless setting can be empirically tested. In particular, it is shown that the order books satisfy both a set of static break-even conditions and a set of dynamic updating conditions. An empirical strategy is formulated based on these conditions and implemented on order flow data from the Stockholm Stock Exchange.

In limit order markets the liquidity at any point in time is characterized by the limit order book. A limit order specifies a fixed order price and quantity and is not assured execution. When a limit order is submitted it enters

This article is based on my dissertation at Carnegie Mellon University. I am indebted to my committee members: Burton Hollifield, Robert Miller, and Duane Seppi. I thank Ulf Axelson, Bruno Biais, Richard Evans, Thierry Foucault, Rick Green, Joel Hasbrouck, Ken Kavajecz, Jonas Niemeyer, Christine Parlour, Matthew Rhodes-Kropf, Chester Spatt, Per Strömberg, and especially the editor, Larry Glosten, and an anonymous referee, as well as seminar participants at the 1999 AFA meetings, Carnegie Mellon, Columbia, LBS, MIT, NYU, UBC, UNC-Chapel Hill, and Wharton for helpful comments. I thank the Stockholm Stock Exchange, Stockholms Fondbörs Jubileumsfond, and Dextel Findata AB for providing the dataset, and the A. P. Sloan and the W. L. Mellon Foundations for generous financial support. I am responsible for any remaining errors. Address correspondence to Patrik Sandås, Finance Department, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104-6367, or e-mail: sandas@wharton.upenn.edu.

[1] Domowitz (1993) documents 35 financial markets that use limit order book systems. Recent developments are discussed in Domowitz and Steil (1999).

[2] Recent studies include Glosten (1994), Biais, Hillion, and Spatt (1995), De Jong, Nijman, and Röell (1995), Foucault (1999), Parlour (1998), Parlour and Seppi (1998), and Hollifield, Miller, and Sandås (1999).

The Review of Financial Studies Fall 2001 Vol. 14, No. 3, pp. 705–734

the centralized order book, where it is stored until it is executed or canceled. A trade occurs when a market order arrives and is executed against the outstanding limit orders. In the book, limit orders are prioritized first by price and then by submission time. In most limit order markets there are no designated market makers and therefore all liquidity is provided by limit orders.

The model used in this article is a version of Glosten (1994) with discrete prices and a time priority rule as in Seppi (1997).[3] In this single-period model, limit orders are first submitted by profit-maximizing market makers. A trader then arrives and submits a market order. It is assumed that the market orders may be submitted by a privately informed trader and that a larger market order implies a greater change in the expected value of the asset. The expected profit for a market maker submitting a limit order depends on the probability of execution and the value of the asset conditional on the order being executed. Since limit orders are executed in succession, the execution probability for a limit sell order is decreasing in its price, and at a given price it is decreasing in the submission time of the order. The resulting price schedule is then, in general, increasing or decreasing for market buy or sell orders, respectively. Each marginal bid and offer that defines this price schedule is characterized by a break-even condition.

By combining the break-even conditions with moment conditions that characterize the distribution of market orders it is possible to identify the underlying adverse selection costs. The adverse selection cost in this setting measures how much the expected value of the asset is revised in response to a given market order quantity. In other words, the adverse selection cost is represented by a "price impact function," which depends on the trade size. A second building block for the empirical analysis is a set of updating restrictions that link the market order flow to the order book dynamics. If updating is rational, the change in the value predicted by the price impact function should on average be correct. By combining these *updating restrictions* with the *break-even conditions*, I obtain a set of internally consistent restrictions on both the "shape" of the limit order book at a point in time and on the order book revisions in response to trades.

The main empirical findings of the article are as follows:

1. The price schedules offered by the order books appear to be too steep, that is, they offer insufficient depth relative to the theoretical predictions.
2. For most stocks the market order flow distribution strongly depends on state variables such as the past trading volume. I also find evidence

[3] Rock (1996) use a similar setup for the order book in a model of the interaction between the specialist and the limit order book at the NYSE. Biais, Martimort, and Rochet (2000) analyze liquidity provision by competing market makers and obtain bid (ask) prices equal to "upper (lower) tail" expectations as in Glosten (1994) in the limit as the number of market makers goes to infinity.

consistent with more (less) severe adverse selection risk in markets with high (low) stock-specific or market-wide volatility and low (high) trading volume, holding everything else constant. A restricted model that does not allow for any state dependence is strongly rejected for all stocks. However, even with the state variables, stock-specific and market-wide volatility and trading volume, a test of overidentifying restrictions rejects for 5 of the 10 sample stocks at the 1% level.

3. The deviations from the model predictions are inversely correlated with the time elapsed since the last transaction. These findings suggest that the order book may adjust with a delay to new information.

This article contributes to the literature on limit order markets. Biais, Hillion, and Spatt (1995) study the order flow dynamics at the Paris Bourse and find evidence consistent with information asymmetries and competition. This article builds on their article by directly testing theoretical predictions about the order books and the order flow. De Jong, Nijman, and Röell (1996) test a version of Glosten (1994). This article extends their analysis by providing a direct test of the model restrictions, incorporating the effects of price discreteness and the time priority rule and utilizing both order book and order flow information.[4]

The article is organized as follows. Section 1 provides a brief description of the market institution and the dataset. The theoretical model is presented in Section 2. Section 3 develops the empirical methodology. Section 4 presents the empirical results. Extensions are discussed in Section 5. Concluding comments are in Section 6.

1. Description of the Market and the Dataset

1.1 The market institution

In 1990 the Stockholm Stock Exchange introduced a computerized limit order market. This trading system is similar to those systems used on the Toronto Stock Exchange and the Paris Bourse. In this mechanism, limit orders are stored in a centralized order book and automatically crossed with market orders, giving priority according to price and time. Investors submitting limit orders have the option to "hide" part of the limit order quantity—that is, to allow only a part of the limit order quantity to be displayed in the book. The hidden part of any limit order has lower priority than all the visible order quantities at a given price level. The data used in this article do not distinguish between hidden and visible order quantities.

[4] Other empirical studies of limit order books with/without a specialist include Coppejans and Domowitz (1998), Hamao and Hasbrouck (1995), Hedvall, Niemeyer, and Rosenqvist (1997), Kavajecz (1999), and Lehmann and Modest (1994). Handa and Schwartz (1996), Harris and Hasbrouck (1996), and Lo, MacKinlay, and Zhang (1999) analyze order submission strategies.

The limit order market is transparent; information about the five best bids and offers with corresponding quantities is instantaneously transmitted to the computer screens in the offices of the exchange members. The transparency and rich information disseminated to all market participants facilitate detection of profit opportunities. The computerized trading system allows market participants to quickly shift their market-making activities from one stock to another. The above-mentioned features of the institution suggest that it is reasonable to consider a model where the order book is determined by value traders, who do not need to trade but submit orders to take advantage of profit opportunities in the market.

1.2 The dataset

The dataset contains histories for all trades and orders that were submitted to the electronic trading system for a sample of 10 actively traded stocks at the Stockholm Stock Exchange.[5] The sample period consists of 59 trading days between December 3, 1991, and March 2, 1992. The sample is representative of the electronic trading but it does not cover all trading in these stocks. Transactions made on London's SEAQ International and on NASDAQ accounted for a significant fraction of the turnover in some companies during this period. In addition, block trades can be settled outside the electronic system. While other trading venues account for a significant proportion of the volume, most transactions are routed to the electronic trading system. The analysis in this article focuses on electronic trading.

Descriptive statistics for the sample stocks are reported in Table 1. Henceforth I will refer to the sample stocks by their respective symbols listed in the first column of Table 1. The total number of transactions over the sample period, reported in the second column, ranges from a low of 667 (PROC) to a high of 8102 (LME). In the third column, similar differences in activity are observed in the total number of limit orders submitted, where the low is 1463 (PROC) and the high is 12,646 (LME). The fourth column reports the average transaction price for the stocks. During this period US $1 was equal to roughly 6.5 Swedish Crowns. Most of the stocks in the sample are traded with a tick size of one Crown (prices greater than 100 Crowns), corresponding to a tick size of 15 cents for a $15 stock. The exception is SEB, and for part of the sample LME and INVE, which traded in the 0.5 Crown tick size range (prices less than 100 Crowns). The daily trading volume for the stocks in the sample is quite variable, with a minimum of 1.1 million Crowns (PROC) and a maximum of 35.2 million Crowns (LME). The last column reports the sample period returns.

[5] The full sample included all companies in the OMX index during the sample period. OMX is a market index that includes the 30 most actively traded securities on the Stockholm Stock Exchange. The index is updated semiannually. A smaller subsample of 10 companies was randomly selected among the stocks with at least 500 trades during the 59-day sample period.

Table 1
Descriptive statistics

Stock	Symbol	Transactions	Limit orders	Average transaction price	Average daily trading volume	Sample period return
Asea	ASEA	1374	2423	309.75	2.983	12.01%
Astra	ASTR	3654	5699	522.29	28.400	15.44%
Electrolux	ELUX	2342	4271	242.19	13.338	18.53%
Investor	INVE	784	1672	124.73	2.131	22.31%
Ericsson	LME	8102	12646	109.35	35.151	9.53%
Procordia	PROC	667	1463	199.33	1.079	1.51%
Skandia	SDIA	956	2125	154.50	2.817	−15.95%
S-E Banken	SEB	2141	3719	48.36	5.598	−23.58%
Skanska	SKAB	1367	2302	123.59	2.727	−23.21%
Stora	STOR	894	1868	266.74	2.737	−3.85%

The table reports descriptive statistics for the sample. The sample period is December 3, 1991, to March 2, 1992. Transaction prices are measured in Swedish Crowns and the trading volume is measured in million Swedish Crowns. During the sample period US $1 was roughly equivalent to 6.5 Swedish Crowns. The sample period return is computed using the first and the last transaction price for each stock.

Table 2
Limit order book statistics

	Order book spreads (Crowns)			Cumulative order quantities offered in the limit order book (100 shares)			
Stock	Ask_2-Ask_1	Ask_1-Bid_1	Bid_1-Bid_2	$AskQ_2+AskQ_1$	$AskQ_1$	$BidQ_1$	$BidQ_2+BidQ_1$
ASEA	3.28	1.72	1.87	23.34	11.51	9.85	21.72
	6.26	1.04	1.99	16.00	11.61	10.17	17.00
ASTR	1.48	1.40	1.58	40.59	18.90	23.55	51.69
	1.21	0.82	1.99	32.10	19.29	28.64	44.60
ELUX	1.42	1.16	1.61	79.75	38.50	48.90	115.36
	1.25	0.49	2.85	57.64	37.87	45.83	76.93
INVE	1.55	1.22	1.52	80.74	40.18	35.35	77.57
	1.20	0.56	1.19	54.12	40.06	31.14	55.33
LME	0.93	0.91	0.91	470.71	200.25	199.91	446.85
	0.23	0.20	0.22	293.99	193.79	178.22	319.53
PROC	2.27	1.69	1.90	20.48	11.13	17.22	38.00
	2.02	1.09	2.18	17.22	13.70	18.69	28.88
SDIA	1.65	1.46	1.76	43.40	21.67	34.75	75.82
	1.17	0.98	1.98	31.43	19.61	30.20	46.23
SEB	0.59	0.53	0.57	393.98	196.08	160.18	326.06
	0.31	0.13	0.23	249.58	187.15	176.92	246.97
SKAB	1.48	1.24	1.54	61.96	29.70	44.64	88.34
	0.96	0.67	1.35	57.44	34.48	45.46	68.89
STOR	3.48	1.82	2.09	39.67	18.75	12.60	26.72
	4.40	1.28	1.79	30.69	21.04	13.90	23.24

Summary statistics for the limit order books. The mean and standard deviation (first and second row for each stock, respectively) of the difference between the second-best and the best ask price, the best ask and best bid price, and the best and second-best bid price for each stock are reported in the first three columns. The mean and standard deviation for the quantities offered in the limit order book for the two best ask and two best bid price levels are reported in the last four columns for each stock. The quantities at the second-best price levels are cumulative, that is, they include the quantities at the respective best price level (bid or ask). The spreads are measured in Swedish Crowns and the quantities in 100 shares.

In Table 2 the average order book spreads and the cumulative order quantities are reported for the two best bid and ask levels in the book.[6] Notice that for some stocks, for example, SEB, the average spreads are very close to 0.5 Crown, which is the minimum tick size for that stock. For other stocks, for example, ASEA, PROC, and STOR, spreads of several Crowns are common. The total order quantities offered at the best and second-best price levels are between 1.8 and 2.4 times the quantities offered at the best price levels. For all stocks, the spreads between the best and the second-best quotes on either the bid or the offer side are greater than the bid-ask spreads.

Table 3 reports the 10th, 30th, 50th, 70th, and 90th percentiles for the market buy and sell order quantities (numbers of shares), as well as the average market order size for each stock. Market orders submitted in a sequence within 30 seconds by the same broker-dealer on the same side of the market (e.g., a market buy order by dealer A directly followed by another market buy order from dealer A within 30 seconds) were aggregated into one market order observation. Note that for all of the stocks the median market order size is smaller than the average market order size, and for most, much smaller.

Table 2 documents that the spread is weakly increasing in the trade size for all stocks. One explanation for this is that larger orders are more likely to be submitted by informed traders. A standard method for decomposing the bid-ask spread into an adverse selection component (which may depend on the

[6] In the data it is rare to observe orders that "walk up the book" beyond the two best price levels. In other words, orders beyond the two best price levels are unlikely to be executed by the next market order.

Table 3
Market order flow statistics

		Market order quantity percentiles					
Stock	Order type	10th	30th	50th	70th	90th	Average order size
ASEA	Sell	50	100	200	400	1000	378
	Buy	50	100	200	500	1000	435
ASTR	Sell	50	100	300	750	2000	759
	Buy	100	200	500	1000	2500	1035
ELUX	Sell	100	200	500	1100	3600	1299
	Buy	100	300	800	1600	4000	1481
INVE	Sell	200	400	800	1400	3600	1388
	Buy	200	400	800	1000	3000	1208
LME	Sell	100	400	1000	2450	7000	2757
	Buy	100	300	600	1500	5000	1997
PROC	Sell	100	102	200	300	700	378
	Buy	100	200	315	600	1825	654
SDIA	Sell	100	200	500	1000	2500	1009
	Buy	200	400	900	1368	3000	1305
SEB	Sell	200	800	1400	3400	9000	3309
	Buy	200	600	1000	2200	8200	3155
SKAB	Sell	100	200	500	1000	2300	1002
	Buy	100	200	400	1000	2100	889
STOR	Sell	100	200	300	600	1500	639
	Buy	100	200	300	600	1800	717

All market order quantities reported in this table are measured in number of shares.

Table 4
Estimation of price impact regressions

	Stock									
	ASEA	ASTR	ELUX	INVE	LME	PROC	SDIA	SEB	SKAB	STOR
a	0.016	0.021	0.018	0.012	0.001	0.006	−0.011	−0.006	−0.011	−0.012
	0.55	1.38	1.33	0.39	0.20	0.13	−0.34	−1.43	−0.71	−0.31
b	0.351	0.265	0.216	0.254	0.198	0.294	0.269	0.114	0.292	0.362
	14.89	21.64	19.69	9.69	56.83	7.70	9.46	29.38	23.54	10.95
c	0.0186	0.0075	0.0039	0.0081	0.0006	0.0328	0.0096	0.0004	0.0048	0.0124
	3.88	7.51	6.33	4.75	6.86	5.33	4.93	5.30	5.00	3.70
R^2	0.20	0.17	0.21	0.19	0.32	0.18	0.16	0.33	0.36	0.17

The table reports estimation results for the following price impact regression:

$$p_{t+1} - p_t = a + b(I_{t+1} - I_t) + cq_{t+1}I_{t+1} + \epsilon_{t+1},$$

where I_t is a trade indicator, q_{t+1} is the order quantity (100 shares), p_t is the marginal price of the time t transaction, and ϵ_{t+1} is an update of the expected fundamental value that is orthogonal to the time t market order. The corresponding t-statistics are reported directly below each parameter estimate. R^2 values are reported in the last row.

trade size) and an order processing component is to estimate a "price impact regression" [see Glosten and Harris (1988)]. This methodology is designed for a different institutional setting and therefore the estimates presented below are best interpreted as benchmarks for the subsequent analysis. However, the empirical strategy presented in this article is related to this method and the key differences are briefly discussed below.

Denote the fixed order processing cost by b and the variable adverse selection cost by c. The expected change in the fundamental value conditional on the order submitted at time t, or its "price impact," is cq_tI_t, where I_t is an indicator that takes value $+1$ and -1 for market buy and sell orders, respectively. The market order quantity q_t is measured in 100 shares. A transaction price for a time t market order of $p_t = X_t + bI_t + cq_tI_t$, where X_t denotes the fundamental value at time t, covers both the order processing and the adverse selection costs. Differencing p_{t+1} and p_t yields

$$p_{t+1} - p_t = a + b(I_{t+1} - I_t) + cq_{t+1}I_{t+1} + \varepsilon_{t+1}, \tag{1}$$

where a is the expected change in value and ε_{t+1} is an update of the expected fundamental value that is orthogonal to the time t market order. The results for this price impact regression are reported in Table 4.[7] For each stock, the parameter estimates and the corresponding t-values are reported in the first six rows, followed by the R^2 values. The fixed order processing cost, b, and the adverse selection cost, c, are both positive and statistically significant at the 1% level for all of the stocks.

Figure 1 plots the mean and the median empirical price schedules as well as an implied price schedule based on the above regression results for two representative stocks, ASEA and ASTR. Each plot in Figure 1 is constructed

[7] When a market order traded at several prices, that is, a market order "walked up the book," the price for the last unit (the marginal price) was used for p_t.

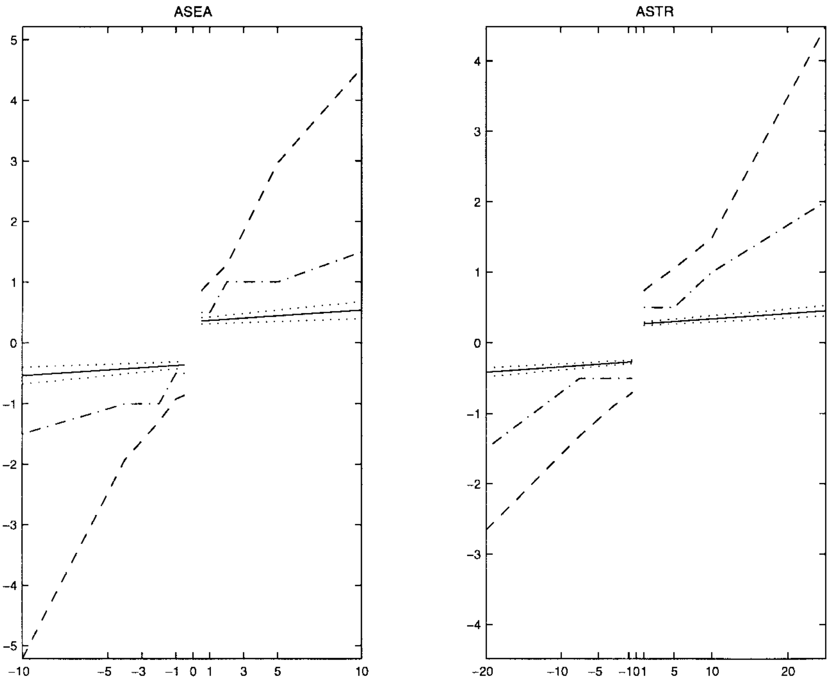

Figure 1
Comparison of Benchmark and Observed Price Schedules
The time-series mean (dashed line) and median (dash-dot line) of the markup or discount (measured in Swedish Crowns on the vertical axis) that a trader submitting hypothetical market orders of different sizes (measured in roundlots [100 shares] on the horizontal axis) would have paid or received relative to the mid-quote are plotted for two stocks, ASEA and ASTR. The markups and discounts are computed for the marginal unit of each market order size as follows. For each limit order book observed before a transaction, the markup (discount) that market orders of various sizes would pay or receive, respectively, are computed. The market order sizes are chosen to match the order flow by percentiles reported in Table 3. The solid line in each subplot corresponds to the implied price schedule based on the price impact regression results reported in Table 4, that is, $bI + cIq$ is plotted for $I = \pm 1$ and various values of q, with a 5% confidence interval (dotted line).

as follows. For each limit order book observed before a transaction, I figured the markup or discount that would be paid or received relative to the midquote by a trader submitting a market order of different sizes. The market order sizes are chosen to match the order flow by percentiles reported in Table 3. The mean (dashed line) and the median (dash-dot line) of these markups and discounts are plotted in Figure 1 as proxies for the typical price schedules observed in the data. The estimates of the cost of liquidity provision in Table 4 are then used to compute a benchmark for the price schedules. A price schedule reflecting the estimated order processing cost, b, and the "price impact," cqI, is plotted (solid line) for the different order size percentiles in Figure 1, that is, $bI + cIq$ is plotted for $I = \pm 1$ and different values for q. It is clear from the graphs that these benchmarks are "too flat" to match either the mean or the median schedule offered by the order book. The difference tends to grow as we move further out in the limit order book. These initial findings suggest that there is not enough depth in the order book or, in other words, the order books appear to offer profit opportunities. It is, however, premature to draw this conclusion at this point.

In this trading institution a trader cannot, as I did in Equation (1), condition on the size of the market order that will trigger the execution of her limit order. This implies that the execution of a limit order can result in a gain or a loss and that an order only breaks even on average [see Rock (1996) and Glosten (1994)]. Thus the transaction price is not, in general, equal to the conditional expectation of the fundamental value (plus or minus the fixed order processing cost) as in Equation (1). In addition, discrete prices combined with the time priority rule allows for nonmarginal limit orders to earn positive profits in the order book. These properties are accounted for in the empirical strategy presented in Section 3.

In the next section, I present a model of the liquidity provision in a limit order market. The model captures important institutional features of this market, including the transparency of the trading system, price discreteness, and the time priority rule. The model imposes testable restrictions on the data that allow us to examine whether or not the discrepancies between the empirical price schedules and the implied price schedules presented above imply that there are unexploited profit opportunities in the market.

2. Model

The model presented here follows Glosten (1994), but it is modified to include discrete prices and a time priority rule as in the limit order market version of Seppi (1997). The theoretical results are reviewed here to provide a framework for the empirical analysis.

2.1 General features of the model

There are two types of agents in the market: market makers and traders. Patient market participants, who supply liquidity by submitting limit orders, are referred to as market makers. Traders, who may have private information, demand liquidity by submitting market orders. Consequently they can be thought of as impatient market participants who do not want to postpone their trading. There are a large number of both types of agents. The market makers are risk neutral and profit maximizing.

The agents trade a risk-free asset with a return normalized to zero and a risky security whose value in period t is given by X_t. This is the fundamental value of the security conditional on all publicly available information. This value changes as new information arrives. The next period's fundamental value, X_{t+1}, is given by

$$X_{t+1} = X_t + \mu + d_{t+1}, \tag{2}$$

where d_{t+1} represents a random innovation in the value and μ is the expected change.

Trading occurs sequentially over periods indexed by t. In each period t, there are three stages. Shortly before time t, market makers are randomly given the opportunity to submit new limit orders. This process is repeated as long as someone wishes to place a new order. At time t, when the new limit order book has been established, a trader arrives and submits a market order. After the trade, the new fundamental value is announced and the process starts over.

The agents trade by submitting market and limit orders to a limit order book. For any order book I will use the following notation for the order prices and quantities. Let the prices $\{p_{+1}, p_{+2}, \ldots, p_{+k}\}$ denote the ask prices in the book ordered from the best ask price, p_{+1}, to the kth best ask price, p_{+k}. The order quantities associated with each price level are denoted by $\{Q_{+1}, Q_{+2}, \ldots, Q_{+k}\}$. The bid side variables are denoted analogously with a negative index. The market order quantity submitted at time t is denoted by m_t. Market buy orders correspond to positive quantities, $m > 0$, and market sell orders to negative quantities, $m < 0$.

2.2 The liquidity demand

The trader arriving in a given period is independently drawn from a population of traders. In principle, the optimal market order quantity submitted by the trader will depend on factors such as the trader's current position in the security and her information about the likely future value of the security. In order to simplify the analysis and concentrate on the decision problem of the market makers I use a reduced form representation for the trader's demand for liquidity. I assume that a trader will be a buyer or a seller with equal probability and that the desired market order quantity is exponentially distributed with expected market order quantities λ and ϕ, respectively, for buyers and sellers. The distribution of market order quantities m can then be summarized by the following density function $f(m)$, where

$$f(m) = \begin{cases} f_{mb}(m) = \frac{1}{2\lambda} e^{-\frac{m}{\lambda}} & \text{if } m > 0 \text{ (market buy)}, \\ f_{ms}(m) = \frac{1}{2\phi} e^{\frac{m}{\phi}} & \text{if } m < 0 \text{ (market sell)}. \end{cases} \tag{3}$$

I can now solve the market-makers' decision problem using this parameterization.

2.3 The market-makers' decision problem and the order book

The traders who demand liquidity may be informed about the innovation in the value of the security, d_{t+1}. The market order is therefore informative about the likely value of this innovation. An alternative explanation, which in this case will lead to the same limit order books, is that some traders observe public information regarding the value of the innovation and "pick off" limit orders that are exposed. In both cases the market order quantity submitted is correlated with the value of the security in the next period. The information link between the market order quantity m and the fundamental value X_t can be represented by a nondecreasing *price impact function*, $h(m)$, which relates the market order quantity m to the next period's fundamental value as follows:

$$E[X_{t+1}|X_t, m] = X_t + \mu + h(m), \tag{4}$$

where μ is the expected change in the fundamental value. The subsequent analysis is based on a symmetric linear specification for $h(\cdot)$ given by

$$h(m) = \alpha m, \tag{5}$$

where α is the per unit price impact of market orders. If market buy (sell) orders tend to signal increases (decreases) in the value of the asset, α is expected to be positive. For a given distribution of market orders, a larger α parameter implies a more severe adverse selection problem.[8]

Consider an arbitrary order book and an arbitrary price level p. A limit order, which is the Qth best sell order unit in the limit order book, is executed whenever a market order that is sufficiently large arrives, that is, $m_t \geq Q$. Multiple limit orders at the same price level do not "split the surplus." Instead, limit orders are executed according to strict time priority. This means that all limit order units with a higher time priority must be executed first.

Market makers are assumed to incur a quantity-invariant order processing cost, which is denoted by γ. I assume that the order processing cost is equal for buy and sell orders. In the absence of adverse selection costs (i.e., $\alpha = 0$), the order processing costs would determine the bid-ask spread, and the order book would offer infinite depth at the best quotes.

Suppose the price level p is the lowest price, above X_t, at which it is profitable to supply a positive quantity. Consider the market-makers' problem of deciding how many units to offer at this price level. The expected profit on the last unit, q, is determined by

$$E[(p_1 - \gamma - E[X_{t+1}|m])I(m \geq q)], \tag{6}$$

where $p_1 - E[X_{t+1}|m]$ is the difference between the price the market maker receives and the expected fundamental value conditional on the market order m, $I(m \geq q)$ is an indicator, which is 1 if the market order submitted is larger

[8] Many specifications of this model, including the normal exponential version, yield a nondecreasing price impact function. See Glosten (1994, p. 1137) for a counterexample.

than q, and γ is the order processing cost. We can rewrite Equation (6) as follows using Equations (3), (4), and (5):

$$\int_q^\infty (p_1-\gamma-X_t-h(u))\frac{1}{2\lambda}e^{-\frac{u}{\lambda}}du = \frac{1}{2}(p_1-\gamma-X_t-\alpha(q+\lambda))e^{-\frac{q}{\lambda}}. \quad (7)$$

Market makers are profit maximizing, and therefore the quantity offered at the price level p must be such that the last unit breaks even. This implies that the quantity Q_1 submitted at the best sell price level p_1 is given by

$$Q_1 = \frac{p_1-\gamma-X_t}{\alpha} - \lambda. \quad (8)$$

Given that the quantity Q_1 is offered at p_1 we can now ask what quantity the market makers will offer at the next price level, p_2. Using similar arguments we find that the quantity offered at the second-best price level, Q_2, is given by

$$Q_2 = \frac{p_2-\gamma-X_t}{\alpha} - Q_1 - \lambda. \quad (9)$$

By recursively following the procedure described above we can construct the whole limit order book on the offer side. By analogy we can construct the order book on the bid side.

The model presented above is static and therefore we cannot make predictions about the timing of limit order submissions. One way to think about the liquidity provision is to allow one market maker to enter and replenish the limit order book until all price levels satisfy the conditions presented above. Alternatively, several market makers could enter and compete for the profit opportunities in the limit order book. Provided that enough market makers regularly observe the status of the limit order book, profit opportunities will tend to be short-lived.

3. Empirical Methodology

The theoretical model discussed here imposes restrictions on the price and quantities in the limit order book as well as on the joint dynamics of the fundamental value and the order books. This section develops these two types of restrictions in more detail. The empirical strategy, which is subsequently outlined, builds directly upon these restrictions.

3.1 Break-even conditions

The model framework presented in Section 2 implies that the limit order book at time t is characterized by a set of break-even conditions. I concentrate on the restrictions that apply to the two best bid and ask quotes. The sharp distinction between the time of limit and market order arrivals that the model is based upon does not account for the randomness in the arrival times of orders. It is also likely that the relevant model parameters change over time, for example, the liquidity demand may depend on market conditions. In order to bridge the gap between the model and data I assume that the break-even conditions hold approximately by introducing an error term in each break-even condition. Let the error terms corresponding to the different price levels be denoted by ϵ_k, $k = \pm 1, \pm 2$. The offer side break-even conditions given in Equations (8) and (9) can then be restated as

$$p_k - X_t - \gamma - \alpha\left(\sum_{i=+1}^{k} Q_i + \lambda\right) = \epsilon_k \qquad k = +1, +2, \tag{10}$$

where $\sum_{i=+1}^{k} Q_i$ is the aggregate quantity offered at the $+k$th price level and all price levels below the kth level on the offer side.[9] Correspondingly, for the bid side we have

$$X_t - p_k - \gamma - \alpha\left(\sum_{i=-1}^{k} Q_i + \phi\right) = \epsilon_k \qquad k = -1, -2, \tag{11}$$

where $\sum_{i=-1}^{k} Q_i$ is the aggregate quantity offered at or above the $-k$th price level on the bid side. The error terms capture the fact that at any point in time the order book may offer profit opportunities or exposed limit orders. Market makers are assumed to frequently enter the market and update their outstanding limit orders in the order book and potentially add new limit orders when the order book exhibits profit opportunities. I assume that on average the limit order book does not exhibit exposed limit orders or unexploited profit opportunities, that is, $E[\epsilon_k] = 0$ for all k.

In order to eliminate the unobserved fundamental value, X_t, I difference the equations above.[10] The sum of the $+k$th and $-k$th break-even conditions, Equations (10) and (11), then result in the following two moment conditions:

$$E[e_1(y_t, \varphi)] = E\big[p_{+1} - p_{-1} - 2\gamma - \alpha(Q_{+1} + \lambda + Q_{-1} + \phi)\big] = 0 \tag{12}$$

and

$$\begin{aligned} E[e_2(y_t, \varphi)] = E\Bigg[& p_{+2} - p_{-2} - 2\gamma \\ & - \alpha\left(\sum_{i=+1}^{+2} Q_i + \lambda + \sum_{i=-1}^{-2} Q_i + \phi\right)\Bigg] = 0, \end{aligned} \tag{13}$$

where y_t denotes the time t vector of data, $\{p_k, Q_k\}_{k=-2}^{k=+2}$, and φ denotes the vector of parameters, $\{\alpha, \gamma, \lambda, \phi\}$.

[9] I could also allow the parameters to vary with time-varying instruments, but the development of the estimation strategy assumes that they are constant.

[10] In principle, I could have developed the break-even conditions based on the expected marginal profit, Equation (7), instead of the marginal profit conditional on execution, as in Equations (8) and (9). The advantage of the latter is that it makes it straightforward to eliminate the unobservable fundamental value, X_t.

3.2 Updating restrictions and order book revisions

The law of motion for the fundamental value given in Equation (2), together with the price impact function $h(m)$ given in Equation (5), provides a link between the changes in the limit order books over time and the market order flow. This link will be used to generate another set of restrictions on the data that complement the information in the break-even conditions.

The change in the fundamental value from time t to time $t+1$ is given by

$$X_{t+1} - X_t = \mu + h(m_t) + \nu_{t+1} = \mu + \alpha m_t + \nu_{t+1}, \tag{14}$$

where ν_{t+1} is an expectation update that is orthogonal to m_t and αm_t is the price impact of the time t market order. The limit order books on the offer side at time t and time $t+1$ are characterized by break-even conditions as in Equation (10):

$$p_{k,t} - X_t - \gamma - \alpha\left(\sum_{i=+1}^{k} Q_{i,t} + \lambda\right) = \epsilon_{k,t} \qquad k = +1, +2 \tag{15}$$

and

$$p_{k,t+1} - X_{t+1} - \gamma - \alpha\left(\sum_{i=1}^{k} Q_{i,t+1} + \lambda\right) = \epsilon_{k,t+1} \qquad k = +1, +2. \tag{16}$$

By subtracting Equation (15) from Equation (16) and using Equation (14) we get

$$\begin{aligned} p_{k,t+1} - p_{k,t} - \alpha\left(\sum_{i=+1}^{k} Q_{i,t+1} - \sum_{i=+1}^{k} Q_{i,t}\right) - \mu - \alpha m_t \\ = \nu_{t+1} + \epsilon_{k,t+1} - \epsilon_{k,t} \qquad k = +1, +2, \end{aligned} \tag{17}$$

and the analogous equations for the bid side are given by

$$\begin{aligned} p_{k,t+1} - p_{k,t} + \alpha\left(\sum_{i=-1}^{k} Q_{i,t+1} - \sum_{i=-1}^{k} Q_{i,t}\right) - \mu - \alpha m_t \\ = \nu_{t+1} - \epsilon_{k,t+1} + \epsilon_{k,t} \qquad k = -1, -2. \end{aligned} \tag{18}$$

I rule out a market order flow m_t that depends on the state of the order book and assume that the market order flow is uncorrelated with the profit opportunities in the order book at every price level. In other words, I assume that $E[m_t \epsilon_{k,t}] = 0$ for all k. Furthermore, I assume that the order book is replenished quickly enough after each trade so that any deviation at time $t+1$ from the break-even conditions is uncorrelated with the previous time t market order, that is, I assume that $E[m_t \epsilon_{k,\,t+1}] = 0$ for all k.[11] Given these assumptions, two additional conditions that apply to the offer side are obtained:

$$E[e_3(y_t, \varphi)] = E\left[\Delta p_{+2} - \alpha\left(\sum_{i=+1}^{+2} Q_{i,\,t+1} - \sum_{i=+1}^{+2} Q_{i,\,t}\right) - \mu - \alpha m_t\right] = 0 \quad (19)$$

and

$$E[e_4(y_t, \varphi)] = E[\Delta p_{+1} - \alpha(Q_{+1,\,t+1} - Q_{+1,\,t}) - \mu - \alpha m_t] = 0, \quad (20)$$

and two restrictions for the bid side:

$$E[e_5(y_t,\, \varphi)] = E[\Delta p_{-1} + \alpha(Q_{-1,t+1} - Q_{-1,\,t}) - \mu - \alpha m_t] = 0 \quad (21)$$

and

$$E[e_6(y_t,\, \varphi)] = E\left[\Delta p_{-2} + \alpha\left(\sum_{i=-1}^{-2} Q_{i,t+1} - \sum_{i=-1}^{-2} Q_{i,t}\right) - \mu - \alpha m_t\right] = 0, \quad (22)$$

where $\Delta p_j = p_{j,t+1} - p_{j,\,t}$.

An additional set of moment conditions is obtained by forming the difference between the expected values of the buy and sell market order quantities and the parameters λ and ϕ, respectively. These restrictions are given by

$$\begin{aligned} E[e_7(y_t, \varphi)] &= E[m_t - \lambda | m_t > 0] = 0, \\ E[e_8(y_t, \varphi)] &= E[m_t + \phi | m_t < 0] = 0. \end{aligned} \quad (23)$$

A vector of moment conditions denoted $g_t(y_{t,}\varphi)$ is obtained by stacking all the moment conditions as $g_t(y_{t,}\varphi)' = [e_1(y_t, \varphi) \cdots e_8(y_t, \varphi)]$. A GMM estimator of the parameter vector φ can then be defined as

$$\varphi_T = \arg\min_{\varphi \in \Psi} \left\{ \left[\frac{1}{T} \sum_{t=1}^{T} g_t(y_{t,}\varphi) \right]' W_T \left[\frac{1}{T} \sum_{t=1}^{T} g_t(y_{t,}\varphi) \right] \right\}, \quad (24)$$

where W_T is a positive definite weighting matrix. Hansen (1982) proves that the GMM estimator of φ is consistent and asymptotically normally distributed. The variance-covariance matrix of φ_T is given by $\Omega_\varphi^T = [D_0' S_0^{-1} D_0]^{-1}$, where $D_0 \equiv E[\frac{\partial g_t(\varphi)}{\partial \varphi}]$ and $S_0 \equiv \sum_{s=-\infty}^{s=+\infty} E[g_t g_{t-s}']$. A consistent first-round estimate is obtained by using the identity matrix as the weighting matrix. In the second round, a heteroscedasticity and autocorrelation-consistent estimate of the variance-covariance matrix S_0 is obtained by using the Newey–West procedure.

[11] I thank the editor for pointing out the importance of these assumptions.

4. Empirical Results

In this section I first report estimation results based on applying the break-even and updating conditions separately and then the results from applying them jointly. The second subsection analyzes the properties of the market order distribution in more detail. The third subsection presents the results for an extended model that allows the parameters to change with time-varying instruments. Finally, the last subsection addresses the importance of time in the dynamic adjustment of the order books.

4.1 Tests of the main model predictions

Table 5 presents the model parameter estimates obtained by applying the break-even and the updating conditions separately. The estimation based on the break-even conditions utilizes Equations (12) and (13), and the mean equations for the market order distributions, Equation (23). The estimates of λ and ϕ, which characterize the distribution of market buy and sell order quantities, are in general very close to the empirical averages of the market buy and sell order quantities (divided by 100) reported in the last column of Table 3. The estimated order processing cost, γ, is negative for all of the stocks. It is, of course, unlikely that the true order processing cost is

Table 5
Estimation using the break-even and the updating conditions separately

	Results using break-even conditions				Results using updating conditions		
Stock	λ	ϕ	γ	α	μ	α	$\chi^2(6)$
ASEA	4.362	3.776	−2.349	0.217	0.021	0.006	2.543
	36.622	40.097	−10.187	13.722	0.737	1.189	0.8636
ASTR	10.348	7.605	−1.155	0.061	0.023	0.003	1.054
	58.193	42.058	−18.332	28.803	1.557	2.293	0.9835
ELUX	14.809	12.993	−1.041	0.028	0.016	0.013	7.518
	48.756	40.262	−14.150	20.664	1.342	18.304	0.2756
INVE	12.023	13.881	−1.270	0.037	0.008	0.011	3.314
	29.340	34.194	−11.001	13.920	0.310	7.197	0.7685
LME	19.968	27.573	−0.341	0.004	0.000	0.002	40.035
	51.705	76.827	−32.391	65.923	0.052	26.844	<0.0001
PROC	6.541	3.779	−1.838	0.139	0.006	0.040	18.893
	32.136	20.885	−9.772	13.820	0.158	3.889	0.0043
SDIA	13.063	10.093	−1.430	0.054	0.001	0.021	5.122
	42.193	26.502	−12.299	18.111	0.036	8.577	0.5283
SEB	31.571	33.092	−0.406	0.003	−0.009	0.001	16.473
	27.371	40.648	−17.545	26.818	−2.164	10.571	0.0114
SKAB	8.902	10.018	−1.230	0.040	0.014	0.013	9.540
	35.702	32.643	−11.006	15.748	0.977	9.370	0.1454
STOR	7.174	6.380	−2.659	0.159	−0.028	0.037	3.063
	19.148	28.549	−8.621	11.773	−0.732	5.612	0.8009

Summary of estimation results obtained by imposing (i) the break-even conditions [Equations (12) and (13), and Equation (23)] and (ii) the updating restrictions [Equations (19)–(22)] separately. Coefficients characterizing the market order flow (λ, ϕ), the order processing cost (γ), and the slope of the price impact function (α) obtained using the break-even conditions are reported in the first four columns with corresponding t-statistics in the second row for each stock. Estimates of the drift (μ) and the slope of the price impact function (α) based on using the updating conditions are reported in the fifth and sixth column, respectively. A constant and the signed market order quantity are used as instruments in the second estimation. Given the four moment conditions we have a total of total of eight orthogonality conditions. A chi-squared statistic for a test of overidentifying restrictions (six degrees of freedom) is reported in the last column with p-values directly below.

negative. However, this could reflect rational behavior in a setting where limit orders are placed by traders with heterogeneous valuations for the asset. It is possible that a limit buy order that is suboptimal for a market maker in this model is optimal for a trader with a high private valuation for holding the asset. Such a trader will not use zero, the payoff from not submitting an order, as the benchmark, but rather will compare the expected payoff on a limit buy order with that of a market buy order.

The estimated slope coefficient, α, that determines the price impact of market orders is positive for all stocks. It is interesting to note that these coefficients are much larger than the coefficients inferred from the standard price impact regression. Note that the data used to obtain the first set of α coefficients reported in Table 5 include only "snapshots" of the limit order book at different points in time as well as the market order series. On average the slope coefficient obtained with the break-even conditions is about nine times greater than the corresponding slope coefficient obtained using the standard price impact regression. Based on these estimates the adverse selection cost for the marginal unit of an average size market order is between 10% and 59% of the quoted bid-ask spread. The large discrepancy between the two sets of slope coefficients is consistent with the large differences between the price schedules and the benchmark presented in Figure 1. Using the updating restrictions, which use time-series information, we would expect a smaller difference.

The last three columns of Table 5 reports the estimation results based on applying the updating restrictions in Equations (19)–(22). The order processing cost, γ, is not identified using these restrictions alone. Furthermore, the updating restrictions do not depend on the order flow parameters λ and ϕ. The drift parameter, μ, is close to zero for most of the stocks. The slope of the price impact function, α, is positive for all stocks. Note that these α estimates are smaller in magnitude than the corresponding estimates obtained using the break-even conditions. There is large variation between the stocks, but on average the magnitude of the α estimates obtained using the updating conditions are only about 28% of the corresponding values obtained using the break-even conditions. The magnitude of the α's for the updating conditions also exceeds the estimates reported in Table 4, but this difference is much smaller. The main difference between the standard price impact regression and the updating conditions is that the effects of price discreteness and time priority, as well as the "upper (lower) tail" expectations, are appropriately accounted for in the latter.

Using a constant and the market order flow at time t as instruments in combination with Equations (19)–(22), I obtain a total of eight orthogonality conditions. The tests of overidentifying restrictions, reported in the last column of Table 5, reject this specification for LME, PROC, and SEB at the 1% level. One possible reason for these results is that not enough time has elapsed for the order book to be replenished to the "equilibrium" levels.

Table 6
Estimation using the break-even and the updating conditions jointly

Stock	λ	ϕ	γ	α	μ	$\chi^2(3)$
ASEA	4.120	3.766	−1.882	0.1836	−0.0710	39.6340
	39.825	45.162	−12.596	17.1498	−1.7013	<0.0001
ASTR	9.944	7.057	−1.076	0.0597	0.0015	52.4118
	62.212	44.352	−18.673	30.1126	0.0744	<0.0001
ELUX	13.800	12.244	−0.852	0.0255	0.0101	70.4384
	51.499	44.399	−14.298	22.8040	0.7073	<0.0001
INVE	11.944	12.521	−1.057	0.0339	−0.0832	42.2795
	31.598	44.822	−12.772	16.1759	−3.2524	<0.0001
LME	20.131	26.836	−0.312	0.0035	−0.0033	144.1210
	55.486	78.368	−32.077	67.8169	−1.0585	<0.0001
PROC	6.185	3.498	−1.588	0.1272	−0.0012	22.8218
	35.695	27.425	−10.950	15.6923	−0.0301	<0.0001
SDIA	12.167	9.452	−1.266	0.0516	0.0564	39.0322
	47.539	27.044	−12.368	19.2479	1.6747	<0.0001
SEB	30.665	34.952	−0.318	0.0028	−0.0090	78.9999
	33.506	45.739	−19.867	30.7691	−1.7532	<0.0001
SKAB	9.203	9.699	−0.973	0.0331	0.0471	54.6851
	40.000	37.477	−12.058	18.0610	2.6456	<0.0001
STOR	5.924	6.220	−1.831	0.1327	0.0304	36.2041
	29.506	31.642	−9.555	14.8480	0.5814	<0.0001

Summary of estimation results obtained by jointly imposing the break-even and updating conditions [Equations (12) and (13) and Equations (19)–(22)] and the market order conditions in Equation (23). The first five columns report the parameter estimates: λ, the market buy order mean, ϕ, the market sell order mean, γ, the order processing cost, α, the slope of the price impact function, and μ, the drift term. T-statistics are reported in the second row for each stock. The last column reports the chi-squared statistic (three degrees of freedom) for the test of overidentifying restrictions.

This possibility is analyzed in the last subsection. The results in Table 5 indicate that the break-even and updating conditions lead us to make different inferences about the slope of the price impact function.

Table 6 reports the results for the combined restrictions. By applying the break-even conditions in Equations (12) and (13), the updating restrictions in Equations (19)–(22), and the moment conditions for the market order flow in Equation (23), I obtain eight moment conditions. No instruments are used in this estimation.[12] The system is overidentified, since only five parameters are estimated. The λ and ϕ estimates are comparable to the estimates reported in Table 5. The estimated order processing costs are negative and significant for all of the stocks. The estimated α coefficients tend to be smaller than the ones reported in Table 5. The combined restrictions yield estimates of α, the price impact coefficient, that differ from those obtained using the break-even conditions alone by roughly 10%.

The tests of overidentifying restrictions strongly reject this model specification for all stocks. The key common parameter is α and it is therefore clear that the discrepancy between the α estimates based on the break-even and updating conditions is at least partially responsible for the rejections. The order books together with the market order distribution suggest that

[12] Note that unlike the estimation based on Equations (19)–(22) alone, we cannot use the market order flow at time t as an instrument here since it is not part of the market-makers' information set at time t. I thank the editor for pointing this out.

quotes should be revised by a certain amount following trades. The actual revisions in quotes and depths are in fact much smaller. Thus, based on the observed revisions, we would expect to see more depth in the books.

The negative γ coefficients persist. As noted above, this could reflect a trade-off between market versus limit orders made by liquidity traders. The discrepancy between the α's implied by the break-even and the updating conditions may be consistent with an order *submission* cost for market makers, as in Seppi (1997), rather than an order *processing* cost. For example, it is plausible that it is costly to monitor limit orders. This cost must be paid whether or not the order is executed and may lead market makers to quote larger markups (discounts) on limit orders with lower probabilities of execution, that is, limit orders further away from the quotes. It is beyond the scope of this study to consider these alternatives. Instead I will consider the present model framework and ask whether, within that framework, there are good explanations for the results above.

To better understand how the model fails to match the data, it is useful to consider how the price schedules we observe differ from those implied by the model estimates. Figure 2 plots the average and median price schedules (dash and dash-dot lines, respectively) based on the observed limit order books as in Figure 1. In addition, two implied price schedules based on the estimation results reported in Table 5 are plotted. The implied price schedules are computed based on the α estimates and the estimates of λ and ϕ, which together determine the "upper tail" and "lower tail" expectations. The implied price schedules are computed for the market order quantities that match the percentiles reported in Table 3. All four price schedules are normalized to have a zero "price" for a trade size of one roundlot. Thus the graphs are designed to capture effects that are related to the market order quantities rather than quantity-invariant effects. The solid lines, which correspond to the α estimates obtained using break-even conditions, are closer to the actual schedules. The dotted lines represent the price schedules implied by the α estimates using the updating conditions. As in Figure 1, these schedules are close to zero. That is, prices respond less to trades than the schedules would suggest.

4.2 Market order distribution

One potential explanation for the rejections reported above is that the parametrization used for the market order distribution, two-sided exponential, provides a poor approximation of the empirical distribution. Alternatively, the conditional market order distribution could depend on the state of the market, that is, the assumption of an exogenous order flow is incorrect.

Figure 3 presents a comparison of the empirical and the implied distributions for ASEA and ASTR.[13] The empirical distribution of market order

[13] The graphs are very similar for the other eight stocks.

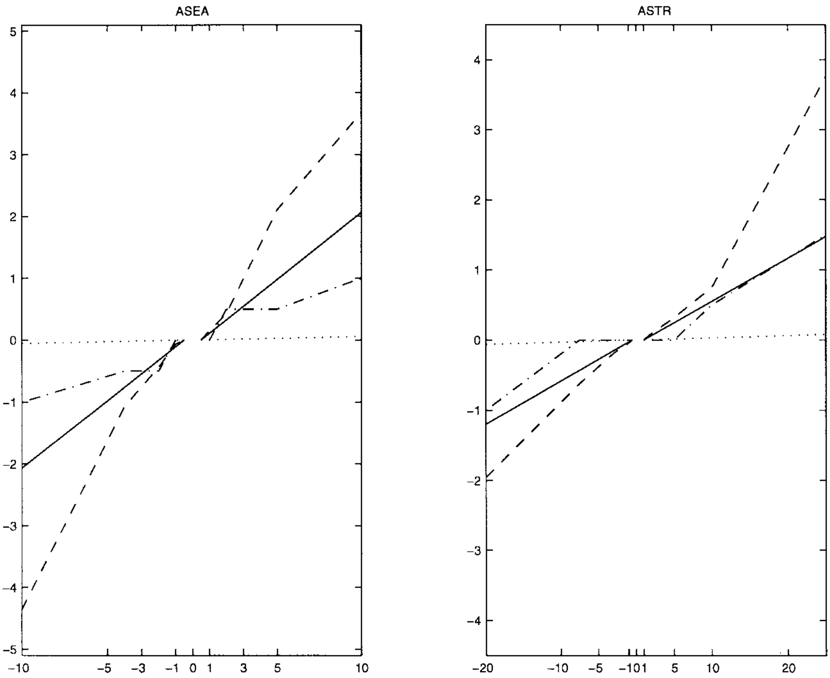

Figure 2
Comparison of Implied (Model) and Observed Price Schedules
The time-series mean (dashed line) and median (dash-dot line) of the markup or discount (measured in Swedish Crowns on the vertical axis) that a trader submitting hypothetical market orders of different sizes [measured in roundlots (100 shares) on the horizontal axis] would have paid or received relative to the mid-quote are plotted for two stocks, ASEA and ASTR. The markups and discounts are computed for the marginal unit of each market order size as follows. For each limit order book observed before a transaction, the markup (discount) that market orders of various sizes would pay or receive, respectively, are computed. The market order sizes are chosen to match the order flow by percentiles reported in Table 3. The solid line in each plot corresponds to the price schedules implied by the α estimates based on the break-even conditions (see Table 5). The dotted line represents the price schedules implied by the α estimates based on the updating conditions (see Table 5). The price schedules are normalized so that a one round-lot trade has a zero markup/discount.

quantities is plotted (bars) against the distribution implied by the exponential function and the parameter estimates for λ and ϕ (solid line) in the subplots on the left. For both stocks there seems to be significant “clustering" that the exponential distribution cannot capture. There appear to be more orders close to zero and far away from zero than we would expect if the true distribution was an exponential with a parameter equal to the average order size. This is potentially a problem as the adverse selection cost in the model is based on “upper tail” and “lower tail” expectations that may be sensitive to these types of deviations.

To further study how the assumed functional form may affect the results, I computed the “upper tail” and “lower tail” expectations implied by the estimated distribution. In Figure 3 the graphs on the right show the implied “upper (lower) tail” expectations for the market order flow, for example,

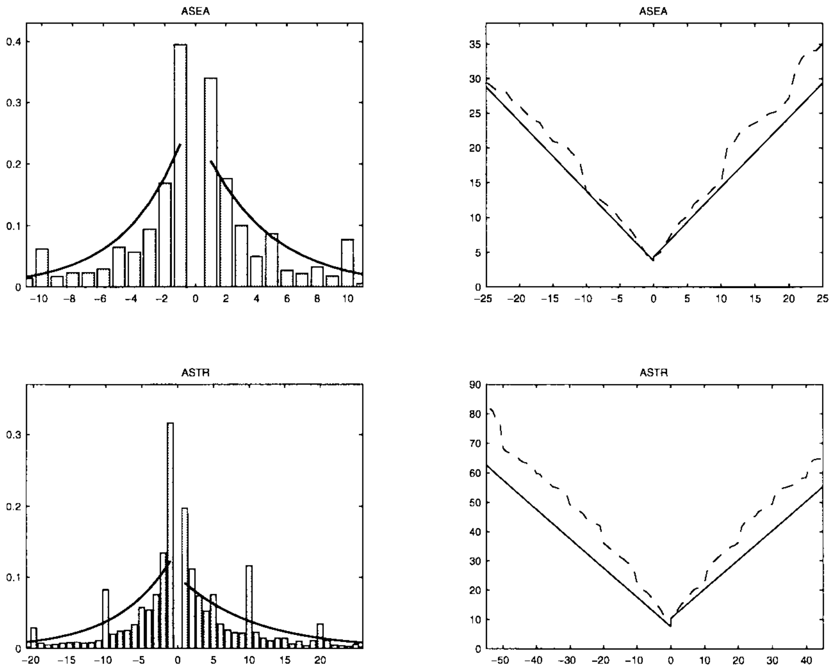

Figure 3
Distribution of Market Orders and "Upper (Lower)" Tail Expectations
The distribution of market order quantities implied by the parameter estimates (λ and ϕ) reported in Table 6 (solid line) and the empirical distribution of market order quantities (bars) are plotted on the left for ASEA and ASTR, respectively. The plots on the right show the "upper tail" expectations and the absolute value of the "lower tail" expectations computed based on (i) the estimated market order distribution (solid line) and (ii) the empirical distribution (dashed). The market sell order quantities appear with negative signs. The units for the market order quantities (x-axis) are 100 shares.

$E[m|m > 500]$, based on the estimated distribution (solid line). The same expectations based on the sample information are indicated with a dashed line. The estimates produce "upper (lower) tail" expectations that are too small. This difference grows as we move toward larger order quantities. This suggests that the functional form assumption is partly responsible for the rejections. Consider the effects of using a linear price impact function but a distribution of market order quantities, which does a better job in matching the "upper (lower) tail" expectations. This would allow us to capture more of the "excess" slope in the order books without violating the updating restrictions. Basically a steeper price schedule could be generated without causing too much overshooting in the dynamic updating. For example, update by αm after a trade of m units, but $E[x|x \geq m] = km + \lambda$, where $k > 1$, so the price schedule is given by $\alpha(km + \lambda)$ rather than $\alpha(m + \lambda)$.

Assuming an exogenous market order flow distribution significantly simplifies the derivation of the equilibrium limit order book. However, it is plausible that in reality the market order flow is a function of the limit order book.

Table 7
Order book quantities and order flow

Stock	ρ_1	p-value	ρ_2	p-value
ASEA	0.0617	0.0956	0.1408	0.0003
ASTR	0.0537	0.0330	0.0507	0.0208
ELUX	0.5485	<0.0001	0.2821	<0.0001
INVE	0.3029	<0.0001	0.5446	<0.0001
LME	−0.0082	0.5664	0.2203	<0.0001
PROC	0.7191	<0.0001	0.2536	<0.0001
SDIA	0.7098	<0.0001	0.2332	<0.0001
SEB	0.1712	<0.0001	0.3132	<0.0001
SKAB	0.3678	<0.0001	0.3402	<0.0001
STOR	0.4709	<0.0001	0.5000	<0.0001

This table reports Spearman rank correlation coefficients for the best order book quantities and the market order flow. The rank correlation coefficients (ρ_1) for the quantity at the best ask level in the order book (AQ_1) and the market buy orders (m, $m > 0$) are reported in the first column with the corresponding p-values in the second column. The corresponding correlation coefficients (ρ_2) for the best bid level in the order book (BQ_1) and the market sell orders ($|m|$, $m < 0$) with p-values are reported in columns three and four, respectively.

For example, it is likely that traders optimally choose between market and limit orders or follow dynamic order submission strategies that involve order splitting. It is beyond the scope of this article to formally consider a model with endogenous market order flow. In Table 7, I report correlation coefficients for the order book quantities (best bid and offer quantities) and the market order quantities (market sell and buy quantities, respectively). With the exception of ASEA and LME, the correlations are positive and significant. This means that larger market buy (sell) orders are observed more frequently when the best ask (bid) quantities in the order books are relatively large. Whether or not this phenomenon can explain the rejections above depends on how the market order distribution interacts with the information asymmetry.

The above findings suggest two sets of empirical facts that could explain the rejections of the main model restrictions. First, a flexible distribution for market orders should do a better job of matching the empirical distribution of market orders and thereby better approximate the "upper tail" and "lower tail" expectations. Second, a richer model that allows for an endogenous order flow would better match the empirical facts. One benefit of endogenizing the order flow might be to match some of the substantial time variation in the order books, documented in Table 2. Without changing the current model framework we can still ask whether changes in market conditions can explain some of the variation in the books.

4.3 Changing market conditions

It seems plausible that the relative intensity of informed and liquidity trading changes over time. Changes in the arrival rate of earnings and macro news could generate this type of behavior. Thus, we would expect that both the distribution of market orders and the price impact of market orders would change with market conditions. Of course, without a model it is hard to know what a relevant state variable is. With that limitation in mind the following analysis focuses on a small set of reasonable state variables.

The variables used are the volatility of the mid-quote over the last 60 minutes (z_1), the trading volume over the last 60 minutes (z_2), and the volatility of the market index over the last 60 minutes (z_3). The state variables are normalized by subtracting and dividing by the average value of the respective state variables. The first and the third state variables are selected to capture changes in the arrival rate of stock-specific and market-wide information. The trading volume variable should capture changes in trading activity and allow us to separate high-volatility periods with thin trading from high-volatility periods with heavy trading. The state variables affect all model parameters except the drift, μ, which is assumed to be constant. The interaction is assumed to be linear and given by

$$\begin{aligned}
\lambda(z) &= \lambda_0 + \lambda_1 z_1 + \lambda_2 z_2 + \lambda_3 z_3,\\
\phi(z) &= \phi_0 + \phi_1 z_1 + \phi_2 z_2 + \phi_3 z_3,\\
\alpha(z) &= \alpha_0 + \alpha_1 z_1 + \alpha_2 z_2 + \alpha_3 z_3,\\
\gamma(z) &= \gamma_0 + \gamma_1 z_1 + \gamma_2 z_2 + \gamma_3 z_3,
\end{aligned} \tag{25}$$

where λ_0, ϕ_0, α_0, and γ_0 may be interpreted as the base case for the "average" state and the other coefficients measure the effects of deviations from this state. Table 8 presents the estimation results based on using the break-even conditions [Equations (12) and (13)], the updating conditions [Equations (19)–(22)], and the market order mean equations [Equation (23)], combined with a set of instruments, which includes a constant, the three state variables (z_1, z_2, z_3), and the first differences of the state variables (Δz_1, Δz_2, Δz_3).

In Table 8 we note that the coefficients that capture the effect of trading volume on the price impact function, λ_2 and ϕ_2, are positive and significant for most of the stocks. Thus, this specification potentially captures some of the endogeneity of the order flow. The coefficients for the stock-specific volatility (λ_1 and ϕ_1) tend to be negative, six of the λ_1 and five of the ϕ_1 estimates are negative and significant at the 5% level. For the market-wide volatility the results are more mixed, with some significant positive and negative coefficients and many insignificant ones.

The estimates of the order processing cost for the "average state" γ_0 are close to the estimates reported in Table 8. The state-dependent components of γ are significant for less than half of the stocks and the estimates tend to have both positive and negative signs. Thus the results for the order processing cost do not follow any general pattern.

The estimated mean slope for the price impact function (α_0) is positive and significant at the 1% level for all stocks. The α_0 estimates reported here are close to the α estimates reported in Table 6. The α_1 and α_3 estimates

Table 8
Changing market conditions

Stock	λ_0	λ_1	λ_2	λ_3	ϕ_0	ϕ_1	ϕ_2	ϕ_3	γ_0	γ_1	γ_2	γ_3	α_0	α_1	α_2	α_3	μ
ASEA	3.840	−0.434	0.939	0.038	3.763	−0.304	1.210	−0.080	−1.813	0.214	0.140	−0.179	0.187	0.003	−0.028	0.009	−0.045
	26.097	−3.134	4.433	0.245	23.047	−2.452	4.541	−0.715	−10.714	3.174	1.213	−1.560	15.588	0.605	−3.723	1.571	−2.492
ASTR	9.123	−0.742	2.404	−0.482	6.318	−0.019	1.140	−0.120	−1.031	0.186	−0.248	−0.044	0.061	0.002	0.001	0.005	−0.003
	32.811	−1.548	5.299	−1.621	30.034	−0.035	3.062	−0.372	−17.893	1.967	−3.047	−0.836	31.466	0.821	0.512	2.610	−0.335
ELUX	13.733	−3.161	3.485	0.373	11.511	−0.310	2.177	−0.392	−0.835	0.109	0.004	−0.009	0.026	0.001	−0.003	0.003	−0.005
	26.709	−3.740	4.764	0.712	27.462	−0.444	2.976	−0.829	−13.445	1.621	0.087	−0.201	22.210	0.945	−4.250	4.568	−0.625
INVE	11.164	−0.004	0.803	−0.942	13.764	−1.412	4.341	1.316	−1.021	0.0004	0.084	−0.070	0.034	0.001	−0.007	0.005	−0.040
	22.206	−0.010	1.685	−1.906	18.238	−1.978	3.804	1.915	−11.665	0.006	1.441	−0.755	17.358	0.877	−6.182	3.365	−3.932
LME	15.254	1.154	4.659	0.433	24.036	−4.912	2.299	2.142	−0.302	0.055	0.033	−0.048	0.004	0.0005	−0.0004	0.0004	−0.006
	32.047	0.905	4.485	0.849	35.293	−2.430	1.521	2.282	−29.940	2.531	2.091	−3.569	68.215	4.790	−6.227	6.766	−5.073
PROC	6.115	−0.488	1.167	−0.508	3.672	0.072	0.678	−0.398	−1.548	0.024	−0.167	−0.310	0.128	0.004	−0.004	0.018	0.023
	15.079	−3.425	3.808	−2.771	19.611	0.696	3.248	−3.279	−10.797	0.415	−2.002	−4.532	16.038	1.925	−1.522	5.678	1.230
SDIA	12.890	−1.040	4.056	−0.254	9.968	0.033	2.242	0.837	−1.281	−0.139	−0.027	0.116	0.052	0.006	−0.005	−0.002	0.073
	23.298	−2.651	4.533	−0.647	23.121	0.090	4.240	1.865	−13.004	−2.214	−0.395	2.263	20.918	4.964	−3.877	−1.782	5.257
SEB	26.269	−7.778	9.392	2.674	28.476	−6.885	8.765	0.955	−0.339	0.063	0.051	0.002	0.003	0.001	−0.001	0.0003	−0.011
	21.959	−4.450	4.083	2.550	24.036	−4.099	5.183	0.794	−18.227	2.987	3.562	0.145	32.098	5.935	−10.838	5.104	−4.918
SKAB	8.279	−0.549	2.133	−0.310	8.990	−0.707	0.961	0.771	−1.078	−0.196	0.251	0.202	0.038	0.008	−0.008	−0.004	0.051
	19.176	−1.610	2.932	−1.003	24.302	−2.100	1.983	2.024	−11.722	−3.725	3.707	3.991	18.889	5.789	−5.577	−3.001	5.210
STOR	6.207	−0.996	3.988	0.157	5.972	0.196	0.167	−0.355	−1.827	−0.055	0.207	−0.189	0.141	0.012	−0.017	0.006	0.040
	16.979	−4.625	6.896	0.550	19.178	1.418	1.017	−1.384	−8.871	−0.684	2.861	−1.810	14.804	2.981	−6.460	1.655	1.767

Summary of estimation results using the break-even conditions [Equations (12) and (13)], the updating conditions [Equations (19)–(22)], the market order conditions [Equation (23)], and allowing the λ, ϕ, α, and γ parameters to depend on the vector of state variables (mid-quote volatility, trading volume, market index volatility), as stated in Equation (25). Estimates of λ and ϕ are reported in the first row and t-statistics in the second row for each stock.

are positive and significant at the 5% level for five and six stocks, respectively. This is consistent with a positive relationship between the adverse selection cost and the stock-specific and market-wide volatility. The eight significantly negative estimates of α_2 indicate that typically adverse selection risk is smaller when the trading volume is higher than normal.

Table 9 reports the results for different specification tests based on this extended model. The first column reports the chi-squared test of overidentifying restrictions for the extended model. The test rejects for five of the stocks at the 1% level. This suggests that the extended model performs better than the specification without state variables, but it still fails to capture important features of the data. On the other hand, many of the coefficients reported in Table 8 are significant and are consistent with economic intuition. Thus it is informative to perform some further specification tests on this model. The last four columns of Table 9 report the chi-squared statistics for tests of no state dependence for different parameters of the model. First, a joint test of no state dependence for all four parameters (λ, ϕ, α, γ) is reported. This test strongly rejects, all p-values are less than 0.0001, for all

Table 9
Specification tests

Stock	Test of over-identifying restrictions $\chi^2(39)$	Test of restrictions for $i = 1, 2, 3$: $\lambda_i = 0$, $\phi_i = 0$, $\alpha_i = 0$, $\gamma_i = 0$ $\chi^2(12)$	$\lambda_i = 0$, $\phi_i = 0$ $\chi^2(6)$	$\alpha_i = 0$ $\chi^2(3)$	$\gamma_i = 0$ $\chi^2(3)$
ASEA	50.712	70.695	34.579	15.239	11.144
	0.0991	<0.0001	<0.0001	0.0016	0.0110
ASTR	121.071	83.930	32.326	12.977	10.467
	<0.0001	<0.0001	<0.0001	0.0047	0.0150
ELUX	89.044	103.357	40.007	38.390	2.792
	<0.0001	<0.0001	<0.0001	<0.0001	0.4248
INVE	60.951	86.977	36.193	46.690	2.373
	0.0138	<0.0001	<0.0001	<0.0001	0.4987
LME	222.585	194.284	42.914	94.197	19.235
	<0.0001	<0.0001	<0.0001	<0.0001	0.0002
PROC	29.653	165.461	68.088	36.152	24.028
	0.8599	<0.0001	<0.0001	<0.0001	<0.0001
SDIA	59.211	79.032	32.873	32.677	14.444
	0.0200	<0.0001	<0.0001	<0.0001	0.0024
SEB	85.531	277.026	66.797	172.805	25.534
	<0.0001	<0.0001	<0.0001	<0.0001	<0.0001
SKAB	64.649	77.606	16.564	57.140	35.934
	0.0061	<0.0001	0.0110	<0.0001	<0.0001
STOR	45.245	135.423	57.047	48.175	10.926
	0.2275	<0.0001	<0.0001	<0.0001	0.0121

Summary of specifications tests for the extended model. The extended model allows the λ, ϕ, α, and γ parameters to depend on the vector of state variables (mid-quote volatility, trading volume, market index volatility), as stated in Equation (25). It is estimated using the break-even conditions [Equations (12) and (13)], the updating conditions [Equations (19)–(22)], and the market order conditions [Equation (23)]. The first column reports a chi-squared test of overidentifying restrictions for the extended model. The corresponding p-values are reported immediately below the chi-squared test statistic. The following four columns report specification tests of no state dependence at all, no state dependence in the market order distribution, no state dependence in the price impact function, and no state dependence in the order processing cost, respectively. The p-values for each of the tests are reported below the chi-squared statistics.

of the stocks. Second, a test of no state dependence for the market order distribution strongly rejects as well. The p-values for this test are less than 0.0001 for nine stocks and the p-value is 0.011 for one stock, SKAB. Third, a restriction of no state dependence in the price impact function is rejected for all stocks at the 1% level. Finally, a test of no state dependence in the order processing cost rejects at the 1% level for five of the stocks. To sum up, the specification tests imply that (i) accounting for changing market conditions appears to be important for all stocks, (ii) changes in market conditions affect the market order distribution and the price impact function in an economically intuitive way, (iii) the order processing cost appears to be less dependent on market conditions, and (iv) tests of overidentifying restrictions still reject for five stocks at the 1% level.

4.4 Time to replenish the book

One aspect of liquidity provision that does not enter the model explicitly is time. It takes time for the market participants to learn about profit opportunities or to find out whether their outstanding orders are exposed. Thus it is possible that some of the observations we use in the estimation represent order books that are not at their "equilibrium" levels. It is also possible that time affects the shape of the "equilibrium" order book along the lines of Easley and O'Hara (1992)—periods of no-trade signal that information-based trades are less likely to arrive and hence depth should increase or the spread decrease. It is beyond the scope of this study to formally model the possible effects of time in this setting. Instead, some indirect evidence on the possible effects of time are reported.

Table 10 reports results for two regressions that were designed to capture some of these effects. The residuals from the break-even conditions given by Equation (12) were computed for the best bid and ask quotes based on the parameter estimates reported in Table 6. In order to focus on the situation in which there are "profit" opportunities in the book, the dependent variable is the maximum of zero and this residual. Note that if the residual is positive, it suggests that according to the model there is either insufficient depth at the quotes, the bid is too low, or the ask is too high.

The first two columns of Table 10 report the results of regressing the measure of "profit opportunities" in the order book on a constant and an indicator that takes on a value of one if the time elapsed from the last transaction is less than the median time between transactions. The eight significantly positive coefficients suggest that typically we see larger deviation or "profit opportunities" when a relatively short period of time has elapsed since the last transaction. This supports the idea that it takes time to replenish the limit order book.

The second regression takes the average time elapsed between the last three trades as an independent variable. This measure captures changes in the arrival rate of market orders. For seven stocks the coefficient on this

Table 10
Time and price schedule deviations

Stock	Constant	Indicator for fast trading	Constant	Average time for last three trades
ASEA	1.0968	0.1846	1.3098	−0.0002
	19.9842	2.3778	23.5681	−3.0314
ASTR	0.6936	0.1516	0.9047	−0.0005
	29.0247	4.4863	36.6724	−7.4914
ELUX	0.5458	0.1050	0.6855	−0.0002
	23.3094	3.1721	29.0429	−5.1589
INVE	0.6297	0.1172	0.7193	0.0000
	13.3710	1.7637	16.6490	−1.1134
LME	0.2606	0.1051	0.3838	−0.0006
	45.5149	12.9993	70.3333	−18.9753
PROC	1.2381	−0.0085	1.2252	0.0000
	15.2642	−0.0743	15.7124	0.1647
SDIA	0.6415	0.2496	0.9507	−0.0002
	12.0365	3.3114	17.0287	−4.4611
SEB	0.2494	0.0791	0.3382	−0.0001
	24.0514	5.4042	34.1080	−7.3081
SKAB	0.7301	0.2555	0.9611	−0.0002
	17.7593	4.3991	22.4054	−3.2684
STOR	1.0788	0.2816	1.2513	0.0000
	14.2062	2.6222	16.4564	−0.5897

The table reports results for two regressions of a measure of price schedule errors on variables that measure the time between arrivals. The price schedule deviation is computed as the maximum of zero and the residual from applying Equation (12), that is, the difference between the actual and the predicted bid-ask spread conditional on the estimated price impact function and order distribution. A positive value indicates that there are "profit" opportunities in the order book. The first two columns report the coefficients and the t-values for a regression of the price schedule deviations on a constant and an indicator for fast trading, that is, whether time since the last trade is below the median time between trades. The last two columns report the results for a regression of the price schedule deviation on a constant and the average time elapsed between the last three trades.

variable is negative and significant at the 1% level. This suggests that longer time periods between trades are associated with books that exhibit fewer or at least smaller "profit opportunities." Thus these regression results are also consistent with the idea that it takes time to replenish the order book. They are also consistent with the idea that periods of no-trade are informative, as in Easley and O'Hara (1992). In their model longer time between trades implies a lower probability of informed trading, ceteris paribus. However, the generally negative relation between α and volume documented in Table 8 would seem to be inconsistent with this explanation.

5. Extensions

The results reported above were based on using the two best price levels in the order books. However, the data used here, like most limit order market datasets, include information about price levels deeper into the book. In principle, we can get better estimates by using these price levels when data are available. On the other hand, as we move further into the book the one-period break-even conditions that guide limit order submissions in the model may be increasingly unrealistic. Limit orders submitted deeper into the book will on average spend a longer time in the book before execution compared to the limit orders around the best quotes. Therefore it is likely that the decision to submit such orders is driven by multiperiod trade-offs between gains from obtaining time priority at prices currently away from the quotes and the costs of monitoring such orders. The current model does not address these types of trade-offs.

In this analysis I used a particular parametrization of the model, a symmetric linear price impact function and a two-sided exponential market order distribution, which greatly simplified the derivations. Neither of these assumptions is necessary in order to implement this estimation strategy. Figure 3 shows that market order quantities tend to cluster on certain numbers of round lots. A general market order distribution that allows clustering could be used in combination with a linear price impact function to match the data more closely. A more general functional form for the price impact function can be estimated by using a richer parametrization or without making any functional form assumptions. An outstanding question is whether it is possible to identify the characteristics of an underlying distribution of information and hedging needs that would lead to a liquidity demand distribution and price impact function consistent with the observed order book data.

6. Conclusions

In this article, I estimate and test a version of Glosten's model of open limit order markets. The empirical strategy is directly based on two types of restrictions that the model imposes on the data: the static break-even conditions and the dynamic updating conditions. Break-even conditions for marginal orders in the limit order book define the slope of the price schedule offered at a given point in time. Updating restrictions determine how the order book or price schedule responds to the information content in the market orders.

In the empirical implementation I find that there is a strong tension between the adverse selection costs implied by the break-even conditions and the updating restrictions. Combining these restrictions generates overidentifying restrictions that are strongly rejected in the data. Additional evidence suggests that one important reason for the rejections is that the limit order books offer too little depth or imply price schedules that are too steep relative to the order book changes in response to trades. Several explanations for this finding are explored in the article: the functional form of the market order distribution, endogenous order flow, time-varying market conditions, and time itself. Allowing for time-varying market conditions, I find evidence suggesting that the market order flow distribution and the price impact function changes with market conditions. In particular, the variance of the market order flow distribution is positively correlated with the past trading volume. The coefficients for the price impact function imply that the adverse selection risk is greater (smaller) in markets with higher (lower) stock-specific and market-wide volatility and lower (higher) trading volume, respectively.

Overall the empirical results offer some support for a specification with time-varying parameters. An interesting question for future work is to derive these types of restrictions theoretically. Such restrictions would also be useful in deriving cross-sectional tests for the order books in different stocks.

References

Biais, B., P. Hillion, and C. 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, B., D. Martimort, and J-C. Rochet, 2000, "Competing Mechanisms in a Common Value Environment," *Econometrica*, 68, 799–838.

Coppejans, M., and I. Domowitz, 1998, "Stock and Order Flow Information as Inputs to Limit Order Book Trading Activity," working paper, Duke University.

De Jong, F., T. Nijman, and A. 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.

De Jong, F., T. Nijman, and A. Röell, 1996, "Price Effects of Trading and Components of the Bid-Ask Spread on the Paris Bourse," *Journal of Empirical Finance*, 3, 1–35.

Domowitz, I., 1993, "A Taxonomy of Automated Trade Execution Systems," *Journal of International Money and Finance*, 12, 607–631.

Domowitz, I., and B. Steil, 1999, "Automation, Trading Costs, and the Structure of the Securities Trading Industry," in R. Litan and A. Santomero (eds.), *Brookings-Wharton Papers on Financial Services*, Brookings Institution Press, Washington, D.C.

Easley, D., and M. O'Hara, 1992, "Time and the Process of Security Price Adjustment," *Journal of Finance*, 47, 577–605.

Foucault, T., 1999, "Price Formation in a Dynamic Limit Order Market," *Journal of Financial Markets*, 2, 99–134.

Glosten, L., 1994, "Is the Electronic Limit Order Book Inevitable?" *Journal of Finance*, 49, 1127–1161.

Glosten, L., and L. Harris, 1988, "Estimating the Components of the Bid-ask Spread," *Journal of Financial Economics*, 21, 123–142.

Hamao, Y., and J. 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, V., and R. Schwartz, 1996, "Limit Order Trading," *Journal of Finance*, 51, 1835–1861.

Hansen, L. P., 1982, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1024–1054.

Harris, L., and J. Hasbrouck, 1996, "Market vs. Limit Orders: The SuperDOT Evidence on Order Submission Strategies," *Journal of Financial and Quantitative Analysis*, 31, 213–231.

Hedvall, K., J. Niemeyer, and G. Rosenqvist, 1997, "Do Buyers and Sellers Behave Similarly in a Limit Order Book? A High Frequency Data Examination of the Finnish Stock Exchange," *Journal of Empirical Finance*, 4, 279–293.

Hollifield, B., R. A. Miller, and P. Sandås, 1999, "Empirical Analysis of Limit Order Markets," Working Paper 029-99, Rodney L. White Center, Wharton School, University of Pennsylvania.

Kavajecz, K., 1999, "A Specialist's Quoted Depth and the Limit Order Book," *Journal of Finance*, 54, 747–771.

Lehmann, B., and D. Modest, 1994, "Trading and Liquidity on the Tokyo Stock Exchange: A Bird's Eye View," *Journal of Finance*, 49, 951–984.

Lo, A.W., A. C. MacKinlay, and J. Zhang, 1999, "Econometric Models of Limit-Order Executions," Working Paper 012-99, Rodney L. White Center, Wharton School, University of Pennsylvania; forthcoming in *Journal of Financial Economics*.

Parlour, C., 1998, "Price Dynamics in Limit Order Markets," *Review of Financial Studies*, 11, 789–816.

Parlour, C., and D. Seppi, 1998, "Liquidity-Based Competition for Order Flow," working paper, Carnegie Mellon University.

Rock, K., 1996, "The Specialist's Order Book and Price Anomalies," working paper, Harvard University; forthcoming in *Review of Financial Studies*.

Seppi, D., 1997, "Liquidity Provision with Limit Orders and a Strategic Specialist," *Review of Financial Studies*, 10, 103–150.