

Trade-time based measures of liquidity

Yashar H. Barardehi
Ohio University

Dan Bernhardt
University of Illinois and University of Warwick

Ryan J. Davies
Babson College

Dramatic microstructure changes in equity markets have made standard liquidity measures less accurate proxies of trading costs. We develop trade-time based liquidity measures that reflect per-dollar price impacts of fixed-dollar volumes. Our measures better capture institutional trading costs and better explain the cross-section of returns than do standard measures, especially in recent years. Despite improvements in measures of market quality, expected trading costs still have explanatory power for the cross-section of expected returns: we obtain monthly liquidity premium estimates of 5.3bp for expected returns and 2.4bp for risk-adjusted returns. Estimated premia rise after the financial crisis and remain high thereafter. (*JEL* G12, G14)

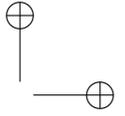
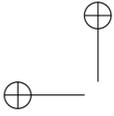
Received XXXX XX, XXXX; editorial decision XXXX XX, XXXX by Editor
XXXXXXXXXXXXXXXX.

1. Introduction

In the new millennium, U.S. equity markets have undergone revolutionary institutional and technological changes. Regulatory changes, including decimalization,¹ combined with massive increases in computer

We are grateful for comments received from Adam Clark-Joseph, Peter Haslag, Tim Johnson, Laurie Krigman, Stefan Negal, Thomas Ruchti, Ronnie Sadka, James Upson, Mao Ye, and seminar participants at the 2017 Financial Management Annual Meeting, the 2016 Midwest Finance Annual Meeting, University of Hong Kong, Hong Kong University of Science and Technology, the University of Illinois, Ohio University, and the Boston Area Finance Symposium. Yashar Barardehi acknowledges support from Paul Boltz Fellowship. Ryan Davies acknowledges support from the Babson Faculty Research Fund.

¹ Goldstein and Kavajecz (2000) and Jones and Lipson (2001) study how decimalization, which reduced the tick size to a penny, affected trading costs.



power have led to algorithmic trading, explosions in sub-second (low-latency) order submissions and cancellations, and large increases in trading volume. These changes transformed the nature of liquidity provision,² and greatly improved many measures of market quality.³ In this paper, we ask: Given these radical changes, what has happened to the liquidity premia that traders demand in return for holding less liquid assets?

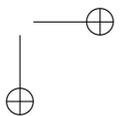
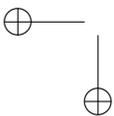
The challenge with addressing this question is that the changes in the trading environment have reduced accuracies of traditional liquidity measures⁴ as proxies of trading costs. Institutional traders now employ sophisticated algorithmic trading strategies that finely split orders over time; thus, traders care about the cumulative cost of the resulting transactions. These trading strategies respond sensitively to variations in extant liquidity, leading to dramatic variations in how fast trading takes place. We develop measures of liquidity that aggregate over trade-times of stock-specific fixed dollar volumes rather than fixed calendar-time intervals. Our measures are easy to construct; and encapsulate the feature that execution algorithms adjust according to available liquidity, relying on more aggressive orders when liquidity is abundant and employing passive orders in less-liquid times. Relative to traditional measures, we establish that: (i) our trade-time based liquidity measures better capture trading costs than standard liquidity measures in a stylized dynamic trading model that incorporates key features of the data; and (ii) our measures are more strongly correlated with estimates of institutional trading costs provided by Investment Technology Group (ITG), a leading provider of institutional trade execution services. In asset pricing tests, our measures outperform traditional liquidity measures in explaining cross-sections of expected raw and risk-adjusted returns, especially in more recent years. We document economically meaningful liquidity premia in 2001–2014, and establish that premia rose significantly after the financial crisis and that they have remained high since.

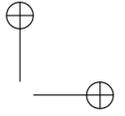
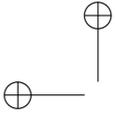
Motivation for our liquidity measures stems from the observation that standard liquidity measures reflect trading environment features that

² Traditional market makers ceased to exist; for example, specialists on the NYSE have been replaced by Designated Market Makers (DMMs) and Supplemental Liquidity Providers (SLPs), who are not subject to negative obligations.

³ Hendershott et al. (2011) show that, especially for large stocks, increases in algorithmic trading improved stock liquidity. Hasbrouck and Saar (2013) find that increased low-latency trading improves market quality measures. Also, O’Hara et al. (2014) and Conrad et al. (2015), among others, find that aspects of market quality such as price discovery and market resiliency, improve as algorithmic trading dominates. Angel et al. (2011) document that improvements in measures of liquidity post-decimalization in 2001 did not continue after implementation of Reg-NMS in 2007.

⁴ For example, spread measures, estimates of Kyle’s λ , or Amihud (2002)’s measure or more generally measures that aggregate trade information over calendar-time intervals.





Trade-time based measures of liquidity

no longer prevail. The minimum penny tick now often binds, making spread measures noisy proxies of trading costs;⁵ and with minimal depth, spreads matter less to investors seeking to establish large positions. The implementation of Regulation ATS in 2000 and Regulation NMS in 2005 led to the growth in alternative lit and dark trading, and a proliferation of order types. Sophisticated algorithms now identify the optimal trading venue, order type (establishing positions by consuming *and* providing liquidity), and order size, breaking “parent” orders into many “child” orders.⁶ The resulting temporal-dependence in individual institutional trades can be exacerbated by high frequency traders (HFTs) who exploit their cumulative price impacts.⁷ Traditional liquidity measures also do not account for the fact that, due to these execution strategies, cumulative trading costs realize over variable time horizons according to how quickly algorithms complete executions of target positions in response to extant liquidity.⁸

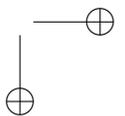
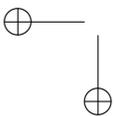
We develop a measure of stock liquidity that (i) addresses temporal dependence of trades, and (ii) accounts for how institutional trading strategies vary with extant liquidity. Our simplest measure, *BBD*, captures average per-dollar absolute returns of fixed-dollar volumes. To construct *BBD*, we first group a stock’s transactions into successive trade sequences whose cumulative dollar volumes correspond to a fixed proportion of stock’s market capitalization at the end of *previous* month. By setting target dollar values proportional to market capitalization, we can aggregate by “enough” to address temporal dependence of transactions, but not by so much as to mix very different levels of trading activity. We also consider a more sophisticated version of *BBD* that accords with the benchmarks set by many execution algorithms. To construct this measure, dubbed *WBBD*, we calculate the price impact of a trade sequence using the difference between the volume-weighted average price (VWAP) of the trades in the sequence and the sequence

⁵ Cross-stock distributions of dollar spreads are truncated at 1c: over 40% of quoted spreads are at the penny minimum. As a result, cross-sectional variations in relative spreads are driven increasingly by cross-stock variation in share prices, and not variations in dollar spreads.

⁶ See O’Hara (2015) who analyzes institutional orders worked by ITG’s order execution algorithms.

⁷ See Hendershott et al. (2011), O’Hara (2015), and Boulatov et al. (2016).

⁸ Moreover, the dominance of low-latency quote submission and cancellation necessitates sub-second time stamps when constructing traditional quote-based liquidity measures such as effective spreads or Kyle’s λ , requiring researchers to deal with massive quoted price data sets. Even if one does this, one must still confront the dilemma that in a highly-fragmented market place, the lack of sub-second connections between trading venues makes it almost impossible to accurately match transaction prices with the most recent prevailing midpoint prices in order to distinguish between liquidity demanders/providers, i.e., to correctly determine “NBBO” at sub-second frequencies (O’Hara (2015)). Moreover, the often “flickering and fleeting” posted quotes only give an indication of the likely price, not a guarantee.



starting price. Constructing *BBD* and *WBBD* uses transaction data that does not need to have sub-second time stamps.

The distinguishing feature of our approach is that, in contrast to traditional measures that aggregate over fixed calendar-time intervals, we aggregate trade information over variable trade-time intervals that correspond to a *fixed dollar value*. We establish that shorter trade-time intervals—indicative of greater trading activity—are associated with higher net order flow imbalances, larger transaction sizes, but *smaller* price impacts (e.g., Kyle’s λ and absolute returns (volatility)). Thus, passive orders are used more often in times of low liquidity, giving rise to low trading activity, balanced order flow, and small transactions when price impacts are large; while aggressive marketable orders are used more in high-liquidity times, producing high trading activity, large order flow imbalances, and large transactions, when price impacts are small. Grouping transactions in this way accords with the feature that the order submission strategies of investors and HFTs adjust to reflect available liquidity. Our method samples price impacts over shorter intervals when liquidity and trading intensity is higher, aligning the sampling of price impacts with the actual intensity of liquidity consumption—trading costs are realized over shorter horizons when liquidity is high, and realized over longer horizons when liquidity is scarce.

BBD can be viewed as a trade-time analogue of Amihud (2002)’s measure that samples price impacts according to intradaily intensities of trading activity. One could instead construct an Amihud measure with price changes and dollar volumes measured over fixed intraday time intervals, e.g., fifteen or thirty minutes. However, such a design assigns equal weight to each fixed time interval regardless of the level of trading activity, and, importantly, it can mix different trading conditions. High trading activity over a fixed time interval may reflect volatile order flow driven by sequential arrivals of large orders on opposite sides that offset each others’ price impacts. This feature manifests itself in the smaller signed trade imbalances when volume is higher in a calendar-time interval. By lumping together orders with opposing price impacts in active markets, calendar-time aggregation may lead to underestimated price impacts. Trade-time aggregation is less prone to such over-aggregation in active times because it is volume-indexed, and samples price impacts more often in more active markets. This reduced over-aggregation manifests itself in the large signed trade imbalances over trade-time intervals in active markets.

To convey why this makes the trade-time cumulative price impacts (absolute returns) used to construct *BBD* less noisy proxies of trading costs than price impacts measured over calendar-time intervals, we develop and simulate a dynamic model of trading. This model

Trade-time based measures of liquidity

also provides insights into the considerations entering the choice of aggregation horizon. In our model (a) institutional trading drives order flow imbalance, and comprises a significant portion of trading volume; (b) there is persistence in order size and sign; (c) supply of liquidity stochastically varies over time; and (d) public information arrival can move prices. In these simulations, *BBD* always outperforms Amihud’s measure—and spreads-based measures—in capturing trading costs. The simulations show that when trading activity is high, calendar-time aggregation tends to bundle together large trades/parent orders in opposing directions within the same window, resulting in offsetting price impacts and under-estimates of trading costs, even in a setting where institutional trade is completely non-strategic. In contrast, when markets are active, trade time measures tend to allocate large parent orders to different trade sequences, separating out the distinct price impacts, and more accurately capturing trading costs.

To underscore that our measures are good proxies of institutional trading costs, we explore the cross-stock correlations between different liquidity measures and estimates of institutional trading costs from ITG. We find that *BBD* and *WBBD* are *highly* correlated with ITG’s ACE (Agency Cost Estimator) estimates of institutional trading costs *regardless* of the parent order size. Averaging ACE across all parent order size bins reveals that *BBD* and *WBBD* have very high correlations of 0.82 and 0.83 with this measure of institutional trading costs that exceed those for other liquidity measures. Relative quoted spread is more strongly correlated with ACE measures for very small, retail-type, order sizes than are our measures, but this correlation drops off sharply for larger orders; and the standard Amihud’s measure has a marginally higher correlation with ACE for very large and rare order sizes than do *BBD* and *WBBD*, but its correlation is lower for smaller and mid-sized orders.

We then investigate which liquidity measures best explain cross-sections of expected returns and risk-adjusted expected returns among NYSE-listed stocks between 2001 and 2014. We first establish that most measures, including ours, explain these cross-sections. That is, consistent with findings for earlier eras (e.g., Amihud and Mendelson (1996), Brennan and Subrahmanyam (1996), Chordia et al. (2000), Amihud (2002), Hasbrouck (2009)), lower characteristic liquidity is still associated with larger expected returns: despite the improvements in market quality, measures of expected trading costs still have explanatory power for expected returns. The estimated premia are economically meaningful: for example, relative to a perfectly liquid idealized security, the median monthly liquidity premium estimates using *WBBD* are 5.3 basis points for expected returns and 2.4 basis points for risk-adjusted expected returns. These results accord with Asparouhova et al. (2010,

2013)’s finding of positive liquidity premia for U.S.-based stocks in the 2001–2006 period; although they contrast with the finding of Ben-Rephael et al. (2015), who were unable to discern *any* liquidity premia on NYSE-listed stocks in recent years.⁹

We then show that *BBD* and *WBBD* better explain cross-sections of (risk-adjusted) expected returns than do traditional measures. To do this, we first obtain the orthogonal component of our two measures with each alternative measure (low- and high-frequency Amihud measures, relative quoted and effective spreads, estimates of Kyle’s λ), and vice versa. We then use these orthogonal components (residuals) separately as the explanatory variable in an asset pricing framework. The orthogonal component of our measures has additional explanatory power, whereas those of the other measures do not. The superior performances of *BBD* and *WBBD* vis à vis other measures are quite pronounced in more recent years, and they are not driven by any single year or two years of observations.

Having established that our measures better capture trading costs and better explain the cross-sections of returns, we explore the evolutions of liquidity premia based on them. We find that liquidity premia rose significantly during and after the financial crisis, and remained consistently high in post-crisis years. Concretely, monthly liquidity premia on expected returns based on *WBBD* rose from 3.8pb in 2001–2007 to 4.6bp in 2008–2014; but focusing on the post-crisis years of 2010–2014, we find liquidity premia as large as 6.3bp. Thus, despite the improvements found in measures of market quality, investors still demand nontrivial compensation in return for holding less liquidity stocks.

The paper is organized as follows. Section 2 summarizes the related literature. Section 3 explains explains the notion of trade time and develops our liquidity measures. Section 4 describes the data. Section 5 details the theoretical and empirical characteristics of our trade-time measures. Section 6 uses these measures in an asset pricing framework. Section 7 concludes. An appendix provides various robustness analyses.

2. Related literature

We now relate our contribution to existing research on trade-time methodologies, and to alternative empirical measures of liquidity.

Trade time. Our methodology contributes to the literature that maps equity market dynamics onto trade time space. Dufour and Engle

⁹ Their finding likely reflects the use of a variant of Amihud’s measure constructed at *annual* frequencies, and that the measure obtained from a year is used to capture expected trading costs for *all* monthly return observations of the following year, resulting in measurement error that biases their estimates toward zero.

(2000) examine the duration of time between successive trades, showing that price adjustments are faster when trading intensity is high (i.e., durations are short). Engle and Russell (1998) and Engle (2000) develop related models of autoregressive conditional duration (ACD), which provide semi-parametric estimates of trade intensities based on trade arrival rates. ACD models are designed primarily to estimate the expected cost and time to execute a single order, submitted within a very short time frame, rather than the cost of a series of orders submitted over longer periods of time.

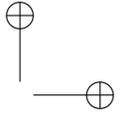
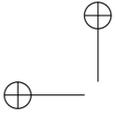
Gouriéroux et al. (1999) use the time to sell or buy a predetermined volume or value of stocks to study time-of-day patterns in liquidity on the Paris Bourse. While Gouriéroux et al. (1999) and our paper both analyze trade sequences of a given value, the focuses differ: they focus solely on the durations (trade time) of trade sequences, whereas we focus on the price changes that occur over these trade sequences. In addition, to deal with small sample issues, they study sequences of trades that begin at every transaction. Thus, there is significant overlap in their trade sequences, whereas our trade sequences are mutually exclusive.

To address temporal dependence of transactions driven by dynamic order-splitting strategies, other researchers have developed alternative trade-time measures. Easley et al. (2012) group consecutive trades across time to obtain a certain number of fixed-volume groups (e.g., 50) per trading day.¹⁰ Within each group, the sign of price changes over one-minute periods is used to infer buyer- versus seller-oriented trades, creating a measure of volume “toxicity”, VPIN. Feldhütter (2012) measures trading activity over periods of time for corporate bonds. He develops a notion of an imputed round-trip trade, based on the empirical feature that trades in corporate bonds often occur infrequently, but when they do trade, there are often several trades of the same size in a short period of time. He argues that these trades are likely linked and should be grouped together when measuring trading costs.

Measures of liquidity. Our study is the first to develop measures of stock liquidity based some notion of trade time. To further motivate the relevance of trade-time based measures, we next overview traditional liquidity measures and their use in empirical asset pricing, highlighting features that make these measures less informative in modern U.S. equity markets.

Quote-based measures have often been used for asset-pricing purposes. Amihud and Mendelson (1996) show that when the standard tick size was one-eighth, bid-ask spreads were priced and were reasonable measures of liquidity. Other quote-based liquidity measures include

¹⁰ Cross-day comparisons/aggregations of such volume buckets may be problematic because daily trading volumes are highly variable.



effective bid-ask spreads, Roll (1984)’s measure, and Hasbrouck’s Gibbs estimate.¹¹ Researchers have also estimated price impacts using variants of Kyle’s λ (Kyle (1985)) as measures of liquidity (see Glosten and Harris (1988), Brennan and Subrahmanyam (1996), or Pástor and Stambaugh (2003)); Bernhardt and Hughson (2002) structurally estimate price impacts using a model whose equilibrium explicitly incorporates that orders are not pooled.

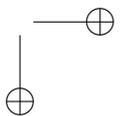
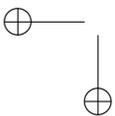
Chordia et al. (2011), Angel et al. (2011), Kim and Murphy (2013), Holden and Jacobsen (2014), and others have raised questions about the accuracy of existing liquidity measures in a low-latency trading world with tiny spreads and little depth.¹² Today, liquidity concerns do not revolve around avoiding minimum price ticks, so spreads do not capture liquidity from that perspective. Algorithmic trading featuring both consumption and provision of liquidity, and dynamic order-splitting combine to make consecutive trades dependent and complicate interpretations of estimates of λ . In fact, recent evidence suggests that estimates of λ may not be consistent with adverse selection interpretations: Collin-Dufresne and Fos (2015) find that smaller price impacts on days where insiders trade, and Barardehi and Bernhardt (2017) find that estimates of Kyle’s λ *fall* both over the trading day, and as trading activity *rises*.¹³ One would expect the opposite if estimates of Kyle’s λ reflected pricing of asymmetric information.

Deuskar and Johnson (2011) propose a liquidity measure that uses information from the cumulative depth in the limit order book, rather than just the inside quote. However, their measure imperfectly captures dynamic order splitting. The use of hidden orders makes it difficult to accurately measure depth, and over 25% of trade occurs off exchange. Today’s high volume also means that instances of zero returns are rare, precluding their use as a liquidity measure as in Mazza (2015).

¹¹ Effective bid-ask spreads measure the spread between the trade price and quote midpoint. Roll (1984) shows how the tick size induces negative correlation in prices. Hasbrouck (2009) proposes a Bayesian platform to generate a Gibbs estimate of Roll’s effective trading cost indicator.

¹² Holden and Jacobsen (2014) document realized dollar spreads and share depth of 0.88¢ and 564 shares, respectively, for a sample of randomly selected stocks in 2008. These quantities have fallen further since then.

¹³ Barardehi and Bernhardt (2017) also find that higher trading activity (shorter time durations of trade sequences) is, on average, associated with (i) *smaller BBD* (and *WBBD*) magnitudes (reduced price impact measures and return volatility), but (ii) *larger* average signed trade imbalances, and (iii) *larger* average trade sizes. These findings indicate that short-term variations in liquidity provision are the primary drivers of variations in trading activity and price impacts. They also find that price impacts of fixed-dollar positions fall steadily over the trading day (by about 50%!), even though trading volume is U-shaped.



Among low-frequency measures, Amihud’s measure remains widely used to study characteristic and systematic liquidity (risk) in cross-sectional and time series settings.¹⁴ Goyenko et al. (2009) show that Amihud’s measure remains useful post decimalization, in part because at the aggregate daily level, dynamic order splitting issues are reduced. This advantage comes at the cost of lost information and reduced accuracy. We regain some of this information and associated explanatory power, by developing a *trade-time* analogue to Amihud’s measure.

3. Trade-time measures of stock-specific liquidity

To start, we explain how we measure price impacts of fixed dollar volumes. For each stock, these price impacts are measured over successive sequences of transactions, with all sequences corresponding to a roughly fixed dollar volume, but spanning different time horizons due to variations in trading activity.

First, we define a **trade sequence**. Each year, we number trades in stock j sequentially, using index n_j . For trade n_j , we use $\tau_j(n_j)$, $Q_j(n_j)$ and $P_j(n_j)$ to denote respectively, (i) its time measured in seconds from the beginning of the year, (ii) its size (in shares), and (iii) its price. A trade sequence consists of consecutive transactions that have an aggregate value of at least $V_{j,t}$ for stock j in month t . Thus, a shorter time duration indicates a more active market. We set $V_{j,t}$ to be proportional to stock j ’s market capitalization at the end of the *previous* month, $M_{j,t-1}$. The first trade sequence begins with the first trade of the year, and each subsequent trade sequence begins with the first trade following the previous sequence. Figure 1 illustrates the patterns of dollar volume, cumulative dollar volume, and prices, identifying trade sequences for a typical stock.

Formally, we iteratively solve for the last trade of the k^{th} trade sequence, $k = \{1, 2, 3, \dots\}$, as:

$$n_j^k = \underset{n^*}{\operatorname{argmin}} \left\{ \sum_{n=n_j^{k-1}+1}^{n^*} P_j^C(n) \times Q_j(n) \mid \sum_{n=n_j^{k-1}+1}^{n^*} P_j^C(n) \times Q_j(n) \geq V_{j,t} \right\}, \quad (1)$$

where $n_j^0 = 0$ and the value of aggregate trades is measured using the *previous* day’s closing price, $P_j^C(n_j)$.¹⁵ We use the previous day’s closing

¹⁴ See Jones (2002), Chordia et al. (2000), Sadka, R. (2006), Acharya and Pedersen (2005), and Akbas et al. (2011), among others.

¹⁵ The last quoted bid-ask midpoint is used when the closing price is not available.

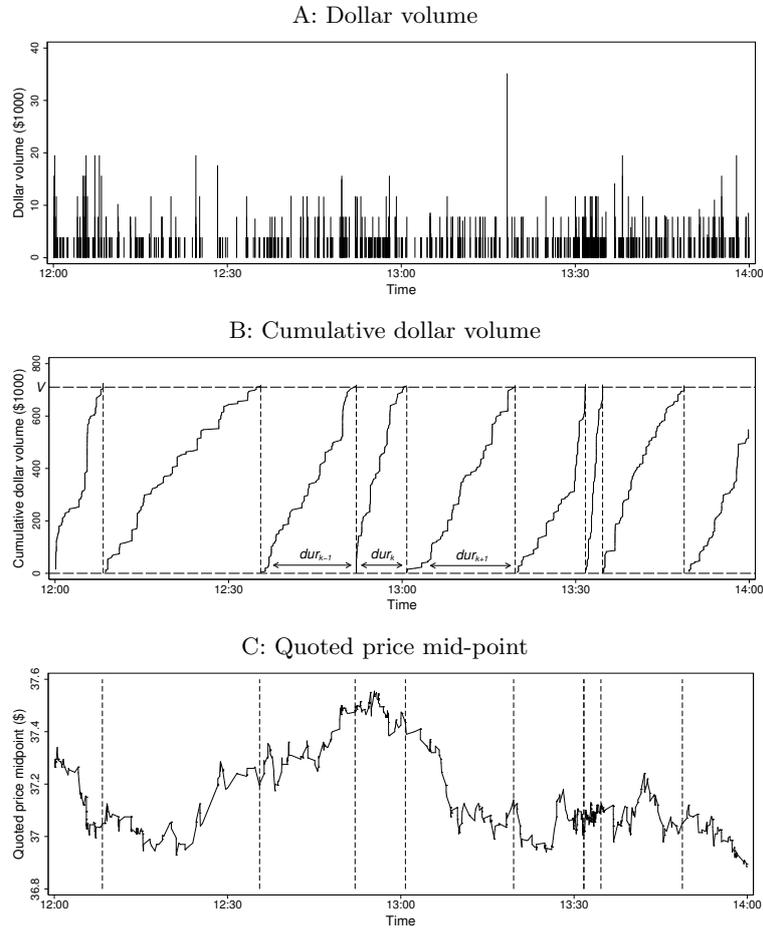


Figure 1
Illustration of trade sequence construction.
 Panel A presents the dollar volume for a mid-sized stock, Acuity Brands Inc. (AYI) on September 22, 2011 over a two hour period. Panel B shows how we aggregate trades into a trade sequence until cumulative dollar volume $V_{j,t}$ is reached. As dollar volumes vary over time in Panel A, durations of the associated trade sequence vary in Panel B. Panel C shows evolutions of quoted price mid-point associated with transactions.

price to calculate dollar volumes to ensure that contemporaneous price movements do not alter identification of trade sequences.¹⁶

¹⁶ Using the previous day’s closing price avoids introducing biases driven by current price movements. For example, with rapidly increasing prices, using today’s prices can give rise to non-trivially growing dollar volumes, causing a downward bias in the time duration of the corresponding trade sequence.

Trade-time based measures of liquidity

We construct trade sequences that span two trading days, but *exclude* them from the analysis. By calculating overnight trade sequences and then excluding them, we: (1) deliver a random starting point for the first trade sequence of a given day, precluding any systematic bias; (2) circumvent issues associated with overnight price adjustments or information arrival;¹⁷ and (3) avoid combining trading activity levels from near close with those just after open, which typically differ.

We then obtain the time duration, i.e., *trade time*, of the k^{th} trade sequence,

$$dur_j(k) = \tau_j(n_j^k) - \tau_j(n_j^{k-1} + 1), \quad (2)$$

the corresponding VWAP,

$$VWAP_j(k) = \frac{\sum_{n=n_j^{k-1}+1}^{n_j^k} \frac{P_j(n) \times Q_j(n)}{\sum_{n=n_j^{k-1}+1}^{n_j^k} Q_j(n)}, \quad (3)$$

return,

$$r_j(k) = \frac{P_j(n_j^k)}{P_j(n_j^{k-1} + 1)} - 1, \quad (4)$$

volume-weighted return,

$$wr_j(k) = \frac{VWAP_j(k)}{P_j(n_j^{k-1} + 1)} - 1, \quad (5)$$

and dollar value of the sequence,

$$DVOL_j(k) = \sum_{n=n_j^{k-1}+1}^{n_j^k} P_j^C(n) \times Q_j(n). \quad (6)$$

The per-dollar price impact for stock j of the k^{th} trade sequence is given by

$$DIMP_j(k) = \frac{|r_j(k)|}{DVOL_j(k)}. \quad (7)$$

We divide by $DVOL_j(k)$ rather than $V_{j,t}$ because the size of the last trade, n_j^k , may marginally exceed the level needed to obtain a total value of $V_{j,t}$.

Let $\mathbf{K}_{j,t}$ be the set of eligible trading sequences for stock j in the 3-month period ending on the last day of month t , and let $\#\mathbf{K}_{j,t}$ be the number of trade sequences in this set. Our simplest **trade-time measure of stock-specific liquidity** for stock j in month t , is:

$$BBD_{j,t} = \frac{\sum_{k \in \mathbf{K}_{j,t}} DIMP_j(k)}{\#\mathbf{K}_{j,t}}, \quad (8)$$

¹⁷ For instance, we do not need to adjust $r_j(k)$ for stock splits or dividend distributions.

i.e., the mean per-dollar price impact for stock j in month t associated with the fixed dollar position, $V_{j,t}$. To capture better the realized cost of trading over the sequence, we also consider the VWAP-based analogue of our per-dollar price impact measure, calculated as the absolute volume-weighted return of a sequence scaled by its dollar volume:

$$WDIMP_j(k) = \frac{|wr_j(k)|}{DVOL_j(k)}. \quad (9)$$

We form our **VWAP-based trade-time measure of stock-specific liquidity**, $WBBD_{j,t}$, by averaging VWAP-based price impacts across all trade sequences for stock j in month t :

$$WBBD_{j,t} = \frac{\sum_{k \in \mathbf{K}_{tj}} WDIMP_j(k)}{\#\mathbf{K}_{tj}}. \quad (10)$$

$WBBD_{j,t}$ modifies $BBD_{j,t}$ to account for the realized cost, relative to the original price, of participating in trades at the actual trade sizes and prices observed over a trade sequence. Thus, it picks up components of commonly-used implementation shortfall measures of institutional trading costs. In this sense, when used in an asset pricing framework, $WBBD$ captures realized costs and price impact, whereas BBD primarily captures price impact.

In our subsequent analysis, we compare BBD and $WBBD$ to five alternative measures of stock-specific liquidity:¹⁸

1. **Amihud:** The standard (low-frequency) Amihud measure ($AML_{j,t}$) for stock j in month t is the rolling average of the daily per-dollar price impact (measured as the absolute value of the stock’s realized daily return divided by its daily dollar volume) during the 3-month period ending on the last day of month t .
2. **High-frequency Amihud:** Because $AML_{j,t}$ aggregates over a trading day, it does not exploit informative intraday variations in trading activity and prices. This leads us to consider a high-frequency variant of AML based on observations of price changes and dollar volumes from thirty-minute intervals, to address possible concerns about excessive temporal aggregation in the low-frequency measure. We denote the high-frequency (intraday) analogue by $AMH_{j,t}$. To calculate it, we first compute the per-dollar price impacts for a stock as the absolute realized returns divided by the corresponding executed dollar trading volumes over

¹⁸ In unreported results, we construct various versions of BBD , using observations from different time-of-day windows (e.g., near market open, middle of the day, and near market close) and trading activity levels (time duration of a trade sequence is inversely related to trading activity). Each version of BBD explains variations in expected returns, and all versions deliver qualitatively similar liquidity premia.

Trade-time based measures of liquidity

thirty-minute intervals.¹⁹ The high-frequency Amihud measure for stock j in month t is given by the rolling average of the per-dollar price impacts in the 3-month period ending on the last day of month t .

3. **Percentage bid-ask spreads:** Average percentage bid-ask spreads ($PSP_{j,t}$) are calculated using a time-weighted average of spreads based on the NBBO at a one-second frequency, on a daily basis. To be consistent with the other measures, the average percentage spread in month t for stock j , denoted by $PSP_{j,t}$, is given by the rolling 3-month average of spreads in months t , $t-1$ and $t-2$.
4. **Kyle’s λ :** The price impact measure λ based on Kyle (1985) is estimated following the approach of Hasbrouck (2009).²⁰ For a given stock, λ is estimated as the slope coefficient from a regression of five-minute stock returns, based on quote midpoints, on the signed square-root of dollar volume over the corresponding five-minute period. Each trade is signed using the Lee-Ready approach. For each stock j , we use an average of the monthly regression coefficients for months t , $t-1$, and $t-2$ to estimate $LAMBDA_{j,t}$.
5. **Dollar-weighted percentage effective spreads:** The effective spread of a trade is twice the absolute value of the difference between the trade price and the prevailing NBBO midpoint. Percentage effective spreads scale the effective spread by the midpoint price.²¹ Our dollar-weighted percentage effective spreads measure for stock j in month t , denoted by $EPSP_{j,t}$, is given by the value-weighted average of percentage effective spreads across all transactions executed during regular market hours for stock j in the 3-month period ending the last day of month t .²²

We scale the *BBD*, *WBBD*, *AML*, and *AMH* measures by a million for presentational purposes. In Appendix 4, we show that qualitatively identical findings obtain when we construct measures of stock liquidity using observations from the current month, rather than using three-month rolling averages.

¹⁹ We exclude the very few thirty-minute intervals with zero trading volume.

²⁰ See Goyenko et al. (2009) for additional details.

²¹ Effective spreads are obtained using a modified version of the SAS code provided by Professor Craig Holden on his website (<https://kelley.iu.edu/cholden/>), following the approach described in Holden and Jacobsen (2014).

²² Similar results obtain using percentage effective spreads, based on simple averages and share-weighted averages. We also find that trade-time measures outperform *dollar* effective spreads (calculated based on simple averages, dollar-weighted averages, or share-weighted averages) in explaining expected returns.

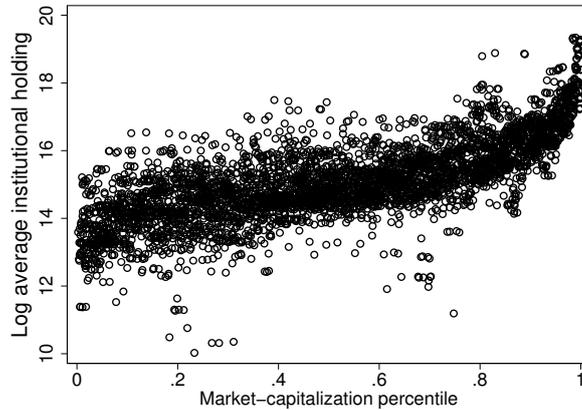


Figure 2
Institutional ownership and firm size.
 Natural log of average institutional ownership of top-10 institutional investors against market-cap percentile. Sample includes common NYSE-listed stocks in 2012, obtained from Thompson Reuters 13-F database.

3.1 Choice of $V_{j,t}$

We focus on the time to execute a sequence of trades that for a firm j have an aggregate value $V_{j,t}$ equal to $\theta\%$ of its market capitalization at the end of the *previous* month, $M_{j,t-1}$. Our base formulation of $\theta=0.04$ generates a median duration (across stocks and years) of about 30 minutes. Several considerations enter our base specification: (i) the positions are large enough to control for dynamic order splitting and division of orders against the book; (ii) the positions are not so large that a single trade sequence spans different activity levels; (iii) the proportional specification facilitates cross-stock comparisons, and it accords with the feature that institutional traders tend to establish larger positions in larger firms (see Figure 2); (iv) the proportional specification flattens out the distribution of trade time across market capitalizations (see Figure 3). As a result, trade sequence durations of small market capitalization stocks are not too much longer, delivering enough observations for our analysis. To verify that $\theta=0.04$ aggregates sufficiently to control for the positive autocorrelation induced by dynamic order splitting, we estimate stock-specific return autocorrelations over successive trade sequences, establishing that the mean autocorrelation is marginally negative and insignificantly different from zero (see Appendix 1).

Trade-time based measures of liquidity

An unreported robustness analysis verifies that our qualitative findings are not driven by the choice of θ : similar findings obtain using $\theta=0.03$, $\theta=0.05$ or $V_{j,t}=0.00025M_{j,t-1}+\$80,000$.²³ To reinforce the fact that our findings are not driven by cross-sectional variation in market-capitalizations, we also construct liquidity measures based on dollar volumes that are the same for a group of stocks in a given month. Using the same fixed dollar-amount based on a medium-sized stock’s market capitalization for *all* stocks would deliver too few observations for small-cap stocks, and too many for large-caps. This leads us to create size portfolios of small, medium and large stocks each month and use $V_{j,t}=0.0003M_{g,t-1}^{med}$ as the target dollar value, where $M_{g,t-1}^{med}$ is the median market-cap in size portfolio g in month $t-1$.²⁴ The high correlations of liquidity measures based on this specification with *BBD* of 0.90, 0.92, and 0.95 for portfolios of small, mid-sized, and large market-capitalizations indicate that it is the trade-time approach, and not the variations in market-capitalizations entering via $V_{j,t}$ or the choice of θ that drive cross-sectional variations of *BBD*.

4. Data

Our sample period runs from April 1, 2001 to December 31, 2014. Each year, we construct a sample of U.S.-based NYSE-listed common shares (CRSP share codes 10 and 11). Our selection of stock characteristics is similar to those in Amihud (2002) and Ben-Rephael et al. (2015). We estimate $\beta_{j,t}^{mkt}$, $\beta_{j,t}^{smb}$, $\beta_{j,t}^{hml}$, and $\beta_{j,t}^{umd}$ of a four-factor Fama-French model for each stock j in month t . We regress excess weekly stock returns against the weekly Fama-French factors, over the previous two-year period ($t-24$ to $t-1$).²⁵ If a stock is not listed over the entire previous 24-month period, we employ the longest time period available, requiring at least one year of weekly observations. Firm j ’s book-to-market ratio, $BM_{j,t-1}$, is given by its most recent reported book value per share, obtained from Compustat, divided by its share price, at the end of month $t-1$.²⁶ We measure stock-specific idiosyncratic volatility, $SD_{j,t-1}$, using the standard deviation of monthly returns over the

²³ The fixed amount of \$80,000 captures a minimum institutional parent order size, and ensures that the positions are large enough to control for dynamic order splitting.

²⁴ To do this, we focus on stocks with market capitalization ranks between 101 and 1000, so that large stocks have rank 101-400, medium stocks have rank 401-700 and small stocks have rank 701-1000. We drop the top 100 stocks because they differ so radically in size and daily trading volume that a fixed dollar volume again delivers too many trade sequences for the largest stocks or too few for the smaller ones.

²⁵ Weekly observations on factors are obtained from Kenneth French’s website.

²⁶ We exclude observations with negative book values (consequently, negative book-to-market ratios).

preceding sixty months, requiring at least one year of data.²⁷ We use the dividend yield, $DYD_{j,t-1}$ over the past twelve months as a predictor of expected returns. We construct two momentum metrics: the compound monthly returns over the preceding four months, $MOM_{j,t-1}^{1-4}$, and the compound returns over the eight months before that, $MOM_{j,t-1}^{5-12}$.

We obtain trade prices, quantities, and time stamps from the consolidated trade history, from January 1, 2001 to December 31, 2014, in the NYSE Monthly TAQ database. We consider all stock trades on U.S.-based trading venues during regular market hours between 9:30AM and 4:00PM (EST), as well as trades flagged as market-on-close orders (which can be recorded after 4:00PM). After constructing trade sequences, we exclude those with absolute (volume-weighted) returns that exceed 10%. We calculate percentage bid-ask spreads using the National Best Bid and Offer (NBBO) series on WRDS. We merge these three databases using the variable NCUSIP from CRSP and the variable CUSIP from TAQ and Compustat. We drop stocks missing this identifying information in all three databases. We use the NYSE Daily TAQ database to obtain trade and quote information at millisecond frequencies for the calendar year 2012 in an analysis of intradaily variations in liquidity and trading costs, featuring estimates of Kyle’s λ , quoted spreads, and effective spreads.²⁸

We apply additional filters in the following order. First, we drop a stock-month observation if it had a closing price below \$1.00 on any trading day in the corresponding year.²⁹ As in Amihud (2002), we trim the top 1% of observations of the liquidity measure(s) each month to avoid spurious results driven by the high variability of extreme observations.³⁰ We also drop observations whose estimates of $\beta_{j,t}^{mkt}$, $\beta_{j,t}^{smb}$, $\beta_{j,t}^{hml}$, or $\beta_{j,t}^{umd}$ take extreme values falling in the top or bottom 0.1% of the monthly cross-sections. Finally, to preserve a similar composition of stocks over time, each monthly sample comprises the 1,050 largest stocks based on market capitalization at the end of

²⁷ In an unreported robustness analysis, we employ four alternative measures of *daily* volatility to control for fundamental volatility. We measure close-to-close daily returns based on both transaction prices and average bid-ask at close to identify possible contribution of bid-ask bounce. We calculate daily volatility based on each type of return, once using observations from the previous month and once using observations from the preceding three months—the latter is more consistent with constructions of *BBD* and *WBBD*, in terms of sampling horizon. Replacing $SD_{j,t-1}$ with any of these alternative measures has virtually *no* impact on our estimates of liquidity premia. These findings are available upon request.

²⁸ Holden and Jacobsen (2014) highlight the value added of using sub-second time stamps over one-second time stamps for such analyses of trading outcomes.

²⁹ The average of the closing bid and ask prices is used if a last trade price is not available.

³⁰ Amihud (2002) also trims the bottom 1%; this additional trimming has no qualitative effects on estimates reflecting the positively-skewed liquidity measure distribution. Similar estimates also obtain with alternative trimming thresholds (e.g., top 2% or top 5%).

Trade-time based measures of liquidity

Table 1
Year-by-year summary statistics of the final sample.

Year	Share Price	Book-to-Market	ln(Market Cap)	\$.spreads	# of Sequences
2001	25.80	0.49	20.98	0.070	1829
2002	25.00	0.50	21.03	0.051	2130
2003	24.90	0.52	21.10	0.033	2451
2004	29.76	0.45	21.33	0.030	2669
2005	33.10	0.43	21.51	0.029	2965
2006	35.57	0.40	21.64	0.030	3570
2007	36.77	0.40	21.74	0.029	4518
2008	27.41	0.51	21.52	0.031	5858
2009	22.06	0.66	21.20	0.023	5802
2010	27.42	0.55	21.48	0.020	4962
2011	31.14	0.50	21.67	0.022	4808
2012	32.15	0.51	21.63	0.024	4079
2013	39.10	0.44	21.86	0.028	3795
2014	43.21	0.40	22.05	0.033	3569

Closing share prices, book-to-market ratios, and natural log of market capitalizations are reported as the medians of their monthly values. \$.spreads are the median time-weighted dollar bid-ask spread at the NBBO. The last column reports the median number of trade sequences (each with an aggregate value of at least $V_{j,t}$) in a year across stocks.

the previous month—this is the largest monthly cross-section size that could be maintained over the period studied. The stable sample size and composition over time facilitates analysis of the temporal evolution of liquidity premia. Were we, instead, to include all firms in a year with “enough” trade sequences, the sample size would grow in recent years due to the increases in trading activity, tilting the sample composition in later years toward less actively traded (and typically smaller) firms.³¹

Table 1 presents annual summary sample statistics of the final sample. Median stock prices largely mimic overall market performance. The median annual number of trade sequences rises more than twofold over the sample period, peaking at 5,858 in 2008 due to the lower market capitalizations after the financial crisis and the higher trading volumes. Even though median market capitalizations are higher in later years than in 2001 (so the median $V_{j,t}$ rises), the median numbers of trade sequences are far higher in later years due to the massive growth in trading volumes. Median dollar-spreads fall from 7¢ in 2001 to 3.3¢ in 2014.

5. Characteristics of Trade-time Measures

We first show how our measure of trading activity—the trade times of trade sequences—vary over time and with firm size. Figure 3 reveals that trade time and firm size have a U-shaped relationship in any given month. Larger stocks are more actively traded, but choosing

³¹ Also, the sample of firms with “enough” observations shrinks during the 2008–2009 financial crisis because many firms were delisted or failed to maintain the minimum daily closing price of \$1. A varying sample size would introduce systematic temporal sample selection patterns, complicating analyses and interpretations of temporal changes in liquidity premia.

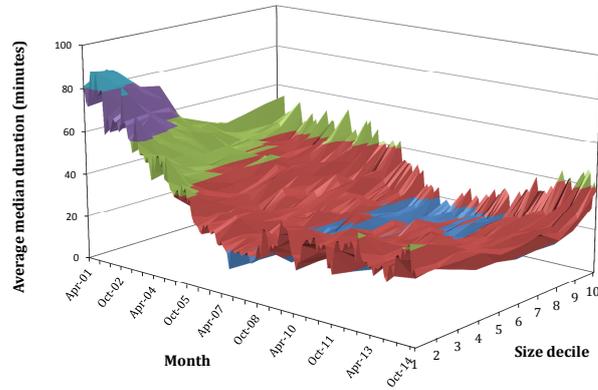


Figure 3
Trading activity by year and by firm size.
 Average median trade time (in minutes) of a sequence of trades with a cumulative aggregate value of at least 0.04% of a firm’s market capitalization. Median durations are calculated on a stock-by-stock basis. The average is then computed for stocks in a size decile each month. Size decile 1 contains the smallest firms; decile 10 contains the largest.

$V_{j,t}$ to be proportional to market-capitalization offsets much of the heterogeneity in trading volumes, flattening the relationship between trade time and market capitalization, and delivering similar numbers of trade sequences for stocks with different market capitalizations. The longer trade times for the largest stocks are driven by the large sizes of their target dollar values. The longer trade times for small stocks reflect their modest trading activities. Overall, except for the earliest years of the sample, the number of trade sequences does not vary greatly with market capitalization.

Trade times tend to fall in early years, reflecting increases in trading activity. The average median trade time for the second size decile of stocks falls from 86 minutes in April 2001 to 30 minutes in December 2014; while that for the largest stocks goes from 56 to 28 minutes. Trade times are shortest in 2008–2009 due to the lower market valuations in those years. This highlights the importance of controlling for fixed month effects by calculating a stock’s relative activity on a monthly basis, instead of targeting a dollar value that is fixed over time.

Does BBD measure liquidity? We begin by providing evidence that *BBD* and *WBBD* do, in fact, measure liquidity. We first document this at the individual stock level. We show that as trading activity for a given stock rises, i.e., trade times shorten—an indication of increased liquidity provision for that stock—the levels of price impact measures fall sharply.

Trade-time based measures of liquidity

We contrast these patterns with their very different calendar-time analogues; and detail what they mean for trade-time and calendar-time measures of trading costs. We then develop and simulate a stylized model of trading to provide deeper intuition as to why trade-time measures of liquidity capture trading costs better than calendar-time measures. We next establish that correlations in the cross-section between our liquidity measures and trading cost measures for institutional investors estimated by ITG significantly exceed those for the commonly-used liquidity measures.³² Finally, we document that the temporal evolutions of *BBD* and *WBBD* over our sample period are very similar to those of other commonly-used measures of liquidity, such as Kyle’s λ or Amihud’s measure, and that the correlations between these liquidity measures are quite high.

Our analysis of intraday stock-level trade data begins by showing that variations in the trade time of a fixed-dollar value capture variations in liquidity provision, and that our cumulative price impact measures strongly and meaningfully vary with extant liquidity. To measure trading activity, we use the inverse of the time duration of a trade sequence with cumulative dollar value $V_{j,t}$ —shorter trade times indicate higher trading activity. For each stock j in month t , we sort trade sequences from longest to shortest (by trade time) into *deciles of trading activity*. We then examine how trading outcomes vary with trading activity level.³³

For stock j , we calculate the following trading outcomes over each trade sequence k :

- absolute return, $|r_j(k)|$: absolute value of the return realized over a trade sequence. With roughly fixed dollar volumes across different trade sequences of a stock j in month t , variations in $|r_j(k)|$ are closely related to variations in *BBD* measure, had we calculated the measure by trading activity level.
- Transaction size, $ts_j(k)$: the average transaction size in a trade sequence.
- Trade imbalance, $imb_j(k)$: the proportion of buyer- or seller-initiated dollar volume, whichever is highest, out of total dollar

³² *BBD* and *WBBD* also evolve more similarly over time with measures of institutional implementation shortfall obtained from Abel Noser Solutions’ institutional trades data. We do not report these results because these data are proprietary, and Abel Noser no longer makes them available to academics.

³³ We restrict our sample to 2012 to exploit the milli-second time stamps in the Daily TAQ database. Milli-second time stamps drastically reduce errors in trade classifications from the high levels documented by Holden and Jacobsen (2014) for the Monthly TAQ database.

volume of a trade sequence. We classify trades using Lee-Ready algorithm.³⁴ Because dollar volumes associated with trade sequences of a stock are roughly fixed in each month, we use the *proportion* of buyer- (or seller-) initiated dollar volume to measure net order flow.

- Quoted spread, $sp_j(k)$: trade-weighted average of bid-ask spreads at the times of transactions.
- Relative effective spread, $eps_pj(k)$: trade-weighted average of the absolute difference between transaction price and the mid-point price at the end of the previous second relative to the mid-point price.

We also estimate for each stock j a trade-time version of Kyle’s λ by trading activity decile. For trade sequence k of stock j , we define net order flow $nof_j(k)$ to be the signed square root of the absolute difference between the proportion of *buyer-initiated* dollar volume and 0.5 (which represents perfectly balanced order flow), assigning a positive sign if the proportion of buyer-initiated dollar volume exceeds 0.5, and a negative sign if the opposite holds.³⁵ We then regress $r_j^m(k)$, the return based on the prevailing mid-point prices at the beginning and end of trade sequence k of stock j , on $nof_j(k)$ by trading activity level, controlling for month-fixed effects, obtaining an estimate of Kyle’s λ at each trading activity level.

Figure 4 reveals that greater trading activity is associated with more abundant liquidity provision. In particular, higher trading activity is associated with both more aggressive trading—transaction sizes and signed order trade imbalances both rise—and reduced price impacts. Concretely, as trading rises from the bottom decile of activity to the top, median trade imbalance rises from 0.56 to 0.63, median transaction size jumps from 130 to 180 shares, while the median absolute return plunges by almost 50% from 23bp to 13bp.³⁶ This combination indicates that abundant liquidity results in high trading activity, driven by the endogenous choices of traders to aggressively and quickly consume the liquidity by submitting larger marketable orders—resulting in short trade sequence time durations, high signed trade imbalances, large transactions and small price movements. Conversely, when liquidity is limited, traders respond by relying more on passive orders, trying to

³⁴ In unreported results, we verify that our findings are robust to use of alternative trade classification algorithms, including those proposed by Ellis et al. (2000) and Chakrabarty et al. (2006).

³⁵ Results are robust if we use levels to calculate net order flow, rather than proportions.

³⁶ Barardehi and Bernhardt (2017) thoroughly examine intradaily patterns in trading activity and various trading outcomes, showing that findings documented here are robust to controlling for time of day and major informational events such as earnings announcements and analyst recommendations.

Trade-time based measures of liquidity

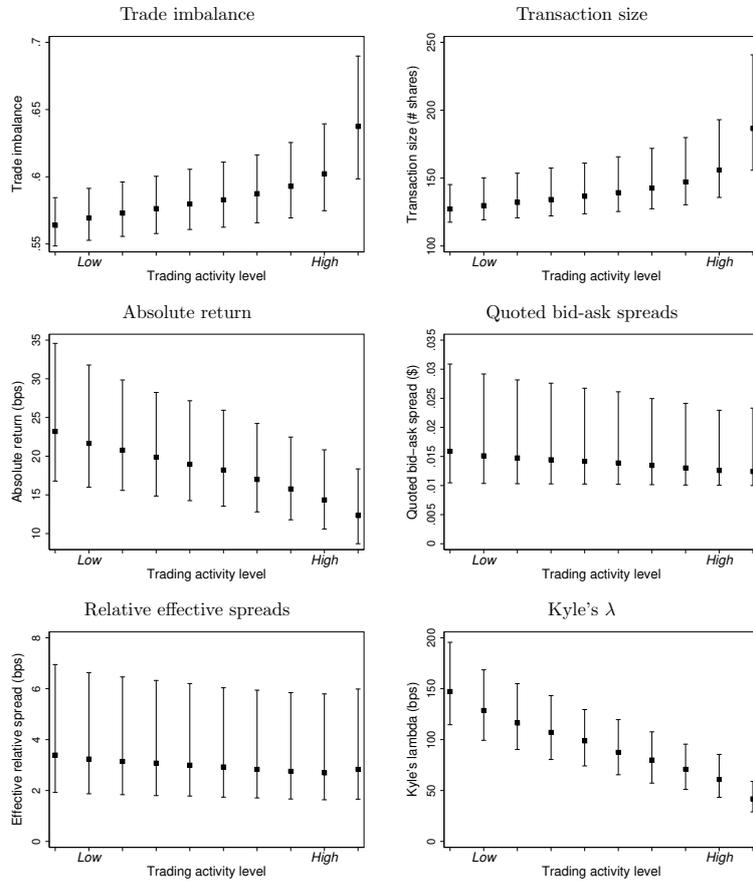


Figure 4
Trading outcomes and measures of trading costs vs. trading activity, 2012.
 The figure presents medians and inter-quartile ranges of trading outcomes and measures of trading costs at different trading activity deciles. For each trading outcome (trade imbalance, transaction size, absolute return, dollar-weighted bid-ask spreads, and dollar-weighted effective spreads) we first calculate the median outcome across trade sequences of stock j in month t at each trading activity decile. We present quartile statistics (medians, 25th percentiles, and 75th percentiles) of these medians, pooling across months and stocks. We also present quartile statistics of stock-activity-level-specific estimates Kyle's λ by trading activity decile.

establish positions via providing liquidity rather than consuming it in order to reduce the otherwise even higher price impacts—resulting in longer trade times for a given fixed dollar volume, small transactions, more balanced order flow, and large price impacts. Figure 4 also provides the reinforcing evidence that trade-time estimates of Kyle's λ drop sharply with trading activity. The steep decline in estimates of Kyle's λ reflects that when there is greater liquidity provision (and hence when

trading activity is higher), traders submit larger marketable orders, resulting in less-balanced (larger) measures of net order flow being associated with smaller price impacts, and hence reduced sensitivity of prices to net order flow.

In sum, variations in trading activity reflect variations in extant liquidity, and variations in the corresponding cumulative price impacts $|r_j(k)|$ capture the relevant pricing implications, underscoring our measure of average per-dollar cumulative price impacts, *BBD*, as a measure of stock liquidity.³⁷ Our method samples price impacts over shorter intervals when liquidity and trading intensity is higher, aligning the sampling of price impacts with the actual intensity of liquidity consumption—trading costs are realized over shorter horizons when liquidity is high, and realized over longer horizons in low-liquidity times. In contrast, calendar-time aggregation approaches that, for example, aggregate over thirty minute periods weight each observation the same regardless of the level of liquidity. We will show that our approach delivers more accurate proxies of trading costs.³⁸

Our findings indicate that variations in trade-time returns are distinct from fundamental volatility. Absolute trade-time returns, i.e., the constituent components of our *BBD* measure, are largest when markets are least active, and smallest when markets are most active. This is the *opposite* of what should happen were fundamental volatility in the form of information arrival the driver of *BBD*. Moreover, Barardehi and Bernhardt (2017) show that trade-time returns over consecutive trade sequences exhibit reversion in less active markets, and momentum in more active markets. Reversals in returns when trading activity is lower indicates that the larger trade-time absolute returns found then do not represent information-driven price movements; and are distinct from fundamental volatility. Rather, these reversals in less active times suggest that liquidity providers are being compensated quickly for their services when markets are less liquid.

Figure 4 also suggests that spread measures fail to capture intradaily variations in liquidity and the corresponding variations in *cumulative* price impacts. Median quoted spreads barely vary with trading activity level. In addition, while the third quartile of quoted spreads does decline slightly trading activity rises, the first quartile is essentially fixed at ¢1, reflecting truncation in the cross-section of quoted spreads due

³⁷ Unreported results confirm qualitatively identical patterns for absolute volume-weighted returns $|wr_j(k)|$ verifying the merits of our *WBBD* measure.

³⁸ Note that a volume-weighted version of a high-frequency Amihud measure would still mix different trading activity levels within a calendar-time period, and, in active conditions, may not identify relevant within time-period price movements that our trade-time approach would identify. That is, volume weighting may not correct for the essential differences between calendar-time and trade-time *sampling*.

Trade-time based measures of liquidity

to the minimum penny tick. Moreover, relative effective spreads are almost uncorrelated with trading activity. This reflects the systematic increases in transaction sizes (and use of larger marketable orders) as trading activity rises—that spread measures do not account for.³⁹ When markets are less active, it takes many more small trades to establish positions, and spread measures do not account for the cumulative costs of these trades—spread measures do not account for serial correlations in transactions and prices driven by dynamic order splitting strategies.

To reinforce these arguments, we contrast the patterns presented in Figure 4 with their calendar-time analogues. We measure trading outcomes over 15-minute intervals, and then examine how trading outcomes vary with 15-minute dollar volumes, i.e., with calendar-time measures of trading activity.⁴⁰ Figure 5 shows that the relationships between trading outcomes and trading activity in calendar-time differ dramatically from those found in trade-time. Importantly, trade imbalances over 15-minute intervals decline as calendar-time dollar volumes rise. This suggests an issue of over-aggregation in more active times that leads to the lumping together of orders from opposing sides, resulting in balanced order flow estimates despite high trading activity. We also find that the median 15-minute absolute return, i.e., calendar-time volatility, rises with calendar-time dollar volume. As Barardehi and Bernhardt (2017) highlight, this reflects high trading volumes, rather than large per-unit-trade price movements, driving the positive association between calendar-time volume and volatility. That is, calendar-time measures of volatility conflate the distinct impacts of per-unit-trade price movements with those of trading volumes. The twin facts that price impacts rise with trading volume, while signed order imbalances fall underlie our next finding that estimates of Kyle’s lambda, using calendar-time observations, rise with trading volume.^{41,42}

We next present and simulate a dynamic model of liquidity provision and institutional trade. The model integrates several of the key empirical findings just established in order to explain why trade-time measures

³⁹ One manifestation of this is that effective spreads rise slightly when trading is most active. This seems to reflect that when depth at best prices is very high it tends to be placed slightly further away from the quote midpoint to protect against adverse moves in the asset value, and the greater costs of sniping.

⁴⁰ For the reasons mentioned earlier, to avoid issues with endogeneity of current dollar volume and price, we use the closing price on the previous day to calculate dollar-volumes. However, findings are not qualitatively altered if we use current price to calculate intradaily dollar volumes.

⁴¹ This can lead to a misplaced conclusion that, because price impacts of order flow rise with trading volume, variations in trading volumes/activity are primarily information-driven.

⁴² As expected, average transaction size, quoted spreads, and effective spreads vary similarly with trading activity regardless of whether one uses calendar- or trade-time measures of activity. These variables primarily reflect transaction-level information, and not the temporal correlation structures of orders and prices.

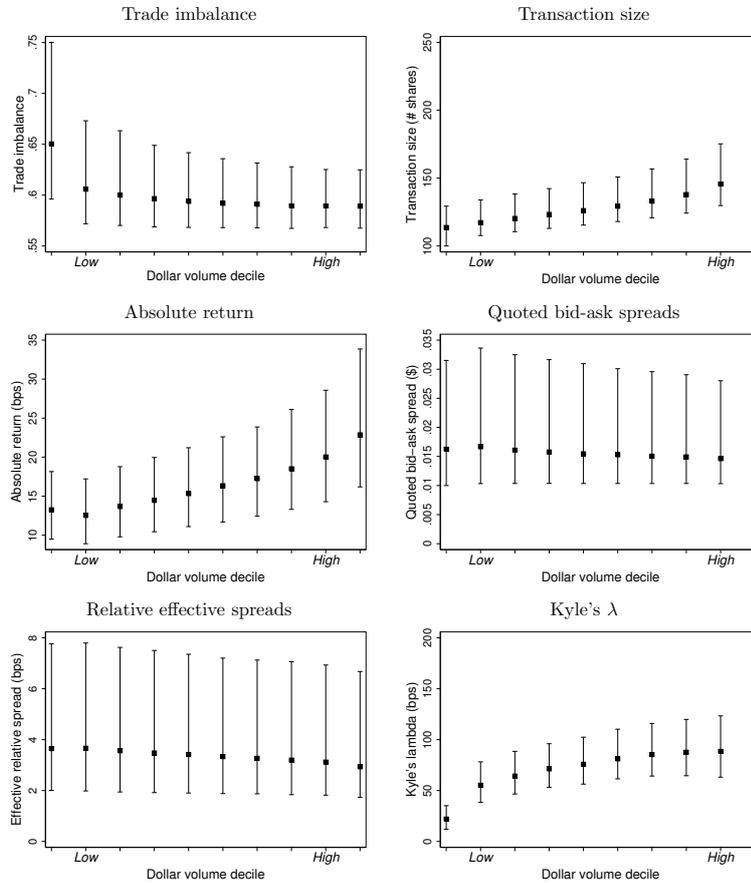


Figure 5
Trading outcomes and price impacts vs. calendar-time dollar volume, 2012.
 The figure presents medians and inter-quartile ranges of trading outcomes and estimates of Kyle's lambda at different dollar volume deciles. For each trading outcome (trade imbalance, transaction size, absolute return, dollar-weighted bid-ask spreads, and dollar-weighted effective spreads) we first calculate the median outcome across 10-minute time intervals of stock j in month t at each stock-specific dollar-volume decile. We present quartile statistics (medians, 25th percentiles, and 75th percentiles) of these medians, pooling across months and stocks. We also present quartile statistics of stock-volume-decile-specific estimates Kyle's λ by dollar volume decile.

better capture trading costs than calendar-time measures, even when institutional trade is completely non-strategic. The model provides deeper intuition and understanding as to why *BBD*, which calculates price impacts over trade-time intervals, better measures trading costs than variants of Amihud's measure, which calculate price impacts over calendar-time intervals. In essence, when trading activity is high, calendar-time aggregation tends to bundle together large trades in

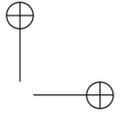
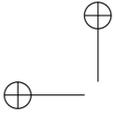
opposing directions within the same window, resulting in offsetting price impacts and under-estimates of trading costs. The model also provides insight into the tradeoffs that enter the choice of the intervals over which price impacts are calculated. We show that very short intervals may lead to overestimated trading costs, but that very long intervals lead to underestimates.

5.1 Trade-time vs. Calendar-time Measures: a Dynamic Trading Model

In our discrete time, dynamic stochastic Markov limit-order book model, (a) the bid-ask spread is always at the minimum tick; (b) there is one round lot of retail trade on each side of the market; (c) order flow imbalances reflect institutional, uninformed trade; (d) order flow imbalance stochastically moves best prices against order flow in the next period; and (e) public information arrival may draw institutional trade creating order flow imbalance.

Each period has one retail buy order, and one retail sell order: period order flow is balanced absent institutional trade. Institutional traders enter the market one at a time, and their trade executions do not overlap. Each institutional “parent” order is split into round lot “child” orders that are executed in successive periods, creating an imbalance of one round lot. Parent orders are summarized by an order size and sign pair, (x, s) with $x \in \{0, 1, 2, 5\}$ denoting size in round lots, and $s \in \{+, -\}$ indexing buy (+) and sell (-) orders. $x=0$ indicates that no institutional order arrived. size and \bar{x} to denote the current size. To minimize the number of free parameters, we suppose that if the previous parent order size is \underline{x} , then the current parent order remains \underline{x} with probability $\alpha_{\underline{x}} + \frac{1-\alpha_{\underline{x}}}{4}$; and with uniform probability $\frac{1-\alpha_{\underline{x}}}{4}$, the current parent order size is $\bar{x} \neq \underline{x}$. Thus, $\alpha_{\underline{x}} \in [0, 1)$ measures the autocorrelation in order size. A non-zero parent order has the same sign as the previous non-zero parent order with probability $\beta + \frac{1-\beta}{2}$, where $\beta \in [0, 1)$ measures the autocorrelation in order sign. Buy and sell parent orders are equally likely after a parent order of 0. Thus, the transition matrix on parent orders is characterized by 5 parameters: $\alpha_0, \alpha_1, \alpha_2, \alpha_5$ and β .

To capture discretionary liquidity provision, we assume that after a positive buy order imbalance, the bid and ask price may shift up. For simplicity, a single parameter ζ describes the dynamics of liquidity provision in the two most recent time periods. Order imbalances on the same side in two successive periods may move quotes by two ticks: liquidity providers fill in the depth at the existing quotes with probability ζ ; but with probability $1-\zeta$, consecutive periods of net buy order flows shift the bid and ask up two ticks, and consecutive sells shift them down two ticks. If order flow is imbalanced only in the most recent period, or order flow imbalances in the two past periods are in opposite



directions, then when liquidity is withdrawn (with probability $1 - \zeta$), order flow imbalances only shift quotes by one tick.

To allow for the possibility that information arrival may draw institutional trade, we assume that in periods in which institutional orders are executed, the fundamental asset value may rise or fall by ϵ ticks, each with probability η , or remain unchanged with probability $1 - 2\eta$.

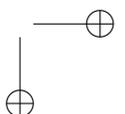
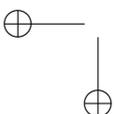
We simulate the economy over horizons that span 1,000 units of time, drawing 100 samples. For each sample, we calculate trading costs, and the analogues of Amihud’s measure, and *BBD* at different aggregation horizons. We then calculate the mean and 95% confidence intervals of these three measures using the quantities from the 100 samples. The 1,000 period horizon captures the limited number of trade sequences or time intervals that enter constructions of measures, which introduces measurement error.

Actual trading costs. The trading cost of an institutional order is given by the total absolute difference between execution prices of the corresponding round lots and the quote mid-point at the period in which the trader enters the market, net of the effect of information arrival.⁴³ For example, consider a two round lot parent order, where the first unit traded caused quotes to shift up two ticks. Then the total trading cost of that order is 2.5 ticks (half of the bid-ask spread plus the two-tick price impact). During this execution, there were also four round lots of trade from retail traders that had a net trading cost of zero (as their trading costs offset).⁴⁴ Thus, the 2.5 ticks reflect the trading cost associated with six round lots of trading volume, i.e., 0.425 ticks per round lot transaction. In time periods when institutional trading is absent, trading costs are zero. The measure of trading costs reflects the volume weighed average across all time units of per round lot trading costs. With information arrival, both asset values and prices shift up or down, and hence measured price impacts shift up or down, but trading costs themselves are unaffected.

Comparison with liquidity measures. By construction, the bid-ask spread is always one tick, regardless of realized trading costs. To contrast estimates of trading costs obtained using *BBD* and Amihud’s measures, we compare them at different comparable aggregation horizons. We first set a time aggregation horizon length, and calculate Amihud’s

⁴³ Equivalently, one can account for information arrival; but with uninformed institutional trade, the expected signed price impact of information arrival on trading costs is zero in our model.

⁴⁴ If the quote midpoint equals the true fundamental value, when orders are filled at the ask, the retail trader with a buy order incurs a trading cost of 0.5 tick, and the retail trader with a sell order gains, incurring a “negative trading cost” of -0.5 tick—in essence, they are trading against each other, and one trader’s gain is the other’s loss.



Trade-time based measures of liquidity

measure as the average per-round lot transaction absolute returns based on transaction prices at the ends of successive aggregation horizons. The *BBD* measure is given by the average per-round lot transaction absolute returns based on transaction prices associated with the trading of the mean calendar-time volume over the time aggregation horizon.⁴⁵ Because period trading volumes are either two or three round lots, realized volumes may exceed the target by one or two round lots.⁴⁶ One may adjust target volumes downward, so that such “overshooting” results in realizations of desired average target volumes—but this is difficult due to sampling variation. Recognizing that this overshooting *attenuates* our findings, we just note its effect.⁴⁷

Figure 6 illustrates outcomes in settings without information arrival ($\eta=0$). In the benchmark Case 1 parameterization, the α parameters are set so that (a) 75% of periods have institutional trade, (b) the sign of order flow is highly persistent ($\beta=0.8$), and (c) markets are relatively illiquid (quoted depth refills with probability $\zeta=0.2$). The figure highlights the general properties that without information arrival, (1) both measures underestimate trading costs; and (2) Amihud’s measure underestimates trading costs by far more than *BBD* for all plausible parameterizations and appropriately-chosen aggregation horizons.

For instance, in Case 1, Amihud’s measure underestimates trading costs by 16% at the shortest aggregation horizon, while *BBD* underestimates by only 8.5%; for aggregation horizons 5 and 10, the extents of underestimation are 18% vs. 14% and 24% vs. 21%, respectively. As mentioned before, *BBD* would underestimate by even less (for a given target trading volume) were the target adjusted to account for instances of “overshooting”.

To understand the economics, it is first useful to recognize that both measures underestimate trading costs by large amounts over long aggregation horizons because long horizons tend to lump buy- and sell-driven order flow together. The result is that the net signed price impact over the entire interval sharply underestimates actual trading costs. *BBD* is far more accurate than Amihud’s measure for shorter aggregation horizons. This reflects that calendar-time aggregation places the same weight on all time periods regardless of trading activity.

⁴⁵ The mean typically is not an integer, so we probabilistically round the target volume to match it: if the mean is 15.6, we randomize between targets of 15 and 16 so that the average target volume is 15.6.

⁴⁶ For example, for Case 1 in Figure 6, while *targeted* trading volume rises from 5.46 to 43.41 round lots going from the shortest to the longest calendar-time aggregation horizon, the corresponding average *realized* trade-time volume goes from 6.05 to 44.11 round lots.

⁴⁷ As Figure 6 shows, *BBD* is decreasing in target trading volume in the model. Thus, the gap between *BBD* and Amihud’s measures would widen were we to adjust for overshooting, re-enforcing our findings.

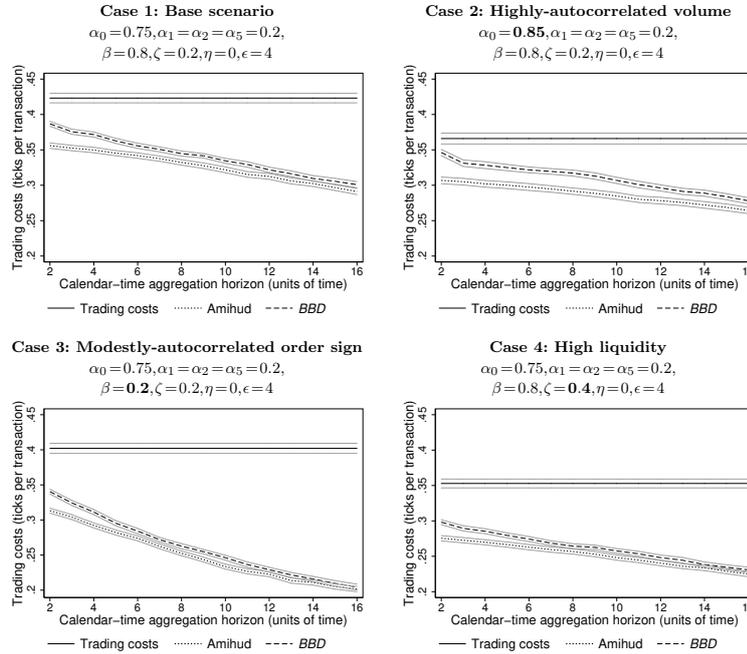


Figure 6
Simulation: calendar time vs. trade time measures (no information arrival).
 The economy is simulated over a time span of 1,000 time units, 100 times. Per round lot trading costs reflect the sum of differences between transaction prices and the mid-point price at the period the trader entered the market, divided by the corresponding trading volume. In each draw, Amihud’s measure is calculated as the average absolute returns (based on transaction prices) over 2- to 16-unit time intervals. The corresponding BBD measures are based on the average absolute returns calculated over trade times of target trading volumes, where target trading volumes are given by the associated mean calendar time trading volumes. The mean and 95% confidence intervals of trading costs, Amihud measures, and BBD using the samples from the 100 draws are plotted against aggregation horizon.

Trade-time aggregation, however, is volume indexed, and samples more often in more active conditions. As a result, *BBD* can better proxy price movements driven by variations in economic activity, which in our model is driven by institutional order flow. Importantly, time aggregation results in far greater volatility in trading activity across observations—under-aggregating in low activity times, and over-aggregating in high. This especially matters when time aggregation bundles together moments of very high trading activity, sometimes grouping large trades in opposing directions in the same window, resulting in offsetting price impacts. Bundling of large offsetting trades is far less with trade time aggregation because high trading volume is likely to result in hitting the volume target, causing subsequent large trades in the opposing direction to be allocated to a different trade sequence.

Trade-time based measures of liquidity

In the data, as Figure 4 and 5 reveal, with trade-time aggregation, order flow becomes increasingly unbalanced as trading activity rises, indicating that aggregation of offsetting trades is modest; but with calendar-time aggregation, order flow becomes more *balanced*, suggesting extensive aggregation of offsetting trades. Our stylized model captures this.

To underscore this, we show how stochastic temporal autocorrelation in trading volume drives the relatively poor performance of calendar-time measures. Parameters in Case 1 are selected so that periods featuring order flow imbalance (due to non-zero institutional orders) are three times more frequent than periods with balanced order flow. Case 2 illustrates the impact of temporal autocorrelation in volume, which delivers variation in trading volume across calendar-time windows even though our stylized model keeps this variation in a period to a single round lot. Case 2 parameters are set so that periods with order flow imbalance are only twice as frequent as those with just retail trade—instances of balanced order flow are more persistent. Figure 6 shows how trade-time measures accommodate the increased persistence because the sampling windows adjust to equate trading volume, but the performance of calendar-time measures falls off sharply.

Case 3 reveals that, as one would expect, weaker positive correlations in *signs* of parent orders cause both measures to underestimate trading costs by more because the aggregation windows are more likely to aggregate trades of opposing signs, but the qualitatively superior performance of trade-time based measures is unaffected.

Case 4 makes the point that with binding minimum ticks, spread-based measures fail to capture trading costs. Case 4 shows that improving liquidity from $\zeta=0.2$ to 0.4 —so that liquidity providers are more likely to refill removed liquidity at best prices—reduces actual and estimated trading costs. Both calendar- and trade-time estimates of trading costs significantly respond to changes in liquidity. However, because the penny tick always binds, the bid-ask spread does not vary with the level of liquidity: both the half-quoted and effective spreads remain half a tick, regardless of liquidity.

Cases 5 and 6 in Figure 7 show the effect of introducing information arrival. Information arrival has only modest effects on the *difference* between *BBD* and Amihud’s measure at different aggregation horizons, but because information has price impacts unrelated to trading costs, it shifts these measures up relative to trading costs. As a result, trading costs are best captured by an intermediate aggregation horizon that accounts for the non-trading cost component of price impacts, but is not so long as to aggregate excessively.

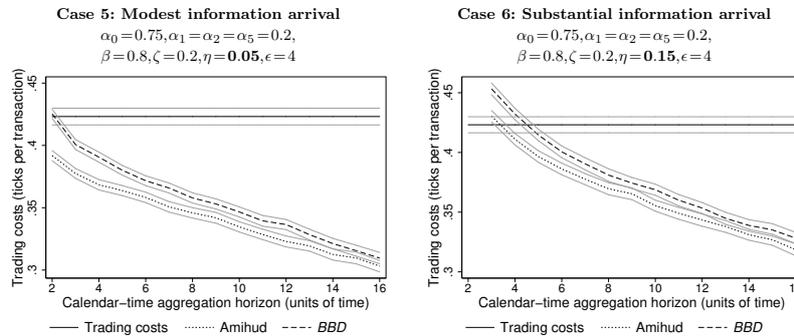


Figure 7
Simulation: calendar time versus trade time measures (information arrival).
 The economy is simulated over a time span of 1,000 time units, 100 times. Per round lot trading costs reflect the sum of differences between transaction prices and the mid-point price at the period the trader entered the market, divided by the corresponding trading volume. In each draw, Amihud’s measure is calculated as the average absolute returns (based on transaction prices) over 2- to 16-unit time intervals. The corresponding BBD measures are based on the average absolute returns calculated over trade times of target trading volumes, where target trading volumes are given by the associated mean calendar time trading volumes. The mean and 95% confidence intervals of trading costs, Amihud measures, and *BBD* using the samples from the 100 draws are plotted against aggregation horizon.

5.2 Comparing trade-time vs. standard liquidity measures

Returning to our empirical analysis, we now show that our measures of liquidity do a better overall job of capturing cross-stock variations in trading costs than do other measures. To establish this, we calculate the cross-stock correlations between each liquidity measure used in our paper and estimated trading costs provided by Investment Technology Group’s Agency Cost Estimator (ITG ACE).⁴⁸ ITG uses these trading costs estimates to highlight the execution services that ITG provides its institutional clients, detailing the relationships between order sizes, trading strategies, and expected trading costs. The trading costs of each stock are estimated according to order sizes. ACE considers 32 order size bins: the first bin features orders of 100 shares and smaller, the 10th bin contains orders of about 12% of median daily trading volume, the 20th bin contains orders that are about 66% of median daily volume, and the 32nd bin features massive orders that on average are ten times median

⁴⁸ ITG (2007) provides detailed descriptions of the estimator. The ITG estimates build on ideas set out in Almgren and Chriss (2001). According to Almgren and Chriss (1999), “ACE computes the expected cost of executing a basket of securities along the multivariate optimal liquidation path... [using] a proprietary parameterization of the temporary and permanent impact functions to compute transaction costs...” For more details see Borkovec and Heidle (2010).

Trade-time based measures of liquidity

Table 2
Cross-stock correlations between liquidity measures and ITG ACE, 2011–2013.

	ACE^{tot}	ACE^{b1}	ACE^{b2}	ACE^{b3}	ACE^{b4}	ACE^{b5}
<i>BBD</i>	0.82	0.82	0.81	0.79	0.76	0.74
<i>WBBD</i>	0.83	0.84	0.83	0.80	0.77	0.75
<i>AML</i>	0.80	0.78	0.80	0.81	0.79	0.77
<i>AMH</i>	0.71	0.74	0.71	0.67	0.62	0.58
<i>PSP</i>	0.82	0.88	0.81	0.74	0.68	0.65
<i>LAMBDA</i>	0.78	0.80	0.77	0.74	0.71	0.69
<i>EPSP</i>	0.75	0.80	0.74	0.67	0.62	0.59

Pairwise cross-stock correlations of trade-time based liquidity measures (*BBD* and *WBBD*), the low frequency Amihud measure (*AML*), the high frequency Amihud measure (*AMH*), percentage bid-ask spreads (*PSP*), estimated Kyle’s λ (*LAMBDA*), and dollar-weighted percentage effective spreads (*EPSP*) with each version of ITG’s estimated trading costs (ACE^{tot} , ACE^{b1} , ACE^{b2} , ACE^{b3} , ACE^{b4} , and ACE^{b5}) are calculated every month. The table reports the cross-month averages of the correlation coefficients.

daily volume.⁴⁹ We focus on ACE’s output based on ITG’s optimal execution strategy for a risk-neutral client in 2011–2013, averaging the per-share estimated trading cost of each stock in each month across order size bin groups. We focus on bins 2–31, dropping the smallest bin that may reflect retail orders, and the largest bin of extremely large and rare institutional orders. We group the remaining bins into five equally-sized groups. ACE_{jt}^{tot} represents the overall average across the thirty order size bins, ACE_{jt}^{b1} captures the average cost of small orders (bins 2–6), ACE_{jt}^{b2} represents the average cost of medium-small orders (bins 7–12), ACE_{jt}^{b3} captures the average cost of medium orders (bins 13–18), ACE_{jt}^{b4} captures the average cost of medium-large orders (bins 19–24), and ACE_{jt}^{b5} captures the average cost of large orders (bins 25–31). We calculate the cross-stock correlation coefficients between each liquidity measure and the six versions of ACE estimates each month, and then compute the cross-month average.

Table 2 reveals that all liquidity measures covary with ITG’s trading cost estimates, but that the trade-time based liquidity measures, *BBD* and *WBBD*, do a somewhat better overall job of matching cross-sectional variations in trading costs. On average, trade-time based liquidity measures are more strongly correlated with ACE^{tot} , i.e., with average trading cost estimates over all parent order sizes. The high average (across stock) correlations of 0.82–0.83 of *BBD* and *WBBD* with ACE^{tot} exceed the correlations of ACE^{tot} with other liquidity measures. More importantly, the correlation of trade-time measures with trading costs is high *regardless* of the sizes of parent orders being considered. In contrast, decomposing estimated trading costs by order size reveals that while the mean correlation of *PSP* with the

⁴⁹ The massive dispersion of sizes in larger bins reflects the fact that very large order sizes are rare, even though such orders account for a comparable share of dollar volume. As such, a wider range of observations is required to maintain statistical power.

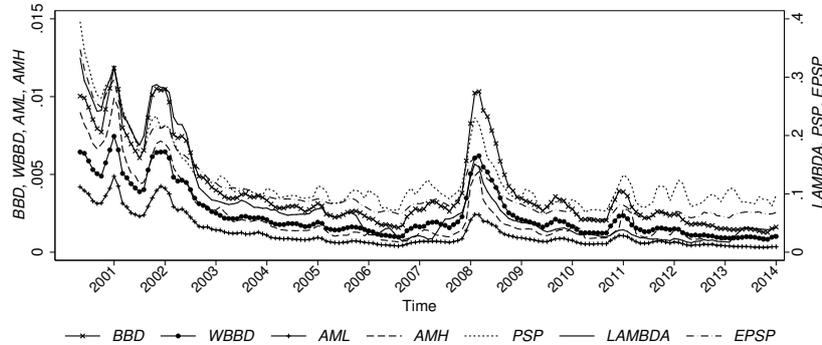


Figure 8
Monthly median values of different liquidity measures over time.
 The month-by-month medians of the trade-time based liquidity measures (*BBD* and *WBBB*), the low frequency Amihud measure (*AML*), the high frequency Amihud measure (*AMH*), percentage bid-ask spreads (*PSP*), estimated Kyle’s λ (*LAMBDA*), and dollar-weighted percentage effective spreads (*EPSP*) are presented.

trading costs of small orders ACE^{b1} is 0.88, the correlation between *PSP* and trading costs drops sharply with order size, falling to 0.65 for the bin group containing the largest orders. This finding suggests that percentage spreads are a viable proxy of trading costs for small retail transactions, but not for large institutional orders. Conversely, correlations with the low-frequency Amihud measure slightly exceed those for trade-time based measures for large parent orders whose executions may span multiple days; but correlations with low-frequency Amihud measure are distinctly lower for smaller parent orders. Overall, *BBD* and, especially, *WBBB* reveal themselves to be reliable proxies for cross-stock variations in trading costs, regardless of the size of the institutional orders.

As a final piece of evidence establishing that *BBD* and *WBBB* capture liquidity, we show that they are highly correlated with the other standard liquidity measures. Table 3 shows the averages of the month-by-month pairwise correlation coefficients between variables used in our analysis.⁵⁰ Correlations of *BBD* and *WBBB* with other liquidity measures are all high—uniformly exceeding 0.6—indicating that they capture similar phenomena.

Figure 8 presents the month-by-month evolution of the medians of these liquidity measures. It shows that *BBD* and *WBBB*, and other commonly-used measures of liquidity for a typical stock display close co-movements over time, falling sharply prior to the financial crisis, rising steeply during the crisis, before returning roughly to the lower

⁵⁰ For each pair of variables, we estimate the correlation coefficient every month, and then find the average of 165 monthly estimates.

Trade-time based measures of liquidity

Table 3: Sample correlations.

	WBBD	BBD	WBD	AML	AMH	PSP	LAMBDA	EPSP	β^{mkt}	β^{mb}	β^{ml}	$\beta^{um,d}$	$\beta^{um,l}$	$\beta^{um,b}$	$\beta^{um,t}$	β^{um}	$\ln(M)$	BM	$\ln(M)$	MOM^{1-4}	MOM^{5-12}	DYD	SD	
WBBD	0.95																							
AMH	0.81	0.79																						
AML	0.82	0.85	0.80																					
PSP	0.74	0.75	0.64	0.75																				
LAMBDA	0.67	0.68	0.47	0.58	0.71																			
EPSP	0.60	0.62	0.51	0.63	0.79	0.58																		
β^{mkt}	-0.09	-0.09	-0.09	-0.10	-0.12	-0.12	-0.09																	
β^{mb}	0.10	0.11	0.04	0.07	0.20	0.15	0.16	0.05																
β^{ml}	0.04	0.04	0.01	0.01	0.05	0.03	0.06	0.07	0.18															
$\beta^{um,d}$	-0.07	-0.07	-0.03	-0.04	-0.07	-0.06	-0.08	-0.12	-0.16	-0.02														
$\beta^{um,l}$	0.18	0.19	0.08	0.13	0.19	0.16	0.17	0.02	0.11	0.16	-0.12													
$\beta^{um,b}$	-0.36	-0.38	-0.24	-0.32	-0.58	-0.45	-0.46	0.03	-0.45	-0.10	0.10	-0.12												
$\beta^{um,t}$	-0.06	-0.07	-0.01	-0.01	-0.01	-0.01	0.00	-0.02	0.01	0.02	0.01	0.02	0.01	0.02	0.01	0.02	0.01	0.02	0.01	0.02	0.01	0.02	0.01	0.02
$\ln(M)$	-0.07	-0.07	-0.02	-0.04	-0.05	-0.05	-0.06	0.00	0.01	0.02	0.01	0.02	0.10	0.10	0.10	0.10	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02
MOM^{1-4}	-0.03	-0.03	-0.01	-0.01	-0.02	-0.02	-0.02	-0.02	-0.05	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07
MOM^{5-12}	0.14	0.16	0.06	0.12	0.22	0.22	0.18	0.20	0.33	0.42	0.09	-0.18	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06
DYD																								
SD																								

The table reports pairwise sample correlation coefficients between variables. Correlations between variables of each pair are calculated every month, and the cross-month averages are reported. The β s reflect factor loading estimates from a four-factor Fama-French model; BM is the book-to-market ratio; $\ln(M)$ is the natural logarithm of stock's market-capitalization; MOM^{1-4} is the compound return over the four-month period prior to month of observation, and MOM^{5-12} is the compound return over the eight months prior to that; DYD is the dividend yield over the prior twelve months; SD is the standard deviation of monthly returns over the previous sixty months; BBD is the average per-dollar return measured over trade-time intervals (time required for 0.04% of market-cap to be traded); WBD is the average per-dollar VWAP-based returns measured over trade-time intervals; AML reflects the average per-dollar daily returns; AMH reflects average per-dollar returns measured over thirty-minute intervals; PSP is the average time-weighted percentage bid-ask spreads; $LAMBDA$ reflects estimates of Kyle's lambda using data from 10-minute windows; and $EPSP$ is dollar-weighted average percentage effective spreads.

pre-crisis levels. Anand et al. (2013) show in their Figure 1 that these common measures of liquidity are, in turn, highly correlated with average implementation shortfalls of institutional orders, a measure of institutional traders’ price impacts—which means that our measures are also highly correlated with average implementation shortfalls. Comparing the patterns in Figure 8 for all liquidity measures with the temporal patterns in trading activity in Figure 3 confirms that increased trading activity is associated with improvements in liquidity of typical stocks.

6. Cross-sections of returns and liquidity premia

We begin by describing the asset pricing models that we use to investigate the role of stock liquidity in explaining cross-sections of expected returns and risk-adjusted expected returns. Our models exploit variations in systematic risk factor loadings and stock characteristics, including liquidity, to explain variations in returns.

Model I (returns). We first use a standard specification, similar to that used by Amihud (2002), Ben-Rephael et al. (2015) and others, to establish the existence of characteristic liquidity premia and to explore their temporal evolution. Model I regresses monthly returns, $R_{j,t}$, on factor loadings and stock characteristics from the previous month, as follows:

$$\begin{aligned}
 R_{j,t} = & \lambda_0 + \lambda_1 \beta_{j,t-1}^{mkt} + \lambda_3 \beta_{j,t-1}^{smb} + \lambda_2 \beta_{j,t-1}^{hml} + \lambda_4 \beta_{j,t-1}^{umd} \\
 & + \lambda_5 BM_{j,t-1} + \lambda_6 \ln(M_{j,t-1}) + \lambda_7 MOM_{j,t-1}^{1-4} \quad (11) \\
 & + \lambda_8 MOM_{j,t-1}^{5-12} + \lambda_9 DYD_{j,t-1} \lambda_{10} SD_{j,t-1} \\
 & + \lambda_{11} LIQ_{j,t-1} + \sum_{m=1}^T \lambda_{12}^m Dum_m + \lambda_{13} \Delta(YIR)_{j,t} + \epsilon_{j,t},
 \end{aligned}$$

where $LIQ \in \{BBD, WBBD, AML, AMH, PSP, LAMBDA, EPSP\}$ denotes the liquidity measure; $\Delta(YIR)_{j,t}$ denotes the change (from month $t-1$) in the difference between the industry and the sample’s average return;⁵¹ and $Dum_m = 1$ if j is observed in month m and 0 otherwise.⁵² The monthly dummies capture any common variation in returns resulting, for example, from temporal changes in the trading environment. The liquidity measures, factor loadings and stock characteristics are all lagged to avoid any impacts of contemporaneous

⁵¹ We exclude stock j when computing the return of the industry to which stock j belongs.

⁵² Our sample period begins in April 2001, due to our 3-month rolling window construction.

covariations between returns and independent variables that could complicate inference.

Our primary empirical approach is generalized least squares (GLS) panel estimation with stock fixed effects in the error term. This estimation strategy helps identify differences in expected returns caused by unobserved temporally-constant firm attributes; it also allows for specification of firm clusters. Such clusters allow us to estimate robust standard errors in the presence of firm-specific serial correlations in the error term.⁵³ Petersen (2009) highlights the merits of this estimation strategy in terms of avoiding biasing estimated standard errors downward. We discuss the error-term structure and alternative estimation approaches in detail in the appendix. In addition to clustering standard errors, we use $\Delta(YIR)$ to control for industry-time-specific error term autocorrelations—i.e., to control for common industry shocks that could lead to contemporaneously correlated errors for stocks in the same industry.

Liquidity premia obtained using GLS are given by the product of the characteristic liquidity coefficient and the cross-month average of month-specific median liquidity. The resulting quantities measure the estimated monthly return premia that investors demand for establishing a position in a typical stock versus a hypothetical perfectly-liquid stock.

We report results for the commonly-used Fama-MacBeth estimation strategy in the appendix. These estimates support the findings derived from GLS estimates.⁵⁴ We do not focus on the Fama-MacBeth results for a few reasons. First, the Fama-MacBeth approach presumes that the factor loadings, i.e., the $\beta_{j,s}$ s, are constant over time—they help serve as substitutes for stock fixed-effects. However, the stock-specific factor loadings that we estimate using time-series data vary substantially from month to month, creating “Errors-In-Variables” (EIV) issues for the Fama-MacBeth approach. Common approaches to circumventing EIV concerns are to use portfolios as test assets (Jensen et al. (1972) and Fama and MacBeth (1973)) or to estimate factor loadings from portfolios of stocks (Fama and MacBeth (1992)). We cannot employ such approaches because our goal is to identify the role of *stock* specific characteristics. EIV is not an issue for the GLS fixed-effects estimator; as such, it facilitates explaining expected returns using *stock* characteristics. The Fama-MacBeth approach also assumes that the

⁵³ For example, because we use observations from the preceding three months to construct liquidity measures, liquidity is necessarily serially-correlated at the stock level. Use of stock-specific clusters addresses this. In Appendix 4, we show that qualitatively identical findings obtain when we only use observations from the previous month to construct our (now noisier) liquidity measures.

⁵⁴ We also present results for the Asparouhova et al. (2010, 2013) modification of Fama-MacBeth to deal with potential cross-sectional correlation between liquidity and microstructure noise. Qualitatively identical findings obtain.

cross-sectional samples are realizations of an *i.i.d.* process. Violations of this assumption in our relatively short time-series give rise to incorrect inferences.⁵⁵

Model II (risk-adjusted returns). We also consider an alternative, more conservative, estimation approach that uses stock characteristics to explain variations in expected risk-adjusted returns. To deal with the EIV issues in Fama-MacBeth estimation when individual stocks serve as test assets, Brennan et al. (1998) propose moving the errors-in-variables to the left-hand side. In the three-step procedure, the risk-free rate and products of risk factors and factor loadings are used to construct risk-adjusted returns. Stock characteristics are then used to explain the variations in expected *risk-adjusted* returns. We estimate

$$\begin{aligned}
 R_{j,t}^* &= \delta_0 + \delta_2 BM_{j,t-1} + \delta_3 \ln(M_{j,t-1}) + \delta_4 MOM_{j,t-1}^{1-4} \\
 &+ \delta_5 MOM_{j,t-1}^{5-12} + \delta_6 DYD_{j,t-1} + \delta_7 SD_{j,t-1} \\
 &+ \delta_8 LIQ_{j,t-1} + \sum_{m=1}^T \delta_9^m Dum_m + \delta_{10} \Delta(YIR)_{j,t} + \epsilon_{j,t},
 \end{aligned} \tag{12}$$

where $R_{j,t}^*$ equals

$$R_{j,t} - R_t^f - \beta_{j,t-1}^{mkt} (r_t^{mkt} - r_t^f) - \beta_{j,t-1}^{hml} HML_t - \beta_{j,t-1}^{smb} SMB_t - \beta_{j,t-1}^{umd} UMD_t.$$

Here, R_t^f is the interest rate on one-month T-Bills, and the monthly Fama-French factors capture systematic risk.⁵⁶ We estimate Model II using both GLS and Fama-MacBeth approaches.

Table 4 shows that for all liquidity measures, except *EPSP*, characteristic liquidity still significantly explains variations in expected returns. In Model I, estimated monthly liquidity premia vary between 0.5 and 6.3 basis points depending on the liquidity measure used. *BBD* and *WBBD* deliver large monthly liquidity premia of 5.2 and 5.3bp, respectively, i.e., annual liquidity premia of 62.3 and 64.2bp.⁵⁷ Estimates of Model II yield liquidity premia for expected *risk-adjusted* returns that are roughly half of those for expected returns.⁵⁸ These lower estimates

⁵⁵ Due to the violations of the *i.i.d.* premise, when implementing the Fama-Macbeth approach, we rely on Newey-West standard errors with two lags to partially account for any heteroskedasticity in the error terms.

⁵⁶ Model II also helps address the possibility that liquidity may be correlated with systematic risk factors or factor loadings; for example, Table 3 shows that the correlation between *BBD* (or *WBBD*) and β^{smb} is 0.24. By employing risk-adjusted expected returns as the dependent variable, Model II offers very *conservative* estimates of liquidity premia.

⁵⁷ For each model, coefficients of “non-liquidity” characteristics are stable, not varying with the different characteristic liquidity measures. We verify that they remain stable in the rest of our analysis, so we only report estimates of the coefficients on characteristic liquidity.

⁵⁸ Monthly liquidity premia estimated by Fama-MacBeth using *BBD* and *WBBD* as liquidity measures are, respectively, 4.4bp and 4.1bp for expected returns, and 4.9bp and 4.6bp for expected risk-adjusted returns.

Trade-time based measures of liquidity

Table 4: Model I (returns) and Model II (risk-adjusted returns) estimation, April 2001–December 2014.

Co-variate	Model I (returns)					Model II (risk-adjusted returns)				
	5.2	5.3	5.5	6.3	6.5	2.2	2.4	2.6	2.8	3.0
BBP_{t-1}	0.131*** (0.027)					0.056** (0.027)				
$WBBD_{t-1}$		0.218*** (0.044)					0.099*** (0.043)			
AML_{t-1}			0.041** (0.019)					0.015 (0.018)		
AMH_{t-1}				0.008** (0.003)					0.004 (0.003)	
PSP_{t-1}					0.025*** (0.004)					0.015*** (0.004)
$LAMBDA_{t-1}$						0.076*** (0.023)				
$EPSF_{t-1}$										0.018 (0.022)
β_{t-1}^{mb}	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	0.049 (0.054)				
β_{t-1}^{amb}	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)				
β_{t-1}^{mf}	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)				
β_{t-1}^{umd}	0.004*** (0.001)	0.004*** (0.001)	0.004*** (0.001)	0.004*** (0.001)	0.004*** (0.001)	0.004*** (0.001)				
BM_{t-1}	0.021*** (0.003)	0.022*** (0.003)	0.022*** (0.003)	0.022*** (0.003)	0.022*** (0.003)	0.023*** (0.003)				0.018*** (0.003)
$\ln(M_{t-1})$	-0.023*** (0.001)	-0.023*** (0.001)	-0.023*** (0.001)	-0.023*** (0.001)	-0.023*** (0.001)	-0.023*** (0.001)				-0.022*** (0.001)
MOM_{t-1}^1	0.012*** (0.002)	0.012*** (0.002)	0.012*** (0.002)	0.012*** (0.002)	0.012*** (0.002)	0.013*** (0.002)				0.017*** (0.002)
MOM_{t-1}^{1-2}	0.002* (0.001)	0.002* (0.001)	0.002* (0.001)	0.002* (0.001)	0.002* (0.001)	0.002* (0.001)				0.018*** (0.002)
DYD_{t-1}	0.066** (0.028)	0.066** (0.028)	0.066** (0.028)	0.066** (0.028)	0.066** (0.028)	0.066** (0.028)				0.018*** (0.003)
SD_{t-1}	0.037*** (0.014)	0.037*** (0.014)	0.037*** (0.014)	0.037*** (0.014)	0.037*** (0.014)	0.037*** (0.014)				-0.022*** (0.001)
Within R^2	0.30	0.30	0.30	0.30	0.30	0.06				0.06
Log-lik	161844.4	161850.2	161785.4	161788.8	161793.3	161837.4				161826.1
Observations							168397			168397

The table reports summaries of panel estimates of Model I (Equation (11)) and Model II (Equation (12)) using five different liquidity measures. Month dummies capture common variation caused by market condition changes. Estimated standard errors are in parenthesis. Standard errors are clustered at the stock level, and stock fixed effects capture any fixed heteroskedasticity in the error term. Symbols **, ***, and *** denote significance at 10%, 5%, and 1% levels, respectively. The liquidity premium for each liquidity measure is the product of the liquidity coefficient and the sample average of the monthly median liquidity measure.

reflect that Model II imposes more restrictions on how risk factors enter than Model I, essentially assuming that the four-factor model is “correct” and inflating error terms when it is not. An indication of such inflation is that the within R^2 s for Model II are only one-fifth of those for Model I.

Our findings of significant and economically meaningful liquidity premia in the 2001–2014 sample that obtain—regardless of the liquidity measure used—contrast with Ben-Rephael et al. (2015)’s failure to find significant liquidity premia in the post decimalization period. Differences in the construction of liquidity measures likely underlie why we uncover liquidity premia, but they do not. Importantly, our liquidity measures are updated each month, so changes in liquidity are incorporated on a month-by-month basis as they occur. In contrast, because Ben-Rephael et al. (2015) must use *daily* observations for their historical analysis (which goes back decades), they measure key independent variables, including stock liquidity, at *annual* frequencies. Thus, they use liquidity measures from the *preceding year* to explain the *monthly* cross-section of returns: they use the same liquidity measurement for a stock in all 12 months of a given year. This limits the variation of these independent variables, and it amplifies measurement error that biases their coefficients toward zero. Moreover, by assuming no adjustment in investors’ assessments of stock-specific liquidity over a year, Ben-Rephael et al. may mix fixed-year effects with those of cross-stock liquidity. These observations, however, do *not* undermine their central finding that liquidity premia are now far lower than in earlier eras, and, in particular, are lower post-decimalization.

Liquidity premia based on *BBD* and *WBBD* exceed those based on standard measures other than estimates of Kyle’s λ . To identify which estimates of liquidity premia one should rely on, one wants to uncover which measures best reflect liquidity considerations of investors. We previously argued that *BBD* and *WBBD*, which aggregate to account for the cumulative nature of trading costs, also better account for the fact that trading costs may realize over shorter/longer horizons due to variations in liquidity. We then showed that *BBD* and *WBBD* are more highly correlated with trading cost estimates, suggesting that they better proxy for trading costs. We now establish that *BBD* and *WBBD* better explain the cross-sections of both expected returns and expected risk-adjusted returns, suggesting that they better capture liquidity concerns of investors, especially in more recent years.

We first use *BBD* and an alternative liquidity measure in the same pricing model to estimate the importance of liquidity for cross-sections of both expected returns and expected risk-adjusted returns. Table 5 presents estimation results, revealing that using data for the entire sample period, the coefficient on *BBD* is positive and significant in

Trade-time based measures of liquidity

Table 5: Performance of *BBD* in the presence of an alternative liquidity measure.

	Model I (returns)					
	April 2001–December 2014					
<i>LIQ</i>	-0.120	0.172***	0.130***	-0.100	0.108**	0.051
<i>BBD</i>	(0.100)	(0.037)	(0.027)	(0.116)	(0.042)	(0.051)
<i>WBBD</i>	0.422**	0.184***	0.130***	0.292	0.043	0.044
	(0.167)	(0.041)	(0.037)	(0.196)	(0.044)	(0.044)
<i>AML</i>	-0.058**	-0.025	-0.024	-0.037	-0.024	-0.024
	(0.025)	(0.025)	(0.024)	(0.196)	(0.024)	(0.024)
<i>AMH</i>	-0.008*	0.012**	0.001	0.001	0.001	0.001
	(0.005)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
<i>FSP</i>	0.012**	0.012**	0.027***	0.027***	0.027***	0.027***
	(0.006)	(0.006)	(0.007)	(0.007)	(0.007)	(0.007)
<i>LAMBDA</i>	-0.042	-0.042	-0.027	-0.027	-0.027	-0.027
	(0.032)	(0.032)	(0.035)	(0.035)	(0.035)	(0.035)
<i>EFSP</i>	-0.002	-0.002	5.102***	5.102***	5.102***	5.102***
	(0.031)	(0.031)	(0.888)	(0.888)	(0.888)	(0.888)

	Model II (risk-adjusted returns)					
	April 2001–December 2014	April 2001–December 2014	April 2001–December 2014	April 2001–December 2014	April 2001–December 2014	April 2001–December 2014
<i>LIQ</i>	-0.180	0.077**	0.070*	-0.103	0.025	-0.020
<i>BBD</i>	(0.107)	(0.036)	(0.039)	(0.121)	(0.044)	(0.051)
<i>WBBD</i>	0.382**	0.028	0.079**	0.180	0.043	0.045
	(0.177)	(0.034)	(0.036)	(0.205)	(0.045)	(0.047)
<i>AML</i>	-0.030	-0.024	-0.018	-0.018	-0.018	-0.018
	(0.024)	(0.024)	(0.024)	(0.024)	(0.024)	(0.024)
<i>AMH</i>	-0.003	0.011**	0.004	0.004	0.004	0.004
	(0.005)	(0.005)	(0.006)	(0.006)	(0.006)	(0.006)
<i>FSP</i>	0.011**	0.011**	0.019***	0.019***	0.019***	0.019***
	(0.005)	(0.005)	(0.007)	(0.007)	(0.007)	(0.007)
<i>LAMBDA</i>	-0.043	-0.043	-0.059	-0.059	-0.059	-0.059
	(0.032)	(0.032)	(0.037)	(0.037)	(0.037)	(0.037)
<i>EFSP</i>	0.005	0.005	3.504***	3.504***	3.504***	3.504***
	(0.032)	(0.032)	(0.841)	(0.841)	(0.841)	(0.841)

Characteristic liquidity coefficients on *BBD* and one other liquidity measure for April 2001 to December 2014, April 2001 to December 2007, and January 2008 to December 2014 periods are estimated for Model I (returns, Equation (11)) and Model II (risk-adjusted returns, Equation (12)). Standard errors are reported in parentheses. Standard errors are clustered at the stock level; stock fixed effects are introduced to capture any fixed heteroskedasticity in the error term, and month dummies are included to capture any common variation. The symbols *, **, and *** denote significance at 10%, 5%, and 1% levels, respectively.

presence of any other standard liquidity measure; similar results obtain when we instead use *WBBD* along with other measures. In contrast, the coefficient on the standard liquidity measure is either negative or insignificant, regardless of the alternative liquidity measure considered.

Decomposing the sample into sub-samples of early versus late years reveals that much of the better performance of our measures is realized in later years. *BBD*, when included along with *AML* or *LAMBDA*, remains positively and significantly correlated with expected returns in both subsamples; whereas estimated coefficients on *AML* and *LAMBDA* become negative. When included along with the spread measures *PSP* and *EPSP*, coefficients on *BBD* become significantly positive in the later years of the sample; whereas *PSP* and *EPSP* are positively and significantly correlated with expected returns in earlier years. Similar results obtain when we estimate the base Model I formulation and Model II, where we employ stock characteristics, including liquidity, to explain expected risk-adjusted returns.⁵⁹ These results suggest that (i) standard transaction-based liquidity measures became noisier measures as algorithmic trading and, more generally, high frequency trading exploded, increasing the temporal dependence in individual trades and increasing the volatility in trading volumes; and (ii) spread-based measures grew noisier as the penny tick became an increasingly-binding constraint on bid-ask spreads, and individual trade sizes and depth fell dramatically. As a result, trade-time liquidity measures now better measure liquidity than traditional measures and hence better capture liquidity concerns.

When interpreting results in Table V, one should recall the **strong** pairwise correlations between the liquidity measures: pricing regressions that include *BBD* and another liquidity measure have multicollinearity issues. To address this, we decompose each measure $F \in \{WBBD, AML, AMH, PSP, LAMBDA, EPSP\}$ into two linearly-orthogonal components with respect to *BBD*. We first estimate the residuals $z_{j,t}^F$ and $\tilde{z}_{j,t}^F$ by fitting (13) and (14):

$$BBD_{j,t} = \hat{\alpha}_1 + \hat{\alpha}_2 F_{j,t} + z_{j,t}^F \quad (13)$$

$$F_{j,t} = \hat{\alpha}'_1 + \hat{\alpha}'_2 BBD_{j,t} + \tilde{z}_{j,t}^F. \quad (14)$$

We then estimate Models I and II using $z_{j,t}^F$ **or** $\tilde{z}_{j,t}^F$ as the liquidity measure, $LIQ_{t,j}$. This approach isolates the information content in

⁵⁹ In unreported analyses, we exclude the years 2008 and 2009 from the sub-sample of later years, verifying that our findings are not sensitive to the inclusion of data from financial crisis period.

Trade-time based measures of liquidity

one measure that is not in the other measure since, by construction, $Cov(z_{j,t}^F, F_{j,t}) = Cov(z_{j,t}^F, BBD_{j,t}) = 0$.⁶⁰

Table 6 reveals that *BBD* better explains the cross-section of returns than do the standard measures: After filtering out the commonality between *BBD* and each alternative measure, what remains (i.e., the *BBD* residual, z_{t-1}^F) is positively and significantly correlated with expected returns; but the converse is not true. The better performance is more pronounced in recent years, suggesting that *BBD* better captures stock liquidity in modern markets. In contrast, in the 2001–2007 period, *PSP* and *EPSP* outperform *BBD*. We also find that *WBBD* outperforms *BBD* in the overall sample period as well as in the later 2008–2014 period—the coefficient on the residual information in *WBBD* (but not *BBD*) is significant and positive, explaining both expected returns and expected risk-adjusted returns. This finding suggests that because *WBBD* accords with VWAP-based measures of price impacts that institutional investors seek to minimize, it more closely reflects their considerations.

This latter result leads us to conduct an analogous orthogonal decomposition analysis for *WBBD*. Table 7 shows that *WBBD* outperforms the alternative measures in the 2001–2014 period as well as in more recent years; but, as with *BBD*, both *PSP* and *EPSP* perform better than *WBBD* in the earlier 2001–2007 period. *WBBD* consistently performs slightly better than *BBD*. For example, only *WBBD* significantly outperforms *AMH* in explaining expected *risk-adjusted* returns (Model II). These results highlight the practical value of using *WBBD* rather than *BBD* to examine cross-sectional variations in stock liquidity, and are consistent with our earlier findings regarding cross-sectional correlations between liquidity measures and ITG’s estimates of institutional trading costs.

Temporal evolution of liquidity premia. We next investigate the evolution of liquidity premia. To do this, we first explore how the cross-stock dispersion of liquidity, based on different measures, evolves over time, separately exploring transaction-based measures (*BBD*, *WBBD*, *AML*, and *AMH*) and quote-based measures (*PSP*, *LAMBDA*, and *EPSP*). We calculate the inter-quartile *ratio* of each measure (i.e., the ratio of the 75th to 25th percentiles) to obtain a standardized metric with which to assess the month-by-month evolution of the dispersion in these measures.⁶¹ Figure 9 shows the evolution of these ratios. The ratios

⁶⁰ We adopt the orthogonal decomposition in a way that is consistent with the estimation method and the sample horizon. For GLS estimation, equations (13) and (14) are estimated using OLS given the data in the sub-sample of interest. For Fama-MacBeth, we fit equations (13) and (14) and store the residuals month by month.

⁶¹ Findings are robust to the use of alternative percentile ratios (e.g., 85th to 15th).

Table 7: Liquidity coefficients for orthogonally-decomposed measures; *WBBD* versus other measures.

	Model I (returns)						Model II (risk-adjusted returns)								
	April 2001–December 2014			April 2001–December 2007			April 2001–December 2014			April 2001–December 2007					
<i>LIQ</i>	<i>AML</i>	<i>AMH</i>	<i>PSP</i>	<i>LAMBDA</i>	<i>EPSP</i>	<i>AML</i>	<i>AMH</i>	<i>PSP</i>	<i>LAMBDA</i>	<i>EPSP</i>	<i>AML</i>	<i>AMH</i>	<i>PSP</i>	<i>LAMBDA</i>	<i>EPSP</i>
z^F	0.196*** (0.048)	0.216*** (0.051)	0.126*** (0.050)	0.248*** (0.054)	0.180*** (0.040)	0.097** (0.044)	0.119** (0.049)	0.036 (0.050)	0.150*** (0.054)	0.075*** (0.040)	0.097** (0.044)	0.119** (0.049)	0.036 (0.050)	0.150*** (0.054)	0.075*** (0.040)
\tilde{z}^F	-0.047** (0.020)	-0.011*** (0.004)	-0.001 (0.006)	-0.121*** (0.033)	-0.037*** (0.026)	-0.033* (0.018)	-0.008** (0.004)	0.003 (0.006)	-0.097*** (0.032)	-0.036*** (0.027)	-0.033* (0.018)	-0.008** (0.004)	0.003 (0.006)	-0.097*** (0.032)	-0.036*** (0.027)
<i>LIQ</i>	<i>AML</i>	<i>AMH</i>	<i>PSP</i>	<i>LAMBDA</i>	<i>EPSP</i>	<i>AML</i>	<i>AMH</i>	<i>PSP</i>	<i>LAMBDA</i>	<i>EPSP</i>	<i>AML</i>	<i>AMH</i>	<i>PSP</i>	<i>LAMBDA</i>	<i>EPSP</i>
z^F	0.093** (0.045)	0.071 (0.051)	-0.098* (0.052)	0.143*** (0.055)	-0.192*** (0.057)	0.033 (0.046)	0.046 (0.050)	-0.113** (0.051)	0.125** (0.055)	-0.178*** (0.064)	0.033 (0.046)	0.046 (0.050)	-0.113** (0.051)	0.125** (0.055)	-0.178*** (0.064)
\tilde{z}^F	-0.032* (0.018)	-0.003 (0.004)	0.018*** (0.005)	-0.071** (0.028)	3.645*** (0.686)	-0.027 (0.017)	-0.006 (0.004)	0.014*** (0.005)	-0.085*** (0.028)	2.679*** (0.625)	-0.027 (0.017)	-0.006 (0.004)	0.014*** (0.005)	-0.085*** (0.028)	2.679*** (0.625)
<i>LIQ</i>	<i>AML</i>	<i>AMH</i>	<i>PSP</i>	<i>LAMBDA</i>	<i>EPSP</i>	<i>AML</i>	<i>AMH</i>	<i>PSP</i>	<i>LAMBDA</i>	<i>EPSP</i>	<i>AML</i>	<i>AMH</i>	<i>PSP</i>	<i>LAMBDA</i>	<i>EPSP</i>
z^F	0.260** (0.117)	0.324*** (0.085)	0.238*** (0.080)	0.323*** (0.096)	0.253*** (0.072)	0.120 (0.093)	0.172** (0.078)	0.094 (0.074)	0.157** (0.091)	0.110 (0.067)	0.120 (0.093)	0.172** (0.078)	0.094 (0.074)	0.157** (0.091)	0.110 (0.067)
\tilde{z}^F	-0.086 (0.105)	-0.027** (0.011)	-0.024* (0.013)	-0.314** (0.129)	-0.051* (0.027)	-0.046 (0.085)	-0.018** (0.010)	-0.007 (0.012)	-0.165 (0.120)	-0.032 (0.028)	-0.046 (0.085)	-0.018** (0.010)	-0.007 (0.012)	-0.165 (0.120)	-0.032 (0.028)

Characteristic liquidity coefficients on z_{tj}^F or \tilde{z}_{tj}^F for subsamples April 2001 to December 2014, April 2001 to December 2007, and January 2008 to December 2014 are estimated using Model I (returns, equation (11)) and Model II (risk-adjusted returns, equation (12)). Standard errors are reported in parentheses. Standard errors are clustered at the stock level; stock fixed effects are introduced to capture any fixed heteroskedasticity, and month dummies are included to capture any common variation. The symbols *, **, and *** denote significance at 10%, 5%, and 1% levels, respectively.

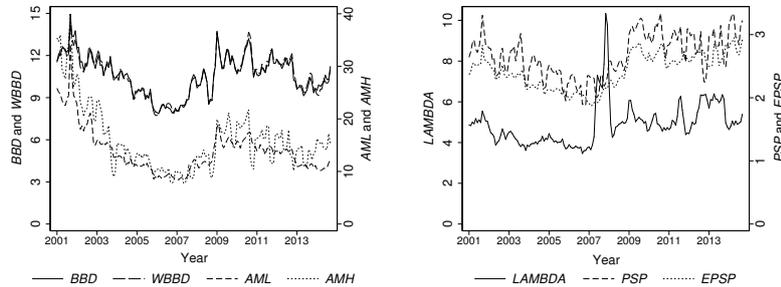


Figure 9
Monthly 75th-to-25th percentile ratios of different liquidity measures over time.
 The month-by-month 75th to the 25th percentile ratios of the trade-time based liquidity measure (*BBD*), the low frequency Amihud measure (*AML*), the high frequency Amihud measure (*AMH*), estimates of Kyle’s λ (*LAMBDA*), average percentage quoted bid-ask spreads (*PSP*), and average percentage effective spreads (*EPSP*).

for the two Amihud measures fall by about 50% in the earlier years of the sample, suggesting that the standard transaction-based measures of expected trading costs converged across stocks. In contrast, the ratios for both *BBD* and *WBBD* and the spread-based measures⁶² are all fairly stable over the entire sample period, with peak-to-trough differences that are less than half of those for the Amihud measures. These results suggest that despite all of the radical market microstructure changes in U.S. equity markets, relative trading costs in the cross-section of stocks remained fairly stable over time—but that standard transaction-based liquidity measures began to underestimate cross-stock differences in expected trading costs.

Having established the superior performances of trade-time based liquidity measures, especially in more recent years, we next investigate the temporal evolution of liquidity premia estimated using these measures, recognizing that shorter windows lead to higher standard errors. We find that the liquidity premia demanded by investors rise over our sample period. The financial crisis years (2008–2009) do not seem responsible for the increased liquidity premia; indeed, liquidity premia are highest in post-crisis years.

Table 8 shows the estimated characteristic liquidity coefficients and the corresponding liquidity premia for the sub-samples of early and late years. *BBD* and *WBBD* explain variations in expected returns both in early and later years. Estimated characteristic liquidity premia rise by about two basis points in the last five years of our sample—the financial crisis years (2008–2009) is not responsible for the increased liquidity premia. For example, liquidity premia on expected returns based on

⁶² Save for the spike in the inter-quartile ratio for *LAMBDA* during the financial crisis, this ratio is stable.

Trade-time based measures of liquidity

Table 8
Characteristic liquidity premia by sub-sample.

Model	LIQ	Apr 2001–Dec 2007	Jan 2008–Dec 2014	Jan 2008–Dec 2009	Jan 2010–Dec 2014
Model I	BBD_{t-1}	0.071** (0.035) 3.441	0.148*** (0.045) 4.612	0.087 (0.057) 4.496	0.295*** (0.067) 6.796
		0.127** (0.057) 3.844	0.247*** (0.074) 4.642	0.150 (0.095) 4.656	0.453*** (0.107) 6.326
	$WBBD_{t-1}$	0.008 (0.816) 0.408	0.060 (0.158) 1.871	0.005 (0.934) 0.248	0.189*** (0.004) 4.358
		0.024 (0.060) 0.729	0.106 (0.070) 2.003	0.018 (0.095) 0.570	0.301*** (0.105) 4.202

The table reports liquidity coefficient estimates from panel estimates of Model I (returns, equation (11)) and Model II (risk-adjusted returns, equation (12)) using the BBD and $WBBD$ liquidity measures. Month dummies capture common variation caused by market condition changes, i.e., by common variation in risk-adjusted returns over time. Standard error estimates are reported in parenthesis. Standard errors are clustered at the stock level, and stock fixed effects capture any fixed heteroskedasticity in the error term. Symbols *, **, and *** denote significance at 10%, 5%, and 1% levels, respectively. For each liquidity measure, the liquidity premium in basis points (presented in bold font) is the product of the liquidity coefficient and the sample average of the monthly median liquidity measure.

BBD goes from 3.4bp in 2001–2007 to 4.6bp in 2008–2014. Segmenting the latter time period into “crisis” and “post-crisis” years yields a statistically insignificant premium of 4.5bp in 2008–2009 (likely due to the small sample size) versus a statistically significant premium of 6.8bp in the 2010–2014 period. BBD and $WBBD$ only significantly explain variations in expected risk-adjusted returns in later sample periods. In each sub-period, estimated liquidity premia are much higher when estimated using expected returns (Model I) than with using expected risk-adjusted returns (Model II). As discussed previously, the four-factor model used to estimate expected risk-adjusted returns imposes more structure on the risk factors, presuming that the four-factor model is “correct”, which inflates error terms and lowers estimated liquidity premia when it is not.

To refine these observations, we explore the temporal evolution of liquidity premia at finer frequencies. To identify the incremental periodic changes in the size and significance of liquidity premia, we employ rolling four-year estimation periods that commence at the beginning of each year.⁶³ We estimate Models I and II using data from each four-year period. We use these estimates to calculate the liquidity premium and the implied 95% confidence interval for each rolling period. To avoid making restrictive assumptions about the sampling distributions of test statistics, we construct the percentile-t bootstrap confidence intervals given random samples *with replacement* of 1000 stocks per month and 999 replications. Figure 10 presents results for both BBD and $WBBD$.

⁶³ The first “four-year” period spans April 2001–December 2004, and contains less than four years of data.

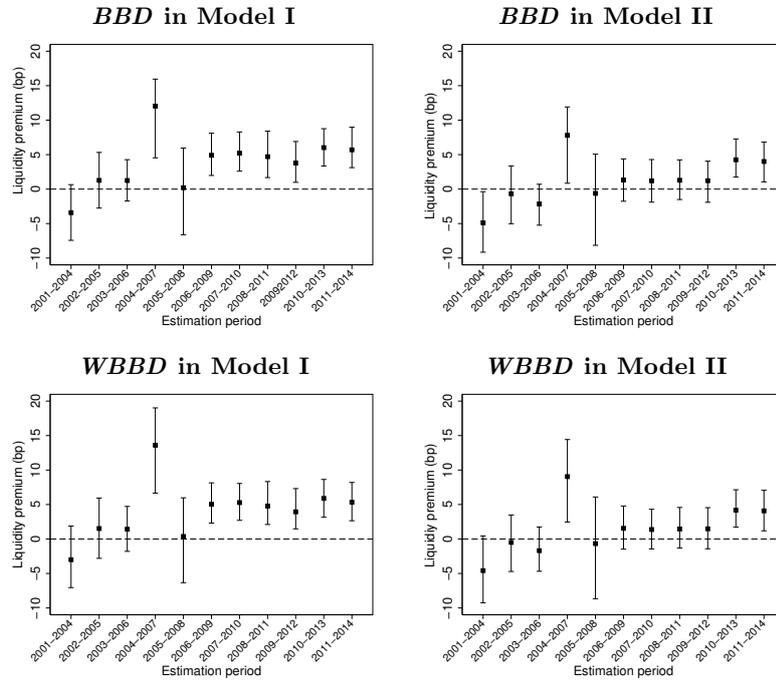


Figure 10
Evolution of liquidity premia for NYSE-listed stocks.
 Model I (returns, equation (11)) and Model II (risk-adjusted returns, equation (12)) are estimated using data from eleven four-year rolling periods, with *BBD* and *WBB* serving as *LIQ*. Each four-year period starts at the beginning of a year. Each period, liquidity premium and the associated 95% interval are defined as the product the characteristic liquidity coefficient (and the associated 95% percentile-t bootstrap confidence interval limits) and the cross-month average of the month-specific median *BBD*. Standard errors are clustered at stock level, and stock-fixed effects control for any fixed heteroskedasticity. Black squares are estimated liquidity premia, and bars span the 95% confidence intervals.

Figure 10 reinforces that liquidity premia have not fallen over the post-decimalization period. Instead, they are close to zero and insignificant in early years, before jumping sharply around 2007, and then remaining positive. The non-overlapping confidence intervals for the 2001–2004 and 2011–2014 periods for both models (i.e., for both expected returns and expected risk-adjusted returns) indicate statistically significant increases in liquidity premia.⁶⁴

7. Conclusion

The nature of liquidity provision has changed. Today, quotes are tiny, fleeting, and have little depth. As a result, quote-based measures of

⁶⁴ Qualitatively identical results obtain for three- and five-year rolling periods.

Trade-time based measures of liquidity

liquidity have become not only practically challenging to construct accurately, but also less relevant to institutional investors who often employ algorithmic order-splitting strategies. These strategies give rise to temporally-dependent transactions, complicating the statistical properties of liquidity measures that rely on trade-by-trade information. Designs of other commonly-used measures, such as Amihud’s (2002) measure, do not reflect that investors time trades according to available liquidity, rendering them noisier proxies of trading costs.

The use of noisier liquidity measures might lead researchers to underestimate the value investors place on liquidity. Our first contribution is to develop simple measures of stock liquidity that control for intradaily variations in liquidity provision and trading activity by measuring the average per-dollar price impacts of fixed-dollar volumes. We show that trade times of such fixed-dollar volumes reflect variations in available liquidity, implying that the corresponding trade-time price impacts capture the feature that trading costs may realize over short or long horizons depending how fast trading takes place. We show that our measures of liquidity are more strongly related to measures of institutional trading costs estimated by ITG, a leading provider of trade execution services, than are traditional liquidity measures.

We find that our liquidity measures better explain the cross-sections of both expected returns and risk-adjusted expected returns of NYSE-listed common stocks in the 2001–2014 period, especially in more recent years. Over this period, we find *monthly* liquidity premia estimates of 5.3 basis points for expected returns and 2.4 basis points for risk-adjusted expected returns. Estimated liquidity premia increase significantly around the years of financial crisis, and remain high thereafter. This highlights how, despite the improvements in measures of market quality post-decimalization, investors still demand non-trivial compensation for holding less liquid stocks.

We conclude by highlighting the broad value of trade-time approaches. The inverse of trade time of a fixed-dollar volume measures stock-specific trading activity. One can use these trade-time intervals in Autoregressive Conditional Duration models (Engle and Russell (1998)) to retrieve an *i.i.d.* error term structure that is violated by today’s temporally-dependent trades. Barardehi, Bernhardt, and Ruchti (2017) show that when one uses the same fixed-dollar value target for each stock in a portfolio, one can test the microstructure invariance hypothesis of Kyle and Obizhaeva (2016) at the stock level. They use this approach to highlight the importance of systematic risk for invariance theories. Barardehi and Bernhardt (2017) use trade-time measures to investigate how trading outcomes vary with the intensity of trading activity, after controlling for time of day; and how outcomes vary with expected vs. unexpected trading activity. So, too, one can shed light on the dynamics

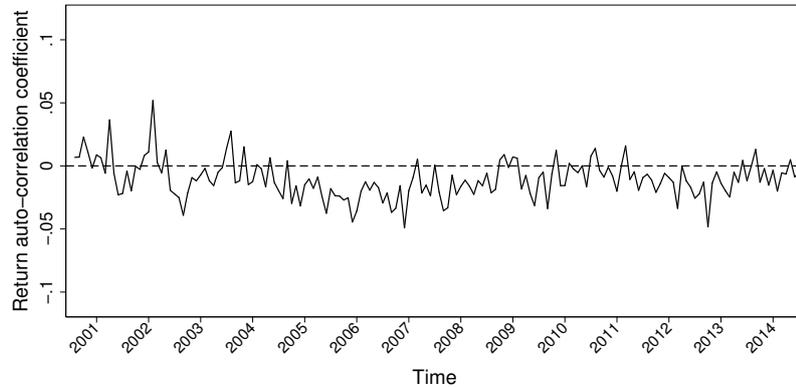


Figure 11
Cross-stock median of trade-time return auto-correlation.
 Each month, $r_j(k) = \rho_j^0 + \rho_j^1 r_j(k-1) + e_j(k)$ is estimated stock-by-stock. Monthly cross-stock medians of $\hat{\rho}_j^1$ are plotted against month of observation.

of trading—uncovering how trading outcomes vary at a given level of trading activity according to how trading activity evolves to *reach* that level. For example, conditional on reaching a given level of trading activity, increases in activity are associated with momentum in returns, and decreases in activity are associated with reversion.

Appendix

1 Temporal dependence of trade-time returns

Our approach demands that we adequately aggregate trades to address the temporal dependence of trades (and returns) driven by dynamic order splitting. We now show that the returns measured over trade-time intervals are barely serially-correlated for a typical stock. To do this we estimate a simple AR(1) of trade-time returns of stock j , $r_j(k)$, every month—note that the structure that allows fixed-month effects accounts for the fact that $V_{j,t}$ is updated on a monthly basis. Thus, each month, we estimate $r_j(k) = \rho_j^0 + \rho_j^1 r_j(k-1) + e_j(k)$ for each stock j : the autocorrelation coefficient $\hat{\rho}_j^1$ measures the temporal dependence in returns over trade-time intervals. Figure 11 shows that the first order autocorrelation of trade-time returns (of fixed-dollar positions) is very close to zero, and seems to be independent of general market conditions. This indicates that we aggregate sufficiently to address the temporal dependence of transactions in modern markets.

2 Fama-MacBeth estimates

We present Fama-MacBeth estimation results for the April 2001–December 2014 sample that includes 165 monthly cross-sections. Table 9 presents Fama-MacBeth estimates of Model I and Model II. These results are consistent with those of fixed-effects GLS (see Table 4). Table 10 presents the Fama-MacBeth estimation results of the orthogonal decomposition exercise, comparing *BBD* against other measures. As expected, and consistent with the fixed-effects GLS estimation results, only *WBBD* seems to perform better than *BBD*.

To obtain liquidity premia from Fama-French estimates, we follow Ben-Rephael et al. (2015). The Fama-MacBeth approach fits the model month-by-month, and draws statistical inference using the empirical distributions of coefficients.⁶⁵ Each month, we multiply the estimated liquidity coefficient by the liquidity measure of each stock ($LIQ_{j,t-1}$) to obtain month-stock-specific liquidity premia. The liquidity premium is then given by the sample-wide average of the month-specific medians of these premia.

We find that the extent of liquidity premia obtained from Fama-MacBeth estimates are barely affected by microstructure noise. Asparouhova et al. (2010, 2013) argue that potential cross-sectional correlation between liquidity and microstructure noise, e.g., bid-ask bounce, may cause OLS estimates to be biased upwards, i.e., estimated liquidity premia may be biased upward. They propose a set of corrections for cross-sectional OLS estimates of stock and portfolio returns. They propose Weighted-OLS rather than OLS, to fit the first-stage monthly cross-sectional regressions in the Fama-MacBeth approach. A commonly-used correction weighs stock j 's observation in month t by this stock's return in the previous month, $(1+r_{j,t-1})$. Table 11 shows that our qualitative findings are robust to adjustments for microstructure noise: estimated premia only decline marginally vis à vis those reported in Table 9.⁶⁶ This finding is expected because (i) our sample consists of very liquid NYSE-listed stocks, and (ii) we focus on the post decimalization era where drivers of microstructure noise, e.g., tick size, have fallen sharply.

⁶⁵ Obviously, month dummies are dropped in the Fama-MacBeth approach.

⁶⁶ The corrections proposed by Asparouhova et al. (2010, 2013) were only investigated in cross-sectional settings; and not in panel data settings, where autocorrelation in error terms is likely. In panel settings the lagged gross returns that serve as weights in this correction may be temporally-correlated. This may introduce complications when calculating clustered standard errors if the extent of temporal correlations in error terms is related to lagged gross returns. With this caveat, we nonetheless implement their proposed correction in our weighted-GLS panel estimation approach, recognizing that the extent to which it impacts clustered standard errors has not been established. In these unreported results, we find that, as with the Fama-MacBeth estimates, there are no meaningful changes in estimates of liquidity premia and clustered standard errors when we correct for microstructure noise in this way.

Table 9: Fama-MacBeth estimation results, April 2001–December 2014.

Co-variate	Model I (returns)					Model II (risk-adjusted returns)						
	0.197*** (0.049)	0.285*** (0.008)	0.151** (0.072)	0.099 (0.007)	0.012*** (0.005)	0.226*** (0.058)	0.326*** (0.078)	0.203** (0.096)	0.010 (0.007)	0.016*** (0.005)	0.197*** (0.061)	1.832*** (0.534)
<i>WBBD</i>												
<i>AML</i>												
<i>PSP</i>												
<i>LAMBDA</i>												
<i>EPSP</i>												
β_{ML}^{ML}	-0.004* (0.002)	-0.004* (0.002)	-0.004* (0.002)	-0.004* (0.002)	-0.004* (0.002)	-0.004* (0.002)	-0.004* (0.002)	-0.004* (0.002)	-0.004* (0.002)	-0.004* (0.002)	-0.004* (0.002)	-0.003** (0.002)
β_{ML}^{ML}	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)
β_{ML}^{ML}	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)
β_{ML}^{ML}	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)
β_{ML}^{ML}	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)
β_{ML}^{ML}	0.005*** (0.002)	0.005*** (0.002)	0.005*** (0.002)	0.005*** (0.002)	0.005*** (0.002)	0.005*** (0.002)	0.005*** (0.002)	0.005*** (0.002)	0.005*** (0.002)	0.005*** (0.002)	0.005*** (0.002)	0.003** (0.002)
$\ln(M_{t-1})$	-0.0004 (0.0001)	-0.0004 (0.0001)	-0.001** (0.0004)	-0.001** (0.0004)	-0.0006 (0.0004)	-0.0006 (0.0004)	-0.0006 (0.0004)	-0.0006 (0.0004)	-0.0006 (0.0004)	-0.0006 (0.0004)	-0.0006 (0.0004)	-0.0002 (0.0003)
MOM_{t-1}^1	0.002 (0.005)	0.002 (0.005)	0.002 (0.005)	0.002 (0.005)	0.002 (0.005)	0.002 (0.005)	0.002 (0.005)	0.002 (0.005)	0.002 (0.005)	0.002 (0.005)	0.002 (0.005)	0.001 (0.005)
MOM_{t-1}^2	-0.003 (0.004)	-0.003 (0.004)	-0.003 (0.004)	-0.003 (0.004)	-0.003 (0.004)	-0.003 (0.004)	-0.003 (0.004)	-0.002 (0.003)	-0.002 (0.003)	-0.002 (0.003)	-0.002 (0.003)	-0.002 (0.003)
DYD_{t-1}	-0.018 (0.071)	-0.019 (0.070)	-0.013 (0.071)	-0.012 (0.071)	-0.014 (0.071)	-0.015 (0.071)	-0.015 (0.071)	0.005 (0.070)	0.011 (0.070)	0.009 (0.071)	0.009 (0.071)	0.002 (0.070)
SD_{t-1}	0.014 (0.021)	0.014 (0.020)	0.017 (0.021)	0.016 (0.020)	0.016 (0.021)	0.016 (0.021)	0.016 (0.021)	0.008 (0.021)	0.008 (0.021)	0.008 (0.021)	0.008 (0.021)	-0.042** (0.017)
# of months	164	164	164	164	164	164	164	164	164	164	164	164

The table reports summaries of Fama-MacBeth estimates of Model I (returns, Equation (11)) and Model II (risk-adjusted returns, Equation (12)) using the five different liquidity measures. Numbers in the parenthesis are estimated Newey-West standard errors given two lags. Symbols *, **, and *** denote significance at 10%, 5%, and 1% levels, respectively. For each liquidity measure, the liquidity premium is the average of the month-specific median product of the liquidity coefficient and level of liquidity.

Trade-time based measures of liquidity

Table 10
Fama-MacBeth estimates of characteristic liquidity coefficients for orthogonally-decomposed measures with respect to BBD.

<i>LIQ</i>	Model I (returns)					
	<i>WBBD</i>	<i>AML</i>	<i>AMH</i>	<i>PSP</i>	<i>LAMBDA</i>	<i>EPSP</i>
z^F	-0.266 (0.314)	0.278*** (0.061)	0.262*** (0.069)	0.180*** (0.052)	0.216*** (0.073)	0.164*** (0.056)
\tilde{z}^F	0.673 (0.504)	-0.175** (0.086)	-0.023** (0.011)	-0.007 (0.007)	-0.130 (0.100)	0.169 (0.835)
<i>LIQ</i>	Model II (risk-adjusted returns)					
	<i>WBBD</i>	<i>AML</i>	<i>AMH</i>	<i>PSP</i>	<i>LAMBDA</i>	<i>EPSP</i>
z^F	-0.341 (0.340)	0.296*** (0.071)	0.292*** (0.079)	0.190*** (0.056)	0.199*** (0.068)	0.193*** (0.062)
\tilde{z}^F	0.810 (0.546)	-0.158 (0.104)	-0.027** (0.013)	-0.005 (0.008)	-0.108 (0.102)	0.288 (0.870)

Characteristic liquidity coefficients on z_{tj}^F or \tilde{z}_{tj}^F for subsamples April 2001 to December 2014 using Model I (returns, Equation (11)) and Model II (risk-adjusted returns, Equation (12)). The orthogonal decomposition is carried out on a month-by-month basis to fit the Fama-MacBeth routine. Numbers in the parenthesis are estimated Newey-West standard errors given two lags. Symbols *, **, and *** denote significance at 10%, 5%, and 1% levels, respectively.

3 Error Structure

We impose a robust standard error structure with clustering at the firm level. Our approach follows Petersen (2009) and accounts for potential cross-firm and cross-time correlations that might be present in the panel data error terms. Failure to account for these correlations would lead to underestimated standard errors. To account for common fixed firm and time effects, we estimate a fixed-effect model using GLS and include month dummies. The other potential error term correlations in either time or cross-section dimensions are non-fixed time and non-fixed firm effects. Non-fixed firm effects are those causing temporal autocorrelations in a given firm’s error term. Non-fixed time effects make error terms of different firms related in a given period of time while the correlation differs for different pairs of firms (see Petersen (2009) or Cameron et al. (2008)).

Many potential sources of non-fixed time and firm effects exist. For example, non-fixed time effects can arise when industry-specific shocks affect most firms in an industry similarly, while leaving firms in other industries unaffected; and non-fixed firm effects can arise when a shock to stock j persists (e.g., the effect of a technology shock on stock’s j ’s performance, which decays over time).

Alternative approaches in finance try to deliver reliable standard errors. Fama and MacBeth (1973) adjust for non-fixed time effects by running cross-sectional regressions period by period. Different point estimates from different cross-sections give the empirical distributions of parameters of interest. Further statistical inference is based on the averages and standard deviations of empirical distributions; and to

Table 11: Fama-MacBeth estimates after adjusting for Microstructure noise, April 2001–December 2014.

Characteristic	Model I (returns)					Model II (risk-adjusted returns)								
	0.196*** (0.049)	0.286*** (0.070)	0.157*** (0.073)	0.010 (0.007)	0.011** (0.005)	0.164*** (0.053)	1.482*** (0.534)	0.292*** (0.068)	0.322*** (0.079)	0.204** (0.095)	0.011 (0.007)	0.015*** (0.005)	0.196*** (0.061)	1.720*** (0.557)
<i>WBBD</i>														
<i>AMH</i>														
<i>PSP</i>														
<i>LAMBDA</i>														
<i>EPSP</i>														
β_{t-1}^{mk}	-0.003 (0.002)	-0.003 (0.002)	-0.004* (0.002)	-0.004* (0.002)	-0.003* (0.002)	-0.004* (0.002)	-0.003 (0.002)	-0.003 (0.002)	-0.003 (0.002)	-0.004** (0.002)	-0.004** (0.002)	-0.004** (0.002)	-0.004** (0.002)	-0.003** (0.002)
β_{t-1}^{amb}	-0.001 (0.001)	-0.001 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	
β_{t-1}^{hml}	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	
β_{t-1}^{vwd}	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	
β_{t-1}^{m1}	0.004 (0.004)	0.004 (0.004)	0.005 (0.004)	0.005 (0.004)	0.005 (0.004)	0.005 (0.004)	0.005 (0.004)	0.005 (0.004)	0.005 (0.004)	0.004** (0.004)	0.004** (0.004)	0.004** (0.004)	0.004** (0.004)	
$\ln(M_{t-1})$	-0.004 (0.0004)	-0.004 (0.0004)	-0.004 (0.0004)	-0.004 (0.0004)	-0.004 (0.0004)	-0.004 (0.0004)	-0.004 (0.0004)	-0.004 (0.0004)	-0.004 (0.0004)	-0.004** (0.0003)	-0.004** (0.0003)	-0.004** (0.0003)	-0.004** (0.0003)	
MOM_{t-4}^1	0.002 (0.002)	0.003 (0.003)	0.003 (0.003)	0.003 (0.003)	0.003 (0.003)	0.003 (0.003)	0.003 (0.003)	0.003 (0.003)	0.003 (0.003)	0.001 (0.001)	0.001 (0.001)	0.000 (0.000)	0.000 (0.000)	
MOM_{t-12}^5	-0.003 (0.004)	-0.003 (0.004)	-0.003 (0.004)	-0.003 (0.004)	-0.003 (0.004)	-0.003 (0.004)	-0.003 (0.004)	-0.003 (0.004)	-0.002 (0.003)	-0.002 (0.003)	-0.002 (0.003)	-0.002 (0.003)	-0.002 (0.003)	
DY_{t-1}	0.013 (0.013)	0.013 (0.013)	0.013 (0.013)	0.015 (0.015)	0.013 (0.013)	0.013 (0.013)	0.013 (0.013)	0.013 (0.013)	0.013 (0.013)	0.013 (0.013)	0.013 (0.013)	0.013 (0.013)	0.013 (0.013)	
SD_{t-1}	0.020 (0.020)	0.020 (0.020)	0.020 (0.020)	0.020 (0.020)	0.020 (0.020)	0.020 (0.020)	0.020 (0.020)	0.020 (0.020)	0.020 (0.020)	0.017 (0.017)	0.017 (0.017)	0.017 (0.017)	0.017 (0.017)	

The table reports summaries of Fama-MacBeth estimates of Model I (returns, equation (11)) and Model II (risk-adjusted returns, equation (12)) using the five different liquidity measures: Numbers in the parenthesis are estimated Newey-West standard errors given two lags. Each observation is weighted by the corresponding gross return from the previous month, and Fama-MacBeth estimates are based on Weighted-OLS estimates of monthly cross-sections. Symbols *, **, and *** denote significance at 10%, 5%, and 1% levels, respectively. For each measure of liquidity, the liquidity premium reflects the average of the month-specific median product of the liquidity coefficient and the liquidity measure.

be valid, this requires temporal independence of cross-sectional point estimates to obtain unbiased estimates. Unfortunately, time dependence in our sample makes the Fama-MacBeth approach problematic in our setting.⁶⁷ In addition, Fama-MacBeth technique was originally designed to deal with *portfolios* as test assets over long time series; our focus was on stocks as test assets over a relatively short time horizon.

Some estimation strategies call for multi-dimensional clustering when a researcher suspects errors to autocorrelate in multiple dimensions (see Cameron et al. (2008)). To the best of our knowledge, however, estimation techniques that allow for multi-dimensional and non-nested clustered error terms have been developed only for ordinary least squares (OLS) estimation. OLS estimation using panel data is equivalent to pooling the information content of different periods about a given individual. Such aggregation is not appropriate here.

It is inappropriate to try to control for non-fixed time effects by clustering by time when we employ fixed-effects GLS, since the unbalanced nature of the panel makes the clusters non-nested. Most significantly, the set of stocks in our sample varies non-trivially over time, e.g., due to mergers and acquisitions. With non-nested clusters it is not clear which periods’ cross-stock error term autocorrelations may cause upward biases in estimated standard errors. These concerns lead us to control for non-fixed time effects by introducing industry-specific shocks. This allows us to account for non-fixed time effects associated with shocks that affect firms within an industry similarly.

Using four-digit GIC industry codes, we identify about 65 industries per month. For each firm j , we compute the associated average industry returns on a monthly basis excluding its own return. To de-trend, we find the excess industry return against the equally weighted monthly average return. The first difference of the monthly industry excess return is employed as proxy of non-fixed time effects (ΔYIR). Adding this variable to the model does not qualitatively alter our findings.⁶⁸

An alternative approach to dealing with non-fixed time effects is to model the serial correlation of error terms parametrically. Accordingly, we estimated the model given error terms that follow an AR(1) process, $\epsilon_{j,t} = \rho\epsilon_{j,t-1} + \nu_{j,t-1}$, with $\nu_{it} \sim iidN(0, \Sigma_\nu^2)$. More specifically, we performed a two stage estimation of standard errors, while equations (11) and (12) are estimated based on a fixed-effects GLS strategy. As when we cluster at the firm level, the autoregressive specification of errors does not affect point estimates. The model is estimated using

⁶⁷ Petersen (2009) identifies downward biases in Fama-MacBeth estimated standard errors in the presence of firm effects.

⁶⁸ Nichols and Schaffer (2007) note that with differently-sized clusters, inference could be problematic. This is the case in our sample; however, by clustering at the firm level we get a large number of clusters, so asymptotic properties should hold.

the fixed-effects GLS strategy, residuals are used to estimate ρ , and the covariance matrix is estimated accordingly. When we do so, we find that regardless of the liquidity measure employed, the associated standard errors are slightly *smaller* than those reported in Table 4.

This analysis supports our allowance in the model for firm clustering, along with stock fixed effects and time dummies. These assumptions are designed to reduce the downward biases in estimated standard errors. Note that it is not necessary to cluster in the time dimension since the inclusion of both time dummies and ΔYIR presumably capture most of the commonalities in the time variation of stock returns.

4 Robustness to construction horizons of liquidity measures

In Section 3, we described that measures of stock liquidity in each month are constructed using observations from the past three months. This design was employed to provide enough data to estimate accurate low-frequency Amihud measures. However, it also makes our liquidity measures serially-correlated. We can and do address this in our regression analysis: we cluster error terms *at the stock level* to account for any stock-specific unobserved autocorrelation. As an additional robustness check, we now replicate the entire analysis using liquidity measures that are constructed using observations from the previous month, rather than from the previous three months. This exercise verifies that our findings are robust to the backward-looking horizon employed to construct measures of stock liquidity.

When liquidity measures are constructed based on the prior month’s observations, the results are almost identical to those obtained when measures are constructed based on the three-month rolling moments (presented throughout the paper).⁶⁹ Comparisons of estimates in Table 12 and Table 4 show that most coefficient estimates change marginally when we use liquidity measures based on observations from the previous month.

The remainder of the analyses is also robust. In unreported results, we confirm that *BBD* and *WBBD* outperform the other measures regardless of whether we construct measures using observations from the previous month or the previous three months. Consistent with our previous findings, Table 13 shows how liquidity premia rises over time and that the increase in the size and statistical significance of premia are not driven by the years of financial crisis.

⁶⁹ Results available upon request.

Trade-time based measures of liquidity

Table 12: Model I and II estimation results, April 2001–December 2014.

Co-variate	Model I (returns)										Model II (risk-adjusted returns)									
	4.2	4.4	4.4	0.4	0.4	3.1	3.0	0.2	1.5	1.8	0.1	0.0	1.6	-1.2	-0.1					
$BBID_{t-1}$	0.108*** (0.025)																			
$WBBD_{t-1}$		0.184*** (0.041)																		
AML_{t-1}			0.038** (0.015)																	
AMH_{t-1}				0.006** (0.003)																
PSP_{t-1}					0.025*** (0.004)															
$LAMBDA_{t-1}$						0.037* (0.021)														
$EPSP_{t-1}$							0.020 (0.026)													
β_{t-1}^{mk}	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)															
β_{t-1}^{mb}	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)															
β_{t-1}^{hml}	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)															
β_{t-1}^{umd}	0.004*** (0.001)	0.004*** (0.001)	0.004*** (0.001)	0.004*** (0.001)	0.004*** (0.001)															
BM_{t-1}	0.021*** (0.003)	0.021*** (0.003)	0.023*** (0.003)	0.022*** (0.003)	0.022*** (0.003)	0.021*** (0.003)	0.022*** (0.003)	0.023*** (0.003)	0.017*** (0.003)	0.017*** (0.003)	0.018*** (0.003)	0.018*** (0.003)	0.017*** (0.003)	0.018*** (0.003)	0.018*** (0.003)					
$\ln(M_{t-1})$	-0.023*** (0.001)	-0.023*** (0.001)	-0.025*** (0.001)	-0.025*** (0.001)	-0.024*** (0.001)	-0.023*** (0.001)	-0.024*** (0.001)	-0.025*** (0.001)	-0.022*** (0.001)	-0.021*** (0.001)	-0.022*** (0.001)	-0.022*** (0.001)	-0.021*** (0.001)	-0.022*** (0.001)	-0.022*** (0.001)					
COM_{t-1}^4	0.013*** (0.002)	0.017*** (0.002)	0.017*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.017*** (0.002)	0.018*** (0.002)	0.018*** (0.002)												
COM_{t-1}^{12}	0.002* (0.001)	0.000 (0.001)	0.000 (0.001)	0.001 (0.001)	0.001 (0.001)	0.000 (0.001)	0.001 (0.001)	0.001 (0.001)												
DYD_{t-1}	-0.110 (0.067)	-0.110 (0.067)	-0.116* (0.066)	-0.116* (0.066)	-0.112* (0.066)	-0.112* (0.066)	-0.116* (0.066)	-0.118* (0.066)	-0.100 (0.063)	-0.100 (0.063)	-0.102 (0.063)	-0.103 (0.063)	-0.100 (0.063)	-0.104 (0.063)	-0.103 (0.063)					
SD_{t-1}	0.038*** (0.014)	0.038*** (0.014)	0.037*** (0.014)	0.037*** (0.014)	0.036*** (0.014)	0.036*** (0.014)	0.037*** (0.014)	0.036*** (0.014)	-0.001 (0.013)	-0.001 (0.013)	-0.002 (0.013)	-0.002 (0.013)	-0.002 (0.013)	-0.003 (0.013)	-0.002 (0.013)					
Within R^2	0.30	0.30	0.30	0.30	0.30	0.30	0.30	0.30	0.06	0.06	0.06	0.06	0.06	0.06	0.06					
Log-lik	161558.7	161565.2	161510.1	161510.0	161555.4	161507.3	161502.9	161544.0	161533.7	161533.8	161534.4	161533.7	161547.8	161534.4	161533.7					
Observations	168110	168110	168110	168110	168110	168110	168110	168110	168110	168110	168110	168110	168110	168110	168110					

The table reports summaries of panel estimates of Model I and Model II using five different liquidity measures. Liquidity measures are constructed using observations from the previous month. Month dummies capture common variations over time. Estimated standard errors are in parenthesis. Standard errors are clustered at the stock level, and stock fixed effects capture any fixed heteroskedasticity in the error term. Symbols *, **, and *** denote significance at 10%, 5%, and 1% levels, respectively. The liquidity premium for each liquidity measure is the product of the liquidity coefficient and the sample average of the monthly median liquidity measure.

Table 13
Characteristic liquidity premia by sub-sample.

Model	LIQ	Apr 2001–Dec 2007	Jan 2008–Dec 2014	Jan 2008–Dec 2009	Jan 2010–Dec 2014
Model I	BBD_{t-1}	0.030	0.146***	0.085*	0.202***
		(0.029)	(0.044)	(0.052)	(0.059)
		1.430	4.497	4.436	4.526
	$WBBD_{t-1}$	0.059	0.250***	0.159*	0.256***
(0.046)		(0.072)	(0.086)	(0.094)	
	1.745	4.655	4.937	3.491	
Model II	BBD_{t-1}	-0.023	0.058	-0.014	0.162***
		(0.029)	(0.041)	(0.052)	(0.058)
		-1.096	1.798	-0.730	3.635
	$WBBD_{t-1}$	-0.020	0.108*	-0.006	0.214**
(0.045)		(0.066)	(0.087)	(0.092)	
	-0.587	2.015	-0.183	2.908	

The table reports liquidity coefficient estimates from panel estimates of Model I (returns, equation (11)) and Model II (risk-adjusted returns, equation (12)) using the BBD and $WBBD$ liquidity measures. Measures of liquidity are constructed using *observations from the previous month*. Month dummies capture common variation caused by market condition changes, i.e., by common variation in risk-adjusted returns over time. Standard error estimates are reported in parenthesis. Standard errors are clustered at the stock level, and stock fixed effects capture any fixed heteroskedasticity in the error term. Symbols *, **, and *** denote significance at 10%, 5%, and 1% levels, respectively. For each liquidity measure, the liquidity premium (presented in **bold font**) is the product of the liquidity coefficient and the sample average of the monthly median liquidity measure.

References

- Acharya, V.V., L.H. Pedersen. 2005. Asset pricing with liquidity risk. *Journal of Financial Economics* 77, 375–410.
- Akbas, F., W.J. Armstrong, R. Petkova. 2011. Idiosyncratic volatility of liquidity and expected stock returns. Working paper, Purdue University.
- Almgren, R. and N. Chriss (1999). Value under liquidation. *Risk* 12(12), 61–63.
- Almgren, R. and N. Chriss (2001). Optimal execution of portfolio transactions. *Journal of Risk* 3, 5–40.
- Amihud, Y. 2002. Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets* 5(1), 31–56.
- Amihud, Y., H. Mendelson. 1986. Asset pricing and the bid-ask spread. *Journal of Financial Economics* 17(2), 223–249.
- Angel, J., L. Harris, C. Spatt. 2011. Equity Trading in the 21st Century. *Quarterly Journal of Finance* 1, 1–53.
- Asparouhova, E., H. Bessembinder, and I. Kalcheva. 2010. Liquidity biases in asset pricing tests. *Journal of Financial Economics* 96, 215–237.
- Asparouhova, E., H. Bessembinder, and I. Kalcheva. 2013. Noisy prices and inference regarding returns. *Journal of Finance* 68(2), 655–714.
- Barardehi, Y.H., D. Bernhardt. 2017. The dynamics of intraday trade time outcomes. Working paper.
- Barardehi, Y.H., D. Bernhardt, T.G. Ruchti. 2017. Systematic risk and microstructure invariance. Working paper.
- Ben-Rephael, A., O. Kadan, A. Wohl. 2015. The diminishing liquidity premium. *Journal of Financial and Quantitative Analysis* 50, 197–229.
- Bernhardt D., E. Hughson. 2002. Intraday trade in dealership markets. *European Economic Review* 46(9), 1697–1732.

Trade-time based measures of liquidity

- Brennan, M.J., T. Chordia, A. Subrahmanyam. 1998. Alternative factor specifications, security characteristics, and the cross-section of expected stock returns. *Journal of Financial Economics* 49, 345–373
- Brennan, M.J., A. Subrahmanyam. 1996. Market microstructure and asset pricing: On the compensation for illiquidity in stock returns. *Journal of Financial Economics* 41, 441–464.
- Borkovec, M., H.G. Heidle (2010). Building and evaluating a transaction cost model: a primer. *Journal of Trading* 5(2), 57–77.
- Boulatov, A., D. Bernhardt, I. Larionov. 2016. Predatory and Defensive Trading in a Dynamic Model of Optimal Execution by Multiple Traders. Working Paper, Higher School of Economics.
- Cameron, A.C., J.B. Gelbach, D.L. Miller. 2008. Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics* 90(3), 414–427.
- Conrad, J., S. Wahal, J. Xing. 2015. High frequency quoting, trading, and the efficiency of prices. *Journal of Financial Economics* 116(2), 271–291.
- Chakrabarty, B., B. Li, V. Nguyen, and R. Van Ness. 2006. Trade classification algorithms for electronic communication networks. *Journal of Banking and Finance* 31, 3806–3821.
- Chordia, T., R. Roll, A. Subrahmanyam. 2000. Commonality in liquidity. *Journal of Financial Economics* 56, 3–28.
- Chordia, T., R. Roll, A. Subrahmanyam. 2011. Recent trends in trading activity and market quality. *Journal of Financial Economics* 101(2), 243–263.
- Collin-Dufresne, P. and V. Fos. 2015. Do Prices Reveal the Presence of Informed Trading? *Journal of Finance* 70(4), 1555–1582.
- Deuskar, P., T.C. Johnson. 2011. Market liquidity and flow-driven risk. *Review of Financial Studies* 24(3), 721–753.
- Dufour, A., R.F. Engle. 2000. Time and the price impact of a trade. *Journal of Finance* 55(6), 2467–2598.
- Easley D., M.M. Lopez de Pardo, M. O’Hara. 2012. Flow toxicity and liquidity in a high frequency world. *Review of Financial studies* 25(5), 1457–1493.
- Ellis, K., R. Michaely, and M. O’Hara. 2000. The accuracy of trade classification rules: Evidence from Nasdaq. *Journal of Financial and Quantitative Analysis* 35, 529–552.
- Engle, R.F. 2000. The econometrics of ultra-high frequency data. *Econometrica* 1, 1–22.
- Engle, R.F., J.R. Russell. 1998. Autoregressive conditional duration: A new model for irregularly spaced transaction data. *Econometrica* 66, 1127–1162.
- Fama, E.F., J.D. MacBeth. 1973. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy* 81(3), 607–636.
- Fama, E.F., K.R. French. 1999. The cross-section of expected stock returns. *Journal of Finance* 72(2), 427–465.
- Feldhütter, P. 2012. The same bond at different prices: Identifying search frictions and selling pressures. *Review of Financial Studies* 25(4), 1155–1206.
- Glosten, L., L. Harris. 1988. Estimating the components of the bid-ask spread. *Journal of Financial Economics* 21, 123–142.
- Goldstein, M.A., K.A. Kavajecz. 2000. Eighth, sixteenth, and market depth: changes in tick size and liquidity provision on the NYSE. *Journal of Financial Economics* 56(1), 125–149.

- Gourièroux, C., J. Jasiak, and G. Le Fol. 1999. Intra-day market activity. *Journal of Financial Markets* (2): 193–226.
- Goyenko, R.Y., C.W. Holden, C.A. Trzcinka. 2009. Do liquidity measures measure liquidity? *Journal of Financial Economics* 92(2), 153–181.
- Hasbrouck, J. 2009. Trading costs and returns for U.S. equities: Estimating effective costs from daily data. *Journal of Finance* 64(3), 1445–1477.
- Hasbrouck, J. and G. Saar. 2013. Low-latency trading. *Journal of Financial Markets* (16), 646–679.
- Hendershott, T., C.M. Jones, A.J. Menkveld. 2011. Does algorithmic trading improve liquidity? *Journal of Finance* 66(1), 1–33.
- Holden, C.W., S.E. Jacobsen (2014). Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions. *Journal of Finance* 69(4), 1747–1785.
- ITG Inc. (2007). ITG ACE–Agency Cost Estimator: A model description. *ITG Inc. White Paper*.
- Jensen, M.C., F. Black, M.S. Scholes. 1972. The Capital Asset Pricing Model: Some Empirical Tests. Michael C. Jensen, *STUDIES IN THE THEORY OF CAPITAL MARKETS*, Praeger Publishers Inc., 1972.
- Jones, C.M. 2002. A century of stock market liquidity and trading costs. Working paper, Columbia University.
- Jones, C.M., M.L. Lipson. 2001. Sixteenths: direct evidence on institutional execution costs. *Journal of Financial Economics* 59(2), 253–278.
- Kim, S., D. Murphy. 2013. The impact of high-frequency trading on stock market liquidity measures. Working paper.
- Kyle, A.S. 1985. Continuous auctions and insider trading. *Econometrica* 53(6), 1315–1335.
- Kyle, A.S. and A. Obizhaeva. 2016. Market microstructure Invariance: empirical hypotheses. *Econometrica* 84(4), 345–1404.
- Mazza P. 2015. Rethinking zero returns in the liquidity puzzle of a limit order market, *Finance* 36, 7–36.
- O’Hara, M., C. Yao, and M. Ye. 2014. What’s not there: odd-lots and market data. *Journal of Finance* 69(5): 2199–2236.
- O’Hara, M. 2015. High frequency market microstructure. *Journal of Financial Economics* 116(2), 257–270.
- Pástor, L., R.F. Stambaugh. 2003. Liquidity risk and price discovery. *Journal of Political Economy* 111(3), 642–685.
- Petersen, M.A. 2009. Estimating standard errors in finance panel data sets: Comparing approaches. *Review of Financial Studies* 22(1), 435–480.
- Roll, R. 1984. A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance* 39, 1127–1139.
- Sadka, R. 2006. Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk. *Journal of Financial Economics* 80, 309–349.