ELSEVIER

Journal of Financial Markets 4 (2001) 113–142

Journal of FINANCIAL MARKETS

www.elsevier.nl/locate/econbase

Predicting VNET: A model of the dynamics of market depth☆

Robert F. Engle[a], Joe Lange[b,*]

[a]*Department of Economics, University of California, San Diego, 9500 Gilman Drive, San Diego, CA 92093, USA*

[b]*Federal Reserve Board, Mail Stop 59, 2000 C Street NW, Washington, DC 20551, USA*

Abstract

The paper proposes a new intraday measure of market liquidity, VNET, which directly measures the depth of the market corresponding to a particular price deterioration. VNET is constructed from the excess volume of buys or sells associated with a price movement. As this measure varies over time, it can be forecast and explained. Using NYSE TORQ data, it is found that market depth varies with volume, transactions, and volatility. These movements are interpreted in terms of the varying proportion of informed traders in an asymmetric information model. When an unbalanced order flow is transacted in a surprisingly short time relative to that expected using the Engle and Russell (Econometrica 66 (1998) 1127) ACD model, the depth is further reduced providing an estimate of the value of patience. The analysis is repeated for 1997 TAQ data revealing that the parameters of the relationships changed only modestly, despite shifts in market volume, volatility, and minimum tick size. A dynamic market reaction curve is estimated with the new data.

JEL classification: C41; D82; G1

Keywords: Market microstructure; Asymmetric information; Stock market liquidity; Market depth; Market reaction curve; ACD model; TORQ data; TAQ data; NYSE

☆We are thankful for support from NSF grant SBR-9422575 and SBR-9730062 and from the NBER Asset Pricing Group. The views expressed in this paper are those of the authors alone and do not reflect the opinions of the Board of Governors or its staff. We are indebted to the referee and editor for thoughtful suggestions.

* Corresponding author. Tel.: + 1-202-452-2628.

E-mail address: joe.lange@frb.gov (J. Lange).

1386-4181/01/$ - see front matter
PII: S1386-4181(00)00019-7

1. Introduction

Over the past decade, equity market activity has increased dramatically in terms of both trading volume and price volatility. From one perspective, the ability of the stock market to handle an increasing number of daily transactions points to greater liquidity. However, the large price fluctuations that accompanied many of the high-volume days indicate that the market did not absorb the additional transactions without some degree of price impact. The net effect on the cost of trading is by no means obvious. Clearly neither volume nor volatility is a direct measure of liquidity, although they are closely connected. Beyond the bid–ask spread, few established measures of market liquidity are available and several are measurable only cross-sectionally.

To the extent that stock market liquidity is a time-varying process, it may be possible to forecast when the market will be most accommodative to incoming trade activity. A tool capable of distinguishing and predicting shifts in market depth would be particularly valuable to institutional traders conducting high-volume trades in a particular stock. In addition, risk managers seeking ways to measure liquidity risk should find the prediction of market reaction curves useful. Not only would this present the possibility of computing price deterioration from a known quantity of portfolio holdings, but it also would offer a menu of liquidation costs depending upon the unwind strategy chosen.

This paper introduces a new, intraday statistic for market depth. Quoted depth reflects the number of shares that can be bought or sold at a particular bid or offer price. The new statistic, VNET, measures the number of shares purchased minus the number of shares sold over a period when prices moved a certain increment, and it is therefore a measure of realized depth for a specific price deterioration. VNET is constructed in event-time, similar to Cho and Frees (1988), and can be measured repeatedly throughout the trading day to capture the short-run dynamics of market liquidity.

Motivated by the asymmetric information models in the market microstructure literature, a predictive model of intraday market depth is developed and estimated for 17 stocks from the NYSE's TORQ data set. As anticipated, VNET is observed to vary both over time and across stocks. The results show VNET to be a function of the magnitude and timing of current and lagged transaction flows. The transactions data used to derive our measure of market depth presumably were themselves optimized according to investor criteria. Thus, time variation in expected VNET must be a result of agents who chose not to completely smooth liquidity over time, such as information-based traders. The prediction of VNET based on a valid conditioning set can only be precisely associated with market depth under the assumption that the contemplated trades are treated by the market in the same way that trades were treated historically. That is, a well-known troubled hedge fund might find that the depth available to it would be less than that forecast because the trades would be identifiable. Conversely, an index fund might find greater depth than predicted.

In the next section, the liquidity concept is specified, then in Section 3 the market microstructure theory is discussed. Section 4 describes the TORQ data, and Section 5 presents the estimation results. Section 6 tests the robustness of these findings using a more current data sample, and Section 7 concludes.

2. Defining stock market liquidity

The concept of liquidity can have a variety of interpretations. Generally, it is the ability to transact at low cost. The divergence between buying and selling prices, referred to as the bid-ask spread, is the most commonly cited facet of liquidity. However, this measure only captures the tightness of the market price for low volume trades. Larger orders almost always face worse execution – the extent of which may be quite substantial for impatient, high-volume traders.

Fig. 1 below shows the hypothetical transaction price to be expected for various size buying or selling orders. This schedule is often called the market reaction curve and may depend on other features of the trades. The slope is sometimes called Kyle's lambda after Kyle (1985). Tightness is depicted by the degree of divergence between the buy and sell curves at the zero share line. Another dimension of liquidity is depth, defined as the maximum number of shares that can be traded at a given price. Looking at Fig. 1, the horizontal distance between the center axis and the market reaction curve, represents the volume that can be traded at a particular price. The posted quote depth, represented by the flat segments near the zero share line, does not provide a comprehensive picture of market depth. Whereas effective spreads are often tighter than the posted bid–ask spread, effective depth may differ from that quoted by the market maker, as well. Regardless, quoted depth can at best provide only a partial view of the market reaction curve.

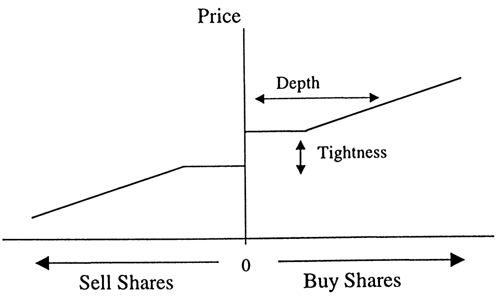

Fig. 1. Hypothetical market reaction curve.

The slope of the reaction function away from the current quotes is important for prospective large trader. While the market maker, termed a "specialist" on the NYSE, can independently influence the bid and ask prices, the shape of the market reaction curve away from these quotes is determined for the most part by the supply of standing limit orders. A steeper curve reflects a shortage of limit orders, implying a larger price impact for a given trade volume. This represents a lack of liquidity in the market. Of course, the true market reaction curve is not likely to be piecewise linear as illustrated in Fig. 1. More importantly, it is not a static schedule. Over time, limit orders are submitted, cancelled, and executed, altering the slope. This paper attempts to uncover the factors that influence the short-run behavior of the price response curve.

A natural approach to estimating the slope of the market reaction curve would be to measure the net trade volume and corresponding price change over a fixed interval of time. The price change per share of excess demand then estimates of the slope of the reaction function. There are several reasons not to follow this strategy.

Since excess demand can be positive or negative, the possibility of dividing by a number close to zero is high and outliers are to be expected. Furthermore, the discreteness of prices means that only a few possible values of the numerator can be anticipated and many zeroes are likely. Both of these problems are most severe if the measurement interval is short. However, the use of long intervals obviously reduces the ability of the statistic to capture short-run dynamics, particularly when the market is very active.

In this paper, we parse the data in a manner that avoids these problems. Market depth is most directly defined as the number of shares that can be bought or sold within a given price range. Therefore, the measurement interval for VNET should be dictated by the price level rather than calendar time. This general approach is used by Cho and Frees (1988) to construct a "temporal" measure of price volatility that eliminates the discreteness bias by focusing on the time takes prices to move a fixed amount. We expand upon this method of event-time analysis by recording the trade flows over price-determined intervals, or "price-durations". From this, market depth can be computed around a price event, often interpreted as an information event.

On some days there may be many price events while on other days there may be very few. The price-duration framework is able to accommodate active episodes by directly linking the frequency of measurement to the volatility of the market. For example, if two distinct news events occur within a short period of time causing the price to first rise by 50 cents then fall by 50 cents, a standard calendar-time approach would record zero price change over the period. However, the price-duration framework would record two observations of VNET, one after each large price movement, giving a more accurate picture of market liquidity over this period.

The number of price-durations that are recorded is determined by the size of the price threshold, which can be adjusted to achieve the desired resolution. The expected length of a price-duration is shown by Engle and Russell (1998) to be inversely proportional to the expected volatility, and in the context of VNET, can be interpreted as the pace at which excess demand flows into the market. The net directional volume, defined as the difference between the volume of buyer-initiated and seller-initiated trades within a price-duration, is the new proposed measure of market depth. Since each price-duration corresponds to a similar price change, the discreteness of prices does not feed through to the distribution of our statistic. By choosing depth as the feature to be measured, the dependent variable becomes the net volume per price change, not the reciprocal, and far better statistical properties are achieved.

3. Market microstructure

The validity of VNET as a measure of market depth hinges upon the assumption that it is the imbalance between buys and sells which causes prices to move. At first glance it may seem that public news presents a major challenge to this notion. If prices adjust purely in response to an announcement rather than underlying trading activity, the net directional transactions before a price move, VNET, will not accurately characterize the depth of the market over that price-duration. We argue that this is rarely, if ever, the case in the continuous-trading specialist system of the NYSE.

First, consider the ambiguity of news. Public announcements relating to a corporation, industry, or macroeconomic event never provide a precise indication of future price levels. Instead, analysts formulate a range of valuations and the market converges to a new price after a period of volatile trading. During this episode of price discovery, each trade is presumed to contain a high degree of information, and consequently, the price impact is large.

However, even if the market could unanimously quantify the impact of a public news event, the internal structure of the exchange mitigates exogenous price jumps. The specialist is explicitly charged with maintaining price continuity. In addition, unless all limit traders are constantly monitoring their orders so that they can cancel them after a news release, there will remain some stale limit orders with which to trade along the path toward the new price. So while it may seem extreme to propose that only trades move prices, in actuality it is quite rare to witness a large price adjustment without any intervening trades.

This is not to say that news does not indirectly influence prices. Information affects both order submissions and the responsiveness of the market to these orders. Asymmetric-information models of market microstructure, such as Easley and O'Hara (1992), suggest that the presence of informed traders in the market tends to amplify the price impact of a trade. These models assume that there is some probability of a *private* news event that is revealed to a subset of the population. If a transaction is known to be initiated by an informed agent, then the equilibrium price of the stock should shift according to the direction of the trade. Because of the anonymity of these "insiders", the price impact of a trade, and thus the depth of the market, is determined by the assumed probability of confronting an informed agent.

In the extreme case of a public news release, the fraction of informed traders (i.e. traders who know that the true valuation is different than the current quote) approaches one. The market becomes extremely responsive to trading activity, and the next trade will likely lead to a permanent price revision. Depending on the number of stale limit orders and the extent of efforts by the market maker to insure a continuous price path, there may be several trades before prices reach their new level. With prices moving on very low volume, realized VNET will be small, appropriately reflecting diminished market depth during this period. While the ability to forecast VNET may seem improbable in the above context, most price-durations do not stem from a public announcement, but instead tend to evolve over a longer time frame. Under these more standard circumstances, liquidity suppliers may use recent transaction patterns to develop a sense of the market's informational distribution.

The notion of heterogeneously informed agents and adverse selection is a well-documented aspect of the uncertainty facing liquidity suppliers. However, intraday variability in this informational asymmetry and any implications of such on time-varying liquidity is less thoroughly noted. If informed and liquidity traders have different trading tendencies, then the distribution of market information may be partially revealed in the nature of transaction activity at any given moment. In that the supply of liquidity is sensitive to informational assumptions, the realized depth of the market may be time-varying in a manner related to trading conditions.

Distinguishing informed from uninformed agents is fundamental to a liquidity provider's risk assessment. A number of studies have looked at this identification issue from a stationary point of view using both the bid–ask spread and the price impact of a trade. Easley and O'Hara (1987) and Hasbrouck (1988) find a positive correlation between trade size and price impact, with the implication that informed agents trade more heavily in order to profit from their fleeting informational advantage. McInish and Wood (1992) reveal that the bid–ask spread tends to widen following large volume orders.

The intensity of trade activity, defined by either the number of shares or the number of transactions per time, may also be a function of the asymmetry of information. The relationship between trading intensity and market depth depends on which type of traders (informed or uninformed) are predominantly responsible for episodes of above average market thickness (i.e. more transactions per time). Because informed agents are often constrained by the time sensitivity of their information, Foster and Viswanathan (1995) suggest that the pace of trading be positively correlated with the proportion of informed agents, as well as price volatility.

Most of the market microstructure literature abstracts from timing issues by constructing fixed trade interval models. Easley and O'Hara (1992) indirectly loosen this assumption by allowing traders the option of not trading during an interval. From this, a longer time between transactions indicates that market participants have abstained from trading. Since the portfolio adjustment needs of liquidity traders should be uniform throughout the day, informed agents likely initiate swings in transaction frequency. Again this supports the notion that high trade intensity is related to greater informational asymmetry, and low liquidity.

With the availability of transaction-by-transaction data for high frequency markets such as the NYSE, the time between trades has become another statistic for the empiricist. Engle and Russell (1998) model durations between trades for IBM, revealing significant autocorrelation or clumping of orders. If the factors which determine the timing of trades or price changes are related to the distribution of information amongst market traders, then forecasts of the time between market events may give added insight into the behavior of liquidity. The extent of the relationship between trading activity, market volatility, and the cost of trading will be explored in the empirical models below.

4. Data

The data for this study is taken from the TORQ (Trades, Orders, Reports, and Quotes) set, compiled by Joel Hasbrouck and the New York Stock Exchange. It contains tick-by-tick data for 144 stocks over the three-month period, November 1, 1990 through January 31, 1991. Trade time, trade size, and the prevailing quotes are extracted for the 17 stocks which traded most frequently on the first day of the sample, November 1, 1990.[1] A minimum level of trading activity is necessary in order to isolate price events within a single day. This abstraction from extremely inactive stocks should not be completely ignored, but the econometric techniques used in this analysis of intraday liquidity fluctuations apply most readily to active investment assets.

During these months, trading was abnormally slow on two dates, November 23rd (the Friday after Thanksgiving) and December 27th. Because VNET is theoretically grounded in a continuous trading environment, these two dates are dropped from the analysis leaving 61 days of data. While it may be interesting in future work to investigate these low-activity days, at present we focus on the

[1] Quotes from regional exchanges are excluded since they often differ from New York. The use of 17 stocks was purely arbitrary.

normal liquidity characteristics of the market. Similarly, overnight episodes are ignored in this purely *intra*-day study.

In determining the prevailing quotes for a given transaction, we implement the 'five second' rule suggested by Lee and Ready (1991). On the NYSE floor, new quotes can be posted more quickly than transactions can be recorded, meaning a quote revision may be time stamped earlier than the instigating trade. Matching transactions with quotes that are at least 5 s old mitigates the concern over mis-sequenced data records.

Along with the prevailing quote, each trade is given a marker according to the initiating party (buyer or seller). Again following Lee and Ready, a modified 'midpoint' rule is used to infer this unrecorded information. If the transaction price is closer to the ask than the bid quote, then it is a buy, otherwise it is labeled a sell. However, if the transaction occurred precisely at the midpoint between the bid and ask, then the 'tick' rule applies. Under this method, an up tick, meaning the current transaction price is greater than the previous price, implies that a buyer must have initiated the trade. Likewise, down ticks indicate sells. Lee and Ready found this process for distinguishing buys from sells to be the most accurate for a variety of simulated scenarios.

From here, the data for each stock are filtered in order to create a consistent set of observations and to isolate the intraday price fluctuations. To account for irregular trading patterns and procedures around the start of each day, the first five minutes of trading are dropped. Although the opening of the session can be both interesting and important, the rate of informational flows and price discovery may be fundamentally different from the rest of the day. This paper hopes to isolate the impact of trading activity on market depth, independent of time-of-day effects. The close can also present problems. The TORQ data set includes a number of transactions time-stamped after the 4 : 00 p.m. bell. While the true timing of these trades may be somewhat unclear, in practicality this is not an issue because none of these post-close trades happen to trigger a price-duration. The filtering procedure used to define a price-duration (described in detail later) ignores overnight activity, meaning that the trades following the last price-duration of a day are effectively excluded from the analysis.

In measuring price movements we use the change in the midpoint of the specialist's quotes. Not only does the mid-quote price provide a more accurate indication of the true market value of the asset, it does not encounter the problem of bid–ask bounce, although discreteness still plays a role. Transaction prices are also difficult to interpret because they often depend upon the size of the trade, even if the equilibrium valuation remains constant.

The models analyzed in this paper rely on a construct called a price-duration. Unlike typical trade-to-trade durations, price-based durations are defined as the time elapsed between significant price movements. Although this aggregation of trades over stable price sequences hides some of the information contained in the individual transaction records, much of the noise stemming from price

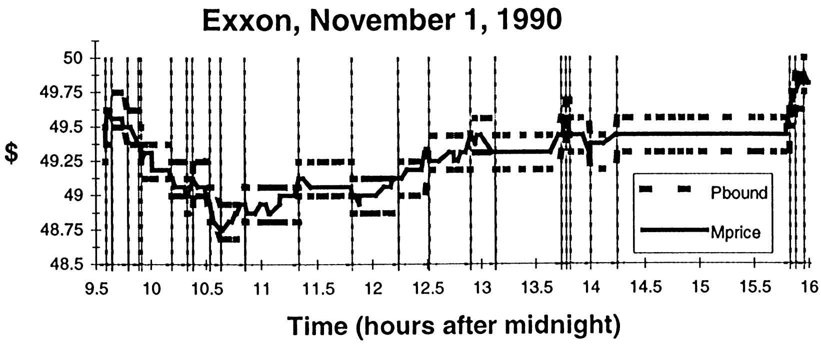

Fig. 2. NYSE quote patterns for Exxon on November 1, 1990. Mprice is the midpoint between the bid and ask, Pbound is our constructed price-duration barrier. The dashed vertical lines mark the end of a price-duration.

discreteness is avoided as well, allowing for a more realistic view of the equilibrium price behavior of the market. To insure we are isolating real price events, and not simply stray data entries, at least two consecutive data points outside the preset threshold are required to signal the end of a duration.

Fig. 2 displays a one-day sample of the time paths of the quote midpoint (Mprice) and the constructed price barriers (Pbound) used to define price-durations. The stock-specific threshold magnitudes are designed to be wider than a random noise jump, yet narrower than a true permanent price adjustment. In this way, the price-duration methodology reaps the benefits of aggregation while maintaining the flexibility of an event-time analysis.

Obviously, distinguishing noise from information is fairly arbitrary. The width of the pre-defined price threshold can be calibrated to suit the particular needs of the analyst. For this study we pick thresholds yielding roughly ten informational events per day. With this in mind, the price level and volatility of each stock determine the absolute price change necessary to achieve an average of ten price-durations per day – for the 17 stocks this ranged from 1/16th to 1/4th of a dollar (see Table 1 below). Despite our aim to equalize the average number of identified price events across stocks, the daily frequency ranged from as few as 5 for California Federal Bank (CAL) to 15 for IBM due to the minimum 1/8th tick size.[2]

The number of price-durations identified over the 61 trading days ranged from 321 for California Federal Bank to 945 for IBM, with corresponding

[2] The overnight period is excluded so no price-durations range across days. It should be noted that our exclusion of the first 5-min of trading each morning will impact the daily sequence of recorded durations.

Table 1
Price-duration statistics and underlying quote volatility (Nov. 1990–Jan. 1991).

Stock	Durations per day	Nominal price threshold ($)	Average midquote price ($)	Percentage price threshold (%)	Annualized half-hour volatility (%)
Boeing (BA)	10	0.1875	45.61	0.41%	33.1
Cal Fed Bank (CAL)	7	0.0625	3.29	1.90%	113.3
Colgate-Palmolive (CL)	8	0.1875	70.56	0.27%	21.3
CPC International (CPC)	15	0.1250	78.73	0.16%	20.6
(DI)	7	0.1250	20.25	0.62%	36.3
FedEX (FDX)	9	0.1250	34.04	0.37%	37.9
Fannie Mae (FNM)	8	0.1875	33.72	0.56%	39.0
FPL Group (FPL)	12	0.0625	28.44	0.22%	17.6
General Electric (GE)	12	0.1875	55.80	0.34%	26.0
Glaxo Well (GLX)	9	0.1250	32.48	0.38%	27.7
Hanson PLC (HAN)	10	0.0625	18.42	0.34%	32.2
IBM (IBM)	15	0.2500	113.30	0.22%	21.1
Philip Morris (MO)	7	0.1875	50.31	0.37%	21.6
Potomac Electric (POM)	8	0.0625	20.15	0.31%	19.6
Schlumberger (SLB)	13	0.1875	55.31	0.34%	28.6
AT&T (T)	9	0.1250	31.28	0.40%	27.4
Exxon (XON)	10	0.1250	50.57	0.25%	14.9

average price-duration times of 2,354 and 1,333 s, respectively. The average volume of trading activity within a price-duration ranged from 8,566 shares in 6 transactions for CPC International to 154,091 shares in 63 transactions for Philip Morris. Of course, these statistics are sensitive to the pre-selected distance that quotes must move to trigger a price-duration.

For each price-duration, a variety of summary measures are compiled. The number of trades, the total volume traded, the actual amount prices moved, the elapsed clock time (PTIME), and the bid–ask spread are the fundamental statistics; average trade size and the average time between trades, as well as interaction effects, are imputed.

The central statistic in this study is VNET, which captures the net directional (buy or sell) volume over price-duration. That is, the imbalance between the number of shares bought and the number of shares sold within a duration depicts the realized depth of the market. This statistic reveals the amount of one-sided volume that was traded before the quotes moved beyond the specified threshold.

$$VNET = \log\left|\sum_i (d_i vol_i)\right|.$$

In the definition above, d is the direction of trade indicator (buy $= 1$ and sell $= -1$) and vol is the number of shares traded. The summation is over all transactions within a given price-duration. As described in the next section, the entire VNET equation is estimated in log levels.

5. Empirical models

In developing an intraday model of liquidity, we hope to clarify the relationship between market activity and price movements. The amount of one-sided volume (VNET) that can be sustained before prices adjust does not appear to be constant over time for a given stock. If this variability relates to market perceptions about the extent of informational asymmetry, then perhaps some signals may be found in trading patterns. Within the price-duration framework, the time between price events should be included in the set of explanatory variables. It is likely that pertinent information may be conveyed in the decision of when to trade as well as in how many shares and at what price. In light of this we first model the timing of price-durations.

5.1. PTIME

The autoregressive conditional duration (ACD) model assumes the time between future events to be a function of the time between past events. The capabilities of these models to forecast time durations was introduced by Engle and Russell (1997). In what can be thought of as an equivalent to an ARMA process for time durations, these models forecast the time between events conditioned on their history.

$$\psi_t = \omega + \alpha_1 X_{t-1} + \beta_1 \psi_{t-1}.$$

The standard ACD(1, 1) specification shown above posits the conditional time (ψ) to be a function of the previous actual duration (X) and the previous conditional duration. As described below, a generalization of this basic model is employed to model PTIME for each stock. As market depth is obviously contingent on the length of the trading interval, a more accurate estimate of the expected rate of time flow should improve these dynamic liquidity models.

In estimating conditional price-duration times, PTIME is first diurnally adjusted with respect to time-of-day effects by dividing by the mean value of PTIME in the relevant hour of the day. Although the intraday pattern does not always display the prominent inverted U-shape found in transaction-based duration times, F-tests confirm the significance of hourly dummy variables in the model.[3] The normalized price-durations are next examined for serial cor-

[3] The inverted U-shaped pattern for intraday duration times reflects heightened market intensity and volatility at the beginning and end of each trading session.

relation. With 15 lags, the Ljung-Box statistics for the null hypothesis of temporal independence exceed the 5% critical value for 8 of the 17 stocks examined, providing evidence of significant clustering of price movements over time. These results corroborate the findings of Engle and Russell (1998). This suggests that an ACD model may indeed be useful in forecasting PTIME.

The ACD model is considered a *conditional point process*. This class of models focuses on the timing of irregularly spaced events. Fundamental to this formulation is the *hazard function*, which is the instantaneous probability of an event. Although the hazard may often depend on the time since the last event, a constant hazard function is a simple initial guess as to the nature of this process. Econometrically this case is rather clean because the standardized durations (ε), defined as the actual duration (PTIME) divided by expected duration (ψ), will be exponentially distributed with a standard deviation equal to their mean of one. To test the validity of this assumption we perform one-tailed T-tests for unit variance. The hypothesis is rejected for 7 of the 17 stocks, revealing significant excess dispersion.

To accommodate the greater volatility apparent in the data we instead estimate the standardized durations with a Weibull distribution. This allows for a monotonically increasing or decreasing hazard function, as well as the central (constant hazard) case that is equivalent to the exponential model. Given that the data appears to have a tendency for long durations (excess dispersion), we may expect a decreasing hazard function ($\gamma < 1$) to work best. The Weibull ACD models appear to adequately capture the excess dispersion of the input series. The hypothesis that the fitted values have unit variance, $H_0{:}\,\sigma_\varepsilon^2 = 1$, can be accepted for all 17 stocks.

Before settling on a final specification for estimating conditional price-durations, a few additional structural choices are necessary. Namely, the number of lags and any exogenous variables to be included. To prevent overnight episodes from entering the analysis we must exclude one observation at the start of each new day for every lag. A simple first order process produces reasonable results and minimizes the number of lost observations. To this base structure we add a lagged value of the nominal bid–ask spread (SPR_NOM) as a predetermined component. The Ljung-Box autocorrelation statistics for the conditional durations produced by this final specification are below the 5% critical value for all 17 stocks examined.

In Eq. (1), EPTIME is the conditional expectation of PTIME. This WACD(1, 1) model uses the lagged conditional expectation and the lagged value of PTIME, along with the predetermined variable SPR_NOM_{t-1} to forecast the time between price changes.

$$EPTIME_t = \omega + \alpha_1 PTIME_{t-1} + \beta_1 EPTIME_{t-1} + \phi SPR\_NOM_{t-1}. \tag{1}$$

Table 2 below lists the ACD parameter estimates for each of the 17 stocks in the sample. As can be seen in the second to last column, lagged SPR_NOM is highly

Table 2
Maximum likelihood estimated coefficients (p-values) for Eq. (1) (Nov. 1990–Jan. 1991)

Eq. (1): $\text{EPTIME}_t = \omega + \alpha\text{PTIME}_{t-1} + \beta\text{EPTIME}_{t-1} + \phi\text{SPR\_NOM}_{t-1}$

	ω	α	β	ϕ	$\gamma(\text{H}_0 : \gamma = 1)$
BA	0.78 (0.0001)	0.16 (0.001)	0.49 (0.0001)	− 0.36 (0.0001)	0.99 (0.69)
CAL	1.98 (0.0001)	0.059 (0.27)	− 0.173 (0.23)	− 0.86 (0.0001)	1.03 (0.45)
CL	1.07 (0.0001)	0.04 (0.42)	0.40 (0.04)	− 0.50 (0.0001)	1.01 (0.80)
CPC	0.29 (0.001)	0.09 (0.001)	0.77 (0.0001)	− 0.15 (0.004)	0.97 (0.19)
DI	1.07 (0.0001)	0.10 (0.04)	0.14 (0.48)	− 0.34 (0.002)	0.91 (0.0002)
FDX	0.33 (0.001)	0.05 (0.11)	0.82 (0.0001)	− 0.21 (0.0001)	0.94 (0.13)
FNM	0.35 (0.0001)	0.14 (0.001)	0.73 (0.0001)	− 0.22 (0.0001)	1.04 (0.32)
FPL	1.64 (0.0001)	− 0.05 (0.05)	− 0.35 (0.03)	− 0.37 (0.0001)	0.89 (0.0002)
GE	0.23 (0.0001)	0.07 (0.002)	0.82 (0.0001)	− 0.13 (0.0001)	0.95 (0.09)
GLX	0.17 (0.06)	0.11 (0.01)	0.73 (0.0001)	− 0.01 (0.76)	1.01 (0.80)
HAN	1.95 (0.0001)	0.02 (0.42)	− 0.51 (0.0001)	− 0.53 (0.0001)	0.95 (0.13)
IBM	0.25 (0.0001)	0.21 (0.0001)	0.64 (0.0001)	− 0.09 (0.0001)	1.03 (0.32)
MO	0.76 (0.0001)	0.11 (0.03)	0.54 (0.001)	− 0.40 (0.0001)	1.05 (0.23)
POM	1.80 (0.0001)	0.06 (0.19)	0.01 (0.92)	− 0.86 (0.0001)	1.05 (0.23)
SLB	0.58 (0.0001)	0.12 (0.003)	0.55 (0.0001)	− 0.24 (0.0002)	1.04 (0.21)
T	0.28 (0.05)	0.14 (0.003)	0.65 (0.0001)	− 0.07 (0.27)	0.92 (0.04)
XON	0.25 (0.05)	0.14 (0.003)	0.67 (0.0001)	− 0.06 (0.32)	0.95 (0.11)

significant in predicting the time between price changes. The negative coefficient supports the theoretical prediction that wider bid–ask spreads are indicative of a more volatile market.

EPTIME incorporates past information on PTIME and SPR_NOM and represents the conditional forecast of the time until the next significant price change. As will be seen, unanticipated shocks to PTIME will be most useful in modeling market depth. PTIME_ERR is defined as actual divided by expected PTIME and is the fraction of PTIME that could not be predicted by the WACD(1,1) model. While these residuals should be independent of our information set, there may still remain some unidentifiable, yet systematic component of the forecast error that is related to the level of liquidity in the market. Surprises in the timing of price changes reflect unanticipated trade flows. To the extent that aggregate market activity is endogenous to an agent's transaction decision, PTIME_ERR can have direct impact on realized market depth.

5.2. *VNET* (*depth*)

VNET measures the net directional volume that can be traded before prices are adjusted. Ex post, the new statistic provides a measure of realized market depth. This section develops a model to forecast market depth over a price-duration. For a variable to potentially explain time-varying liquidity it must be related to the extent of informational asymmetry in the market. As discussed earlier, market microstructure theory provides many candidates. The explanatory variables tested in the various formulations for VNET are:

SPREAD = log(ASK/BID) at the final trade of the price-duration[4]
NUMBER = the log of the number of trades during the price-duration
VOLUME = the log of the total volume traded during the price-duration
NUM_SPR = the log of the number of trades occurring at large spreads[5]
VOL_SPR = the log of the aggregate volume transacted at large spreads
PJUMP = the log of the absolute price change over the duration
EPTIME = the conditional expectation of PTIME[6]
LEPTIME = log(EPTIME)
PTIME_ERR = log(PTIME / EPTIME)

A number of different specifications are examined, with the search working from general down to more specific. The most general formulation tested on each of the 17 stocks is included in Appendix A. Because we are looking for a single specification that explains liquidity in all 17 of the individually modeled stocks, our choice for the "best" model is somewhat subjective. The set of regressors that display statistical significance for a majority of the stocks turns out to be fairly concise. Eq. (2) below (with all variables in logs) is the preferred model. For 14 of the 17 stocks, F-tests can not reject the hypothesis that the four variables dropped from the most general model are insignificant.

$$VNET = \beta_0 + \beta_1 SPREAD(-1) + \beta_2 VOLUME(-1) + \beta_3 NUMBER(-1) + \beta_4 LEPTIME + \beta_5 PTIME\_ERR. \quad (2)$$

The lagged dependent variable never enters significantly in any of the specifications within our search process. However, estimation of an AR(1) model of VNET found all stocks to have positive autocorrelations, with 13 of the 17 statistically significant. The insignificance of lagged VNET in Eq. 2 implies that the right-hand side variables must adequately represent the past depth of the market.

Looking at the regression results in table 3, the coefficients on SPREAD(− 1) appear to qualitatively fit our expectations. The bid–ask spread immediately preceding a price-duration is negatively related to VNET for 14 of the 17 stocks, although the confidence level for the estimates is above 95% for only 5. The

[4] Interestingly, the end-of-duration bid–ask spread out-performed the average spread in all specifications. Perhaps in forecasting liquidity it is important to include evidence of the informational concerns at the instant closest to the upcoming forecast period.

[5] Large spreads are any bid–ask deviation greater than the minimal one eighth. Minimal spreads are present in nearly half of all transactions for most stocks.

[6] Expected PTIME is the one-step forecast taken from the WACD(1, 1) model.

spread also impacts VNET indirectly through the expected PTIME in the ACD model. This effect is also generally negative. Since the bid-ask spread and depth are both aspects of liquidity, this relationship is not surprising.

The number of trades per duration depicts the transaction intensity of the market. If periods of unusually rapid trading reflect an influx of informed traders, asymmetric information models would predict the number of trades within a price-duration to negatively impact liquidity. Indeed, the coefficient on NUMBER(− 1) is negative for all but one stock, with 8 of these statistically significant.

While the aggregate number of shares traded within a duration (VOLUME) may be another indication of transaction intensity, it also provides perspective for the relative imbalance between buys and sells associated with a given level of VNET. Since VNET is an absolute measure of one-sided trading, higher VOLUME implies a smaller percentage imbalance in orders, all else equal. In Table 3, the coefficients on VOLUME are uniformly smaller than one. This less than proportional response of VNET to VOLUME may reflect the heightened risk of informed trading associated with higher volume trades.

The error in forecasting the time length of a price-duration, PTIME_ERR, has an unambiguously significant positive impact in all of the stocks tested. Although this is a contemporaneous variable, a trader can influence this shock term by trading on one side of the market. The positive coefficient on PTIME_ERR shows that the market interprets "impatience" to reflect a high likelihood of asymmetric information. In this way, rapid trading reduces the volume that could otherwise be traded at a particular price. From Table 3 it can be seen that this coefficient is about 0.4 so that a trader who spreads his trades over twice the expected time, all else equal, would face market depth 40% greater.

Eq. (2) is estimated assuming PTIME_ERR to be a weakly exogenous variable, which may seem somewhat tenuous. However, if the parameter of interest is the expected value of VNET *conditional* on the time allowed to trade, PTIME_ERR will be weakly exogenous. That is, given an expectation of the time between price movements (EPTIME), the error in this forecast (PTIME_ERR) is determined by actual PTIME. And in this model, contemporaneous PTIME is under the control of the agent since trading activity instigates price movements. We envision a trader who contemplates exercising his market power at a particular speed and who wants to know how many shares can be traded in that time with less than a specified price impact. The answer is the expectation of VNET conditional on PTIME.

The anticipated duration, LEPTIME, also enters positively in Eq. (2) and is statistically significant for 12 stocks. The expected time for prices to move a fixed amount is simply the reciprocal of an expected volatility measure. Since the model is estimated in logs, the coefficient is interpreted as the negative of a volatility effect. It is therefore not surprising that increased volatility leads to decreased market depth since high volatility is associated with news and the potential for informed trading.

Table 3
OLS estimated coefficient (*p*-values) for Eq. (2) (Nov. 1990–Jan. 1991).

Log(VNET) = c + …	LOG(SPREAD(− 1))	LOG(VOLUME(− 1))	LOG(NUMBER(− 1))	LOG(EPTIME)	LOG(PTIME_ERR)
BA	− 0.058 (0.71)	**0.183 (0.03)**	− 0.152 (0.14)	0.135 (0.27)	**0.334 (0.0001)**
CAL	− 0.612 (0.07)	0.120 (0.14)	− 0.195 (0.13)	**0.756 (0.001)**	**0.250 (0.0004)**
CL	− 0.178 (0.29)	**0.173 (0.03)**	− 0.169 (0.14)	**0.470 (0.0004)**	**0.314 (0.0001)**
CPC	**−0.481 (0.001)**	**0.198 (0.0001)**	**−0.183 (0.02)**	**0.229 (0.04)**	**0.314 (0.0001)**
DI	− 0.015 (0.95)	**0.299 (0.0001)**	**−0.325 (0.001)**	**0.462 (0.001)**	**0.438 (0.0001)**
FDX	− 0.240 (0.17)	**0.291 (0.0001)**	**−0.392 (0.0002)**	**0.267 (0.02)**	**0.351 (0.0001)**
FNM	− 0.041 (0.78)	**0.377 (0.0001)**	**−0.215 (0.05)**	**0.203 (0.06)**	**0.347 (0.0001)**
FPL	**−0.835 (0.0001)**	0.067 (0.25)	− 0.036 (0.65)	**0.519 (0.0001)**	**0.579 (0.0001)**
GE	− 0.176 (0.12)	**0.386 (0.0001)**	**−0.205 (0.02)**	**0.305 (0.001)**	**0.329 (0.0001)**
GLX	**−0.391 (0.02)**	**0.274 (0.0001)**	**−0.266 (0.002)**	**0.288 (0.02)**	**0.440 (0.0001)**
HAN	**−2.27 (0.0001)**	0.098 (0.06)	− 0.138 (0.14)	− 0.060 (0.68)	**0.432 (0.0001)**
IBM	− 0.172 (0.11)	**0.268 (0.004)**	− 0.186 (0.07)	**0.334 (0.0001)**	**0.311 (0.0001)**
MO	− 0.214 (0.23)	**0.391 (0.002)**	**−0.258 (0.04)**	0.183 (0.29)	**0.373 (0.0001)**
POM	**−0.801 (0.01)**	− 0.005 (0.93)	− 0.036 (0.73)	**0.388 (0.03)**	**0.393 (0.0001)**
SLB	0.089 (0.44)	**0.182 (0.003)**	**−0.220 (0.01)**	**0.388 (0.0003)**	**0.385 (0.0001)**
T	0.079 (0.63)	**0.302 (0.0002)**	− 0.155 (0.06)	0.232 (0.08)	**0.405 (0.0001)**
XON	0.268 (0.07)	0.024 (0.75)	0.008 (0.92)	**0.253 (0.02)**	**0.410 (0.0001)**

5.3. Pooled estimates

The estimates for the individual stocks presented above are similar in character but not uniformly significant. This is unsurprising because the sample period is rather short and there are many obvious sources of noise in VNET that inflate standard errors. To obtain estimates that summarize the behavior of VNET across these 17 stocks, a pooled regression is computed. As the typical volume and the threshold for price-durations are different for each stock, it is important to allow stock-specific effects, so the intercept is now a 17-element vector. Because the regression variables are already in log levels, this set of additive dummies control for the cross-sectional heterogeneity assuming elasticities are constant across stocks.

The OLS estimated pooled equation (standard errors in parentheses) is

$$VNET_t = -\underset{(0.04)}{0.33}\cdot SPREAD_{t-1} + \underset{(0.02)}{0.15}\cdot VOLUME_{t-1} - \underset{(0.02)}{0.15}\cdot NUMBER_{t-1} + \underset{(0.03)}{0.39}\cdot LEPTIME_t + \underset{(0.01)}{0.38}\cdot PTIME\_ERR_t. \quad (3)$$

All regressors now show significant influence at greater than 99.9% confidence, which is to be expected as we increase the sample size. The signs and the magnitudes are similar to the 17 individual regressions and support our earlier conclusions.

6. Robustness of the model

To validate the usefulness of VNET and the price-duration framework for measuring and forecasting intraday market depth, we next test the robustness of the models with respect to both the sample period and the width of the price threshold. As a by-product of the latter exercise, an estimate of the shape of the market reaction curve will be derived for a representative stock.

6.1. The new sample (August–December 1997)

For the five months between August and December 1997, transaction and quote records are extracted from the TAQ (Trades and Quotes) data set for 16 of the 17 stocks (California Federal Bank (CAL) is no longer traded on the NYSE). This period is desirable for several reasons. First, none of the 16 stocks split during this time. Second, the minimum tick size was lowered from 1/8th to 1/16th on June 24, 1997, providing an interesting comparison to the original sample. Finally, the market is much less bullish in this more recent period with an average monthly rate of return of 0.3%, as compared to 4.1% in the 1990–91 TORQ data.

To maintain consistency with the original analysis, VNET should be measured over an identical percentage price change. However, price-durations are defined by a nominal price change in order to insulate VNET from the distortions caused by price discreteness.[7] Unfortunately, this complicates cross-sample comparisons. Since nominal share prices may have shifted substantially between 1990 and 1997, the nominal price threshold must also be adjusted in order to generate the desired relative threshold. If the percentage threshold for each stock could be matched precisely across samples, and if the volatility of a stock was unchanged between 1990 and 1997, then we would expect the average number of price-durations per day to be roughly similar as well.[8] However, for many stocks, daily volume and intraday volatility have increased from their earlier levels (see Fig. 3). Consequently, the average number of identified price-durations per day is greater in the 1997 sample for 11of the 16 stocks.

In comparing only a few months of data separated by nearly seven years, it is difficult to make a definitive assessment of market trends. Any noted differences between the two periods could very well be unrelated to their temporal spacing. The increase in trading activity and price volatility between the original period of November 1990–January 1991 and the new sample of August 1997–December 1997 may or may not reflect a consistent trend in the NYSE, and for this study it is not important.

Of more direct interest to this research is the difference in market depth as measured by VNET. Normalized by the percentage price change implied by the threshold size, the average value of VNET is greater in 1997 for 11 of 16 stocks. VNET measures depth in units of volume per price change. It may be expected that average VNET will rise with volume, and fall with volatility. Inspection of Fig. 3 reveals that this does seem to be the case. The five stocks that have smaller VNET in the current period than in the original period also have some of the smallest gains in volume relative to their increase in volatility. As both volume and volatility are coarse measures of liquidity, this relationship is reassuring.

6.2. Re-estimating Eq. (*1*) (*PTIME*)

Similar to the original sample, the new deseasonalized PTIME series reveal significant autocorrelation. At 15 lags, the Ljung-Box statistics are greater than

---

[7] The use of a percentage price threshold scheme could potentially create large jumps in the implied absolute price change as the number of discrete price ticks changes with the underlying asset price.

[8] Discreteness in prices limits the choice of possible nominal thresholds making it often impossible to precisely match the relative threshold of the original sample.

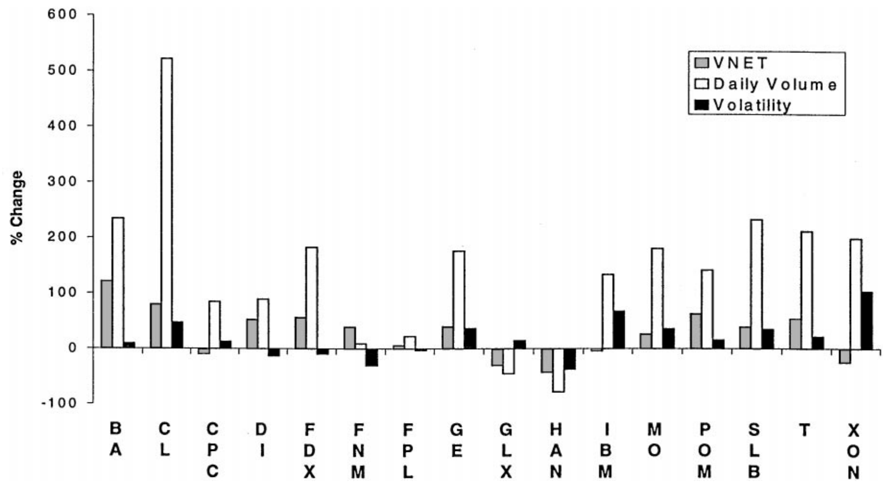

Fig. 3. The percentage changes in average VNET (normalized by the percentage price threshold), daily volume, and annualized half-hour volatility from Nov. 1990–Jan. 1991 to Aug. 1997–Dec. 1997.

the critical value of 25 for 15 of 16 stocks. In addition, the excess dispersion of PTIME (i.e. standard deviation greater than one) noted in Section 5 is also apparent in the new sample for 11 of 16 stocks.

The maximum likelihood estimates of the WACD (1, 1) model are displayed in Table 4. The model seems to do a reasonable job of capturing the autocorrelation, with only 3 residual series still showing significant autocorrelation. Unfortunately, the excess dispersion is not so effectively remedied. As the final column of Table 4 reveals, the γ parameter, which controls the rate of increase/decrease of the hazard function, is significantly greater than one for 13 stocks. This implies an increasing hazard, which does not help account for the abundance of long durations in the raw data. Indeed, the residuals for 11 stocks retain excess dispersion. Obviously this WACD (1, 1) specification does not completely characterize the distribution of this new sample. However, for consistency with the original estimates, this model will still be used to produce expectations of the time between price changes (EPTIME).

Again in accordance with the earlier results, the effect of the bid-ask spread on expected PTIME is negative. Comparing Table 4 below to the original estimates in Table 2, the coefficients for SPR_NOM (the nominal bid-ask spread) are nearly an order of magnitude smaller for the new sample. This may reflect the 1/16th minimum tick regime of this more recent period, as opposed to 1/8th in 1990–91. Reduced tick size has led to smaller average bid–ask spreads, at least for the 16 stocks in this study over the two periods being examined. The dynamics of the bid–ask spread and its relationship to volatility may also have changed since the switch to 1/16ths. However, this paper will not investigate this issue beyond its immediate impact on PTIME.

Table 4
Maximum likelihood estimated coefficients (p-values) for Eq. (1) (Aug. 1997–Dec. 1997)

Eq. (1): $\text{EPTIME}_t = \omega + \alpha\text{PTIME}_{t-1} + \beta\text{EPTIME}_{t-1} + \phi\text{SPR\_NOM}_{t-1}$
(August 1997–December 1997)

	ω	α	β	ϕ	$\gamma(H_0 : \gamma = 1)$
BA	0.04 (0.002)	0.09 (0.0001)	0.89 (0.0001)	– 0.013 (0.002)	0.99 (0.76)
CAL	NA	NA	NA	NA	NA
CL	0.13 (0.0001)	0.07 (0.0001)	0.86 (0.0001)	– 0.062 (0.0001)	1.05 (0.01)
CPC	0.11 (0.004)	0.10 (0.0001)	0.85 (0.0001)	– 0.053 (0.02)	1.17 (0.0001)
DI	0.07 (0.0001)	0.05 (0.0001)	0.93 (0.0001)	– 0.039 (0.0001)	1.06 (0.01)
FDX	0.14 (0.0001)	0.13 (0.0001)	0.79 (0.0001)	– 0.066 (0.0001)	0.99 (0.55)
FNM	0.15 (0.0001)	0.07 (0.0001)	0.88 (0.0001)	– 0.084 (0.0001)	1.20 (0.0001)
FPL	0.06 (0.001)	0.04 (0.0001)	0.93 (0.0001)	– 0.023 (0.01)	1.07 (0.001)
GE	0.05 (0.0001)	0.17 (0.0001)	0.81 (0.0001)	– 0.012 (0.0001)	1.17 (0.01)
GLX	0.33 (0.0001)	0.11 (0.0001)	0.69 (0.0001)	– 0.132 (0.0001)	1.05 (0.02)
HAN	0.26 (0.01)	0.06 (0.01)	0.78 (0.0001)	– 0.087 (0.01)	1.13 (0.0001)
IBM	0.08 (0.0001)	0.13 (0.0001)	0.83 (0.0001)	– 0.031 (0.0001)	1.08 (0.0002)
MO	0.10 (0.0001)	0.12 (0.0001)	0.81 (0.0001)	– 0.027 (0.0001)	1.06 (0.02)
POM	0.07 (0.0002)	0.08 (0.0001)	0.88 (0.0001)	– 0.027 (0.0003)	1.10 (0.0001)
SLB	0.09 (0.0001)	0.14 (0.0001)	0.80 (0.0001)	– 0.022 (0.003)	1.09 (0.0001)
T	0.08 (0.0001)	0.09 (0.0001)	0.86 (0.0001)	– 0.028 (0.0001)	0.97 (0.16)
XON	0.13 (0.0001)	0.11 (0.0001)	0.82 (0.0001)	– 0.045 (0.0001)	1.12 (0.0001)

A likelihood ratio breakpoint test is used to test the stability of the ACD model of PTIME across the two sample periods, stock by stock. The independent estimates presented above represent the unrestricted model. For the restricted model, the two samples are stacked and deseasonalized by a single set of time-of-day constants. Eq. (1) is then estimated for this combined sample. The coefficients and T-statistics are listed in Table 5. Comparing the maximum of the log-likelihood function for these restricted estimates to the unrestricted estimates, the likelihood ratio test accepts the stability of the model across the two periods for 7 of the 16 stocks. However, given the nearly seven year span between these two periods and the qualitative similarities in the parameter estimates and large sample sizes, a lack of universal statistical stability does not necessarily invalidate the original findings.

6.3. Re-estimating Eq. (2) (VNET)

The estimates of Eq. (2) for the new sample are displayed in Table 6. The sign, magnitude, and significance levels of the coefficients are fairly similar to those of the original sample presented in Table 3. Standard errors are a bit smaller in the new sample, which generally has a greater number of observations per stock. As in the previous section, a Chow breakpoint test is used to examine the model's

Table 5
Maximum likelihood estimated coefficients (*p*-values) for Eq. (1) (joint sample).

Eq. (1): $\text{EPTIME}_t = \omega + \alpha\text{PTIME}_{t-1} + \beta\text{EPTIME}_{t-1} + \phi\text{SPR\_NOM}_{t-1}$

	ω	α	β	ϕ	$\gamma(H_0: \gamma = 1)$
BA	0.03 (0.0001)	0.08 (0.0001)	0.90 (0.0001)	− 0.013 (0.001)	1.00 (0.84)
CAL	NA	NA	NA	NA	NA
CL	0.02 (0.0002)	0.07 (0.0001)	0.92 (0.0001)	− 0.009 (0.03)	1.06 (0.0001)
CPC	0.01 (0.04)	0.05 (0.0001)	0.94 (0.0001)	0.001 (0.69)	1.08 (0.0001)
DI	0.12 (0.004)	0.12 (0.0001)	0.79 (0.0001)	− 0.020 (0.06)	1.08 (0.01)
FDX	0.07 (0.0001)	0.18 (0.0001)	0.76 (0.0001)	− 0.005 (0.42)	1.00 (0.92)
FNM	0.36 (0.0001)	0.19 (0.0001)	0.61 (0.0001)	− 0.131 (0.001)	1.13 (0.0001)
FPL	0.01 (0.06)	0.02 (0.002)	0.97 (0.0001)	− 0.001 (0.48)	0.98 (0.27)
GE	0.03 (0.0001)	0.12 (0.0001)	0.86 (0.0001)	− 0.004 (0.0001)	1.06 (0.0003)
GLX	0.28 (0.0001)	0.19 (0.0001)	0.62 (0.0001)	− 0.08 (0.001)	1.04 (0.09)
HAN	0.25 (0.001)	0.05 (0.01)	0.81 (0.0001)	− 0.11 (0.0003)	0.99 (0.69)
IBM	0.05 (0.0001)	0.15 (0.0001)	0.83 (0.0001)	− 0.020 (0.0001)	1.07 (0.0001)
MO	0.05 (0.0001)	0.14 (0.0001)	0.83 (0.0001)	− 0.014 (0.0001)	1.04 (0.01)
POM	0.01 (0.07)	0.02 (0.002)	0.97 (0.0001)	0.001 (0.76)	1.04 (0.07)
SLB	0.01 (0.001)	0.07 (0.0001)	0.92 (0.0001)	0.001 (0.84)	1.10 (0.0001)
T	0.04 (0.0001)	0.11 (0.0001)	0.87 (0.0001)	− 0.009 (0.04)	0.95 (0.01)
XON	0.02 (0.0001)	0.09 (0.0001)	0.90 (0.0001)	− 0.004 (0.05)	1.04 (0.0002)

stability over time. The combined sample estimation (see Table 7) uses EPTIME derived from the restricted PTIME model in the previous section. Chow tests find the VNET model to be stable across samples for 9 of the 16 stocks. Given the significant differences in the level of trading, volatility, and minimum tick size between the two samples, these results present a relatively promising statement as to the integrity of the relationships between liquidity, prices, and market activity presented in Section 5.

6.4. Sensitivity to the price threshold

Another direction in which to test the robustness of the relationship between market liquidity and trading activity is with respect to the width of the aggregation window. Our original thresholds were designed to generate roughly 10–15 durations per day. However, this choice was somewhat arbitrary. While a certain base number of observations are required for an intraday study, the optimal price range over which to measure market depth ultimately depends on the tolerance of the trader.

The duration-based statistic, VNET, measures the depth of the market away from the posted quotes. VNET estimates a single point along the market reaction curve from which we can infer only the average slope, not the true shape

Table 6
OLS estimated coefficients (*p*-values) for Eq. (2) (Aug. 1997–Dec. 1997).

LOG(VNET) = c + …	LOG(SPREAD(− 1))	LOG(VOLUME(− 1))	LOG(NUMBER(− 1))	LOG(EPTIME)	LOG(PTIME_ERR)
BA	0.011 (0.86)	**0.0635 (0.0001)**	**−0.498 (0.0001)**	**0.329 (0.0001)**	**0.421 (0.0001)**
CAL	NA	NA	NA	NA	NA
CL	**−0.136 (0.01)**	**0.352 (0.0001)**	**−0.318 (0.0001)**	**0.326 (0.0001)**	**0.452 (0.0001)**
CPC	**−0.145 (0.01)**	**0.211 (0.0001)**	**−0.211 (0.0002)**	**0.356 (0.0001)**	**0.424 (0.0001)**
DI	0.001 (0.99)	**0.500 (0.0001)**	**−0.367 (0.001)**	0.120 (0.32)	**0.384 (0.0001)**
FDX	− 0.095 (0.23)	**0.451 (0.0001)**	**−0.340 (0.0001)**	0.037 (0.67)	**0.390 (0.0001)**
FNM	− 0.021 (0.85)	**0.234 (0.04)**	**−0.238 (0.05)**	**0.464 (0.0001)**	**0.209 (0.0001)**
FPL	**−0.224 (0.01)**	**0.215 (0.0001)**	**−0.217 (0.01)**	**0.539 (0.0001)**	**0.556 (0.0001)**
GE	0.083 (0.12)	**0.416 (0.0001)**	**−0.356 (0.0001)**	**0.323 (0.0001)**	**0.312 (0.0001)**
GLX	**−0.240 (0.05)**	**0.281 (0.0001)**	**−0.228 (0.02)**	**0.359 (0.004)**	**0.492 (0.0001)**
HAN	**−0.492 (0.03)**	0.141 (0.10)	− 0.073 (0.60)	**0.446 (0.03)**	**0.477 (0.0001)**
IBM	0.053 (0.22)	**0.369 (0.0001)**	**−0.304 (0.0001)**	**0.248 (0.0001)**	**0.321 (0.0001)**
MO	0.030 (0.62)	**0.256 (0.0001)**	**−0.157 (0.01)**	**0.275 (0.0001)**	**0.410 (0.0001)**
POM	**−0.258 (0.01)**	**0.190 (0.001)**	**−0.193 (0.02)**	**0.750 (0.0001)**	**0.552 (0.0001)**
SLB	0.049 (0.30)	**0.351 (0.0001)**	**−0.224 (0.0001)**	**0.149 (0.01)**	**0.428 (0.0001)**
T	0.104 (0.23)	**0.485 (0.0001)**	**−0.372 (0.0001)**	**0.373 (0.0001)**	**0.380 (0.0001)**
XON	− 0.051 (0.27)	**0.241 (0.0001)**	**−0.208 (0.0001)**	**0.311 (0.0001)**	**0.328 (0.0001)**

Table 7
OLS estimated coefficients (*p*-values) for Eq. (2) (joint sample).

LOG(VNET) = c + …	LOG(SPREAD(− 1))	LOG(VOLUME(− 1))	LOG(NUMBER(− 1))	LOG(EPTIME)	LOG(PTIME_ERR)
BA	**−0.216 (0.0001)**	**0.634 (0.0001)**	**−0.465 (0.0001)**	**0.137 (0.01)**	**0.398 (0.0001)**
CAL	NA	NA	NA	NA	NA
CL	**−0.204 (0.0001)**	**0.346 (0.0001)**	**−0.273 (0.0001)**	**0.180 (0.0001)**	**0.416 (0.0001)**
CPC	**−0.185 (0.0001)**	**0.215 (0.0001)**	**−0.205 (0.0001)**	**0.293 (0.0001)**	**0.407 (0.0001)**
DI	**−0.256 (0.0001)**	**0.376 (0.0001)**	**−0.263 (0.001)**	0.163 (0.12)	**0.399 (0.0001)**
FDX	**−0.265 (0.0001)**	**0.309 (0.0001)**	**−0.186 (0.01)**	**0.143 (0.04)**	**0.384 (0.0001)**
FNM	**−0.192 (0.001)**	**0.419 (0.0001)**	**−0.326 (0.0001)**	**0.382 (0.0001)**	**0.309 (0.0001)**
FPL	**−0.220 (0.0001)**	**0.124 (0.003)**	**−0.138 (0.02)**	**0.462 (0.0001)**	**0.611 (0.0001)**
GE	**−0.100 (0.01)**	**0.503 (0.0001)**	**−0.391 (0.0001)**	**0.272 (0.0001)**	**0.340 (0.0001)**
GLX	− 0.006 (0.95)	**0.283 (0.0001)**	**−0.167 (0.01)**	0.091 (0.31)	**0.508 (0.0001)**
HAN	**−0.674 (0.01)**	0.012 (0.80)	− 0.129 (0.12)	− 0.180 (0.08)	**0.630 (0.0001)**
IBM	0.023 (0.54)	**0.364 (0.0001)**	**−0.288 (0.0001)**	**0.213 (0.0001)**	**0.319 (0.0001)**
MO	− 0.031 (0.54)	**0.267 (0.0001)**	**−0.184 (0.001)**	**0.245 (0.0001)**	**0.406 (0.0001)**
POM	**−0.414 (0.0001)**	**0.188 (0.0001)**	**−0.222 (0.0004)**	**0.509 (0.0001)**	**0.568 (0.0001)**
SLB	**−0.078 (0.04)**	**0.306 (0.0001)**	**−0.188 (0.0001)**	**0.038 (0.34)**	**0.420 (0.0001)**
T	− 0.047 (0.39)	**0.489 (0.0001)**	**−0.355 (0.0001)**	**0.316 (0.0001)**	**0.373 (0.0001)**
XON	− 0.025 (0.52)	**0.188 (0.0001)**	**−0.159 (0.0001)**	**0.330 (0.0001)**	**0.367 (0.0001)**

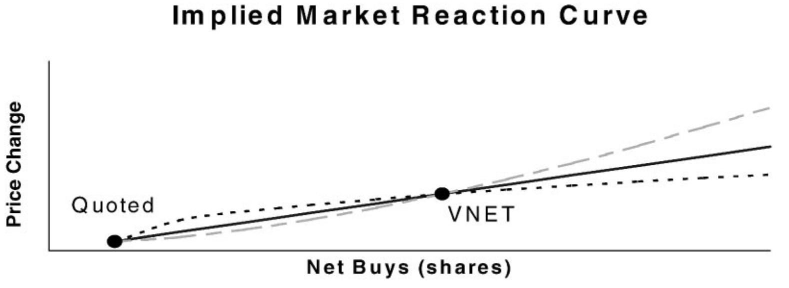

Fig. 4. Some possible shapes of the market reaction curve.

of the curve. As seen in Fig. 4, linear, convex, and concave functions can all be fit through the two known points (quoted depth and VNET).[9]

By adjusting the size of the threshold from which price-durations are defined, one can estimate other points along the market reaction curve. So although the discussion to this point has abstracted from any non-linearity issues, it may be possible to empirically examine the plausibility of this assumption. For this exercise, the 1/16th regime is advantageous in that it permits a greater number of nominal price increments within a given range. Because price-durations are triggered by movements in the mid-point of the quotes, the smallest feasible change is half of the minimal tick, or 1/32nd, as the quotes may move one at a time.

Over the 1997 sample, price-durations are recalculated for Fannie Mae (FNM) using eight different price increments. FNM is a good representative stock because its share price is high enough to support a wide range of threshold widths. The narrowest threshold examined is 3/32nds ($0.0935) and the maximum price change is 10/32nds ($0.3125) – the widest threshold that created an average of at least three observations per day.[10] The number of price-durations at each threshold increment over the 102-day sample is presented in the first column of Table 8.

Average VNET increases, as expected, with the price threshold (see Fig. 5 below). More over, this upward sloping market reaction schedule is fairly linear.[11] The first point on the curve is not measured VNET, but rather the average quoted depth of the market. The price change associated with this value is labeled 1/32nd since fully depleting the quoted volume would necessitate

[9] Beyond the limit order book possibly looking convex in price-volume space, the authors are not aware of any theoretical foundation for the appropriate shape of the market reaction curve.

[10] Even at narrow price thresholds, some days would be expected to contain no price-durations.

[11] Estimated market reaction curves for the other stocks showed similar linear patterns.

Table 8
Number of price-durations (n) and OLS estimated coefficients from Eq. (2) for FNM at various price thresholds (Aug. 1997–Dec. 1997).

Eq. (2): VNET = $\alpha + \beta1$*SPREAD(-1) + $\beta2$*VOLUME(-1) + $\beta3$*NUMBER(-1) + $\beta4$*EP-TIME + $\beta5$*PTIME_ERR

threshold	n	α	$\beta1$	$\beta2$	$\beta3$	$\beta4$	$\beta5$
0.09375	3496	4.682	− 0.187	0.198	− 0.173	0.341	0.471
0.125	2135	3.401	− 0.140	0.346	− 0.348	0.442	0.456
0.15625	1457	4.641	− 0.113	0.383	− 0.313	0.234	0.349
0.1875	1054	5.533	0.111	0.402	− 0.331	0.297	0.318
0.21875	765	4.829	− 0.038	0.474	− 0.460	0.233	0.320
0.25	585	5.033	0.098	0.337	− 0.283	0.436	0.277
0.28125	442	5.226	− 0.021	0.234	− 0.238	0.464	0.209
0.3125	345	5.671	0.144	0.396	− 0.465	0.413	0.240

a move in the quote to at least the next tick (i.e. a half-tick move in the mid-quote). To the extent that the quoted volume can be thought to reflect the depth of the market at small price changes, the coincidence of this point with the rest of the curve provides further support for the plausibility of VNET as a measure of market depth. The rest of the points depict the average, unconditional levels of VNET over the five-month sample.

By replicating the estimation of Eq. (1) and (2) at each of the feasible price thresholds, the *conditional* market reaction curve can be derived, as well. These PTIME and VNET equations are re-estimated on data generated by each of the eight feasible thresholds for Fannie Mae (FNM). The parameters of the VNET model at the various price bands are displayed in Table 8. While there is a good deal of consistency, the relationships do adjust somewhat as the aggregation threshold changes. Bilateral F-Tests confirm the stability of the coefficients across adjacent price thresholds, but a joint test rejects the hypothesis that all eight models are statistically identical. Of course, these tests do not have standard distributions as the estimates are not independent.

From this set of parameter estimates, the behavior of the full market reaction curve can be determined, conditional on the five explanatory variables in Eq. (2). The solid lines in Figs. 6a–e represent the fitted values of VNET using the means of the right-hand side variables over the five-month period. This is the baseline market reaction curve. Each of the figures also plots the response of the market reaction curve to a 50% positive or negative shock to each of the explanatory five factors.

As can be seen in Fig. 6a, a 50% increase the average volume traded within a price-duration (VOLUME) boosts expected depth and pushes the market

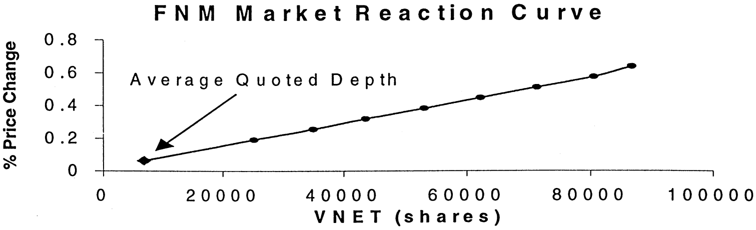

Fig. 5. Average market reaction curve for FNM (Fannie Mae), Aug.–Dec. 1997.

reaction curve to the right, all else equal. Similarly, a decrease in VOLUME pulls the curve left.

Charts 6b–d reveal that shocks to NUMBER, EPTIME, and PTIME_ERR also shift the market reaction curve in a manner consistent with the original findings. Although several of the simulated response curves are somewhat distorted and non-linear, the overall positioning of the curves tends to reinforce our earlier qualitative assessments of Eq. 2. Beyond describing the depth of the market at a specific price change, these simulations represent estimates of market depth, conditional on various aspects of trading behavior, across a continuum of feasible price movements.

Fig. 6e, which displays the sensitivity of the conditional market reaction curve to changes in the bid–ask spread, is far less informative than then the earlier figures. The total impact of a shock to the bid–ask spread includes the direct influence of SPREAD in Eq. (2), plus the indirect influence of SPREAD on EPTIME in Eq. (1). These two effects should work in the same direction, with greater values of SPREAD lowering expected VNET. However, the indirect impact is relatively small, and for Fannie Mae, the direct impact is statistically insignificant and inconsistently signed, as indicated by the coefficient on SPREAD in Table 8. This results in the ill-shaped curves above. Looking back at Table 6, SPREAD is the least significant explanatory variable in the model across the entire group of stocks. This may reflect the incompleteness of the bid–ask spread as a measure of liquidity, particularly when studying the depth of the market for larger price deviations.

Fig. 6. (a) The sensitivity of the FNM conditional market reaction curve to VOLUME. (b) The sensitivity of the FNM conditional market reaction curve to NUMBER. (c) The sensitivity of the FNM conditional market reaction curve to EPTIME. (d) The sensitivity of the FNM conditional market reaction curve to PTIME_ERR. (e) The sensitivity of the FNM conditional market reaction curve to SPREAD.

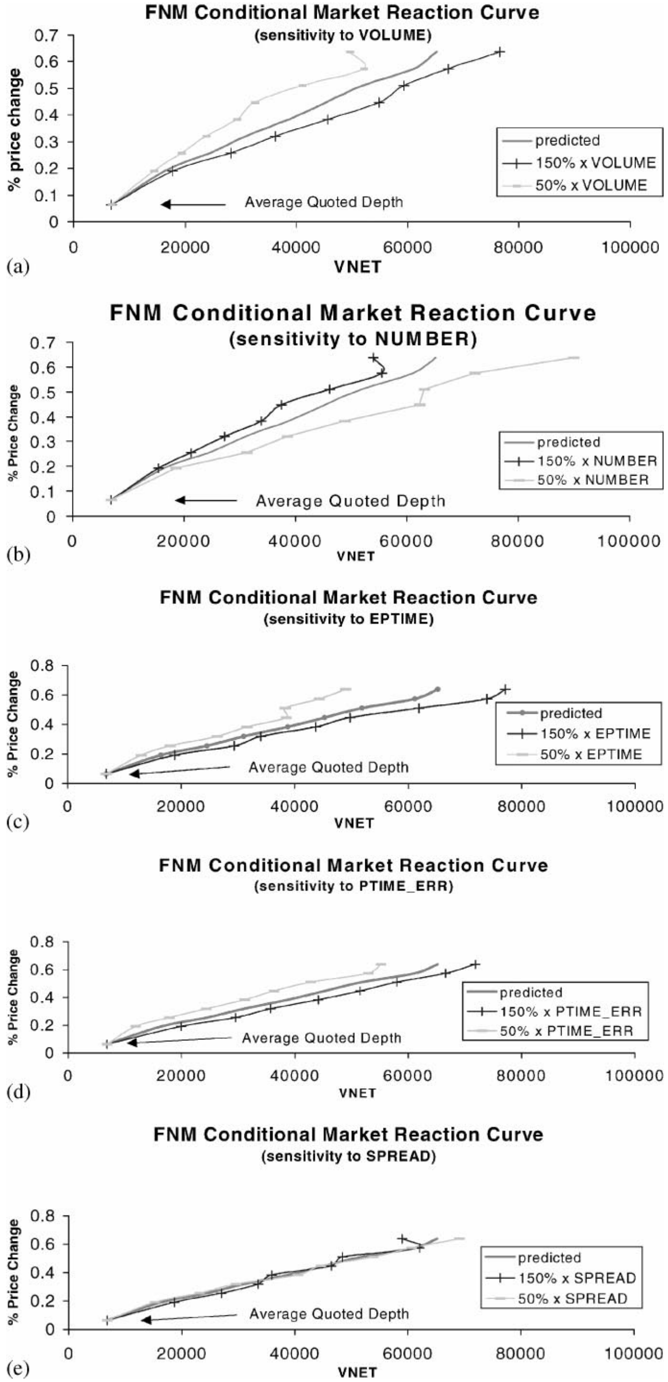
FNM Conditional Market Reaction Curve
(sensitivity to VOLUME)
% price change
predicted
150% x VOLUME
50% x VOLUME
Average Quoted Depth
VNET
(a)
FNM Conditional Market Reaction Curve
(sensitivity to NUMBER)
% Price Change
predicted
150% x NUMBER
50% x NUMBER
Average Quoted Depth
VNET
(b)
FNM Conditional Market Reaction Curve
(sensitivity to EPTIME)
% Price Change
predicted
150% x EPTIME
50% x EPTIME
Average Quoted Depth
VNET
(c)
FNM Conditional Market Reaction Curve
(sensitivity to PTIME_ERR)
% Price Change
predicted
150% x PTIME_ERR
50% x PTIME_ERR
Average Quoted Depth
VNET
(d)
FNM Conditional Market Reaction Curve
(sensitivity to SPREAD)
% Price Change
predicted
150% x SPREAD
50% x SPREAD
Average Quoted Depth
VNET
(e)

7. Conclusions

NYSE transaction and quote data are used to identify, measure, and model intraday variations in the market depth of individual stocks. Our models for the length of a price-duration (PTIME) and the net directional volume traded within a price-duration (VNET) are estimated over two distinct time periods, producing roughly similar parameter estimates. The relative stability of the relationship between quote revisions and trading behavior suggests that our *price-duration* based microstructure approach may indeed touch upon some of the fundamental determinants of equity market liquidity and volatility.

The empirical analysis explores the depth of financial markets, which to this point has been difficult to quantify. By defining price-durations as the time between substantial adjustments in the midpoint of the quotes, a measure of the one-sided trading behind price movements can be obtained. With this new statistic, VNET, we are able to estimate the shape of the market reaction curve, both ex ante and ex post. Our models of VNET reveal that the realized depth of the market varies according to internal trading conditions. In general, the market traits associated with a higher likelihood of price adjustment following a given amount of one-sided volume (small VNET) are similar to those corresponding to low liquidity as represented by tightness (wide bid–ask spread) in earlier studies. This result is important in that it unifies our definition of depth with more traditional views of liquidity.

The models propose some strategies of how to trade large volume at the least cost. First, it may seem obvious that the greater the overall trading volume, the more of a nominal imbalance will be accepted by the market. However, the percentage imbalance between buys and sells sufficient to move prices declines with the total number of shares traded. The number of transactions per duration also appears to reduce the depth of the market. This supports the notion that market thickness is generally a consequence of informed traders flooding the market after a semi-private news event.

The empirical models find that movements in VNET are negatively correlated with movements in the bid–ask spread. Along with providing evidence that the new statistic is a valid measure of liquidity, this relationship adds another trading strategy component, albeit an obvious one: when the market is tight it will also lack depth. In addition, the positive impact of the expected duration length on expected VNET suggests that when the market is volatile it will offer less depth. Finally, unanticipated shocks to the length of a price-durations, represented by PTIME_ERR, also increase realized depth. With respect to large volume trading strategies, this result carries the implication that patience may greatly reduce transaction costs. These results carry the implication that trading behavior may play a significant role in shaping and predicting the intraday liquidity of the stock market.

Appendix A. General model estimation results

Dependent variable: VNET

	VNET (−1)	SPREAD (−1)	VOLUME (−1)	VOLSPR (−1)	NUMBER (−1)	NUMSPR (−1)	PJUMP (−1)	EPTIME	PTIMEERR
BA	−0.047	0.731	0.161	0.112	−0.039	−0.241	−0.059	0.109	**0.330**
CAL	0.062	**1.592**	0.102	**−0.234**	−0.122	0.086	−0.064	**0.496**	**0.223**
CL	−0.012	−0.059	0.190	0.004	−0.145	−0.057	−0.263	**0.474**	**0.312**
CPC	−0.003	**−0.790**	0.158	0.042	−0.166	−0.018	0.105	**0.227**	**0.314**
DI	−0.096	**0.969**	**0.531**	−0.149	**−0.552**	0.240	−0.357	**0.476**	**0.426**
FDX	0.023	−0.493	**0.319**	−0.055	**−0.670**	0.425	−0.092	**0.280**	**0.353**
FNM	0.068	0.906	**0.317**	−0.035	−0.113	−0.098	−0.755	0.210	**0.334**
FPL	−0.030	1.085	0.139	−0.147	0.051	−0.044	0.207	**0.603**	**0.579**
GE	0.064	−0.529	**0.234**	0.107	−0.010	**−0.264**	0.158	**0.305**	**0.333**
GLX	−0.019	**−1.62**	**0.253**	0.091	**−0.284**	−0.008	**0.798**	**0.299**	**0.451**
HAN	0.186	−0.847	−0.048	**−0.174**	−0.159	**0.321**	0.110	−0.084	**0.436**
IBM	0.063	−0.133	0.124	0.073	0.024	−0.225	0.184	**0.341**	**0.316**
MO	0.040	−1.29	**0.363**	−0.017	**−0.403**	0.273	0.840	0.169	**0.392**
POM	0.182	0.909	−0.135	−0.158	0.007	0.073	−0.024	**0.420**	**0.387**
SLB	0.081	−0.113	0.018	0.084	−0.016	−0.220	−0.072	**0.369**	**0.387**
T	0.036	1.38	**0.326**	−0.137	**−0.235**	0.175	−0.534	0.231	**0.401**
XON	−0.006	0.859	0.046	−0.034	0.023	−0.017	0.088	**0.252**	**0.411**

References

Easley, D., O'Hara, M., 1987. Price, trade size, and information in securities markets. Journal of Financial Economics 19, 69–90.

Easley, D., O'Hara, M., 1992. Time and the process of security price adjustment. The Journal of Finance 47, 577–606.

Engle, R., Russell, J., 1997. Forecasting the frequency of changes in quoted foreign exchange prices with the autoregressive conditional duration model. Journal of Empirical Finance 4, 187–212.

Engle, R., Russell, J., 1998. Autoregressive conditional duration: a new model for irregularly spaced data. Econometrica 66, 1127–1162.

Foster, F.D., Viswanathan, S., 1995. Can speculative trading explain the volume-volatility relation? Journal of Business and Economic Statistics 13, 379–408.

Hasbrouck, J., 1988. Trades, quotes, inventories, and information. Journal of Financial Economics 22, 229–252.

Kyle, A.S., 1985. Continuous auctions and insider trading. Econometrica 53, 1315–1365.

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

McInish, T.H., Wood, R.A., 1992. An analysis of intraday patterns in bid/ask spreads for NYSE stocks. The Journal of Finance 47, 753–764.