

An Index-Based Measure of Liquidity¹

George Chacko Sanjiv Das
Santa Clara University Santa Clara University

Rong Fan
Gifford Fong Associates

August 2010

¹We thank Yakov Amihud for valuable suggestions. We also thank attendees at presentations at the Risk Conference (New York 2009), the Portfolio Management Conference (Frankfurt 2010), Neuberger Berman, ICE Capital, and the Auda Alternative Investment Roundtables in Munich and Cologne for their comments. The paper has improved greatly from this input.

Abstract

An Index-Based Measure of Liquidity

The liquidity shocks of '08-'09 revealed that the measures of liquidity risk being used in most financial institutions - measures constructed from work in leading academic papers - turned out to be woefully inadequate. While many papers calculate liquidity measures of securities, the process used by these papers - the construction of long-short portfolios based on liquidity proxies - introduces errors such as extraneous risk factors and hedging error. In this paper we develop a new measure for liquidity risk using exchange-traded funds (ETFs) that attempts to minimize this error, and thus produce as pure a liquidity measure as possible. We form long-short portfolios that are long ETFs and short the underlying components of that ETF. Consequently, we are long and short the exact same set of underlying securities with the same weights. We show, as a result, that pricing discrepancies between the long and short positions are due to liquidity differences between the ETF and its underlying components. We use this trait to construct liquidity risk factors in a number of markets. We then use these liquidity factors to show that financial institutions (broadly defined) carry substantial amounts of liquidity risk in their balance sheets.

1 Introduction

The eight months beginning in September of 2008 demonstrated to the world the enormous destructive effects that a global liquidity shock can have on not only the financial markets but the real economy. One of the most widely used quotes during this time period was provided by the “sage of Omaha,” Warren Buffet, who said “You only find out who is swimming naked when the tide goes out.”¹ In writing this, Buffett essentially meant that an investor can only tell how much risk he was really running in his balance sheet when the downside of those risks manifested itself. This was certainly true for the case of liquidity risk. Many financial institutions during this period (including banks, insurance companies, funds, endowments, and pension funds) discovered to their detriment just how much liquidity risk they were carrying on their balance sheets. While liquidity risk has become a topic of greater interest for academics during the last decade, most measures of liquidity risk are confounded by the fact that they mix liquidity risk with varying amounts of other risk factors. As a result, even those investors who were utilizing measures of liquidity risk in their risk management processes discovered that their measures did not adequately capture the true level of liquidity risk in their balance sheets.

In this paper, we take on three tasks. First, we develop a new liquidity risk measure which attempts to completely isolate liquidity risk from all other risk factors, thereby producing a relatively pure liquidity risk factor. Second, we utilize this measure to answer the question many investors have been asking: “How much am I getting compensated for taking liquidity risk?” Third, we demonstrate that financial institutions need to be especially cognizant of liquidity risk as their balance sheets contain significant amounts of liquidity risk due to the nature of their normal business activities.

1.1 Liquidity Risk

Liquidity is a fundamental need of all investors at some point of time. The need for liquidity can be equated to the need for immediacy in doing a transaction, whether buying or selling. Liquidity volatility, or risk, is essentially the risk that an investor may need transaction immediacy at a particularly convenient or inconvenient time in the markets, i.e., when the price of transaction immediacy is low or high, respectively.

¹Chairman’s Letter, 2001 Berkshire Hathaway Annual Report

However, not all investors face the same degree of liquidity risk. Investors with long-dated liabilities, for example, face less risk of suddenly needing transaction immediacy in the short-term. Investors like these, who face little risk of requiring sudden liquidity, should then be able to collect a premium for providing liquidity to those investors who do.

Investors with long investment horizons (such as those with long-dated liabilities) are perfectly positioned to serve as providers of liquidity. Their long investment horizon assures that the risk of requiring short-term transaction immediacy, i.e., liquidity risk, is minimal. Thus, they can earn a premium for providing liquidity services, and thereby bearing liquidity risk. Because of their long horizons, this liquidity risk is minimal to them and therefore, they can be viewed as earning a premium (assuming that bearing liquidity risk earns a premium) for bearing a risk that is not really a risk to them. Thus, it can be potentially a source of “true alpha.”

Many financial institutions utilize this concept and structure their balance sheets to essentially provide these liquidity services (and thereby bear liquidity risk) to other investors in return for a premium—a liquidity premium. Consider an endowment fund as an example.² An endowment fund’s liabilities, essentially its annual contributions to the university’s operating budgets, are both short- and long-dated.³ An endowment fund can therefore use its long-dated liabilities to hold very illiquid assets against them—and most do. Many endowments invest in illiquid assets such as private equity, real estate, and infrastructure. By doing so, they are using their long investment horizon stemming from their long-dated liabilities to hold assets that other investors with shorter investment horizons are not capable of holding. In this example, the liquidity risk that the endowment fund is bearing is not really risky to the fund due to its long-dated liabilities, and therefore the liquidity premium can potentially be viewed as alpha to the fund.

On the other hand, many financial institutions structure their balance sheets to take liquidity risk by intentionally mismatching their assets with some of their liabilities along the liquidity risk dimension. Consider a bank for example. The largest source

²Merton (1993) shows that the optimal investment strategy for university endowments should take into account each university’s unique characteristics, one of which is the liability structure of the endowment, i.e., the structure of the university’s expected operating budget in the future.

³The long-dated liabilities stem from the fact that the endowment is responsible for contributing to the university’s operating budgets in perpetuity.

of a bank's liabilities stem from its deposits, any of which has a high degree of liquidity risk—essentially, any deposit can be withdrawn by a depositor at any time. Banks, however, pool deposits together and thereby decrease the liquidity risk of the total pool. Any single deposit may be withdrawn at any time and therefore have a high degree of liquidity risk, but it is unlikely that all depositors will simultaneously withdraw their deposits at the same time. Therefore, part of the pool of deposits has a very low degree of liquidity risk, which the bank then uses to match up against illiquid assets, which pay a substantial liquidity premium. Of course, there is still a risk that there is a “run on the bank” and a substantial fraction of the deposits are demanded simultaneously. In such a scenario, the bank would have to dispose off illiquid assets quickly, and this need for transaction immediacy would entail the bank taking substantial losses. A bank holds equity capital to mitigate this risk and any resulting losses, which can never be completely eliminated. Thus, ultimately the bank bears this risk resulting from a liquidity mismatch and earns a liquidity risk premium for doing so. This lowering of liquidity risk by forming a pool of illiquid liabilities, can be viewed as a form of liquidity risk diversification. The liquidity diversification of the liability pool allows the bank to mismatch the liquidity risk of some of its assets with the liquidity risk of some of its liabilities. Ultimately however, the bank is taking some amount of risk due to this mismatch, and therefore the premium it earns is not alpha as in the endowment example above, but simply a fair return for taking that risk.

Hedge funds are similar to banks in the structure of their balance sheets. Hedge funds' liabilities consist of the equity capital provided by investors, who are typically locked in contractually to specific liquidity terms.⁴ In addition, there can be debt capital in the form of margin borrowing or short positions in securities. The debt capital usually has a shorter term than the equity capital—normally, it can be called instantly as, for example, in the case of a margin call. Then, many hedge funds mismatch the liquidity of the investments they take on the asset sides of their balance sheets with the liquidity terms of the equity capital and the leverage they employ. This mismatch in turn results in the fund taking liquidity risk and, therefore, at least part of the return the fund earns is simply compensation for bearing this risk.

Finally, it should be noted that endowment funds, banks, and hedge funds are merely examples. Virtually all financial institutions invest in relatively illiquid securities,

⁴For example, a “monthly, 30-day” fund allows its investors to withdraw their capital at the end of any month, but they must provide the fund with 30 days notice of their intention to withdraw.

either to take advantage of the long-dated nature of their liabilities or they run liquidity mismatches between their assets and liabilities—in either case, they are trying to collect a liquidity premium. Other institutions such as insurance companies, private equity funds, and pension funds also use similar strategies for some or all of their balance sheets.⁵ The approach we take in this paper for measuring liquidity risk utilizes precisely this same concept of mismatching the liquidity of assets and liabilities. We measure the level of liquidity risk by the difference in price between two assets which are otherwise identical except for the liquidity of the two. Essentially, we calculate the value of a long-short portfolio where the long and short positions are identical but their prices are not because the long position is more liquid than the short—in the theoretical development that follows, we will show how these positions may be interpreted as call and put options on trading immediacy. Our methodology is very general and applies to any market in which an exchange-traded fund (ETF) is traded.

With liquidity risk so widespread, it is important for any candidate liquidity measure to be capable of detecting and quantifying liquidity risk exposure in financial institutions, which we have thus far only described qualitatively. In this paper one of our goals is to demonstrate that our liquidity risk measure can be useful in detecting and measuring the level of liquidity risk residing on the balance sheets of financial institutions.

1.2 Prior Literature on Liquidity Risk

The academic papers that have been written to measure and analyze liquidity risk have used this same concept of mismatching asset and liability liquidity risk to create net liquidity risk on a balance sheet. The difficulty with liquidity risk is that it is a latent risk factor—it cannot be directly observed. Therefore, academics have used characteristics about securities to essentially instrument for liquidity risk, and then created long-short portfolios where the long and short positions are mismatched on these characteristics. This, in turn, gives rise to a portfolio with liquidity risk whose return can therefore be viewed as a liquidity premium—essentially a liquidity index.

⁵One effect of the prevalence of liquidity risk on the balance sheets of financial institutions and the interconnections of these institutions (the liabilities of many financial institutions are the assets of other financial institutions) is the widespread transmission of liquidity shocks around the world—such as the one we just experienced from mid-2007 thru 2009.

The problem of course is that because liquidity risk is not observable, especially situations of high liquidity risk, one is never sure whether the characteristics being used as instruments are also instrumenting for other risk factors. If they are, then the resulting long-short portfolio is a mix of liquidity risk and other risk factors.

While liquidity risk has received considerable attention recently in the academic literature, owing primarily to financial market disruptions in late Summer 2007 caused by liquidity problems in the mortgage securitization market, there had already been a growing literature addressing this issue well before the recent market turbulence. Many papers have investigated the importance of liquidity for explaining returns, using data from the equity markets. Amihud and Mendelson (1986), Brennan and Subrahmanyam (1996), Brennan, Chordia and Subrahmanyam (1998), Datar, Naik and Radcliffe (1998), Chordia, Roll and Subrahmanyam (2002) have all found positive relationships between stock returns and overall liquidity as measured by spreads, depth, and volume. However, Chordia, Subrahmanyam and Anshuman (2001) find a negative relationship between liquidity and expected returns, while Hasbrouck and Seppi (2001) find no relationship. Finally, Huberman and Halka (2001) and Pastor and Stambaugh (2003) examined the question of whether liquidity risk is systematic. Both papers find substantial systematic components in liquidity risk. More recent work, for example from Acharya and Pedersen (2005), Sadka (2006), Korajczyk and Sadka (2008), Li, Mooradian and Zhang (2007), and Das and Hanouna (2010), seem to find more positive results for the pricing of liquidity risk, though in many cases the pricing is small.

The mixed results on liquidity pricing in the equity markets is likely the result of performing liquidity tests in a market where liquidity is typically pervasive and therefore, an unimportant characteristic. In fact, next to the Treasury market, the US equity market is the most liquid in the world. Work done in markets where the effects of illiquidity are pronounced, seems to indicate that liquidity risk is in fact priced. For example, Chacko (2009), utilizing holdings data instead of trading data, analyzes the question of whether liquidity risk is priced in the US corporate bond market.⁶ Using data for the corporate bond market (excluding convertible corporate bonds), he finds strong evidence for a systematic liquidity risk factor. Longstaff, Mithal and Neis (2005), Ericsson and Renault (2002), and Chen, Lesmond and Wei (2007) try to relate corporate bond liquidity to yield spreads as a way of ascertaining the pricing

⁶See Goodhart and O'Hara (1997), and Edwards, Harris and Piwowar (2007).

of liquidity risk—and all find some evidence that liquidity risk is priced.⁷

1.3 An ETF-based Measure of Liquidity Risk

In all of the papers on liquidity, a common factor is that security characteristics such as turnover, bid-ask spread, etc., are used to proxy for liquidity, as liquidity is unobservable directly. These papers directly or indirectly define liquidity risk as being long a set of securities that have a “high” characteristic measure and short a set of securities that have a “low” characteristic measure. Such an approach introduces other risk factors into this long-short portfolio, and hence the true risk factor that is measured is a blend of liquidity and other risk factors. Usually, authors are careful to control for other risk factors by hedging, but even in this case extraneous risk factors are still present due to our inability to measure all of the other possible risk factors as well as measurement error in determining contemporaneous exposure using historical data. The best evidence for this is the empirical observation that if one calculates the correlations of various risk factors produced by these papers in the same markets and time periods, they typically range from 0.1 to 0.4. Assuming the papers are all taking reasonable approaches for calculating liquidity, i.e., they are in fact picking up a true liquidity signal, one would expect these correlations to be much higher. Our interpretation is that there is material error being introduced into the liquidity calculation process.

In this paper we introduce a new way to calculate liquidity risk that is not based on characteristic-based measures. The goal of this approach is to minimize the error introduced through extraneous factors and, as a result, produce a liquidity measure that is as uncontaminated as possible. Our approach for measuring liquidity risk in a particular market is to form a portfolio that is long an ETF in that market and short the individual securities of that ETF—the securities of the short portfolio are weighted in a manner that is identical to the weighting scheme of those securities in the ETF. So, essentially we are going long and short the exact same securities except that the long positions are obtained through an ETF while the short positions are obtained through direct shorting of the individual securities.

⁷Elton, Gruber, Agrawal and Mann (2001) and Huang and Huang (2003) on the other hand find evidence that the tax effects play a greater role than liquidity, while Campbell and Taksler (2003) find equity volatility to be more important than liquidity. Das and Hanouma (2009) find that equity liquidity is important even in explaining credit default swap spreads.

Absent any frictions, this long-short portfolio should produce zero return always. However, we know that by almost any liquidity measure, ETFs are generally more liquid than the individual securities in the ETF. Therefore, in practice we observe systematic pricing discrepancies between the ETF and its underlying components. Because we are long and short the same securities, there can reasonably be only two explanations for systematic pricing discrepancies: 1) the market is inefficient and these pricing discrepancies represent arbitrages, or 2) the long and the short positions differ on liquidity. As we will assume, like other papers in this area, that widespread, systematic arbitrages are not possible, the pricing discrepancy must be due to the liquidity-based explanation. Essentially, the more liquid security has a higher price due to a lower liquidity premium, while its illiquid counterpart has a lower price due to a higher liquidity premium.

The structure of this paper is as follows. In the next section, Section 2, we define the data and process that we use to construct sample ETF-based long-short portfolios. We present an options-based analysis that generates a formula for liquidity, and we also connect this formula to the literature on differences in price adjustment speeds. In Section 3, we then run a number of empirical tests to demonstrate that the time-series of these long-short portfolios are in fact related to liquidity risk. As our methodology for calculating liquidity risk hedges out all other risk factors and frictions, we argue that liquidity is the only remaining risk factor—and the empirical results provide evidence for this argument. Finally, in Section 4 we use these liquidity measures to test our hypothesis that liquidity risk is related to the performance of financial institutions - we use hedge funds and an index of financial services firms (including banks and insurance companies) as representative financial institutions. We find that liquidity risk is indeed a systematic risk factor in explaining the returns of financial institutions. The final section then concludes.

2 Methodology and Data

In this section we lay out the basic methodology for calculation of our liquidity measure and, subsequently, our liquidity factor. We then provide details about the data on which we apply this methodology to generate our empirical findings.

2.1 Methodology

Most liquidity factors involve the utilization of liquidity measures that inadvertently include other risks beyond just liquidity. This results in a noisy liquidity measure. Our methodology for calculating liquidity in a market is fairly simple and avoids the problem of residual risk factors. Essentially, we take a traded index, an ETF, and compare the price of the index to the price of a portfolio (NAV) containing the same underlying components as the ETF. Because the index and the portfolio contain identical securities with identical weightings, the difference between their prices can only be attributed to liquidity. Essentially, the portfolio of ETF components is less liquid than the ETF itself and, therefore, investors face liquidity risk when buying the portfolio. These investors therefore are only willing to purchase these securities at a discount that compensates for the higher liquidity risk. Similarly, the ETF is much more liquid than the portfolio, and therefore investors are willing to pay a premium for the ETF relative to the portfolio of its components.

There are an infinite number of mathematical functions that we can use to compare the ETF price and the value of the portfolio of its components, or NAV. We look to the theoretical market microstructure literature to guide us in the determination of this mathematical function. A large literature exists that explores the inventory cost model of transaction costs (Garman (1976), Stoll (1978), Amihud and Mendelson (1986), and Ho and Stoll (1981)). An interesting line of this literature (Copeland and Galai (1983) and Harris (2003)) has utilized contingent claims to model bid and ask prices. Recently, Chacko, Jurek and Stafford (2008) showed that bid and ask prices could be modeled specifically as American calls and puts. The intuition behind their approach is to recognize that when a bid or an ask price is put out into the market the agent quoting the bid or ask price⁸ is essentially writing an American option. For example when a sell limit order is put out by an investor, the investor has written an American call option to the market maker. The call option is written on the fundamental value of the security, and the exercise price is the limit price

⁸Bid and ask prices could be quoted by either a market maker or an investor. For example in bond markets the market maker, or bond dealer, puts out bid and ask quotes for the bonds in which he is willing to deal. On the other hand in equity markets investors can put out bid and ask quotes through the mechanism of limit orders. Thus, equity investors could actually become de facto market makers if their buy and sell limit order prices are inside the quotes of the market makers, or specialists. The label of market maker versus investor is simply determined by who has the least need for transaction immediacy or lower inventory costs. This is how we proceed in this paper.

set by the investor. When the fundamental value of the security exceeds the limit price, or exercise price, by a sufficient, optimal amount, the option is exercised by the market maker, and a transaction takes place. Therefore, the value of the call option essentially represents the bid price of a security, and the value of the corresponding put option represents the ask price. Because a bid-ask spread is essentially a proxy for the level of illiquidity in a market, we can use the value of the call plus the value of the put as a measure of the liquidity in the market.

In applying a contingent-claims approach to our setting, we can think of illiquidity as the value of an option to exchange the ETF for the NAV, i.e., the difference between the ETF and the NAV upon exercise of the option is the effective transaction cost, and therefore the value of this option is the value of liquidity. If the transaction is initiated by a buyer of the ETF—say, by placing a limit order—then, an option to exchange the ETF for the NAV is written by the prospective buyer, i.e., a seller into the limit order would exercise a put of the ETF at strike NAV. On the other hand, if the transaction is initiated by a seller of the ETF, then an option to exchange the NAV for the ETF is written by the prospective seller, i.e., a buyer would exercise a call of the ETF at strike NAV. Empirically whenever a transaction occurs, we do not know if the transaction was buyer or seller initiated. However when a transaction occurs, we do know that one of the options has non-zero value because it is being exercised.⁹ Therefore, we use the sum of the call and the put as our measure of liquidity because only one of them will have a positive value:

$$Illiquidity = Call(ETF, NAV) + Put(ETF, NAV) \quad (1)$$

where $Call(x, y)$ denotes an American call option with an exercise price of y written on an underlying security with a value of x , and $Put(x, y)$ similarly denotes an American put option with a exercise price of y written on an underlying security with a value of x . Because illiquidity should not depend on the level of the ETF and NAV, we modify our measure above by dividing through by the level of NAV:

$$Illiquidity^* = Call\left(\frac{ETF}{NAV}, 1\right) + Put\left(\frac{ETF}{NAV}, 1\right) \quad (2)$$

⁹Actually, it is being *early* exercised because these are both perpetual American options in the contingent-claims framework.

In order to simplify the interpretation of our liquidity measure, we now use the following transformation to define a new liquidity measure in a way that is analogous to a bond yield:

$$\begin{aligned}
BILLIQ &= -10,000 \times \log \left[\frac{1}{1 + Illiquidity^*} \right] \\
&= -10,000 \times \log \left[\frac{1}{1 + Call\left(\frac{ETF}{NAV}, 1\right) + Put\left(\frac{ETF}{NAV}, 1\right)} \right] \quad (3)
\end{aligned}$$

where $\log[\cdot]$ is the natural logarithm function.

The *BILLIQ* measure will be the bond illiquidity measure that we use throughout this paper. The main advantage of the form of this liquidity measure is that *BILLIQ* can be interpreted as a continuously-compounded rate (quoted in basis points). When the financial markets are perfect and there is zero illiquidity, the value of the call and put are zero and *BILLIQ* takes on a value of 0. As illiquidity increases in the market, *BILLIQ* increases in value and is unbounded from above.

When a buy or sell transaction occurs, the call option or put option is exercised and the other is worthless. Therefore we can write

$$\begin{aligned}
Illiquidity^* &= Call\left(\frac{ETF}{NAV}, 1\right) + Put\left(\frac{ETF}{NAV}, 1\right) \\
&= \max\left[\frac{ETF}{NAV} - 1, 0\right] + \max\left[1 - \frac{ETF}{NAV}, 0\right] \\
&= \left| \frac{ETF}{NAV} - 1 \right| \quad (4)
\end{aligned}$$

Substituting this result into Equation 3, we get the following equation:

$$BILLIQ = -10,000 \times \log \left[\frac{1}{1 + \left| \frac{ETF}{NAV} - 1 \right|} \right] \quad (5)$$

We then re-arrange to derive the formula we use to calculate *BILLIQ* in the paper.

$$BILLIQ = -10,000 \times \log \left[\frac{NAV}{NAV + |ETF - NAV|} \right] \quad (6)$$

As one would expect, the core of our illiquidity measure is the difference between the ETF value and NAV—essentially long the ETF and short the underlying components of the ETF. This measure implies that the greater the difference between the ETF and NAV, the higher the value of this metric, i.e., the greater the illiquidity. When there is a higher disparity between the ETF and portfolio NAV in a market, it means that investors are commanding a larger premium in that market for holding the portfolio of components rather than the NAV. Therefore, there is a large liquidity premium and our illiquidity measure takes on a high value. On the other hand, if the ETF and NAV are identical, then investors are pricing the ETF and its components the same and there is no liquidity premium in that market. In this case, the *BILLIQ* measure takes on a value of zero.¹⁰

Finally, in illiquid markets, it is often the case that pricing is based on quoted prices rather than traded prices. In such a case, if the market experiences a sudden liquidity shock, the ETF usually experiences a quick negative shock, but the underlying portfolio NAV will not experience this shock—essentially, the price adjustment process for the less liquid NAV is slower than that of the more liquid ETF. The same occurs for positive shocks as well. Hence, the absolute difference between ETF and NAV in the denominator of our measure reflects this fact. Because the price-adjustment of the NAV is slower than that of the ETF, the difference in rates of price adjustment is picked up by our formula, irrespective of the direction of the price move—see Amihud and Mendelson (1989) for a formal derivation of the difference in price adjustment processes between an index and underlying securities. We believe that our derivation here is the first to connect the price of immediacy using option’s arguments with the literature on differences in price-adjustment speeds.¹¹

¹⁰In passing, we also note that the *BILLIQ* measure is calculated daily using only NAV and ETF prices for that day. Unlike several other liquidity measures in the literature, *BILLIQ* does not need a time series of data to produce any one day’s illiquidity score. It does not require trading volume data, nor data on bid-ask spreads. This makes it extremely easy to produce for real-world applications.

¹¹Based on Amihud and Mendelson (1989), the autocovariance of the NAV should be higher than that of the ETF because the latter is liquid and has faster price adjustment. We found that in nine of the ten ETFs in our sample (barring CIU), the NAV had a higher return autocovariance than that of the ETF.

2.2 Data

For our empirical testing, we use ten bond ETFs along with an equity ETF.¹² Table 1 provides a description of these ETFs. The main differences among the various bond ETFs are primarily in the durations of the ETFs and their average credit qualities. We try to use ETFs which span different parts of the duration/credit quality matrix such as LQD, which represents the investment grade, medium-duration space, and HYG, which represents the below investment grade, medium-duration space. We have AGG which covers the entire investment grade space and is therefore the most comprehensive ETF. We include ETFs that capture the mortgage-backed securities and emerging-market debt markets, MBB and EMB, respectively. Finally, we utilize an equity ETF, the IVV, which represents the S&P 500, as a way to compare the liquidity measures generated in the fixed-income markets with one generated in the equity markets and see how much of liquidity risk is systemic across markets.

In subsequent sections, we utilize the time series of the *BILLIQ* measured above as calculated with these ETFs and their underlying components. Figure 1, Panel A presents a graph of these *BILLIQ* time series as examples to help readers visualize the liquidity measures that are generated by our procedure. The graphs show that during 2007–2009, the financial markets experienced two liquidity shocks. One occurred in the Summer of 2007, and slowly receded during the remainder of 2007 and the first half of 2008. Then, the second liquidity shock hit, driven by the collapse of Lehman Brothers. It is also interesting to see that the financial market volatility and downturn in 2007–2009 was very different than that of 2001–2003. The 2001–2003 period was not driven by a liquidity shock, while the 2007–2009 period clearly was. In Figure 1, Panel B, we present the *BILLIQ* series for LQD and HYG only, for the period 2007–2009, smoothed by taking a monthly moving average—as expected, the illiquidity of the HYG series is greater than that of LQD. In times of greater market stress, the relative illiquidity of HYG to LQD increases.

In Table 2 we provide descriptive statistics of the ETFs. Because the ETFs were created and launched at different times, we have quite a long time series with some of the ETFs, like the iShares Investment Grade Bond Fund that was launched in July


¹²Note that while our procedure works for any ETF, we chose primarily corporate bond funds due to the illiquidity of that market (relative to equity markets, for example). If one is doing work on illiquidity risk, it is best to utilize markets where that risk is substantial, though not so substantial that no trading occurs in the market.

Table 1: Data Description

This table provides a description of each of the ETFs used in the study to compute bond illiquidity measures. All the ETF time series run up to April 1, 2009.

Ticker	Full Title	Issuer/Industry	
LQD	iShares iBoxx Investment Grade Bond Fund	Corp/Pref-Inv Grade	
HYG	iShares iBoxx High Yield Corporate Bond Fund	Corp/Pref-High Yield	
CSJ	iShares Barclays 1-3 Year Credit Bond Fund	Government/Corporate	
CFT	iShares Barclays Credit Bond Fund	Government/Corporate	
CIU	iShares Barclays Intermediate Credit Bond Fund	Corp/Pref-Inv Grade	
AGG	iShares Barclays Aggregate Bond Fund	Government/Corporate	
GBF	iShares Barclays Government/Credit Bond Fund	Government/Corporate	
GVI	iShares Barclays Intermediate Government/Credit Bond Fund	Government/Corporate	
MBB	iShares Barclays MBS Bond Fund	Asset Backed Securities	
EMB	iShares JP Morgan USD Emerging Markets Bond Fund	Emerging Market-Debt	
IVV	iShares S&P500 Index (NYSE)	Equity	
Ticker	Rating Focus	Maturity Focus	Start Date
LQD	Investment Grade	Intermediate Term (3-10 yr)	7/26/02
HYG	Speculative Grade/High Yield	Intermediate Term (3-10 yr)	4/11/07
CSJ	Investment Grade	Short Term (1-3 yr)	1/11/07
CFT	Investment Grade	Short/Intermediate Term	1/11/07
CIU	Investment Grade	Intermediate Term (3-10 yr)	1/11/07
AGG	Investment Grade	No Restriction	9/26/03
GBF	Investment Grade	Short/Intermediate Term	1/11/07
GVI	Investment Grade	Short/Intermediate Term	1/11/07
MBB	Investment Grade	No Restriction	3/16/07
EMB	Mixed	Intermediate/Long Term	12/19/07
IVV	N/A	N/A	5/15/00

Panel A



Panel B




Figure 1: *BILLIQ* Time Series

This graph presents the time series of the various *BILLIQ* measures extracted from the ETFs. All series are presented in the upper graph. See Table 1 for a description of the different series. The lower graph show the illiquidity series for the investment grade sector (LQD) and high-yield sector (HYG) from May 2007 to March 2009. Illiquidity is expressed in basis points. This is a 30-day moving average of the daily series.

Table 2: Descriptive Statistics of ETF Series

The data is daily. The number of observations is the number of days for which it was possible to compute *BILLIQ* for each ticker. *BILLIQ* is represented in basis points. Price and NAV are in dollars. Volume is the number of units of the ETF that are traded per day. The second panel displays the correlations of returns amongst the ETFs. The lower triangle of the correlation matrix below comprises correlations computed for only the period over which all series had complete observations. In contrast, the upper triangle of the correlation matrix presents the pairwise correlations for all observations where complete data is available for the pair of ETFs. Hence, the upper triangle comprises correlations computed off more observations than the lower triangle. Bold font values are significant at the 5% level.

Ticker	Mean Price	Mean NAV	Mean Volume	Mean BILLIQ	No of Obs
LQD	106.62	106.13	259395	59.1259	1669
HYG	91.77	90.34	215750	198.7806	497
CSJ	100.71	99.80	55293	96.6358	557
CFT	97.67	96.74	13445	109.8060	552
CIU	98.78	97.84	30376	106.0145	553
AGG	100.74	100.43	307023	40.1194	1382
GBF	101.38	100.97	12096	48.3448	546
GVI	102.08	101.61	14064	51.8541	556
MBB	101.74	101.58	65428	18.9880	505
EMB	93.24	92.20	16637	160.3290	321
IVV	118.67				

	LQD	HYG	CSJ	CFT	CIU	AGG	GBF	GVI	MBB	EMB
LQD	1	0.5261	0.4453	0.0380	0.0249	0.9784	-0.0093	0.0305	0.0330	0.0420
HYG	0.5599	1	0.3045	0.0243	0.2337	0.3334	-0.0121	0.0889	-0.0032	0.0617
CSJ	0.4365	0.3194	1	0.0628	0.0294	0.5148	0.0054	0.0396	0.0231	0.0631
CFT	0.0648	0.0462	0.0762	1	-0.0042	0.0572	-0.0104	0.0001	0.0016	0.0062
CIU	0.4307	0.2497	0.4479	0.0608	1	0.0298	0.3786	-0.0016	0.0478	0.0667
AGG	0.4349	0.3506	0.5101	0.0697	0.5379	1	-0.0025	0.0428	0.0442	0.0695
GBF	0.2706	0.0227	0.3615	0.0795	0.5170	0.4558	1	-0.0053	-0.0063	0.0682
GVI	0.3358	0.0990	0.4170	0.0547	0.3365	0.4563	0.5438	1	0.0574	0.0344
MBB	0.2762	-0.0153	0.2470	0.0541	0.4026	0.4337	0.4916	0.4873	1	0.0146
EMB	0.0419	0.0617	0.0631	0.0062	0.0668	0.0697	0.0683	0.0345	0.0147	1

2002, and shorter time series with others, like the iShares Credit Bond Fund, CFT, which was launched in January 2007. For all of the empirical tests we do subsequently, when we need to use multiple data series we use only the time periods for which all of the series overlap. The iShares Aggregate Bond Fund ETF, AGG, has the most number of observations and the largest volume, as expected, because it covers the entire investment grade space. Because AGG has the broadest coverage of all of the fixed income ETFs, we will use the liquidity measure generated from this ETF to make statements about liquidity in the fixed income markets as a whole. The lower panel of the table shows the return correlations between the ETFs. We see that the correlation of most pairs of ETFs is not significant when taken over as long a sample as possible.

Another item to note from Table 2 is that the ETFs represent various levels of trading volume, from an average of \$30 million per day to around \$1 million per day. We chose a diversity of trading volumes specifically to check how trading volume correlated with our liquidity measure. From the mean BILLIQ measure, there appears to be little correlation between our measure and trading volume, which indicates that even though there are frequent large trades occurring in some markets (e.g., HYG), the value of immediacy is still high. Of course, this is simply a levels test, and we will conduct a time series test, controlling for levels, in subsequent tables.

3 Empirical Results

Because we do not have a universally agreed-upon definition of liquidity, one of the difficulties with any kind of liquidity measure is verifying that the measure is indeed displaying volatility due to liquidity exposure as opposed to exposure to other risks. In this section, we conduct a number of empirical tests to provide evidence that the risk measure that we are generating is primarily due to liquidity risk exposure.

In Table 3, we calculate correlations between the liquidity measures calculated from each of the fixed income ETFs. These correlations are calculated over periods where the two ETF return series for which we are calculating a correlation overlap (in the upper triangle of the correlation matrix) or periods where all of the ETF return series overlap (the lower triangle of the correlation matrix). As expected, the correlations in Table 3 reveal that all of the fixed income liquidity measures are highly correlated with each other. This seems to suggest that there is a strong systematic component to

Table 3: Bond Illiquidity Series Correlations

The lower triangle of the correlation matrix below comprises correlations computed for only the period over which all series had complete observations. In contrast, the upper triangle of the correlation matrix presents the pairwise correlations for all observations where complete data is available for the pair of ETFs. Hence, the upper triangle comprises correlations computed off more observations than the lower triangle. Only two correlations, the values in italics, are not significant at the 5% level; all other correlations are significant, mostly at the 1% level. The last line and column of the table shows the correlations of the various liquidity series with the Treasury-Eurodollar (TED) spread. Only two of these correlations are not significant (in italics) at the 5% level.

	LQD	HYG	CSJ	CFT	CIU	AGG	GBF	GVI	MBB	EMB	TED
LQD	1	0.6147	0.7061	0.7069	0.7224	0.4307	0.5859	0.6373	0.5832	0.3787	0.5820
HYG	0.6210	1	0.5989	0.5575	0.6088	0.4393	0.5124	0.6014	0.5618	0.5697	0.3588
CSJ	0.6158	0.6143	1	0.8240	0.8873	0.4129	0.7523	0.8485	0.5697	0.4954	0.3098
CFT	0.6182	0.5428	0.7473	1	0.8333	0.5942	0.7784	0.7993	0.5300	0.4165	0.4108
CIU	0.6325	0.6140	0.8351	0.7525	1	0.3621	0.7508	0.8414	0.5598	0.4415	0.2566
AGG	0.3532	0.4124	0.3332	0.5727	0.2635	1	0.5855	0.4898	0.3545	0.2975	0.4433
GBF	0.4773	0.4843	0.6773	0.7148	0.6691	0.5500	1	0.8052	0.5457	0.4373	0.3358
GVI	0.5300	0.6008	0.7923	0.7281	0.7777	<i>0.4288</i>	0.7527	1	0.6213	0.4889	0.3526
MBB	0.5607	0.5706	0.5326	0.4839	0.5180	<i>0.3152</i>	0.4943	0.5953	1	0.4152	0.4486
EMB	0.3779	0.5709	0.4942	0.4165	0.4406	0.2972	0.4368	0.4882	0.4141	1	0.2018
TED	0.3299	0.2435	<i>0.0969</i>	0.1824	<i>-0.0331</i>	0.4784	0.1555	0.1737	0.4474	0.2020	1

the individual liquidity measures—we will explore this further in subsequent empirical testing. The liquidity measures from the emerging markets bond index, EMB, and the mortgage-backed bond index, MBB, have very low correlations with other liquidity measures. This is probably because these two markets are quite different from the other markets, which are composed entirely of corporate and US government bonds. It is interesting to note that the liquidity measure from the aggregate bond index, AGG, has the lowest correlations with all other indices. This seems unusual, but it may be due to AGG being substantially more liquid, as measured by trading volume (see Table 2), than the other ETFs.

We also examined whether the Treasury–Eurodollar (TED) spread was related to our illiquidity measures. The TED spread is known to widen when there is systematic illiquidity in the interest-rate markets, and we see from Table 3 that there is a statistically significant relationship between the TED spread and most of our measures of liquidity.

Virtually every theoretical model of liquidity predicts that illiquidity should be correlated with volatility, i.e., as volatility increases so should illiquidity—see O’Hara (2004). The first test we run is to calculate the time series of BILLIQ for each of

Table 4: Relationship of *BILLIQ* to the VIX

Results of each ETF *BILLIQ* on the VIX. The regression equation is $BILLIQ = b_0 + b_1 VIX + \epsilon$. We report the coefficients, t-statistics, adjusted R-squares, F-statistics, and the P-values of the F-statistics.

Ticker	b_0	b_1	b_0 tstat	b_1 tstat	Adj R-sq	F-statistic	P-value
LQD	-21.8823	3.9037	-7.97	33.64	40.25%	1131.42	0.0000
HYG	55.6994	4.8875	3.26	9.38	14.89%	87.96	0.0000
CSJ	-37.0638	4.8654	-5.52	22.69	48.02%	514.69	0.0000
CFT	-4.4416	4.1435	-0.72	21.10	44.64%	445.28	0.0000
CIU	-16.5115	4.4389	-2.23	18.81	39.00%	353.85	0.0000
AGG	12.0919	1.4425	6.64	17.96	18.87%	322.67	0.0000
GBF	2.3442	1.6603	0.63	14.11	26.63%	199.21	0.0000
GVI	-10.4570	2.2650	-2.91	19.75	41.21%	390.06	0.0000
MBB	1.5644	0.6052	1.09	13.69	26.67%	187.53	0.0000
EMB	26.2721	3.8556	1.14	6.35	10.95%	40.36	0.0000

our ETFs and then to run a time-series regression of each ETF *BILLIQ* against a volatility index. We use the VIX, the implied volatility of at-the-money S&P 500 options, as our proxy for systemic volatility. Table 4 shows that our *BILLIQ* measure is related to volatility for all of the ETFs. The t-stats on the VIX coefficients are all strongly significant. It is interesting to note that the EMB and AGG indices, which we saw in Table 3 to have the lowest correlations to the other indices, also have low correlations to the VIX. Therefore, one potential explanation is that if we think of the VIX as proxying for systemic liquidity the EMB and AGG indices have the least liquidity effects, i.e., are least likely to be affected by liquidity shocks. It is widely felt in the bond markets and verified in many fixed income papers that bonds of lower credit quality have greater liquidity than higher credit quality bonds—see Alexander, Edwards, and Ferri (2000). This result appears to corroborated by this table as well. Another high-yield ETF, HYG, also has a low correlation with the VIX, similar to EMB and AGG. However it is interesting to note that in the correlations in Table 3, the correlations of HYG with the other ETFs were not as low as EMB and AGG to the others. This may be due to the different sampling period and size of HYG relative to all the other indices.

In Table 5 we look directly at how our liquidity measure performs relative to other liquidity measures created by academics. Specifically, we compare the liquidity measure that is generated by our approach, i.e., the time series of the *BILLIQ* mea-

Table 5: Relationship of *BILLIQ* to Another Bond Illiquidity Measure

Results of regressing *BILLIQ* for each ETF on the reciprocal of the Chacko illiquidity measure. We denote this measure *ILLIQFAC*. The regression equation is $BILLIQ = b_0 + b_1 ILLIQFAC + \epsilon$. We report the coefficients, t-statistics, adjusted R-squareds, F-statistics, and the P-values of the F-statistics. The construction of *ILLIQFAC* is described in Section 3 of the paper. The *ILLIQFAC* series is monthly and so we have regressed the average *BILLIQ* for each month on *ILLIQFAC*.

Ticker	b_0	b_1	b_0 tstat	b_1 tstat	Adj R-sq	F-statistic	P-value
LQD	-703.50	1034.90	-4.89	5.30	26.55%	28.12	0.0000
HYG	-1249.65	1986.23	-2.85	3.32	32.23%	10.99	0.0035
CSJ	-1574.87	2298.91	-9.75	10.41	82.99%	108.34	0.0000
CFT	-1211.66	1815.46	-9.84	10.81	85.29%	116.96	0.0000
CIU	-1453.80	2146.86	-6.59	7.13	70.38%	50.89	0.0000
AGG	-372.95	566.86	-5.96	6.61	41.59%	43.73	0.0000
GBF	-537.83	805.40	-6.40	7.05	70.87%	49.65	0.0000
GVI	-702.23	1036.76	-8.21	8.87	77.91%	78.59	0.0000
MBB	-144.99	224.62	-4.35	4.93	51.48%	24.34	0.0001
EMB	-1279.17	1915.09	-2.58	2.93	38.73%	8.59	0.0137

sure, against the liquidity measure that is generated in Chacko (2009), denoted *ILLIQFAC*. The Chacko measure is based on security holdings rather than trade-based characteristics, and that paper applied its liquidity measure in the corporate bond market to generate a liquidity factor for that market. Our measure has the advantage that it does not use proprietary data like that paper, but of course it is forced to confine itself to the more liquid parts of markets (where we observe trading). In Table 5 we see that despite the fact that the two measures are calculated using very different approaches, they are in fact highly correlated. The correlations are much higher for the measures derived from ETFs that cover the broad credit space, and this is likely due to the fact that *ILLIQFAC* uses the aggregate credit market. As we isolate the market to specific spaces, like below investment grade, such as with EMB, HYG, and AGG, or investment grade only, such as LQD, the correlations to *ILLIQFAC* drop. Nevertheless, in all cases the t-statistic of *ILLIQFAC* is significant (the F-statistics for the regressions are also highly significant) indicating that their measure of liquidity is important in explaining some part of the liquidity measures that we are generating, with the amount explained depending on how closely the ETF we use matches the Chacko (2009) paper in terms of coverage of the US credit markets.

Table 6: The Principal Component and Volatility

This table presents three regressions relating the first principal component of *BILLIQ* to the VIX. A principal components analysis is undertaken using all the ETF *BILLIQ* measures. The data is daily, and there is a common period spanning all the ETFs resulting in 319 dates with data for all ETFs in the sample. The first principal component extracted is denoted PC1. We undertake three regressions here: (1) shows the strong contemporaneous relationship of PC1 and VIX. (2) shows that PC1 is predictable using lagged values of both PC1 and VIX. (3) shows that VIX is predictable using its own lagged value but not that of PC1. Hence regressions (2) and (3) are Granger regressions suggesting that it is volatility that causes illiquidity. The numbers below the coefficient values are t-statistics.

	(1)	(2)	(3)
Dependent Variable	PC1(t)	PC1(t)	VIX(t)
Intercept	-327.1448 -9.08	-53.4317 -2.67	0.9029 1.68
VIX(t)	9.3876 9.942		
PC1(t)			
VIX(t-1)		1.5272 2.853	0.9759 67.85
PC1(t-1)		0.8616 30.99	-0.0002 -0.263
Adjusted R^2	0.2353	0.8137	0.9500
F-stat p -value	0	0	0




Figure 2: Principal Component Analysis

A principal components analysis is undertaken using all the ETF *BILLIQ* measures. The data is daily and there is a common period spanning all the ETFs resulting in 319 dates with data for all ETFs in the sample. This figure presents the percentage of common variation explained by each principal component.

Each of the ETFs generate a different liquidity measure—this is clear from the correlations in Table 3. However, a natural question arises as to whether there are common, or global, components to the individual liquidity measures. The global component that is the “most important” by some metric may be considered as a global liquidity factor for the fixed income markets. Therefore, we perform a principal components analysis (PCA) and rotate the individual liquidity measures generated from each of the ETFs in order to extract a set of orthogonal components driving these individual liquidity measures. Figure 2 shows the percentage of the common variation of the individual liquidity measures is explained by each of the orthogonal components generated by the PCA. This figure shows that almost two-thirds of the variation of the individual liquidity measures is explained by a single orthogonal component. In Table 6, we take this primary orthogonal component and test how it performs in the VIX regressions that we ran in Table 4. From the first regression we see that the VIX is significantly correlated with the primary orthogonal component. The second and third regressions allow us to conduct a Granger causality test. Because lagged values of the primary orthogonal component do not seem to explain the VIX, while lagged values of the VIX do seem to explain the primary orthogonal component, it seems that volatility causes illiquidity—as predicted by the Chacko, Jurek and Stafford (2008) model. Thus, this table lends evidence to the primary orthogonal component being a pure or true liquidity factor. We conduct a similar set of regressions with the Chacko liquidity measure and the primary orthogonal component in Table 7. The primary orthogonal component is significantly and highly correlated to the Chacko measure indicating that both seem to be capturing the same economic factor. The Granger regressions do not reveal that one measure drives the other, indicating that perhaps both are metrics of fundamental liquidity.

Another obvious measure of illiquidity is the transaction cost an investor pays in buying or selling a security—the higher the illiquidity, the higher the transaction cost. This occurs for one of two reasons. First, inventory costs for market makers are higher during illiquid periods due to longer expected round-trip transaction times (and therefore longer expected holding times in inventory for a security). Second, if information asymmetry is a problem, the greater the percentage of better-informed investors in the market when markets are illiquid. Of course, measuring a transaction cost is not possible in an actual transaction because the true, or fundamental, value of the security is never observed (the transaction cost is the difference between the transaction price and the fundamental price of a security). However when bid and ask quotes are observable, we can look at the spread between the bid and ask quotes for a

Table 7: The Principal Component and the Bond Illiquidity Measure

This table presents three regressions relating the first principal component of *BILLIQ* to the *ILLIQFAC* factor. A principal components analysis is undertaken using all the ETF *BILLIQ* measures. The data is daily and there is a common period spanning all the ETFs resulting in 319 dates with data for all ETFs in the sample. The first principal component extracted is denoted PC1. The value of PC1 is then averaged monthly because the *ILLIQFAC* data is monthly - this results in 16 monthly observations. We undertake three regressions here: (1) shows the strong contemporaneous relationship of PC1 and *ILLIQFAC*. (2) shows that PC1 is predictable using lagged values of *ILLIQFAC*. (3) shows that *ILLIQFAC* is predictable using its own lagged value but not that of PC1. The numbers below the coefficient values are t-statistics.

	(1)	(2)	(3)
Dependent Variable	PC1(t)	PC1(t)	ILLIQFAC(t)
Intercept	-3552.0000 -4.54	-3501.9506 -2.82	-0.0477 -0.53
ILLIQFAC(t)	4725.0000 4.54		
PC1(t)			
ILLIQFAC(t-1)		4687.0947 2.83	1.0720 8.89
PC1(t-1)		0.1074 0.41	-0.0003 -1.46
Adjusted R^2	0.5666	0.6341	0.9185
F-stat p -value	0.0005	0.0010	0.0000

Table 8: Relationship of *BILLIQ* to Quoted Spread (QS)

Results of regressing *BILLIQ* for each ETF on the QS. The regression equation is $BILLIQ = b_0 + b_1 QS + \epsilon$. The table contains the coefficients, t-statistics, the adjusted R-squareds, F-statistics and the P-values of the F-statistics.

Ticker	b_0	b_1	b_0 tstat	b_1 tstat	Adj R-sq	F-statistic	P-value
LQD	51.39	267.98	26.47	2.28	0.26%	5.21	0.0225
HYG	178.06	925.70	20.41	5.13	5.37%	26.31	0.0000
CSJ	49.42	17.51	13.03	4.14	3.73%	17.13	0.0000
CFT	62.65	28.05	19.79	8.71	15.63%	75.86	0.0000
CIU	47.58	30.13	13.46	9.03	16.59%	81.54	0.0000
AGG	37.09	105.71	37.34	7.59	4.16%	57.67	0.0000
GBF	27.99	11.83	20.75	7.35	12.06%	54.08	0.0000
GVI	26.76	12.39	24.91	9.18	16.96%	84.31	0.0000
MBB	16.19	6.14	22.41	1.17	0.09%	1.37	0.2432
EMB	125.73	-9.22	7.47	-0.94	-0.08%	0.88	0.3499
Pooled	55.40	19.19	49.83	10.44	1.77%	109.00	0.0000

security as a proxy for the transaction cost: the wider the quoted bid-ask spread, the greater the distance between actual transaction prices and fundamental value is likely to be. In Table 8, we run regressions of the quoted spreads for each of the underlying ETFs against the illiquidity measure generated from each ETF. This table shows that all but two of the illiquidity measures are significantly and positively related to the quoted spread. The only regressions where the relationship between the quoted spread and the *BILLIQ* measures were not significant were for asset-backed securities (MBB) and emerging market debt (EMB). Both of these markets were amongst the hardest hit on a percentage basis by the recent liquidity crisis stemming from Lehman Brothers' bankruptcy. One of the things that happened during the recent liquidity crisis is that conditions became so bad in the markets that dealers quoted artificially wide spreads because they were not willing to do any transactions. Furthermore, these dealers left their quotes unchanged at these wide levels for a long period of time during the crisis, which would explain the lack of correlation. If we had a significant amount of data prior to the crisis we could eliminate the crisis data, but these two series are also among the shorter time series that we have, so that throwing out the crisis period would eliminate more than half of the data.

Finally, we do a check of whether our illiquidity measure detects liquidity risk in financial markets. We use the the US fixed-income markets as our test market. It

is well known in practice that these markets have significant amounts of illiquidity—fewer than 50% of the securities in the credit markets trade in any given year, and less than 1% of the securities are traded every day. Therefore, we speculate that there is considerable liquidity risk in these markets. In Table 9 we regress various indices, representing different sectors of the credit markets, against the bond illiquidity measure from the aggregate bond index, AGG, and several other commonly used risk factors in benchmark performance regressions. In Panel *A* of the table, we use the difference in long-term and short-term Treasury returns to proxy for a risk-less bond market factor, and in Panel *B*, we use pure long-term Treasury returns. The results are similar across specifications.

As one would expect, interest-rate related factors, as proxied by the short-term Treasury return (R_f), long-term Treasury return (Try LT Ret), or the excess long-term Treasury return (Try LT-ST) are important for explaining the returns of high credit quality corporate bonds. However, all of the interest-rate related factors become significantly less important as credit quality decreases. As credit quality decreases, the role of the equity markets ($R_m - R_f$) becomes more important in explaining corporate bond returns—with high yield bonds, the long-term risk-free rate becomes significantly less important than with investment grade bonds; meanwhile equity market performance, while small and statistically insignificant in explaining investment grade bond returns, becomes large in magnitude and statistically significant in explaining high yield bond returns. This is not surprising as all theoretical models of credit risk, beginning with the Merton (1974) model, predict that equities are a proxy for default risk. As one would expect, the decrease in the role of interest rate risk relative to credit risk is muted a bit as we move to longer maturity bonds, where longer durations would imply that interest-rate risk should matter more.

The most interesting result in the table is the role of liquidity risk in explaining bond returns. Our illiquidity measure in Table 9, $BILLIQ_{AGG}$, is statistically significant in all of the regressions for all sectors of high-yield bonds, and also for long-term investment-grade bonds. Its sign indicates that when illiquidity increases, the credit markets produce negative returns—as expected. The fact that the signs of the coefficients do not change—they are all consistently negative when significant—reinforces the notion that this measure is picking up something separate from simple equity or bond market performance (if not the sign would change as we go from investment-grade to high yield, as with the $R_m - R_f$ coefficient). The statistical significance of the $BILLIQ_{AGG}$ coefficients in most corporate bond returns is an important result—it

is an indication that liquidity risk is priced, i.e., that it is a systematic risk factor.

It is also interesting to see that though liquidity risk decreases in importance as default risk decreases, it remains a statistically significant component of the returns of *all long-maturity* corporate bonds. This is likely because many financial institutions such as insurance companies and pension funds that buy credit instruments do so at the long end of maturity structure, for asset-liability matching reasons (in fact, these institutions are the majority of buyers in long-maturity instruments). Furthermore, because most of their liabilities are long-dated, these institutions tend to be buy-and-hold investors—thereby reducing liquidity in long-maturity bonds, even in investment-grade ones. Therefore, we see some evidence that the long investment horizon nature of the participants in certain markets tend to make those markets more susceptible to liquidity risk.

Table 9: Regression of Bond Index Returns on Pricing Factors.

Using the bond illiquidity measure for the aggregate bond market ETF (ticker: AGG), we regressed the Barclays (Lehman) bond index returns for six series on the standard Fama-French factors, and the bond illiquidity series (denoted $BILLIQ_{AGG}$). The factors are the short-term risk free bond return (R_f), the excess performance of the equity market ($R_m - R_f$), the performance of small size firms relative to large size firms (SMB), the performance of high book-to-market ratio firms relative to low book-to-market ones (HML), and a bond market factor, Try LT-ST or Try LT Ret— Treasury long-term minus short-term return in Panel A, or Treasury long-term returns in Panel B. The regression is run on daily returns for the entire period spanned by the AGG ETF. The F-statistic for the regression is significant at the 99% level. The F-statistic for the incremental explanatory power of the $BILLIQ_{AGG}$ series over the Fama-French factors is also significant at the same level. T-statistics are reported below the parameter estimates. We use six bond return series, all pertaining to US Corporates, investment grade and high-yield, further broken down into intermediate maturity and long-term bonds. The significant negative coefficient on $BILLIQ_{AGG}$ indicates that bonds perform poorly in illiquid markets.

PANEL A	Dependent Variables					
	USCorp Inv grade All	USCorp Inv grade Intermediate	USCorp Inv grade Long Term	USCorp High Yield All	USCorp High Yield Intermediate	USCorp High Yield Long Term
Intercept	-0.0202	-0.0234	-0.0072	0.0710	0.0720	0.0630
T-stat	-1.66	-1.97	-0.44	2.99	3.12	1.90
R_f	1.8686	1.8467	1.8068	-0.0786	-0.1123	0.2290
T-stat	2.46	2.49	1.76	-0.05	-0.08	0.11
$R_m - R_f$	0.0028	-0.0051	0.0277	0.0670	0.0640	0.0894
T-stat	0.69	-1.29	5.07	8.52	8.37	8.14
SMB	-0.0002	0.0008	-0.0049	-0.0366	-0.0366	-0.0372
T-stat	-0.02	0.10	-0.43	-2.25	-2.32	-1.64
HML	0.0123	0.0187	-0.0078	0.0349	0.0324	0.0591
T-stat	1.36	2.12	-0.64	1.98	1.89	2.40
$BILLIQ_{AGG}$	-0.0001	0.0001	-0.0006	-0.0016	-0.0016	-0.0015
T-stat	-0.38	0.80	-2.98	-5.86	-6.08	-4.07
Try LT-ST	54.7892	39.4565	101.6129	1.0474	-0.9559	15.9349
T-stat	59.00	43.51	80.72	0.58	-0.54	6.29
Adj R-sq	73.54%	60.50%	83.74%	10.69%	10.85%	10.28%

PANEL B Independent Variables	Dependent Variables					
	USCorp Inv grade	USCorp Inv grade	USCorp Inv grade	USCorp High Yield	USCorp High Yield	USCorp High Yield
	All	Intermediate	Long Term	All	Intermediate	Long Term
Intercept	-0.0160	-0.0218	0.0050	0.0713	0.0721	0.0651
T-stat	-1.41	-1.98	0.30	3.01	3.13	1.97
R_f	1.2426	1.4307	0.5404	-0.0958	-0.1047	0.0261
T-stat	1.75	2.08	0.52	-0.06	-0.07	0.01
$R_m - R_f$	0.0117	0.0025	0.0408	0.0670	0.0638	0.0913
T-stat	3.08	0.68	7.30	8.45	8.27	8.25
SMB	-0.0082	-0.0054	-0.0186	-0.0367	-0.0364	-0.0393
T-stat	-1.06	-0.71	-1.62	-2.26	-2.30	-1.74
HML	0.0092	0.0160	-0.0122	0.0349	0.0324	0.0584
T-stat	1.09	1.96	-0.98	1.98	1.89	2.38
$BILLIQ_{AGG}$	-0.0001	0.0001	-0.0008	-0.0016	-0.0016	-0.0016
T-stat	-1.06	0.54	-4.10	-5.89	-6.09	-4.18
Try LT Ret	49.1168	36.1815	88.6376	0.8165	-0.9353	13.7991
T-stat	64.61	49.02	79.20	0.51	-0.61	6.22
Adj R-sq	76.85%	65.85%	83.22%	10.69%	10.85%	10.23%

As our *BILLIQ* measure captures illiquidity, an interesting question arises as to how liquidity in the bond markets is related to liquidity in other markets. As a test, we calculate a liquidity measure in the equity markets in the exact same way as the *BILLIQ* measure and then compare the equity market liquidity measure against the bond market liquidity measure. We call the equity market liquidity measure *EILLIQ*, and we compute it using the S&P 500 iShares ETF. Table 10 contains regressions of the liquidity measures generated using each of our fixed income ETFs on the liquidity measure generated by the equity ETF. For all but CIU—the intermediate-maturity, investment grade bond ETF—there is a statistically strong relationship between the equity market liquidity measure and bond market liquidity measures. However, the R^2 s for all of the regressions are small indicating that much of liquidity risk is market specific and only a small component of liquidity measures are common across the bond and equity markets. An important implication of this is that while in theory the

Merton (1974) model of credit risk predicts a tight relationship between equities and corporate bonds¹³, market-specific liquidity conditions in each market—particularly the corporate bond market—are one important reason that the empirical relationship is not as tight as one might expect.

In the second panel of Table 10, we conduct vector autoregressions (VARs) using the equity market liquidity measure and the bond market liquidity measure $BILLIQ_{AGG}$. We see from the VAR that there is a strong relationship between lagged values of both $BILLIQ_{AGG}$ and $EILLIQ$ and the current value of $BILLIQ_{AGG}$. This relationship is also present and statistically significant between lagged values of both $BILLIQ_{AGG}$ and $EILLIQ$ and the current value of $EILLIQ$; however, the R^2 for the latter regression is very small, so it appears that the bond market liquidity measure is driven by the equity market liquidity measure and not so much the other way around. The structural VAR confirms this result as the lagged values of the equity market liquidity measure are important in explaining the current bond market liquidity measure, even when we include the current values of the equity market liquidity measure as an independent variable.

4 Liquidity Risk and Financial Institutions

Many financial institutions structure their balance sheets to take on liquidity risk by mismatching their assets with some of their liabilities along the liquidity risk dimension. For example, a bank's liabilities stem from its deposits, any of which have a high degree of liquidity risk—essentially, any deposit can be withdrawn by a depositor at any time. Banks, however, pool deposits together and thereby decrease the liquidity risk of the total pool, a form of liquidity risk diversification. The bank then uses this liability pool to match up against illiquid assets, which in turn pay a substantial liquidity premium.

In this section we test whether financial institutions are indeed exposed to liquidity risk. To the extent that there is a liquidity mismatch between the assets and liabilities of a financial institution, the shareholders of that financial institution ultimately bear

¹³The Merton (1974) model predicts that equities and the credit portion of corporate bonds are simply a call option and a short put option, respectively, on the assets of the issuing firm. Because both the call and put have the same underlying—the enterprise value of the firm—they should be highly correlated.

Table 10: Bond Market Liquidity and Equity Market Liquidity

Results of regressing *BILLIQ* for each ETF on equity market illiquidity as represented by *EILLIQ*. Equity market illiquidity is computed in the same way as *BILLIQ*, but the ETF used was the iShares S&P500 Equity Index (ticker IVV). The regression equation is $BILLIQ = b_0 + b_1 EILLIQ + \epsilon$. We report the coefficients, t-statistics, and adjusted R-squareds. These regression results are shown in Panel A. In Panel B, vector autoregressions are performed using $BILLIQ_{AGG}$ and $EILLIQ$ as endogenous variables. Coefficients are reported, with t-statistics shown below each respective coefficient.

<i>Panel A: Regressions of $BILLIQ$ on $EILLIQ$</i>						
Ticker	b_0	b_1	b_0 tstat	b_1 tstat	R-square	
LQD	49.0544	0.8972	23.05	7.70	3.41%	
HYG	179.2383	1.7730	17.41	3.20	2.03%	
CSJ	87.8586	0.8217	15.95	2.67	1.27%	
CFT	93.7928	1.4939	19.57	5.61	5.41%	
CIU	100.2317	0.5400	17.87	1.73	0.54%	
AGG	33.5958	0.6211	26.78	8.88	5.40%	
GBF	40.6264	0.7199	16.26	5.19	4.71%	
GVI	45.2118	0.6227	16.47	4.07	2.90%	
MBB	16.4400	0.2338	17.63	4.61	3.99%	
EMB	138.1929	1.9926	11.76	3.44	3.57%	

<i>Panel B: Vector Auto-Regression</i>				
Indep Variables	Dependent Variables			
	VAR		Structural VAR	
	$BILLIQ_{AGG}$	$EILLIQ$	$BILLIQ_{AGG}$	$EILLIQ$
Intercept	11.1236 9.73	8.0202 13.63	8.0639	6.899
$BILLIQ_{AGG}$				0.1008
$EILLIQ$			0.3815	
$BILLIQ_{AGG}$ Lag1	0.6791 34.17	0.0403 3.94	0.6637	-0.0281
$EILLIQ$ Lag1	0.1643 3.09	0.0826 3.02	0.1327	0.0660
Adj R-square	43.85%	2.12%		
F-statistic	647.90	15.94		
F-stat P-value	0.0000	0.0000		

that risk first and foremost. We therefore regressed the returns of a portfolio of equities of various financial institutions on our liquidity measures. For our portfolio of financial institution equities, we used the Financial Select Sector SPDR Fund (XLF). This ETF provides investment results that correspond to the price and yield performance of the Financial Select Sector of the S&P 500 index. The index includes a broad range of financial institutions including commercial and investment banks as well as insurance companies.¹⁴

We regressed the returns of XLF on each of our ten *BILLIQ* measures created from the ten bond ETFs. Table 11 contains the results of these regressions. The coefficients on our liquidity measures are all negative, as one would expect. Financial institutions are long illiquid assets and short liquid ones; therefore, when illiquidity increases, the values of illiquid assets drop more than the values of liquid ones. As a result, the assets side of the balance sheets of these institutions decrease more than the liability side. This results in a net loss, and the institutions' share prices drop. The negative coefficients in our regressions confirm that this is indeed happening. Additionally our liquidity measures are statistically significant in eight out of the ten regressions. The regression against AGG, our liquidity measure from the aggregate bond market—and therefore our closest candidate to a systematic liquidity risk factor, produced the strongest results. The two regressions which had the least explanatory power and no statistical significance were against HYG and EMB, liquidity measures from the high yield and emerging market sectors, respectively. As financial institutions generally hold very few of instruments from these sectors, they have very little exposure to liquidity movements in these markets, and so it is not surprising that these came out insignificant. Overall, the results of Table 11 demonstrate that financial institutions do indeed have considerable exposure to liquidity risk, just as we theorized.

Hedge funds are another type of financial institution that earn a major portion of their returns from taking liquidity risk. For example, convertible arbitrage and capital structure arbitrage involve taking a long position in corporate bonds and short positions in equities. As a result, the asset and liability sides of a convertible arbitrage fund's or a capital structure arbitrage fund's balance sheet is heavily mismatched on liquidity—the corporate bonds are orders of magnitude more illiquid than equities.¹⁵ Thus the fund bears significant liquidity risk, which is borne by the fund's investors,

¹⁴For example, its top 5 holdings at the end of March, 2010 were, in order, JP Morgan Chase, Bank of America, Wells Fargo, Citigroup, and US Bancorp.

¹⁵See, for example, Edwards, Harris and Piwowar (2007) and Chacko (2009).

Table 11: Relationship between *BILLIQ* and Financial Institution Index (XLF)

Results of regressing XLF, an index comprised of financial institutions, on *BILLIQ* for each ETF. The regression equation is $Return_{XLF} = b_0 + b_1 BILLIQ + \epsilon$. The table contains the coefficients, t-statistics, the adjusted R-squareds, F-statistics and the P-values of the F-statistics.

Ticker	b_0	b_1	b_0 tstat	b_1 tstat	Adj R-sq	F-statistic	P-value
LQD	0.0242	-0.00058	2.49	-4.91	23.38%	24.10	0.0000
HYG	-0.0280	-0.00012	-0.76	-0.77	2.62%	0.59	0.4495
CSJ	-0.0092	-0.00039	-0.38	-2.23	16.57%	4.97	0.0351
CFT	0.0154	-0.00058	0.54	-2.66	22.08%	7.08	0.0134
CIU	-0.0101	-0.00035	-0.39	-1.97	13.47%	3.89	0.0597
AGG	0.0472	-0.00148	3.26	-4.83	26.41%	23.33	0.0000
GBF	0.0151	-0.00131	0.57	-2.96	26.02%	8.79	0.0066
GVI	-0.0051	-0.00081	-0.20	-2.12	15.20%	4.48	0.0444
MBB	0.0084	-0.00311	0.25	-2.09	15.99%	4.38	0.0476
EMB	-0.0119	-0.00037	-0.23	-1.36	11.63%	1.84	0.1961

and a portion of the fund’s return is fair compensation for bearing this risk. As another example, consider long-short equity funds. Most such funds are careful to have significantly more liquidity in their short positions than their long positions. This is because a short position can be called at any time, or there may be a sudden short squeeze in the markets, requiring the fund to liquidate a short position quickly. The risk of this need for transaction immediacy in turn compels funds to keep very liquid short positions - much more liquid, on average, than the liquidity of the median stock in the equity markets. As a result, the long positions of long-short equity funds are on average less liquid than the short position of these funds. As with convertible arbitrage, this liquidity mismatch in the balance sheet results in liquidity risk, which is borne by the equity capital, i.e, the fund’s investors, and at least some portion of the fund’s return is compensation for bearing this risk.

In fact, this liquidity mismatch is present in all funds that take long and short positions. Unless the long and short positions are carefully constructed so as to equate the liquidity of the long positions with the liquidity of the short positions—something very few funds do¹⁶—there will inevitably be a liquidity mismatch between the two. Therefore, virtually all funds taking long and short positions will have liquidity risk in their balance sheets, and a portion of all such funds’ returns is simply fair com-

¹⁶One reason for this is the lack of good liquidity risk measures.

Table 12: Hedge Fund Returns and Liquidity Risk

Results of regressing different hedge fund strategies, represented by the CS Tremont Hedge Fund sub-indices, against the $BILLIQ_{AGG}$ measure, which captures aggregate bond market liquidity. The specific regression equations are of the form $Return_{Strategy} = b_0 + b_1 BILLIQ_{AGG} + \epsilon$. The table contains the coefficients, t-statistics, adjusted R-squareds, F-statistics, and the P-values of the F-statistics.

Strategy	b_0	b_1	b_0 tstat	b_1 tstat	R^2	F-stat	P-value
Convertible Arbitrage	0.02	-0.00064	2.63	-3.70	17.39%	13.68	0.0004
Dedicated Short Bias	-0.01	0.00022	-1.04	1.01	1.55%	1.02	0.3160
Emerging Markets	0.04	-0.00091	4.66	-4.99	27.69%	24.89	0.0000
Equity Market Neutral	0.02	-0.00057	3.15	-4.24	21.64%	17.95	0.0001
Event Driven	0.01	-0.00023	2.86	-2.44	8.53%	5.97	0.0173
Fixed Income Arbitrage	0.03	-0.00090	5.30	-7.22	44.53%	52.18	0.0000
Global Macro	0.02	-0.00043	2.92	-3.36	14.98%	11.27	0.0013
Long Short Equity	0.02	-0.00045	3.37	-3.58	16.45%	12.80	0.0007
Managed Futures	-0.01	0.00027	-0.84	1.75	4.49%	3.06	0.0852
Multi-Strategy	0.02	-0.00056	4.41	-5.25	29.76%	27.53	0.0000
Tremont Blue Chip Index	0.02	-0.00047	4.24	-4.98	27.58%	24.76	0.0000

pensation for bearing this risk.

To test this idea, we regressed the returns from eleven hedge fund strategies¹⁷ on the wide bond market illiquidity measure, i.e., $BILLIQ_{AGG}$, to see whether hedge fund returns reacted adversely to liquidity shocks. The results are reported in Table 12. Nine out of the eleven hedge fund strategy returns show significant negative reactions to increases in illiquidity. The two types of hedge funds that do not evidence sensitivity to liquidity are dedicated short bias funds and managed futures funds. Dedicated short bias funds do not hold long positions; therefore, there can be no liquidity mismatch between the long and short positions on their balance sheets. As a result, the lack of significance in the liquidity coefficient in this strategy is very much expected, and further confirms that our liquidity measure is functioning as it should. The managed futures strategy generally confines itself to the very liquid parts of financial markets for both its long and short positions, and therefore does not result in significant liquidity mismatching between its long and short positions. As with the dedicated short strategy, the insignificant liquidity coefficient on this strategy is expected and further corroborates our liquidity measure.

¹⁷We chose hedge funds because most hedge funds hold both long and short positions.

In conclusion, our liquidity measures indicate that several large classes of financial institutions owe at least part of their returns to liquidity risk exposure. They mismatch liquidity on the asset and liability sides of their balance sheets, which creates the liquidity exposure, and they seem to earn a liquidity premium from doing so.

5 Conclusion

Most papers that have generated and analyzed liquidity measures suffer from the problem that they are not able to fully filter out non-liquidity market risk factors. In fact, the signal-to-noise ratio for most liquidity measures are probably quite low, given the low correlations amongst the various liquidity factors generated by different papers.

In this paper, we created a liquidity measure in a way that tries to deal as effectively as possible with the noise created by non-liquidity factors. Our methodology involved creating a portfolio by going long and short the exact same set of securities. We created a liquidity measure by going long a traded index of securities and simultaneously going short the same individual securities that make up that index. Because we were going long and short the exact same set of securities, almost all market risk factors had effectively been hedged away from the time series of the portfolio's equity value. We therefore interpreted the portfolio's equity value, which is the price difference between the very liquid index and the relatively illiquid components of the index, to have been a pure liquidity measure—it was essentially a measure of the liquidity gap between the index and its underlying components.

We performed a non-linear transformation of this long-short portfolio to create a liquidity measure that could be interpreted in a manner similar to an interest rate. We then applied this measure to a number of bond ETFs to create liquidity measures for various segments of the credit markets. While it is difficult to determine whether a risk factor truly represents liquidity risk, we conducted a number of checks. The liquidity measures were strongly related to volatility—as volatility increased, so did illiquidity—as most theoretical models and practitioners would expect. The liquidity measures were also strongly related to a liquidity measure generated from another paper covering the fixed-income space using a different metric for calculating liquidity (a sort of out-of-sample test). We also found that the liquidity measures were strongly related to bid-ask spread quotes, another commonly used measure of liquidity. The

BILLIQ measures may be computed using only NAV and ETF prices on a given day, not bid-ask spreads or trading volume, and are therefore easily generated for practical use.

Finally, we used our liquidity risk measures to test whether financial institutions are exposed to liquidity risk and found compelling evidence that the balance sheets of financial institutions are indeed significantly exposed to liquidity mismatch through the liquidity mismatch between assets and liabilities. Liquidity risk is an important risk factor in financial institutions and needs to be assessed carefully in order to get an accurate measure of the risk/health of any financial institution. Our liquidity measure offers one way to quantify this liquidity exposure.

In conclusion, we have found that the formula derived in this paper appears to have substantial merit for measuring liquidity and generating a liquidity measure. There are close to a thousand ETFs covering many markets, and, as a consequence, this approach may be applied very widely.

References

- Acharya, V., and L. Pedersen (2005): "Asset Pricing with Liquidity Risk," *Journal of Financial Economics*, 77, 375-410.
- Alexander, G., A. Edwards, and M. Ferri (2000): "The Determinants of Trading Volume of High-Yield Corporate Bonds," *Journal of Financial Markets* v3, 177-204.
- Amihud, Y. and H. Mendelson (1986): "Asset Pricing and the Bid-Ask Spread," *Journal of Financial Economics*, v17, 223-249.
- Amihud, Y. and H. Mendelson (1989): "Index and Index-Futures Returns," *Journal of Accounting, Auditing and Finance* 4, 415-431.
- Brennan, M., T. Chordia, and A. Subrahmanyam (1998): "Alternative Factor Specifications, Security Characteristics and the Cross-Section of Expected Stock Returns," *Journal of Financial Economics*, v49, 345-373.
- Brennan, M. and A. Subrahmanyam (1996): "Market Microstructure and Asset Pricing: On the Compensation for Illiquidity in Stock Returns," *Journal of Financial Economics*, v41, 441-464.

- Campbell, J. and G. Taksler (2003): "Equity Volatility and Corporate Bond Yields," *Journal of Finance*, v58, 2321-2349.
- Chacko, G. (2009): "Liquidity Risk in Corporate Bond Markets," Working Paper, Santa Clara University.
- Chacko, G., J. Jurek, and E. Stafford (2008): "The Price of Immediacy," *Journal of Finance*, 63(3), 1253-1290.
- Chen, L., D. Lesmond, and J. Wei (2007): "Corporate Yield Spreads and Bond Liquidity," *Journal of Finance*, v62, 119-149.
- Chordia, T., R. Roll, and A. Subrahmanyam (2002): "Commonality in Liquidity," *Journal of Financial Economics*, v56, 3-28.
- Chordia, T., A. Subrahmanyam, V. Anshuman (2001): "Trading Activity and Expected Stock Returns," *Journal of Financial Economics*, v59, 3-32.
- Copeland, T., and D. Galai (1983): "Information Effects on the Bid-Ask Spread," *Journal of Finance*, 38(5), 1453-69.
- Das, S., and P. Hanouna (2009). "Hedging Credit: Equity Liquidity Matters," *Journal of Financial Intermediation* 18(1), 112-123.
- Das, S., and P. Hanouna (2010): "Run Lengths and Liquidity," *Annals of Operations Research*, 176(1), 127-152.
- Datar, V., N. Naik, R. Radcliffe (1998): "Liquidity and Stock Returns: An Alternative Test," *Journal of Financial Markets*, v1, 203-219.
- Edwards, A., L. Harris, and M. Piwowar (2007): "Corporate Bond Market Transparency and Transaction Costs," *Journal of Finance*, v62, 1421-1451.
- Elton, E., M. Gruber, D. Agrawal, and D. Mann (2001): "Explaining the Rate Spread on Corporate Bonds," *Journal of Finance*, v56, 247-277.
- Ericsson, J. and O. Renault (2002): "Liquidity and Credit Risk," Working Paper, McGill University.
- Garman, M. (1976): "Market Microstructure," *Journal of Financial Economics*, v3, 257-275.

- Goodhart, C. and M. O'Hara (1997): "High Frequency Data in Financial Markets: Issues and Implications," *Journal of Empirical Finance*, v4, 73-114.
- Harris, L. (2003): "Trading and Exchanges: Market Microstructure for Practitioners," Oxford University Press.
- Hasbrouck, J. and D. Seppi (2001): "Common Factors in Prices, Order Flows, and Liquidity," *Journal of Financial Economics*, v59, 383-411.
- Ho, T. and H. Stoll (1981): "Optimal Dealer Pricing Under Transactions and Return Uncertainty," *Journal of Financial Economics*, v9, 47-73.
- Huang, J. and M. Huang (2003): "How Much of the Corporate-Treasury Yield Spread is Due to Credit Risk," Working Paper, Stanford University.
- Huberman, G. and D. Halka (2001): "Systematic Liquidity," *Journal of Financial Research*, v24, 161-178.
- Korajczyk, R., and R. Sadka (2008): "Pricing the Commonality across Alternative Measures of Liquidity," *Journal of Financial Economics*, 87(1), 45-72.
- Li, J., R. Mooradian, and W. Zhang (2007): "Is Illiquidity a Risk Factor? A Critical Look at Commission Costs," *Financial Analysts Journal*, 63(4) (July/August), 28-39.
- Longstaff, F. S. Mithal, and E. Neis (2005): "Corporate Yield Spreads: Default Risk of Liquidity? New Evidence from the Credit-Default Swap Market," *Journal of Finance*, v60, 2213-2253.
- Merton, R. (1974): "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, v29.
- Merton, R. (1993): "Optimal Investment Strategies for University Endowment Funds," in *Studies of Supply and Demand in Higher Education*, edited by C. Clotfelter and M. Rothschild. University of Chicago Press, Chicago.
- O'hara, M. (2004): *Market Microstructure Theory*, Blackwell Publishing, Oxford, UK.
- Pastor, L. and R. Stambaugh (2003): "Liquidity Risk and Expected Stock Returns," *Journal of Political Economy*, 642-685.

- Sadka, R. (2006): "Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk," *Journal of Financial Economics*, 80, 309-349.
- Stoll, H. (1978): "The Supply of Dealer Services in Securities Markets," *Journal of Finance*, v33, 1133-1151.