

Dynamics of Liquidity in an Electronic Limit Order Book Market

Mark Coppejans*
Ian Domowitz**
and
Ananth Madhavan†

Current Version: March 6, 2002

* Department of Economics, Social Science Building, Duke University, Durham, NC 27708, Tel: (919)-660-1804; E-mail: mtc@econ.duke.edu

** ITG Inc., 380 Madison Avenue, New York, NY 10017, Tel: (212)-444-6279; E-mail: Domowitz@itginc.com

† ITG Inc., 380 Madison Avenue, New York, NY 10017, Tel: (212)-444-6361; E-mail: Amadhava@itginc.com

We are grateful to Lester Loops for making the data available and providing useful insights with respect to the market structure examined here. We thank Pete Kyle, Frank Hatheway, Patrick Sandås, Tim Simin, and seminar participants at Yale University, University of Maryland, Pennsylvania State University, and the Atlanta Federal Reserve Bank's Financial Markets Conference for helpful comments. Any errors are entirely our own.

Dynamics of Liquidity in an Electronic Limit Order Book Market

Abstract

We take advantage of special features of a limit order book to construct measures of liquidity. It is first shown that liquidity varies substantially over the trading day. This is of interest, in part, because we provide strong evidence that discretionary traders trade in high liquidity periods. We also model the comovements of liquidity, returns, and volatility of returns. In doing so, several results emerge. Our findings are consistent with models where liquidity is a factor in expected returns, but also suggest more complicated dynamics consonant with supply and demand imbalances in the market. Furthermore, while increases in liquidity substantially reduce volatility, volatility shocks reduce liquidity over the short-run, impairing price efficiency. These effects dissipate quickly, however, and their magnitudes are small, indicating a high degree of market resiliency.

JEL Classification: G10, G34

Keywords: Electronic trading, liquidity, execution costs, return dynamics

1. Introduction

The automated auction has transformed securities markets. Advantages of speed, simplicity, and low costs drive the rapid adoption of automated auctions to trade equities, bonds, foreign exchange, and derivatives worldwide.¹ Unlike traditional markets, trading in an automated auction is through an electronic limit order book without the need for a physical exchange floor or intermediaries such as market makers. But in the absence of intermediaries, an automated auction is dependent on public limit orders for liquidity. If the limit order book is thin, even small trades can induce large price movements, increasing trading costs and volatility. Such time-variation in liquidity can exhibit complex dynamics that are not well understood. For example, discretionary traders may pool together, providing a virtuous circle where liquidity begets liquidity. The converse is also possible. In a vicious circle, market illiquidity and high volatility discourage the use of limit orders, reinforcing a low-liquidity equilibrium.² Intertemporal variation in liquidity is also important given evidence that liquidity affects expected returns.

This paper examines liquidity in an automated auction and the resulting dynamics of overall market liquidity, volatility, and returns. The analysis in this paper follows three steps. We first document that liquidity varies substantially across the trading day. To show that this result is not just of academic interest, we next provide strong evidence that changes in liquidity affects traders' behavior. Given this, the next logical step is to determine, in part, why liquidity is variable. By examining the comovements of liquidity, returns, and volatility, we shed some light on this latter issue.

We use intraday order-level data obtained from the electronic market for Swedish stock index futures (henceforth OMX). Our dataset is one of the few complete limit order books in ex-

¹ Outside the US and a handful of emerging markets, virtually all equity and derivative trading systems are automated. A partial list of major automated markets includes, for equities, the Toronto Stock Exchange, Euronext (Paris, Amsterdam, Brussels), Borsa Italiana, National Stock Exchange (India), London Stock Exchange, Tradepoint, SEATS (Australian Stock Exchange), Copenhagen Stock Exchange, Deutsche Borse, and Electronic Communication Networks such as Island. Fixed income examples include eSpeed, Euro MTS, BondLink, and BondNet. Foreign exchange examples are Reuters 2002 and EBS. Derivative examples include Eurex, Globex, Matif, and LIFFE. Domowitz (1993) provides a taxonomy of automated systems and updates are contained in Domowitz and Steil (1999).

² Such issues relating to viability and stability are reviewed from a theoretical perspective in O'Hara (1995, ch. 7).

istence and is ideally suited for our study in several respects.³ In particular, we observe the instantaneous demand and supply curves at every point in time. These yield natural metrics for liquidity in terms of market depth, i.e., order flow necessary to move price by a given amount. The cross-border automated limit order book system studied here is typical of many markets, including the Toronto Stock Exchange and Paris Bourse, allowing for some confidence that our results are not artifacts of special institutional arrangements. An unusual, but valuable, feature of our database is that it identifies orders arising from the so-called “upstairs” market where large-block trades are negotiated and crossed. Failure to distinguish these trades from regular trades biases any assessment of the real costs of trading and true underlying liquidity of the market.⁴ Finally, the index futures contracts traded represent claims to the entire equity market, so that our analysis is one of aggregate liquidity.

Beyond any interest in the operation of a pure limit order market itself, limit order book data permit the investigation of broad questions concerning the microstructure of markets. For example, measures of liquidity are constructed here, based on the instantaneous demand and supply curves inherent in the limit order book. We show that the variation in liquidity over time is economically and statistically significant. This is consistent with strategic models, in which discretionary traders trade in high liquidity periods, in turn reinforcing the concentration of volume and liquidity at certain points in time.

The results suggest that traders can add value by strategic order placement behavior. We present evidence in favor of this hypothesis. In particular, the actual execution costs incurred by traders are significantly lower than the costs that would be incurred under a naïve strategy that ignores time-variation in liquidity. The cost differences are especially pronounced for larger trades, even after excluding off-exchange crossed trades. One implication of this result is that institutional traders who simply partition their orders mechanically over the day with the objective of trading at the “value weighted average price” could benefit from attempting to time their

³ The Paris Bourse data, for forty stocks, is described and analyzed by Biais, Hillion, and Spatt (1995), and Gouriéroux, Le Fol, and Meryer (2000) provide a factor analysis of the order queue for a single stock. Hollifield, Miller, and Sandås (1999) and Sandås (1999) use OM data, but for a selection of 10 stocks traded on the equities order book. Some data also are available for trading on the Australian SEATS automated system, Toronto Stock Exchange, and Tel Aviv Stock Exchange.

⁴ For example, if the market is very illiquid, institutions may send large orders upstairs market so that more crosses are observed. Without explicit identifiers, such mid-quote executions would falsely suggest a high level of liquidity.

trades to take advantage of periodic liquidity surpluses while avoiding liquidity deficits. The nonlinear nature of the demand and supply schedules, together with systematic intraday variation in liquidity, generally implies that the optimal dynamic trading strategy is not uniform.

Finally, we analyze the dynamic relation between measures of liquidity, volatility, and short-horizon expected returns using structural vector autoregressive models, where the change in the midquote and the absolute value of the change in midquote are used as proxies for returns and volatility, respectively, throughout this paper. A growing literature suggests that there is a relation between liquidity and expected returns. Amihud and Mendelson (1986, 1991) find cross-sectional evidence of a positive relation between asset returns and bid-ask spreads, for example. Amihud, Mendelson, and Lauterbach (1997) document large changes in asset values for stocks moving to more liquid trading systems on the Tel Aviv Stock Exchange. Brennan and Subrahmanyam (1996) and Brennan, Chordia, and Subrahmanyam (1999) show that liquidity can explain the cross-sectional variation in returns while Hasbrouck and Seppi (2000) examine commonality in liquidity. All such findings are cross-sectional in nature. Our time-series analysis is complementary, suggests some complicated and unexplored dynamics, and provides support for microstructure models such as that of Spiegel and Subrahmanyam (1995), where liquidity is a factor in expected returns. In particular, volatility shocks reduce liquidity, a fact that supports arguments for trading halts following sharp market movements. Shocks to liquidity dissipate quickly, indicating a high degree of resiliency. This self-correcting ability turns out to be an attractive feature of the automated auction, mitigating doubts with respect to the resilience of that form of market structure under pressure.

The paper proceeds as follows: institutions and data underlying the analysis are laid out in section 2; section 3 contains the formal definition of liquidity and results showing that liquidity varies across the trading day; section 4 provides evidence of discretionary trading; section 5 presents the autoregressive model for joint analysis of liquidity and returns; the dynamic relation between liquidity and volatility is examined in section 6, and some concluding remarks are offered in section 7.

2. Institutions and Data

2.1. Market Architecture

Trading in Swedish stock index futures contracts takes place via a consolidated automated trade execution system, including activity from Sweden, the U.K., Denmark, and the Netherlands. We refer to the overall market as OMX, given the complete integration of trading across countries.⁵

The electronic system functions as a continuous pure limit order book market. Trading on the order book is in round lots of 10 contracts. Orders are prioritized on the book in terms of price, then time. There are two ways in which a trade may be executed. Counterparty limit orders may match on the book in terms of price, in which case the maximum feasible size is filled.⁶ Alternatively, a trader may “hit the bid” or “lift the offer,” taking up to as much quantity as advertised on the book. This is accomplished by executing a single keystroke and submitting desired volume. Once a trade is completed, unexecuted volume at the trade price remains on the order book, until cancelled. Cancellations of orders are possible at any time.

The trading day is six hours, beginning at 9:00 AM and ending at 3:00 PM, GMT. Unlike many automated markets, such as the Paris Bourse, there is no opening algorithm or batch auction at the beginning of the day. With that exception, the design and mechanics of the OMX market are quite similar to that described by Biais, Hillion, and Spatt (1995) for the CAC system, and by Domowitz (1993) for generic price/time priority continuous limit order systems.

There are some additional features that are relevant to the analysis to follow. Block transactions are allowed, in the form of “crosses.” Crosses are arranged “upstairs” or off-exchange, and the two sides are not listed on the order book. Nevertheless, crosses, described in terms of price and quantity, are displayed in the continuous time transaction record observed by traders. Unlike the practice in some other markets, there is no interference with a cross from activity on the limit order market.⁷ A small amount of odd-lot trading also takes place. A separate

⁵ Clearing is conducted on a local basis. The Swedish contract originated on OM Stockholm in 1985, and OMLX, the London Securities and Derivatives Exchange, was established in 1989, with the additional links following thereafter.

⁶ “Locked markets” do not result if an entered bid price is higher than an offer price on the book. A transaction occurs based on time priority, at the offer price in this example.

⁷ The Swiss SOFFEX derivatives system, for example, exposes arranged trades to the limit order book, similar to the practice on the NYSE for upstairs blocks.

facility exists for this activity, but such trading is integrated with the main book. For example, an odd lot of 3 contracts and one of 7 contracts automatically matches with a round lot of 10 contracts on the main book.

Order and trade information are distributed directly from the trading system, making the OMX highly transparent.⁸ Specifically, market participants observe a transactions record (price and volume) and the five best bids and offers on the book, with aggregate volume at each price.⁹ No “indicative” prices or other non-price expressions of trading interest are provided. A trader may view information through OM’s interface or accept a real-time feed, which allows for customized screens and data processing. Although this seems to be a small detail, it proves relevant in the analysis of trading cost management to follow.

2.2. Data

Our database comprises the complete limit order book for Swedish stock index futures contracts from the period 7/31/95 through 2/23/96. The data are obtained from a trading house that chose the real-time feed, permitting the collection of some historical information for analysis.¹⁰ Prices are denominated in Swedish currency (SEK), and volume is given in number of contracts. Information is time-stamped to the second. Transactions files and order information are matched. The order book is reconstructed from the raw data and completely consistent with transactions reported.¹¹ Odd-lot trades are identified, but constitute only about three percent of all trades, and average less than five contracts per trade. Crosses are isolated, and matched in time with limit order book trading activity.

Activity for near-term contracts is analyzed in what follows, since there is little liquidity in contracts for which expiration is further away. Some data is eliminated at the end of expira-

⁸ Transparency refers to the quantity and quality of information provided to market participants during the trading process. Limit order markets are typically highly transparent because they provide relevant information before (quotes, depths) and after (actual prices, volumes) trade occurs. By contrast, foreign exchange and corporate junk bond markets rely heavily on dealers to provide continuity but offer very little transparency while other dealer markets, such as Nasdaq, offer moderate degrees of transparency.

⁹ There is some facility for so-called “hidden orders” that are unobserved by traders. As in the analyses of Biais, Hillion, and Spatt (1995) and Hollifield, Miller, and Sandås (1999), we cannot ascertain the effects of such unobservable orders, but their importance in automated systems is generally very limited as discussed by Irvine, Benston, and Kandel (2000).

¹⁰ We thank Lester Loops, who provided the raw numbers and some assistance with issues involved in merging the order and transactions records.

¹¹ Irregularities, initially constituting about one percent of trading activity, were uncovered, but all are reconciled with the assistance of the trading house that provided the original data.

tion cycles, mitigating liquidity effects stemming from lack of trading due to rollover effects. The daily average number of orders, cancels, and transactions in the data analyzed below are 1941, 1334, and 177, respectively.

3. Liquidity

In what follows, we define market liquidity or depth as the number of contracts offered for sale at up to k ticks from the midquote. We distinguish between liquidity on the buy and sell sides, denoted by $D_b(k)$ and $D_a(k)$, respectively. These measures are natural in that they can be interpreted as the volume necessary to move the price by k ticks. More liquid markets are deeper in that they can accommodate larger trades for a given price impact.

Table 1 contains summary statistics relating to the depth of the order book, in number of contracts, by time of day, averaged over 105 trading days. Data for the bid side appears in Panel A, and data for the offer side appears in Panel B. Column headings indicate the number of ticks away from the midpoint of the best quote in the market at the time. The figures reported are the number of contracts available at or below that number of ticks away from the midquote. Essentially, these figures constitute the instantaneous supply and demand curves, averaged across days. For example, from Panel A, at 10:15 AM, there are (on average) 58 contracts bid at up to 8 ticks below the midquote. Market depth at any distance from the midquote is lowest at the opening session mainly because there is no opening algorithm or batch auction at the beginning of the day.

The instantaneous demand and supply schedules are of considerable interest in themselves, because their shape provides some clue with respect to strategic trading activity. Linear schedules suggest that large orders are broken up into equal size blocks for submission over the trading day in a uniform manner. Nonlinearity suggests departures from such a uniform strategy. We investigate the potential nonlinearity of schedules by estimating polynomial approximations to the bid and offer curves.¹² The regressions relate average depth to the number of ticks away from the midquote. The approximations are graphed in Figure 1 for bid and offer schedules. A linear approximation is also illustrated. The bid and offer functions are roughly S-shaped, with

¹² A fifth-order polynomial is used for the results reported.

slight convexity at prices close to the spread midpoint, and considerable departure from linearity starting at about eight ticks (approximately 0.16 percent of value) away from the midquote.

Comparing Panels A and B, the bid and offer sides of the book are roughly symmetric in terms of depth and execution probabilities. There appears to be little difference between the demand and supply schedules, on average, which also is evident from Figure 1. This suggests that trading behavior and patterns arising from order imbalances are likely to be short-lived, a topic investigated further in section 5.

Casual inspection of depth by time of day suggests little time variation in liquidity, except for the open. This is incorrect. First, the standard deviation (not provided) of depth is large, even relative to the mean, in all cases reported in table 1. In other words, there is considerable variation in observed depth at different times of day even though on average they are roughly equivalent. Second, first-order autoregressive models of depth suggest a moderate degree of mean reversion in liquidity, and a large residual variance relative to mean depth.¹³ Such results also suggest substantial time variation, but not necessarily that which would be captured by simple time-of-day analysis. In fact, models such as that of Admati and Pfleiderer (1988) do not predict time-of-day effects, although they are often associated with empirical phenomena at the open or close. Rather, they predict that patterns in liquidity and trading occur over time, with no statement as to the clock, as pointed out by O'Hara (1995, p. 139).

4. Discretionary Trading

Discretionary timing of trades involves several underlying hypotheses and predictions. In Admati and Pfleiderer (1988), it is optimal for discretionary uninformed traders to trade at the same time, for example. This in turn implies liquidity clustering, in an environment in which informed trading further exaggerates the clustering effect. In Scharfstein and Stein (1990), large order flows, observable here through the book, encourage entry by traders, suggesting that greater liquidity should be correlated with more and larger trades. A similar herding effect in the case of discretionary timing is predicted by Spiegel and Subrahmanyam (1995). An even sharper result is obtained by Mendelson and Tunca (2000). In their model, discretionary liquidity traders

adjust order sizes along with changing market depth, equalizing trading costs across size of transaction.

4.1. Price Impact Functions

We begin by summarizing in a simple manner the expected trading costs facing a trader at any point in time based on the prevailing demand and supply schedules. In particular, consider a market order of size Q (with the sign convention that $Q > 0$ represents a purchase and $Q < 0$ a sale) that, given the extant book, is executed at k different prices, with q_k shares executing at a price p_k , where $\sum q_k = Q$. The price impact of the trade is then defined in terms of the appropriately signed percentage difference between the weighted-average execution price and the pre-trade midpoint:

$$p(Q) = \ln\left(\frac{\sum p_k q_k}{Q p_0}\right) \text{sign}(Q), \quad (1)$$

where p_0 is the midpoint of the bid-ask spread at the time of the trade. The price impacts thus defined are inversely related to the depth measures defined above. So, for example, if $D_b(k) = Q$,

the total price movement associated with a buy order of size Q is k .¹⁴

Table 2 contains the expected price impact of trades, reported in percentage terms relative to the quote midpoint, by time of day. Calculations are done for hypothetical trades of 10 to 100 contracts in increments of 10, compared with the observed order book at a specific time of day, averaged over 105 trading days. Figures in the row marked “average” are computed based on computations at 15-minute intervals over the trading day, averaged over intervals and trading days. Panel A contains data for transactions at the bid, and Panel B contains figures for transactions at the offer. As one might expect, the price impact of the trade is strictly increasing in order size, ranging from 7 to 15 basis points overall. Consistent with table 2, the price impacts are much higher at the open, but do not vary by whether the order is a market buy or a market sell.

¹³ First order serial correlation coefficients for depth at 6 ticks away from the midquote, for example, are 0.66 and 0.60 on the bid and offer side, respectively, with estimated residual standard deviations of between 22 and 23, relative to means of between 24 and 50 in Table 1, and constant terms of 13 to 15 in the autoregressions.

¹⁴ The actual percentage price impact depends on the distribution of limit orders on the price grid. Suppose $p_0 = 100$, $Q=50$, and at $p_1 = 110$ there are 20 contracts, and at $p_2=120$ there are 40 contracts. Then $p(50)=\ln[(110 \times 20 + 120 \times 30) / 100 \times 50] = 0.148$. If there were 50 contracts at p_1 , then $p(50) = 0.095$.

In equity market studies, it is increasingly common to model the price impact of a trade as a strictly concave function of size. Hasbrouck (1991), for example, advocates the use of square-root transformations for order size. Similar results are obtained by Madhavan and Smidt (1991), among others. By contrast, the price impacts here are nearly linear functions of size.

The difference between our results and those based on NYSE or Nasdaq data might be the result of market structure. On the NYSE, for example, the trading crowd and specialist may step in to provide liquidity for large orders, while Nasdaq dealers may offer volume discounts to their customers. On an automated auction like the OMX, however, traders are unwilling to offer large quantities at prices far away from the current price. Such limit orders constitute free options to the market, options that will be taken if the market moves by a large amount. The absence of depth at far prices implies that the price impact function is convex, because large trades incur proportionately greater costs.

It is also possible that the difference in the shape of the price impact function reflects upstairs trades. The data used to test models of the U.S. equity markets do not identify large-block trades executed upstairs. These trades typically occur within the bid-ask spread, possibly biasing the estimated costs of execution for large orders downward. This is not an issue for us, since the computations in table 2 use the current limit order book.

4.2. Realized Price Impact Costs

We view the costs in table 2 as a benchmark, being produced from a completely naïve trading strategy. Table 3 contains the actual price impact of trades, reported in percentage terms relative to the quote midpoint, by time of day. We use equation (1) to compute these impacts except that we use the *realized* executions from an incoming market order in computing the trade price. Calculations are done for actual trades of 10 to 100 contracts in increments of 10, compared with the observed order book at the time of trade, over 105 trading days.

In contrast to table 2, the realized impacts in table 3 are virtually constant across order sizes. This pattern is true for both trades on the bid and offer sides. It also is true for off-exchange crosses. The results are strongly supportive of the Mendelson-Tunca (2000) model of discretionary trading equilibrium. It also is evident that traders obtain substantially lower costs than they would through a naïve order submission strategy, especially for large orders, even ignoring crosses. For example, the hypothetical price impact of a trade on the bid side for 100

contracts is 275 percent larger than realized price impact costs. These findings support the predictions of the Admati-Pfleiderer (1988) model.

Constancy of price impact across size has an immediate practical implication. Many institutional managers use the value-weighted average price (VWAP) as their benchmark price in evaluating trade performance. Consistent with this, some traders attempt to realize VWAP by using a simple break-up strategy, submitting equal-size pieces over the trading day. Our findings suggest that this strategy is suboptimal; efforts to take advantage of time-varying liquidity may result in substantially better executions. These results are precisely what were expected given the evidence on nonlinearity of the demand and supply schedules illustrated in Figure 1.

Interestingly, many crosses do not go down at the midpoint. The crosses are often at the bid or offer, as is obvious from the nonzero price impacts reported in table 3. Crossing away from the midquote does not save much money relative to doing the trade directly with the book, except for large-block trades of 90 contracts or more. Crosses are often done in the morning, and the relatively thin book at the opening provides an obvious rationale for off-exchange dealing. On the other hand, an even greater number of crosses are executed towards the close, with a very thick book, perhaps because traders are concerned that they might not be able to execute a large block trade with little time remaining in the trading day. This is consistent with the evidence on the proportion of block transactions in the US equity market, which also diminishes sharply at the end of the day.

4.3. Strategic Order Placement Behavior

The difference between hypothetical and actual price impacts confirms the existence of discretionary timing, and is consistent with strategic behavior on the part of traders. All theories relating to discretionary trading then predict that traders time purchases and sales for periods when the market is especially deep, avoiding those periods when market depth is low. If so, the pooling of liquidity should result in markets in which depth is associated with more trades and larger trade size.

Confirmation that trading activity is indeed positively related to liquidity requires some control for other factors that may affect activity. With regards to trading activity, a natural object of interest is trading frequency. Since this variable is discrete and can take on the value 0, we model trading activity using a Poisson model. Let N denote the number of trades in a five

minute interval and \mathbf{X} denote a vector of explanatory variables. Then, with $\ln(\lambda) = \beta\mathbf{X}$, the Poisson model is:

$$\Pr[N = n | \mathbf{X}] = \frac{e^{-\lambda} \lambda^n}{n!}; \quad n = 0, 1, 2, \dots, \quad (2)$$

We estimate this model separately for buys and sells. Table 4 contains coefficients and standard errors (in parentheses) for Poisson models of trade arrivals for buyer-initiated and seller-initiated trades. Estimates are computed by maximum likelihood techniques, based on 5-minute intervals over 105 trading days. The vector \mathbf{X} includes a constant, the number of trade arrivals on the opposite side of the market (“side”), returns (measured as change in the midquote), open and close dummies, depth of the market up to six ticks away from the midquote, and the effective spread, computed for trade sizes of 20 contracts. All estimated coefficients are statistically significant for both sides, and are of the expected sign.

Trading activity is positively related to depth and negatively related to spreads; both have economically and statistically significant effects. Taken together as measures of liquidity, both results reinforce the hypotheses stemming from discretionary entry into the market. An increase in order arrivals on the *opposite side of the market* implies greater activity. The finding highlights the theoretical prediction of Scharfstein and Stein (1990), that high contraside order flows generate entry on the other side of the market, consistent with greater pressures to trade quickly.

The coefficient estimates for returns are consistent with the hypothesis that traders place buy orders following market dips and sell following price upturns. Further, as traders observe upwards price pressure, they tend to place more sell-side orders at prices away from the best quotes, accounting for part of the result. Open and close dummies are positive. There is nothing new about this result, since it is consistent with the well-known U-shaped volume pattern observed in other markets. Taken together many other studies over many markets, the finding suggests that market structure has little influence on the informational and behavior influences leading to U-shaped activity over the course of the trading day.

5. Dynamics of Liquidity and Returns

We now turn to an investigation of the dynamics of market liquidity and its time-varying effect on returns, and vice versa. The method of analysis is reminiscent of Hasbrouck's (1991)

examination of specialist quote setting. The goal in that paper is to relate specialist quote revisions to trades, modeled as empirically signed volume. In doing so, Hasbrouck identifies the effects of random trade innovations on quote revisions, and interprets a measure of the expected cumulative quote revision as an index of private information. The measure used is the impulse response function of a bivariate vector autoregression.

There is no specialist in a limit order book market, and the changes in midquote prices used by Hasbrouck (1991) are more naturally interpreted here as returns to trading activity. Our interest centers upon the interplay between liquidity and prices, and in the dynamic relationship between liquidity characteristics on opposite sides of the market.

Throughout our discussion we use market depth as our measure of liquidity and change in midquotes as a proxy for returns. Our conclusions also hold for other metrics of liquidity including price impacts. As in Hasbrouck (1991), a generalized vector autoregression framework is used, and we first turn to the assumptions underlying our particular model.

5.1 Identification and the Statistical Model

We begin with the following complete dynamic system or structural model,

$$RY_t = \sum_{s=1}^q B_s Y_{t-s} + v_t, \quad (3)$$

where Y_t and v_t are vectors and R and B_s , $s=1, \dots, q$, are matrices. This is closely related to a reduced form model,

$$Y_t = \sum_{s=1}^q A_s Y_{t-s} + \eta_t, \quad (4)$$

where $A_s = R^{-1}B_s$ and $R^{-1}v_t$.

Use of the complete dynamic system, as opposed to simply the reduced form, has two main advantages. First, estimates of the complete model also include *contemporaneous* influences, permitting description of current period effects on market liquidity itself. Second, it permits explicit delineation of the identification conditions required to isolate shocks to market liquidity. These conditions often are hidden in the estimation of the reduced form alone, confusing inference with respect to the shocks of interest.¹⁵

¹⁵ There is a large literature devoted to this point, starting with Sims (1986) and explicated in more detail in Hamilton (1994).

In terms of estimation, the difficulty often encountered in structural estimation is that there are more parameters than moments. Therefore, we have to make meaningful restrictions in order to identify R and B_s . The identification conditions chosen here are expressed in terms of the variance-covariance matrix of v_t and the elements of the matrices R and B_s . Identification is similar to that of a Wold causal chain.¹⁶ In our case, the covariance matrix of the structural error is block diagonal, restrictions are imposed on R such that the matrix is block triangular, and returns follow a unit root process by a restriction imposed on $B_l \equiv B$. We make the latter assumptions explicit below, once the elements of Y have been specified.

5.2 Specification and Estimation of Market Liquidity Dynamics

Our primary interest, beyond a characterization of the dynamics of liquidity, is in the dynamic relationship of returns with depth. We therefore specify the vector Y_t as $(D_{bt}, D_{at}, \Delta m_t)'$, where Δm_t is the change in the quote midpoint, and depth on the bid and sell side, D_{bt}, D_{at} , are six ticks away. A variety of additional elements of Y suggest themselves, but we found that depth at six ticks away to be the most representative.

Theoretical treatments of the relationship between liquidity and returns are essentially static in nature. Our approach to identification is therefore empirical, using elements of the techniques in Swanson and Granger (1997) and Sims (1986). The combination of techniques involves the use of different identification schemes, each allowing the assessment of the strength of various correlations among the variables. The scheme below represents a choice based on this procedure, but also is intuitively plausible in nature.

The variance-covariance matrix of the structural error vector is taken to be block diagonal. In particular, it is assumed that shocks to liquidity on the bid and offer sides of the market are contemporaneously correlated. Returns are assumed to be uncorrelated, which is supported by the data. Lag lengths are truncated at $s = 1$. The matrix of contemporaneous effects, R , is specified as

$$R = \begin{bmatrix} 1 & 0 & -\rho_{13} \\ 0 & 1 & -\rho_{23} \\ 0 & 0 & 1 \end{bmatrix}. \quad (5)$$

¹⁶ See, for example, Sims (1986).

The matrix of lagged effects, B , is unrestricted, with the exception of the coefficient on lagged returns, which is set to zero.

The combination of restrictions has the following economic intuition. Neither bid nor offer side depth contemporaneously affect returns. This has some intuitive appeal, in that depth is a function of bids and offers, which naturally precede transactions. As such, bid and offer depth should affect returns in the next period, if at all, which is allowed by the specification. Similarly, depth on one side of the market does not contemporaneously affect depth on the other side, but does so with a lag. Identification schemes that permit estimation of contemporaneous effects of depth on returns and side of the market yield economically and statistically insignificant R -matrix coefficients.¹⁷ On the other hand, the model assumes that shocks to depth on the bid and offer sides of the market are correlated, since such shocks may derive from the same source of market information.

The specification permits a contemporaneous effect of returns on depth in both sides of the market. Price movements influence the current submission of bids, offers, and cancellations, reflected in the depth measures. Prior returns also have an influence on current depth in the specification. The inclusion of both contemporaneous and lagged effects permits a test as to whether discretionary behavior, manifested through returns, has any instantaneous or lagged feedback into liquidity provision. The relative strength of the contemporaneous and lagged influences of returns on liquidity is an empirical question.

Based on the above identification conditions, equation (3) is estimated by method of moments, and the standard errors are computed using the usual GMM form. Results are reported in table 5 for liquidity measured in terms of number of contracts available at six ticks away from the quote midpoint.

The liquidity clustering predicted by Admati and Pfleiderer (1988) is clearly evident from the estimates and standard errors. The first-order serial correlation of depth with lagged depth ranges from 0.33 to 0.38 and is very precisely estimated. The correlation of depth on the offer side with lagged buy side liquidity is 0.04, and statistically significantly different from zero. Although this coefficient, and that relating lagged sell side liquidity to current buy side depth, are

¹⁷ Hasbrouck (1991) maintains a timing convention where trades contemporaneously influence quote revisions, but not vice versa. The cited test suggests that the same interpretation cannot be used here, and we use the opposite timing convention for liquidity and returns.

economically small, the results do suggest not only that the liquidity clustering hypothesis holds even across buy and sell sides, but also that the entry predictions of Scharfstein and Stein (1990) and Spiegel and Subrahmanyam (1995) appear to hold. We investigate the last point further in the context of the impulse response functions.

The contemporaneous impacts of returns on market depth are symmetric and different from zero at any reasonable level of statistical significance. As returns rise, liquidity increases on the offer side of the market and falls on the bid side. Lagged returns are both economically and statistically insignificantly different from zero in terms of their effect on liquidity.

These results clearly do not derive from the mechanics of a limit order book market. Simple mechanics would imply that buying pressure increases depth on the buy side, at least for prices at or very near the best quote, for example. Such results would be expected only for depth measured in terms of number of contracts available very close to the quote midpoint. In fact, this empirical phenomenon is observed only for depth measured at two ticks away from the midpoint in our sample.¹⁸

The findings have an interpretation consistent with the results on management of transactions costs. An increase in prices occurs due to pressure on the buy side of the market. Some sellers may simply hit the bid in a rising market, reducing depth at the top of the book on the bid side, but this is relatively costly. Generally, buying pressure implies that buyers must pick contracts off the offer curve in order to achieve execution. Stale bids below the best quote are cancelled, further reducing bid-side liquidity. The response of sellers is to put in offers at prices higher than the prevailing best offer quote in the market. As a result, liquidity on the offer side rises, as returns go up. In a rising market, this order placement behavior achieves savings in transactions costs due to price impact.

Conversely, decreases in liquidity on the bid side, and increases in liquidity on the offer side, are associated with larger returns, but with a lag. The observed relationship suggests that the effect of liquidity shocks upon returns is dynamic and potentially persistent, and we now turn to an analysis of the interplay between the two over time.

¹⁸ An example of when this phenomenon could be mechanical is the following. Suppose the best offer was at price 100 for 10 contracts, but an order came in wanting 20 contracts at price 100. In this case, 10 contracts would be traded, and the best bid would now be 10 contracts at 100. Given this, it is possible to construct a scenario where the spread actually decreases, depth on the sell side increases, and depth on the bid side decreases. But we found little evidence of this occurring here.

5.3. Impulse Response Functions

The dynamic responses of returns to market liquidity shocks, and those of depth on one side of the market to shocks on the other side, are computed based on the estimated version of equation (3) specified by full simultaneous equations model,

$$Y_t = \sum_{s=1}^q \hat{R}^{-1} \hat{B}_s Y_{t-s} + \hat{R}^{-1} \hat{v}_t. \quad (6)$$

This autoregression is transformed into its infinite order vector moving average representation, through the device of matching moments.¹⁹ The moving average representation is then used to generate the impulse response functions.

Table 6 contains results for shocks to liquidity and returns, illustrated graphically in Figure 2. Results are presented for shocks to liquidity on the bid side (panel A), on the offer side (panel B), and for shocks to midquote returns (panel C). Dynamic responses are given for the first five minutes, as well as average responses over time periods following the initial shock, up to 60 minutes. Shocks to market liquidity consist of an increase in depth of 30 contracts. Shocks to returns are in units of 5 ticks.²⁰ Responses for liquidity are measured in terms of number of contracts; those for spreads and returns are given in terms of ticks.

Shocks to liquidity on the offer side of the market tend to lower returns, while increases in liquidity on the bid side raise them. Given that the market structure explicitly mimics the interaction between supply and demand, the result is intuitively plausible. A positive shock to liquidity on the bid side produces the analogue of a demand-side imbalance. If the market equilibrates itself through an implicit calculation of a Walrasian equilibrium, prices rise, and returns are positive. This effect has been observed in experimental studies of limit order markets.²¹

The effects are short-lived, in that virtually all of the impact occurs during the first 10 minutes following the liquidity event. The response of returns with respect to a liquidity shock also is moderate in magnitude. A simple calculation shows, for example, that an increase in bid depth of 62 contracts, or roughly two standard deviations, is required to increase returns by a sin-

¹⁹ See Hamilton (1994, chapter 11).

²⁰ The precise scaling is immaterial, given the linearity of the system. A shock of 90 contracts to depth, for example, results in a response that is 3 times what is given in the table. The size of the shocks illustrated here was chosen to be approximately one standard deviation.

²¹ Unpublished results, gathered from personal communication with Charles Plott at the California Institute of Technology. See <http://eeps.caltech.edu/> under “animations” for description of the experiments and results.

single standard deviation.²² On an annualized return basis, these impacts are much larger, of course.

An increase in liquidity on one side of the market leads to a rise in liquidity on the other side. Interpreted as a form of liquidity clustering, the result is confirmatory of the predictions of Admati and Pfleiderer (1988) with respect to discretionary timing of trading activity. Alternatively, the findings support the predictions of Spiegel and Subrahmanyam (1995) and Scharfstein and Stein (1990). In those papers, herding behavior also involves entry on the opposite side of the market, given increases in order flow activity.

6. Volatility

It is generally assumed that increased market liquidity is associated with lower volatility, and vice versa. Such a prediction also follows naturally from the theories relating to discretionary timing of trades. On the other hand, there is little direct empirical evidence on this point, to the best of our knowledge. Rather, trading volume and the absolute value of price changes are commonly found to be positively correlated, and there is some evidence that the volatility/volume correlation extends to common factors in prices and volumes.²³ We now extend the investigation of the last section to include an analysis of the dynamic interactions between depth on the order book, effective spreads, and volatility.

6.1. Regression Results

Volatility is easily captured in our present framework. We redefine the vector Y_t in equation (3) as $(D_{bt}, D_{at}, |\Delta m|_t)'$, where $|\Delta m|_t$ is the absolute value of the change in the quote midpoint. The same identification scheme is employed as before. The correlation of current and lagged absolute returns is left unrestricted, however, following the large literature on volatility clustering. Results are reported in table 7 for liquidity measured in terms of number of contracts available at six ticks away from the quote midpoint.

²² The standard deviation of returns is 5.185 ticks, the measured response is 2.513, and $(5.185/2.513) = 2.063$, times 30 contracts is 62. Other calculations summarized in text are done similarly.

²³ See, for example, Karpoff (1987), Gallant, Rossi, and Tauchen (1992), and Hasbrouck and Seppi (1999).

Volatility has a contemporaneous, statistically significant negative effect on liquidity, regardless of side of market.²⁴ The result stands in sharp contrast to the typically trading volume/volatility relationship, in which the positive correlation between variables typically is attributable to information effects (e.g., Blume, Easley, and O'Hara (1994)). In an open limit order book system, higher volatility increases the value of the free option stemming from liquidity provision to the order book. Periods of higher information intensity and concomitant higher volatility increase the likelihood of adverse selection, and adverse selection effects have been found to be large in electronic markets.²⁵ In both cases, the incentive to provide liquidity to the book in the form of limit orders decreases, and market liquidity falls. The good news is that the effects on liquidity are relatively short-lived, so that the market self-corrects. In other words, while our results on liquidity might be taken as an argument for a trading halt, the natural resiliency of the market obviates this measure.

Conversely, increases in market liquidity lower future price volatility. The result is intuitively plausible, and consistent with the findings of Bollerslev and Domowitz (1991) in their investigation of the relationship between volatility dynamics and generic order book systems. The effects are economically larger, and statistically significant, on the bid side of the market, relative to the offer side. The difference might be thought to represent variability in this particular sample, since there is no obvious reason for a disparity. On the other hand, the literature on trading costs suggests that costs are substantially higher for sells than for buys in both traditional market structure (Keim and Madhavan (1998)) and electronic venues (Domowitz and Steil(1999)). Evidence from these cost studies is consistent with the fact that volatility does not respond significantly to offer-side depth, remaining relatively high even when the market is relatively deep on the sell side.

6.2. The Dynamic Relationship Between Liquidity and Volatility

The dynamic responses of shocks to liquidity and volatility are summarized in table 8 and Figure 3, for liquidity defined in terms of number of contracts 6 ticks away from the midquote. As in the previous analysis, we report the initial 5-minute effect, as well as averages over subperiods within the hour following the shocks. The magnitude of the shocks to liquidity is as dis-

²⁴ The effects of lagged volatility on depth are economically negligible and statistically insignificantly different from zero.

²⁵ See Kofman and Moser (1997) and Coppejans and Domowitz (1999).

cussed previously. Shocks to volatility represent an increase of 5 ticks, or about 0.1 percent of contract value.²⁶

Increases in market liquidity lower volatility. The volatility impacts of the liquidity shocks die away quickly, with the responses over the 15 to 25 minute interval being only 16 to 18 percent of the average impacts over the first 10 minutes. The standard deviation of volatility is about 3.5, and the 5-minute impact is -1.619 , so a shock of $(3.5/1.619) \times 30$ or 65 contracts to depth is required to move volatility by one standard deviation. Shocks to liquidity on one side of the market move the other side of the market in the same direction as the initial shock. These results are unsurprisingly similar to those obtained using the structural VAR system incorporating midquote returns.

Shocks to volatility not only have a contemporaneous effect on liquidity, but also a more persistent effect over time. Higher volatility clearly decreases liquidity over the hour following the shock. The effects are especially strong only in the first 10 minutes following the volatility event, consistent with our overall findings of high natural market resiliency. Further, the magnitude of the effect of a volatility shock is relatively small. An increase in volatility of 5 percent of value decreases bid depth by only 14 contracts, for example, less than the average trade size.

7. Conclusion

The rapid adoption of electronic limit order book systems (or automated auctions) for equities, derivatives, and bonds worldwide has generated considerable practitioner and academic interest in the operation of such markets. In particular, many questions concern the nature and characteristics of liquidity in automated systems because of their reliance on public limit orders. This paper analyzes the dynamic links between market liquidity, order placement behavior, and returns in an electronic limit order book market. The study also examines questions of interest beyond those relating to the precise structure of an automated auction. In particular, the nature of data availability has allowed us to confirm several theoretical predictions, previously unexplored, while raising theoretical issues regarding the interaction between liquidity and discretionary trading activity, that are as yet untreated.

²⁶ Average 5-minute volatility over the estimation period is 3.67 ticks, with a standard deviation of 3.6 ticks. A move of two standard deviations is approximately the size of the average bid-ask spread.

We construct measures of liquidity and market depth using natural metrics from the limit order book. These measures vary widely over time, suggesting that traders can add value by strategic order placement behavior. We document evidence in favor of this hypothesis. Specifically, the actual execution costs are significantly lower than the costs that would be incurred under a naïve strategy that fails to account for time-variation in liquidity. Cost differences are especially pronounced for larger trades. A simple trade frequency model supports our view that discretionary traders trade in high liquidity periods, reinforcing the concentration of volume and liquidity at certain points in time. Taken as a whole, the findings confirm the predictions of several models of discretionary trading behavior. Even the sharp result of Mendelson and Tunca (2000), that realized trading costs should be equalized over trade size, is confirmed by the data.

We examine the dynamic relation between measures of liquidity and short-horizon expected returns using structural vector autoregressive models. We find a high degree of autocorrelation in liquidity, consistent with liquidity clustering predicted by theoretical models. Time series dynamics are more complex than in previous analyses. There is, for instance, dynamic feedback from trading activity to liquidity, a feature missing from theoretical models and suggesting a new line of research. The dynamic interaction between volatility and liquidity is investigated using a similar dynamic framework. The use of liquidity metrics based on supply and demand curves inherent in the market structure allows for results that go beyond those available for inside-spreads, for example, and that stand in contrast to well-known volume-volatility relationships. Positive liquidity shocks reduce volatility. While short-lived, this effect is relatively large in magnitude. Volatility shocks reduce liquidity, supporting arguments for trading halts following sharp market movements. On the other hand, impulse response functions show that shocks to liquidity dissipate quickly, indicating a high degree of resiliency. Natural market resiliency also is suggested by the small magnitudes of the impact of large volatility shocks on market depth. This “self-correcting” ability of the automated auction is an important element of this mechanism’s success, and belies arguments that the “free option” problem is potentially fatal with respect to automated market viability.

References

- Admati, A. and Paul Pfleiderer, 1988, A theory of intraday patterns: Volume and price variability, *Review of Financial Studies* 1, 3-40.
- Amihud, Yakov, and Haim Mendelson, 1986, Asset pricing and the bid-ask spread, *Journal of Financial Economics* 17, 223-250.
- Amihud, Yakov, and Haim Mendelson, 1991, Liquidity, maturity and the yields on U.S. government securities, *Journal of Finance* 46, 1411-1426.
- Amihud, Yakov, Haim Mendelson, and Beni Lauterbach, 1997, Market microstructure and securities values: Evidence from the Tel Aviv stock exchange, *Journal of Financial Economics* 45, 365-390.
- 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, Pierre Hillion, and Chester Spatt, 1999, Price discovery and learning during the pre-opening period in the Paris Bourse, forthcoming, *Journal of Political Economy*.
- Blume, Lawrence, David Easley, and Maureen O'Hara, 1994, Market statistics and technical analysis: The role of volume, *Journal of Finance* 49, 153-181.
- Bollerslev, Tim, and Ian Domowitz, 1991, Price volatility, spread variability, and the role of alternative market mechanisms, *Review of Futures Markets* 10, 78-102.
- Brennan, Michael J., and Avanidhar Subrahmanyam, 1996, Market microstructure and asset pricing: On the compensation for illiquidity in stock returns, *Journal of Financial Economics* 41, 441-464.
- Brennan, Michael J., Tarun Chordia, and Avanidhar Subrahmanyam, 1998, Alternative factor specifications, security characteristics, and the cross-section of expected stock returns, *Journal of Financial Economics* 49, 345-373.
- Coppejans, Mark, and Ian Domowitz, 1999, Pricing behavior in an off-hours computerized market, *Journal of Empirical Finance* 6, 583-607.
- Domowitz, Ian, 1993, A taxonomy of automated trade execution systems, *Journal of International Money and Finance* 12, 607-631.
- Domowitz, Ian, and Benn Steil, 1999, Automation, trading costs, and the structure of the securities trading industry, *Brookings-Wharton Papers on Financial Services*, 33-82.
- Foster, F. Douglas, and S. Viswanathan, 1990, A theory of the interday variations in volume, variance, and trading costs in securities markets, *Review of Financial Studies* 3, 593-624.
- Foster, F. Douglas, and S. Viswanathan, 1992, Variations in trading volume, return volatility, and trading costs: Evidence on recent price formation models, *Journal of Finance* 48, 187-212.
- Gallant, A. Ronald, Peter E. Rossi, and George Tauchen, 1992, Stock prices and volume, *Review of Financial Studies* 5, 199-242.

- Glosten, Lawrence, 1994, Equilibrium in an electronic open limit order book, *Journal of Finance* 49, 1127-1161.
- Gouriéroux, Christian, Gaëlle Le Fol, and Bertrand Meyer, 2000, Analysis of Order Queues, CREST working paper.
- Hamilton, James D., 1994, *Time Series Analysis* (Princeton University Press, Princeton, NJ).
- Harris, Lawrence E., 1990, Liquidity, trading rules, and electronic trading systems, Monograph 1990-4 (Leonard N. Stern School of Business, New York University, New York).
- Hasbrouck, Joel, 1991, Measuring the information content of stock trades, *Journal of Finance* 46, 179-207.
- Hasbrouck, Joel, and Duane J. Seppi, 1999, Common factors in prices, order flow, and liquidity, working paper, Stern School of Business, NYU.
- Hollifield, Burton, Robert A. Miller, and Patrik Sandås, 1999, An empirical analysis of limit order markets, Working paper, GSIA, Carnegie Mellon University.
- Irvine, Paul, George J. Benston, and Eugene Kandel, 2000, Liquidity beyond the inside spread: Measuring and using information in the limit order book, working paper, Emory University.
- Karpoff, Jonathan M., 1987, The relation between price changes and trading volume: A survey, *Journal of Financial and Quantitative Analysis* 22, 109-126.
- Keim, Donald, and Ananth Madhavan, 1998, The cost of institutional equity trades: An overview, *Financial Analysts Journal* 54: 50-69.
- Kofman, Paul, and James T. Moser, 1997, Spreads, information flows, and transparency across trading systems, *Applied Financial Economics* 7, 281-294.
- O'Hara, Maureen, 1995, *Market Microstructure Theory*, Cambridge, MA: Blackwell Publishers.
- Mendelson, Haim, and Tunay I. Tunca, 2000, Transactions costs, information asymmetry, and dynamic market design, Working paper, Stanford University.
- Sandås, Patrik, 1999, Adverse selection and competitive market making: Evidence from a limit order market, Working paper, University of Pennsylvania.
- Scharfstein, David, and Jeremy Stein, 1990, Herd behavior and investment, *American Economic Review* 80, 465-479.
- Sims, Christopher A., 1986, Are forecasting models usable for policy analysis?, *Quarterly Review*, Federal Reserve Bank of Minneapolis, Winter, 1-16.
- Spiegel, Matthew, and Avanidhar Subrahmanyam, 1995, On intraday risk premia, *Journal of Finance* 50, 319-340.
- Subrahmanyam, Avanidhar, 1991, A theory of trading in stock index futures, *Review of Financial Studies* 3, 37-71.
- Swanson, Norman R., and Clive W.J. Granger, 1997, Impulse response functions based on a causal approach to residual orthogonalization in vector autoregressions, *Journal of the American Statistical Association* 92, 357-367.

Table 1
Average Depth of the Book by Tick, Time and Side

This table contains summary statistics relating to the depth of the order book, in number of contracts, by time of day, averaged over 105 trading days. Data for the bid side appears in Panel A, and data for the offer side appears in Panel B. Column headings indicate the number of ticks away from the midpoint of the best quote in the market at the time. The figures reported are the number of contracts available at or below that number of ticks away from the midquote. Numbers in parentheses are the probability, in percent, of observing volume at the indicated number of ticks away from the midquote. Numbers in brackets are the standard deviation, in number of contracts, of depth over the full sample.

Panel A: Bid side of the book

Time	4	6	8	10	12	16	20
9:15	12	25	37	48	58	77	86
10:15	26	42	58	84	109	140	143
12:15	21	37	56	80	103	129	137
14:15	25	38	58	77	102	130	137
15:00	31	50	63	82	95	117	124

Panel B: Offer side of the book

Time	4	6	8	10	12	16	20
9:15	12	24	33	47	60	79	92
10:15	28	42	60	81	108	139	145
12:15	18	33	50	68	92	126	132
14:15	26	40	56	77	102	127	133
15:00	27	46	58	77	94	113	122

Table 2
Hypothetical Price Impacts by Time of Day

This table contains the price impact of trades, reported in percentage terms relative to the quote midpoint, by time of day. Calculations are done for hypothetical trades of 10 to 100 contracts in increments of 10, compared with the observed order book at a specific time of day, averaged over 105 trading days. Figures in the row marked “average” are computed based on computations at 15 minute intervals over the trading day, averaged over intervals and trading days. Panel A contains data for transactions at the bid, and Panel B contains figures for transactions at the offer. Trades at the bid are necessarily negative, and the absolute value is reported here.

Panel A: Bid Transactions

Time	10	20	30	40	50	60	70	80	90	100
9:15	0.08	0.09	0.10	0.12	0.13	0.14	0.15	0.16	0.17	0.19
10:15	0.06	0.07	0.08	0.09	0.10	0.11	0.12	0.12	0.13	0.14
12:15	0.07	0.08	0.09	0.10	0.11	0.12	0.12	0.13	0.14	0.15
14:15	0.06	0.07	0.08	0.09	0.10	0.11	0.12	0.13	0.14	0.15
15:00	0.06	0.06	0.07	0.08	0.09	0.10	0.11	0.12	0.13	0.14
Average	0.07	0.07	0.08	0.09	0.10	0.11	0.12	0.13	0.14	0.15

Panel B: Offer Transactions

Time	10	20	30	40	50	60	70	80	90	100
9:15	0.08	0.10	0.11	0.12	0.13	0.14	0.15	0.17	0.18	0.19
10:15	0.06	0.07	0.08	0.08	0.10	0.11	0.12	0.12	0.13	0.14
12:15	0.07	0.08	0.09	0.10	0.11	0.12	0.13	0.14	0.15	0.15
14:15	0.06	0.07	0.08	0.09	0.10	0.11	0.12	0.13	0.14	0.15
15:00	0.05	0.06	0.07	0.08	0.09	0.10	0.11	0.12	0.13	0.14
Average	0.07	0.07	0.09	0.10	0.10	0.11	0.12	0.13	0.14	0.15

Table 3
Actual Price Impacts by Time of Day

This table contains the price impact of trades, reported in percentage terms relative to the quote midpoint, broken down by time of day, by side (bid or offer), and for regular trades and crosses. Calculations are done for actual trades of 10 to 100 contracts in increments of 10, compared with the observed order book at the time of trade, over 105 trading days. Trades at the bid are necessarily negative, and the absolute value is reported here.

	10	20	30	40	50	60	70	80	90	100
Bid Side	0.04	0.04	0.04	0.04	0.05	0.05	0.04	0.04	0.06	0.04
Offer Side	0.04	0.04	0.05	0.05	0.05	0.05	0.05	0.05	0.04	0.07
Cross Bid	-----	-----	0.05	0.04	0.05	0.05	0.05	0.07	0.03	0.05
Cross Offer	-----	-----	0.04	0.05	0.05	0.04	0.06	0.06	0.02	0.05

Table 4
Poisson Models of Trade Arrivals

This table contains coefficients and standard errors (in parentheses) for Poisson models of trade arrivals for buyer-initiated and seller-initiated trades. Estimates are computed based on 5-minute intervals over 105 trading days. The specification of the conditional mean is $E[y|\mathbf{X}] = \exp(\beta' \mathbf{X})$, where y is the number of trades in a five minute period and \mathbf{X} denotes the vector of explanatory variables. The vector includes a constant, the number of trade arrivals on the opposite side of the market (“side”), returns (measured as change in the midquote), open and close dummies, depth of the market up to six ticks away from the midquote, and the effective spread, computed for trade sizes of 20 contracts.

	Buy-side	Sell-side
Constant	0.401 (0.058)	0.456 (0.055)
Side	0.054 (0.005)	0.066 (0.005)
Return	-0.099 (0.003)	0.106 (0.004)
Open	0.373 (0.047)	0.356 (0.045)
Close	0.227 (0.040)	0.310 (0.041)
Depth	0.024 (0.005)	0.028 (0.005)
Effective Spread	-0.006 (0.002)	-0.006 (0.002)
R ²	0.147	0.165

Table 5
Coefficient Estimates for the Model of Depth and Returns

This table contains estimates of the dynamic simultaneous equations model,

$$RY_t = BY_{t-1} + v_t,$$

in which $Y_t = (D_{bt}, D_{at}, \Delta m_t)'$, where Δm_t is the change in the quote midpoint, D_{bt} is depth of market, measured in lots of 10 contracts on the bid side of the order book at 6 ticks away from the quote midpoint, and D_{at} is the same measure, computed for the offer side of the book. The matrix, R , is given by

$$R = \begin{bmatrix} 1 & 0 & -\rho_{13} \\ 0 & 1 & -\rho_{23} \\ 0 & 0 & 1 \end{bmatrix}$$

Figures in the table are coefficient estimates (GMM robust standard errors in parentheses) for the regression of each of the elements of Y_t (column headings) on the variables in the left hand column. Estimation is based on 5-minute intervals.

	Bid depth	Offer depth	Δ midquote
Constant	2.344 (0.065)	2.345 (0.067)	-0.192 (0.131)
Δ midquote _t	-0.027 (0.007)	0.033 (0.007)	-----
Bid depth _{t-1}	0.384 (0.017)	0.035 (0.012)	0.084 (0.021)
Offer depth _{t-1}	0.016 (0.012)	0.326 (0.018)	-0.058 (0.022)
Δ midquote _{t-1}	0.008 (0.005)	-0.003 (0.006)	-----

Table 6
Dynamic Responses to Shocks in Depth and Returns

This table contains the dynamic responses (impulse response function estimates) of bid-side depth, offer-side depth, and midquote returns, to shocks to market depth on the buy side (Panel A), market depth on the sell side (Panel B), and returns (Panel C). Depth of market is measured in number of contracts on bid and offer sides of the order book at 6 ticks away from the quote midpoint. Calculations are based on five-minute intervals, and use coefficient estimates of a complete dynamic simultaneous equations model, also estimated over 5-minute periods. Figures in the first row, labeled “5 minutes” are responses to the initial shock. The remainder of the rows give figures for average effects over the interval indicated (e.g., 15-25 minutes is the response calculated for five minute periods, starting at 15 minutes and ending at 25 minutes, averaged over the period). Depth responses are given in number of contracts. Return responses are given in number of ticks.

Panel A: 30 Contract Shock to Depth on Bid Side

	Bid depth	Offer depth	Δ midquote
5 minutes	11.46	1.122	2.513
5-10 minutes	7.920	0.951	1.704
15-25 minutes	0.878	0.240	0.162
30-60 minutes	0.030	0.009	0.005

Panel B: 30 Contract Shock to Depth on Offer Side

	Bid depth	Offer depth	Δ midquote
5 minutes	0.534	9.720	-1.734
5-10 minutes	0.450	6.452	-1.125
15-25 minutes	0.108	0.503	-0.071
30-60 minutes	0.003	0.009	-0.000

Panel C: 5 Tick Shock to Midquote Returns

	Bid depth	Offer depth	Δ midquote
5 minutes	-0.008	0.036	-0.021
5-10 minutes	-0.006	0.024	-0.012
15-25 minutes	-0.000	0.002	-0.000
30-60 minutes	-0.000	0.000	-0.000

Table 7
Coefficient Estimates for the Model of Depth and Volatility

This table contains estimates of the dynamic simultaneous equations model,

$$RY_t = BY_{t-1} + v_t,$$

in which $Y_t = (D_{bt}, D_{at}, |\Delta m_t|)'$, where $|\Delta m_t|$ is volatility, measured as the absolute value of the change in the quote midpoint, D_{bt} is depth of market, measured in lots of 10 contracts on the bid side of the order book at 6 ticks away from the quote midpoint, and D_{at} is the same measure, computed for the offer side of the book. The matrix, R , is given by

$$R = \begin{bmatrix} 1 & 0 & -\rho_{13} \\ 0 & 1 & -\rho_{23} \\ 0 & 0 & 1 \end{bmatrix}$$

Figures in the table are coefficient estimates (GMM robust standard errors in parentheses) for the regression of each of the elements of Y_t (column headings) on the variables in the left-hand column. Estimation is based on 5-minute intervals.

	Bid depth	Offer depth	\Delta m
Constant	2.730 (0.084)	2.626 (0.087)	3.247 (0.114)
\Delta m _t	-0.085 (0.010)	-0.070 (0.010)	-----
Bid depth _{t-1}	0.373 (0.017)	0.032 (0.012)	-0.054 (0.015)
Offer depth _{t-1}	0.146 (0.012)	0.321 (0.017)	-0.021 (0.015)
\Delta m _{t-1}	-0.006 (0.008)	0.001 (0.008)	0.193 (0.019)

Table 8
Dynamic Responses to Shocks in Depth and Volatility

This table contains the dynamic responses (impulse response function estimates) of bid-side depth, offer-side depth, and volatility, measured as the absolute value of midquote returns, to shocks to market depth on the buy side (Panel A), market depth on the sell side (Panel B), and volatility (Panel C). Calculations are based on five-minute intervals, and use coefficient estimates of a complete dynamic simultaneous equations model, also estimated over 5-minute periods. Figures in the first row, labeled “5 minutes” are responses to the initial shock. The remainder of the rows give figures for average effects over the interval indicated (e.g., 15-25 minutes is the response calculated for five minute periods, starting at 15 minutes and ending at 25 minutes, averaged over the period). Depth responses are given in number of contracts. Volatility responses are given in number of ticks.

Panel A: 30 Contract Shock to Depth on Bid Side

	Bid depth	Offer depth	$ \Delta\text{midquote} $
5 minutes	11.33	1.068	-1.619
5-10 minutes	7.833	0.915	-1.283
15-25 minutes	0.855	0.234	-0.230
30-60 minutes	0.023	0.009	-0.008

Panel B: 30 Contract Shock to Depth on Offer Side

	Bid depth	Offer depth	$ \Delta\text{midquote} $
5 minutes	0.492	9.678	-0.639
5-10 minutes	0.426	6.411	-0.497
15-25 minutes	0.111	0.495	-0.081
30-60 minutes	0.005	0.009	-0.002

Panel C: 10 Tick Shock to Volatility

	Bid depth	Offer depth	Δmidquote
5 minutes	-0.280	-0.191	0.996
5-10 minutes	-0.206	-0.184	0.604
15-25 minutes	-0.029	-0.018	0.023
30-60 minutes	-0.001	-0.000	0.000




Figure 1: Linear and nonlinear estimates of average depth, based on a fifth degree polynomial approximation. The darker line is the bid side, and the lighter line is the ask side.

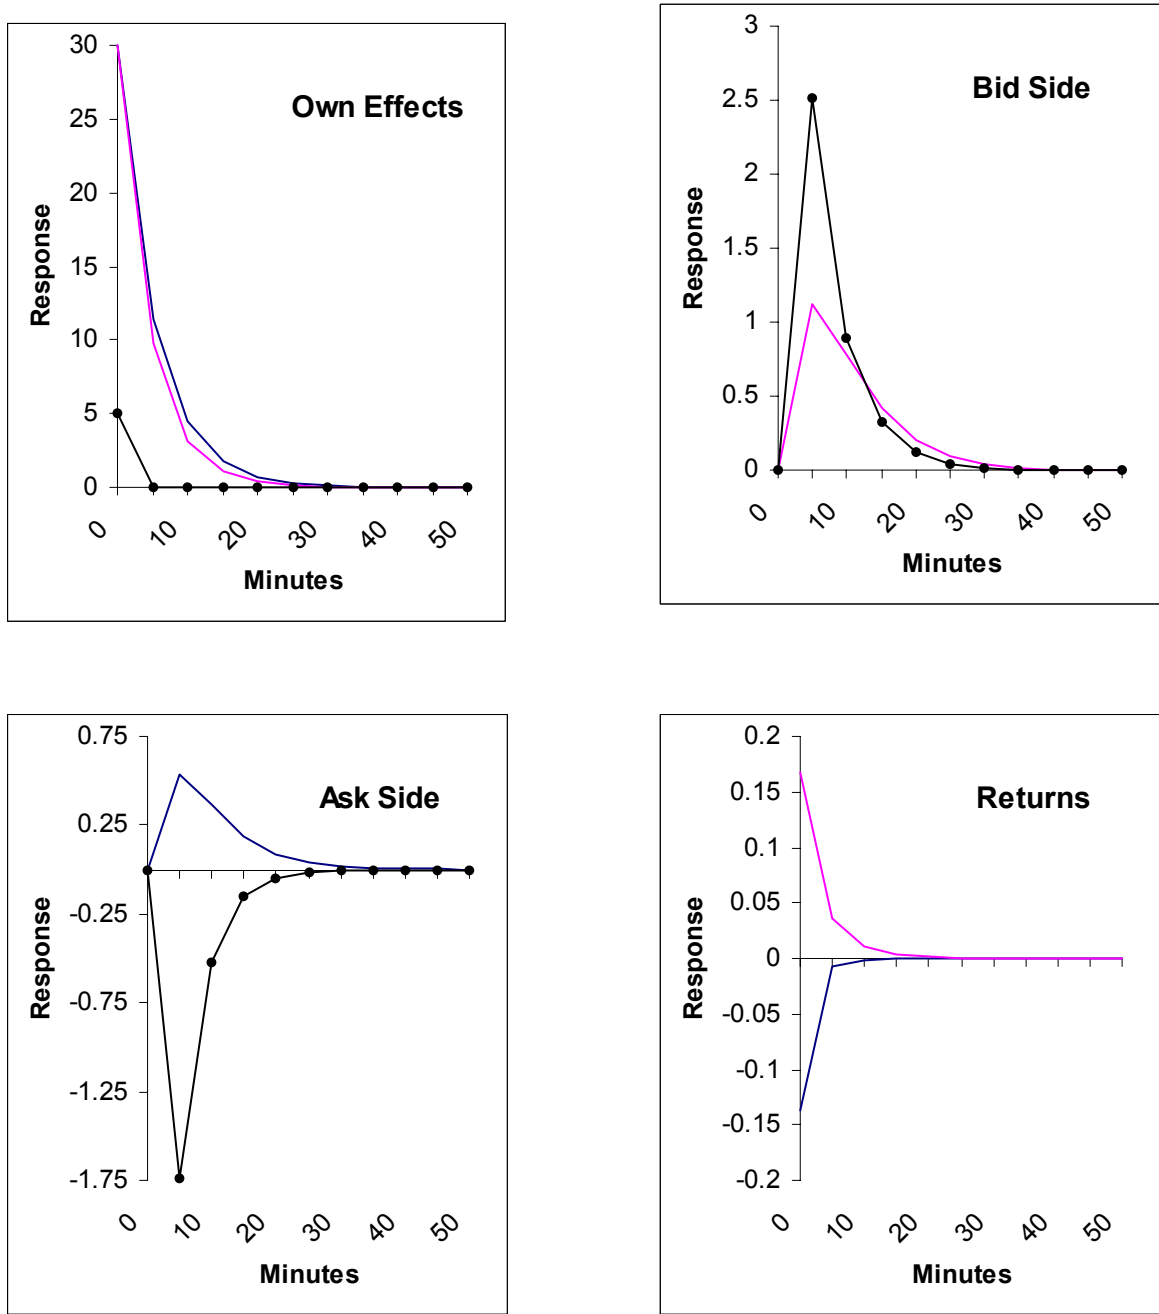


Figure 2: Impulse Responses to Shocks in Depth and Returns. Dotted line is returns, darker solid line is bid side depth, and lighter solid line is ask side depth. Own effects represents the effect of a shock on bid side depth, ask side depth, and returns on bid side depth, ask side depth, and returns, respectively. The three other plots capture the remaining responses. For example, the plot Bid Side represents the effects on ask side depth and returns given that bid side depth has been shocked. The plots Ask Side and Returns are defined analogously. Note that the dots and the dashes on the horizontal axis are in increments of five minutes. Shocks to depth are in units of 30 contracts; shocks to returns are in units of 5 ticks.

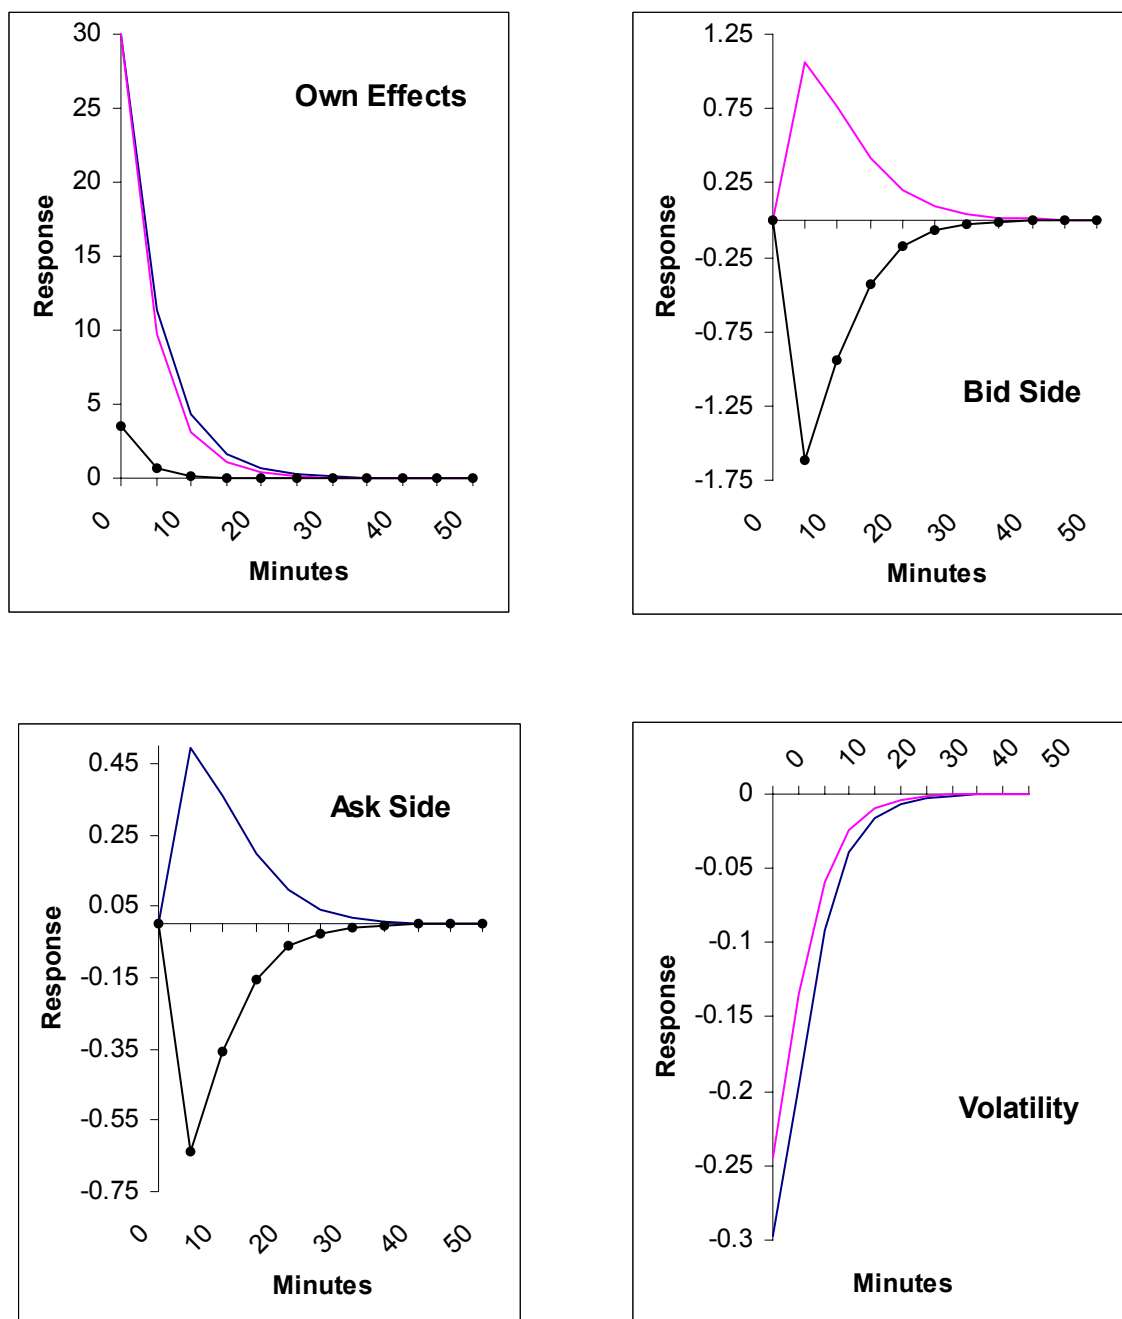


Figure 3: Impulse Responses to Shocks in Depth and Volatility. Dotted line is volatility, darker solid line is bid side depth, and lighter solid line is ask side depth. Own effects represents the effect of a shock on bid side depth, ask side depth, and volatility on bid side depth, ask side depth, and volatility, respectively. The three other plots capture the remaining responses. For example, the plot Bid Side represents the effects on ask side depth and volatility given that bid side depth has been shocked. The plots Ask Side and Volatility are defined analogously. Note that the dots and the dashes on the horizontal axis are in increments of five minutes. Shocks to depth are in units of 30 contracts; shocks to volatility are in units of 5 ticks.