The Liquidity of the Market Portfolio

Prachi Deuskar and Timothy C. Johnson*

May 2009

Abstract

Asset pricing theory suggests that liquidity only affects prices if claims to the market portfolio display time-varying liquidity. Most empirical research on aggregate liquidity has had to rely on indirect measures constructed from liquidities of individual stocks. These measures may differ significantly from the liquidity of an claim on the index itself. We directly study the liquidity of the market portfolio via the price impact of order flow in the S&P 500 futures. Using identification through heteroskedasticity to address simultaneity of order flow and returns, we find that flow strongly and permanently affects prices. We construct a directly observable *ex-ante*, real-time measure of illiquidity via the slope of the limit order book. This measure is a highly significant predictor of subsequent price impact, with a coefficient that is typically close to one. From its dynamics we find that *(i)* the nonflow component of return volatility decreases liquidity, but the volatility of order flow does not; *(ii)* trading volume has only a transient effect on liquidity; and *(iii)* liquidity varies positively with order flow itself. Our results point to limited risk bearing capacity, rather than asymmetric information or temporary price pressure, as a primary determinant of market illiquidity.

KEYWORDS: LIQUIDITY DYNAMICS, SYSTEMATIC RISK, LIMIT ORDER BOOK.

JEL CLASSIFICATIONS: G12, G14

---

*University of Illinois at Urbana-Champaign. We are grateful to the Chicago Mercantile Exchange for making the data in this study available. Chen Yao provided excellent research assistance for the project. We thank seminar participants at UIUC for many many helpful comments and suggestions.

Asset pricing theories ascribe a central role to the market portfolio – a hypothetical claim to all productive processes in the economy – because the logic of diversification dictates, in many settings, that investors will hold no other risky securities. The expected risk and return of such a portfolio then become the dominant quantities for financial decision making. For the same reason, the *liquidity* of such a portfolio – defined as the cost of executing an informationless trade – is the dominant quantity in theories that attempt to explain why liquidity matters to investors.[1]

Following this logic, measuring aggregate liquidity and quantifying its risk has become an important goal, both in academic research and in investment practice. However, the theoretical nature of the portfolio itself means that one must rely on indirect measurement strategies.

A natural approach is to use microstructural data on individual stocks, and then form some kind of cross-sectional average or common factor.[2] This has often seemed the most feasible approach. Yet it has two problems. As a practical matter, the constructed measures for individual securities often differ widely from each other at any point, and tend to be only weakly correlated.[3] In principle, too, it is not clear that aggregating such measures is appropriate. It is not obviously the case that the liquidity of a claim to a portfolio should be equal to an average of the constituent securities' liquidities. Indeed, as highlighted by Subrahmanyam (1991) and Gorton and Pennacchi (1993), an important feature of index claims is precisely that they may eliminate what may be the dominant component of individual securities' liquidity: asymmetric information. Averages of asymmetric information are quite different from asymmetric information about the average.

An alternative approach is to try to assess the liquidity of claims to large segments of the market. The last decade has seen widespread growth in the market for exchange traded funds (ETFs), such as SPDRs, which represent shares in the S&P 500 index. At the same time, the market for stock index futures has continued to deepen, and the adoption of electronic limit order books by the major exchanges has greatly enriched the available data. In this paper, we use high-frequency data on the S&P 500 e-mini futures contracts to study aggregate liquidity. We argue that the size of this market as well as its

[1] See, for example, Pástor and Stambaugh (2003) and Acharya and Pedersen (2004).

[2] See, for example, Chordia, Roll, and Subrahmanyam (2000), Hasbrouck and Seppi (2001) and Amihud (2002).

[3] See, for example, Table 2 in Hasbrouck (2006) or Table 2 in Goyenko, Holden, and Trzcinka (2008).

direct link (via index arbitrage) to the underlying index make its liquidity characteristics precisely those that a well-diversified investor would care about.[4]

Besides data on prices and order flow, we also have real-time observations of the electronic limit order book for the S&P 500 futures. We establish empirically that the slope of this limit order book is a highly significant estimator of the market impact of subsequent trades. This result means, in effect, that we have direct observations of the illiquidity of the market. This allows us to study its level and its variation in detail and without need of a further estimation step.

We use our data to directly address some fundamental issues, including:

1. *How illiquid is the market?*

2. *Why is the market illiquid at all?*

3. *How risky is market liquidity?*

4. *Why does market liquidity change over time?*

Regarding the first and third question, we do find that the market is much less than perfectly liquid and that there is wide variation in illiquidity over time. On average, a trade of 5000 index units (or 100 e-mini contracts) moves the S&P 500 by 0.03 points. At the end of our sample, the index level was approximately 885, so this corresponds to a "price pressure" of 0.34 basis points for a trade of 4.4 million dollars. This impact, however, can often differ by a factor of two within a trading day. It is also subject to substantial persistent shocks, having swung from about 0.01 to 0.10 during the course of our sample.

As to the "why" questions (the second and the fourth above), we do not claim to provide complete answers, yet our findings do shed some new light on both.

The issue of why the market portfolio should be less than perfectly liquid goes to the heart of our understanding of price formation. Many explanations of the illiquidity of individual securities (or even market sectors) would seem weaker when applied to aggregate claims. As noted, the impact of asymmetric information should be mitigated

[4]Other papers studying trading costs of market index claims include Hegde and McDermott (2004), and Berkman, Brailsford, and Frino (2005).

since it would require non-public information on the entire economy.[5] Search frictions, limited competition, and capital immobility should also be minor concerns as the market in question is monitored constantly by traders around the globe and participation requires almost no fixed cost.

An important finding on this question is the simple fact that the price impact of orders is not transitory. In fact, our specifications imply that the long-term impact of a given trade hardly decreases from the initial impact. This casts doubt on stories of liquidity determination based on temporary "price pressure" due to market frictions or irrationality.

Turning to dynamics, we find that although illiquidity is extremely variable, it is also highly predictable. Lagged predictors account for over 95 percent of its variation. This has potentially important implications for pricing theories, because only the unspanned component of liquidity risk should affect expected returns.

While we document substantial predictability, there are some interesting negative results. We find no relationship between illiquidity and market returns, which differs from results obtained with average liquidities of individual stocks. We also find that liquidity does not decrease with *order flow uncertainty*, which it might be expected to if asymmetric information were a primary determinant of illiquidity. Last, the level of overall market activity (measured by volume) appears to have only transient effects on liquidity. This poses a challenge to the view that liquidity is driven by the profits of intermediaries or specialist providers.

On the other hand, we document significant predictive roles for both return volatility and signed order flow. The former result is in line with previous research, to which we contribute by isolating the non-flow driven component of volatility. The latter result is new. Both findings are consistent with the model of Pagano (1989), which views liquidity as primarily driven by the equilibrium risk-bearing capacity of the market as a whole.[6]

In summary, this paper fills an important gap in the empirical literature on liquidity and liquidity risk by providing detailed analysis of a direct measure of the liquidity of

[5]Theories of systematic variation in asymmetric information have been proposed by Eisfeldt (2004) and Evans and Lyons (2004). A structural model of market wide private information is developed and estimated in Albuquerque, De Francisco, and Marques (2008).

[6]The model of Johnson (2007) formalizes these relationships and predicts that signed flow, not trading activity, should drive liquidity.

the market portfolio. The findings include some novel results that have implications for our understanding of the causes of aggregate illiquidity and its risks.

The remainder of the paper is organized as follows. The next section introduces our data sample and documents the strong sensitivity of index prices to order flow. Section II connects the illiquidity of the market to a particular directly observable function of the state of the limit order book. Section III analyzes the time-series variation in this measure. The final section summarizes our contribution.

I. Illiquidity of a Market Index

This study attempts to directly measure the illiquidity of claims to broad portfolios of assets, as opposed to aggregating measures of component illiquidity. In this section, we describe a market for one such index claim and outline how we think about its liquidity. We separately identify structural shocks to order flow and to prices and quantify the extent to which flow shocks move the market.

A The S&P 500 e-mini Futures Market

The S&P 500 index of U.S. stocks is the most widely used proxy for the market portfolio, both in research and in practice. The underlying stocks are almost all liquid and easily shortable, and collectively account for approximately 75 percent of the value of U.S. equities.[7] Huge sums of money are committed to tracking funds that explicitly mimic the index, and even more are benchmarked against it. Investors can also directly trade the index via exchange traded funds, stock index futures and options, and a panoply of over-the-counter derivatives, such as return swaps. Arbitrage links all of these markets, both to each other and to the underlying stocks. The arbitrage relationships are well-understood and actively enforced. As a result, prices of any of these claims accurately represent the level of the index. Further, the price *responses* of any of these markets are extremely highly correlated with each other.

We use the latter fact to justify studying the liquidity of the market via the liquidity of a particular subset of derivatives: the electronically-traded S&P 500 futures. Specifically,

[7] See S&P 500 index fact sheet by Standard and Poor's.

we examine the most active contract of the most active version of these futures, the near-month small-denomination futures known as the e-mini contract.[8] These contracts have grown to dominate the index futures market. (See Panel A of Table I.) Moreover, prior research has established that, among the different equity-linked markets, index futures are the quickest to react and impound new information.[9] Recently, comparing the e-mini contracts to other index contracts (including the traditional S&P futures) and to the most liquid exchange traded funds (SPDRs), Hasbrouck (2003) established that e-mini futures were the most important venue for price discovery. The dominance of these contracts, in terms of volume and liquidity, has continued to increase from the time of that study. It thus seems justifiable to identify price changes in this contract as true innovations in the market portfolio.[10]

A further maintained hypothesis is that order flows in the futures market – buy orders and sell orders – are representative of aggregate order flows. While we do not have information on order flow in all the competing products, Panel B of Table I shows that the correlations of daily dollar volume across the main markets for the index are around 0.90. Intuitively, one would expect order flow to be most representative in the market with the lowest trading costs, and certainly the e-mini futures are among the lowest-cost vehicles for adjusting market exposure. Panel C of Table I presents a measure of average illiquidity, *Illiq*, suggested in Amihud (2002), for different markets.[11] The numbers certainly reinforce the idea that the e-mini futures are the lowest cost market to trade the index. While SPDRs are comparable, *Illiq* for the futures is the lowest of the markets shown and is three orders of magnitude smaller than *Illiq* for the individual stocks. Theoretically, if one imagines liquidity providers as having a preferred mix of derivatives and stock positions, then a customer sale, say, in any one market ought to result in adjustment sales by the intermediary in the other linked markets.

Having stated the case for regarding (1) price changes and (2) order flow for the e-mini futures market as valid measures for the market portfolio as a whole, our immediate focus of investigation is the response of the first quantity to the second.

[8] Each e-mini futures contract has a notional value of 50 dollars multiplied by the index level. These contracts are traded on CME Globex, an electronic platform. Other than denomination and trading mechanism, they are identical to the traditional pit-traded S&P 500 futures contract.

[9] Many studies, starting with Kawaller, Koch, and Koch (1987) have documented that index futures lead cash stocks. Fleming, Ostdiek, and Whalley (1996) extended this conclusion to index options.

[10] Note that this is tantamount to the assertion that changes in the cash-futures "basis" represent transient error in the (measured) stock prices.

[11] *Illiq* for day t is calculated as $Illiq_t = \frac{|return_t|}{Dollar\ Volume_t}$.

The sample period for this study is February 2006 through January 2009. The Chicago Mercantile Exchange generously provided us with real-time data on the limit orders and trades for e-mini futures. Figure 1 shows the time series of prices and volume over the sample. The turbulent nature of this period offers econometric benefits, in that we observe interactions over a wide range of conditions.[12] At the same time, we are sensitive to the possibility that the interactions we document may not correspond to those in more stable times.

The limit order data consists of snapshots of the order book whenever there is a change in any limit order. It provides the time of the observation and the total limit order quantity at each price on bid and ask side, up to a maximum depth of five price levels on each side. The trade data consists of time of the trade, quantity traded, and price. The dataset has limit orders and trades for all the e-mini futures contracts that have any activity. However, at each time only one contract is actively traded. Typically, this is the contract with shortest maturity. Roughly ten days before the maturity date, this contract almost stops trading and the contract with the next shortest maturity becomes most active. We focus only on the most active contract at any time. While our data cover overnight Globex trading, activity in the market is sparse outside U.S. trading hours. Our tests use data from 8:30 AM to 3:14 PM, Central time.

Using the limit order data, for each minute, we take the midpoint of the latest best bid and ask quotes before the end of the minute as the observation price. Per-minute price changes are calculated as a simple difference between midpoint prices for the current minute and the previous minute. Per-minute returns are the price change divided by mid-price for the previous minute. Trading volume is calculated as the sum of quantity (number of contracts) traded during a minute. To construct order flow, we assign trades as buy or sell using the algorithm suggested in Lee and Ready (1991) as follows: if the trade price is higher than the most recent mid-price, we treat the trade as a "buy". If the trade price is lower than this mid-price, we treat it as a "sell". For every minute, we define order flow as the difference between total quantity for buy trades and total quantity for sell trades.

[12] The futures market does include temporary price limits that can be triggered on extreme moves. We have confirmed with the CME that no such instances occurred during the sample period.

B Joint Determination of Returns and Flows

Figure 2 shows the relationship between order flow and concurrent price changes for our sample period. The figure measures price changes in both raw and percentage terms and order flow in both contract numbers and monetary value. Regardless of the scaling, there is clearly a strong positive correlation between order flow and price changes. Flow apparently moves prices. Such a relationship has been well documented in the foreign exchange markets (see Evans and Lyons (2002)) and the U.S. Treasury market (for example, Brandt and Kavajecz (2004)). We are not aware that this effect has been previously reported for aggregate market claims such as stock index futures.

A causal interpretation of the flow-return relation requires caution, however. We address identification in our formal estimation below. For now, note simply that the response of prices to flow is the sense in which the market is illiquid, and the degree of this illiquidity can be measured by the slope of the plots. From the lower right panel of the plot, the raw magnitude of the response is roughly one percent of the index value per billion dollars worth of purchases or sales. On absolute terms, a market in which the price impact of a 10 million dollar order is only one basis point is a very liquid one. On the other hand, the market capitalization of the S&P 500 stocks is of the order of magnitude of 10 trillion dollars (it changes by a factor of two over our sample). Relative to this, a market in which an order of size equal to 1/1000th of total value of the securities moves prices by one percent is significantly less than perfectly liquid.

Before going on, we note one other salient feature of our plots. The hour-glass shape of the relation is highly suggestive of a relationship that, while approximately linear, has a changing slope. Figure 3 shows the same scatter plot over sub-samples of the data. Not surprisingly, the market was substantially more illiquid during the turmoil of 2008. Sections 3 and 4 of the paper are explicitly concerned with modelling the time-variation in illiquidity. The present section will, however, estimate an unconditional linear relationship in order to address more basic measurement issues. In particular, we now develop a structural bivariate model of flow and returns that allows us to identify their simultaneous relationship.

To truly gauge illiquidity, we must isolate the component of flows that is exogenous to returns. Returns and flows may comove for purely mechanical reasons, such as the presence of stale limit orders. More deeply, as discussed by Obizhaeva (2008), causation can run from returns to flows whenever agents employ price-contingent trading strategies.[13] For purposes of measuring illiquidity, we are interested in the response of the market to the component of order flow that cannot itself be viewed as caused by returns.

Econometrically, we wish to identify the simultaneous effect parameters in the joint system:

$$BY_t = AX_{t-1} + \epsilon_t \tag{1}$$

where X is a vector of exogenous regressors and

$$Y = \begin{bmatrix} \text{return} \\ \text{flow} \end{bmatrix}. \qquad B = \begin{bmatrix} 1 & -b_r \\ -b_f & 1 \end{bmatrix} \tag{2}$$

and

$$\epsilon_t \sim \mathcal{N}(0, \Sigma), \qquad \Sigma = \begin{bmatrix} \sigma_r & 0 \\ 0 & \sigma_f \end{bmatrix}.$$

To do this, we employ the method of identification through heteroskedasticity (ITH) introduced by Rigobon (2003). As is well known, the matrix B above cannot be directly identified due to the simultaneity bias in the ordinary least squares (OLS) coefficients. However, if we have two separate samples in which the innovation variances change, then the coefficients in B are restricted by the application of the orthogonality condition, $\mathrm{E}[\epsilon_r \epsilon_f] = 0$, to both regimes. This can be seen by noting that the variance-covariance matrix of the reduced-form residuals $B^{-1}\epsilon$ differs across regimes and its elements define a system of six equations in the six parameters $b_r, b_f, \sigma_{r,1}, \sigma_{f,1}, \sigma_{r,2}, \sigma_{f,2}$. Equating these to sample moments then yields consistent estimates of the response coefficients.

To illustrate intuitively, suppose σ_r is constant but that there are periods of high and low σ_f. (Also assume X is not present for simplicity.) If f were truly exogenous (i.e. $b_f = 0$) then there is nothing wrong with OLS and the plot of r versus f would exhibit the same slope in both periods. On the other hand, if $b_f \neq 0$, OLS is biased, and the magnitude of the bias depends on the ratio σ_r/σ_f. So if we observe the slope of the plots changing from one regime to the next, we can infer the amount of the bias from the degree of change.

[13] At the level of executions, this includes the use of market-limit orders or stop-loss orders. In the former case, positive return shocks would lower (signed) flow as buy orders are cancelled when the market moves up. In the latter case, the opposite happens.

As a general matter, there may be many volatility regimes (rather than two), in which case, the moment conditions provide more restrictions than the minimum necessary for identification. Applying a method of moments minimization to the full collection then yields estimators of increasing efficiency. In the present case, the two data series exhibit wide variations in volatility over time. (See Figure 4.) Our identification scheme posits that the ratio of the two series switches levels every p days, where p is a tuning parameter that we vary. It is not necessary, for this procedure, to specify a particular stochastic law for these volatility changes across these regimes. The identification will be stronger if there is wider variation in the volatility ratio between regimes (and likewise weaker if there is substantial variation within regimes). However, Rigobon (2003) establishes consistency of the method even under conditions of regime misspecification.

While identifying simultaneous effects is crucial for causal inference, lagged effects are also extremely important in our high-frequency series. Allowing a sufficiently flexible lag structure is necessary both to avoid misestimating simultaneous effects and to accurately model longer horizon evolution. In our baseline specifications, we include 30 lags of flow and returns. Note that our procedure directly estimates the structural form of the lag coefficients (the matrix A above) not the reduced form.[14] For brevity, we do not report all these coefficients below, instead describing the findings qualitatively and summarizing their quantitative impact via impulse response functions.

Table II shows the results of our estimation of the impact of flows on returns. The table uses each of the pairs of variables plotted in Figure 2 and several choices of the regime identification parameter p. The top panel displays the estimated simultaneous impact parameter, which in all cases is highly significant.[15] The bottom panel reports coefficients b_r and b_f when returns are measured either as price change or proportional returns and order flow is measured in terms of number of contracts or dollars. In all cases, both flow and returns are scaled by their respective standard deviations so that the results are comparable across specifications. As can be seen, the results are very consistent irrespective of how we measure the variables.

The first rows in each panel show the OLS coefficients of price impact. They are

[14]Due to the large number of variables, we do not treat the lag coefficients as free parameters in our estimation of the simultaneous effects. They are estimated by OLS within each minimization step. This is equivalent to a two-stage GMM procedure. The standard errors that we report account for the joint dependence of the two stages.

[15]The table reports asymptotic standard errors computed from the general covariance matrix for extremum estimators. See the Appendix for details.

always substantially larger than the ITH estimates. This implies that there is significant endogeneity of flows: returns also appear to positively cause flows. The coefficient estimates for b_f are always positive and statistically significant. Thus simultaneity bias is a factor in assessing price impact. However, economically, the scale of the estimated coefficients is not that different from the slopes of the scatter plots discussed above. The bias does not change the sense of the initial conclusions as to the degree of illiquidity in the market.

Another way of judging the economic significance of the flow effect on prices is to note how much of measured price variation is in fact due to flow. The fourth column of the first panel shows that this fraction is about 40 percent.[16] This a huge number. Not only does flow move prices, it accounts for nearly half of measured return variance. Even without modelling conditional variation, almost half of the risk in the market portfolio (at least at this short horizon) is due to order flow uncertainty.[17]

Do the results tell us anything about the causes of market illiquidity? If, as we suggest, price changes and flows in the futures market are representative of aggregate quantities in the primary market, as well as in the other derivatives markets, then the illiquidity we document is unlikely to be due to details of market microstructure in the e-mini futures limit order book.[18] We also observe that deep pools of arbitrage capital are available across related markets, making local market segmentation an unlikely explanation.

Nevertheless, probably the most common interpretation of illiquidity is just as "price pressure." This notion carries the implication that, since flow pushes prices away from "fundamentals", price impact should be largely temporary. Our estimation allows us to speak to this supposition directly. The results do not support it.

Table III examines the persistence of the price impact under our estimated specifications. For each, we report I_0, the initial price response, in units of own standard deviation, to a unit standard deviation shock to flow. Also shown is I_∞, the infinite horizon cumulative impulse response. The long-horizon number tells us that, far from being temporary price pressure, the responses of price to flow shocks are long-lasting.

[16]The table also shows a very large R^2 for the flow equation. Understanding the causes of variation in order flow is an important topic, both theoretically and empirically. However, it is beyond the scope of the present work.

[17]The explanatory power does include some contribution from the lagged variables. But this is small.

[18]We note that simultaneous trading takes place in equivalent contracts via open outcry; the component stocks and ETFs trade in specialist markets; and most indexed-linked derivatives trade in dealership markets.

On average, about 90 percent of the initial response will not reverse according to these estimates.

Thus far we have documented a highly significant response of prices to the exogenous component of order flows. This is the sense in which we view the market portfolio as illiquid. Our findings are qualitatively robust to several (unreported) variations of our basic specification, such as the choice of lag length and sample sub-period. We emphasize, however, that the constant-impact model employed so far is undoubtedly misspecified. We introduce it first to establish some basic facts about market illiquidity and our estimation methodology. There is, in fact, rich variation in price impact over time, which involves measurement issues of its own. We now turn to these.

II. A Direct Measure of Market Illiquidity

Having analyzed unconditional illiquidity of the S&P 500 futures, we now construct an *ex ante* observable measure of price impact using limit order book data for the e-mini contract. We establish that this measure is highly relevant for predicting subsequent price impact of order flow. Thus, this high-frequency real-time measure can be a direct proxy of market illiquidity, allowing us to study its dynamics in a rich setting.

A The Limit-Order Book

The limit order book data, described in Section A, allows us to continuously gauge, not just the depth or quantity of orders, but the sensitivity of that depth to price. Thus it resembles the market's demand curve (on the bid side) and supply (on the ask side) for units of the security. Of course, it is not a true supply/demand schedule: it is contaminated both by the exclusion of the orders of agents not actively participating in the market, and by the inclusion of stale orders that no longer represent a willingness to trade at the indicated price.[19] Thus it is an empirical question whether summary statistics computed from the book do accurately predict the subsequent price response to traded quantities. That is what we will establish in Section B.

We first describe our procedure for summarizing the information in the limit order

[19]Note that CME Globex does not allow for hidden limit orders.

book.[20]

As an initial step, we construct a dataset that gives snapshots of the book every minute.[21] We match the data on returns, volume and order flow for each minute to the latest limit order book prior to that minute.

Next, for each limit order book observation, we construct a slope measure by fitting a line through cumulative quantities against limit order prices. Specifically, the *inverse limit order book slope* (ILOBS) is calculated as follows:

$$ILOBS = \frac{\sum_{i=1}^{K} Mdist_i \cdot Mdist_i}{\sum_{i=1}^{K} Mdist_i \cdot CQ_i} \tag{3}$$

where K is the total number of limit order prices on bid as well as ask side, $Mdist_i$ is the distance of the ith limit order price from the mid-price, and CQ_i is the cumulative quantity in thousands of contracts at the ith limit order price. Mid-price is the midpoint of the best bid and best ask quotes for this limit order book. We treat bid side quantities as negative values, in line with the convention used for order flow calculation. Figure 5 graphically depicts the construction of ILOBS.[22]

ILOBS is designed to capture the effect of trade quantity on price and hence is a measure of price impact of (potential) trades. Its units quantify the effect of a market order of one thousand contracts on the index level, holding the limit orders fixed. As we will show later, ILOBS is indeed a good predictor of the impact of subsequent orders. We view price impact as economically the most important component of illiquidity in this market. Liquidity can also be measured by the bid-ask spread and by the speed of

[20] Liquidity proxies have been constructed from limit order data in other markets by Irvine, Benston, and Kandel (2000), Coppejans, Domowitz, and Madhavan (2004) and Gomber, Schweickert, and Theissen (2004).

[21] At any one time, there may be less than five orders on the bid or offer sides. Thus there is a trade off between getting the most recent snapshot (at the end of every minute) and getting most the complete limit order book. Within each minute we look for the latest limit order book that has five limit order prices on each side. If there are no such limit order book observations, we gradually look for fewer number of limit order prices on each side. If we do not find any observations within a minute, we go back minute-by-minute, using the same procedure, until we find an observation. This very rarely occurs during U.S. trading hours.

[22] We also tried alternate specifications of the slope estimation, including weighting each price-quantity in the book by the inverse of absolute value of $Mdist$ or the inverse of square of $Mdist$ to give more weight to the observations that are closer to the mid-price. Slope estimates calculated using these specifications are very highly correlated with ILOBS calculated as per equation (3) (see Panel B of Table IV) and yield nearly identical inferences. In the subsequent analysis, we focus only on ILOBS calculated as per equation (3).

reversal of price impact. We have presented some evidence above on the duration of price impact in this market. Below, we will also address the informativeness of the bid/ask spread.

Panel A of Figure 6 plots the daily median of ILOBS in our sample. As can be seen, ILOBS shows substantial variation in this period. During the volatile phase of late 2008 the limit order book is substantially flatter (its inverse slope is bigger).

Panel B of Figure 6 shows the intraday pattern of ILOBS. It plots median ILOBS for every minute of the trading day. Market illiquidity, as measured by ILOBS, shows substantial variation even at high frequency. When we study the dynamics of market illiquidity in the next section, we control for time of the day effects.

Panel A of Table IV presents descriptive statistics for ILOBS for the entire sample as well as for subsamples. In the latter subsamples, ILOBS is not only higher, it is also more volatile. Thus liquidity risk is changing over time. In the next section, we study dynamics of ILOBS in more detail. We now turn to the question of its usefulness in predicting the price impact of trade.

B ILOBS as a measure of Price Impact

As can be seen from Figure 3, price impact of order flow is not constant over time. If ILOBS is a good proxy of market illiquidity, order flow during an interval times ILOBS at the *start* of the interval should yield a good estimate of the price impact of that order flow. Thus we can measure the ability of ILOBS to capture variation in illiquidity by assessing its usefulness as a conditioning variable in our price change regressions.

In fact, if the limit order book did perfectly measure the market's supply/demand schedule, we would find a predictive coefficient on the product of ILOBS with flow that was near unity. A somewhat weaker test is that the coefficient should be the same in different subsamples. Figure 7 suggest that both conditions, in fact, hold to a good approximation. Panel A of the figure shows a scatter plot of price change against order flow times ILOBS for the entire sample. We can see that the slope of this scatter plot is indeed remarkably close to one. Panels B to D show the same scatter plot for three different subsamples. In all the subsamples, the slope of the scatter is also close to one. Thus, based on visual examination, it is apparent that ILOBS is nearly an unbiased measure of price impact of trade.

Next, we formally estimate the price impact of order flow conditional on ILOBS, using identification through heteroskedasticity. We estimate a system similar to (1), but instead of estimating an unconditional parameter b_r given in (2), we estimate a conditional b_r that captures the effect of ILOBS on the relationship between order flow and returns. Table V presents the results for conditional price impact estimation using a linear specification, where

$$b_r = b_0 + b_1 ILOBS. \tag{4}$$

Panel A presents results for returns calculated as price change and order flow in terms of thousands of contracts. The coefficient for the interaction of order flow and ILOBS b_1 is always highly significant, establishing the importance ILOBS. We also find that the coefficient b_0 for order flow is much smaller in Table V than the corresponding coefficient b_r in Panel A of Table II. Including a meaningful conditioning variable such as ILOBS reduces the unconditional impact of order flow. Panel B presents results for proportional returns and dollar order flow. Here, not only is the coefficient b_1 significant, but it also substantially closer to one, in line with the visual evidence.

The different levels of the interaction coefficient estimated in these specifications is largely due to the effect of outliers. Nonlinearities may matter for observations with extremely large flow relative to the size of the order book. Table VI shows the effect of outliers via changing the specification slightly. Panels A and B present a winsorized specification, where

$$b_r = b_0 + b_1 min(ILOBS, c). \tag{5}$$

The cut-off c for $ILOBS$ is set at the 99th or 98th percentile of ILOBS. Panel C presents a piecewise specification that conditions on outliers of the interaction of ILOBS and flow. Here we fit separate coefficients when the absolute value of this interaction is greater than 1.72 index points which corresponds to its 1st and 99th percentile. In particular, we estimate

$$b_r = (b_0 + b_1 ILOBS)(I_{abs(ILOBS \cdot Flow) \leq C}) + (b_2 + b_3 ILOBS)(I_{abs(ILOBS \cdot Flow) > C}). \tag{6}$$

We find that the coefficients b_1 in Table VI, in particular in Panel C, are indeed closer to the idealized value of one. Panel C also shows much smaller b_0 estimates, further reducing the unconditional price impact and reinforcing the conclusion that ILOBS captures illiquidity.

We next conduct various robustness checks, focussing on our piecewise linear specification, using ITH regime length of one day. Panel A of Table VII presents the results for the entire sample as well as for three subsamples. Here b_1 in the second and third subsamples covering 2007 and 2008 is close to one. The coefficient in the first subsample is smaller, yet still extremely statistically significant. The broad pattern of results is similar across subsamples.

In Panel B of Table VII, we vary the number of lags as well as the time interval over which price changes and order flows are calculated. The first two rows present results over 1-minute interval for 30 and 15 lags respectively. The last two rows show parameters estimated using 5-minute intervals for 30 and 6 lags respectively. Changing the number of lags does not change the results. The magnitude of b_1 is smaller for five-minute intervals. Looking at the intraday pattern of ILOBS shown in Figure 6, it is clear that ILOBS shows a lot of variation even at high frequency. Thus ILOBS five minutes ago, does not do as good a job of predicting price impact as ILOBS one minute ago. This explains the lower magnitudes of b_1 over five-minute intervals. But, again, in all specifications b_1 is highly significant.

Is ILOBS actually providing more information about liquidity than just using the bid-ask spread? Panel A of Table VIII present descriptive statistics for the bid-ask spread. This quantity shows very little variation over time. Even though ILOBS and bid-ask spread are, on average, about the same order of magnitude, the bid-ask spread has much smaller standard deviation. Essentially the spread is virtually a constant in our sample, and equal to the minimum tick-size for the futures contract. Panel B presents price impact regression results for linear, winsorized and piecewise specifications for ILOBS when including bid-ask spread times flow as an additional independent variable. Comparing Panel B with Tables V and VI, we find that the coefficients for ILOBS times flow do not change. The coefficient for bid-ask spread times flow is actually negative suggesting a lower price impact when the bid-ask spread is higher. Also, the statistical significance of ILOBS is undiminished. Thus, the bid-ask spread appears to convey little useful information about variation in market liquidity.

To summarize, the analysis in this section points to ILOBS being a nearly unbiased measure of the slope of the market's demand curve for index units, and hence of the price impact function facing investors. This conclusion holds at least for trade sizes up to the

98th or 99th percentile. It was not obvious that this would be so. The mapping from actual demand curves to limit order placement and cancellation strategies is not at all straightforward, depending on, among other things, the strategies of other participants and the details of the trading mechanism. (For a summary of the issues, see the survey of Parlour and Seppi (2008).) The results tell us that, for measuring liquidity, this complexity can be summarized to a good approximation by a single observable number.

III. The Dynamics of Market Liquidity

Having an observable *ex ante* proxy for the illiquidity of the market portfolio enables us to study the dynamic properties of illiquidity more directly than has been previously possible. Since ILOBS is not determined simultaneously with returns, volume, volatility, and order flow, we do not require further complicated identification schemes to separate exogenous and endogenous variation. ILOBS is also available at very high frequency providing us with a wealth of information, even in a relatively short sample. Using this data, we now investigate both the level and the drivers of aggregate liquidity risk.

To begin, compare the time-series plots in Panels A and B of Figure 6 that show, respectively, the evolution of market illiquidity over our sample period at daily frequency and the evolution over a typical day when observed at one minute intervals.

The lower-frequency data show that there is indeed substantial variability to market impact costs over long periods. The events of 2008 raised average impact by a factor of four or more. While the plot shows some mean-reversion in early 2009, certainly the shocks have been very persistent, lasting for periods of several weeks. On the other hand, the high-frequency data reveal substantial transient variability as well.

How risky, then, is aggregate illiquidity? Economically, the answer depends on the time-horizon of one's exposure, and one's required trading frequency.

The intraday standard deviation of the log of ILOBS is 0.24 and the standard deviation of log differences is 0.19. This means that illiquidity can change by about 20 percent *every minute*. This variability would add very significantly to the risk of an investor attempting to execute a continuous trading strategy at a short horizon (as might be required for hedging a barrier option, for example).

On the other hand, this risk is almost entirely transient. Longer horizon risk will be dominated by the persistent changes. Interday differences in log ILOBS have a standard deviation of about 0.17 and a first order autocorrelation of 0.95, implying that these shocks are the dominant component of unconditional liquidity risk. The unconditional standard deviation of log ILOBS is 0.59, of which 0.54 comes from the variation in intraday levels. A buy-and-hold investor with 1000 contracts in the market portfolio would have an expected liquidation cost of approximately 0.03 index points with a one standard deviation range of about 0.017 to 0.054. Clearly this risk becomes a smaller and smaller fraction of total portfolio risk (which scales with time) as the investor's horizon increases.

A second economically relevant question is how much of the liquidity variation is predictable. From an asset pricing perspective, it is only the covariance of the unpredictable component with marginal utility that would make agents demand compensation for liquidity risk.

This brings us to the search for the drivers of liquidity risk. Not only are we interested in the degree to which liquidity risk is spanned by other exogenous factors, but we also wish to shed light on the fundamental question of why the market is illiquid at all. Different hypotheses suggest different explanatory variables.

Dynamic theories of limit order book trading highlight the rewards to providing liquidity and imply that illiquidity should fall as market activity rises. Theories of asymmetric information imply that the risk of predicting the direction of informed demand should raise illiquidity. Theories that emphasize the equilibrium risk-bearing capacity of the market as a whole imply that signed order flow should decrease illiquidity. Our tests are among the first to address these *a priori* hypotheses using measures of liquidity of the market portfolio.

Moreover our data offer another advantage in that our proxy for market illiquidity is directly observable at fixed points in time. This means that we do not need to be concerned about the simultaneous determination of illiquidity with variables such as volume or volatility, which are measured over *intervals* of time. If we measure ILOBS at point t, for example, then the volatility (or volume or return) from $t-1$ to t is truly exogenous. Hence our proxy obviates the need for a more complex identification strategy, which would be required if illiquidity had to be estimated at simultaneous intervals with other variables.

Our investigation is, at the same time, limited to the extent that our sample is short. Our contribution is to exploit a large amount of high-frequency information. But we have little ability to detect causal relationships that unfold over weeks and months rather than minutes and hours. We also lack data on non-market quantities at the same time-scale as our market quantities. So, for example, we cannot correlate our observations with high-frequency data on interest rate changes or credit market conditions.[23]

We now describe our primary variables. The last section explained in detail the construction of the illiquidity measure, ILOBS. Here we employ the log of ILOBS. We truncate the variable at its 99.95th percentile, but do not otherwise transform large values. This is the dependent variable in our regressions.

While our data is updated continuously, in this section we resample at discrete intervals, typically over five minutes. The dependent variable is measured at the end of these intervals. Our primary independent variables are all effectively rates of flow (or changes) over the preceding interval.

Prior research has identified a negative association between returns and liquidity. We thus include as a control changes in the market index, which we measure in index points. (Recall that the units of our illiquidity measure are in index points per 1000 contracts.) We next include both net order flow and volume (gross order flow), both measured in contracts. We take the log of volume, which by itself is highly right skewed.

Our volatility controls are logs of the sum of squared changes over each time interval. We exploit the results of the last section that separately identified the non-flow-driven component of index innovations and the non-return-driven component of flow innovations. We use the volatility of both of these series.[24] Recall that the latter innovations feed through to returns and account for as much as 40-50 percent of high frequency return variance. As an alternative specification, we also employ the raw return volatility itself, which is effectively a weighted sum of our two separately identified components.

Table IX shows the correlation matrix of these variables. Some relationships are worth remarking on. Despite the large negative returns over our sample and the large increase in illiquidity, there is no evidence of an unconditional relationship between these

[23] High frequency data may be available on eurodollar and treasury futures contracts. Future work will explore the role of these quantities in liquidity determination.

[24] The precise specification used for this decomposition employs one-day regimes in our identification-through-heteroskedasticity methodology.

two quantities. Surprisingly, the unconditional correlation between illiquidity and volume is strongly *positive*. This is misleading, however, because of the well-known strong contemporaneous correlation between volume and volatility. We note that the relationship between volume and *flow* volatility is partially mechanical. However the same is not true for the non-flow component of return volatility. The latter is strongly positively correlated with illiquidity, whereas the former is not.

Table X contains our primary regression results. The table shows OLS estimated coefficients and Newey and West (1987) t-statistics (which adjust for residual autocorrelation and heteroskedasticity) for two specifications. Both specifications include time-of-day fixed effects. We normalize both dependent and independent variables by their standard deviation to put the coefficients into consistent units.

There are four primary findings that we wish to highlight.

1. *Flow and returns.* A "stylized fact" about illiquidity is that it increases in down markets. Our results do not support this: there is no significant relationship between illiquidity and contemporaneous or lagged returns. Instead, the data show a previously undocumented negative relation with order flow. Since we have already established the strong positive correlation between flows and returns, the results suggest that it is the flow-driven component of returns *only* that is responsible for the apparent relationship between illiquidity and returns.

2. *Volume.* We find that there is a significant negative contemporaneous response of illiquidity to volume, but that there is roughly equal *positive* lagged response. Measured by the sum of contemporaneous and lagged coefficients,[25] there is essentially no relationship between activity levels and liquidity.

3. *Volatility.* We do find a significant positive relation between market volatility and market illiquidity, in line with existing literature. We can go a step further, however, and determine that there is no such relationship with order-flow volatility. In fact, the latter relationship is significantly *negative.* (This is shown in the second panel of the table.) Hence the uncertainty that affects illiquidity is that of the non-

[25]The sum of the coefficients quantifies the cumulative response to a one-time level shift in volume holding the other independent variables fixed. To quantify the impulse response to a one-time structural shock to volume would require a full dynamic model of first and second moments (as well as an identification scheme), which are beyond the scope of the present work.

flow-induced component of returns, i.e. "fundamental" uncertainty not uncertainty about orders.

4. *Explained variation.* The regression residuals have a standard deviation of between 11 and 12 percent. The dependent variable itself has a standard deviation of about 60 percent. (Recall this is the logarithm of the price impact slope.) In terms of equivalent R^2, more than 96 percent of unconditional liquidity risk is spanned by its own lags, other market variables, and time-of-day effects.

We view our results as challenging some of the theories of liquidity determination discussed above. The volume findings are consistent with *changes* in activity – but not the level of activity itself – affecting levels of liquidity. This could be problematic for models in which activity drives the rewards to, and hence the amount of, liquidity provision.

Next, order flow uncertainty would seem to be a natural proxy for the probability of facing informed trade. But then the finding that this uncertainty does not increase (and actually decreases) illiquidity appears inconsistent with the idea that the degree of asymmetric information drives liquidity provision.

On the other hand, our results appear largely consistent with the view that market illiquidity is driven by fundamental risk. Moreover, the order-flow finding supports the conjecture that liquidity rises and falls with increased net participation, and hence with aggregate risk-bearing capacity.

The findings in Table X are robust to a number of alternative specifications and variable definitions. Adding more lags, splitting the sample in time, dropping the intraday fixed-effects, altering the variable transformations, and employing different first-stage specifications in the estimation of the two volatility components all had minor and economically insignificant effects on the conclusions. Inferences are also insensitive to the choice of lag length in the computation of standard errors. For brevity, we omit results for these cases.

One might also wonder whether our results are sensitive to the units of measurement of illiquidity. Recall that ILOBS is in units of price change per unit order flow. (And, for compatibility, the "return" variable in the regressions is also the raw price change, with the corresponding "volatility" variables constructed from the second moment of these.)

If illiquidity were instead expressed in units of *return* per unit flow, then we would be understating its rise over the course of our sample due to the large fall in prices. That, in turn, might account for our failure to detect a negative association between returns and illiquidity. However, this is not the case. We also perform our regressions scaling ILOBS by price level (and measuring returns as log-differences and likewise for volatility) and find the same relationships described above. In fact, the model fit (as measured by residual standard error) is improved slightly. Yet there is still no trace of a significant relationship of illiquidity with returns.

Another natural concern is the extent to which our methodology is picking up only ultra-high frequency effects that may be driven more by institutional or technological mechanisms rather than by deeper economic forces. We have noted above our limited ability to assess our findings over longer time horizons. However we can expand our sampling intervals within a day without losing too many observations. In addition, we can attempt to control for lower frequency variation in the independent variables.

Table XI shows two specifications that vary the time parameters. The top panel aggregates the variables to 15 minute intervals, and again includes six lags. The main findings are unchanged. The negative contemporaneous volume coefficient is now quite a lot smaller in magnitude than the sum of the lags, which are positive. The nonflow variance and flow variances effects are both considerably stronger, with the latter component retaining its puzzling negative sign.

To check for further influence of lower frequency evolution, the bottom panel of the table samples the data at hourly intervals and, in addition to three lags, includes averages of all variables over the prior day, week, and month. The sample size drops steeply, further weakening statistical significance.[26] However the simultaneous responses to flow, volume, and nonflow variance are all still significant and of the same sign.

To summarize, this section has shown that, while there is an enormous amount of variation in levels of market illiquidity, a very substantial proportion of this is locally predictable, and hence does not constitute a distinct source of risk to investors. Our flow data and our decomposition of return variance into flow-induced and nonflow-induced fluctuations has revealed some novel findings. Signed order flow itself predicts liquid-

[26]Recall that we do not include observations whose lags occur overnight (or on a previous day), and we do not use overlapping observations. So, with the lags, there are effectively only three hourly observations available per day.

ity, but order flow uncertainty does not. We have offered some interpretations of our estimation results in terms of differing paradigms for understanding aggregate liquidity determination. Going forward, the availability of almost continuous observation of market liquidity opens the door for direct estimation of structural models.

IV. Conclusion

In this paper, we study the liquidity of a claim to the market portfolio. Aggregate market liquidity is an economically important quantity, reflecting the resilience of asset markets as a whole and thus their ability to reallocate risk. As a practical matter, anyone engaging in dynamic portfolio strategies bears the burden of illiquidity as a cost and as a source of risk. Yet not much is actually known about the degree of market illiquidity, its variation, or the drivers of that variation.

One reason for this is measurement: until recently, little information has been available about trading in broad market claims. Researchers have thus relied on proxies built from averages of the liquidities of component stocks, or extracted common factors from these. However, these quantities are different, both in principle and in practice, from aggregate illiquidity. Moreover, different averaging approaches yield very different estimates, the units of which may be difficult to compare or understand economically.

We employ a rich, high-frequency data-set on orders and trades for the S&P 500 e-mini futures, which we argue represents the best available proxy for the aggregate portfolio. We construct and study simple, direct measures of illiquidity in this market.

We use identification through heteroskedasticity to simultaneously estimate the effects of order flow and returns on each other. We find that order flow has a significant and lasting impact on prices. In this sense, the market is illiquid, and illiquidity is not caused by temporary price pressure.

Market illiquidity is not constant. To study its dynamic properties, we construct a measure of price impact based on limit order book slopes. This is a direct *ex ante* measure of market illiquidity and hence does not suffer from simultaneity problems of regression-based price impact measures. It is available in real time and thus is useful for studying dynamics of market illiquidity at high frequency. We find that this measure is nearly unbiased and highly significant in predicting price impact of subsequent order flow.

We examine dynamics of this measure in the hope of shedding light on the causes of variation in market illiquidity. We find that, after controlling for order flow, market return has no effect on illiquidity. But positive order flow itself reduces illiquidity. Thus previously documented asymmetric nature of market liquidity appears to be due to the flow-driven component of return. Volume has only a transitory negative effect on illiquidity, most of which is reversed within a few hours. Volatility of order flow does not reduce illiquidity, as it might be expected to if asymmetric information is a primary driver of market illiquidity. However, volatility of the non-flow component of returns has a significant positive effect on illiquidity. This result, along with effect of flow itself on illiquidity, suggests that fluctuations in the aggregate risk bearing capacity of the market may play an important role in liquidity determination. As flow increases, more investors enter the market, increasing risk bearing capacity and hence reducing illiquidity. As return variance increases, risk increases, and agents' willingness to accommodate others' desired reallocation declines.

Our data span only three years, and thus limit our ability to study low frequency movements in market illiquidity. However, the sample does include some very interesting times, featuring wide variation in all the quantities under investigation. While none of our results is driven by the turbulence of 2008, the inclusion of this period suggests that the levels and relationships that we report hold quite generally.

Appendix

This appendix provides details of our identification-through-heteroskedasticity (ITH) estimation and the computation of asymptotic standard errors.

The estimation procedure identifies the simultaneous response coefficients by minimizing M moment conditions, where M is the number of distinct volatility regimes. The problem also requires estimation of the lagged response coefficients (the exogenous variables). Since there are a very large number of these (e.g. 30 lags of 3 variables for each of 2 equations), we do not include orthogonality conditions for them in the minimization. For this reason, the overall procedure is not a pure generalized method-of-moments (GMM) estimation.

Our algorithm consists of the following steps.

1. For a given partition into regimes, specify an estimate $\left\{H_m^{(i)}, G_m^{(i)}\right\}_{m=1}^{M}$ of the variances of the data, Y, whose components are the returns, r, and the flow, f (c.f. equation (1) in the text).

2. Specify a moment weighting matrix $W^{(i)}$.

3. For a candidate estimate $\hat{B}$ of the simultaneous response coefficient matrix:

  (a) Form the "right-hand side" residuals $u = Y\,\hat{B}$.

  (b) Estimate the coefficient matrix $\hat{A}$ from the generalized least squares (GLS) regression of u on the exogenous variables X, using the inverse roots of $\left\{H^{(i)}, G^{(i)}\right\}$ as weights. Call these "true" residuals ϵ.

  (c) Form the M-vector of within-regime cross products $g = (\epsilon_m^{r\prime}\epsilon_m^{f})_{m=1}^{M}$.

  (d) Form the ITH criterion $Q = g'\,(W^{(i)})^{-1}\,g$.

4. Find the minimizer, $\hat{B}^{(i)}$, of the criterion Q, repeating the inner loop for each candidate value.

5. Update $\left\{H^{(i+1)}, G^{(i+1)}\right\}$ to be the sample variances of the residuals ϵ in each regime.

6. Update $W^{(i+1)}$ to be the sample variances of the moment errors g in each regime.

7. Loop to step 3 until estimates converge.

The procedure mimics iterative GMM with asymptotically optimal choice of weighting matrix. As noted above, the inner GLS step departs from standard GMM. As a result, GMM efficiency results do not follow and GMM asymptotics are not applicable.

Also unlike standard GMM problems, the moment conditions we are minimizing pertain to distinct, non-overlapping subsamples (the regimes). Since these regimes are specified as having fixed duration, the large sample properties of the estimator do not derive from the number of observations per regime going to infinity. (We do not assume the regimes recur.) Instead, we derive standard errors as the number of regimes grows large.

To this end, we treat the ITH criterion function as an extremum estimator and apply the general theory of Amemiya (1985) (c.f. Chapter 4). Under conditions given there, *(i)* the second derivative of the criterion function with respect to the parameters converges in probability to a matrix E; *(ii)* the root-M scaled mean of the first derivatives of the criterion increments converges to a normal vector with covariance matrix, F; and *(iii)* the root-M scaled error of the parameter vector minimizing the criterion converges to a normal vector with mean zero and variance $E^{-1}\ F\ E^{-1}$. We estimate these matrices by their sample counterpart using numerical derivatives. Our asymptotic standard errors are the square roots of the diagonal entries of $M^{-1}\,\hat{E}^{-1}\ \hat{F}\ \hat{E}^{-1}$.

Note that this computation effectively treats an entire regime as a single observation and thus is analogous to clustering by regime. Since there are up to several thousand observations per regime, this is quite conservative.

To verify the applicability of this asymptotic analysis, we also conducted a bootstrap analysis for some of the specifications with regimes of one-day length. In particular, we randomly drew days, with replacement, from each of the three subsamples: February 2006 to January 2007, February 2007 to January 2008, and February 2008 to January 2009. Thus, in the bootstrap sample each subsample had same number of days as the original sample. We stratified the sample into subsamples so that each bootstrap sample had enough variation in volatility ratio σ_r/σ_f across regimes. Identification is stronger if there is wider variation in the volatility ratio across regimes. (See Section B for a discussion of role played by the volatility ratio in ITH estimation. See Figure 4 for a time-series of the volatility ratio.) This stratification is consistent with the bootstrap procedure employed in Rigobon (2003). We then included all the observations for the selected days in the bootstrap sample. This was done to preserve the lag structure of the variables. We ran the ITH estimation procedure on 500 such bootstrap samples and calculated bootstrap standard error for each coefficient of interest as

$$S.E._{boot} = \sqrt{\frac{\sum_{n=1}^{N} (b_{boot,n} - b_{OS})^2}{N}}.$$

Here N is the number of bootstrap iterations (500 in this case), $b_{boot,n}$ is the coefficient of interest estimated in the nth bootstrap iteration, and b_{OS} is the corresponding coefficient estimated from the original sample. We found that bootstrap standard errors were very similar to the asymptotic standard errors. All the conclusions based on the asymptotic standard errors were preserved using bootstrap standard errors.

References

Acharya, Viral V., and Lasse Pedersen, 2004, Liquidity risk and asset pricing, *Journal of Financial Economics* 77, 375–410.

Albuquerque, Rui, Eva De Francisco, and Luis B. Marques, 2008, Marketwide private information in stocks: Forecasting currency returns, *Journal of Finance* pp. 2297–2337.

Amemiya, Takeshi, 1985, *Advanced Econometrics* (Harvard University Press: Cambridge).

Amihud, Yacov, 2002, Illiquidity and stock returns: Cross-section and time-series effects, *Journal of Financial Markets* 5, 31–56.

Berkman, Henk, Tim Brailsford, and Alex Frino, 2005, A note on execution costs for stock index futures: Information versus liquidity effects, *Journal of Banking and Finance* 29, 565–577.

Brandt, Michael W., and Kenneth A. Kavajecz, 2004, Price discovery in the U.S. treasury market: The impact of orderflow and liquidity on the yield curve, *Journal of Finance* 59, 2623–2654.

Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam, 2000, Commonality in liquidity, *Journal of Financial Economics* 56, 3–28.

Coppejans, Mark T., Ian H. Domowitz, and Ananth Madhavan, 2004, Resiliency in an automated auction, BGI working paper.

Eisfeldt, Andrea, 2004, Endogenous liquidity in asset markets, *Journal of Finance* 59, 1–30.

Evans, Martin D. D., and Richard K. Lyons, 2002, Orderflow and exchange rate dynamics, *Journal of Political Economy* 110, 170–180.

——— , 2004, A new micro model of exchange rate dynamics, Unpublished working paper. University of California-Berkeley.

Fleming, Jeff, Barbara Ostdiek, and Robert E. Whalley, 1996, Trading costs and the relative rates of price discovery in stock, futures, and options markets, *Journal of Futures Markets* 16, 353–387.

Gomber, Peter, Uwe Schweickert, and Erik Theissen, 2004, Zooming in on liquidity, University of Bonn working paper.

Gorton, Gary B., and George G. Pennacchi, 1993, Security baskets and index-linked securities, *Journal of Business* 66, 1–27.

Goyenko, Ruslan Y., Craig W. Holden, and Charles A. Trzcinka, 2008, Do liquidity measures measure liquidity?, *Journal of Financial Economics* forthcoming.

Hasbrouck, Joel, 2003, Intraday price formation in U.S. equity index markets, *Journal of Finance* 58, 2375–2399.

——— , 2006, Trading costs and returns for U.S. equities: Estimating effective costs from daily data, NYU-Stern working paper.

——— , and Duane J. Seppi, 2001, Common factors in prices, order flows, and liquidity, *Journal of Financial Economics* 59, 383–411.

Hegde, Shantaram P., and John B. McDermott, 2004, The market liquidty of DIAMONDs, Qs, and their underlying stocks, *Journal of Banking and Finance* 28, 1043–1067.

Irvine, Paul, George Benston, and Eugene Kandel, 2000, Liquidity beyond the inside spread: Measuring and using information in the limit order book, Emory University working paper.

Johnson, Timothy C., 2007, Volume, liquidity, and liquidity risk, *Journal of Financial Economics* 87, 388–417.

Kawaller, Ira G., Paul D. Koch, and Timothy W. Koch, 1987, The temporal price relationship between S&P 500 futures and the S&P 500 index, *Journal of Finance* 42, 1309–1329.

Lee, Charles M. C., and Mark J. Ready, 1991, Inferring trade direction from intraday data, *Journal of Finance* 46, 733–746.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708–256.

Obizhaeva, Anna, 2008, Price impact and spread: Bias-free estimates from portfolio transition data, University of Maryland working paper.

Pagano, Marco, 1989, Trading volume and asset liquidity, *Quarterly Journal of Economics* 104, 255–274.

Parlour, Christine A., and Duane J. Seppi, 2008, Limit order markets: A survey, in Arnoud Boot, and Anjan V. Thakor, ed.: *Handbook of Financial Intermediation and Banking* (Elsevier: Amsterdam).

Pástor, L̂uboŝ, and Robert Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642–685.

Rigobon, Roberto, 2003, Identification through heteroskedasticity, *Review of Economics and Statistics* 85, 777–792.

Subrahmanyam, Avanidhar, 1991, A theory of trading in stock index futures, *Review of Financial Studies* 4, 17–51.

Panel A: S&P 500 Futures Price

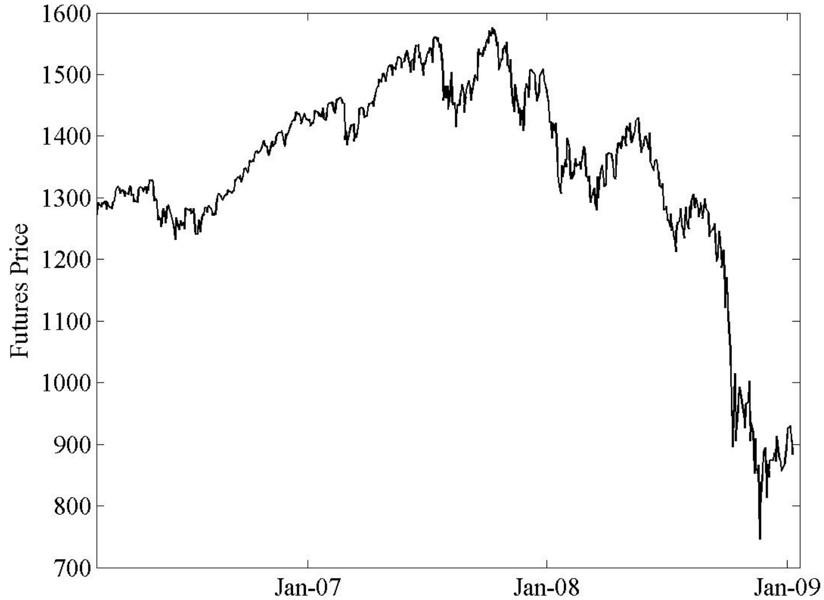

Panel B: S&P 500 E-mini Futures Trading Volume

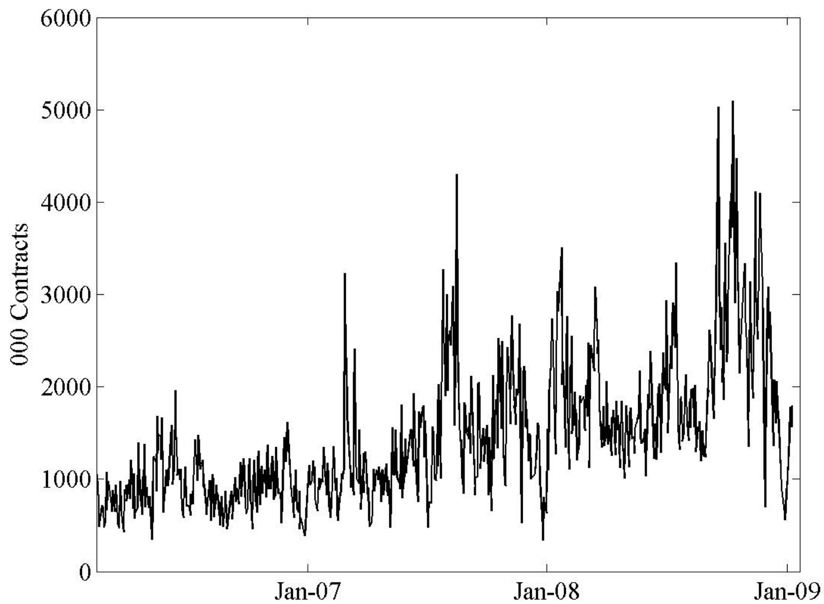

Figure 1: **Time series of price and trading volume of the S&P 500 e-mini futures.** Data are for the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. Panel A shows daily closing price in S&P 500 index points. Panel B shows daily total trading volume in thousands of e-mini contracts. Each e-mini contract is worth USD 50 times the index level.

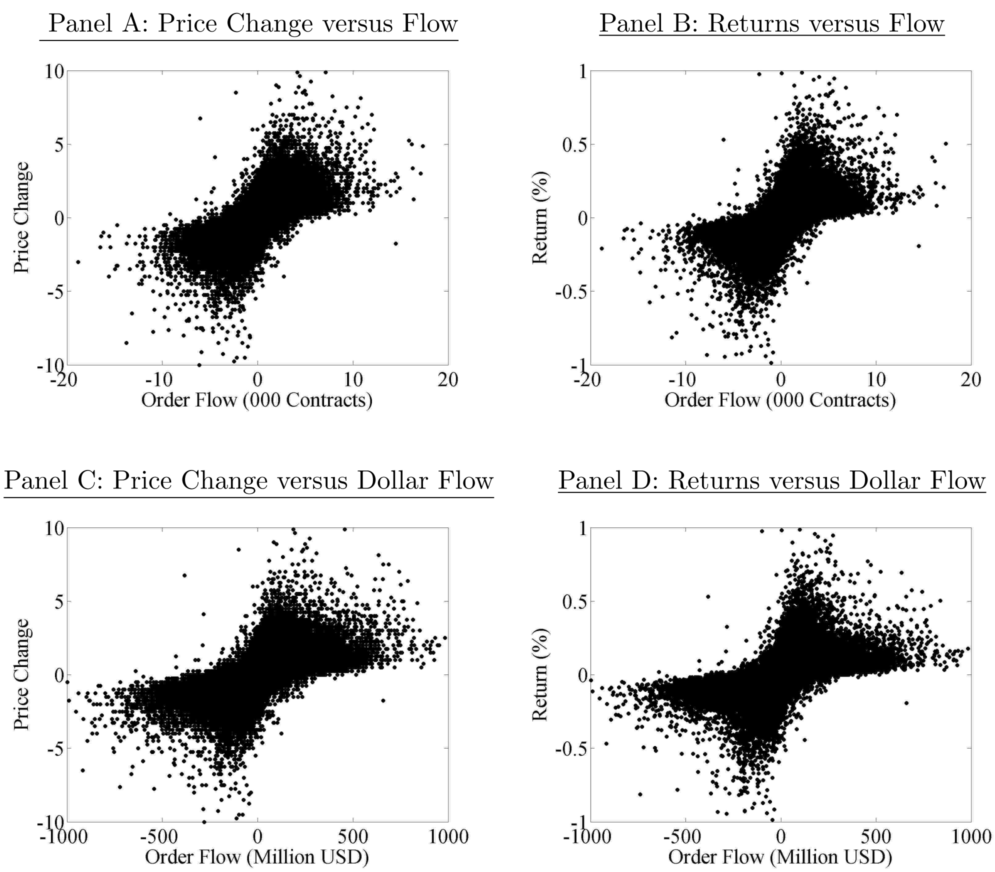

Figure 2: **Scatter plots of price change against order flow.** Price change, shown on the vertical axis, and order flow, on the horizontal axis, are from the S&P 500 e-mini futures market and are calculated over 1-minute intervals from the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. Order flow denotes the quantity traded at a price higher than the mid-price minus the quantity traded at a price lower than the mid-price. Panel A measures price change in index points and order flow in thousands of e-mini contracts. Each e-mini contract is worth USD 50 times the index level. Panel B plots percent returns against order flow in thousands of contracts. Panel C shows price change in index points against order flow in millions of dollars. Panel D has percent returns against dollar order flow.

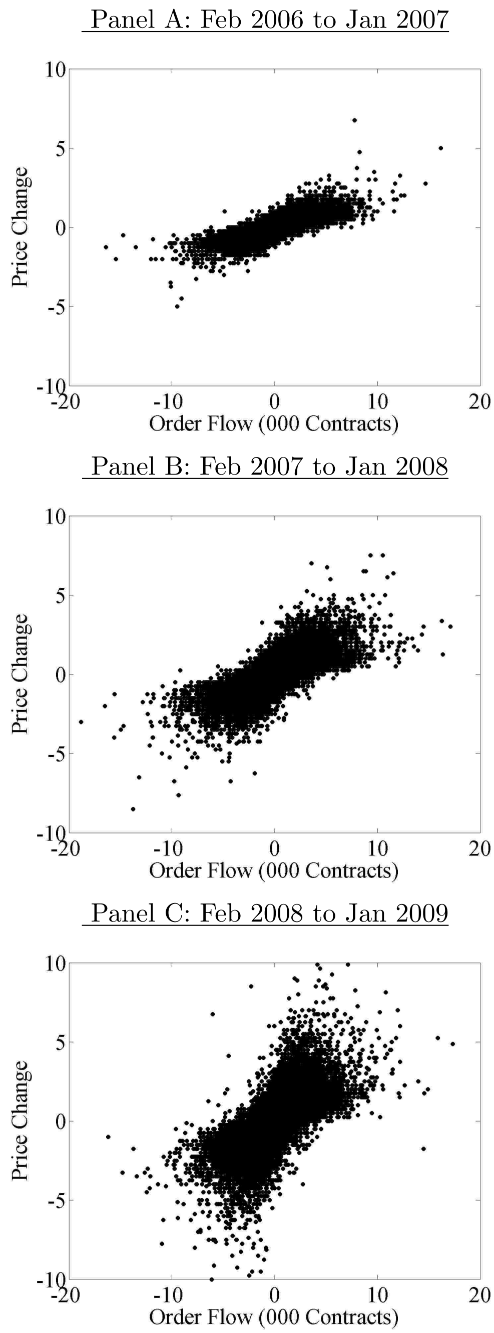

Figure 3: **Scatter plots of price change against order flow in different subsamples**. Price change, shown on the vertical axis, and order flow, on the horizontal axis, are from the S&P 500 e-mini futures market and are calculated over 1-minute intervals from the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. Order flow denotes the quantity traded at a price higher than the mid-price minus the quantity traded at a price lower than the mid-price. Price change is in index points and order flow is in thousands of e-mini contracts. Each e-mini contract is worth USD 50 times the index level.

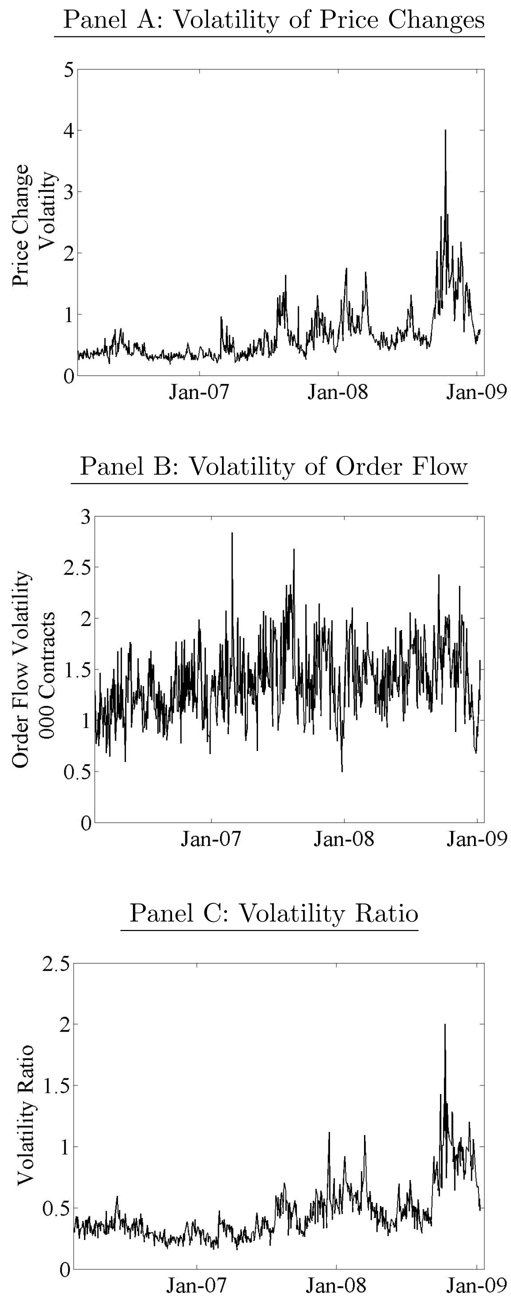

Figure 4: **Time series of volatility of price change and volatility order flow.** Price change and order flow are from the S&P 500 e-mini futures market from the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. Order flow denotes the quantity traded at price higher than the mid-price minus the quantity traded at a price lower than the mid-price. Panel A shows daily volatility of 1-minute price changes in index points. Panel B plots daily volatility of 1-minute order flow in thousands of e-mini contracts. Each e-mini contract is worth USD 50 times the index level. Panel C shows the ratio of the two volatilities.

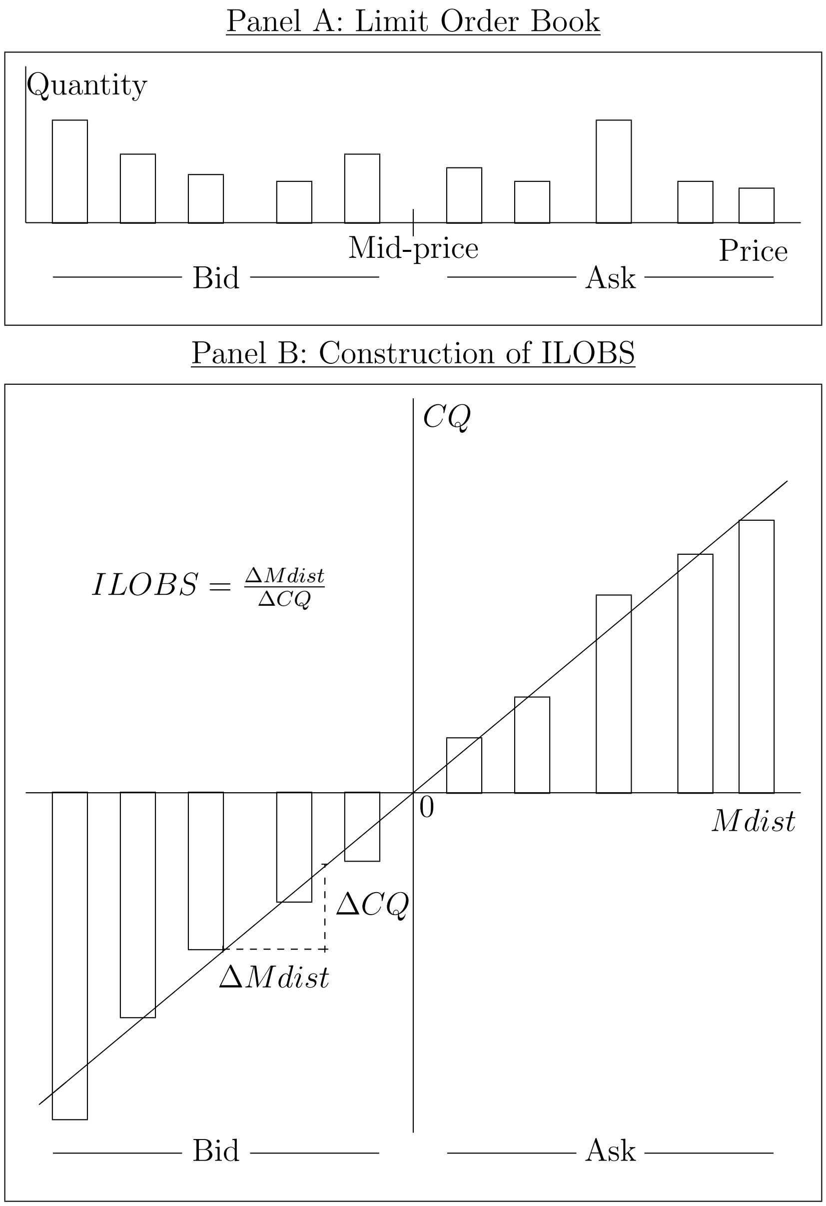

Figure 5: **Construction of the inverse slope of the limit order book (ILOBS)**. Panel A depicts the limit order book at a point in time. The horizontal axis shows price and the vertical axis shows quantity. Each bar represents the total limit order quantity at a particular price. Panel B shows construction of ILOBS associated with the limit order book in Panel A. The horizontal axis shows $Mdist$, the difference between a limit order price and the mid-price. The vertical axis shows CQ, the cumulative limit order quantity at a given limit order price. Bid side quantities are treated as negative values. Change, along the fitted line, in CQ is termed as ΔCQ and in $Mdist$ is termed as $\Delta Mdist$.

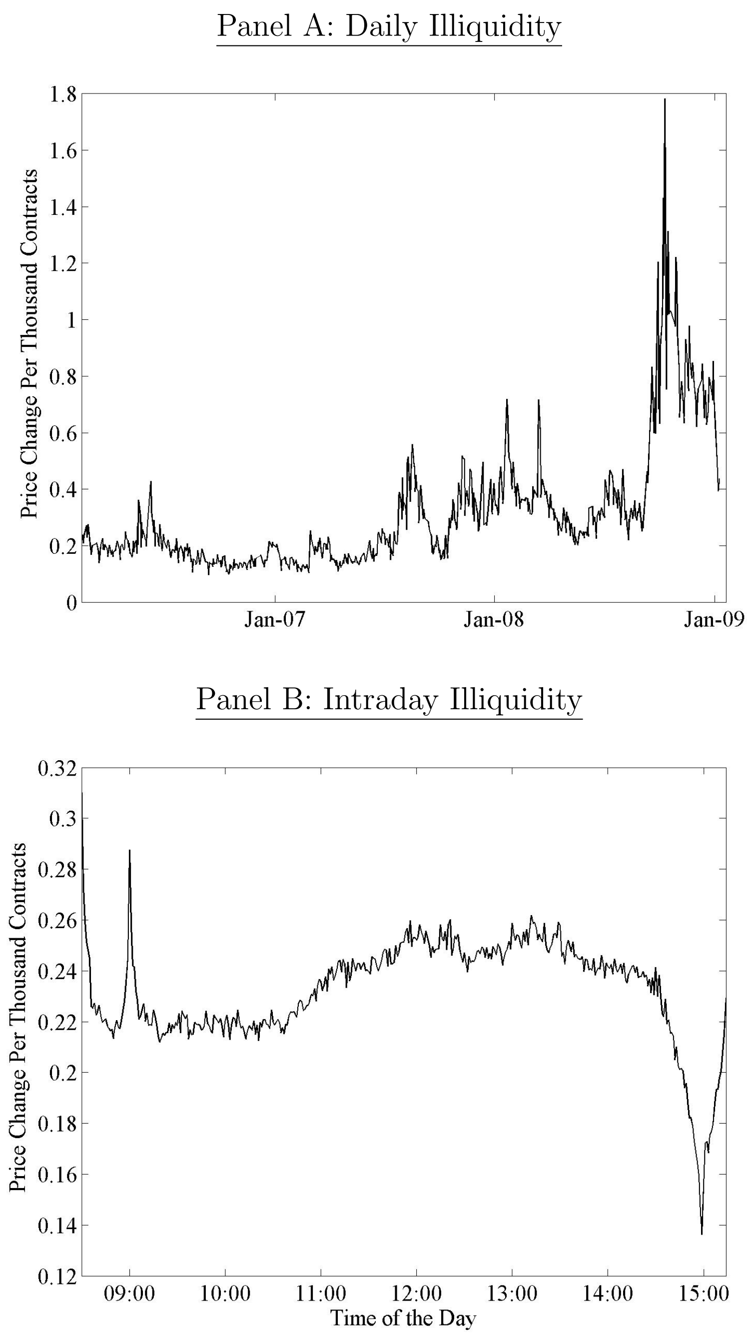

Figure 6: **Time series of illiquidity.** Illiquidity is given by the inverse slope of the limit order book (ILOBS) from the S&P 500 e-mini futures market and is measured in index points per thousand e-mini contracts. Each e-mini contract is worth USD 50 times the index level. Subsection B and Figure 5 give details of the construction of ILOBS. Data are sampled at 1-minute intervals from the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. Panel A shows the time series of daily median ILOBS. Panel B shows median ILOBS every minute of the trading day as per Central time.

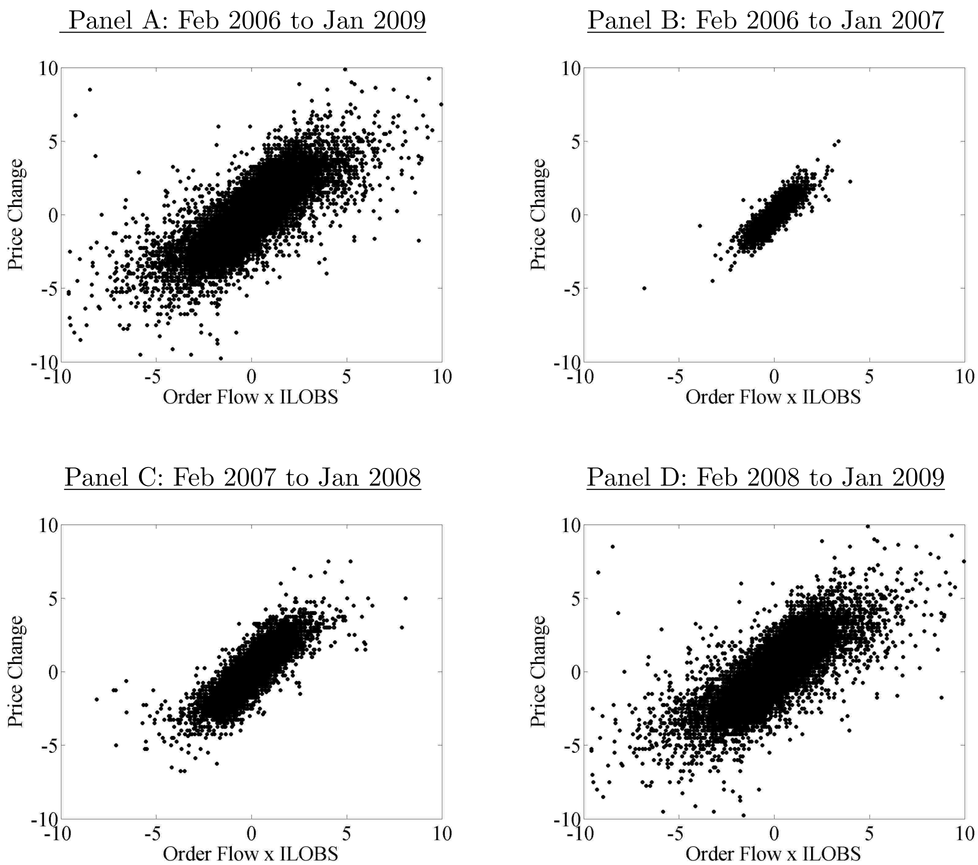

Figure 7: **Scatter plot of price change against order flow times the inverse slope of the limit order book (ILOBS).** Price change, order flow and ILOBS are from the S&P 500 e-mini futures market. Order flow denotes the quantity traded at price higher than the mid-price minus the quantity traded at a price lower than the mid-price. Price change and order flow are calculated over 1-minute intervals from the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. Price change is in index points and order flow is in thousands of e-mini contracts. Each e-mini contract is worth USD 50 times the index level. ILOBS, a proxy of illiquidity, is measured in index points per thousand e-mini contracts. Subsection B and Figure 5 give the details of the construction of ILOBS. Price change, in index points, is shown on the vertical axis, and order flow times ILOBS, again in index points, is on the horizontal axis. Panel A shows the scatter plots for the entire sample. Panels B to D present the plots for three subsamples.

Table I: **Trading in the Market Portfolio**

This table provides some descriptive statistics showing importance of the S&P 500 e-mini futures as a trading venue for the index. SPDRs refer to the S&P 500 Depository Receipts. The sample goes from 13 Feb 2006 to 31 Dec 2008. Panel A shows average daily volume (in billions of dollars) for different markets for trading the S&P 500 index. Panel B shows correlations of daily dollar volume in different venues for index trading. Panel C shows average daily illiquidity, measured by *Illiq*, on the different trading venues. *Illiq* for day t is calculated as $Illiq_t = \frac{|return_t|}{Dollar\ Volume_t}$. *Illiq* for a portfolio of stocks is either value-weighted or equal-weighted average of *Illiq* for individual stocks.

Panel A: Average Daily Volume (Billions of dollars)

S&P 500 e-mini futures	96
S&P 500 pit-traded futures	17
SPDRs	23
S&P 500 stocks	109

Panel B: Correlations of Daily Dollar Volume

	S&P 500 E-mini futures	SPDRs	S&P 500 stocks
S&P 500 E-mini futures	1.00	0.92	0.89
SPDRs	0.92	1.00	0.88
S&P 500 stocks	0.89	0.88	1.00

Panel C: Illiq (Basis points per billion dollars)

S&P 500 E-mini futures	1
SPDRs	4
S&P 500 stocks - Value-weighted	1418
NYSE stocks - Value-weighted	6271
NYSE stocks - Equal-weighted	36326

Table II: **Unconditional Price Impact Estimation**

This table provides results for unconditional price impact of over flow using the S&P 500 e-mini futures. Returns and order flow are measured over 1-minute intervals from the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. Order flow denotes the quantity traded at price higher than the mid-price minus the quantity traded at a price lower than the mid-price. b_r (effect of order flow on return) and b_f (effect of return on order flow) are the structural coefficients from simultaneous estimation using identification through heteroskedasticity (ITH). t-stats based on asymptotic standard errors are given in parentheses. R_r^2 and R_f^2 give one minus the ratio of sum of squares of residuals and sum of squares of the dependent variable for return and flow regressions respectively. Each row gives results for ITH estimation using different length of a regime in days. Obs gives total number of observations and the last column gives number of heteroskedasticity regimes used in the estimation. The first row provides OLS estimates. Panel A measures return as price change (ΔP) in index points and order flow in thousands of e-mini contracts (Q). In Panel B returns are measured either as price change (ΔP) in index points or proportional returns (r) and order flow is measured in terms of thousands of contracts (Q) or dollars (PQ). In each case, both the endogenous variables are scaled by their standard deviation to make the results comparable across specifications.

Panel A: Simultaneous Response Coefficients

Estimation	b_r	R_r^2	b_f	R_f^2	Obs	Regimes
OLS	0.38	0.51	-	-	254830	-
	(511.98)					
ITH 1 day	0.21	0.40	0.96	0.48	254830	684
	(73.79)		(45.65)			
ITH 7 days	0.22	0.41	0.90	0.47	254830	153
	(38.59)		(22.76)			
ITH 30 days	0.21	0.41	0.93	0.48	254830	36
	(18.55)		(12.43)			
ITH 91 days	0.21	0.41	0.89	0.47	254830	13
	(13.55)		(8.96)			

Panel B: Different Measures of Orderflow and Returns

	$\Delta P, Q$		r, Q		$\Delta P, PQ$		r, PQ	
Estimation	b_r	b_f	b_r	b_f	b_r	b_f	b_r	b_f
OLS	0.73	-	0.65	-	0.69	-	0.60	-
ITH 1 day	0.40	0.49	0.38	0.43	0.46	0.36	0.42	0.26
ITH 7 days	0.42	0.46	0.40	0.39	0.48	0.32	0.43	0.24
ITH 30 days	0.41	0.48	0.40	0.39	0.48	0.32	0.43	0.25
ITH 91 days	0.41	0.46	0.41	0.34	0.49	0.29	0.44	0.22

Table III: **Impulse Response**

This table provides estimates of instantaneous (I_0) and long run (I_∞) price impact of one-standard-deviation shock to the innovation in order flow in the S&P 500 e-mini futures using identification through heteroskedasticity (ITH). Returns and order flow are measured over 1-minute intervals from the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. Order flow denotes the quantity traded at price higher than the mid-price minus the quantity traded at a price lower than the mid-price (in thousands of contracts). Returns are measured either as price change (ΔP) in index points or proportional returns (r) and order flow is measured in terms of thousands of e-mini contracts (Q) or dollars (PQ). In each case, both the endogenous variables are scaled by their standard deviation to make the results comparable across specifications. Each row gives results for ITH estimation using different length of a regime in days.

	$\Delta P, Q$		r, Q		$\Delta P, PQ$		r, PQ	
Estimation	I_0	I_∞	I_0	I_∞	I_0	I_∞	I_0	I_∞
ITH 1 day	0.46	0.41	0.43	0.38	0.56	0.49	0.49	0.43
ITH 7 days	0.49	0.44	0.46	0.41	0.59	0.53	0.50	0.45
ITH 30 days	0.48	0.43	0.46	0.42	0.59	0.54	0.50	0.46
ITH 91 days	0.48	0.43	0.47	0.43	0.61	0.55	0.52	0.46

Table IV: **Descriptive Statistics of ILOBS**

This table presents descriptive statistics of the inverse slope of the limit order book (ILOBS) for the S&P 500 e-mini futures. ILOBS, a proxy of illiquidity, is measured in S&P 500 index points per thousand e-mini contracts. Subsection B and Figure 5 give the details of the construction of ILOBS. Data are sampled at 1-minute intervals from the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. Panel A presents the descriptive statistics for the whole sample and for subsamples. Panel B presents correlation between different versions of the inverse slope of the limit order book. $ILOBS_{wt1}$ is calculated by weighting every observation in the limit order book by the inverse of absolute value of the distance of the limit order price from the mid-price. $ILOBS_{wt2}$ is calculated by weighting every observation in the limit order book by the inverse of the square of such distance.

Panel A: Descriptive Statistics

Sample	Obs	Mean	Std Dev	Median	5th Percentile	95th Percentile
Feb '06-Jan '09	276030	0.31	0.42	0.23	0.11	0.79
Feb '06-Jan '07	53542	0.18	0.07	0.17	0.10	0.29
Feb '07-Jan '08	82008	0.26	0.16	0.22	0.11	0.53
Feb '08-Jan '09	140480	0.50	0.68	0.37	0.20	1.14

Panel B: Correlation between different versions of ILOBS

	$ILOBS$ and $ILOBS_{wt1}$	$ILOBS$ and $ILOBS_{wt2}$	$ILOBS_{wt1}$ and $ILOBS_{wt2}$
	0.996	0.980	0.993

Table V: **Conditional Price Impact Estimation**

This table presents price impact of order flow conditional on the inverse slope of the limit order book (ILOBS) estimated using a simultaneous system of return and flow and identified through heteroskedasticity (ITH). Flow denotes the quantity traded at a price higher than the mid-price minus the quantity traded at a price lower than the mid-price. Subsection B and Figure 5 give the details of the construction of ILOBS. Returns, order flow and ILOBS from the S&P 500 e-mini futures are calculated over 1-minute intervals from the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. The table presents a linear specification where $return = (b_0 + b_1 ILOBS)Flow$ and $Flow = b_f return$. t-stats based on asymptotic standard errors are in parentheses. R^2_r and R^2_f give one minus the ratio of sum of squares of residuals and sum of squares of the dependent variable for return and flow regressions respectively. Each row gives results for ITH estimation using different length of a regime in days. The first row in each panel provides OLS estimates. In panel A, returns are price changes (ΔP) in index points, order flow is in thousands of e-mini contracts (Q) and ILOBS is in index points per thousand contracts. In panel B, returns are proportional returns (r) in basis points, order flow is in billions of dollars (PQ) and ILOBS is in basis points per billion dollars.

Panel A: Price Change and Number of Contracts ($\Delta P, Q$)

Estimation	b_0	b_1	R^2_r	b_f	R^2_f
OLS	0.28	0.30	0.58	-	-
	(327.77)	(200.46)			
ITH 1 day	0.11	0.69	0.46	0.58	0.37
	(32.48)	(38.96)		(21.33)	
ITH 7 days	0.10	0.71	0.44	0.50	0.34
	(18.13)	(25.41)		(9.88)	
ITH 30 days	0.10	0.76	0.42	0.42	0.30
	(11.14)	(14.77)		(2.89)	
ITH 91 days	0.09	0.87	0.33	0.10	0.11
	(8.54)	(16.19)		(0.66)	

Panel B: Return and Dollar Volume (r, PQ)

Estimation	b_0	b_1	R^2_r	b_f	R^2_f
OLS	11.81	0.48	0.53	-	-
	(218.52)	(315.49)			
ITH 1 day	4.08	0.97	0.33	0.002	0.12
	(27.83)	(70.93)		(6.564)	
ITH 7 days	4.07	0.97	0.33	0.002	0.10
	(16.04)	(40.55)		(2.262)	
ITH 30 days	3.99	0.98	0.33	0.000	0.05
	(9.72)	(27.14)		(0.328)	
ITH 91 days	3.54	1.02	0.29	-0.002	-0.03
	(9.92)	(28.59)		(-1.917)	

Table VI: **Conditional Price Impact: Effect of Outliers**

This table presents price impact of order flow conditional on the inverse slope of the limit order book (ILOBS) estimated as a part of simultaneous system of return and flow and identified through heteroskedasticity controlling for outliers. Flow denotes the quantity traded at price higher than the mid-price minus the quantity traded at a price lower than the mid-price. Subsection B and Figure 5 give the details of the construction of ILOBS. Returns, order flow and ILOBS from the S&P 500 e-mini futures are calculated over 1-minute intervals from the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. Returns are price changes (ΔP) in index points, order flow is in thousands of e-mini contracts (Q) and ILOBS is in index points per thousand e-mini contracts. b_f gives effect of return on flow. t-stats based on asymptotic standard errors are given in parentheses. R_r^2 and R_f^2 give one minus the ratio of sum of squares of residuals and sum of squares of the dependent variable for return and flow regressions respectively. Each row gives results for ITH estimation using different length of a regime in days. The first row in each panel provides OLS estimates. Panels A and B present a winsorized specification where $return = (b_0 + b_1 min(ILOBS, c))Flow$. c, the cut-off for $ILOBS$, is set at 99th or 98th percentile of $ILOBS$. Panel C presents a piecewise specification where $return = (b_0 + b_1 ILOBS)(I_{abs(ILOBS \cdot Flow) \leq C})Flow + (b_2 + b_3 ILOBS)(I_{abs(ILOBS \cdot Flow) > C})Flow$. C, the cut-off for $ILOBS \cdot Flow$ is set at 99th percentile of the interaction.

Panel A: Winsorized Specification (99th Percentile)

Estimation	b_0	b_1	b_2	b_3	R_r^2	b_f	R_f^2
OLS	0.11	0.84	-	-	0.66	-	-
	(115.92)	(338.85)					
ITH 1 day	0.10	0.71	-	-	0.65	0.54	0.35
	(31.70)	(37.83)				(18.44)	
ITH 7 days	0.10	0.75	-	-	0.65	0.45	0.31
	(18.27)	(23.91)				(7.72)	
ITH 30 days	0.09	0.83	-	-	0.66	0.32	0.24
	(10.50)	(12.58)				(1.98)	
ITH 91 days	0.08	0.88	-	-	0.66	0.17	0.16
	(8.48)	(13.25)				(0.79)	

Panel B: Winsorized Specification (98th Percentile)

Estimation	b_0	b_1	b_2	b_3	R_r^2	b_f	R_f^2
OLS	0.10	0.91	-	-	0.66	-	-
	(93.94)	(340.98)					
ITH 1 day	0.10	0.74	-	-	0.65	0.52	0.34
	(31.01)	(37.33)				(17.18)	
ITH 7 days	0.10	0.80	-	-	0.65	0.41	0.29
	(17.53)	(23.11)				(6.51)	
ITH 30 days	0.09	0.84	-	-	0.66	0.32	0.24
	(9.79)	(12.09)				(2.15)	
ITH 91 days	0.08	0.94	-	-	0.66	0.11	0.12
	(8.22)	(13.24)				(0.52)	

Panel C: Piecewise Specification

Estimation	b_0	b_1	b_2	b_3	R_r^2	b_f	R_f^2
OLS	0.05	1.15	0.43	0.16	0.64	-	-
	(41.10)	(258.14)	(247.56)	(102.63)			
ITH 1 day	0.06	0.97	0.34	-0.14	0.51	0.63	0.39
	(15.70)	(39.71)	(16.31)	(-2.83)		(26.82)	
ITH 7 days	0.06	0.96	0.42	-0.29	0.44	0.62	0.39
	(9.14)	(21.70)	(8.60)	(-2.73)		(19.57)	
ITH 30 days	0.06	0.89	0.52	-0.49	0.29	0.65	0.40
	(4.10)	(7.62)	(4.35)	(-1.86)		(7.44)	
ITH 91 days	0.06	0.98	0.43	0.03	0.62	0.41	0.29
	(3.02)	(7.56)	(4.25)	(0.31)		(6.10)	

Table VII: **Price Impact For Different Time Intervals**

This table presents price impact of order flow conditional on the inverse slope of the limit order book (ILOBS) estimated as a part of simultaneous system of return and flow and identified through heteroskedasticity (ITH). Flow denotes the quantity traded at price higher than the mid-price minus the quantity traded at a price lower than the mid-price. Subsection B and Figure 5 give the details of the construction of ILOBS. Returns, order flow and ILOBS are from the S&P 500 e-mini futures the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. Returns are calculated as price change (ΔP) in index points, order flow is in thousands of e-mini contracts (Q) and ILOBS is in index points per thousand e-mini contracts. Length of heteroskedasticity regime is one day. b_f gives effect of return on flow. t-stats based on asymptotic standard errors are given in parentheses. R^2_r and R^2_f give one minus the ratio of sum of squares of residuals and sum of squares of the dependent variable for return and flow regressions respectively. A piecewise linear specification is estimated where $return = (b_0 + b_1 ILOBS)(I_{abs(ILOBS \cdot Flow) \leq C})Flow + (b_2 + b_3 ILOBS)(I_{abs(ILOBS \cdot Flow) > C})Flow$. C, the cut-off for $ILOBS \cdot Flow$ is set at 99th percentile of the interaction. Panel A presents results for different subsamples for returns and flow over 1-minute interval including 30 lags of dependent and independent variables as controls. Panel B presents the result for the entire sample by varying number of lags and time interval for returns and flow calculation.

Panel A: Subsample Analysis

Sample	b_0	b_1	b_2	b_3	R^2_r	b_f	R^2_f
Feb '06-Jan '09	0.06	0.97	0.34	-0.14	0.51	0.63	0.39
	(15.70)	(39.71)	(16.31)	(-2.83)		(26.82)	
Feb '06-Jan '07	0.07	0.64	0.58	-1.25	0.54	1.25	0.47
	(11.47)	(11.39)	(4.80)	(-2.76)		(26.23)	
Feb '07-Jan '08	0.06	0.93	0.27	0.05	0.67	0.78	0.45
	(8.42)	(18.24)	(16.73)	(22.35)		(15.56)	
Feb '08-Jan '09	0.04	0.97	0.34	-0.03	0.52	0.55	0.43
	(2.42)	(13.82)	(10.34)	(-0.56)		(21.71)	

Panel B: Different Time Intervals and Lags

Interval, Lags	b_0	b_1	b_2	b_3	R^2_r	b_f	R^2_f
1-Minute, 30	0.06	0.97	0.34	-0.14	0.51	0.63	0.39
	(15.70)	(39.71)	(16.31)	(-2.83)		(26.82)	
1-Minute, 15	0.06	0.98	0.35	-0.18	0.49	0.63	0.39
	(15.34)	(39.91)	(16.87)	(-3.52)		(27.23)	
5-Minute, 30	0.08	0.68	0.24	0.04	0.61	0.88	0.45
	(25.88)	(38.00)	(22.87)	(27.04)		(32.52)	
5-Minute, 6	0.07	0.70	0.23	0.04	0.63	0.85	0.43
	(25.76)	(40.30)	(25.29)	(10.78)		(32.66)	

Table VIII: **Effect of Bid Ask Spread**

This table presents results for bid ask spread from the S&P 500 e-mini futures. The spread is measured in index points. Panel A presents descriptive statistics for the bid ask spread. Panel B presents conditional price impact estimation after controlling for bid ask spread. Variables other than bid-ask spread are defined in the caption of Table VI. All the variables are measured over 1-minute intervals from the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. Length of heteroskedasticity regime is one day. b_f gives effect of return on flow. t-stats based on asymptotic standard errors are given in parentheses. R_r^2 and R_f^2 give one minus the ratio of sum of squares of residuals and sum of squares of the dependent variable for return and flow regressions respectively. A linear specification refers to $return = (b_0 + b_1 ILOBS + b_{BAS} BidAskSpread)Flow$. A winsorized specification refers to $return = (b_0 + b_1 min(ILOBS, c) + b_{BAS} BidAskSpread)Flow$ and c, the cut-off for $ILOBS$, is set at 99th or 98th percentile of $ILOBS$. A piecewise specification refers to $return = (b_0 + b_1 ILOBS)(I_{abs(ILOBS \cdot Flow) \leq C})Flow + (b_2 + b_3 ILOBS)(I_{abs(ILOBS \cdot Flow) > C})Flow + b_{BAS} Flow$. C, the cut-off for $ILOBS \cdot Flow$ is set at 99th percentile of the interaction. The linear and winsorized specifications control for 30 lags, whereas the piecewise specification controls for 15 lags, of the dependent and the independent variables.

Panel A: Descriptive Statistics for Bid Ask Spread

Sample	Obs	Mean	Std Dev	Median	5th Percentile	95th Percentile	Correlation with ILOBS
Feb '06-Jan '09	276030	0.26	0.08	0.25	0.25	0.25	0.17
Feb '06-Jan '07	96463	0.27	0.12	0.25	0.25	0.25	0.13
Feb '07-Jan '08	95070	0.26	0.07	0.25	0.25	0.25	0.11
Feb '08-Jan '09	84498	0.25	0.05	0.25	0.25	0.25	0.52

Panel B: Conditional Price Impact controlling for Bid Ask Spread

Specification	b_0	b_1	b_2	b_3	b_{BAS}	R_r^2	b_f	R_f^2
Linear	0.35	0.67	-	-	-0.00	0.46	0.55	0.35
	(9.37)	(36.28)			(-6.46)		(19.98)	
Winsorized	0.33	0.69	-	-	-0.00	0.64	0.52	0.34
	(9.17)	(34.79)			(-6.25)		(17.67)	
Piecewise	0.11	0.96	0.40	-0.18	-0.00	0.49	0.64	0.39
	(4.61)	(38.71)	(13.24)	(-3.45)	(-2.32)		(27.14)	

Table IX: **Correlation Matrix**

Correlations coefficients are shown for the primary variables employed in the illiquidity regressions. Data are sampled at 5-minute intervals from the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. Illiquidity is the log of the inverse slope of the limit order book (ILOBS). Subsection B and Figure 5 give details of the construction of ILOBS. Flow denotes the quantity traded at price higher than the mid-price minus the quantity traded at a price lower than the mid-price (in thousands of contracts). Volume is the log of the sum of those two quantities. Nonflow and flow variance denote the log of the sum of squared residuals from a specification of the simultaneous system estimated in Section 3. Correlations are estimated after removing time-of-day fixed effects.

	Illiquidity	Price Change	Order Flow	Volume	Nonflow Variance	Flow Variance
Illiquidity	-	-0.0033	-0.0098	0.4093	0.6536	0.0281
Price Change		-	0.7521	-0.0011	-0.0117	-0.0153
Order Flow			-	-0.0004	-0.0220	-0.0239
Volume				-	0.5148	0.2613
Nonflow Variance					-	0.2921
Flow variance						-

Table X: **Market Illiquidity Regressions**

The table shows coefficients from ordinary least squares regressions of market illiquidity, measured by the log of the limit order book slope, on its own lags and lags of six explanatory variables. The regression uses data sampled at 5-minute intervals from the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. The variables are as described in the caption to Table IX. All variables are standardized by their full-sample standard deviation. The dependent variable is truncated at its 99.95th percentile. All specifications use time-of-day fixed-effects. Newey and West (1987) t-statistics are shown in parentheses.

Panel A: 5-minute intervals

Variable:	Lag=0	1	2	3	4	5	6
Dependent	-	0.3388	0.1718	0.1233	0.1027	0.0874	0.0986
		(51.33)	(27.88)	(20.13)	(16.40)	(15.09)	(18.10)
Price Change	-0.0004	0.0000	-0.0019	0.0011	-0.0002	0.0015	-0.0077
	(0.112)	(0.013)	(0.546)	(0.319)	(0.066)	(0.402)	(2.492)
Order Flow	-0.0126	-0.0119	-0.0014	-0.0028	-0.0010	-0.0017	0.0051
	(4.046)	(4.048)	(0.483)	(0.989)	(0.356)	(0.580)	(1.990)
Volume	-0.1284	0.0196	0.0174	0.0197	0.0178	0.0239	0.0095
	(37.46)	(6.172)	(5.644)	(6.328)	(5.654)	(7.929)	(3.379)
Return Variance	0.0781	0.0246	0.0046	-0.0035	-0.0066	-0.0083	0.0002
	(23.28)	(7.973)	(1.501)	(1.125)	(2.130)	(2.703)	(0.058)
σ_{resid}	0.1131						
N	48008						

Panel B: Separate Components of Return Variance

Variable:	Lag=0	1	2	3	4	5	6
Dependent	-	0.3192	0.1649	0.1118	0.0915	0.0742	0.0896
		(51.12)	(28.70)	(19.54)	(15.42)	(13.38)	(17.39)
Price Change	0.0018	0.0018	-0.0007	-0.0016	-0.0031	0.0015	-0.0048
	(0.490)	(0.597)	(0.211)	(0.527)	(0.983)	(0.417)	(1.602)
Order Flow	-0.0111	-0.0099	-0.0008	-0.0004	-0.0003	-0.0027	0.0039
	(3.737)	(3.748)	(0.298)	(0.151)	(0.115)	(0.958)	(1.570)
Volume	-0.0642	0.0241	0.0153	0.0126	0.0126	0.0164	0.0109
	(15.75)	(9.296)	(6.061)	(4.928)	(4.996)	(6.661)	(4.629)
Nonflow Variance	0.0120	0.0126	0.0188	0.0245	0.0277	0.0511	0.0177
	(4.882)	(5.095)	(7.394)	(9.623)	(10.39)	(18.08)	(6.202)
Flow Variance	-0.0077	-0.0046	-0.0098	-0.0091	-0.0133	-0.0560	-0.0198
	(4.702)	(2.734)	(5.737)	(5.390)	(7.764)	(31.24)	(7.303)
σ_{resid}	0.1045						
N	47003						

Table XI: **Illiquidity Regressions for Different Time Intervals**

The table shows coefficients from regressions of market illiquidity, measured by the log of the limit order book slope, on its own lags and lags of six explanatory variables. The regression uses data from the hours of 8:30 to 3:14 Cental time over the period 13 Feb 2006 to 9 Jan 2009. The top panel uses observations every 15 minutes. The bottom panel uses hourly observations. The variables are as described in the caption to Table IX. The bottom panel includes averages of the variables over the prior day, week, and month as additional predictors. All variables are standardized by their full-sample standard deviation. The dependent variable is truncated at its 99.95th percentile. All specifications use time-of-day fixed-effects. Newey and West (1987) t-statistics are shown in parentheses.

Panel A: 15-minute intervals

Variable:	Lag=0	1	2	3	4	5	6
Dependent	-	0.3922	0.1447	0.1037	0.1043	0.0783	0.0758
		(26.28)	(13.49)	(9.768)	(9.676)	(7.913)	(8.508)
Price Change	0.0028	0.0073	0.0113	-0.0010	-0.0013	-0.0087	-0.0071
	(0.359)	(1.132)	(1.714)	(0.164)	(0.196)	(1.486)	(1.242)
Order Flow	-0.0152	-0.0131	-0.0090	-0.0005	-0.0051	0.0073	0.0060
	(2.382)	(2.366)	(1.608)	(0.009)	(0.936)	(1.369)	(1.210)
Volume	-0.0884	0.0348	0.0290	0.0354	0.0263	0.0076	0.0199
	(7.460)	(3.186)	(2.707)	(3.183)	(2.428)	(1.325)	(3.609)
Nonflow Variance	0.0559	0.1697	0.0789	-0.0111	-0.0366	-0.0370	-0.0451
	(7.140)	(13.36)	(8.110)	(1.248)	(4.527)	(4.073)	(5.489)
Flow Variance	-0.0234	-0.0661	-0.0442	0.0023	-0.0122	-0.0085	-0.0019
	(6.272)	(17.00)	(6.127)	(0.335)	(1.828)	(1.212)	(0.276)
σ_{resid}	0.1285						
N	12934						

Panel B: 1-hour interval

Variable:	Lag=0	1	2	3	1 day	1 week	1 month
Dependent	-	0.1885	0.0674	0.1578	0.0977	0.2445	0.0861
		(6.993)	(2.784)	(6.799)	(2.053)	(2.696)	(0.904)
Price Change	0.0260	-0.0143	0.0106	0.0012	0.0160	-0.0084	-0.0009
	(1.562)	(0.876)	(0.757)	(0.066)	(0.939)	(0.641)	(0.071)
Order Flow	-0.0391	-0.0014	-0.0015	0.0048	-0.0187	0.0189	-0.0026
	(2.281)	(0.093)	(0.115)	(0.277)	(1.275)	(1.333)	(0.241)
Volume	-0.0697	0.0121	-0.0150	-0.0174	0.0263	0.0239	0.1502
	(2.638)	(0.533)	(0.691)	(1.024)	(0.411)	(0.262)	(1.948)
Nonflow Variance	0.3919	0.0286	0.0158	-0.0310	-0.0905	-0.0358	-0.2281
	(11.97)	(0.725)	(0.516)	(1.024)	(1.546)	(0.387)	(2.580)
Flow Variance	-0.1369	0.0075	0.0167	0.0273	-0.0165	0.0223	-0.0217
	(10.48)	(0.449)	(1.109)	(2.052)	(0.529)	(0.603)	(0.678)
σ_{resid}	0.1327						
N	1950						