Foundations and Trends® in Finance
Vol. 8, No. 4 (2013) 263–365

DOI: 10.1561/0500000044

the essence of knowledge

The Empirical Analysis of Liquidity

Craig W. Holden
Kelley School of Business
Indiana University, USA
cholden@indiana.edu

Stacey Jacobsen
Cox School of Business
Southern Methodist University, USA
staceyj@cox.smu.edu

Avanidhar Subrahmanyam
Andersen School of Management,
University of California, Los Angeles, USA
subra@anderson.ucla.edu

Contents

Abstract

We provide a synthesis of the empirical evidence on market liquidity. The liquidity measurement literature has established standard measures of liquidity that apply to broad categories of market microstructure data. Specialized measures of liquidity have been developed to deal with data limitations in specific markets, to provide proxies from daily data, and to assess institutional trading programs. The general liquidity literature has established local cross-sectional patterns, global cross-sectional patterns, and time-series patterns. Commonality in liquidity is prevalent. Certain exchange designs enhance market liquidity: a limit order book for high volume markets, a hybrid exchange for low volume markets, and multiple competing exchanges. Automatic execution increases speed, but increases spreads. A tick size reduction yields a large improvement in liquidity. Providing ex-post transparency to an otherwise opaque market dramatically improves liquidity. Opening up the limit order book improves liquidity. Regulatory reforms that increase the number of competitive alternatives, move toward linking them up, and level the playing field between exchanges improves liquidity. High-frequency traders trade in both a passive, liquidity-supplying manner and an aggressive, liquidity-demanding manner. Their overall impact improves both liquidity and price efficiency, but concerns remain regarding occasional trading glitches, order anticipation strategies, and latency arbitrage at the expense of slow traders. The liquidity and corporate finance literature provides abundant evidence that liquidity is beneficial in many corporate settings: liquidity increases the power of governance via exit, reduces the cost of governance via intervention, facilitates the entrance of informed traders who produce valuable information about the firm, enhances the effectiveness of equity-based compensation to managers, reduces the cost of equity financing, mitigates trading frictions investors encounter when trading in the market to recreate a preferred payout policy, and lowers the immediate transaction costs and subsequent liquidity costs for firms conducting large share repurchases. Further, the influence goes both ways. There is evidence that firms influence their own liquidity through a broad range of corporate decisions including internal governance standards, equity issuance form and pricing, share repurchases, acquisition targets, and disclosure timeliness and quality. The literature on liquidity and asset pricing demonstrates that both average liquidity cost and liquidity risk are priced, liquidity enhances market efficiency, and liquidity strengthens the arbitrage linkage between related markets. We conclude with directions for future research.

C. W. Holden, S. Jacobsen and A. Subrahmanyam. *The Empirical Analysis of Liquidity*. Foundations and Trends® in Finance, vol. 8, no. 4, pp. 263–365, 2013.

DOI: 10.1561/0500000044.

1

Introduction

This literature survey reviews the empirical analysis of liquidity. We start with an overview of how liquidity is measured and specialized issues in liquidity measurement. Next, we review what is known about cross-sectional and time-series patterns in liquidity, commonality in liquidity, the impact of exchange design, the impact of exogenous policy shifts (such as the reductions in the minimum tick size and changes in transparency of trade reporting) on liquidity, and the impact of high-frequency traders on liquidity. We then review how liquidity relates to the corporate finance literature, including to governance, executive compensation, capital structure, and payout policy. We next review how liquidity influences the asset pricing literature, including return differentials due to average liquidity cost, liquidity premia for systematic liquidity risks, the impact of liquidity on market efficiency, and the impact of liquidity on the law of one price. Finally, we discuss open questions and opportunities for future research.

What is *market liquidity*? A simple definition is the ability to trade a significant quantity of a security at a low cost in a short time.[1] Thus, liquidity is a multi-dimensional concept encompassing quantity, cost, and time dimensions. We discuss liquidity measures of each dimension separately and in combination.

The modern theory of market microstructure formulates the trading process as an interaction between liquidity suppliers and liquidity demanders. *Liquidity suppliers* offer to buy a particular security (e.g., stock, bond, option, futures, currency, etc.) at a *bid* price or sell it at an *offer* price. Then *liquidity demanders* agree to buy the security at the offer price or sell it at the bid price and a trade is born. Liquidity matters because it represents the cost, quantity, and time of a trade to the liquidity demander. Equivalently, it represents the profit, quantity, and time of a trade to the liquidity supplier.

In a *pure limit order book exchange*,[2] each trader can decide moment-by-moment if they want to supply liquidity by submitting a non-marketable limit order[3] to replenish the limit order book or demand liquidity by submitting a market order or a marketable limit order[4] to deplete the limit order book. In a *pure dealer exchange*, dealers supply liquidity by quoting bid and offer prices and other traders demand liquidity by submitting a market buy (sell) order to trade at the current offer (bid) price. In a *hybrid exchange*, both non-marketable limit orders and dealers supply liquidity and other traders demand liquidity. In a *search market*, a liquidity demander seeks potential liquidity

[1] Market liquidity is also called the *transactional* liquidity of a securities market. Market liquidity is different concept than the *funding* liquidity of market makers or the *cash flow* liquidity of a bank.

[2] For simplicity, we use the word exchange to refer to any type of trading venue.

[3] A limit order is an offer to buy or sell a specified quantity at a specified limit price. A non-marketable limit order is a limit buy (sell) order with a limit price below the current offer price (above the current bid price). It cannot execute immediately and must wait on the limit order book for a counterparty to trade with.

[4] A market order is a request to buy or sell a specified quantity at currently available price(s). It will execute in full immediately. A marketable limit order is a limit buy (sell) order with a limit price greater than or equal to the current offer price (less than or equal to the current bid price). It will execute immediately up to (down to) and including the limit price.

suppliers, who offer to buy or sell at a particular price, then decides whether to trade at the quoted price.

Twenty-first century trading has been transformed and continues to change. Electronic trading has almost entirely replaced floor-based trading on a global basis and across all asset classes [Jain, 2005, Johnson, 2010]. Algorithmic trading increasingly dominates manual trading on a global basis and across all asset classes [Johnson, 2010, Boehmer et al., 2014]. Trading has become much faster and continues to accelerate [Angel et al., 2011]. In its ever evolving form, trading still comes down to the interaction between liquidity suppliers ("makers") and liquidity demanders ("takers").

We find that the liquidity measurement literature has established standard measures of liquidity that apply to broad categories of market microstructure data. Specialized measures of liquidity have been developed to deal with data limitations in specific markets (e.g., futures, U.S. corporate bonds, U.S. equity), to provide proxies from daily data, and to assess institutional trading programs.

We find that the liquidity literature has established local cross-sectional patterns (liquidity is positively related to dollar volume and price level and negatively related to volatility and size), global cross-sectional patterns (liquidity is positively related to judicial efficiency, accounting standards, and political stability) and time-series patterns (liquidity exhibits seasonality, declines during crisis periods, and varies around macroeconomic announcements). Commonality in liquidity[5] is prevalent. Certain exchange designs enhance market liquidity: limit order book for high volume markets, hybrid for low volume markets, and multiple competing exchanges. Automatic execution increases speed, but increases spreads. A tick size reduction yields a large improvement in liquidity as measured by average trade-weighted effective spread. These benefits are concentrated in small trades, but large trades are typically not harmed even net of the reduction in depth. Institutional traders have adapted their trading strategies to smaller tick sizes. Adding ex-post transparency to an otherwise opaque

[5]Commonality in liquidity is a common component in liquidity variation across securities markets.

market dramatically improves liquidity. Adding ex-ante limit order book transparency to relatively transparent market causes a more modest improvement in liquidity. Regulatory reforms that increase the number of competitive alternatives, move toward linking them up, and level the playing field between exchanges have improved liquidity on both the cost and speed dimensions. High-frequency traders trade in both a passive, liquidity-supplying manner and an aggressive, liquidity-demanding manner. Their overall impact improves both liquidity and price efficiency, but concerns remain regarding occasional trading glitches, order anticipation strategies, and latency arbitrage at the expensive of slow traders.

We find that the literature on liquidity and corporate finance provides abundant evidence that liquidity is beneficial in many settings: liquidity increases the power of governance via "exit," reduces the cost of governance via intervention, facilitates the entrance of informed traders who produce valuable information about the firm, enhances the effectiveness of equity-based compensation to managers, reduces the cost of equity financing, mitigates trading frictions investors encounter when trading in the market to recreate a preferred payout policy, and lowers the immediate transaction costs and subsequent liquidity costs for firms conducting large share repurchases. Further, the influence goes both ways. There is evidence that firms influence their own liquidity through a broad range of corporate decisions including internal governance standards, equity issuance form and pricing, share repurchases, acquisition targets, and disclosure timeliness and quality. Overall, equity market liquidity can lead to firm value gains via both increases to the cash flows of the firm and decreases in the discount rate.

We find that the literature on liquidity and asset pricing demonstrates that both average liquidity cost and liquidity risk are priced, liquidity enhances market efficiency, and liquidity strengthens the arbitrage linkage between related markets.

This review is organized as follows. In Section 2, we consider the approaches taken to measure liquidity. Section 3 considers cross-sectional and time-series patterns in liquidity, commonality in liquidity, the impact of exchange design, the impact of exogenous policy shifts (such as the reductions in the minimum tick size and changes in transparency on trade reporting requirements) on liquidity, and the impact of high-frequency traders. Section 4 analyzes the relation between liquidity and corporate financial decisions. Section 5 explores the impact of liquidity on asset pricing, and Section 6 concludes with directions for future research.

2

How Liquidity is Measured

In this section, we provide an overview of how liquidity is measured. We begin with standard measures of liquidity that apply to two broad categories of market microstructure data: (1) trade and quote data (e.g., NYSE Trade And Quote (TAQ) for U.S. markets, Thompson Reuters Tick History (TRTH) for global markets, etc.) and (2) limit order book data (e.g., NASDAQ ITCH, Xetra Order Book, etc.). Then we turn to specialized issues in liquidity measurement that address specific issues or apply to specific markets.

2.1 Standard measures of liquidity

First, we examine how the cost dimension is measured (i.e., quoted spread, effective spread, realized spread, and price impact). Second, we examine how the quantity dimension is measured (i.e., quoted depth, slope of the limit order book, and slope of the price function). Third, we examine how the time dimension is measured (i.e., execution speed, partial fill rate, complete fill rate, cancellation rate, and resilience).

Liquidity is often measured relative to the Best Bid and Offer (BBO), where *best* is evaluated from the liquidity demander's point of view. For a single exchange, the *best offer* is the lowest offer price (i.e., lowest limit sell price and/or lowest dealer offer price) and the *best bid* is the highest bid price (i.e., highest limit buy price and/or highest dealer bid price). Similarly, when multiple exchanges trade the same security, the *consolidated best offer* is the lowest offer price across exchanges and the *consolidated best bid* is the highest bid price across exchanges.

2.1.1 The cost dimension

We begin with standard measures of the cost dimension. Our first measure is the percent quoted spread at time t, which is defined as

$$Percent\ Quoted\ Spread_t = Ln(O_t) - Ln(B_t), \tag{2.1}$$

where O_t is the best offer price at time t and B_t is the best bid price at time t.[1] This measure can be applied to a single exchange or to a consolidated market containing multiple exchanges. It is often aggregated over a period of time (e.g., an hour, a day, a month, etc.) by computing the *time-weighted average* of the percent quoted spread, where each quote observation is weighted by the amount of time that the quote observation is in-force. The percent quoted spread can be viewed as a liquidity demander's cost of trading for a hypothetical *round trip* trade in which the liquidity demander buys at the current offer price and simultaneously sells at the current bid price. Equivalently, it can be viewed as a liquidity supplier's profit from a hypothetical *round trip* (i.e., inventory offsetting) trade in which the liquidity supplier sells at the current offer price and simultaneously buys at the current bid price.

Let P_k be the price of the k^{th} trade. A trade is said to be *at the BBO* when $P_k = O_k$ or $P_k = B_k$, where O_k and B_k are the best bid and offer prevailing at the time of the k^{th} trade. If all trades took place at the BBO, then percent quoted spread might be a sufficient

[1]The log representation in this and other measures yields an *approximate* percent difference. Two other conventions that are widely used are: (1) the $Dollar\ Quoted\ Spread_t = O_t - B_t$ and (2) the $Simple\ Percent\ Quoted\ Spread_t = (O_t - B_t)/M_t$, where $M_t = (B_t + O_t)/2$ is the midpoint of the BBO at time t. For convenience, we will often use the dollar as a currency unit, but all concepts could just as easily be denominated in any other currency unit.

cost measure. A trade is defined as *outside the BBO* when $P_k > O_k$ or $P_k < B_k$, which might happen if a large market buy (sell) exhausts the quantity available at the best offer (bid) and the remaining market order walks up (down) the book to trade at the second best offer (bid), the third best offer (bid), etc. A trade is defined as *inside the BBO* when $O_k > P_k > B_k$, which might happen if there is a *hidden order*[2] at a better price than the displayed BBO. Alternatively, an *inside the BBO* trade might happen if a dark pool matches a buy order and a sell order at the midpoint of the displayed BBO.

Given that trades sometimes happen inside or outside the BBO, an especially useful concept is the percent effective spread for the k^{th} trade, which is defined as

$$Percent\ Effective\ Spread_k = 2 \cdot D_k(Ln(P_k) - Ln(M_k)), \tag{2.2}$$

where D_k is an indicator variable that equals $+1$ if the k^{th} trade is a liquidity demander's buy and -1 if the k^{th} trade is a liquidity demander's sell[3] and M_k is the midpoint of the consolidated BBO at the moment of the k^{th} trade.[4] This measure is often aggregated over a period of time (e.g., an hour, a day, a month, etc.) by computing the *dollar-volume-weighted average* of the percent effective spread, where each trade observation is weighted by the dollar-volume of that trade

[2] An *iceberg order* is a type of hidden order in which part of the order is displayed. For example, an iceberg order for 5,000 shares might specify that 500 shares be displayed. After the first 500 shares execute, the second 500 is displayed. After the second 500 shares execute, the third 500 shares are displayed. And so on. A *fully hidden order* is when the full quantity is hidden. A fully hidden order executes after any displayed quantity at the same price is exhausted.

[3] There are three popular *trade-typing* conventions for determining whether a given trade is a liquidity-demander's buy or liquidity-demander's sell. The first convention is Lee and Ready [1991], where a trade is a buy when $P_k > M_k$, a sell when $P_k < M_k$, and the tick test is used when $P_k = M_k$. The tick test specifies that a trade is a buy (sell) if the most recent prior trade at a different price was at a lower (higher) price than P_k. The second convention is Ellis et al. [2000], where a trade is a buy when $P_k = A_k$, a sell when $P_k = B_k$, and the tick test is used otherwise. The third convention is Chakrabarty et al. [2006], where a trade is a buy when $P_k \in [0.3B_k + 0.7O_k, O_k]$, a sell when $P_k \in [B_k, 0.7B_k + 0.3O_k]$, and the tick test is used otherwise.

[4] Two other conventions that are widely used are: (1) $Dollar\ Effective\ Spread_k = 2 \cdot D_k(P_k - M_k)$ and (2) $Simple\ Percent\ Effective\ Spread_k = 2 \cdot D_k(P_k - M_k)/M_k$.

observation. This measure can be applied to analyze a whole sample of trades or any subsamples, such as trades broken out by trade size, order type, time of day, etc.

The midpoint of the consolidated BBO can be viewed as a benchmark for the perfect, frictionless price. So the term $D_k(Ln(P_k) - Ln(M_k))$ is the deviation of the k^{th} trade from the perfect, frictionless price. It is the *one-way* cost of a single trade. Multiplying by 2 converts this one-way cost into the round trip equivalent cost. Thus, the round-trip-equivalent, percent effective spread is on the same scale and can be compared to the hypothetical round trip of the percent quoted spread.[5]

The percent quoted spread is based on displayed quotes, so it represents the hypothetical cost of trading. By contrast, the percent effective spread is based on the actual trade price, so it represents the *actual*, *round-trip-equivalent*, cost of trading to the liquidity demander. Or equivalently, it represents the *actual*, *round-trip-equivalent*, profit to the liquidity supplier.

Market microstructure theory establishes three reasons why the effective spread in a competitive market must be greater than zero: (1) adverse selection [Glosten and Milgrom, 1985, Kyle, 1985, Easley and O'Hara, 1987], (2) order processing costs [Roll, 1984], and (3) compensation for bearing inventory risk [Amihud and Mendelson, 1980, Ho and Stoll, 1981, 1983]. Huang and Stoll [1996] provide an empirical strategy for separating the adverse selection component from the other two components. They consider a world in which all liquidity suppliers are uninformed, some liquidity demanders are uninformed, and other liquidity demanders are informed (i.e., have private information about the future value of the security). In such a world, uninformed liquidity suppliers lose money to informed traders, who systematically trade in the right direction. That is, informed traders buy (sell) when their private information is good news (bad news). In this context, the Percent

[5] For any cost measure, we have the general relationship: *round trip cost* $\equiv 2 \cdot$ (*one way cost*). Given that the round trip convention is for the full spread (percent quoted spread, percent effective spread, etc.), then the one-way convention is often called the half-spread (percent quoted half-spread, percent effective half-spread, etc.).

Effective Spread becomes the *gross* percent profit margin of uninformed liquidity suppliers. To determine the *net* percent profit margin, Huang and Stoll define a two-way decomposition of Percent Effective Spread as

$$\begin{aligned} \textit{Percent Effective Spread}_k = {} & \textit{Percent Price Impact}_k \\ & + \textit{Percent Realized Spread}_k. \end{aligned} \tag{2.3}$$

Percent price impact is the adverse selection component of the k^{th} trade, which is given by

$$\textit{Percent Price Impact}_k = 2 \cdot D_k(Ln(M_{k+n}) - Ln(M_k)), \tag{2.4}$$

where M_{k+n} is the consolidated midpoint of the BBO an arbitrary amount of time (n seconds) after the k^{th} trade. Intuitively, it refers to the increase (decrease) in asset value following a liquidity demander's buy (sell). This component is the uninformed liquidity supplier's loss to informed liquidity demanders, or equivalently, the informed liquidity demanders' profit.

Percent realized spread combines the order processing cost component and the inventory risk component of the k^{th} trade as given by

$$\textit{Percent Realized Spread}_k = 2 \cdot D_k(Ln(P_k) - Ln(M_{k+n})). \tag{2.5}$$

Intuitively, percent realized spread is the *net* percent profit margin of uninformed liquidity providers.[6] Huang and Stoll tested both 5 minutes and 30 minutes after the trade and most subsequent researchers have analyzed 5 minutes after the trade. However, as trading has become radically faster in recent years, the use of intervals of 1 minute or less would be justified.[7] Both percent realized spread and percent price impact would typically be aggregated on a dollar-volume-weighted basis.

[6] Rearranging (2.3), we obtain: *Percent Realized Spread* = *Percent Effective Spread* − *Percent Price Impact* = *Gross Percent Profit Margin* − *Percent Loss to Informed Traders.*

[7] Bessembinder and Kaufman [1997] suggest an alternative convention for assessing the security's post-trade economic value. Let P_{k+n} be the first trade price that is more than an arbitrary amount of time (n seconds) after the k^{th} trade. They substitute P_{k+n} in place of M_{k+n} in equations (2.4) and (2.5). In principle, in a large sample the two conventions should yield the same result.

2.1.2 The quantity dimension

Next, we examine how the quantity dimension is measured. A liquidity supplier always specifies a quantity (e.g., number of shares, number of bonds, number of contracts, etc.) that they are willing to trade at a particular price. In other words, it is always a price-quantity pair. The *offer depth* is the specified quantity that a liquidity supplier is willing to sell at the offer price. The *bid depth* is the specified quantity that a liquidity supplier is willing to buy at the bid price. Our first quantity measure is the average BBO depth at time t, which is defined as

$$Average\ BBO\ Depth_t = \frac{OD_t + BD_t}{2}, \tag{2.6}$$

where OD_t is the offer depth associated with the best offer price at time t and BD_t is the bid depth associated with the best bid price at time t.[8] This measure can be applied to a single exchange or to a consolidated market containing multiple exchanges. It is typically aggregated on a time-weighted basis, where each depth observation is weighted by the amount of time that the depth observation is in-force. Intuitively, this characterizes the average quantity that a trader can trade at the best prices.

To measure the quantity available beyond just the BBO, another quantity measure is the average cumulative depth at time t

$$Average\ Cumulative\ Depth_t = \frac{COD_t + CBD_t}{2}, \tag{2.7}$$

where COD_t is the cumulative offer depth up to a cutoff offer price at time t and CBD_t is the cumulative bid depth down to a cutoff bid price at time t. For example, Boehmer et al. [2005] examine cumulative depths to the following four cutoffs: 0.166%, 0.833%, 3.333%, and 16.67% above or below the BBO midpoint. For a \$30.00 stock, these cutoffs correspond to 5 cents, 25 cents, \$1, and \$5 above or below the BBO midpoint. Again, this measure can be applied to a single

[8] Two other conventions that are widely used are: (1) Average BBO Dollar Depth = (Offer Dollar Depth + Bid Dollar Depth)/2, and (2) Relative Average BBO Depth = Average BBO Depth/Quantity Outstanding. For example in the equity market, Dollar Depth = (Share Depth) $*$ (Share Price) and Quantity Outstanding = Shares Outstanding.

exchange or to a consolidated market and is typically aggregated on a time-weighted basis. Intuitively, this represents the average quantity that a trader can trade at prices that are within a certain distance of the BBO midpoint.

Our next measure combines cost and quantity. It is the slope of the price function, often called λ in reference to the slope of the price function in the well-known adverse selection model of Kyle [1985]. λ is measured by the slope coefficient of the following regression

$$r_n = \lambda \cdot S_n + u_n, \tag{2.8}$$

where r_n is the security's log price change in the n^{th} five-minute period, S_n is the signed square-root of dollar volume in the n^{th} five-minute period, and u_n is the error term. S_n is defined by

$$S_n = \sum_k Sign(v_{kn})\sqrt{|v_{kn}|}, \tag{2.9}$$

where v_{kn} is the signed dollar volume of the k^{th} trade in the n^{th} five-minute period. It is typically aggregated on an equally-weighted basis over all five-minute periods. Following Hasbrouck [2009] and Goyenko et al. [2009], the regression above assumes a square-root functional form. Alternatively, one might estimate a linear functional form. The λ represents the marginal cost of trading an extra unit of quantity.

2.1.3 The time dimension

Next, we examine the time dimension, which provides a dynamic view of liquidity. The first set of measures examine the speed of obtaining a particular outcome. The possible outcomes of a submitted order are: no execution, partial execution, or complete execution. An order that has not been completely executed may expire, be cancelled, or continue in effect.[9] For the k^{th} order that obtains a particular outcome, the speed of that outcome is given by

$$\textit{Speed of partial execution}_k = pt_k - st_k, \tag{2.10}$$

[9] Any remaining unexecuted portion of a *day* limit order automatically expires at the end of a trading day. By contrast, a *good-until-cancelled* limit order never expires. Both types of orders may be cancelled at any time.

$$Speed\ of\ complete\ execution_k = ct_k - st_k, \tag{2.11}$$

$$Speed\ of\ cancellation_k = cant_k - st_k, \tag{2.12}$$

where pt_k is the earliest moment that the k^{th} order partially or completely executes, st_k is the submission time of the k^{th} order, ct_k is the final moment at which the k^{th} order completes executing, and $cant_k$ is the moment at which the k^{th} order is cancelled. All of these speed measures would typically be aggregated on a dollar-volume-weighted basis over all orders in the relevant set. For liquidity-demanding orders (i.e., market orders or marketable limit orders) on electronic exchanges, the *speed of partial execution*$_k$ is nearly always less than one second. By contrast, liquidity-supplying orders (non-marketable limit orders) must wait for a counterparty to trade with, which may take minutes, hours, or longer.

A limit buy (sell) price is said to be more *aggressive* if it has a higher (lower) price. For example, marketable limit orders are more aggressive than non-marketable limit orders. As a general characterization, the more aggressive a limit order price is, then the faster will be the speeds of partial and complete execution.

The next set of measures combine quantity and time. They examine which outcome happens to a set of submitted orders over a fixed time interval (e.g., an hour) or over a remaining time interval (e.g., until the end of the day). For a set of submitted orders, the outcome rates over a fixed or remaining time interval are given by

$$Partial\ fill\ rate = \frac{np}{ns}, \tag{2.13}$$

$$Complete\ fill\ rate = \frac{nc}{ns}, \tag{2.14}$$

$$Cancellation\ rate = \frac{ncan}{ns}, \tag{2.15}$$

where np is the number of orders that partially or completely execute, ns is the number of orders submitted, nc is the number of orders that completely execute, and $ncan$ is the number of orders that are cancelled. Alternatively, the inputs to these rates could be measured in quantity amounts (e.g., number of shares) or dollar amounts. Higher partial and complete fill rates mean more liquidity, but cancellation rates are trickier to interpret. Higher cancellation rates mean that a given set of orders was more frequently cancelled rather than filled (implying less liquidity), but the cancelled orders may have been part of a larger trading program which may have achieved either a higher or lower fill rate overall. See Section 2.2.7 for a discussion of a trading program by an institutional trader.

The final time measure examines how cost or quantity measures of liquidity evolve over time. Kempf et al. [2008] analyze the resiliency of liquidity over time. Let L_t be a cost or quantity measure of liquidity at time t. They postulate that a liquidity measure follows a mean-reverting process given by

$$\Delta L_t = \kappa(\theta - L_{t-1}) + \varepsilon_t, \tag{2.16}$$

where ΔL_t is the change in liquidity over a time period, κ is speed of adjustment to the long-run mean, θ is the long-run mean level of liquidity, and ε_t is the noise term. Then they run the following regression

$$\Delta L_t = \alpha + \kappa L_{t-1} + \sum_{\tau=1}^{p} \gamma_\tau \Delta L_{t-\tau} + \varepsilon_t, \tag{2.17}$$

where $\alpha = \kappa\theta$ is the intercept, γ_τ are lag coefficients, and p is the number of lags. The speed of adjustment κ is their measure of the resiliency of a market.[10] They estimate the resiliency of various liquidity measures (percent quoted spread, dollar quoted spread, bid depth, and offer depth) over one-minute intervals with 20 lags. Intuitively, resiliency measures how fast a market's liquidity that has been shocked returns to the long-run mean level of liquidity.

2.2 Specialized issues in liquidity measurement

So far we have examined standard measures of liquidity that apply generally to the trading of any asset class, anywhere around the world.

[10] Alternative measures of resiliency include: (1) time to recover from an aggressive order as estimated by Degryse et al. [2005], (2) the serial correlation of order type over time intervals as estimated by Ellul et al. [2007], (3) the negative of the serial covariance of prices as suggested by Vayanos and Wang [2011], and (4) the limit order refill rate following a trade as suggested by Obizhaeva and Wang [2012].

In this section, we turn to specialized issues in liquidity measurement that address specific issues or apply to specific markets.

2.2.1 Spread components

Huang and Stoll [1997] perform a two-way decomposition of the dollar effective spread by estimating the following regression

$$\Delta P_t = \frac{S}{2}(Q_t - Q_{t-1}) + \lambda \frac{S}{2} Q_{t-1} + e_t. \tag{2.18}$$

where ΔP_t is the change in price of the t^{th} trade, S is the dollar effective spread, Q_t is the liquidity demander's buy/sell indicator for trade t, λ is a constant, and e_t is the error term. The two components are: (1) λ, which is the portion of the spread due to the adverse selection and inventory risk components, and (2) $1 - \lambda$, which is the portion of the spread due to the order processing cost component. Huang and Stoll's model assumes a perfectly competitive market. However, if the real-world, empirical spread contains a fourth component based on market maker rents due to their market power and/or due to the tick size being greater than the competitive spread size, then those rents will be empirically indistinguishable from the order processing cost component (i.e., $1 - \lambda$ will reflect both order processing cost and market maker rents).

Huang and Stoll [1997] also perform a three-way decomposition of the dollar effective spread by simultaneously estimating the following two equations

$$\Delta P_t = \frac{S}{2} Q_t + (\alpha + \beta - 1)\frac{S}{2} Q_{t-1} - \alpha \frac{S}{2}(1 - 2\pi) Q_{t-2} + e_t, \tag{2.19}$$

$$Q_{t-1} = -\pi Q_{t-2} + u_{t-1}, \tag{2.20}$$

where π is the serial correlation in trade sign and u_{t-1} is another error term. The three components are: (1) α, which is the portion of the spread due to the adverse selection component, (2) β, which is the portion of the spread due to the inventory risk component, and (3) $1-\alpha-\beta$, which is the portion of the spread due to the order processing cost (plus any market maker rents) component.[11]

[11] Huang and Stoll show that they their generalized spread decomposition model incorporates many earlier, more limited models as special cases. Specifically, their

Implicit in the three-way decomposition, the adverse selection component is identified by the *permanent impact* it has on prices, whereas the inventory risk and order processing components only have *temporary impacts* on prices. Similarly, in the two-way decomposition discussed in Section 2.1.1, the adverse selection component (namely percent price impact) is identified by its permanent impact on price, whereas the combined inventory risk and order processing component (namely percent realized spread) is identified by its temporary impact on prices.

2.2.2 The futures market

A challenging issue in the futures market is that intraday data is relatively limited. For example, the Chicago Mercantile Exchange (CME) has volume-tick files which provide the price, volume, and time of all trades. However, the CME data does not provide bid-offer quotes, buy/sell indicators, or order data. So none of the standard liquidity measures discussed in Section 2.1 can be computed with CME data.

Hasbrouck [2004] develops a novel Bayesian solution. He estimates the price dynamics of a variant of the Roll [1984] model[12]

$$\Delta p_t = c\Delta q_t + u_t, \tag{2.21}$$

where Δp_t is the change in price for the t^{th} transaction, c is the effective half-spread, Δq_t is the change in a buy/sell indicator variable, and u_t is an error term, which is assumed to be i.i.d. and normally distributed with mean zero and variance σ_u^2. The model implies that the $var(\Delta p_t) = \sigma_u^2 + 2c^2$ and $cov(\Delta p_t, \Delta p_{t-1}) = -c^2$. He uses a Gibbs sampler to estimate the latent variables c, σ_u, and $\{q_1, q_2, \ldots, q_T\}$, where the latter are liquidity-demander buy/sell indicators for the T trade observations. This yields an estimate of the effective spread $2c$.

model incorporates the covariance spread models of Roll [1984], Choi et al. [1988], Stoll [1989], and George et al. [1991] and the trade indicator spread models of Glosten and Harris [1988] and Madhavan et al. [1996].

[12] Huang and Stoll [1997] show that Roll [1984] model is strictly based on the order processing component of the spread (i.e., the adverse selection and inventory risk components are zero). This implies that the Hasbrouck [2004] model is also based on the order processing component only.

He also estimates the slope and intercept of the price function by estimating

$$\Delta m_t = \sum_{j=0}^{J} q_{t-j}\lambda_j v_{t-j} + u_t, \tag{2.22}$$

where Δm_t is the change in "efficient price" of the t^{th} transaction,[13] $v_t = [1\sqrt{volume_t}]$, λ_j is a (1×2) coefficient vector, and J is the number of lags.

2.2.3 The U.S. corporate bond market

The U.S. corporate bond market is predominantly an over-the-counter dealer market. Traditionally, it was a very opaque market with very little information being reported. On July 1, 2002, the National Association of Securities Dealers (NASD) implemented a new rule that required the public reporting of all trades in most corporate bonds called the Trade Reporting and Compliance Engine (TRACE). In 2003, coverage was expanded to all corporate bonds and the delay in public reporting was reduced to a few minutes. Thus, the corporate bond market became very ex-post transparent.[14]

TRACE reports the date, time, volume and price of all corporate bond trades. Since November 2008, a buy/sell indicator has also been reported.[15] However, the TRACE data does not include bid-offer quotes. This makes it impossible, for example, to compute the percent effective spread, which requires the bid-offer midpoint.

Bessembinder et al. [2006] develop a novel way of estimating the effective half spread using a two-stage model

$$\Delta P = a + wX_t + \gamma SQ_t^* + \alpha S\Delta Q + \omega_t, \tag{2.23}$$

where ΔP is the change in TRACE bond price, X_t is a vector of public information variables (the change in on-the-run, maturity-matched

[13] The efficient price is implied bid-offer midpoint given by $p_t - c$ for buys ($q = +1$) and $p_t + c$ for sells ($q = -1$).

[14] Ex-post transparency is the rapid availability of information about trades. Ex-ante transparency is the rapid availability of information about bid-offer quotes and/or the limit order book.

[15] Dick-Nielsen [2009] reports procedures to clean the TRACE data.

Treasury yields, the firm's stock return, and the spread between long-term BAA-rated corporate bonds and U.S. Treasuries), S is the effective half-spread, $Q_t^* \equiv Q_t - E_{t-1}(Q_t)$ is the surprise in order flow, ΔQ is the change in the liquidity-demander's buy/sell indicator, and ω_t is the error term. The surprise in order flow Q_t^* is estimated as the residual from a first-stage regression

$$Q_t = a + bQ_{t-1} + \varepsilon_t. \tag{2.24}$$

Edwards et al. [2007] estimate a model that is similar in spirit, but which adds special attention to estimating percentage trade costs as a function of trade size. They find that the percentage trade costs are a nonlinear function of trade size with small trades being 25 times more expensive than large trades.

Mahanti et al. [2008] analyze proprietary data from a large custodial bank, State Street Corporation (SSC), including corporate bond transactions and fund holdings from January 1994 to June 2006. They develop a new liquidity measure based on the accessibility of a security, which they call Latent Liquidity. For bond i in month t, it is defined as

$$L_t^i = \sum_j \pi_{j,t}^i T_{j,t}. \tag{2.25}$$

where $\pi_{j,t}^i$ is the fund j percent holding of bond i out of the bond's total outstanding amount in the SSC database at the end of month t and $T_{j,t}$ is the portfolio turnover of fund j in month t. The latter is measured as the dollar trading volume of fund j from month $t-12$ to month t divided by the value of fund j at the end of month t. Since this liquidity measure is weighted average of *fund turnovers*, it can be computed for any corporate bond even if there are few or zero transactions by that bond in a given month.

Due the lack of ex-ante transparency, investors must search multiple dealers to determine available prices. Jankowitsch et al. [2011] show theoretically that in the presence of search costs for investors and inventory risk for dealers, trade prices will exhibit price dispersion. Based on their model, they develop an empirical estimator of price dispersion

$d_{i,t}$ of the i^{th} bond on day t

$$d_{i,t} = \sqrt{\frac{1}{\sum_{k=1}^{K_{i,t}} v_{i,k,t}} \cdot \sum_{k=1}^{K_{i,t}} (p_{i,k,t} - m_{i,t})^2 \cdot v_{i,k,t}}. \tag{2.26}$$

where $p_{i,k,t}$ and $v_{i,k,t}$ are the price and volume of the k^{th} trade in the i^{th} bond on day t and $m_{i,t}$ is the widely-recognized, Markit Group end-of-day composite bond price, which Markit aggregates from more than 30 bond dealers. Their measure is the square root of the volume-weighted sum-of-squared price deviation from an end-of-day consensus valuation. Intuitively, the price dispersion measure captures the search cost component of transaction costs.

In the 1990s, investors used to have to sequentially call each dealer to get their bid and offer quotes. Over the last decade, many electronic platforms have been introduced (Tradeweb, Bloomberg BondTrader, Reuters RTFI, MarketAxess, etc.) which allow an investor to request quotes from a large number of dealers and get firm quotes back in less than a minute. This would suggest that search costs have greatly decreased over time and therefore price dispersion would be predicted to decrease as well.

2.2.4 The U.S. equity market

In recent years, the U.S. equity market has become much faster (i.e., the average speed of partial and complete fills has transformed from human speed in tens of seconds to computer speed in milliseconds) and more competitive. On the competition dimension, from 2005 to 2009 the NYSE's market share in NYSE-listed stocks dropped from 80% to 25% and NASDAQ's market share in NASDAQ-listed stocks dropped from 53% to 30% [see Angel et al., 2011]. While it was once acceptable for researchers to rely on the BBO quotes of a dominant exchange, from at least 2009 forward, researchers must use composite BBO quotes. In the U.S. equity market, the composite BBO quotes are called the National Best Bid and Offer (NBBO) quotes.

The NYSE sells the *Monthly* Trade And Quote (MTAQ) database, which is widely-used in academic research and is updated monthly.

For a three- to four-times higher price, the NYSE also sells the *Daily* Trade And Quote (DTAQ) database, which is aimed at commercial users and is updated daily. The two datasets are identical, except for two key differences. First, trades and quotes are time-stamped to the second in MTAQ vs. to the millisecond in DTAQ. Second, a "NBBO" file containing most of the official NBBO quotes is available in DTAQ, but not in MTAQ.

Holden and Jacobsen [2014] show that standard liquidity measures based on the NBBO as computed from DTAQ are enormously different than the same measures based on the NBBO as computed from MTAQ. They identify three sources of MTAQ NBBO errors: (1) withdrawn quotes, where an exchange or market maker momentarily quotes nothing, (2) millisecond versus second timestamps, and (3) other causes, including cancelled quotes where a limit sell (buy) setting the current offer (bid) is cancelled and the exchange or market maker's quote is updated in the DTAQ NBBO file, but not in MTAQ. All three sources are found to be statistically and economically significant in causing liquidity measurement problems.

Holden and Jacobsen find that using the expensive DTAQ database is the first-best solution, because it yields a much lower frequency of negative quoted spreads and trades outside the NBBO than any MTAQ alternative. If a researcher is financially constrained, then their second-best solution is to use the cheaper MTAQ database and make three adjustments. First, correctly adjust for withdrawn quotes, which are directly observable in MTAQ. Second, they develop an "Interpolated Time" technique that makes an educated guess about the millisecond in which trades and quotes take place, so as to more accurately match trades and quotes. Third, delete both trades and NBBO quotes, whenever the NBBO is zero or negative because this is an economically nonsensical state.

Finally, Holden and Jacobsen consider whether these solutions change research inferences. First they reexamine the impact NYSE's 2006 to 2007 "Hybrid Market" reform that greatly increases exchange automation. They show that the conventional MTAQ treatment (NBBO across all markets, no adjustments for withdrawn quotes, etc.)

yields an incorrect inference of no change in percent effective spread, whereas their first-best and second-best solutions yield the correct inference of an increase in percent effective spreads. Next, they reexamine exchange performance when firms are sorted into quintiles based on percent effective spread. They find that the conventional MTAQ treatment yields biased conclusions about which exchanges have superior versus inferior performance compared to the first-best solution DTAQ. These biases are reduced by using the second-best solution with MTAQ, but not eliminated. Finally, they conduct a firm trading costs sort of the type that is common in the corporate finance and asset pricing literature. They find that using the conventional MTAQ treatment, the majority of dollar effective spread quintiles differ from our first-best solution, whereas using their second-best solution, the vast majority are the same as the first-best solution. Thus, their first-best and second-best solutions affect research inferences in a wide literature.

Henker and Wang [2006] re-examine the Lee and Ready [1991] recommendation that trades be matched to consolidated BBO quotes in-force 5 seconds prior to the trade. The Lee and Ready recommendation is based on an analysis of 1988 data; they document a significant delay in recording trades relative to quotes. Henker and Wang find that the delay in recording trades vanished by the start of MTAQ in 1993. Therefore, trades in MTAQ should be matched to the consolidated BBO quotes in-force immediately prior to the trade, as is done in the Interpolated Time technique discussed above.

2.2.5 Liquidity proxies calculated from daily data

Trade and quote data have grown exponentially over time. For example, Angel et al. [2011] document that the frequency of quote updates in S&P 500 stocks grew at a compound annual rate of 59% per year from 2003 to 2009. Chordia et al. [2011] document that the value-weighted frequency of NYSE trades grew at a compound annual rate of 91% per year from 2003 to 2008. Fong et al. [2014] document that global trades plus quotes have grown at a compound annual rate of 32.8% per year for 37 exchanges from 1996 to 2007. Over the same period, Hennessy and Patterson [2012] report the compound annual growth rate in computer

CPU performance as 31.0% per year. Thus, the exponential increase in trades and quotes has at least kept pace with the exponential increase in computing power, such that it will continue to be very difficult to compute standard measures of liquidity from intraday data for the foreseeable future.

A rapidly growing literature proposes liquidity proxies that can be calculated from daily data. The key advantage of using daily data is a large savings in computational time compared to using intraday data. For their global data, Fong et al. document a 42-fold computational savings when using 1996 data and growing to a 962-fold computational savings when using 2007 data. Projecting this pattern into the future, computational savings will reach the 10,000-fold level by 2016 as intraday data continues to grow at an exponential pace, whereas daily data grows at a linear pace.

This raises the question of how good these liquidity proxies calculated from daily data are at capturing standard liquidity benchmarks calculated from intraday data. Goyenko et al. [2009] investigate this question for U.S. equity markets by comparing a large number of monthly and annual liquidity proxies calculated from the Center for Research in Security Prices (CRSP) daily stock data to monthly and annual liquidity benchmarks computed from MTAQ data and Rule 605 data.[16] They conclude that the Effective Tick, Holden, and LOT Y-split proxies do a good job of capturing percent effective and realized spread benchmarks with both strong correlations and low root-mean-squared prediction errors. They find that the Amihud proxy and a new class of extended-Amihud proxies are successful in obtaining reasonably good correlations with the price impact benchmarks (e.g., lambda), but that all of the proxies fail to capture the correct scale of the price impact benchmarks.

Corwin and Schultz [2012] develop a High-Low spread proxy and show that it does a better job of capturing percent effective spread in U.S. data than any other proxy they test. Chung and Zhang [2014] suggest computing the Closing Percent Quoted Spread from closing bid

[16]S.E.C. Rule 605 requires that all U.S. exchanges and other market centers disclose monthly performance statistics by stock, order type, and order size.

and offer prices available in CRSP for certain time periods[17] and show that it does a better job of capturing percent effective spread in U.S. data than any other proxy they test. However, neither paper tests these two proxies against each other.

Fong et al. [2014] test a large number of monthly liquidity proxies (including both High-Low and Closing Percent Quoted Spread) calculated from Datastream daily stock data against monthly liquidity benchmarks computed from Thomson Reuters Tick History (TRTH) intraday stock data for 43 exchanges around the world. They also test *daily* liquidity proxies against *daily* liquidity benchmarks. They find that for both monthly and daily frequencies that Closing Percent Quoted Spread strongly dominates all other percent-cost proxies for global research. They find that it has much higher correlations with intraday percent-cost benchmarks (percent effective spread, percent quoted spread, percent realized spread, and percent price impact) than any other percent-cost proxy. At both daily and monthly frequencies, they find that it does the best job of capturing the level of percent effective spread and percent price impact. At both frequencies, they find that High-Low does the best job of capturing the level of percent realized spread and percent price impact. Interestingly, they find that *daily* liquidity proxies do incredibly well. For example, on a global basis daily Closing Percent Quoted Spread has an average cross-sectional correlation of 0.691 with daily percent effective spread (keep in mind that the latter is computed from intraday data) and a portfolio time-series of 0.809.

Turning to cost-per-volume proxies, Fong et al. [2014] find that the five best (and nearly equivalent) monthly cost-per-volume proxies are Closing Percent Quoted Spread Impact, LOT Mixed Impact, High-Low Impact, FHT Impact, and Amihud. They find that the best daily cost-per-volume proxy is the daily version of Amihud. At both frequencies, they find that the best cost-per-volume proxies are strongly correlated with lambda, but none of them captures the level of lambda.

[17] The CRSP daily stock database includes closing bid and offer prices for NYSE/AMEX stocks from 1925 to 1942 and from 1993 to present, for the NASDAQ Global Market and Global Select Market (formerly National Market) from 1982 to present, and for the NASDAQ Capital Market (formerly SmallCap) from 1992 to present.

Turning to commodity markets, Marshall et al. [2012] test a large number of monthly liquidity proxies calculated from Datastream daily commodity futures data against monthly liquidity benchmarks computed from Thomson Reuters Tick History (TRTH) intraday commodity futures data. They find that Amihud has the largest correlation with three liquidity benchmarks and that Amivest and Effective Tick are next best. However, none of the liquidity proxies captures the level of the liquidity benchmarks.

Schestag et al. [2013] test a large number of monthly and annual liquidity proxies calculated from Bloomberg daily corporate bond data against monthly and annual liquidity benchmarks computed from TRACE intraday corporate bond data. They find that the High-Low, Roll, and Gibbs proxies are highly correlated with percent effective spread benchmarks. They find that the Amihud and High-Low Impact proxies are highly correlated with price impact benchmarks.

2.2.6 A matched sample to test for a difference in liquidity

Davies and Kim [2009] analyze the best way to form a matched sample in order to test for a difference in liquidity. They perform a large-scale Monte Carlo simulation in which N_A stocks are assigned to exchange A and $1,000-N_A$ stocks are assigned to exchange 0. The bid-offer spread of the stocks assigned to exchange A is artificially increased by θs_y, where θ is a constant and s_y is the standard deviation of the sample spreads. Each of the stocks assigned to exchange A is matched with one or more of the stocks assigned to exchange 0 using a particular weighting scheme and various statistical tests are performance to detect the difference in spreads. Each test is replicated 20,000 times to determine its statistical properties.

Davies and Kim find the following: (1) a Wilcoxon Test has more power than a pairwise t-test or a regression with a dummy, (2) it is generally better to construct one-to-one matches to minimize bias, (3) the best matching criteria for NYSE/AMEX stocks are size (market capitalization) and price, (4) the best matching criteria for NASDAQ stocks are size, price, and volume, (5) matching should *not* be restricted to firms in the same industry, (6) poor matches should *not* be dropped from the sample, (7) when the control sample is relatively small, firms should be matched without replacement, (8) an event study has more power than a matched sample for a given sample size, but a large matched sample might have more power than a small event study, and (9) these results are robust by subsamples and time periods.

2.2.7 A trading program by an institutional trader

An institutional trader will often take a large trading request and divide into many small orders that execute over time. What appears in the TAQ dataset as many small trades, is often a single trading program. Ancerno Ltd. (formerly the Abel/Noser Corporation) sells a database that provides a detailed trading history of a large number of institutional investors. Importantly, it provides information on trading programs, such as what trading requests ("tickets") are sent to which brokers and what trades in which stocks result. It includes identifiers for both the institution and the broker, so the total cost of a complete trading program can be tracked over time.

The total cost of a trading program includes both explicit and implicit costs. *Explicit costs* are out-of-pocket expenses, such as broker commissions, taxes, fees, etc. Since the institution directly pays explicit expenses, they are easily identified and tracked by the institution's accounting system.

Implicit costs are price effects of trading, such as spreads, price impact, delay costs, and opportunity costs. By definition, implicit costs are the difference between the actual trade price and a benchmark price. When there is a failure to purchase some of the requested shares, the definition is extended to include the difference between a potential trade price and a benchmark. In practice, several alternative benchmarks are used and these alternatives have various pros and cons.

Perold [1988] defines the trading program cost, *Implementation Shortfall* (*IS*), as follows:

$$\mathrm{IS} = \underbrace{\sum_{j=1}^{J} x_j (p_j - m_d)}_{\text{Execution Cost}} + \underbrace{\left(X - \sum_{j=1}^{J} x_j \right) (p_N - m_d)}_{\text{Opportunity Cost}} \tag{2.27}$$

where x_j is the size of the j^{th} trade out of J total trades, p_j is the price of the j^{th} trade, m_d is the bid-offer midpoint at the *decision time* (defined as when the trading program was decided upon), X is the total size requested by the trading program, and p_N is the price of the last trade of the day. The first term is the one-way, execution cost of the J trades that actually executed under the trading program. The second term is the one-way, opportunity cost (i.e., foregone profits) of requested shares that failed to be purchased. In both terms, the benchmark is the bid-offer midpoint at the decision time. This is the price in a frictionless market of a hypothetical "paper portfolio" that is executed immediately for the total size. Intuitively, Implementation Shortfall is the value difference between a real portfolio and the corresponding paper portfolio. It represents the trading program cost of real-world frictions.

Alternative measures of trading program cost use the same functional form for execution cost

$$\text{One-Way, Execution Cost} = \begin{cases} \sum_j x_j(p_j - b_j) & \text{for buys} \\ \sum_j x_j(b_j - p_j) & \text{for sells} \end{cases} \tag{2.28}$$

where b_j is the benchmark price for the j^{th} trade. Four popular benchmarks are the

$$\text{Time Weighted Average Price (TWAP)} = \frac{\sum_n p_n}{N}, \tag{2.29}$$

$$\text{Volume Weighted Average Price (VWAP)} = \frac{\sum_n v_n p_n}{\sum_n v_n} \equiv \sum_n w_n p_n, \tag{2.30}$$

$$\text{Closing Price} = p_N, \tag{2.31}$$

$$\text{Bid-Offer Midpoint} = m_j, \tag{2.32}$$

where TWAP and VWAP are summed over all trades $n = 1, 2, \ldots, N$ in that security on a given day (or other time period), v_n is the size of the n^{th} trade, $w_n \equiv v_n / \sum_n v_n$ is the volume-weight of the n^{th} trade, and m_j is the bid-offer midpoint at the time of the j^{th} trade. TWAP is the simple average price over the time period. In effect, it treats all trades as equal. VWAP recognizes different trade sizes and weights each trade accordingly. Closing price is popular as a benchmark because it is synchronous with the end-of-day value assessed for mutual funds and the end-of-day prices reported by financial media. Using the bid-offer midpoint as the benchmark yields the one-way effective spread.

2.2.8 Implementation shortfall components

An institutional investor's trading program cost, as given by Implementation Shortfall, can be decomposed into the following four components

$$\text{Delay Cost} = X(m_0 - m_d), \tag{2.33}$$

$$\text{Change in Midpoint Cost} = \sum_j x_j(m_j - m_0), \tag{2.34}$$

$$\text{Effective Spread Cost} = \sum_j x_j(p_j - m_j), \tag{2.35}$$

$$\text{Opportunity Cost} = \left(X - \sum_j x_j\right)(p_N - M_0), \tag{2.36}$$

where m_0 is the bid-offer midpoint at the dispatch time. The *dispatch time* is when an institutional investor sends a trading request (a "ticket") to a broker (or alternatively submits it to a trading algorithm). The institutional investor bears responsibility for the delay cost, whereas the broker or trading algorithm bears responsibility for the remaining three components. The change in midpoint cost represents the cumulative price impact of multiple trades in the program and/or any trend over time in midpoints. The one-way, effective spread cost represents the deviation of each trade price from the contemporaneous bid-offer midpoint. The opportunity cost is the foregone profit that accumulated from the dispatch time to the end of the period due to requested shares that failed to be purchased.[18] The four components

[18]Note that the opportunity cost in Equation (2.36), which starts accumulating at the *dispatch time*, is slightly difference from the opportunity cost in the second term of Equation (2.27), which starts accumulating at the *decision time*.

help assign responsibility and determine the reasons for the full trading program cost given by implementation shortfall.

In summary, we find that specialized measures of liquidity have been developed to deal with data limitations in specific markets (e.g., futures, U.S. corporate bonds, U.S. equity), to provide proxies from daily data, and to assess institutional trading programs.

3

Patterns in Liquidity

Next, we review what is known about patterns in liquidity in the cross-section, over time, due to exchange design, due to regulatory impacts, and due to high-frequency traders.

3.1 Cross-sectional and time-series patterns

Stoll's [2000] A.F.A. Presidential Address examines the cross-sectional patterns of a wide variety of liquidity measures.[1] He finds that percent quoted spread and the percent effective spread are 99% correlated with each other. In the cross-section, both of them are negatively related to dollar volume and stock price, positively related to volatility and average absolute order imbalance between buys and sells, positively related to size (market capitalization) and number of trades on NYSE/AMEX, and insignificantly related to size and number of trades on NASDAQ.

Next, Stoll decomposes the effective spread into the real friction (the order processing cost and inventory risk components) and the informational friction (the adverse selection component). Interestingly, both the real friction and the informational friction have nearly the same

[1] See also Stoll [1978] and Benston and Hagerman [1974].

cross-sectional patterns as the effective spread. Specifically, both are negatively related to dollar volume and stock price, positively related to volatility and average absolute order imbalance between buys and sells, positively related to size (market capitalization) on NYSE/AMEX, and insignificantly related to size on NASDAQ. The main difference is that he finds inconsistent results for the cross-sectional relationship to number of trades.

Lesmond [2005] examines cross-sectional patterns in 22 emerging markets. After controlling for the standard firm-level, cross-sectional determinants discussed above, he finds that the LOT spread proxy is negatively related to judicial efficiency and political stability and the Amihud price impact proxy is negatively related to political stability. Similarly, Eleswarapu and Venkataraman [2006] examine the cross-sectional patterns of NYSE-listed American Depository Receipts (ADRs) from 44 countries. After controlling for the standard firm-level, cross-sectional determinants, they find that both effective spread and price impact are negatively related to home country ratings for judicial efficiency, accounting standards, and political stability and significantly higher for ADRs from French civil law countries than from common law countries. Both papers conclude that improvements in legal and political institutions will lower liquidity costs in financial markets.

Chordia et al. [2005b] examine the time-series properties of stock and bond liquidity. They find stock and bond quoted spreads, order imbalance, and depth exhibit seasonal patterns by day-of-the-week, month, holidays, crisis periods (i.e., the Russian crisis, the Asian crisis, and the Bond crisis), and before and the-day-of macroeconomic announcements. After removing the seasonal patterns, they estimate a vector autoregression of stock and bond liquidity. They find cross-market dynamics flowing from volatility to liquidity and common influences in both markets. They also find that shocks to net borrowed reserves (i.e., monetary loosening) are associated with increased liquidity.

Goyenko et al. [2014] examine the determinants of equity option effective spreads. They obtain data on all option trades and the corresponding National Best Bid and Offer (NBBO) option quotes covering all U.S. option exchanges at the time of each trade from a commercial vendor LiveVol. This data is derived from the trade and quote data reported by options exchanges to the Options Price Reporting Authority (OPRA). Their sample spans all options on S&P 500 stocks from 2004 to 2012. They find that the major spread determinants are: (1) the hedging costs of option market makers including both delta-hedging and gamma-rebalancing costs, (2) option demand pressures as measured by order imbalances, and (3) information asymmetry. They also find that option effective spreads significantly increase around earnings announcement dates, which provides support to the evidence of informed trading in options.

Deuskar and Johnson [2011] estimate the net market order flow in S&P 500 E-Mini futures over a three-year period. E-Mini futures is the lowest cost method to trade a broad index of investor wealth and Hasbrouck [2003] establishes that is the most important venue for price discovery compared with other index contracts and exchange traded funds. They find that this flow-driven component of returns accounts for between 40% and 70% of market volatility. In other words, they argue that liquidity demand has a first-order effect on aggregate risk.

Nagel [2012] estimates the returns of a day-to-day reversal strategy as a proxy for the returns of liquidity provision.[2] He finds that reversal returns are highly predictable from the VIX index. That is, reversal expected returns and Sharpe ratios spike during periods of financial turmoil. Thus, part of the reason why market liquidity declined during the financial crisis is that liquidity providers demanded higher expected returns during this period. This is consistent with Brunnermeier and Pedersen [2009] who show that funding liquidity dries up when volatility is high.

Baele et al. [2014] empirically specify "flight to safety" (FTS) episodes (a.k.a., "flight to quality" or "flight to liquidity" episodes) as happening during periods of high equity market volatility and entailing a large and positive bond return, a large and negative equity

[2]In principle, a reversal strategy picks up the temporary component of the spread (i.e., the part based on the inventory risk and order processing components), but not the permanent component of the spread (i.e., the adverse selection component).

return, and negative high-frequency correlations between bonds and stocks (which are positively correlated in normal times). They analyze Datastream data on 23 developed countries from January 1980 through January 2012. They find that FTS episodes comprise less than 3% of the sample, are mostly local (as opposed to global), and mostly last three days or less. FTS episodes coincide with increases in VIX and the TED spread, decreases in consumer sentiment indicators in the United States, Germany, and the OCED, appreciations of "safe-haven" currencies (the yen, Swiss franc, and the US dollar), and a major decrease in bond liquidity.

Ben-Rephael [2014] examines "flight to liquidity" (FTL) crisis periods defined by a significant positive jump in VIX. He finds that during such crisis periods there are large withdrawals from mutual funds. These withdrawals force fund managers to disproportionately reduce their holdings of illiquid stocks, which contributes to the relatively large decline of illiquid stock prices.

In summary, the literature has established local cross-sectional patterns (liquidity is positively related to dollar volume and price level and negatively related to volatility and size), global cross-sectional patterns (liquidity is positively related to judicial efficiency, accounting standards, and political stability) and time-series patterns (liquidity exhibits seasonality, declines during crisis periods, and varies around macroeconomic announcements).

3.2 Commonality in liquidity

While early studies of the determinants of liquidity focused principally on its cross-sectional variations (see Section 3.1), more recent work has focused on the time-series properties of liquidity. Thus, Chordia et al. [2000, 2011], Hasbrouck and Seppi [2001], and Huberman and Halka [2001] consider whether liquidity variations share a common component in the equity markets. Chordia et al. [2000] and Huberman and Halka [2001] find that liquidity has a significant common component at both the market as well as the industry levels. Chordia et al. [2005b] find that daily aggregate spreads and depths across equity and

U.S. Treasury bond markets share common variation. Subrahmanyam [2007] demonstrates liquidity commonality between traditional equity and real estate markets (more specifically, Real Estate Investment Trusts).

Gromb and Vayanos [2002] build a model in which liquidity depends on the capital of financial intermediaries; the intuitive notion is that when intermediary wealth levels are high (i.e., funding constraints do not bind) intermediaries are better able to absorb demand shocks from investors. Using this link between funding constraints and liquidity, Brunnermeier and Pedersen [2007] rationalize common variations in liquidity. The idea is that market makers' margin requirements depend on market liquidity (defined as price deviation from fundamental value) so that negative shocks to liquidity may constrain margins, causing liquidity to constrict further, and so on, leading to liquidity "spirals." Common variations in liquidity follow directly from the notion that the funding constraint applies to total dealer wealth and the dealer manages a portfolio of securities. Comerton-Forde et al. [2010] test the idea that liquidity depends on market makers' credit constraints by showing that when NYSE specialist revenues are low, liquidity is also low. Cespa and Foucault [2014] argue that liquidity-providers can learn about the value of one asset from another asset. This interlinks price informativeness to liquidity. For example, if liquidity drops in one asset, this affects the price informativeness in that asset, which reduces the liquidity of a related asset. This phenomenon also causes co-movement in liquidity. In related work Hertrich [2014] argues that liquidity variations may occur due to variations in fundamentals. For example, he shows using Granger causality tests that an adverse shock to credit risk (as measured by credit default swap spreads) forecasts a reduction in liquidity, possibly because risk perceptions increase, thereby increasing the risk of holding an inventory in the asset. If credit risk has a pervasive impact, this phenomenon can also cause common variations in liquidity. Finally, Chung and Chuwonganant [2014] show that index option implied volatility (VIX), a proxy for marketwide uncertainty, has a pervasive impact on liquidity, pointing to another source of common variations in liquidity.

Recent work also considers liquidity commonality in a global context. Brockman et al. [2009] examine liquidity commonality using intraday data for 47 exchanges in 38 countries. After controlling for standard cross-sectional determinants, they find significant commonality in both quoted spread and depths on the large majority of exchanges, in both developed and emerging countries, and by size quintiles. They also identify significant simultaneous contributions to commonality by four equally-weighted, liquidity indices at: (1) the local exchange level, (2) the industry, (3) the region, and (4) the world. Karolyi et al. [2012] show that there is liquidity commonality in several countries and argue that it varies across time because of demand side reasons (e.g., correlated trading activity) rather than supply side reasons (funding liquidity). Mancini et al. [2013] document commonality across liquidity measures for several major foreign currency pairs, and U.S. stock and (Treasury as well as corporate) bond markets, indicating that systematic variation in liquidity is pervasive across asset classes and markets. These authors also show that systematic variations in liquidity risk factor (formed by going long in the most illiquid and short in the most liquid currencies) help explain returns on the popular carry trade strategy (a zero-net-investment strategy where an investor sells a currency with a low interest rate and purchases a different currency yielding a higher interest rate, and thus attempts to arbitrage interest-rate differentials). This suggests that the liquidity risk factor drives at least part of the return on the strategy, possibly also indicating that liquidity risk is priced in foreign exchange markets.

In summary, commonality in liquidity is prevalent.

3.3 Exchange design impacts

Jain [2003] examines how liquidity differs by exchange design features at 51 stock exchanges around the world. He finds that percent quoted spread, percent effective spread, and percent realized spread are more than double in dealer markets compared to the corresponding spreads in pure limit order book markets or in hybrid markets (with liquidity provided by both limit orders and dealers). Comparing the latter two exchange designs, he finds that percent quoted spread, percent effective spread, and percent realized spread are little different in the exchanges of developed countries, but significantly lower in hybrid markets compared to pure limit order books in emerging countries. After controlling for standard firm-level, cross-sectional determinants, he finds that percent quoted spread, percent effective spread, and percent realized spread are lower when there is a limit order book (as opposed to dealers), when there is a market maker (i.e., hybrid as opposed to a pure limit order book), when the limit order book is transparent, when provision is made for automatic execution, when the exchange is incorporated (as opposed to mutually owned), when order flow is centralized to one exchange (as opposed to fragmented), and when the relative tick size is smaller.

Mayhew [2002] examines options trading under alternative exchange designs and competitive conditions. He finds that options traded in the presence of a Designated Primary Marketmaker (DPM) have lower quoted spreads than in the absence of a DPM (i.e., in a traditional open outcry crowd). He also finds that options listed on multiple exchanges have lower quoted spreads than options listed on a single exchange. Intuitively, additional competition lowers spreads.

Anand and Venkataraman [2014] compare a market structure that relies on liquidity suppliers with no obligations to maintain markets (i.e., purely voluntary market makers) versus a market structure that relies on "official" liquidity suppliers who are obligated to maintain markets at all times (i.e., DPMs). They use audit-trail data from the Toronto Stock Exchange over 245 days in 2006. Interestingly, they find that voluntary market makers supply *more* liquidity when stocks are volatile, which is contrary to the prior concerns expressed by regulators. Voluntary market makers supply less liquidity when investor interest in a stock is low or when the order flow is one-sided. In the latter cases, having DPMs improves the investor's certainty of order execution.

Hendershott and Moulton [2011] examine the NYSE's "Hybrid Market" reform, which greatly expanded the scope of its automatic execution system by including market orders, including larger orders (up to one million shares), allowing orders to walk-up-the-book, and eliminating the wait time between sequential orders. They find that this reform reduced execution time for market orders from 10 seconds to less than 1 second, cut floor broker participation in half, cut specialist participation in half, increased quoted and effective spreads, and made prices more efficient. In effect, this reform transitioned the NYSE from predominantly floor-based to predominantly electronic.

In summary, we find certain exchange designs enhance market liquidity: a limit order book for high volume markets, a hybrid exchange for low volume markets, and multiple competing exchanges. The evidence on automatic execution is a little more mixed with significant gains on the speed dimension, but increases in quoted and effective spreads.

3.4 Tick size reduction impacts

New regulatory policies often have both intentional and unintentional impacts on market liquidity. One of the most important regulatory policies is a reduction in tick size.[3]

Bacidore [1997] examines the 1996 tick size reduction on the Toronto Stock Exchange. Specifically, for stocks above $5.00, the tick sized dropped from 12.5 cents to 5 cents; for stocks from $3.00 to $4.99, the tick sized dropped from 5 cents to 1 cent; and for stocks below $3.00, there was no change in tick size, where all numbers are in Canadian Dollars.[4] For stocks above $5.00, he finds an immediate and permanent reduction in quoted spreads, effective spreads, and depths. For stocks from $3.00 to $4.99, he finds no change in effective spread and a slight reduction in depths, suggesting that the prior tick size had not been a binding restriction. For stocks below $3.00, he finds an immediate and permanent *increase* in quoted and effective spread and no change in depth, suggesting that specialist cross-subsidies from higher priced stocks had been reduced. Importantly, for stocks above $5.00 he finds

[3] The tick size is the minimum price increment. For example, when the tick size is $1/8, then allowed prices are $30, $30 1/8, $30 2/8, $30 3/8, etc. When the tick size is $0.01, then allowed prices are $30.00, $30.01, $30.02, $30.03, etc.

[4] See also Porter and Weaver [1997] and Ahn et al. [1998] and Harris [1997] for an early survey.

that effective spreads decline for all trade size categories, suggesting that even the largest trades benefited despite the reduction in depths. He finds no increase in volume at the lower spreads and thus concludes that liquidity suppliers have suffered a reduction in profits.

In 1997, U.S. equity markets reduced their tick size from \$1/8 to \$1/16.[5] Bollen and Whaley [1998] examine the impact on the New York Stock Exchange (NYSE) and find a reduction in quoted spread and depth in all price quintiles and a reduction in effective spread in all but the largest price quintile. Ronen and Weaver [1998] examine the impact on the American Stock Exchange (AMEX) and find a reduction in quoted spread and effective spread and volatility for most price quintiles, but no change in depths and specialist profits. Van Ness et al. [2000] examine the impact on the NYSE, AMEX, and NASDAQ and find a reduction in quoted spread, effective spread, and volatility in nearly all price quartiles on all three exchanges. They find a drop in depth on the NYSE and AMEX, but an increase in depth on NASDAQ.

Goldstein and Kavajecz [2000] examine the impact on the NYSE using limit order book data and find a reduction in cumulative limit order book depth many ticks into the book for the full sample and for four price/volume subsamples. They find a reduction in both limit order and specialist contribution to depth. They also analyze hypothetical orders executing against the limit order and displayed floor interest (but excluding non-displayed floor interest) and find a decrease in effective spread for small trades, no change in effective spread for large trades in frequently traded stocks, and an increase in effective spread for large trades in infrequently traded stocks.

Jones and Lipson [2001] examine the impact on the NYSE using Plexus data on the trading history of a large number of institutional investors. They begin by documenting a reduction in effective spread for all trade size categories. Turning to the Plexus data, they analyze the implementation shortfall cost and total cost (adding commissions) by the total size requested by the trading program. They find a large increase in costs when the total size is 100,000 shares are more, a

[5] The last day of \$1/8 tick-size trading is June 23, 1997 for NYSE, May 6, 1997 for AMEX, and June 1, 1997 for NASDAQ.

small increase when the total size is 10,000–99,999 shares, no change in costs when the total size is 1,000–9,999 shares, and a small reduction in costs when the total size is less than 1,000 shares. Further, these cost increases are concentrated in stocks that had the smallest spreads before the tick-size reduction, where the tick size was most likely to have been binding. The Plexus data distinguishes between trading programs that are worked, defined as executed by multiple brokers or over multiple days, vs. programs that are not worked. They find that costs increase by a large amount for not worked programs, but no change in costs for worked programs. Finally, they find that that aggressive liquidity demanders, particularly momentum traders and those submitting large orders, bear most of the increased costs.

In 2001, U.S. equity markets reduced their tick size from \$1/16 to 1 cent.[6] Bessembinder [2003b] examines the impact on both the NYSE and NASDAQ. He finds an immediate and permanent reduction in quoted and effective spreads for the full sample and for nearly all size groups on both exchanges. He finds a reduction in effective spread for large trades on both exchanges and for most small and medium trade size groups on both exchanges. He finds a reduction or no change in realized spread for all trade size groups on both exchanges. He finds a reduction in primary market depth, NBBO depth, and intraday volatility for the full sample and all size groups on both exchanges, which verifies most of the predictions made by Harris [1999].[7] Importantly, he analyzes the ratio of six-hour variance compared to six times hourly variance. He finds either no change in the variance ratio or a change in the direction of getting closer to the random-walk ideal of 1.0 for the full sample and all size groups on both exchanges. This indicates no tendency to reverse hourly quote changes on either exchange as would be expected if liquidity supply had been seriously damaged.

Bacidore et al. [2003] examine the impact on the NYSE using limit order book data and find no evidence that traders switch away from limit orders or toward floor orders or market orders. They find a drop

[6]The last day of \$1/16 tick-size trading is January 28, 2001 for NYSE and AMEX and March 31, 2001 for NASDAQ.

[7]See also Furfine [2003] and Chakravarty et al. [2004].

in the average size of non-marketable limit orders and an increase in the cancellation rate. They find that depth starts at a smaller quoted spread, but the cumulative displayed depth decreases about 10 cents into the limit order book. Despite the reduction in depth, they find that effective spreads decline for small orders and do not increase for large orders.

In summary, tick size reduction yields a large improvement in liquidity as measured by average trade-weighted effective spread. These benefits are concentrated in small trades, but large trades are typically not harmed even net of the reduction in depth. Institutional traders have adapted to the reduced tick size environment by slicing trading programs into smaller order sizes, submitting limit orders when feasible, and dynamically updating orders through time.

3.5 Transparency impacts

Recall that NASD implemented a new rule on July 1, 2002 requiring the public reporting of all trades in most corporate bonds called TRACE. This transformed the corporate bond market from being ex-post opaque to being ex-post transparent.

Bessembinder et al. [2006] examine the impact of TRACE using six months of corporate bond data before and after the implementation date. For the full sample, they find a 50% reduction in their two-stage estimate of effective spread (see estimation details in Section 2.2.3). They find similar large reductions in effective spreads in every volume category and in every credit rating category. Interesting, they find a spillover effect of a 20% reduction in effective spread for non-TRACE-eligible corporate bonds.

Edwards et al. [2007] examine the impact of TRACE after it was fully rolled out by the end of 2002 using a large sample of corporate bonds from January 2003 to January 2005. They find that the percentage trade costs are a nonlinear function (specifically, a reverse "S" function) of trade size with small trades costing 75 basis points and large trades costing 3 basis points. While controlling for trade size, they find that (1) higher transparency lowers transaction costs, (2) bonds with a higher credit rating have lower transaction costs than bonds with a lower credit rating, and (3) bonds in a large issue size class have lower transaction costs than bonds in a small issue size class.

Biais and Green [2005] examine the reasons why the preponderance of municipal bonds shifted in the late 1920s from trading on the transparent NYSE to the opaque over-the-counter (OTC) market and why the same thing happened with regard to corporate bonds in the 1940s. They find that municipal bonds were squeezed off the NYSE trading floor by the NYSE's decision to reallocate scare trading space to hotly trading stocks in early 1929. They hypothesize that institutional traders would be relatively more interested in trading large blocks OTC and find supporting evidence that there was a large rise in the institutional ownership of corporate bonds in the 1940s at the time of the OTC switch. Interestingly, they estimate the total round-trip, trading cost (spreads + commissions) of trading New York municipal bonds on the transparent NYSE from 1926 to 1930 at 0.90%. By contrast, they estimate the total round-trip, trading cost of trading New York municipal bonds in the opaque, over-the-counter market of 2002 (pre-TRACE) at 2.45%. In other words, they find that transaction costs for retail investors in New York municipal bonds in the 1920s were *less than half* of present-day, pre-TRACE costs!

Madhavan [2000] provides an excellent review of many subjects in market microstructure, including transparency. He summarizes that "Greater transparency is generally associated with more informative prices. Second, complete transparency is not always 'beneficial' to the operation of the market. Indeed, many studies demonstrate that too much transparency can actually reduce liquidity because traders are unwilling to reveal their intentions to trade. Third, there is also general agreement that some disclosure — as opposed to no disclosure whatsoever — can improve liquidity and reduce trading costs. Finally, changes in transparency are likely to benefit one group of traders at the expense of others."

Boehmer et al. [2005] examine the launch of NYSE "Openbook," which provides a real-time view of the NYSE limit order book to subscribers. This was a large increase in ex-ante transparency. They find that effective spreads decline for the full sample and for small trades, do not increase for large trades, and that depths increase slightly. They find a shift in trading strategies such that limit orders become smaller in size and are cancelled more often and more quickly. Despite this shift, limit orders *increase* in their total market share of shares traded, whereas floor traders and specialists decrease. They interpret this as an increase in institutional trader "self-management" through limit orders at the expense of relying on floor traders to work a large trading request. Specialists suffer a decline in both trade participation rate and commitment to the quoted depth. Information efficiency improves slightly.

In summary, adding ex-post transparency to an otherwise opaque market dramatically improves liquidity as measured by effective spreads. Adding ex-ante limit order book transparency to a relatively transparent market causes a more modest improvement in liquidity.

3.6 Other regulatory impacts

Christie and Schultz [1994] and Christie et al. [1994] find evidence that NASDAQ dealers engage in implicit collusion to avoid quoting and trading on odd-eighth prices, which has the effect of maintaining artificially-wide quoted and effective spreads. In response, the SEC adopted the Order Handling Rules, which required that limit orders held by a dealer must be reflected in the dealer's quote and that ECN[8] quotes must be reflected in the quotes of the primary market. In other words, the SEC lost confidence that having multiple dealers alone would generate enough competition and so it decided to bolster the strength of two competitive alternatives: limit orders and ECNs. Barclay et al. [1999] examine the impact of these rules, which were implemented in 1997. They find that the order handling rules reduced quoted and effective spreads by approximately 30%, reduced effective spreads by all

[8] Functionally, an *Electronic Communication Network (ECN)* is an electronic exchange. Legally, an ECN is an *Alternative Trading System (ATS)*, which is a special category created by the SEC under Regulation ATS to facilitate the competitive entry of new exchange-like entities. ATSs are not technically exchanges and thus avoid most of the regulatory burden of exchanges.

trade size categories, caused no change in depths, and greatly increased the frequency of odd-eighth prices.

Battalio et al. [2004] examine the impact of an impending SEC plan to electronically link all of the U.S. options exchanges on market quality. This follows the events of August 1999, when strong political, legal, and competitive pressures led to an explosive surge in options cross-listing by the four, then-existing, options exchanges. Battalio, Hatch, and Jennings compare options trading a short time later in June 2000[9] versus January 2002 when an "inevitable" SEC plan to electronically link all of the options exchange was pending a roll out in the near future. They find that an 85% decline in the amount of time that option quotes are crossed and a 47% decline in the amount of time that option quotes are locked or crossed.[10] They estimate that the apparent arbitrage profits available when the quotes are crossed decline by 90%. They find that trade-through rates[11] decline by two-thirds and quoted and effective spreads decline by more than half.

Chung and Chuwonganant [2012] examine the impact of the SEC Regulation National Market System (NMS), which was implemented in 2007. Specifically, they examine the impact of: (1) the order protection rule (OPR), which requires that fast (electronic) exchanges avoid trading-through better prices available on other fast (electronic) exchanges,[12] but allows trading-through better prices on slow

[9]Immediately after a new electronic options exchange, the International Securities Exchange, has entered.

[10]The bid-offer quotes are *crossed* when the offer of one exchange is less than the bid of another exchange. The bid-offer quotes are *locked* when the offer of one exchange is equal to the bid of another exchange. Equivalently, the former is a negative NBBO and the latter a zero NBBO. Crossed quotes are particularly bad condition, because an arbitrageur could profit by buying at the lower offer price and selling at the higher bid price. By contrast, there is no arbitrage opportunity when the quotes are merely locked.

[11]A *trade-through* is a buy (sell) trade that happens at a higher (lower) price than is contemporaneously available on another exchange. A trade-through hurts the liquidity demander, because they are trading a worse price than they could have gotten on the other exchange.

[12]The new OPR replaced the prior OPR which required that orders in NYSE-listed stocks be exposed to the NYSE floor for 30 seconds to see if a better price could be found before they can be executed in order markets. The prior rule was criticized as conferring monopoly status on the NYSE trading floor.

(floor-based) exchanges, and (2) the access rule, which caps the access fee charged to liquidity-demanding orders at 30 cents per 100 shares. For the full sample, they find that Regulation NMS increased quoted and effective spreads, reduced depth, increased pricing error, and raised volatility.

However, when they examine the trading of NYSE/AMEX listed stocks on different venues they find sharply different results for the floor-based NYSE versus the electronic NASDAQ market. The floor-based NYSE suffered due to increased effective spreads, increased price impact, and slower execution speed, but the electronic NASDAQ benefited by decreased effective spreads, decreased price impact, and faster execution speed. Before Regulation NMS, NASDAQ had been much worse than the NYSE on effective spread and price impact, but pulled ahead of the NYSE on both measures afterwards and increased its speed advantage. Not surprisingly, they document a large shift in market share of NYSE/AMEX listed stocks away from the NYSE and toward NASDAQ and especially toward other exchanges. In other words, Regulation NMS leveled the playing field on the cost dimension of liquidity, which left speed as the main differentiator between the floor-based NYSE and electronic exchanges, leading to a major shift toward electronic exchanges. As mentioned previously, the long-run impact was a decline in NYSE's market share of NYSE-listed stocks down to 25% by 2009 and a significant decline in effective spreads [Angel et al., 2011]. Thus, the long-run impact of Regulation NMS appears to be pro-competitive and cost-reducing.

Battalio et al. [2014] examine the relationship between make/take fees and limit order execution quality. Most U.S. equity exchanges charge traders a small fee for liquidity-demanding orders (market orders or marketable limit orders) and provide a small rebate for liquidity-supplying orders (non-marketable limit orders) when they execute.[13] They analyze the routing decisions of eleven national brokerages and find that four of them, Ameritrade, E*Trade, Fidelity, and Scott Trade,

[13]Interestingly, a few exchanges have an *inverted* structure that does the opposite: pays a small rebate for liquidity-demanding orders and charges a small fee for liquidity-supplying orders that execute.

consistently route liquidity-supplying orders to whomever provides the highest rebate at any moment in time. Next, they analyze a unique dataset from a major investment bank in which identically priced limit orders to trade shares of the same stock are routed simultaneously to multiple venues. When these "identical" orders are routed to two venues, they find that the venue with the higher rebate has lower fill rates and provides lower realized spreads. In other words, brokers are serving their own interests by maximizing the rebates that they get to keep and this is done at the expense of customers who suffer lower execution quality (lower fill rates and lower realized spreads). Finally, they examine a broad sample of TAQ data and confirm the negative relationship between take fees for at-the-quote trades and realized spread.

In summary, regulatory reforms that increase the number of competitive alternatives, move toward linking them up, and level the playing field between exchanges have improved liquidity on both the cost and speed dimensions.

3.7 High-frequency trader impacts

Electronic markets with ever faster processing and ever lower latency communication have facilitated the rise of high-frequency traders (HFT). HFT use sophisticated trading algorithms to trade a large number of times per day and may hold a security for just milliseconds. A diversity of HFT employ a diversity trading strategies. The overall impact of HFT depends on what trading strategies they use, how often, and under what circumstances.

The most basic starting point in understanding HFT activity is the split between passive vs. aggressive trading strategies. On the one hand, HFT engage in market making through the use of passive, liquidity-supplying orders in order to earn the spread. On the other hand, they engage in a variety of other trading strategies through the use aggressive, liquidity-demanding orders.[14] This is a speed race against other

[14] The HFT literature tends to refer to trading strategies based on liquidity-demanding orders as *aggressive* trading strategies. Equivalently, such strategies are called *active*, liquidity-demanding trading strategies, by contrast to passive, liquidity-supplying strategies.

aggressive HFTs to exploit opportunities as they arise and also a speed race against passive HFTs, who rush to update liquidity-supplying orders to avoid being exploited. The aggressive category is generally thought to include four types of strategies. First, HFT engage in statistical arbitrage across multiple securities. Second, they rush to trade on both systematic and idiosyncratic news announcements. Third, they engage in order anticipation strategies that forecast the arrival of large buy or sell orders and then trade in the same direction in the anticipation of price movement in that direction. Fourth, they engage in latency arbitrage to profit by picking off slow traders at stale prices.

It is a non-trivial problem even to define who is a HFT. The U.S. Securities and Exchange Commission (SEC) [2014] defines a HFT as a proprietary trading firm that trades over an intraday time horizon. They take an expansive view of who is a proprietary trading firm, including the proprietary trading desk of a multi-service broker-dealer, hedge funds, and allow firms that are not broker-dealers and not members of FINRA. The intraday time horizon excludes algorithmic trades done by institutional traders to establish or liquidate positions with longer than intraday horizons.

SEC [2014] provides an excellent survey of empirical research on HFT which uses one particular type of identification strategy, namely, uses various non-public datasets that report market activity by trading firm. Various firms may then be classified as HFT manually based on human judgment or automatically based on various quantitative criteria.

One example of this approach is the NASDAQ HFT dataset. NASDAQ identifies 26 firms as HFT in a sample of 120 randomly-selected corporate stocks from 2008 to 2009. The stock sample is stratified into 40 large-caps, 40 mid-caps, and 40 small-caps and each capitalization group is evenly divided between NYSE-listed and NASDAQ-listed stocks. NASDAQ manually classifies trading firms as HFT based on the following considerations: (1) how often the firm's net trading crosses zero, (2) its order duration, (3) its order to trade ratio, and (4) NASDAQ's knowledge of customers. For each trade, the NASDAQ dataset identifies whether the liquidity demander is a HFT or not and whether the liquidity supplier is a HFT or not. NASDAQ makes this dataset available to academics. Many studies have used the NASDAQ HFT dataset, including Brogaard [2011], Brogaard et al. [2014b], Carrion [2013], and Hasbrouck and Saar [2013].

Brogaard et al. [2014b] begin with a simple description of HFT activity based on the NASDAQ HFT dataset. They find that in large-caps, HFTs represent 42% of the passive liquidity-supplying side of all trades and 42% of the aggressive liquidity-demanding side of all trades. In mid-caps this drops to 28% of passive trading and 19% of aggressive trading and drops further in small-caps to 25% of passive trading and 11% of aggressive trading. Thus, HFTs have a larger role in large-caps and play a significant role in both passive liquidity-supplying and aggressive liquidity-demanding. They find that HFTs make profits (net of take fees and plus make rebates) in both passive and aggressive trading in all three capitalization sizes. Their main result is that HFT aggressive trading facilitates price efficiency by trading in the direction of permanent price changes and in the opposite direction of transitory pricing effects. By contrast, HFT passive trading suffers some degree of adverse selection.

Carrion [2013] also examines the NASDAQ HFT dataset. For the full sample, he finds that HFTs represent 41% of passive trading and 42% of aggressive trading. He finds that HFTs tend to provide liquidity when spreads are wider and take liquidity when spreads are narrower. He also finds greater price efficiency with respect to order flows and market-wide returns on days when HFT participation is high.

Chaboud et al. [2014] examine computer-based vs. human traders in the FX market using Electronic Broking Services (EBS) data in three currency-pairs from 2003 to 2007. They find that the introduction and growth of algorithmic trading (and implicitly of HFT) caused two improvements in price efficiency: (1) a large reduction in triangular arbitrage opportunities and (2) a reduction in the autocorrelation of high-frequency returns. They show that the former is due to the aggressive trading of algorithmic traders, while the latter is due to their passive trading.

Other studies automatically classify trading firms as HFT based on various quantitative criteria. A tricky aspect of this approach is that different classification criteria may lead to different results. For example, Malinova et al. [2013] use high message-to-trade ratios and high total messages as their classification criteria, where messages are defined as order submissions plus order cancellations plus trades. In their Canadian equity sample, they find that only 26% of HFT trades were on the aggressive side. Considering that passive market making requires frequent quote updates, their focus on message rate criteria resulted in relatively more passive activity being classified as HFT. By contrast, Kang and Shin [2012] use a high number of daily messages above 1,000, a small end-of-day inventory, and a low median cancellation time for limit orders cancelled in less than 2 seconds as their classification criteria. In their Korean futures sample, they find that 74% of HFT trades were on the aggressive side. Their use of limit cancellation speed criteria probably contributed to relatively more aggressive activity being classified as HFT.

Kirilenko et al. [2011] examine the May 6, 2010 "flash crash" in which a broad cross-section of U.S. financial markets declined sharply for 30 minutes before nearly fully recovering. They use audit trail data for all trades in the June 2010 E-Mini S&P 500 futures contract from May 3–6, 2010. This data is supplied by Chicago Mercantile Exchange to its regulator, the Commodity Futures Trading Commission. They classified trading accounts into five categories based on various quantitative criteria: (1) high frequency traders (16 accounts), (2) intermediaries (179 accounts), (3) fundamental traders (1,263 accounts), (4) small traders (6,880 accounts), and opportunistic traders (5,808 accounts). SEC (2014) points out that this particular classification scheme likely misses one-third of HFT trading activity during the period. The opening event of the flash crash was a decision by a large mutual fund complex to sell 75,000 E-Mini contracts using an algorithm that completed execution in just 20 minutes. They find that: (1) HFT initially passively provided liquidity to the large selling algorithm, (2) then they aggressively sold their inventory as the market declined, and (3) then they aggressively bought as the market recovered. HFT holdings stayed relatively close to zero throughout and they were consistently profitable.

Jones [2013] provides an excellent survey of recent theoretical and empirical research on HFT. He notes that while studies that are based on various HFT datasets "yield considerable insight into overall HFT trading behavior," they are "less well-suited to identify the causal effects of HFT on market quality" due to endogeneity and reverse causality concerns. He says that the best way to measure the incremental effect of HFT is "to isolate market structure changes that facilitate HFT." His summary of the literature which takes this approach is that "virtually every time a market structure change results in more HFT, liquidity and market quality have improved."

An example of using a market structure change as an identification strategy is Boehmer et al. [2014], who examine the first date that each of 43 exchanges around the world permits co-location. Co-location is when fast traders are allowed to locate their server very close to an exchange's data center server so as to minimize the two-way data transmission time. Co-location is used as an exogenous shock that facilitates algorithmic trading (and implicitly HFT), but doesn't otherwise directly affect market quality. They use Thomson Reuters Tick History (TRTH) intraday stock data for 43 exchanges around the world. They find that co-location improves liquidity (lowers percent quoted spread, lowers percent effective spread, and lowers percent realized spread), improves price efficiency (lowers 10 minute autocorrelation), and raises volatility as measured seven different ways.

Jovanovic and Menkveld [2011] and Menkveld [2012] analyze the July 2007 entry of a major HFT with 75%–80% of its trades on the passive side. It quickly became a major player that was involved in roughly every third trade on Chi-X and every twelfth trade on Euronext. They do a difference-in-difference analysis and find a 35% reduction in quoted spreads.

Riordan and Storkenmaier [2012] study a Deutsche Boerse upgrade on April 23, 2007 which reduced Xetra system latency from 50 milliseconds to 10 milliseconds. They find that both quoted and effective spreads dropped 8.9% and 8.8%, respectively. They also find that the contribution of quotes to price discovery doubles to 90%.

Gai et al. [2013] examine the impact of April and May 2010 NASDAQ upgrades that reduced the minimum time between order messages from 950 to 200 nanoseconds. These upgrades had no impact on quoted spread, effective spread, volume, short-term variance ratio, or depth at the best bid or offer. However, it increased the cancellation/execution ratio by 26%, slightly increased short-term volatility, and reduced cumulative depth on the limit order book.

Brogaard et al. [2014a] examine multiple technology upgrades that reduce latency on the London Stock Exchange from 2007 to 2011. They find that reduced latency leads to increased HFT trading, but causes no change in institutional trader execution costs.

Another identification strategy is examine various proxies for HFT, such as high message rates, bursts of order cancellation and modifications, high order-to-trade ratios, etc. SEC (2014) suggests caution in interpreting such studies, as proxies may be associated with the broader phenomenon of algorithmic trading by all types of traders and/or may tend to select passive trading at the expense of aggressive trader (or vice versa).

One example of using a proxy is Hasbrouck and Saar [2013]. They examine HFT activity with the NASDAQ TotalView-ITCH data, which includes millisecond-timestamped order arrivals, cancellations, and executions. They impute HFT algorithmic "strategic runs" by linking a limit order submission message to its subsequent cancellation or execution message and further to any related limit order submission or execution messages that occur within 100 milliseconds on the same side for the same size. They construct a measure of HFT activity ("*RunsInProcess*") as the time-weighted average of the number of strategic runs of 10 messages or more that a stock experiences during a short time interval. They show that this measure has a 0.812 correlation with HFT executions from the NASDAQ HFT dataset. Their key finding is that higher HFT activity increases market liquidity (i.e., decreases the quoted spread, decreases the effective spread, and increases the displayed depth) and decreases short-term volatility.

Another example of using a proxy is Hendershott et al. [2011]. They examine the NYSE's 2003 switch from manual to electronic quote updating. This facilitated algorithmic trading by both HFT and non-HFT. Their proxy for algorithmic trading is based on messages relative to dollar volume, where messages is order submissions plus order cancellations plus trades. Their difference-in-difference findings are that an increase in algorithmic trading causes a decrease in quoted spread, effective spread, depth, and adverse selection for large-cap stocks, has little impact on small and medium cap stocks, and increases the amount of price discovery that takes place via quotes compared to trades.

Goldstein et al. [2014] provide an excellent survey of recent research on algorithmic trading and HFT. They summarize that "while early evidence suggests that under 'normal conditions' high-frequency traders appear to provide liquidity and enhanced market efficiency by acting as market-makers or statistical arbitrageurs across markets, more recent evidence and theoretical work has called into question the benefits of high-speed trading." They point to concerns about "errant or poorly designed HFT programs without necessary risk controls can lead to occasional shocks or disruptive events" and "certain HFT strategies raises concerns about their fairness, given the availability of certain tools or exchange rights to high-frequency traders that are not widely available to other types of investors."

Lewis [2014] tells the story of trader Brad Katsuyama at the Royal Bank of Canada who experimented with different trading strategies. When he sent a single large market order to a single exchange, it executed in full at the ex-ante displayed price. When he simultaneously sent multiple large market orders to multiple exchanges, the first exchange to receive an order along the communications route executed it in full at the displayed price, but other exchanges who received orders just milliseconds later frequently executed them at worse-than-displayed prices. However, when he fine-tuned the millisecond submission time of multiple larger market orders such that all orders would arrive at the destination exchanges at the same millisecond, then all orders executed in full at the displayed price.

This was evidence of HFT order anticipation strategies. Specifically, the HFT could see the first buy trade on the first exchange, forecast that additional large market buys were currently in route to other exchanges that currently match the best offer price, and then sent its own orders to those exchanges using a faster communication line to buy ahead of those buys orders hoping for a price bump when those buys subsequently arrive. Alternatively, if the HFT already had a resting limit order on those exchanges at the best offer price, then it could cancel that limit order and submit a new limit order at a slightly higher price so that the arriving buy orders would walk up the book to execute against the new limit at a slightly higher price. In this case, the HFT would be passively supplying liquidity, but at a wider spread than was displayed ex ante.

Hirschey [2013] examines the issue of order anticipation strategies by HFT. He uses the NASDAQ HFT dataset. He sorts stocks by HFT net aggressive buying. He finds that each additional ten shares purchased by the highest decile of HFT net aggressive buying predicts one additional share aggressively purchased by non-HFT trades in the following by 30 seconds. He estimates that this strategy would allow a HFT to earn a gross return (not considering transaction costs) of 1 basis point.

Ding et al. [2014] examine the issue of latency arbitrage by HFT. They compare the slower official National Best Bid and Offer (NBBO), as computed by the data feed consolidators who are called Securities Industry Processors (SIPs), to the faster direct feeds from each exchange which must be consolidated into an updated NBBO by the feed recipient. They examine 24 securities for 16 days in May 2012. Using Apple (AAPL) trading to illustrate, they find 54,734 price differences on May 9, 2012, which corresponds to 2.34 differences per second on average. They estimate that HFT could make an aggregate profit of $32,510 from latency arbitrage in Apple on one day.

Later in the Lewis [2014] story, Mr. Katsuyama and others leave the Royal Bank of Canada to found a new securities exchange called Investors Exchange (IEX). An important feature of IEX is that it adds a delay of 350 milliseconds to all arriving orders by routing them through a 38-mile length of fiber optic cable. This allows IEX the time to receive direct feeds from all U.S. exchanges, compute the updated NBBO, and then prevent midpoint-pegged limit orders and hidden limit orders from executing based on a stale SIP NBBO whenever the direct feed NBBO has changed. This allows certain types of liquidity-suppliers to avoid losing money to HFT due to latency arbitrage.

Holden and Jacobsen [2014] and Angel [2014] predict that the fundamental legal and economic concept of the NBBO will ultimately break down. They note that current trading algorithms are able to observe a market event (i.e., a limit order submission), process the information, and respond in about two to three milliseconds. Moore [1965] ("Moore's Law") and other similar formulations provide evidence of an exponential increase in computer power (i.e., CPU speed per dollar, memory capacity per dollar, etc.) and network power (i.e., internet backbone bandwidth, network latency, etc.). This has fueled a competitive "arms race" by HFT to increase processing speed and reduce network latency in order to leapfrog the competition. Thus, response times will likely accelerate into microseconds (10^{-6} seconds) in the late 2010s and into nanoseconds (10^{-9} seconds) in the 2020s. If bid and offer update messages could travel arbitrarily fast, then it would be possible to maintain a common NBBO for all economic agents in all locations. However, bid and offer update messages cannot travel faster than the speed of light (186,282 miles per second) and so HFT in different locations may simultaneously observe different "best" prices. The other words, the "National Market System" vision of providing all traders with integrated, time-synchronized information across all venues will break down. Information integration on a global scale could be reestablished with one-second batch auctions as discussed below.

Budish et al. [2013] argue that the continuous limit order book market design runs into difficulties at high-speed. They find that tightly-related securities, such as the E-Mini S&P 500 futures contract (ES) and the SPDR S&P 500 ETF (SPY), are highly correlated on the human time-scales of days, hours, or minutes, but have very low correlations on the computer time-scale of milliseconds. This low correlation creates purely technical arbitrage opportunities that the fastest HFT could exploit to earn approximately $75 million per year in ES-SPY arbitrage and billions more in the full universe of similar arbitrage opportunities. This creates a costly "arms race" by both HFT arbitrageurs and HFT market makers. The cost of this arms race plus the

HFT's arbitrage profit is initially borne by the HFT market makers, who recover this cost by widening their bid-offer spread and passing it on to non-HFT. The HFT market makers also respond by reducing their depths, which passes a large share of the cost to non-HFT institutional traders that trade in large amounts. They propose an alternative market design: frequent batch auctions at, say, one-second intervals. This would eliminate the purely technical, millisecond-scale arbitrage opportunities and avoid an expensive, socially-wasteful arms race.

Budish et al. [2014] analyze the implementation details for frequent batch auctions. They suggest that each round of trading include three intervals: (1) an order submission interval, (2) an auction computation interval, and (3) an auction outcome reporting interval. Orders that arrive during the order submission interval would not be displayed to avoid gaming. During the auction computation interval, the exchange would compute supply and demand schedules from the submitted orders and determine the market clearing price and quantity. Orders that carry over from prior periods would have time priority, but there would be no difference in time priority among orders that arrive in the same order submission interval.

Next, they address how to modify this setup to account for market fragmentation and Regulation NMS. If a given exchange's internal batch auction price is at the NBBO or inside the NBBO, then the internal batch auction's execution results would be implemented using strictly orders from that exchange. However, if the internal batch auction price is outside the NBBO, then buy or sell orders from other exchanges would be included in the auction, such that the combined batch auction price would be at the NBBO. The combined batch auction execution results would be implemented using both internal orders and orders from other exchanges that are accessed using intermarket sweep orders.

Finally, they analyze the shortest time interval that would be consistent with doing all of the steps above. Most steps can be done at fast computer speeds. By far the longest interval is the amount time it takes for information to travel completely around the world at the speed of light, which is 0.135 seconds.

In summary, high-frequency traders trade in both a passive, liquidity-supplying manner and an aggressive, liquidity-demanding manner. Their overall impact improves both liquidity and price efficiency, but concerns remain regarding occasional trading glitches, order anticipation strategies, and latency arbitrage at the expensive of slow traders.

4

Liquidity and Corporate Finance

The theoretical relationship between liquidity and corporate finance is a new and rapidly developing area of research. Current theories have focused on: (1) agency-based explanations, (2) feedback explanations, (3) discount rate explanations, and (4) market friction explanations.

In the agency models, equity liquidity can alter the large shareholder's incentive to acquire firm-specific information and to monitor the firm. Moreover, a more liquid equity market can impact the effectiveness of stock-based compensation for executives. Under the feedback theories, liquidity yields the entrance of informed market participants who produce information incremental to the manager's information set; managers can use this informative stock price as a guide in firm decision making. To the extent that liquidity results in a lower equity premium and thus a lower cost of capital, liquidity should impact firms' issuance and capital structure decisions. Further, a reduced discount rate may impact investment decisions by expanding the set of positive NPV projects. Finally, to the extent that illiquidity creates market frictions such as transaction costs, the trading environment should impact a firm's capital structure and payout policy.

To summarize, existing theories predict that liquidity may impact firm value via both cash flows and the discount rate.

The quantity of empirical research examining the relation between corporate finance decisions and the market microstructure environment has grown rapidly in recent years. This section will primarily focus on the more recent research which is predominantly focusing on the relation of liquidity to agency issues and the feedback impact of liquidity, although the recent research (which has been less active) relating liquidity to capital structure and payout policy decisions will also be discussed. Moreover, this section will focus mainly on how the liquidity channel impacts firm value. Of course, liquidity is not exogenous and cannot simply be included as a right-hand side variable; this discussion will begin with a summary of how this issue has been handled in the literature. Finally, to the extent that liquidity impacts firm value, executives should consider how corporate decisions affect liquidity. As a result, this discussion will conclude with recent research that examines how firms impact their own equity liquidity.

4.1 Identification issues

As with most corporate finance research, the relation between liquidity and corporate finance decisions poses some empirical hurdles in terms of issues related to endogeneity and causality. First, liquidity and corporate finance variables may be jointly determined by some unobservable omitted variable. For example, as described in Fang et al. [2009], it may be the case that high quality managers (an unobservable characteristic) both manage more liquid firms and yield better performance outcomes. Second, causality may run from the corporate finance variable to liquidity and/or from liquidity to the corporate finance variable. As an example, while there are many theories as to why increases to liquidity may affect an institution's propensity to monitor, it can also be argued that institutions are attracted to better governed firms in the first place. This increase in institutional trading will directly impact the liquidity of the firm. Moreover, firms that have better governance may be more transparent, resulting in lower adverse selection risk and bid-offer spreads. While the approaches to deal with these issues is not new (e.g., the use of firm fixed effects, instrumental variable techniques, or the difference-in-differences approach), this literature has identified specific mechanisms to overcome the issue of endogenous liquidity.

Fang et al. [2009], Gerken [2009], Edmans et al. [2013], Fang et al. [2013], Bharath et al. [2013], and Norli et al. [2014] utilize the tick size reductions in U.S. stock markets as exogenous shocks to liquidity (Section 3.4 discusses tick size reductions). Tick size reductions not only have a large impact on liquidity, but they are also appealing settings as there is cross-sectional variation in changes to liquidity (e.g., by volume, by price, etc.) that can be analyzed.

Bharath et al. [2013] discuss how liquidity shocks due to tick size reductions are "expected" and may be different from unexpected shocks if stock prices react to the announcement of a future liquidity shock event, rather than the actual occurrence of the liquidity shock event.[1] They therefore consider liquidity shocks that are unanticipated and in which the duration is unknown such as financial crises. They focus on the Russian default crisis in 1998 and the Asian financial crisis in 1997 (and to a lesser extent the U.S. financial crisis in 2008). They argue that such events are appealing because due to the foreign nature of these events, the main impact of the shock is transmitted through international capital markets (the liquidity effect) rather than a direct shock to firm fundamentals.

Although the 2001 tick size reduction resulted in a clear positive shock to liquidity, the time period is confounded with other governance events (such as Regulation FD which was ratified in October 2000) that may directly impact firm decisions independent of liquidity. To minimize the risk of confounding contemporaneous events, Back et al. [2013] and Balakrishnan et al. [2014] focus on natural experiments that are staggered through time: brokerage closures, market maker closures,

[1] Jain [2005] provides an example of this. He shows that there is a large, positive abnormal return to stocks when the exchange they are trading on announces a future switch from floor trading to electronic trading (a liquidity-improving event), but zero abnormal return when the switch is actually implemented.

and mergers of retail with institutional brokerage firms.[2] As discussed in Kelly and Ljungqvist [2012] and Balakrishnan et al. [2014], brokerage and market maker closures are not likely related to individual firm fundamentals, but do result in a substantial loss of liquidity. Alternatively, the acquisition of a brokerage that services retail clients by a brokerage firm that services institutional clients should have a significant positive impact on liquidity. As described in Kelly and Ljungqvist [2012], in this setting, research that is only available to institutional clients becomes available to retail clients, making previously private information now public. Of course, as discussed by Back et al. [2013], one downfall with the use of such events is that broker and market maker closures and mergers may also cause shocks to the firm's information environment that is independent of liquidity changes.

A couple of papers focus on exogenous liquidity measures at the firm level. Jayaraman and Milbourn [2012] focus on S&P 500 index additions and stock splits. Additions to the S&P 500 index affects a firm in two ways: given the attraction of index funds that trade for non-informational and short-term purposes, the stock may become more liquid, but prices may become less informative. However, Jayaraman and Milbourn [2012] argue that, as index decisions are not made at the firm level, such events should not serve as a signal regarding the firm's future prospects. Further, stock splits often result in the addition of new investors and rebalancing by existing investors; in principle stock splits should positively impact firm liquidity yet have no impact on underlying firm fundamentals. Lipson and Mortal [2009] use the location of the firm's headquarters as an instrument for liquidity. Given that liquidity is increasing in the number of potential investors in the firm, firms in densely populated areas should have higher liquidity (see Loughran and Schultz, 2005 for evidence). Further, following the findings by Christie and Schultz [1994] and Christie et al. [1994] that prior to reform in 1997 the avoidance of odd-eighth quotes by NASDAQ market makers inflated spreads, Lipson and Mortal [2009]

[2] Back et al. [2013] focuses on all three, while Balakrishnan et al. [2014] focuses only on brokerage closures.

use the proportion of odd-eighth quotes for NASDAQ-listed stock as an instrument for liquidity prior to 1997.

4.2 Liquidity and agency

Traditional theory linking liquidity, firm ownership, and governance has yielded conflicting predictions as to how liquidity influences monitoring by a large shareholder. Early papers focus on intervention as the primary governance mechanism for the large shareholder. Coffee [1991] and Bhide [1993] argue that, following poor performance, liquidity reduces the cost of exit for the large shareholder, increasing the attractiveness of selling shares relative to direct intervention. In this setting, weaker governance comes at the expense of liquidity. Conversely, Kyle and Vila [1991], Kahn and Winton [1998], Maug [1998], and Noe [2002] show that liquidity facilitates the formation of blockholders in the first place; in a more liquid market investors can recoup the cost of monitoring by purchasing shares at a price that does not yet fully reflect the value of intervention. In these papers, liquidity can improve incentives to monitor by reducing the cost of acquiring large positions. So combining these two literatures, theory predicts that liquidity may have opposing effects on firm governance, because it enables both block acquisition and block disposition. Alternatively, in Faure-Grimaud and Gromb [2004], liquidity facilitates monitoring via intervention due to increased price efficiency. In the event of a liquidity shock, a more informative stock price will partly reflect the efforts of the intervening shareholder prior to the date his activities are publicly observed.

While the earlier papers examining the relation between equity market liquidity and governance have focused on "control" by a large blockholder, we rarely see overt forms of intervention activities empirically. As discussed in Edmans et al. [2013], many institutions face legal restrictions (such as diversification requirements or "prudent person" rules), lack expertise, or face conflicts of interest, and as such, avoid activist activities. As a result of these barriers, only a small portion of U.S. firms have a controlling blockholder: La Porta et al. [1999]

document that only 10%–20% of firms have a blockholder with at least 20% equity ownership. However, a large fraction of firms have multiple small blockholders: using data from Dlugosz et al. [2006], Edmans and Manso [2011] show that if defining a blockholder as a 5% equity owner, 70% of firms have multiple blockholders, and 26% have at least four blockholders.

Given the hindrances in governance via intervention discussed above, recent theoretical research has focused on how liquidity can foster governance via exit [see e.g., Admati and Pfleiderer, 2009, Edmans, 2009, Edmans and Manso, 2011].[3] In these papers, it is the threat of exit itself that results in more efficient governance. The selling of shares by informed blockholders reduces the share price and hurts the manager ex-post (to the extent that the manager's wealth is tied to the stock price), causing the manager to make value maximizing decisions ex-ante. For the threat to be credible, the firm must have a liquid equity market that allows the large shareholder to sell at a low cost.

In Edmans [2009], liquidity and block ownership reduces myopia and encourages managers to make value maximizing long-term investments. In this setting, blockholders induce managers to choose the value maximizing long-term investment over investments creating short term profits, by collecting information on the value of the project and impounding it into the firm's stock price. In this paper, liquidity aids in improving firm decisions, by increasing the credibility of the exit threat and thus making "loyalty" (holding on to shares) more informative. Admati and Pfleiderer [2009] find that the effectiveness of the threat of exit depends on the nature of the agency problem. Specifically, governance via exit is efficient in preventing managers from undertaking non-value maximizing projects that increase private benefits, but less effective in persuading managers to undertake "good" projects that are privately costly to the manager. In Edmans and Manso [2011], the threat of exit is not credible unless the firm has multiple blockholders. In the presence of multiple blockholders, blockholders cannot coordinate, so they trade aggressively to compete for profits. As a result,

[3]See Edmans [2014] for an excellent review of the literature on mechanisms through which blockholders can exert governance.

prices more closely reflect fundamental value and punish the manager via reduced compensation; this threat disposes the manager to take the value maximizing action ex-ante.

Overall, the empirical evidence linking liquidity to efficient governance via intervention is somewhat mixed. Brav et al. [2008] focus on activism by hedge funds.[4] The benefit of examining hedge funds is that, unlike mutual funds and pension funds, they are not subject to regulations or conflicts of interest that may impede them from intervening in a firm. They identify activist events by identifying hedge funds that file a Schedule 13D with the SEC. Under federal securities law, investors who acquire more than a 5% stake and whose intentions are to force changes or seek control must file a 13D. The 13D filer must also declare its motive for the acquisition (i.e., to sell assets of the firm, change capitalization, or engage in merger and acquisition activity). The authors find that hedge funds target more liquid firms, which allows them to acquire shares in the open market quickly and at a low cost. Their results indicate that activism by hedge funds is an efficient governance mechanism: target firms benefit from increased payout, improved operating performance, and higher CEO turnover following the activist event.

Interestingly, Brav et al. [2008] find that in most instances, the intervening institution lacks control of the target firm (the median ownership stake is 9.1%) and 22% of firms targeted experience 13D filings by multiple blockholders. These results indicate that control by one large shareholder is not necessary — a more liquid equity environment consisting of multiple small blockholders can also facilitate efficient governance. In a related paper, Gerken [2009] shows that although liquidity increases the likelihood of blockholder formation, blockholders choose smaller stakes in more liquid firms so as not to pre-commit to monitoring. Alternatively, Collin-Dufresne and Fos [2014b] argue that the benefit of activism depends on the size of the stake the blockholder has accumulated. In their model, the activist's optimal effort level is increasing in the size of the acquired position, which itself depends on market liquidity.

[4] See Brav et al. [2010] for a survey on hedge fund activism.

Norli et al. [2014] use contested proxy events to understand how liquidity impacts the choice between costly intervention and no-governance. They find stock liquidity increases the probability of shareholder activism, although the positive effect of liquidity on the voice mechanism is diminished when firms are overvalued. Following Winton and Li [2006], in the setting of overvaluation, the profits to an informed trader from exiting may exceed the potential profits from intervention, with the incentive to trade especially strong when liquidity is high. Moreover, they link liquidity to activists' pre-event trading: activists actively and profitability trade prior to the announcement date of activism and liquidity has a direct positive effect on the percent of target shares accumulated in the pre-event period. Overall, their results provide evidence that the profits from informed trading — buying shares at a price that does not yet incorporate the value-add from intervention — are a substantial driver of the positive impact liquidity has on the incentive to engage in costly activism.

In a related paper, Collin-Dufresne and Fos [2014] examine the trading behavior of eventual activist investors. Using Schedule 13D filings (which report the date, price, and quantity of all trades in the target company the last 60 days), they find that although days these informed traders accumulate shares are characterized by positive and significant market-adjusted returns, such days are also characterized by lower measures of adverse selection and stock illiquidity. In fact, the measured price impact is almost 30% lower relative to the sample average on days that Schedule 13D filers trade. The results indicate that Schedule 13D filers (investors likely to become active) choose to trade when available liquidity is high.

Gantchev and Jotikasthira [2014] provide additional evidence that the expected profits from informed trading have a direct role in a shareholder's decision to accumulate shares and take an active role in the firm. They find that institutional selling volume significantly increases the probability of being targeted by an activist hedge fund, and at a daily frequency, institutional selling and hedge fund purchasing are synchronous in time. To distinguish informed institutional trading from exogenous liquidity shocks, they focus on institutional trading in non-target stocks outside of a target's industry. Taken together, their results indicate that target choice and intervention timing are both determined, in part, by market conditions — activist hedge funds are able to camouflage large purchases when institutional price pressure is high.

Fos [2013] argues that it is not the actual intervention (which we observe rarely), but rather the threat of intervention that serves as a governance mechanism. Moreover, liquidity is directly related to this threat of intervention as it reduces the cost of acquiring a large number of shares in the open market for the investor. Using both firms that experience proxy contests and those companies that do not experience a proxy contest between 1994 and 2009, he shows that when the likelihood of a proxy contest is high, firms respond with decreases in investment and compensation and increases in leverage, dividends, and CEO turnover, which translates into improved operating performance. In summary, liquidity renders the threat of intervention sufficiently high so that execution is not needed to mitigate agency issues.

Unlike the research discussed above, Back et al. [2013] find that liquidity is harmful for governance via intervention — liquidity increases the likelihood of the large investor "taking the Wall Street walk". They proxy for intervention by looking at hedge fund activism and shareholder proposals and find that blockholder activity increases with two exogenous shocks that reduce liquidity (brokerage and market maker closures) and decreases with an exogenous shock that increases liquidity (mergers of retail with institutional brokerage firms). They argue their results may differ from research that finds liquidity has a positive effect on activism due to their use of exogenous shocks to liquidity that have both a time and cross-sectional variation.

Although, research examining whether liquidity harms or benefits governance via blockholder intervention is not conclusive, research examining the relation between liquidity and governance via exit generally agrees this is a very successful monitoring mechanism. In the governance via exit setting, liquidity plays a dual role in encouraging monitoring: it allows blockholdings to form in the first place, but also stimulates information acquisition and increased trading in response to acquired information. This trading causes some of the private information about the firm's fundamentals (and the manager's effort) to be impounded in the firm's stock price.

Edmans et al. [2013] link these two literatures by focusing on a type of blockholder that has the ability to both intervene (govern via "voice") and to trade (govern via "exit"). As described by the authors, activist hedge funds have the full "menu" of governance mechanisms to choose from — they do not face legal barriers to intervention, have few business ties, and have a high performance-based fee structure to induce the optimal choice of governance activity. Using a sample of activist hedge funds that made block acquisitions of at least 5% between 1995 and 2010, the authors first show that liquidity increases the probability of block acquisition by 0.2%–0.5% versus the unconditional probability of 1.3%. To identify the governance mechanism chosen by the hedge fund (voice or exit), Edmans et al. [2013] differentiate between Schedule 13D filers (a required filing for those intending to engage in activism) and Schedule 13G filers (a required filing for passive investors who do not intend to change or influence control over the issuer). Conditional on block formation, they find a negative relation between liquidity and voice — a one standard deviation increase in liquidity reduces the likelihood of a 13D filing (yet increases the likelihood of a 13G filing) by 5%–7%.[5] This propensity to file a 13G in more liquid firms is concentrated in firms having high managerial wealth to stock performance sensitivity. Moreover, they find that governance via exit is regarded by the market as an efficient mechanism — compared to firms in the below median liquidity group, firms in the above median liquidity group experience 1.7% higher returns surrounding the 13G filing announcement date and significantly higher holding period returns. Further, firms targeted by hedge funds that file a 13G experience improved operating performance subsequent to the 13G filing — firms in the above liquidity subsample enjoy improvements in EBITDA/Assets of 1.5% higher than matched firms, whereas firms in the below median liquidity sample show no improvements relative to matched control firms. Overall,

[5] Results are robust when using decimalization as an exogenous shock to liquidity.

the authors show that liquidity has a small but positive effect on governance via voice (by increasing the unconditional incidence of activism) and a positive and significant effect on governance via exit.

Bharath et al. [2013] argue that even if exit is not observed, blockholders could still be governing effectively — to the extent that blockholders are informed traders, managers who have equity interest in the firm will take actions to improve firm value and thus to discourage blockholders from exit. Given the "threat" of exit is unobservable, they focus on stock liquidity, which is the mechanism that facilitates the power of exit threats. Using foreign financial crises as exogenous and unexpected liquidity shocks, the authors find that firms having large blockholdings experienced more significant declines in firm value — a one standard deviation increase in block ownership resulted in a 4.1% reduction in Q during the crisis period.[6] Of course, to the extent that liquidity may promote both intervention and exit, the authors attempt to isolate the mechanism that translates to improved firm value by focusing on a prediction that is specific to governance via exit. Specifically, in the blockholder exit models, the manager's sensitivity of wealth to stock price should make the exit threat more effective, whereas manager incentives play no role in the intervention models. They find that Q reductions for firms having large blockholdings were significantly more pronounced for firms providing high equity incentives to managers. Moreover, the authors provide evidence consistent with the predictions of Admati and Pfleiderer [2009] that governance via exit should be more effective in overcoming the "bad" agency problem — the results are stronger for cash rich firms in which the agency problems between shareholders and managers may be more severe. Overall, this paper takes an important step in linking the combination of block ownership and liquidity to firm value and shows that this relation is not only positive but economically large in magnitude.

Rather than look at firm value outcomes directly, Roosenboom et al. [2014] look at an important channel that can influence firm value: merger and acquisition decisions. Recognizing that theory has

[6]The authors also find the results are robust to using decimalization as an exogenous shock to liquidity.

conflicting predictions on the relation between liquidity and governance, the authors examine the premise that if liquidity results in improved institutional monitoring, then this should translate to better acquisition decisions. The acquisition setting offers some advantages: given that acquisition announcements are largely unexpected, reverse causality (where causality runs from the acquisition decision to liquidity) is less of a concern. Moreover, the acquisition is a setting where agency issues are especially a concern.

Unlike the papers described above that find a solely positive effect of liquidity on governance, Roosenboom et al. find that there is a tradeoff between liquidity and monitoring. They find that, in the full sample of acquisitions announced between 1998 and 2008, liquidity is significantly and negatively related to acquisition announcement returns. However, when the disciplining effect from the threat of exit is particularly high, this effect is mitigated. Specifically, the negative relation between liquidity and announcement returns is moderated in acquisitions of public targets using equity — an instance where agency concerns are particularly high and the exit threat should be most effective (as predicted by Admati and Pfleiderer [2009]). Moreover, this negative relation is mitigated when the firm has multiple blockholders (as predicted by Edmans and Manso, 2011, the threat of exit is more effective when multiple blockholders trade aggressively to compete for profits) and when CEO equity incentives are high. Similar results are found when considering other deal outcomes — the propensity to withdraw and CEO turnover following value reducing deals are also negatively related to liquidity.

Thus far we have seen that the effect of liquidity depends on the ownership structure of the firm. Monitoring via the threat of exit is effective when multiple blockholders are informed and can credibly signal firm fundamentals via selling or remaining loyal. In this setting, the blockholders care about long-run firm value and therefore induce the firm to choose the investment that is profitable in the long-term even though it may depress current earnings. However, liquidity may also attract a group of investors having a more short-term trading horizon and lower incentives to monitor.

Fang et al. [2013] find a negative relation between liquidity and innovation output (measured by patents and citations-per-patent): firms experiencing increases in liquidity in the top tercile following decimalization produce 18.5% fewer patents in the first three years following decimalization relative to a matched sample of firms that had liquidity changes in the bottom tercile. They find that following decimalization, firms experiencing large liquidity increases also experienced a large increase in holdings by non-dedicated institutional investors. They argue that the negative relation between liquidity and innovation is partly due to the attraction of more transient investors who pressure managers to cut long-term intangible investment. Further, they argue their findings could also be driven by liquidity increasing the likelihood of hostile takeover; as a result, managers have reduced incentive to invest in innovation.

Kang and Kim [2013] also examine how liquidity and institutional investment horizon interplay with firm monitoring. They find that subsequent to decimalization, the sensitivity of CEO turnover to short-term performance measures significantly decreased for firms with high dedicated institutional ownership and significantly increased for firms with high transient institutional ownership. Overall, these two papers show that the effectiveness of governance via trading varies with institutional investor heterogeneities.

To synthesize the literature that links liquidity to the effectiveness of monitoring by investors to curtail agency issues between managers and shareholders, most research points to liquidity as beneficial to firm governance — if liquidity is not effective at improving governance via intervention (the research is mixed on this relation), at a minimum liquidity can effectively facilitate governance via trading. However, both mechanisms are sensitive to the ownership structure of the firm — the incentives and regulatory requirements of the large shareholder, the size of the position acquired by the blockholder, and the number of blockholders acquiring a stake in the firm. Moreover, in the governance via exit literature, the key ingredient is the combination of investors becoming informed and producing valuable information about the firm and the manager responding to this information via choosing actions that maximize shareholder value. This latter step requires that the manager have sufficient interest in the equity of the firm.

4.3 Liquidity, stock price informativeness, and compensation

Although the research described in the previous section points to liquidity as a potential means to induce the manager to choose the value maximizing action over the action that produces private benefits for the manager, liquidity may also influence firm decisions even if agency issues are absent. To the extent that liquidity facilitates the acquisition of more private information by lowering the cost of exploiting it, which increases the amount of information impounded into the stock price, then liquidity can improve real investment decisions of the firm by facilitating information production that is incremental to the manager's information set.[7] Further, to the extent that a more informative stock price impounds more information about the manager's effort, then liquidity can yield more effective performance-sensitive managerial compensation contracts.

Fang et al. [2009] document a positive relation between liquidity and firm performance. Following decimalization, firms experiencing large increases in liquidity also had better performance as measured by Tobin's Q. Decomposing the market-to-book ratio into a price-to-operating earnings ratio, leverage ratio, and operating return on asset ratio, they find that more liquid stocks have higher operating return on assets and more equity in their capital structure than less liquid stocks, though price-to-operating earnings ratios are similar to less liquid stocks. This decomposition provides evidence that liquidity may impact firm value via improvements to operating performance rather than reductions in the discount rate.

[7]There is a spawning literature including Chen et al. [2007], that provides evidence that managers incorporate this private information produced by informed market participants into investment decisions. Given that Bond et al. [2011] already provide a comprehensive discussion of both the theoretical and empirical literature documenting the real effects of the financial markets, we do not include a discussion of this important literature in this article.

To establish the foundation of the causal relationship between liquidity and firm operating performance, Fang, Noe, and Tice examine both agency and non-agency based explanations. They find the impact of liquidity on firm performance is stronger for firms having high business risk. Given that firms having high cash flow uncertainty benefit the most from market-based information production, this result indicates that stock market feedback may be one mechanism driving the liquidity-performance relation. Moreover, they find the effect of liquidity on operating performance is magnified for firms having high pay-for-performance sensitivity, indicating that liquidity enhances the effectiveness of performance-based contracting. They find no evidence that agency-based explanations drive this relation — firms having a higher percentage of blockholders do not enjoy higher market-to-book ratios when liquidity is high. Further, liquidity has a similar effect on firms having low and high antitakeover measures. To summarize, this paper makes a strong case for a positive liquidity-firm performance relation. But unlike Edmans et al. [2013] and Bharath et al. [2013] who link the relation to a reduction in agency issues, the authors provide evidence that the underlying mechanism for this relation may be through valuable information produced by informed traders that is utilized by insiders and more effective incentive compensation schemes.

Jayaraman and Milbourn [2012] provide evidence that firms realize the benefit of performance based contracting when liquidity is high — both in the cross-section and time series, greater stock liquidity results in increased use of equity-based compensation relative to cash and the use of stock price over earnings in pay-for-performance measures. The authors argue that to the extent that liquidity reduces the cost of selling a manager's equity holdings, managers may have more preference for stock-based compensation relative to cash-based compensation when liquidity is high. Moreover, as described in Holmstrom and Tirole [1993] and Chordia et al. [2008], liquidity can facilitate a more efficient stock price, and as a result, prices are more informative about the manager's actions. When liquidity is high, firms can offer steeper equity-based incentives to managers. They find that liquidity is an economically important determinant of compensation design — a one standard deviation increase in liquidity reduces the proportion of cash pay relative to total pay by 7% and increases pay-for-performance-sensitivity of total wealth with respect to stock prices by 4% (relative to the mean). In related papers, Garvey and Swan [2002] and Kang and Liu [2008] document a positive link between stock price informativeness and CEO pay-for-performance sensitivity.

Ferreira et al. [2011] provide additional evidence on the value of stock price efficiency by linking price informativeness to board independence. On the one hand, a more informative stock price may enhance board effectiveness by producing information that can be used as an input in board member's monitoring activities. However, to the extent that the stock market plays a valuable monitoring role, board activities may not produce monitoring gains incremental to those produced by efficient prices. The author's results support this substitution effect: they find a negative relation between stock price informativeness and board independence. This substitution effect is stronger when institutional ownership and pay-for-performance sensitivity is high. In sum, this paper's findings support the research described in this section by providing evidence that the combination of an efficient stock price and pay-for-performance sensitivity is valuable enough to reduce the need for internal monitoring.

4.4 Liquidity and capital structure

Although the research linking liquidity to a lower equity premium (see Section 5.1 for a discussion of this literature) provides the natural starting point in understanding how liquidity should impact capital structure decisions, the literature examining how liquidity impacts leverage, equity issuance costs, and equity issuance timing is relatively new.

Bharath et al. [2009] evaluate the core assumption of the pecking order model — that information asymmetry is the main determinant of capital structure decisions. Specifically, they argue that to the extent that adverse selection in the microstructure sense (which measures information asymmetry between informed and uninformed traders) is correlated with information asymmetry between managers and shareholders, then microstructure measures of adverse selection should be a good proxy for the information advantage of managers (especially if we think of insiders as a subset of informed traders).

Bharath et al. [2009] estimate seven measures of the adverse selection component of the bid-offer spread, then use the first principal component of either levels or changes in these proxies to form an annual firm-level composite index of adverse selection. They argue that this index of adverse selection is a superior proxy — those used in previous literature may suffer from being static and persistent and often have multiple interpretations. Over the period from 1973 to 2002, they find that for every dollar of financing deficit, firms in the highest adverse selection index decile average 30 cents more debt than firms in the lowest adverse selection index decile. The ability of their adverse selection proxy to explain cross-sectional variation in leverage choice remains significant in regressions that include other known determinants of capital structure decisions (e.g., tangibility, Q, sales, profitability).

Lipson and Mortal [2009] consider liquidity more broadly (rather than focus on the adverse selection component) and document a negative relation between liquidity and leverage. After sorting firms into size and liquidity quintiles, the average debt-to-asset ratio for firms in the most liquid decile is 38%, while the average ratio for firms in the least liquid decile is 55%. After controlling for other factors that may influence capital structure decisions and using an instrumental variable for liquidity, the authors still document a negative relation between liquidity and leverage. In a related paper, Frieder and Martell [2006] consider reverse causality from leverage to liquidity. Using two-staged least squares, they find that controlling for the bi-directional relation between liquidity and leverage attenuates the impact of liquidity — a one standard deviation increase in spread yields a 3% increase in leverage (1.5% lower than when endogeneity is not considered).

Butler et al. [2005] examine how liquidity impacts a firm's cost of capital by linking liquidity to the direct cost of raising external capital — flotation costs in a seasoned equity offering. In playing the role of the intermediary in SEOs, investment banks are similar in spirit to traditional market makers. As a result, investment banks may face inventory and adverse selection risk, as well as incur transaction processing costs and search costs in identifying potential investors. They find that liquidity is a significant determinant of total (gross) investment banking fees — the average difference in gross fees for liquid versus illiquid stocks is 101 basis points per share issued, representing 21% of the average gross fee in the sample. Moreover, the impact of liquidity on fees is stronger for large equity issues where the market making role of the investment bank is most important. Overall, this paper points to liquidity impacting firm value via decreased costs of raising capital and is able to link liquidity to the cost of capital in a setting that does not rely on asset pricing model assumptions.

While the papers described above establish a link between stock liquidity and leverage, and Butler et al. [2005] document a link between underwriting fees and liquidity, Stulz et al. [2013] directly link liquidity to equity issuance decisions. To overcome identification issues that arise when relating idiosyncratic liquidity shocks to firm-level equity issuance, the authors focus on the relation between aggregate liquidity and aggregate equity issuance using a sample of equity issues from 36 countries between 1995 and 2008. In regressions of an aggregate country-level equity issuance count measure on lead, contemporaneous, and lagged liquidity innovations, the authors find that the coefficients on contemporaneous and lagged liquidity measures is positive and significant. The coefficients on lagged liquidity innovation measures are sizeable for three of the four quarters prior, yet smaller in magnitude than the contemporaneous measure. Overall, the cumulative effect on equity issuance over the next five quarters associated with a one standard deviation improvement in liquidity corresponds to an increase of roughly 40% of the unconditional mean of the quarterly equity issuance count measure. Liquidity innovations explain as much variation in equity issuance as proxies for market timing. Moreover, the relation between equity issuance and liquidity innovations remains with the inclusion of proxies for capital and general market conditions, growth prospects, asymmetric information, and investor sentiment. The authors argue that their results support the view that the demand for shares is less than perfectly elastic, and that firms take this into account when making financing decisions (see also Gao and Ritter [2010] who argue that in the SEO underwriting method decision, firms often choose the higher cost fully marketed offers to enhance demand elasticity).

Although the literature examining liquidity as a determinant of capital structure decisions is in its early stage, the existing literature points to a negative relation between liquidity and leverage, a negative relation between liquidity and direct equity issuance costs, and a positive relation between market liquidity and the timing of equity issuances. However, future research should better understand whether the impact of liquidity on capital structure is of appreciable magnitude and how it stacks up against other, more well-understood determinants of capital structure.

4.5 Liquidity and payout policy

Under Miller and Modigliani [1961], in a frictionless market investors can maintain liquidity preferences by investing or liquidating shares in the market — investors do not incur direct trading costs and their trades do not alter the price of the security. However, as we incorporate trading frictions into the market, cash distributions from the firm may be a less costly mechanism to achieve liquidity needs than selling shares in the market. As a result, dividends and repurchases may be valuable to investors of less liquid firms where trading costs are large.

Banerjee et al. [2007] find that less liquid firms are significantly more likely to pay dividends than more liquid firms. Over the period from 1993 to 2003, they find that a one standard deviation increase in liquidity results in a decrease in the probability of dividends from 48.59% to 30.82% (results are qualitatively similar in models that includes periods as far back as 1963). Further, the inclusion of liquidity variables dramatically improves the predictive power of a model of the proportion of dividend payers — they argue that the dramatic increase in the liquidity of the U.S. stock market in the past decades can explain the significant reduction in dividend payers over the years. Moreover, in the spirit of Pastor and Stambaugh [2003], they find that after firms initiate dividend payments, their stock returns become less sensitive to aggregate market liquidity. The authors conclude that investors perceive cash dividends and stock market liquidity as substitutes.

Given the existence of trading frictions, an implication of Miller and Modigliani [1961] is that less liquid firms should be more inclined to make cash distributions than more liquid firms. However, this yields no predictions as to the form of distribution (dividends versus repurchases). Of course liquidity may matter not only to investors but also to the firm itself. Large repurchases become costly when transaction costs are high and may impact subsequent liquidity by reducing the number of shares traded in the market. Moreover, repurchases may impact post-announcement liquidity negatively if the market views the firm as an informed trader or positively if the entrance of the firm as a trader induces competition amongst liquidity suppliers.

Barclay and Smith [1988] and Brockman and Chung [2001] find that spreads widen around repurchase periods; Singh et al. [1994], Wiggins [1994], Miller and McConnell [1995], find no evidence for widening spreads surrounding repurchase announcements days; and Cook et al. [2004] find evidence for reduced spreads on actual repurchase days. So the evidence as to how repurchases affect the firm's information environment is mixed. Brockman et al. [2008] argue that liquidity is a decision variable in corporate payout policy — firms trade off the tax and flexibility advantages of repurchases with the higher information asymmetry costs of repurchases. As information asymmetry concerns decline in periods of high liquidity, they argue that increased stock market liquidity has played a key role in the changing composition of payouts, where repurchases are claiming an increasingly larger portion of total payouts. Brav et al. [2005] provide survey evidence that managers consider liquidity in repurchase decisions, especially if repurchases reduce liquidity beyond some critical level.

Overall, improvements to U.S. equity market liquidity seem to be a significant contributor to both the increases in the levels of repurchases and in the proportion of repurchases relative to total cash distributions.

4.6 Corporate decisions' impact on liquidity

So far, we have seen many benefits of liquidity (improved governance via exit, reduced cost for intervention strategies, more information production, improved effectiveness of managerial compensation, etc.). Holmstrom and Tirole [1993] suggest that liquidity has potential costs as well (i.e., reduced intervention by blockholders, managerial myopia, and direct and indirect costs of disclosure). In their theory, firms should trade off the costs and benefits of liquidity and choose the level of liquidity that maximizes firm value.

Amihud and Mendelson [2008] note that managers who are concerned about increasing the liquidity of their firm can do so through corporate policies. Pham et al. [2003] and Ellul and Pagano [2006] provide evidence that managers consider liquidity outcomes in decisions regarding the initial formation of the firm. Pham et al. [2003] find that firms underprice IPO shares to attract a large number of small investors in order to create a more dispersed ownership structure. Firms that benefit from monitoring by concentrated owners need to underprice less. Ellul and Pagano [2006] argue that firms consider after-market illiquidity that arises from information asymmetries that persist following the IPO in the pricing decision. Firms with lower expected after-market liquidity and higher liquidity risk must underprice more in order to compensate investors for expected losses from trading with informed investors. In Mantecon and Poon [2009], when the benefit of post-market liquidity is high, firms will hire more reputable underwriters, underprice more, and make more price revisions in order to avoid a costly IPO failure. While the above papers link after-market liquidity considerations to IPO underpricing decisions, Eckbo and Norli [2005] show explicitly that a liquidity risk factor reduces expected returns to IPO stocks over the 5 years following the IPO month. Further, there is a large literature that provides evidence that liquidity increases following public offerings, and liquidity may explain the use of publicly underwritten offerings rather than private placements or rights offerings [e.g., see Tripathy and Rao, 1992, Denis and Kadlec, 1994, Kothare, 1997, Eckbo et al., 2000, Qian, 2011].

Massa and Xu [2013] show that the liquidity of an acquisition target affects merger and acquisition outcomes. They find that deals involving more liquid targets are more likely to include a public market bidder, more likely to be completed, and have higher premiums and announcement period returns. These results indicate a willingness by firms to pay a premium for liquidity and shows that M&A is one channel through which firms can gain liquidity.

Beginning with the theoretical work of Diamond and Verrecchia [1991], there is a literature that examines a firm's propensity to disclose information as an attempt to lower the information related part of trading costs. Balakrishnan et al. [2014] provide evidence that firms actively shape market liquidity. In response to an exogenous reduction in public information production (the loss of an analyst), firms disclose more timely and informative earnings guidance. For such firms, the benefit of liquidity outweighs the costs of disclosure (e.g., legal costs and costs to missed forecasts). This evidence supports the predictions of Diamond and Verrecchia, that firms can utilize disclosure activities to reduce information asymmetries and ultimately improve liquidity and firm value.

Coller and Yohn [1997] show that the decision to issue a management earnings forecast is related to widening bid-offer spreads for the firm. They find that spreads decline following a forecast, indicating that management earnings forecasts are an effective mechanism to reduce information asymmetry surrounding the firm.

Dass et al. [2013] focus on the endogenous choice of stock liquidity for innovative firms. Innovative firms may benefit from liquidity increases more than the average firm for multiple reasons. First, innovative firms often rely on equity rather than debt financing. Second, equity-based compensation contracts are often used to monitor valuable human capital inputs. Finally, given the nature of the firm, innovative firms benefit more from stock market feedback and can learn about the value of their output from market prices. They find that following an exogenous shock that reduced the cost of disclosure (a 1994 legislative change that strengthened intellectual property rights), innovative firms increased the frequency of earnings guidance and experienced increases in liquidity.

Chung et al. [2010] link operating and financial transparency to liquidity improvements; they argue firms can adopt internal governance measures to mitigate information asymmetry. There is also a literature that links country-level disclosure attributes and equity market liquidity [see, e.g., Bailey et al., 2006, Eleswarapu and Venkataraman, 2006, Lang et al., 2012]. Overall, this literature shows that firms can take deliberate steps to improve equity market liquidity and that liquidity is a channel through which improved transparency can positively impact firm value.

To summarize the literature on liquidity and corporate finance, abundant evidence points to liquidity as beneficial in many settings: liquidity increases the power of governance via exit, reduces the cost of governance via intervention, facilitates the entrance of informed traders who produce valuable information about the firm, enhances the effectiveness of equity-based compensation to managers, reduces the cost of equity financing, mitigates trading frictions investors encounter when trading in the market to recreate a preferred payout policy, and lowers the immediate transaction costs and subsequent liquidity costs for firms conducting large share repurchases. Further, the influence goes both ways. There is evidence that firms influence their own liquidity through a broad range of corporate decisions including internal governance standards, equity issuance form and pricing, share repurchases, acquisition targets, and disclosure timeliness and quality. Overall, equity market liquidity can lead to firm value gains via both increases to the cash flows of the firm and to decreases in the discount rate.

5

Liquidity and Asset Pricing

5.1 Liquidity premia

We begin by discussing the evidence on liquidity premia. The underpinning asset pricing argument is that stock prices reflect a premium that investors require for holding shares in more illiquid companies. The seminal work of Amihud and Mendelson [1986] demonstrates such a premium for the bid-offer spread measure of liquidity [see also, Jones, 2002]. Follow-up studies such as Brennan and Subrahmanyam [1996], and Amihud [2002] have used price-impact-based measures of liquidity to empirically document the role of liquidity as a determinant of expected returns in equity markets. Datar et al. [1998], and Brennan et al. [1998] suggest measuring liquidity by share turnover and find that this measure is negatively related to average returns. In recent work, Chordia et al. [2009] show that an illiquidity measure derived from Kyle's [1985] theory is positively related to future returns. Their measure incorporates parameters such as return volatility and volume into the illiquidity measure according to Kyle's expression for illiquidity in equilibrium.

Brennan et al. [2012] show that the pricing of illiquidity emanates principally from the sell-side. Allowing for differential price impacts on the buy- and sell-sides, they show that it is the sell-side price impact that is related to future expected returns. The idea is that agents seldom face needs to buy stock urgently, but unexpected needs for cash may force them to sell stock suddenly. These arguments suggest a larger premium for sell-side illiquidity, which is confirmed empirically.

Bali et al. [2014] show that liquidity shocks are not properly incorporated into stock prices, possibly because of investors' limited attention. They show that unanticipated shocks to the Amihud [2002] measure of liquidity have an impact on stock prices in subsequent time periods. Specifically, a shock indicating increased liquidity has a positive impact on subsequent period returns, indicating that the market does not impound the effect of the shock contemporaneously. This phenomenon seems destined to receive more attention in future research.

Asparouhova et al. [2010] argue that when microstructure noise are correlated with explanatory variables like liquidity, estimated premia for those variables are upward biased. After correction for this bias, they find smaller, but still significant liquidity premia in U.S. equity markets.

There also have been attempts to price the risk of trading with investors who have superior information. Easley et al. [2002] show that a structural measure of information asymmetry, PIN, is priced in the cross-section of returns. Duarte and Young [2009] decompose the PIN into components due to information and liquidity trading and find that it is the latter component that is priced, thus raising questions about whether PIN is a valid measure of information asymmetry. Sadka [2006] shows that systematic time-variations in an empirical estimate of the illiquidity parameter in a Kyle [1985]-type model of information asymmetry are priced in equity markets. Specifically, he decomposes illiquidity into a fixed cost component and a component that varies with the size of a trade and shows that unexpected variations in the systematic component command a premium in the cross-section of returns.

We now turn to studies on the effect of liquidity on the bond market. The classic scenario here is the yield differential commanded by different instruments of similar coupon and maturity. Thus, Amihud and Mendelson [1991] find such a differential between Treasury bills and notes with the same time to maturity. Krishnamurthy [2002] finds a price difference between the on-the-run and the most recent off-the-run 30-year Treasury bond and concludes that this differential results from a demand for liquid assets. Longstaff [2004] compares the yield differential between zero-coupon Treasury and Refcorp bonds and also finds a large liquidity premium. Friewald et al. [2012] examine the yield spread between corporate bonds and a similar duration Treasury bond or swap rate curve. They find that liquidity effects account for approximately 14% of the explained market-wide corporate yield spread changes. They also find that liquidity effects are significantly larger in periods of crisis and for speculative grade bonds.

If there are systematic shocks to liquidity, this indicates that aggregate liquidity may be a priced risk factor in asset markets. In two celebrated studies, Pastor and Stambaugh [2003] and Acharya and Pedersen [2005] use different liquidity measures to document evidence that systematic liquidity risk is related to expected stock returns. Pastor and Stambaugh use a construct that measures how much prices reverse per unit volume, whereas Acharya and Pedersen [2005] use the price-impact-based construct of Amihud [2002]. Both authors argue that both average illiquidity and systematic liquidity risk command premia in asset markets.

Lee [2011] tests the Liquidity-adjusted CAPM of Acharya and Pedersen [2005] in a global setting using the same Amihud [2002] price-impact-based proxy. He uses Datastream data for 50 countries from 1988 to 2007. He finds that liquidity factors in the LCAPM are priced. Specifically, after controlling for market risk, average liquidity level, size, and book-to-market factors, he finds that average returns depend on: (1) the covariance of individual stock liquidity with national liquidity, and (2) the covariance of individual stock liquidity with national and global returns. Further, he shows that U.S. liquidity risk is the predominant driving force of global liquidity risk. Liu [2006] uses a measure of liquidity based on the proportion of days with zero volume. He finds that a two factor model (that includes a market factor and a liquidity factor), explains returns well, and in fact accounts for the well-known book-to-market effect (wherein high book/market stocks earn high average returns and vice versa).

Eckbo and Norli [2002] examine six candidate factor representations of aggregate liquidity risk and find evidence that all but one of candidate factors is priced. Korajczyk and Sadka [2008] extract the common component of liquidity by analyzing the aggregate time-series of several measures, including share turnover, the effective spread, and the Amihud [2002] metric. They find evidence that liquidity risk in the form of variation in this component is priced in equity markets, and show that the level of liquidity also commands a premium, after controlling for liquidity risk.

5.2 Liquidity and pricing efficiency

We now review studies documenting the relation of liquidity to pricing efficiency. In an efficient market, return predictability from past information should be short-lived and minimal. However, evidence on the imbalance-return relation indicates that order flows are positively related to contemporaneous as well as future returns at daily and intraday horizons [viz, Chordia and Subrahmanyam, 2004, Chordia et al., 2005a]. This evidence has been attributed to the positive autocorrelations in order flows, which carries over to a predictive relation between returns and lagged order flows when the risk-bearing capacity of market makers is limited.

Chordia et al. [2008] explore how the predictive relation between returns and order flow varies through time and across different liquidity regimes. They show that intraday return predictability from order flows has declined substantially over time with reductions in the NYSE minimum tick size. Such predictability is markedly diminished during liquid periods within each tick regime. Prices are closer to random walk benchmarks during the more recent decimal tick size regime than in earlier ones in that first order return autocorrelations have declined as the minimum tick size was reduced.

The overall evidence is consistent with the hypothesis that increased arbitrage activity during more liquid periods enhances market efficiency.

In other work, Lesmond et al. [2004] and Korajczyk and Sadka [2005] argue that the momentum effect of Jegadeesh and Titman

[1993] is economically insignificant after accounting for illiquidity-related transaction costs. However, Frazzini et al. [2012] challenge these results using evidence from a large institutional money manager who traded nearly one trillion dollars across 19 developed equity markets from 1998 to 2011. This money manager profitably arbitrages the size, value, momentum, and short-term reversal anomalies. Critical to the profitability of their strategies, this money manager achieves blended transaction costs that are one-tenth the size of those found in the studies above by combining liquidity-supplying and liquidity-demanding trades. Similarly, Avramov et al. [2006] document that monthly cross-sectional reversals in stock returns [Jegadeesh, 1990] are present mostly in illiquid stocks (the implicit argument is that they are arbitraged away in the more liquid ones). Sadka and Scherbina [2007] examine the dispersion anomaly of Diether et al. [2002], that stocks with high dispersion of analysts' earnings forecasts earn lower returns, in depth. They show that dispersion anomaly is weaker in more liquid stocks and that increases in aggregate market liquidity have led to reductions in the statistical strength of the anomaly. Overall, the recent literature strongly supports the notion that liquidity enhances market efficiency.

Liquidity also has implications for the Law of One Price which states that two traded or synthesized instruments with the same future cash flows should trade at the same price due to arbitrage forces. The effectiveness of arbitrage in financial markets should depend on liquidity. Roll et al. [2007] test this notion in the context of the index futures/cash markets. They find that the speed of reversion of the basis toward zero (i.e., the cash/futures pricing discrepancy) is positively related to liquidity over a futures contract's lifetime. Further, innovations to the absolute basis and spreads are positively correlated, and spread innovations forecast shifts in the basis. The results suggest that liquidity plays a significant role in moving markets toward an efficient outcome.

In another interesting paper, however, Bakshi et al. [2000] consider the frequency with which index option prices violate theoretical comparative statics. They show that this frequency is associated with bid-offer spreads in the options market, suggesting that liquidity contributes to the arbitrage link between stock and options markets.

Bakshi, Cao, and Chen use data that span about three months, while Roll et al. [2007] use about 15 years of data.

To summarize, the literature on liquidity and asset pricing demonstrates that both average liquidity cost and liquidity risk are priced, liquidity enhances market efficiency, and liquidity strengthens the arbitrage linkage between related markets. These findings should serve as impetus to continue research in liquidity for many years to come.

6

Suggestions for Future Research

Research on liquidity has progressed by leaps and bounds in recent years. However, more needs to be done.

In the general liquidity literature, how do hidden orders and other new order types affect liquidity? How does the growth of dark pools affect liquidity? To increase liquidity more, should the U.S. tick size be reduced further down to a mil (i.e., one-tenth of a penny)? If there is a sudden plunge in available depth (i.e., the beginnings of a flash crash), how should exchanges adjust their trading processes until depth can be replenished? How do different trading mechanisms (fragmented vs. centralized, dark pools, continuous trading versus batch processing, etc.) affect liquidity? Should trading in foreign exchange, mortgage-backed securities, etc. be made more transparent and what impact would this have on liquidity? What can be done to reduce barriers to global exchange competition and what impact would this have on liquidity?

Given that the interaction of the liquidity and corporate finance literature is relatively new, we need a better understanding as to what settings liquidity has a first-order impact on corporate finance decisions. The existing literature points to liquidity as a significantly important, but relatively ignored, determinant of governance and managerial compensation effectiveness, and to a lesser extent, a relatively important determinant of capital structure and payout policy. To appropriately understand the magnitude of liquidity on corporate decisions, future research will need to identify creative ways to tackle the endogeneity issues this literature faces. Most importantly, more work is necessary on the liquidity-firm performance relation. While the few papers on this topic suggest that liquidity has a positive impact on firm outcomes, the evidence as to the driver of this relation is inconclusive. A better understanding of the mechanisms that drive this relation can help firms think about how to take steps to shape their own equity market liquidity. Moreover, to the extent that liquidity is beneficial, what are the least costly actions firms can take to achieve greater liquidity?

In the liquidity and asset pricing area, the time variation in liquidity premia embedded within asset prices needs further analysis. On what do such premia depend? For example, do they increase during crises, and do they vary across the business cycle? Thus, it might be that agents value the ability to liquidate assets cheaply more during crises and recessions, but we do not have a clear understanding of whether this simple intuition is borne out empirically. We also need a better understanding of how adverse selection due to information-based trading (a prime determinant of liquidity) varies across time. For example, do such costs have a systematic component? Do they also vary across the business cycle? What kind of private information is more relevant for liquidity? What about macroeconomic announcements or firm-specific events? The interaction of asset pricing and corporate finance also could be explored further in the context of liquidity. For example, it may be that managers are more prone to consider liquidity in their corporate financial decisions when their cost of capital is more likely to be influenced by liquidity (i.e., when liquidity premia are more likely to be high). These and other issues are likely to keep liquidity research at the forefront of investigations in the field of finance.

Acknowledgements

We thank Sheridan Titman (the editor), an anonymous referee, Yakov Amihud, Hank Bessembinder, Utpal Bhattacharya, Bidisha Chakrabarty, Kee Chung, Alex Edmans, Andrew Ellul, Joel Hasbrouck, Sahn-Wook Huh, Petko Kalev, Andrew Karolyi, Bryan Kelly, David Lesmond, Maureen O'Hara, Lubos Pastor, and Frederik Schlingemann for their helpful comments.

References

V. Acharya and L. Pedersen. Asset pricing with liquidity risk. *Journal of Financial Economics*, 77:385–410, 2005.

A. Admati and P. Pfleiderer. The "wall street walk" and shareholder activism: Exit as a form of voice. *Review of Financial Studies*, 22:2645–2085, 2009.

H. Ahn, C. Cao, and H. Choe. Decimalization and competition among exchanges: Evidence from the Toronto Stock Exchange cross-listed securities. *Journal of Financial Markets*, 1:51–87, 1998.

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

Y. Amihud and H. Mendelson. Dealership market: Market making with inventory. *Journal of Financial Economics*, 8:31–53, 1980.

Y. Amihud and H. Mendelson. Asset pricing and the bid-ask spread. *Journal of Financial Economics*, 17:223–249, 1986.

Y. Amihud and H. Mendelson. Liquidity, maturity, and the yields on U.S. treasury securities. *Journal of Finance*, 46:1411–1425, 1991.

Y. Amihud and H. Mendelson. Liquidity, the value of the firm, and corporate finance. *Journal of Applied Corporate Finance*, 20:32–45, 2008.

A. Anand and K. Venkataraman. Market conditions, obligations, and the economics of market making. Working paper, Southern Methodist University, 2014.

J. Angel. When finance meets physics: The impact of the speed of light on financial markets and their regulation. *The Financial Review*, 49:271–281, 2014.

J. Angel, L. Harris, and C. Spatt. Equity trading in the 21st century. *Quarterly Journal of Finance*, 1:1–53, 2011.

E. Asparouhova, H. Bessembinder, and I Kalcheva. Liquidity biases in asset pricing tests. *Journal of Financial Economics*, 96:215–237, 2010.

D. Avramov, T. Chordia, and A. Goyal. Liquidity and autocorrelations in individual stock returns. *Journal of Finance*, 61:2365–2394, 2006.

J. Bacidore. The impact of decimalization on market quality: An empirical investigation of the Toronto Stock Exchange. *Journal of Financial Intermediation*, 6:92–120, 1997.

J. Bacidore, R. Battalio, and R. Jennings. Order submission strategies, liquidity supply, and trading in pennies on the New York Stock Exchange. *Journal of Financial Markets*, 6:337–362, 2003.

K. Back, T. Li, and A. Ljungqvist. Liquidity and governance. National Bureau of Economic Research, No. w19669, 2013.

L. Baele, G. Bekaert, K. Inghelbrecht, and M. Wei. Flights to safety. Working paper, Tilburg University, 2014.

W. Bailey, A. Karolyi, and C. Salva. The economic consequences of increased disclosure: Evidence from international cross-listings. *Journal of Financial Economics*, 81:175–213, 2006.

G. Bakshi, C. Cao, and Z. Chen. Do call prices and the underlying stock always move in the same direction. *Review of Financial Studies*, 13:549–584, 2000.

K. Balakrishnan, M. Billings, B. Kelly, and A. Ljungqvist. Shaping liquidity: On the causal effects of voluntary disclosure. *Journal of Finance*, 2014. forthcoming.

T. Bali, L. Peng, Y. Shen, and Y. Tang. Liquidity shocks and stock market reactions. *Review of Financial Studies*, 27:1434–1485, 2014.

S. Banerjee, V. Gatchev, and P. Spindt. Stock market liquidity and firm dividend policy. *Journal of Financial and Quantitative Analysis*, 42:369–398, 2007.

M. Barclay and C. Smith Jr. Corporate payout policy: Cash dividends versus open-market repurchases. *Journal of Financial Economics*, 22:61–82, 1988.

M. Barclay, W. Christie, J. Harris, E. Kandel, and P. Schultz. Effects of market reform on the trading costs and depths of Nasdaq stocks. *Journal of Finance*, 54:1–34, 1999.

R. Battalio, B. Hatch, and R. Jennings. Toward a national market system for U.S. exchange-listed equity options. *Journal of Finance*, 59:933–962, 2004.

R. Battalio, S. Corwin, and R. Jennings. Can brokers have it all? On the relation between make take fees and limit order execution quality. Working paper, Indiana University, 2014.

A. Ben-Rephael. Flight-to-liquidity, market uncertainty, and the actions of mutual fund investors. Working paper, Indiana University, 2014.

G. Benston and R. Hagerman. Determinants of bid-asked spreads in the over-the-counter market. *Journal of Financial Economics*, 1:353–364, 1974.

H. Bessembinder. Issues in assessing trade execution costs. *Journal of Financial Markets*, 3:233–257, 2003a.

H. Bessembinder. Trade execution costs and market quality after decimalization. *Journal of Financial and Quantitative Analysis*, 38:747–777, 2003b.

H. Bessembinder and H. Kaufman. A comparison of trade execution costs for NYSE and NASDAQ-listed stocks. *Journal of Financial and Quantitative Analysis*, 32:287–310, 1997.

H. Bessembinder, W. Maxwell, and K. Venkataraman. Market transparency, liquidity externalities, and institutional trading costs in corporate bonds. *Journal of Financial Economics*, 82:251–288, 2006.

S. Bharath, P. Pasquariello, and G. Wu. Does asymmetric information drive capital structure decisions? *Review of Financial Studies*, 22:3211–3243, 2009.

S. Bharath, S. Jayaraman, and V. Nager. Exit as governance: An empirical analysis. *Journal of Finance*, 68:2515–2547, 2013.

A. Bhide. The hidden costs of stock market liquidity. *Journal of Financial Economics*, 34:31–51, 1993.

B. Biais and R. Green. The microstructure of the bond market in the 20th century. Working paper, Toulouse University, 2005.

E. Boehmer, G. Saar, and L. Yu. Lifting the veil: An analysis of pre-trade transparency at the NYSE. *Journal of Finance*, 60:783–815, 2005.

E. Boehmer, K. Fong, and J. Wu. International evidence on algorithmic trading. Working paper, University of New South Wales, 2014.

N. Bollen and R. Whaley. Are "teenies" better? *Journal of Portfolio Management*, 25:10–24, 1998.

P. Bond, A. Edmans, and I. Goldstein. The real effects of financial markets. National Bureau of Economic Research, Working Paper No. 17719, 2011.

A. Brav, J. R. Graham, C. R. Harvey, and R. Michaely. Payout policy in the 21st century. *Journal of Financial Economics*, 77:483–527, 2005.

A. Brav, W. Jiang, F. Partnoy, and R. Thomas. Hedge fund activism, corporate governance, and firm performance. *Journal of Finance*, 63:1729–1775, 2008.

A. Brav, W. Jiang, and H. Kim. Hedge fund activism: A review. *Foundations and Trends in Finance*, 4:185–246, 2010.

M. Brennan, T. Chordia, and A. Subrahmanyam. Alternative factor specifications, security characteristics, and the cross-section of expected stock returns. *Journal of Financial Economics*, 49:345–373, 1998.

M. Brennan, T. Chordia, A. Subrahmanyam, and Q. Tong. Sell-order liquidity and the cross-section of expected stock returns. *Journal of Financial Economics*, 105:523–541, 2012.

P. Brockman and D. Chung. Managerial timing and corporate liquidity: Evidence from actual share repurchases. *Journal of Financial Economics*, 61: 417–448, 2001.

P. Brockman, J. Howe, and S. Mortal. Stock market liquidity and the decision to repurchase. *Journal of Corporate Finance*, 14:446–459, 2008.

P. Brockman, D. Chung, and C. Perignon. Commonality in liquidity: A global perspective. *Journal of Financial and Quantitative Analysis*, 44:851–882, 2009.

J. Brogaard. The activity of high frequency traders. Working paper, University of Washington, 2011.

J. Brogaard, T. Hendershott, S. Hunt, and C. Ysusi. High-frequency trading and the execution costs of institutional investors. *The Financial Review*, 49:345–369, 2014a.

J. Brogaard, T. Hendershott, and R. Riordan. High-frequency trading and price discovery. *Review of Financial Studies*, 27:2267–2306, 2014b.

M. Brunnermeier and L. Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22:2201–2238, 2009.

E. Budish, P. Cramton, and J. Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. Working paper, University of Chicago, 2013.

E. Budish, P. Cramton, and J. Shim. Implementation details for frequent batch auctions slowing down markets to the blink of an eye. *American Economic Review*, 104:418–424, 2014.

A. Butler, G. Grullon, and J. Weston. Stock market liquidity and the cost of issuing equity. *Journal of Financial and Quantitative Analysis*, 40:331–348, 2005.

A. Carrion. Very fast money: High-frequency trading on the NASDAQ. *Journal of Financial Markets*, 16:680–711, 2013.

G. Cespa and T. Foucault. Illiquidity contagion and liquidity crashes. *Review of Financial Studies*, 27:1615–1660, 2014.

A. Chaboud, B. Chiquoine, E. Hjalmarsson, and C. Vega. Rise of the machines: Algorithmic trading in the foreign exchange market. Forthcoming in the *Journal of Finance*, 2014.

B. Chakrabarty, B. Li, V. Nguyen, and R. Van Ness. Trade classification algorithms for electronic communication networks. *Journal of Banking and Finance*, 31:3806–3821, 2006.

S. Chakravarty, R. Wood, and R. Van Ness. Decimals and liquidity: A study of the NYSE. *Journal of Financial Research*, 27:75–94, 2004.

Q. Chen, I. Goldstein, and W. Jiang. Price informativeness and investment sensitivity to stock price. *Review of Financial Studies*, 20:619–650, 2007.

J. Choi, D. Salandro, and K. Shastri. On the estimation of bid-ask spreads: Theory and evidence. *Journal of Financial and Quantitative Analysis*, 23: 219–230, 1988.

T. Chordia and A. Subrahmanyam. Order imbalance and individual stock returns. *Journal of Financial Economics*, 72:485–518, 2004.

T. Chordia, R. Roll, and A. Subrahmanyam. Commonality in liquidity. *Journal of Financial Economics*, 56:3–28, 2000.

T. Chordia, R. Roll, and A. Subrahmanyam. Evidence on the speed of convergence to market efficiency. *Journal of Financial Economics*, 76:271–292, 2005a.

T. Chordia, A. Sarkar, and A. Subrahmanyam. An empirical analysis of stock and bond market liquidity. *Review of Financial Studies*, 18:85–129, 2005b.

T. Chordia, R. Roll, and A. Subrahmanyam. Liquidity and market efficiency. *Journal of Financial Economics*, 87:249–268, 2008.

T. Chordia, S. Huh, and A. Subrahmanyam. Theory-based illiquidity and asset pricing. *Review of Financial Studies*, 22:3629–3668, 2009.

T. Chordia, R. Roll, and A. Subrahmanyam. Recent trends in trading activity and market quality. *Journal of Financial Economics*, 101:243–263, 2011.

W. Christie and P. Schultz. Why do NASDAQ market makers avoid odd-eighth quotes? *Journal of Finance*, 49:1813–1840, 1994.

W. Christie, J. Harris, and P. Schultz. Why did NASDAQ market makers stop avoiding odd- eighth quotes? *Journal of Finance*, 49:1841–1860, 1994.

K. Chung and C. Chuwonganant. Regulation NMS and market quality. *Financial Management*, 41:285–317, 2012.

K. Chung and C. Chuwonganant. Uncertainty, market structure, and liquidity. *Journal of Financial Economics*, 113:476–499, 2014.

K. Chung and H. Zhang. A simple approximation of intraday spreads using daily data. *Journal of Financial Markets*, 17:94–120, 2014.

K. Chung, J. Elder, and J. Kim. Corporate governance and liquidity. *Journal of Financial and Quantitative Analysis*, 45:265–291, 2010.

J. Coffee. Liquidity versus control: The institutional investor as corporate monitor. *Columbia Law Review*, 91:1277–1368, 1991.

M. Coller and T. Yohn. Management forecasts and information asymmetry: An examination of bid-ask spreads. *Journal of Accounting Research*, 35: 181–191, 1997.

P. Collin-Dufresne and V. Fos. Do prices reveal the presence of informed trading? *Journal of Finance*, 2014. forthcoming.

P. Collin-Dufresne and V. Fos. Moral hazard, informed trading, and stock prices. Working paper, University of Illinois, 2014b.

C. Comerton-Forde, T. Hendershott, C. Jones, P. Moulton, and M. Seasholes. Time variation in liquidity: The role of market-maker inventories and revenues. *Journal of Finance*, 65:295–331, 2010.

D. Cook, L. Krigman, and J. Leach. On the timing and execution of open market repurchases. *Review of Financial Studies*, 17:463–498, 2004.

S. Corwin and P. Schultz. A simple way to estimate bid-ask spreads from daily high and low prices. *Journal of Finance*, 67:719–760, 2012.

N. Dass, V. Nanda, and S. Xiao. Innovative firms and the endogenous choice of stock liquidity. Working paper, Georgia Institute of Technology and Rutgers, 2013.

V. Datar, N. Naik, and R. Radcliffe. Liquidity and stock returns: An alternative test. *Journal of Financial Markets*, 1:203–219, 1998.

R. Davies and S. Kim. Using matched samples to test for differences in trade execution costs. *Journal of Financial Markets*, 12:173–202, 2009.

H. Degryse, F. De Jong, M. Van Ravenswaaij, and G. Wuyts. Aggressive orders and the resiliency of a limit order market. *Review of Finance*, 9: 201–242, 2005.

D. Denis and G. Kadlec. Corporate events, trading activity, and the estimation of systematic risk: Evidence from equity offerings and share repurchases. *Journal of Finance*, 49:1787–1811, 1994.

P. Deuskar and T. Johnson. Market liquidity and flow-driven risk. *Review of Financial Studies*, 24:721–753, 2011.

D. Diamond and R. Verrecchia. Disclosure, liquidity and the cost of capital. *Journal of Finance*, 46:1325–1360, 1991.

J. Dick-Nielsen. Liquidity biases in TRACE. *Journal of Fixed Income*, 19: 43–55, 2009.

K. Diether, C. Malloy, and A. Scherbina. Differences of opinion and the cross-section of stock returns. *Journal of Finance*, 57:2113–2141, 2002.

S. Ding, J. Hanna, and T. Hendershott. How slow is the NBBO? A comparison with direct exchange feeds. *The Financial Review*, 49:313–332, 2014.

J. Dlugosz, P. Fahlenbrach, P. Gompers, and A. Metrick. Large blocks of stock: Prevalence, size, and measurement. *Journal of Corporate Finance*, 12:594–618, 2006.

J. Duarte and L. Young. Why is PIN priced? *Journal of Financial Economics*, 91:119–138, 2009.

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

D. Easley, S. Hvidkjaer, and M. O'Hara. Is information-based risk a determinant of asset returns? *Journal of Finance*, 57:2185–2221, 2002.

B. Eckbo and Ø. Norli. Pervasive liquidity risk. Working paper, Dartmouth College, 2002.

B. Eckbo and Ø. Norli. Liquidity risk, leverage and long-run IPO returns. *Journal of Corporate Finance*, 11:1–35, 2005.

B. Eckbo, R. Masulis, and Ø. Norli. Seasoned public offerings: Resolution of the 'new issues puzzle'. *Journal of Financial Economics*, 56:251–291, 2000.

A. Edmans. Blockholder trading, market efficiency, and managerial myopia. *Journal of Finance*, 64:2481–2513, 2009.

A. Edmans. Blockholders and corporate governance. *Annual Review of Financial Economics,* forthcoming, 2014.

A. Edmans and G. Manso. Governance through trading and intervention: A theory of multiple blockholders. *Review of Financial Studies*, 24:2396–2428, 2011.

A. Edmans, V. Fang, and E. Zur. The effect of liquidity on governance. *Review of Financial Studies*, 26:1443–1482, 2013.

A. Edwards, L. Harris, and M. Piwowar. Corporate bond market transaction costs and transparency. *Journal of Finance*, 62:1421–1451, 2007.

V. Eleswarapu and K. Venkataraman. The impact of legal and political institutions on equity trading costs: A cross country analysis. *Review of Financial Studies*, 19:1081–1111, 2006.

K. Ellis, R. Michaely, and M. O'Hara. The accuracy of trade classification rules: Evidence from Nasdaq. *Journal of Financial and Quantitative Analysis*, 35:529–552, 2000.

A. Ellul and M. Pagano. IPO underpricing and after-market liquidity. *Review of Financial Studies*, 19:381–421, 2006.

A. Ellul, C. Holden, P. Jain, and R. Jennings. Order dynamics: Recent evidence from the NYSE. *Journal of Empirical Finance*, 14:636–661, 2007.

V. Fang, T. Noe, and S. Tice. Stock market liquidity and firm value. *Journal of Financial Economics*, 94:150–169, 2009.

V. Fang, X. Tian, and S. Tice. Does stock liquidity enhance or impede firm innovation? *Journal of Finance,* forthcoming, 2013.

A. Faure-Grimaud and D. Gromb. Public trading and private incentives. *Review of Financial Studies*, 17:985–1014, 2004.

D. Ferreira, M. Ferreira, and C. Raposo. Board structure and price informativeness. *Journal of Financial Economics*, 99:523–545, 2011.

K. Fong, C. Holden, and C. Trzcinka. What are the best proxies for global research? Working paper, Indiana University, 2014.

V. Fos. The disciplinary effects of proxy contests. Working paper, University of Illinois, 2013.

A. Frazzini, R. Israel, and T. Moskowitz. Trading costs of asset pricing anomalies. Working paper, University of Chicago, 2012.

L. Frieder and R. Martell. On capital structure and the liquidity of a firm's stock. Working paper, Purdue University, 2006.

N. Friewald, R. Jankowitsch, and M. Subrahmanyam. Illiquidity or credit deterioration: A study of liquidity in the US corporate bond market during financial crises. *Journal of Financial Economics*, 105:18–36, 2012.

C. Furfine. Decimalization and market liquidity. *Economic Perspectives*, 27: 2–12, 2003.

J. Gai, C. Yao, and M. Ye. The externalities of high-frequency trading. Working paper, University of Illinois, 2013.

N. Gantchev and C. Jotikasthira. Activist hedge funds: Do they take cues from institutional exit? Working paper, University of North Carolina at Chapel Hill, 2014.

X. Gao and J. R. Ritter. The marketing of seasoned equity offerings. *Journal of Financial Economics*, 97:33–52, 2010.

G. Garvey and P. Swan. Agency problems are ameliorated by stock market liquidity: Monitoring, information and the use of stock-based compensation. Working paper, University of New South Wales, 2002.

T. George, G. Kaul, and M. Nimalendran. Estimation of the bid-ask spreads and its components: A new approach. *Review of Financial Studies*, 4:623–656, 1991.

W. Gerken. Blockholder ownership and corporate control: The role of liquidity. Working paper, Auburn University, 2009.

L. Glosten and L. Harris. Estimating the components of the bid-ask spread. *Journal of Financial Economics*, 21:123–142, 1988.

L. Glosten and P. Milgrom. Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14:71–100, 1985.

M. Goldstein and K. Kavajecz. Eighths, sixteenths and market depth: Changes in tick size and liquidity provision on the NYSE. *Journal of Financial Economics*, 56:125–149, 2000.

M. Goldstein, P. Kumar, and F. Graves. Computerized and high-frequency trading. *The Financial Review*, 49:177–202, 2014.

R. Goyenko, C. Holden, and C. Trzcinka. Do liquidity measures measure liquidity? *Journal of Financial Economics*, 92:153–181, 2009.

R. Goyenko, C. Ornthanalai, and S. Tang. Trading cost dynamics of market making in equity options. Working paper, McGill University, 2014.

D. Gromb and D. Vayanos. Equilibrium and welfare with financially constrained arbitrageurs. *Journal of Financial Economics*, 66:361–407, 2002.

L. Harris. Decimalization: A review of the arguments and evidence. Working paper, University of Southern California, 1997.

L. Harris. Trading in pennies: A survey of the issues. Working paper, University of Southern California, 1999.

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

J. Hasbrouck. Liquidity in the futures pits: Inferring market dynamics from incomplete data. *Journal of Financial and Quantitative Analysis*, 39:305–326, 2004.

J. Hasbrouck. Trading costs and returns for US equities: Estimating effective costs from daily data. *Journal of Finance*, 46:1445–1477, 2009.

J. Hasbrouck and G. Saar. Low-latency trading. *Journal of Financial Markets*, 16:646–679, 2013.

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

T. Hendershott and P. Moulton. Automation, speed, and stock market quality. *Journal of Financial Markets*, 14:568–604, 2011.

T. Hendershott, C. Jones, and A. Menkveld. Does algorithmic trading improve liquidity? *Journal of Finance*, 66:1–33, 2011.

T. Henker and J. Wang. On the importance of timing specifications in market microstructure research. *Journal of Financial Markets*, 9:162–179, 2006.

J. Hennessy and D. Patterson. *Computer Architecture: A Quantitative Approach*. Elsevier Morgan Kaufmann, Waltham, 2012.

M. Hertrich. Does credit risk impact liquidity risk? Evidence from credit default swap markets. Working paper, Center for Economic Science, University of Basel, 2014.

N. Hirschey. Do high-frequency traders anticipate buying and selling pressure? Working paper, London Business School, 2013.

T. Ho and H. Stoll. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, 9:47–73, 1981.

T. Ho and H. Stoll. The dynamics of dealer markets under competition. *Journal of Finance*, 38:1053–1074, 1983.

C. Holden and S. Jacobsen. Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions. *Journal of Finance*, 69: 1747–1785, 2014.

B. Holmstrom and J. Tirole. Market liquidity and performance monitoring. *Journal of Political Economy*, 101:678–709, 1993.

R. Huang and H. Stoll. Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE. *Journal of Financial Economics*, 41:313–357, 1996.

R. Huang and H. Stoll. The components of the bid-ask spread: A general approach. *Review of Financial Studies*, 10:995–1034, 1997.

G. Huberman and D. Halka. Systematic liquidity. *Journal of Financial Research*, 24:161–178, 2001.

P. Jain. Institutional design and liquidity at stock exchanges around the world. Working paper, University of Memphis, 2003.

P. Jain. Financial market design and the equity premium: Electronic versus floor trading. *Journal of Finance*, 60:2955–2985, 2005.

R. Jankowitsch, A. Nashikkar, and M. Subrahmanyam. Price dispersion in OTC markets: A new measure of liquidity. *Journal of Banking and Finance*, 35:343–357, 2011.

S. Jayaraman and T. Milbourn. The role of stock liquidity in executive compensation. *Accounting Review*, 87:537–563, 2012.

N. Jegadeesh. Evidence of predictable behavior of security returns. *Journal of Finance*, 45:881–898, 1990.

N. Jegadeesh and S. Titman. Returns to buying winners and selling losers: Implications for market efficiency. *Journal of Finance*, 48:65–92, 1993.

B. Johnson. *Algorithmic Trading and DMA: An Introduction to Direct Access Trading Strategies*. 4Myeloma Press, London, 2010.

C. Jones. A century of stock market liquidity and trading costs. Working paper, Columbia University, 2002.

C. Jones and M. Lipson. Sixteenths: Direct evidence on institutional execution costs. *Journal of Financial Economics*, 59:253–278, 2001.

B. Jovanovic and A. Menkveld. Middlemen in limit-order markets. Working paper, New York University, 2011.

C. Kahn and A. Winton. Ownership structure, speculation, and shareholder intervention. *Journal of Finance*, 53:99–129, 1998.

J. Kang and J. Shin. The role of high frequency traders in electronic limit order markets. Working paper, KAIST, 2012.

M. Kang and Y. Kim. Stock market liquidity and short-termism-driven CEO turnover. Working paper, Nanyang Technological University, 2013.

Q. Kang and Q. Liu. Stock market information production and CEO incentives. *Journal of Corporate Finance*, 14:484–498, 2008.

A. Karolyi, K. Lee, and M. van Dijk. Understanding commonality in liquidity around the world. *Journal of Financial Economics*, 105:82–112, 2012.

B. Kelly and A. Ljungqvist. Testing asymmetric-information asset pricing models. *Review of Financial Studies*, 25:1366–1413, 2012.

A. Kempf, D. Mayston, and P. Yadav. Resilience in limit order book markets: A dynamic view of liquidity. Working paper, Oklahoma University, 2008.

A. Kirilenko, A. Kyle, M. Samadi, and T. Tuzun. The flash crash: The impact of high frequency trading on an electronic market. Working paper, University of Maryland, 2011.

R. Korajczyk and R. Sadka. Are momentum profits robust to trading costs. *Journal of Finance*, 59:1039–1082, 2005.

R. Korajczyk and R. Sadka. Pricing the commonality across alternative measures of liquidity. *Journal of Financial Economics*, 87:45–72, 2008.

M. Kothare. The effects of equity issues on ownership structure and stock liquidity: A comparison of rights and public offerings. *Journal of Financial Economics*, 43:131–148, 1997.

A. Krishnamurthy. The bond/old-bond spread. *Journal of Financial Economics*, 66:463–506, 2002.

A. Kyle. Continuous auctions and insider trading. *Econometrica*, 53:1315–1335, 1985.

A. Kyle and J. Vila. Noise trading and takeovers. *RAND Journal of Economics*, 22:54–71, 1991.

R. La Porta, F. Lopez-de-Silanes, and A. Shleifer. Corporate ownership around the world. *Journal of Finance*, 54:471–517, 1999.

M. Lang, K. Lins, and M. Maffett. Transparency, liquidity, and valuation: International evidence on when transparency matters most. *Journal of Accounting Research*, 50:729–774, 2012.

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

K. Lee. The world price of liquidity risk. *Journal of Financial Economics*, 99:136–161, 2011.

D. Lesmond. Liquidity of emerging markets. *Journal of Financial Economics*, 77:411–452, 2005.

D. Lesmond, M. Schill, and C. Zhou. The illusory nature of momentum profits. *Journal of Financial Economics*, 71:349–380, 2004.

M. Lewis. *Flash Boys: A Wall Street Revolt*. W. W. Norton & Company, New York, 2014.

M. Lipson and S. Mortal. Liquidity and capital structure. *Journal of Financial Markets*, 12:611–644, 2009.

W. Liu. A liquidity-augmented capital asset pricing model. *Journal of Financial Economics*, 82:631–671, 2006.

F. Longstaff. The flight to liquidity premium in U.S. treasury bond prices. *The Journal of Business*, 77:511–526, 2004.

T. Loughran and P. Schultz. Liquidity: Urban versus rural firms. *Journal of Financial Economics*, 78:341–374, 2005.

A. Madhavan. Market microstructure: A survery. *Journal of Financial Markets*, 3:205–258, 2000.

A. Madhavan, M. Richardson, and M. Roomans. Why do security prices change? A transaction-level analysis of NYSE stocks. *Review of Financial Studies*, 10:1035–1064, 1996.

S. Mahanti, A. Nashikkar, M. Subrahmanyam, G. Chacko, and G. Mallik. Latent liquidity: A new measure of liquidity, with an application to corporate bonds. *Journal of Financial Economics*, 88:272–298, 2008.

K. Malinova, A. Park, and R. Riordan. Do retail traders suffer from high frequency traders? Working paper, University of Toronto, 2013.

L. Mancini, A. Ranaldo, and J. Wrampelmeyer. Liquidity in the foreign exchange market: Measurement, commonality, and risk premiums. *Journal of Finance*, 68:1805–1841, 2013.

T. Mantecon and P. Poon. An analysis of the liquidity benefits provided by secondary markets. *Journal of Banking & Finance*, 33:335–346, 2009.

B. Marshall, N. Nguyen, and N. Visaltanachoti. Commodity liquidity measurement and transaction costs. *Review of Financial Studies*, 25:599–638, 2012.

M. Massa and M. Xu. The value of (stock) liquidity in the market for M&A. *Journal of Financial and Quantitative Analysis,* forthcoming, 2013.

E. Maug. Large shareholders as monitors: Is there a tradeoff between liquidity and control? *Journal of Finance*, 53:65–98, 1998.

S. Mayhew. Competition, market structure, and bid-ask spreads in stock option markets. *Journal of Finance*, 57:931–958, 2002.

A. Menkveld. High frequency trading and the new-market makers. *Journal of Financial Markets*, 16:712–740, 2012.

J. Miller and J. McConnell. Open-market repurchase programs and bid–ask spreads on the NYSE: Implications for corporate payout policy. *Journal of Financial and Quantitative Analysis*, 30:365–382, 1995.

M. Miller and F. Modigliani. Dividend policy, growth and the valuation of shares. *Journal of Business*, 34:411–433, 1961.

G. Moore. Cramming more components onto integrated circuits. *Electronics*, 38:114–119, 1965.

S. Nagel. Evaporating liquidity. *Review of Financial Studies*, 25:2005–2039, 2012.

T. Noe. Investor activism and financial market structure. *Review of Financial Studies*, 15:289–315, 2002.

Ø. Norli, C. Ostergaard, and I. Schindele. Liquidity and shareholder activism. *Review of Financial Studies*, 2014. forthcoming.

A. Obizhaeva and J. Wang. Optimal trading strategy and supply/demand dynamics. *Journal of Financial Markets*, 9:201–242, 2012.

L. Pastor and R. Stambaugh. Liquidity risk and expected stock returns. *Journal of Political Economy*, 111:642–685, 2003.

A. Perold. The implementation shortfall: Paper versus reality. *Journal of Portfolio Management*, 14:4–9, 1988.

P. Pham, P. Kalev, and A. Steen. Underpricing, stock allocation, ownership structure and post-listing liquidity of newly listed firms. *Journal of Banking and Finance*, 27:919–947, 2003.

D. Porter and D. Weaver. Decimalization and market quality. *Financial Management*, 26:5–26, 1997.

H. Qian. Liquidity changes around seasoned equity issuance: Public offerings versus private placements. *Financial Review*, 46:127–149, 2011.

R. Riordan and A. Storkenmaier. Latency, liquidity and price discovery. *Journal of Financial Markets*, 15:416–437, 2012.

R. Roll. A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance*, 39:1127–1139, 1984.

R. Roll, E. Schwartz, and A. Subrahmanyam. Liquidity and the law of one price: The case of the futures/cash basis. *Journal of Finance*, 62:2201–2234, 2007.

T. Ronen and D. Weaver. Teenies' anyone? *Journal of Financial Markets*, 4: 231–260, 1998.

P. Roosenboom, F. Schlingemann, and M. Vasconcelos. Does stock liquidity affect incentives to monitor? Evidence from corporate takeovers. *Review of Financial Studies*, 27:2392–2433, 2014.

R. Sadka. Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk. *Journal of Financial Economics*, 80:309–349, 2006.

R. Sadka and A. Scherbina. Analyst disagreement, mispricing, and liquidity. *Journal of Finance*, 62:2367–2403, 2007.

R. Schestag, P. Schuster, and M. Uhrig-Homburg. Measuring liquidity in bond markets. Working paper, Karlsruhe Institute of Technology, 2013.

A. Singh, M. Zaman, and C. Krishnamurti. Liquidity changes associated with open market repurchases. *Financial Management*, 23:47–55, 1994.

H. Stoll. Inferring the components of the bid-ask spread: Theory and empirical tests. *Journal of Finance*, 44:115–134, 1989.

H. Stoll. Friction. *Journal of Finance*, 55:1479–1514, 2000.

H. R. Stoll. The supply of dealer services in securities markets. *Journal of Finance*, 33:1133–1151, 1978.

R. M. Stulz, D. Vagias, and M. A. van Dijk. Do firms issue more equity when markets are more liquid? Working paper, The Ohio State University and Erasmus University, 2013.

A. Subrahmanyam. Liquidity, return and order-flow linkages between reits and the stock market. *Real Estate Economics*, 35:383–408, 2007.

N. Tripathy and R. Rao. Adverse selection, spread behavior, and over-the-counter seasoned equity offerings. *Journal of Financial Research*, 15:39–56, 1992.

U.S. Securities and Exchange Commission. Equity market structure literature review: Part ii: High frequency trading. http://www.sec.gov/market structure/research/hft_lit_review_march_2014.pdf, 2014.

B. Van Ness, R. Van Ness, and S. Pruitt. The impact of the reduction in tick increments in major U.S. markets on spreads, depth, and volatility. *Review of Quantitative Finance and Accounting*, 15:153–167, 2000.

D. Vayanos and J. Wang. Theories of liquidity. *Foundations and Trends in Finance*, 6:221–317, 2011.

J. Wiggins. Open market stock repurchase programs and liquidity. *Journal of Financial Research*, 17:217–229, 1994.

A. Winton and S. Li. When do institutional investors become activists? Trading incentives and shareholder activism. Working paper, University of Minnesota, 2006.