Macroeconomic Patterns in System-Wide Liquidity Regimes

Mark D. Flood
Office of Financial Research
Washington, DC

John C. Liechty
Pennsylvania State U.
State College, PA, and
Office of Financial Research
Washington, DC

Thomas Piontek
Office of Financial Research
Washington, DC

November 2, 2014

VERY PRELIMINARY. PLEASE DO NOT QUOTE WITHOUT PERMISSION.

The authors thank Tobias Adrian, Greg Feldberg, Pete Kyle, and workshop participants at Goethe University, Massachusetts Institute of Technology, the Federal Reserve Bank of New York, and the Open Financial Data Group, for numerous helpful comments and suggestions. Any remaining errors are the responsibility of the authors alone.

Abstract

This research seeks statistical commonalities in market liquidity across a range of corporate equities and bond markets and commodity futures markets, with an aim to better understanding system-level patterns in aggregate or funding liquidity that arise from broad patterns in market liquidity. We present a Bayesian estimation of hidden Markov chain (HMC) models to measure the latent structure of liquidity. Starting with granular measures of market liquidity for equity, bond, and futures markets, we use Markov chain Monte Carlo (MCMC) analysis to estimate the latent structure governing liquidity at the system-wide level. Our input measures of market liquidity take advantage of recent work by Kyle and Obizhaeva [2014], who propose "invariant" price-impact measures of responses to order flow that are readily comparable across a broad range of financial markets and conditions. We find that three latent liquidity regimes—corresponding to high, medium, and low price-impact—are adequate to describe all of the markets we consider. In a more focused examination of the equities markets alone, we test the ability of several macroeconomic time series to recover the estimated liquidity dynamics. This exercise has significant explanatory power and allows for an economically meaningful attribution of the estimated latent liquidity states.

1 Introduction

We present a new approach to the study of liquidity that exploits daily data to identify broad-based and relatively high-frequency patterns for individual markets to help explain aggregate, system-wide liquidity conditions. Although "funding liquidity" in the wholesale markets for institutional funding is the most immediate concern for system-wide conditions, there are vastly many more individual markets for financial assets than for intermediaries' liabilities. It is an empirical question whether there is additional information in the asset markets that might help to explain liquidity conditions in the funding markets. Previous studies of "commonalities in liquidity" (e.g., Chordia et al. [2000] and Karolyi et al. [2012]) find that there are indeed significant patterns in the detailed data. We extend the commonalities approach by analyzing a range of asset classes, including equities, bonds, and financial futures.[1] We also extend the commonalities approach with a novel methodology for connecting aggregate liquidity patterns to a panel of macroeconomic time series.

To extract aggregate patterns from the detailed data, we apply Bayesian estimation of hidden Markov chain (HMC) models in order to measure the structure of liquidity in the financial system. Starting with granular measures of market liquidity, we use Markov chain Monte Carlo (MCMC) methods to estimate the latent structure of liquidity at the aggregate, system-wide level. Our primary input measures of market liquidity come from applying recent work by Kyle and Obizhaeva [2014], who propose an "invariant" price-impact measure of responses to order flow that are readily comparable across a broad range of financial markets and conditions. We begin by exploring daily market data from a range of financial markets to create a collection of daily price impact measures. In our initial implementation, we consider volatility index futures, oil futures, and various

[1]Prior studies of commonality have focused on equities markets alone. An exception is the recent working paper by Marra [2013] which pairs individual equities with their matching credit default swaps (CDS); her emphasis, however, is on firm-level interactions between the securities rather than system-wide liquidity.

portfolios of equities and corporate bonds. We assume that the dynamics of each of these daily price impact measures (33 in all) are determined by an underlying variable that alternates between one of several liquidity states and that switches between these states drives sudden changes in the observed levels of price impact. These underlying states are "latent"—i.e., not directly observable—but inferred from the dynamics of daily price impact measurements. In our initial analysis, although we estimate each price impact series independently – assuming no coordination between the dynamics of the latent liquidity states across markets – we find surprising consistency in the dynamics of market liquidity across all of these markets. First, we find that just three liquidity regimes are adequate to describe each market: high, intermediate, and low. Second, we find that the low liquidity regime afflicts all of the markets during the financial crisis of 2008.

As a point of comparison, we also look at the dynamics of two more traditional measures of liquidity, the daily average bid-ask spread and daily turnover, for a limited subset of these markets. While these alternative measures capture general patterns in liquidity dynamics similar to the patterns from the price-impact analysis, they are subject to clear structural changes which reflect the broad changes in the market intermediation infrastructure.

Returning to the price-impact analysis, despite these common features in the price-impact liquidity dynamics, we also find interesting differences across the various markets in the lead-up to the recent crisis and in its aftermath. We build on our initial findings and on earlier studies of liquidity commonalities by using a model that links a collection of latent liquidity states (from multiple markets) together using a multinomial Probit model that is driven by a collection of macro variables. This provides a framework that allows us to assess the usefulness of macro variables as potential indicators of financial stress (at least in terms of market liquidity). Although we restrict our focus in this phase of the analysis to liquidity in the U.S. equities markets for reasons of data consistency, the model reveals that a number of macro variables, such as the Dow Jones Real Estate Index, the

Treasury-Eurodollar (TED) spread, the VIX® and the S&P 500 price-to-book (P/B) ratio are statistically significant in explaining when these markets experience different levels of liquidity.

The remainder of the paper is structured as follows. First, in the remainder of Section 1 we discuss the general issue of liquidity measurement, and provide a rationale for our approach. In Section 2, we describe the models and sampling strategies used in the MCMC analysis for both the univariate models (i.e., one market at a time) and the hierarchical model (multiple markets at a time). Section 3 describes the data and specific formulas for measuring price impact. Section 4 reports our findings from aggregating both the market specific analyses and the multiple market analysis. We conclude with a discussion of potential future work, limitations of the current approach, and ways that these limitations might be overcome to allow these tools to be used to expand our ability to forecast liquidity.

1.1 The Challenges of Liquidity Measurement

Conceptually, "liquidity" is the ease with which participants in the financial system can convert their claims to cash. The settlement obligations of the vast majority of financial contracts are stated in terms of cash: payors must deliver cash, and payees must accept it. As a result, the ability of market participants to "get to cash" has important implications for the overall functioning of the system. Liquidity is particularly important for the study of financial stability, as sudden shifts in liquidity have historically been one of the defining characteristics of financial crises.[2]

For a financial asset, liquidity arises because there are agents who stand ready to

[2] Kindleberger [1993, ch. 15], for example, recounts the history of crises in Western economies, with a special focus on the lender of last resort as a provider of backstop liquidity to the system. Reinhart and Rogoff [2009], and Schularick and Taylor [2012] consider a similar sample, with the latter focusing on the role of "credit booms gone wrong" as a precursor to crisis events. Because aggregate credit growth is typically facilitated by expanding bank balance sheets, there is an important empirical connection linking credit cycles, leverage cycles, and liquidity cycles.

purchase the instrument. In efficient markets, arbitrage implies that there is money to be made by those with the robust valuation models and accurate information to feed them, so that counterparties should be easy to find for a modest price concession. That is, liquidity providers will step in opportunistically if the asset is offered at an acceptable discount, but have little incentive to pay more than the current market price. The practical upshot for liquidity measurement is that we measure the extent of illiquidity in the system as one-sided deviations from an ideal benchmark of "perfect" liquidity.[3]

Ideally, measures of liquidity to support financial stability monitoring would be both timely (available at high frequency to track developments in near real time) and forward-looking (possessing some forecasting power to serve as an early warning signal). These goals are often defeated in practice by certain fundamental challenges. In particular, liquidity exhibits three interrelated characteristics that present special complications to measurement: latency, nonlinearity, and endogeneity. Each of these challenges has ramifications for both funding liquidity and market liquidity.

1.1.1 Latency

Latency means that much of the most interesting liquidity behavior is ex-ante unobservable. At the microstructural level, we typically most wish to know not merely the prices of recent trades or the current best bid and offer, but how deep or resilient the market will be. Many trading mechanisms designed to attract liquidity providers (i.e., buyers and sellers) do so by restricting information availability, for example, with closed limit-order books, hidden or "iceberg" orders, anonymous brokerage to conceal trader identities, and limited-access

[3] It is important to consider what constitutes a "reasonable" price in this context. Most financial instruments have limited liability, so there is typically some nonnegative price at which buyers come forward; however, a market is not liquid if buyers only appear for fire-sale offers. The techniques of financial engineering mean that plausible mark-to-model valuations are usually available even for contracts that trade infrequently. The recent Presidential Address to the Econometric Society by Holmström [2012] emphasizes the lengths markets will endure to achieve liquidity by forcing assets to be "informationally insensitive." See also Dang et al. [2012].

"upstairs" trading venues for large trades. Moreover, traders are never under compulsion to reveal their intentions by actually issuing (or canceling) a limit order prior to the moment of truth; some version of a market order is typically available. Conversely, private visibility into their own customers' positions and order flow can be a valuable information source for dealers.

To hedge against these and other liquidity surprises, firms frequently arrange for contingent liquidity in the form of lines of credit or derivative contracts. However, hedgers must still worry about the wrong-way risk that their supposed guarantors will themselves fail under the precisely the event being insured against. Liquidity is never purely localized to a single transaction or market. At the broadest levels, global liquidity depends in part on the responses of firms (and policy makers) to situations they may have not have planned for explicitly. For example, it is difficult to know in advance whether an initial deleveraging event will generate enough selling pressure to create a fire-sale feedback loop. Geanakoplos [2003], for example, emphasizes that collateral margins are likely to bind in a crisis, unexpectedly depriving participants of flexibility at the crucial moment. Similarly, the repeated clearing crises of the nineteenth century (see Calomiris and Gorton [1991]) were invisible to bankers in the system until it was too late. In general, liquidity measurement is "more honored in the breach," in the sense that it is easier to assess the depth or resilience of the market when conditions are stressful enough to violate the "perfectly liquid" ideal.

We address the challenge of latency by adopting a MCMC approach to estimate the latent regime structure governing the observed price impact series. The maintained assumption is that, while the markets' liquidity behavior is indeed largely latent, these hidden patterns will reveal themselves in a broad cross-section of markets observed at relatively high-frequency (daily). In the results below, we are indeed able to identify three meaningful latent liquidity states (high, medium, and low price-impact) that seem to govern the observed liquidity behavior.

1.1.2 Nonlinearity

Nonlinearity in the response of liquidity to significant market changes compounds the problem of unobserved behavior. Numerous studies have documented the empirical regularity that price response to order flow tends to be concave function of the transaction size. Intuitively, order flow can move the price significantly before additional liquidity providers arrive to dampen the effect.[4] For example, limit orders (and other contingent liquidity) may crowd behind the best posted quote, so that an order flow impulse large enough to work through this initial phalanx may expose gaps in the book, provoking an abrupt shift in prices. Such unevenness in market depth may be an important source of fat-tailed returns distributions. Duffie [2010] suggests that for some relatively illiquid markets, it may be weeks before support arrives in the form of additional order flow. Nonlinearity is a challenge for liquidity measurement because it hampers our ability to extrapolate from small-scale, localized effects to the larger, out-of-sample effects that are often of greatest concern.

Nonlinearity in liquidity is an even greater worry for systemic stability, where the stakes are correspondingly higher. At this level, interactions among nodes in the system can conspire to produce self-amplifying feedback loops. Tirole [2011] provides a tour of the catalog of systemic pathologies related to illiquidity, including contagion, fire sales, and market freezes. He underscores the central fact that one of the basic services provided by the banking (and shadow banking) sector—namely maturity transformation—render it especially vulnerable to runs and other liquidity surprises. Brunnermeier [2009] provides a good overview of how these forces played out in practice, at least through the early (and most severe) phases of the recent crisis. He highlights four specific channels: (a) delever-

[4] Much of the literature identifies a "square-root rule" that posits price impact to be proportional to the square root of the transaction size; see for example, Gabaix et al. [2006], Bouchaud et al. [2008], Hasbrouck and Seppi [2001], and Toth et al. [2011]. In contrast, recent work by Kyle and Obizhaeva [2011a,b] argues for a cube-root rule, and provides a theoretical justification for why this regularity should emerge.

aging spirals driven by erosion in capital levels and increases in lending standards and margins; (b) a credit crunch motivated by funding concerns; (c) outright runs, exemplified by Bear Stearns and Lehman Brothers; and (d) netting failures due to real or perceived counterparty credit risks. All of these modalities involve liquidity. Adrian and Shin [2010] emphasize the role played by institutional leverage in both the expansion and contraction of the system. Figure 1, adapted from similar illustrations in Adrian et al. [2013a], illustrates clearly that institutional leverage expands and contracts via adjustments to assets and liabilities—not equity—suggesting that increases in leverage correspond to increases in overall liquidity, since bank deposits and other liabilities are a key component of liquid assets for other participants in the system. Leverage is strongly procyclical; by increasing global liquidity, aggregate balance sheet expansion encourages investment spending and tends to ease margin constraints. Figure 1 also indicates a correlation between bank leverage and market volatility, at least through the course of the most recent business cycle: periods of low volatility (green and yellow markers) tend to correspond to increases in bank leverage, and episodes of high volatility (orange and red) tend to match decreases in leverage. Unfortunately, the volatility-leverage-liquidity spiral works in reverse as well, leading to debt overhangs as the system contracts, with associated increases in institutional risk aversion and liquidity hoarding.

We address the challenge of non-linearity by agnostically allowing the data determine the "correct" number of liquidity states for each time series. Notably, for all 33 of our univariate series, three liquidity states are adequate to explain the observed variation in the price-impact statistics. Because the expected price impact is allowed to vary idiosyncratically for each of the three regimes (high, medium, and low price-impact), this model naturally captures non-linearities in price impact.

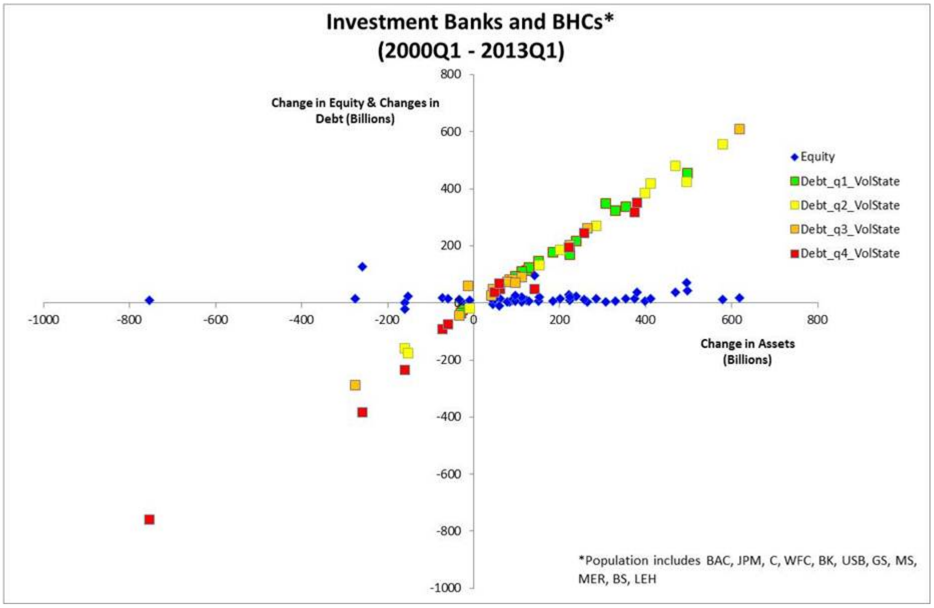

Figure 1: The Leverage-Liquidity CycleSources: Bloomberg, OFR Analysis

1.1.3 Endogeneity

Endogeneity means that liquidity is partly a network effect that emerges through the interactions of many market participants. Thus, active markets should have less need for identified liquidity providers who convert investments to cash by purchasing or rediscounting others' financial assets. Endogenous liquidity creates a straightforward network externality in the sense of Pagano [1989] and Economides [1996]: investors are naturally more willing to enter markets where there are already many other traders and large transaction volumes, because this provides an implicit assurance that counterparties will be easy to find when needed. A familiar example of this phenomenon is the contrast between trading for on-the-run and off-the-run Treasuries (see Barclay et al. [2006]). Similarly, Bessembinder et al. [2006] find that liquidity externalities are consistent with the significantly reduced trade execution costs that followed the introduction of the TRACE feed, which increased transparency in the corporate bond market. Liquidity externalities are also often touted as a benefit of high-frequency trading.

Liquidity externalities operate at the level of the system as well.[5] They have long been a central concern of financial stability supervisors. Elliott et al. [2013], for example, document this multi-dimensional regulatory history for the U.S. One of the core functions of central banking is to provide a lender of last resort in a crisis, an idea first flirted with as an expedient in London's Panic of 1825, "codified" in Bagehot [1873], and institutionalized in the U.S. with the creation of the Federal Reserve in 1913. Recourse to a potentially unlimited source of liquid cash from outside the network of banks and financial firms is important in a crisis, when information asymmetries and other constraints prevent firms from liquidating their assets to meet withdrawals. On the other hand, some have argued that this explicit promise of effectively unlimited contingent liquidity creates a moral hazard—that too-big-to-fail banks undertake excessive leverage and maturity transformation, comfortable that the Fed's emergency backstop provides them with a free liquidity put.[6] Moreover, in spite of discount-window access, Cornett et al. [2011] find that banks in the recent crisis were forced onto a more defensive liquidity posture, in part because Lehman Brothers' failure diverted commercial paper borrowers to draw on banks' liquid assets via backup lines of credit, and partly because wholesale funding sources suddenly shrank. A net result was a restriction in commercial lending. In the wake of the crisis, macroprudential supervisors have focused renewed attention on liquidity buffers, including the new Basel III requirements for banks' to maintain net stable funding and liquidity coverage ratios.[7]

We address endogeneity by estimating a hierarchical model that searches for common

[5] There are many discussion of this endogenous systemic externality. See, for example, Morris and Shin [2004], Dang et al. [2010], and Adrian and Shin [2010] and the references therein.

[6] Goodhart [2008] and Farhi and Tirole [2009] have made this argument. The term "liquidity put" is a metaphor for a commonly used recourse covenant that allows investors in a partially debt-funded structured investment vehicle (SIV) to put back their shares in the SIV to the sponsoring bank if the SIV is unable to roll over its short-term debt; see, for example, Entwistle and Beemer [2008]. However, during at least one episode—the Y2K millennium date change—the Fed literally sold liquidity put options; see Sundaresan and Wang [2006].

[7] The new liquidity framework for banks is discussed in Basel Committee on Banking Supervision [2013, 2010], Adrian et al. [2013b], and Bank of England [2011].

liquidity structure throughout the cross-section of observed price-impact series. Although this work is in preliminary stages, and is currently limited to the cross-section of equity markets, we are able to identify significant patterns and attribute them statistically to particular macroeconomic time series, providing some economic interpretation for the estimation.

1.2 Liquidity Measurement in Practice

Research contributions generally focus either on liquidity in the narrow context of a particular financial market—so-called "market" or "microstructural" liquidity—or at the aggregate level of the financial system as a whole—so-called "global" or "funding" liquidity. How one measures liquidity depends in part on the portion of the financial system under examination. Brunnermeier and Pedersen [2009] connect these two strands of the literature, using bank balance sheets as an organizing device, as depicted in Figure 2. In this framework, the distinction between "market liquidity" and "funding liquidity" is based essentially on which side of a financial firm's balance sheet is involved. Market liquidity refers to the ease with which financial institutions can convert securities or loans from their asset portfolio to cash.[8] Financial institutions can participate as both suppliers and demanders of liquidity in these markets, as indicated by the bidirectional arrows on the left side of Figure 2. Funding liquidity refers to the ease with which institutions can obtain cash by borrowing in funding markets. Figure 2 underscores that one of the central functions of banks and similar intermediaries is to convert relatively long-maturity, low-liquidity commitments on the asset side to relatively short-maturity, high-liquidity obligations on the

[8] Harking back to Moulton [1918], Mehrling [2010] refers to this sort of liquidity as the "shiftability" of bank or dealer assets—i.e., the ability to shift them into cash. It is worth noting that the asset side of Figure 2 represents a primary point of contact between the financial system and the real economy. Commitments like corporate debt or mortgage loans typically translate directly into real activity such as workforce expansions and home improvements investments. Cornett et al. [2011] analyze market and funding liquidity empirically, along with their net effect on overall credit supply.

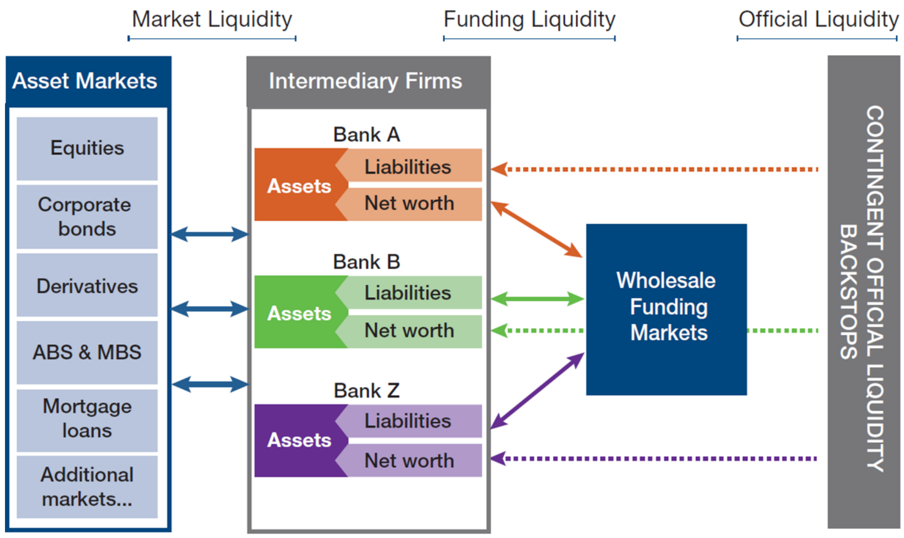

Figure 2: Liquidity Transformation via Financial IntermediationSource: OFR Analysis

liability side of intermediaries' balance sheets. Official liquidity represents the range of short-term cash resources available in a financial crisis—that is, when the wholesale funding markets fail—from central banks and other agencies, enterprises, and programs with explicit or implicit taxpayer backing.[9] Because these liquidity resources typically come into play only in unusual but important occasions, they appear as dashed arrows in the figure, which flow in only one direction.

1.2.1 Aggregate Liquidity Measures

An institution's market liquidity needs are typically handled by drawing on its own cash reserves or by selling assets for cash. If the ordinary give and take of trading in the asset markets does not net out, the firm can turn to the funding markets to borrow or lend the difference. Wholesale funding markets thus aggregate much of the endogenous net supply and demand of liquidity overall, and prices in these markets provide a bellwether for the

[9]In the U.S., these resources include deposit insurance, the Fed's discount window, Federal Home Loan Bank advances, as well as the numerous emergency facilities created as expedients in the recent crisis. Official liquidity is sometimes called "outside" liquidity, meaning it comes from outside the markets themselves. For a definition and model of inside and outside liquidity, see Holmström and Tirole [2013].

state of system. Figure 3 depicts several commonly used measures of aggregate liquidity conditions derived from prices in wholesale funding markets. Following Brunnermeier [2009], Boudt et al. [2013], and Boyson et al. [2010], we proxy the Treasury-Eurodollar (TED) spread as the difference between 3-month T-bill yields and 3-month LIBOR.[10] A frequently cited alternate spread measure for funding liquidity conditions is the LIBOR-OIS (London interbank offered rate vs. overnight index swap) spread; see Gefang et al. [2011], Michaud and Upper [2008], and Taylor and Williams [2009]. Both spreads capture deviations of borrowing conditions in the interbank markets from "pure" credit-risk-free borrowing, and thus reflect a variety of sources of reticence to lend to banks, including credit risk and aggregate liquidity anomalies. The VIX® is a traded index of market volatility. Its gradual downward drift during the pre-crisis period (roughly 2002 to 2007) is symptomatic of the so-called "volatility paradox": market risk as measured by price volatility was dropping while overall risk exposures were simultaneously (and not coincidentally) building across the system (Brunnermeier and Sannikov [2012]).

The first major foreshock of the crisis came in August 2007, triggered by an absence of liquidity that prevented BNP-Paribas from marking to market several of its investment funds backed by subprime mortgages. The surprise provokes a sharp but very temporary drop in the 3-month repo rate and a simultaneous more permanent jump in both the TED and LIBOR-OIS spreads. Market turmoil persists through the failure of Bear Stearns in March 2008, until the failure of Lehman Brothers in September 2008 raises liquidity problems to a new level. Interest rate spreads provide a high-frequency glimpse into aggregate liquidity conditions, but price moves can only hint at the underlying changes in cash holdings. As financial firms withdrew, the aggregate endogenous liquidity supply in the wholesale funding markets became inadequate to satisfy current cash obligations, and

[10] Futures traders at the Chicago Mercantile Exchange first noted the TED spread in the early 1980s, where they tracked the pricing differential between the 3-month T-bill futures and 3-month Eurodollar futures contracts, which traded in neighboring pits McCauley [2001].

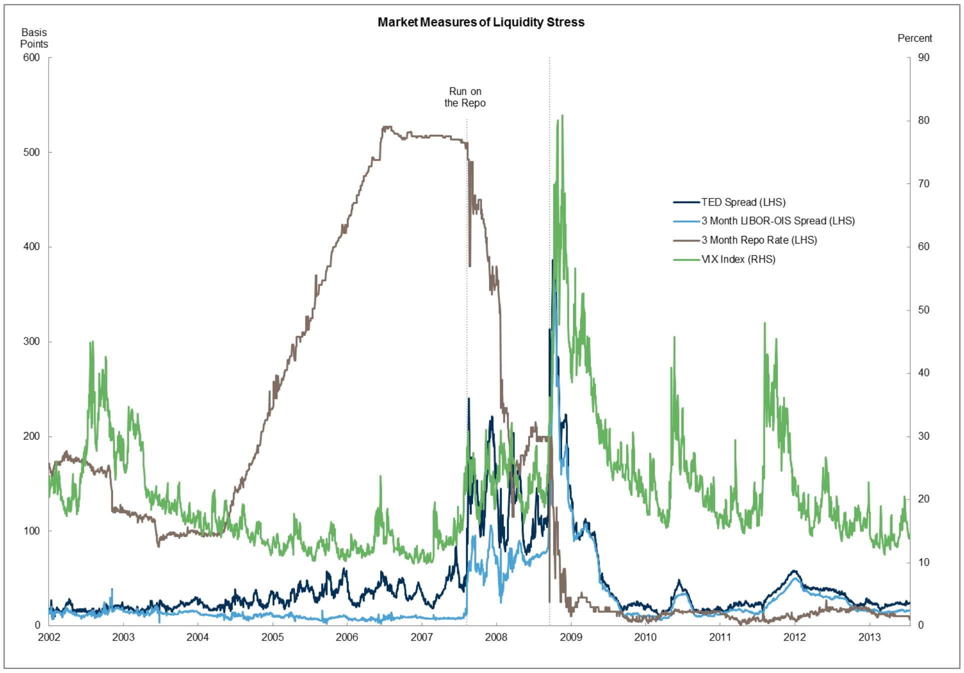

Figure 3: Bellwether Liquidity Measures from Wholesale Funding Markets Source: Bloomberg

official liquidity providers were forced to step in. The Federal Home Loan Banks (FHLBs), provided a first line of defense, via advance funding to member institutions, which include most large commercial banks (Ashcraft et al. [2010]. Over the first few months after August 2007, FHLB advance funding increased by more than $200 billion. Although this recourse to official liquidity was modest compared to what was to come, it was a significant departure from business as usual at the time. The full-blown crisis emerges with the Lehman failure in September 2008. Wholesale markets collapsed, and financial firms proceeded en masse to the Fed's backstop liquidity programs.[11] Surprisingly, the Fed's signature lender-of-last-resort facility, the discount window, plays a miniscule role throughout

[11] The Bear Stearns failure necessitated recourse to the Fed, including the creation of two brand new liquidity vehicles, the Primary Dealer Credit Facility (PDCF) and the Term Securities Lending Facility (TSLF). Following the Lehman failure, the brunt of the wave of new funding demands was borne initially by the Fed's swaps facility, along with FHLB advances, the PDCF and TSLF again, and the newly created ABCP MMMF liquidity facility (AMLF). Much of this funding subsequently transitioned to other programs, including the Treasury's new Troubled Asset Relief Program (TARP), created in October 2008. For an analysis of the Fed's various large-scale asset purchase programs, see Chen et al. [2012a], D'Amico and King [2012], and D'Amico et al. [2012].

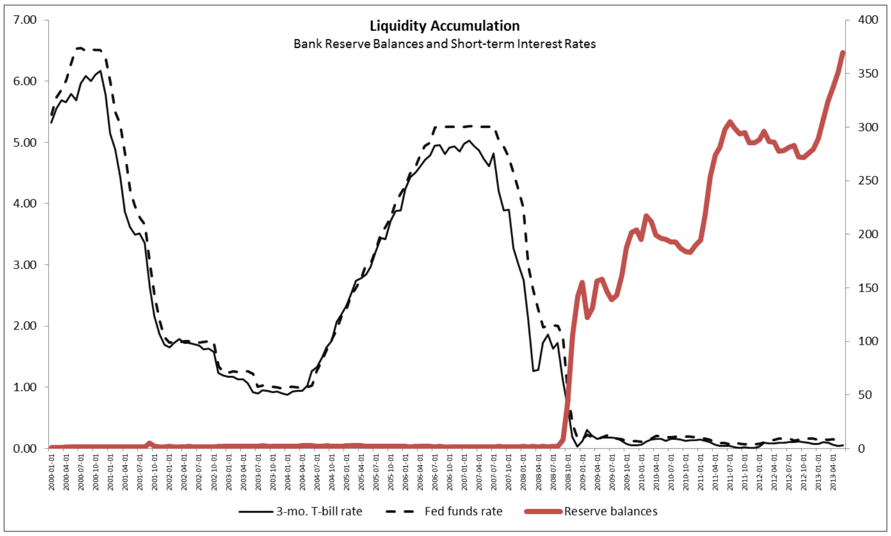

Figure 4: Bank Reserve Balances and Short-term Interest Rates
Source: St. Louis Federal Reserve Economic Data

the episode.

The sharp shift in supply and demand for wholesale liquidity also shows up in the management of banks' cash reserves, as depicted in Figure 4. Prior to the Lehman shock, aggregate bank reserve balances (which paid no interest prior to July 2013) hovered near zero; banks continued to rely on wholesale funding markets to meet short-run cash contingencies. After the Lehman failure, banks begin to hold precautionary reserve balances in significant quantities; this practice of reserving has continued essentially unabated. At the same time, the Federal Reserve has flooded markets with liquidity, driving yields on overnight Fed funds and T-bills to near zero. The persistence of the short-term riskless rate near the zero lower bound while loanable funds pool up—potential lenders always have the alternative of holding cash instead of accepting a negative return—suggests a market failure.

1.2.2 Granular Liquidity Measurement

At the aggregate level of funding liquidity, a primary concern is whether the financial system has the internal flexibility to satisfy all of its immediate funding needs. Economic equilibrium means that every borrower in the wholesale funding markets should be able to find a willing lender, but the long and painful history of systemic crises demonstrates that this equilibrium is not reliable. This is the purview of central banking, macroprudential regulation, and systemic supervision. At the other extreme, market liquidity focuses on the intricacies of specific markets. The availability of granular, high-frequency data on transaction prices, bid-ask quotes, trading volumes and customer order flows facilitates detailed modeling of the behavior of market participants. Market liquidity is typically measured by the price concession required to unload (or acquire) a given investment position in a very short time interval.

By design, liquidity measures based on prices and volumes in the wholesale funding markets aggregate information from across the financial system. There are only a relative handful of funding products traded in these markets, which in turn are dominated by a relatively small set of very large institutions. Yet the issue of liquidity also applies to the many thousands of markets for equities, bonds, indexes, commodities, and derivatives. Although these smaller markets are not as immediately connected to system-level stability, they are much more numerous than the interbank funding markets. Liquidity in these markets may therefore carry additional information about overall liquidity that is lost in the aggregation to the wholesale level. In particular, asset markets offer a smorgasbord of different industries, product types, geographic concentrations, maturity habitats and credit grades that is not available in the short-term, interbank funding markets. It is an empirical question whether this diversity generates measurable cross-sectional patterns in liquidity, and whether this cross-sectional information is helpful in understanding systemic behavior.

This work builds on earlier studies that look for aggregate liquidity patterns. Chordia et al. [2000] was the first in a series of papers to search for "commonalities in liquidity" in the cross-section of equity markets. They perform time-series regressions of liquidity, measured as market depth and bid-ask spreads, for individual stocks on cross-sectional average measures of liquidity. The data are noisy (R^2s are low), but there is strong evidence of contemporaneous correlation between individual stocks and the aggregate. Karolyi et al. [2012] extend the analysis to an international comparison of thousands of stocks in 40 countries. Again, commonalities in liquidity exist, and unsurprisingly differ significantly across countries and over time.[12] Recent research by Chen et al. [2012b] combines price information from financial markets with quarterly quantity information from the Flow of Funds data in an effort to distinguish the differential impact of shifts in liquidity demand versus liquidity supply. They distinguish between "core" and "noncore" liquidity, with the primary difference arising from the inclusion in noncore of financial firms' liabilities held by other financial institutions.

Our contribution differs in scope from these earlier studies. The goal is to exploit fine-grained information—such as daily data on prices and volumes—from a wide range of asset markets, including equities, bonds, and commodities. By casting a wide net across many diverse instrument types, the approach should have a better chance of detecting emerging liquidity anomalies and identifying key liquidity indicators and important patterns among the markets being monitored. We exploit recent work by Kyle and Obizhaeva [2014] on generalized "price impact" measures of market liquidity. Price impact refers to the change in market prices caused by a one-directional order flow (buy or sell) of a given size.[13] For example, it is always a challenge for an investor or dealer to sell a large block of stock, but especially so in stressful market conditions when crowded trades can conspire to generate fire-sale price drops. By design, price impact measures reveal the non-linearity of market

[12]Karolyi et al. [2012] define commonality by the R^2 of each stock's daily price-impact measure, per Amihud [2002], on the average price impact for all other stocks in the country. Individual stock commonality measures are averaged to get a country-level commonality index. Karolyi et al. [2012, p. 99] attribute the time-series variation to both supply- and demand-side proxies in funding markets via regression analysis, noting that "demand-side explanations are more reliably significant."

[13]While a number of alternative metrics exist (e.g. big ask spreads, price impact measures, trading volumes), these tools are typically restricted or customized to a single market.

depth in the face of transaction pressure. Kyle and Obizhaeva [2014] posit a scaling of trading activity by "business time" that allows for the construction of price impact measures that are invariant to both the institutional details of diverse market microstructures as well as the pace of trading activity. This invariant approach facilitates a systematic analysis by allowing the repeated application of the same calculation methodology across a range of markets.

2 Model Description

A central goal of this research is to identify broad patterns or commonalities in market liquidity that might offer clues to system-wide liquidity. A practical implication is the need for measures of liquidity that are consistent and comparable across a broad spectrum of financial markets. Ideally, such measures would be available at daily (or higher) frequencies to support market monitoring and decision making. The price-impact measure of Kyle and Obizhaeva [2014], described in the next subsection, is well suited to these requirements. While plausible, the usefulness of such an approach is ultimately an empirical question. In particular, for system-wide market liquidity monitoring, a successful measure should be responsive to changes in local liquidity conditions (which can shift abruptly) at a daily frequency, should demonstrate a reliable statistical connection to the aggregate liquidity conditions, and should enable attribution of systemic liquidity disruptions to specific market sectors to help focus attention for subsequent analysis. As a robustness check, we also consider several alternative daily measures of market liquidity.

The primary assumption underlying our analysis is that the liquidity in a specific financial market (as defined by a portfolio of publicly traded securities) switches between distinct states, jumping from one state to another (e.g., low to high liquidity) and then staying that new state for a random period of time. The subsequent observed level of liquidity is a random deviation from the level related to the average liquidity for the current state. While the observed liquidity sometimes falls between the average liquidity for two different states, ultimately it is the persistence of the observed liquidity that identifies an underlying state. When liquidity from multiple financial markets is considered, we augment the liquidity models for each individual market with an "add-on" hierarchical model which can explain, in part, periods of coordination where a large subset of the financial markets exhibit similar liquidity patterns. The hierarchical portion of the model, which links individual liquidity models together, can be considered an "add on" because it does not feedback into the individual level liquidity models. Instead it allows us to determine whether macro variables based on economic summaries of the broader financial markets and economy are related to liquidity states across multiple markets. Uncovering such a relationship offers a framework for understanding and potentially predicting (by predicting the underlying dynamics of the macro variables) when system-wide liquidity stress might occur.

2.1 Market Liquidity Measures

We consider several alternative measures of market liquidity. The "microstructure-invariant" measure of Kyle and Obizhaeva [2014] is particularly well suited to our goals of cross-sectional applicability and daily availability.[14] In addition, we consider several more traditional market liquidity measures, including proportional bid-ask spreads and market turnover to confirm the robustness of our results and our focus on the microstructure-invariant approach.

The basic intuition of the microstructure-invariant measure is that market illiquidity is revealed by a market's resilience (or lack thereof) to net speculative order flow. Such

[14] Kyle and Obizhaeva [2014] is a combination and refinement of two earlier papers, Kyle and Obizhaeva [2011b,a]. In addition, they have applied their approach to a historical study of large liquidity events Kyle and Obizhaeva [2012].

speculative "bets" tend to arrive at different rates in different markets, creating a phenomenon of market-specific "business time" defined by the pace of speculative trading. In equilibrium, liquidity provision will adapt to the bet-arrival intensity in each market, so that an appropriate normalization of measured price-impact can place diverse markets on a comparable (microstructure-invariant) scale. Kyle and Obizhaeva [2014, equations (70) and (71)] are able to reduce their theoretical approach to a closed-form liquidity measure, which they calibrate empirically under both a linear and a square-root specification for the shape of the price-impact response. Defining σ as the expected daily volatility of returns, V as expected daily trading volume (in shares or analogous units), X as a "typical" order size for a specific market, and W as the level of betting activity (measured as price times expected volatility times expected volume), their linear calibration is:

$$C(X) = \frac{\sigma}{0.02}\left[\frac{8.21}{10^4}\left(\frac{W}{(0.02)(40)(10^6)}\right)^{-1/3} + \frac{2.50}{10^4}\left(\frac{W}{(0.02)(40)(10^6)}\right)^{1/3}\frac{X}{(0.01)V}\right] \quad (1)$$

and the square-root calibration is:

$$C(X) = \frac{\sigma}{0.02}\left[\frac{2.08}{10^4}\left(\frac{W}{(0.02)(40)(10^6)}\right)^{-1/3} + \frac{12.08}{10^4}\left(\frac{X}{(0.01)V}\right)^{1/2}\right] \quad (2)$$

In practice, their is no unambiguously "right" way to set the typical order size, X. An important consideration is to normalize X by trading activity in each market, to measure price-impact responses on a comparable scale across markets. A corollary requirement is to calibrate equations (1) and (2) to be consistent with the definition of X. The particular calibrations in (1) and (2) assume that order size is a constant fraction of average daily volume. This implies, for example, that the dollar size of the order should move in the same direction as dollar volume. A plausible alternate is to hold the dollar size of an order constant over time, so that the relative size of the order (as a fraction of volume) moves inversely with volume. As a robustness check, without re-estimating the parameters in (1) and (2), we recalibrated order size as a constant dollar value. The price-impact results were similar in magnitude, but noisier than for the calibration of X as a constant fraction of daily volume; the results presented below use the constant-fraction specification.[15] We follow Kyle and Obizhaeva [2014] in using the average trading volume over the preceding month (20 trading days) as a proxy for expected volume in (1) and (2). Similarly, we use the average realized volatility of daily returns over the preceding month as a proxy for expected volatility.[16]

Our primary motivation for using the microstructure-invariant metric is to measure liquidity in a comparable way across a wide spectrum of markets, covering a range of different asset classes. For example, if one cannot assert *ex ante* where liquidity issues will first emerge during an episode of financial instability, one should monitor liquidity conditions as broadly as possible. However, there are numerous other metrics available in the literature. As a robustness check, we consider two other measures of market liquidity, turnover and bid-ask spreads, that are widely used and available at a daily frequency.[17] Turnover (sometimes called velocity) is defined as the simple average of instrument-level

[15] Another possibility is to allow the size of the orders to adjust to market liquidity changes, in a manner more rigorously consistent with the equilibrium arguments in Kyle and Obizhaeva [2014]. For example, speculative order flow should be directly proportional to both overall liquidity and to the cube root of expected dollar volume, per Kyle and Obizhaeva [2014, equation (8)], so that order size should increase with liquidity, while decreasing as a fraction of daily volume. We are grateful to Pete Kyle for a helpful clarifying conversation around this issue. Because we are interested here in applying rather than testing or extending the Kyle and Obizhaeva [2014] model, we restrict attention to the calibrations in equations (1) and (2), together with a fixed dollar size for the order impulse, X.

[16] There are obviously more sophisticated ways to estimate conditional expected volatility. However, Kyle and Obizhaeva [2014, p. 31] note that using a more exacting ARIMA model in log volatilities produces quantitatively similar results. Again, because we are interested in applying rather than refining their model, we apply the simpler specification.

[17] See Hameed et al. [2010] or Cumming et al. [2011] for examples of these measures in practice. See Gabrielsen et al. [2011] for a brief survey of a number of alternatives. We also considered the index of Martin [1975], for the simple reason that it is easy to calculate. This index is based on a simple Brownian motion model in which the rate of variation of price changes is assumed proportional to business time as defined by trading volume, $\Delta P_t \sim N\left(0, \sigma^2 \Delta t\right) = N\left(0, \sigma^2 V\right)$. The index value is given by $M_t = \frac{1}{N_i} \sum_{j \in J_i} \left[\frac{ln(P_{jt}/P_{jt-1})^2}{V_{jt}} \right]$, where P_{jt} is the closing price of asset j on day t, and V_{jt} is the number of units transacted. The index is dominated by noise in our data, and we do not report the results.

turnover ratios across a portfolio, J_i containing N_i positions:

$$T_{it} = \frac{1}{N_i} \sum_{j \in J_i} \frac{V_{jt}}{\Lambda_{jt}}$$

where V_{jt} is daily trading volume and Λ_{jt} is the number of shares (or analogous units) outstanding. The quantity outstanding is not well defined for futures markets, so the application of this measure has natural limitations in a system-wide analysis. The bid-ask spread is the cross-sectional average (across the portfolio J_i) of percentage spreads:

$$S_{it} = \frac{2}{N_i} \sum_{j \in J_i} \frac{A_{jt} - B_{jt}}{A_{jt} + B_{jt}}$$

where B_{jt} and A_{jt} are the bid and ask prices, respectively, for instrument j at time t; $(A_{jt} + B_{jt})/2$ is the spread midpoint.

2.2 Univariate Models of Latent Structure

We begin the analysis of latent structure by proposing a univariate, hidden Markov chain model for each financial market where liquidity is a random deviation from a latent value associated with each state of the hidden Markov chain. We consider two variations of random deviations, independent deviations around an average level and deviations around a value that mean reverts around an average level. Initially we assume that the dynamics of these models are unrelated, then we propose a hierarchical (multiple market) model where macro variables are used to explain the states identified by the collection of univariate hidden Markov models.

Liquidity measurement over T periods, $y_i = (y_{i1}, ..., y_{iT})^T$, for market i are assumed to be a random normal deviation around a dynamic, latent level of liquidity $\theta_i = (\theta_{i1}, ..., \theta_{iT})$, or

$$y_i = \theta_i + \epsilon_i,$$

where

$$\epsilon_i \sim N(0, \sigma_i^2 I_T)$$

and I_T is a T dimensional identity matrix. For the first version of the model, the hidden Markov chain (HMC) only version, the latent level θ_i is one of K levels, each of which represents a different level of liquidity or state for each market specific, discrete-time hidden Markov chain D_i, or

$$\theta_i = F_i \bar{\theta}_i,$$

where $\bar{\theta}_i$ is a $K \times 1$ vector and each element represents the average level of liquidity for the k^{th} state of D_i, or

$$F_i(t, k) = I\{D_{it} = k\}$$

and $I\{\}$ is an indicator function equaling either 0 or 1.

Typically the HMC version is sufficient to identify structural shifts in liquidity patterns, however, there are some markets where the local variation in the level of liquidity supports an overly large number of hidden states. In these cases, we use a mean reverting version of the hidden Markov chain model. For the second version of the model, the Mean Reverting, hidden Markov chain (MRHMC) version, the latent level θ_i mean reverts around the average level associated with the state of D_i, or for $t = 2, ..., T$,

$$\Delta\theta_i = \gamma_i((\theta_i)_{-T} - (F_i\bar{\theta}_i)_{-1}) + (\xi_i)_{-1}, \tag{3}$$

where $()_{-1}$ indicates that the first element and $()_{-T}$ indicates that the last element of the vector $()$ has been removed, and

$$\Delta\theta_{it} = \theta_{it} - \theta_{it-1}.$$

For $t = 1$ we let

$$\gamma_i\theta_{i1} = \gamma_i\bar{\theta}_{i1} + \xi_{i1},$$

where

$$\xi_i \sim N(0, w_i I_T).$$

We require $0 < \gamma_i \leq 1$, which ensures that θ_i is stationary and increases the variance of θ_{i1}, allowing the starting value of θ_i to be relatively vague. Alternatively, (3) can be rewritten as,

$$L_i\theta_i = \gamma_i F\bar{\theta}_i + \xi_i, \tag{4}$$

where L_i is a sparse $T \times T$ matrix with zeros except for the following elements, $L_i(j, j) = 1$ and $L_i(j, j-1) = \gamma_i - 1$ for $j > 1$ and $L_i(1, 1) = \gamma_i$.

For both versions of the model, the dynamics of D_i are given by an initial probability density ν_i, a $K \times 1$ vector, and a transition probability density P_i, a $K \times K$ matrix. Given a realization of D_i, its density is given by

$$f(D_i) = \nu(D_{i0}) \prod_{t=1}^{T} P_i(D_{it-1}, D_{it}).$$

We assume conjugate priors for σ_i^2 and w_i (inverted Gamma), ν_i and each row of P_i (Dirichlet) and $\bar{\theta}_i$ and γ_i (truncated Normal). In addition, we use subjective priors based on initial conditional maximum likelihood estimates of summaries of the data, in order to ensure that the filtered hidden Markov chain model is able to clearly define between the dynamics of the hidden Markov chain and the dynamics of the latent value θ_i.

Full Conditional Distributions HMC Model

We use Markov chain Monte Carlo (MCMC) analysis to infer parameter values for both of these univariate models and multivariate model built on these univariate models. For a description of MCMC methods see Brooks, Gelman, Jones and Meng (2011) and Gelman, Carlin, Stern and Rubin (2000). The full conditional densities used in the MCMC analysis for the HMC model are as follows:

$$\bar{\theta}_i|- \sim N\left(\left(\frac{1}{\sigma_i^2}F_i^T F_i + \frac{1}{\tau_{\bar{\theta}_i}^2} I_K\right)^{-1}\left(\frac{1}{\sigma_i^2}F_i^T y_i + \frac{\mu_{\bar{\theta}_i}}{\tau_{\bar{\theta}_i}^2}\bar{\bar{\theta}}_i\right), \left(\frac{1}{\sigma_i^2}F_i^T F_i + \frac{1}{\tau_{\bar{\theta}_i}} I_K\right)^{-1}\right) I\{\bar{\theta}_{i1} < \ldots < \bar{\theta}_{iK}\};$$

were $-$ represents everything else remaining in the model and

$$\frac{1}{\sigma_i^2}|- \sim Gamma\left(shape_{\sigma_i^2} + \frac{T}{2}, scale_{\sigma_i^2} + \frac{1}{2}(y_i - F_i\bar{\theta}_i)^T(y_i - F_i\bar{\theta}_i)\right).$$

Realizations of the hidden Markov chain D_i, conditional on the remaining parameters and Data, are generated following the Filter Forward, Sample Backwards approach commonly used with discrete-time Hidden Markov chains, see Bahm, Petrie, Soules and Weiss (1970) and described in a more general continuous-time framework in Cappe, Moulines and Ryden (2005). For completeness the filter forward equations for the HMC model are given below,

$$f\left(y_{it}|-,\mathcal{F}_{it-1}\right)=\sum_{k=1}^{K}f\left(y_{it}|D_{it}=k,-,\mathcal{F}_{it-1}\right)f\left(D_{it}=k|-,\mathcal{F}_{it-1}\right), \tag{5}$$

were $\mathcal{F}_{it}=\{Y_{i1},...,Y_{it}\}$ and by

$$f\left(D_{it}=k|-,\mathcal{F}_{it}\right)=\frac{f\left(y_{it}|D_{it}=k,-,\mathcal{F}_{it-1}\right)f\left(D_{it}=k|-,\mathcal{F}_{it-1}\right)}{f\left(y_{it}|-,\mathcal{F}_{it-1}\right)}. \tag{6}$$

Specifying a vague initial state probability, e.g.,

$$f(D_{i0}=k|-,\mathcal{F}_{i0})=\frac{1}{K},$$

completes the forward recursion. The key equation for the backward sampling is given by the density of the hidden Markov chain, conditional on all of the data, or

$$\begin{aligned}&f\left(D_{iT-t}=k|-,\mathcal{F}_{iT}\right)=\\&\textstyle\sum_{j=1}^{K}\frac{f(D_{iT-t+1}=j|D_{iT-t}=k,-,\mathcal{F}_{iT-t})f(D_{iT-t}=k|-,\mathcal{F}_{iT-t})}{f(D_{iT-t+1}=j|-,\mathcal{F}_{iT-t})}f\left(D_{iT-t+1}=j|-,\mathcal{F}_{iT-t+1}\right).\end{aligned} \tag{7}$$

Given these formula, generating a realization is straightforward: i) calculate the forward filter; ii) generate a sample for D_{iT} from (6), with $t=T$; and iii) recursively calculate $f\left(D_{iT-t}=k|-,\mathcal{F}_{iT}\right)$, conditional on all of the draws $\left(D_{iT},...,D_{iT-t+1}\right)$ using (7) and use this to generate a sample for D_{iT-t}. Given a realization of D_i the full conditional distribution for each row of the transition probability is given by

$$P_i(j,:)|-\sim Dirichlet\left(\alpha_{i1}+n_{ij1},...,\alpha_{iK}+n_{ijK}\right),$$

where α_{ijk} is the prior associated with D_i jumping from state j to k and n_{ijk} is the actual number of times that the current realization of D_i jumps from state j to state k. A similar full conditional density exists for ν_i, but this is inconsequential as the backwards recursion dominates the *a priori* initial state. It is important to note that the MRHMC model is disentangling two dynamics, the dynamics of D_i and θ_i. In practice we found that the model required a strong priors on the dynamics of D_i in order to obtain a meaningful distinctions between these two dynamics. Setting α_{ikk} to a sufficiently large value, suggesting *a priori* that the hidden chain is persistent results in a clean separation of these two competing dynamics.

Full Conditional Distributions MRHMC Model

There are similarities between some of the full conditional densities of the MRHMC model and the HMC model. The full conditional density for P_i is unchanged, while the full conditional density for $\bar{\theta}_i$ is obtained by replacing y_i with $\frac{1}{\gamma_i} L_i \theta_i$ and σ_i^2 with $\frac{w_i}{\gamma_i^2}$. The full conditional density for $\frac{1}{\sigma_i^2}$ is obtained by replacing $F_i \bar{\theta}_i$ with θ_i and the full conditional density for D_i is obtained by replacing the likelihood $f\left(y_{it}|D_{it}=k,-,\mathcal{F}_{it-1}\right)$ used in (5) and (6) with $f\left(\theta_{it}|D_{it}=k,-,\theta_{it-1},...,\theta_{i1}\right)$. The remaining full conditional densities for the MRHMC model are as follows:

$$\frac{1}{w_i}|- \sim Gamma\left(shape_{w_i}+\frac{T}{2}, scale_{w_i}+\frac{1}{2}(L_i\theta_i-\gamma_i F_i\bar{\theta}_i)^T(L_i\theta_i-\gamma_i F_i\bar{\theta}_i)\right);$$

$$\gamma_i|- \sim N\left(\Sigma_i\left(\frac{1}{w_i}A_i^T\Delta\theta_i+\frac{\mu_{\gamma_i}}{\tau_{\gamma_i}^2}\right),\Sigma_i\right)I\{0<\gamma_i\leq 1\};$$

where

$$\Sigma_i = \left(\frac{1}{w_i} \left(A_i^T A_i + (\theta_{i1} - \bar{\theta}_{iD_{i1}})^2 \right) + \frac{1}{\tau_{\gamma_i}^2} \right)^{-1} \text{ and } A_i = \left((\theta_i)_{-T} - (F_i \bar{\theta}_i)_{-1} \right),$$

and

$$\theta_i | - \sim N \left(\left(\frac{1}{w_i} B_i + \frac{1}{\sigma_i^2} I_K \right)^{-1} \left(\frac{\gamma_i}{w_i} B_i \left(L_i^{-1} F_i \bar{\theta}_i \right) + \frac{1}{\sigma_i^2} y_i \right), \left(\frac{1}{w_i} B_i + \frac{1}{\sigma_i^2} I_K \right)^{-1} \right),$$

where B_i is a $T \times T$ matrix given by

$$B_i = \left(L_i^{-1} (L_i^{-1})^T \right)^{-1}.$$

An alternate approach for sampling θ_i and γ_i, conditional on D_i, is to treat them as a discrete-time, dynamic linear model and use a filter forward sample, backwards strategy like the Kalman Filter, see Kalman (1960) for the original reference and Carter and Kohn (1994) and DeJong and Shephard (1995) for MCMC based inference approaches. Although we explored a filter forward, backwards sample approach, we found that this approach was not as stable as the regression based approach detailed above. Obviously one disadvantage of the regression approach is the need to calculate B_i, which requires the inversion of a $T \times T$ matrix, something that can become computationally prohibitive as T becomes large. Fortunately, the form of L_i results in a banded matrix for B_i where every element, except for the diagonal and the rows just next to the diagonal, are zeros. In addition, the non-zero elements are functions of γ_i; to be explicit

$$B_i(j,j) = \begin{cases} 1, & \text{if } j = T \\ \frac{1}{2} + 2\left(\frac{1}{2} - \gamma_i\right)^2, & \text{otherwise} \end{cases}$$

and

$$B_i(j, j-1) = B_i(j-1, j) = \gamma_i - 1.$$

2.3 Hierarchical Model

The hierarchical "add on" model runs a collection of HMC or MRHMC models in parallel, one for each of the N markets under consideration. Each sweep of the MCMC algorithm generates a realization of the latent hidden Markov chain for each market, resulting in a collection of realizations $(D_1), ..., (D_N)$. For each realization every individual time point can be viewed as a draw from a multinomial distribution that is driven by a set of time-varying covariates x_t – the macro variables. Conditional on a current realization of the hidden Markov chain, the "add on" portion of the model is a multinomial Probit model, where

$$f(D_{it} = k) = f(\tilde{z}_{itk} > \tilde{z}_{itl}, l \neq k),$$

and $\tilde{z}_{it}$ is multivariate normal or

$$\tilde{z}_{it} \sim N\left(\tilde{\beta} x_t, \tilde{\Sigma}\right).$$

We follow McCulloch and Rossi's (1994) approach, which builds on Albert and Chib (1993), for dealing with the identification issues that arise in using a Bayesian approach for the multinomial Probit model. A related, alternative approach is discussed in McCulloch, Polson and Rossi (1998). The additive identification is overcome by forcing the latent value for state 1 to always be zero. This is done by defining z_{it} as follows,

$$z_{itk} = \tilde{z}_{itk} - \tilde{z}_{it1},$$

which results in (8) becoming

$$f(D_{it} = k) = \begin{cases} f(0 > z_{itl} l \neq k), & \text{if } k = 1 \\ f(z_{itk} > \max(0, z_{itl}), l \neq k, l > 1), & \text{if } k > 1 \end{cases}$$

where

$$z_{it} \sim N(\beta x_t, \Sigma).$$

and β is a $(K-1) \times p$ matrix, where p is the number of macro variables, including an intercept. The scale identification is overcome by restricting

$$\Sigma_{1,1} = 1.$$

We assume conjugate priors for β and Σ. Following McCulloch and Rossi (1994) we sample β and Σ from the unconstrained full conditional densities using Gibb Samplers and then rescale by dividing these draws by $\Sigma_{1,1}$, which enforces the above constraint.

The hierarchical portion of the model is considered to be an "add on" to the model because the distribution of the hidden Markov chains $(D)_1, ..., (D)_N$ does not depend on the multinomial Probit probabilities or, more to the point, they do not depend on the macro variables. In order for the distribution of the hidden Markov chains to depend on the macro variables, we would need to model the transition between the latent liquidity states (as opposed to modeling the states themselves) as multinomial random variables conditional on the macro variables, which is a task that we leave for future research. Instead, the "add on" model summarizes the relationship between latent states and the macro variables, acting as a supplemental analysis that gives us insights into how the latent liquidity states related to the macro variables, but makes no assumption about nor gives any insights into how the macro variables impact the dynamics of the latent states. A less sophisticated approach to getting the same insights would be to save a realization of each hidden Markov chain from the MCMC analysis, and then calibrate a multinomial Probit model for this collection of realizations. Repeat this multiple times, each time with a different set of realizations obtained by stopping the MCMC analysis at a random time, which would result in a set of multinomial Probit parameter estimates, one for each set of realizations, and then average the parameters estimates from all of these analysis. Our approach is more elegant as it updates the parameters of the multinomial Probit model with each sweep of the MCMC analysis. The basis of the relationship between the hidden states and the macro variables is determined by the portions of the hidden Markov chains which are relatively stable (i.e., have a high probability of being in one of the states), which holds for large portions of time over each of the markets that we are considering. Portions of the hidden Markov chain that tend to switch states (have a probability that is distributed between two or more states) have less impact as the hidden Markov chains alternate between these competing states during the analysis and require estimates of β, which can reasonably accommodate this oscillation.

The fact that the state of a hidden Markov chain can switch states during the analysis presents a technical challenge. When the hidden chain changes state, the latent values from the multinomial Probit model, the z_{it} have to change to match their likelihood function (e.g., assume that chain i at time t changes from state 2 to state 3, then z_{it3} must become positive and z_{it2} must be less than z_{it3}). From a practical standpoint, we found that when a hidden Markov chain changes state, we can sample from the truncated, full conditional density of each latent variables in order to impose the new ordering, but doing this once is typically not sufficient to provide a stability for the estimate of β and Σ. This stability issues can be overcome by drawing a small number (on the order of a few dozen) samples of all of the related latent variables (e.g., draw repeatedly from z_{it} when a new ordering constraint imposed by the change in state).

3 Data

We measure market liquidity on a daily basis across 33 "markets," covering thousands of individual securities in four different asset classes. One important goal of casting a wide net across a diverse sample is to improve the chances of identifying emerging risks in liquidity, since it is difficult to assert a priori which market sector(s) will be affected first in an episode of illiquidity. Similarly, a broad panel should help in discerning significant patterns among the markets being monitored, as we map between local markets and system-level conditions. Finally, we hope that our mixing of several distinct asset classes in this analysis serves as an example of how to further expand the scope of the sample in subsequent research.

Specifically, our initial dataset includes the following instruments:

- All U.S. equities, Jan. 1986 – Mar. 2014, from the Center for Research in Security Prices (CRSP), which provides comprehensive coverage of security price, return, and volume data for the NYSE, AMEX, and NASDAQ stock markets; CRSP also includes Standard Industry Classification (SIC) for each security,

- All U.S. corporate bonds, Jul. 2002 – Mar. 2014, from the Trade Reporting and Compliance Engine (TRACE), which is the Financial Industry Regulatory Authority's (FINRA) real-time price dissemination service for the over-the-counter bond market, providing transaction data for all eligible corporate bonds, which include investment grade and high yield debt; we use the public TRACE database in this analysis,[18]

[18] We apply the heuristics of Dick-Nielsen [2009] to scrub the TRACE data. There is a separate "enhanced" version of the TRACE database, FINRA [2009], which does not truncate large trades, but which FINRA publishes only with a lag. We map TRACE bond identifiers (6-digit CUSIP codes) to the issuing firm's SIC code, derived from CRSP, Mergent, or Bloobmerg; approximately 2% of the bonds in sample could not be mapped, and were dropped from the analysis.

- West Texas Intermediate (WTI) light sweet crude oil futures, Jan. 1986 – Mar. 2014, from the New York Mercantile Exchange, which is the world's largest-volume futures contract traded on a physical commodity; we collected data for contracts with expirations from one-month to six-months from Bloomberg, and

- S&P 500 market volatility index (VIX®) futures, Apr. 2004 – Mar. 2014, from the Chicago Board Options Exchange, which is a pure play contract on implied volatility designed to reflect investors' view of future (30-day) expected stock market volatility; we collected data for contracts with expirations from one-month to nine-months from Bloomberg.

Our primary analysis of the liquidity measures starts in 2004, when all series are available. We also provide some secondary comparisons of the longer-term performance of price-impact measures for equities and WTI futures, extending back to 1986. We grouped both the CRSP equities and TRACE corporate bonds data into portfolios based on one-digit SIC codes. For both bonds and equities the SIC portfolios cover SIC codes 0 through 8.[19] This clustering into portfolios reduces the dimensionality of the analysis and presentation of results. In the case of corporate bonds, the combination into portfolios is a practical necessity for the calculation of returns and returns, because the trading of individual issues in this market is far too thin.[20]

We track the VIX® and WTI futures at the level of their relative maturity date, starting with the front-month contract. Actual calendar maturities follow a sawtooth pattern, as expiry dates gradually approach and abruptly transition to the next contract as expiration occurs. For both VIX® and WTI, and for the futures market generally, the near-dated

[19]The miscellaneous category (SIC 9, government establishments) is very lightly populated for both TRACE and CRSP, and did not provide sufficient observations for reliable analysis.

[20]Even with grouping into portfolios, there are numerous missing values in the time series of corporate bond activity. For calculating returns and volatility, we require that the most recent lagged observation be no older than a week (5 trading days). We experimented with higher and lower thresholds, out to 20 trading days, without a significant qualitative impact on the results.

contracts are usually more actively traded than the longer-maturity futures. There is no "official" longest maturity, but many possible long-dated contracts that simply never trade. For the VIX® futures, we draw the line at nine different securities from the front month out to nine months forward. For the WTI futures, we use six different securities from the front month out to six months forward.

4 Liquidity Regimes

In our initial analysis we estimated each price impact series independently, using both the hidden Markov chain (HMC) and the mean reverting hidden Markov chain (MRHMC) models. Although there is no coordination between the dynamics of the latent liquidity states across markets for this initial analysis, we find surprising consistency in the dynamics of market liquidity across all of these markets. Despite these common features, we also find interesting differences across the various markets in the lead-up to the recent crisis and in its aftermath. We formally explore these difference using the hierarchical model, which allows us to link the latent liquidity states (from multiple markets) together with a collection of macro variables. This provides a framework that allows us to asses the value of macro variables.

4.1 Individual Market Liquidity

We start by considering the performance of the two competing univariate models and provide evidence for our finding that there are essentially three different liquidity regimes across these different markets; then we report aggregate summaries based on these models.

General Findings

One of the benefits of analyzing detailed, high-frequency data is that it allows for a vi-

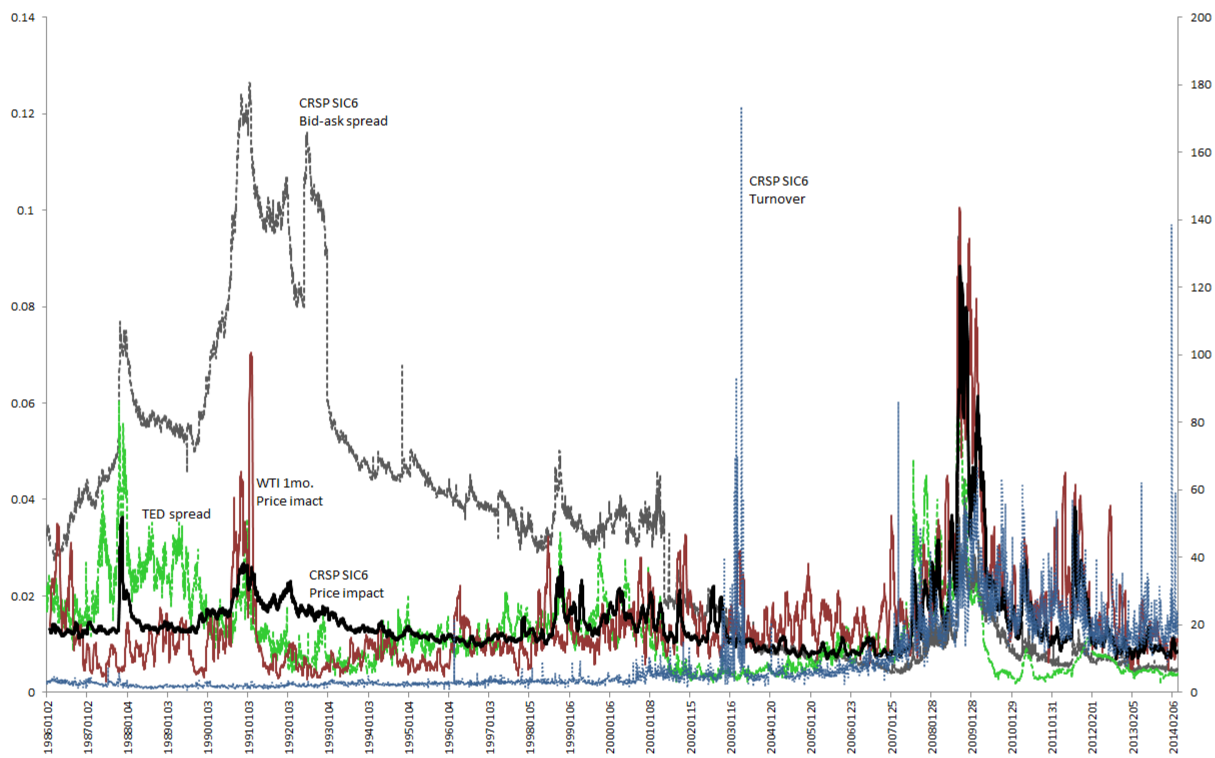

Figure 5: Examples of market liquidity measures, Jan. 1986 – Mar. 2014,
Sources: CRSP, Bloomberg, Mergent, WRDS, St. Louis Federal Reserve Economic Data, OFR Analysis

sual identification of basic cross-sectional liquidity patterns. This is valuable in developing research hypotheses, e.g., for the hierarchical model, or for general situational awareness in a policy context. Figure 5 shows a relatively long view (nearly 30 years, 1986-2014) for a variety of liquidity metrics. The turnover and bid-ask spread series are for our SIC6 equities portfolio (financial stocks); Kyle-Obizhaeva price impacts are shown for both SIC6 equities and the front-month WTI contract; the TED spread is a funding liquidity measure. Although Figure 5 is just a sampler, this picture is broadly consistent across all of the portfolios.

Several key facts are immediately apparent. Most importantly, all of the measures surge during the financial crisis, and especially after the failure of Lehman Brothers in September 2008, offering some face validity for all of them. Turnover, however, is plagued by occasional spikes in trading volume, which introduce significant noise in the form of large, transient outliers. The bid-ask spread appears to have undergone a significant structural shift, most likely attributable to institutional changes such as NASDAQ's introduction of decimal quotes in April 2001 and the large expansion of high-frequency trading in equities during the 2000s. The TED spread is a more stable series, but it is an aggregate measure and cannot support granular attribution of liquidity dynamics to specific sectors or markets. For example, illiquidity (as measured by price impact) surges in the oil futures market during the first Gulf war in 1991, but this effect is much less pronounced in the other liquidity series.

This overview of the metrics provides some initial validity for our focus on the Kyle-Obizhaeva price impact metric. It is available at a daily frequency and granular resolution for any market with published price and volume data. Moreover, it is statistically stable over time and comparable across markets, facilitating cross-sectional attribution and the search for system-wide patterns such as those described below.

Performance of Models

We found that both the HMC and MRHMC models could identify interesting liquidity regimes within the Kyle-Obizhaeva price-impact data over the range of different markets that we considered, but that their relative performance varied depending on the amount of 'local variability' of the liquidity in each state. The simpler HMC model can readily identify the three liquidity regimes, see Figure 6, for all of the equity markets. In some cases, such as the front month of the WTI contract we found that the MRHMC model performed slightly better than the HMC model using standard Bayesian model choice tools, see Figure 7.

We felt that it was important to have a parsimonious model, with respect to the number of states, and as a result adopted a prior on the model space, based on the number of hidden states, which gave a high penalty for increased complexity and resulted in three states being

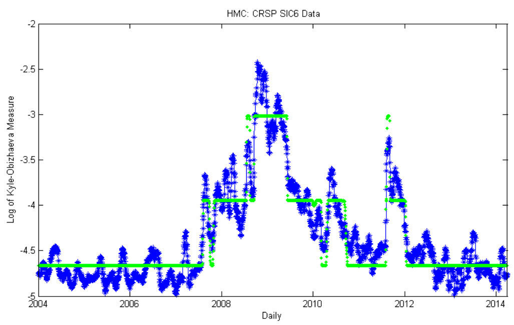

(a) Price-impact measure & level of attraction

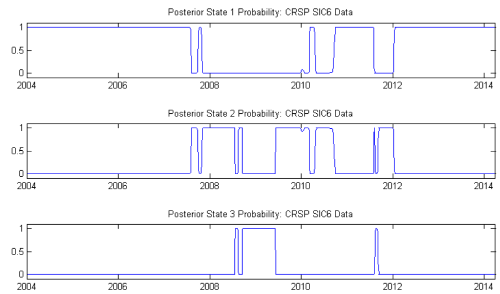

(b) HMC state probabilities

Figure 6: Equities, SIC6, Kyle-Obizhaeva Measure and HMC Estimates
Sources: CRSP, WRDS, OFR Analysis

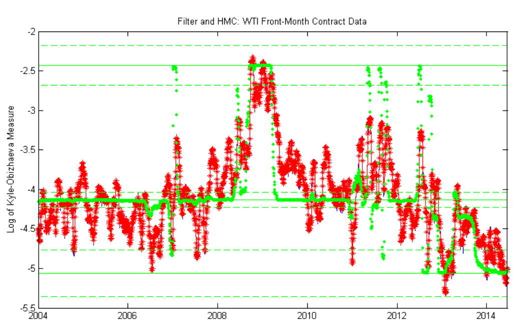

(a) Price-impact measure, mean-reversion level, & level of attraction

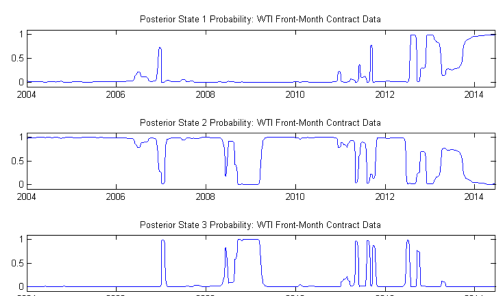

(b) HMC state probabilities

Figure 7: WTI futures, front month contract, Kyle-Obizhaeva Measure and MRHMC Estimates
Sources: Bloomberg, OFR Analysis

preferred for most markets. Alternative priors, which did not have as high of a penalty for complexity would support a high number of latent states (typically in excess of 10 to 15), where a visual inspection of these higher state models indicated that increasing the number of states essentially broke up the mid-liquidity state into a large number of sub-states.

Although the summary of the price-impact offers some insight, we find that in practice there is enough variation in these measure across the markets that a direct comparison of the levels from each market is helpful. The estimated states, however, of the hidden Markov chains from the univariate MRHMC model, is invariant to differences in the price levels across markets and can clearly identify varying levels of liquidity. We labeled these states the: (1) low, (2) intermediate, and (3) high price-impact states for each series, where high price impact means low liquidity, and vice versa. This analysis resulted in a daily estimate of the probability that each market was in each of these three unobserved states. Figure 8 presents the cross-sectional averages across the 33 series of these three probabilities, which must add up to one. Red indicates the likelihood of high price impact, and blue indicates low price impact; yellow is the intermediate state.

While there was diversity in market liquidity for these 33 series, there were also periods of common behavior. For example, in August 2011, the downgrade of U.S. Treasury debt by Standard & Poors coincided with ongoing fiscal weakness in several Eurozone countries and the initiation of the Occupy Wall Street movement to produce a sharp, but ultimately transient, spike in the probability of the low-liquidity (high price-impact) state. Similarly, the liquidity crisis after the failure of Lehman Brothers is plainly visible as the deep and more persistent spike in September 2008, preceded by a series of pronounced foreshocks over the course of the year.

Figure 10 shows stacked "ribbons" of daily data, using the same (red, blue and yellow) color to indicate relative levels (with black for missing data), where series are grouped vertically by instrument type. Figures 10 and 9 illustrate that the equity markets and VIX®

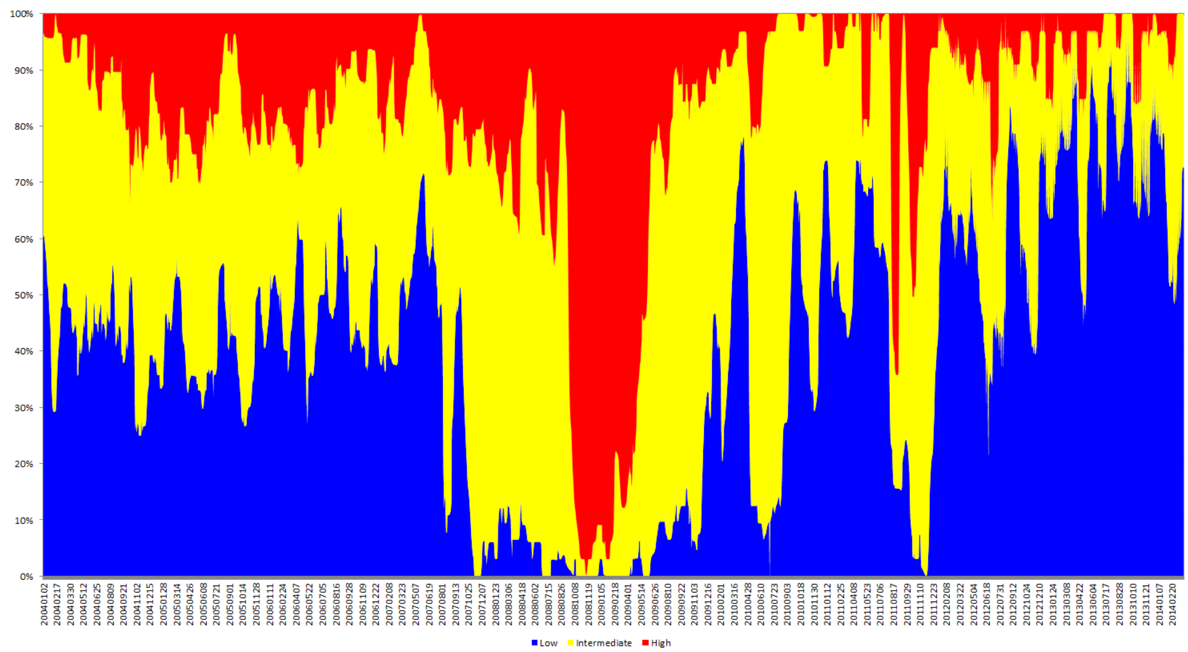

Figure 8: Probabilities of price-impact states, averaged across 33 markets, Apr. 2004 – Mar. 2014
Sources: CRSP, Bloomberg, Mergent, WRDS, FINRA, OFR Analysis

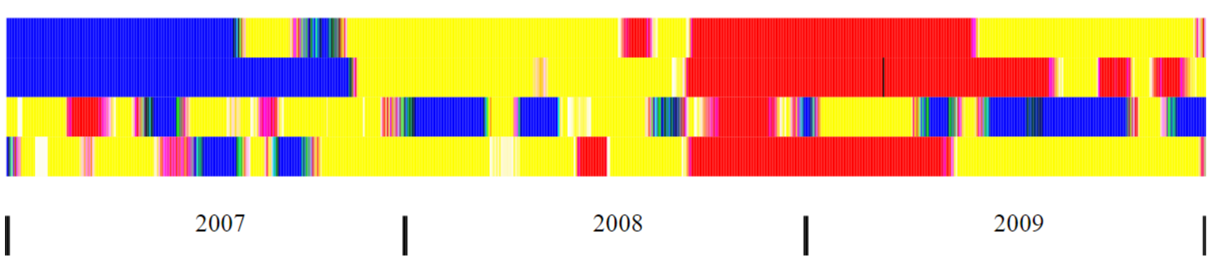

Figure 9: Daily Price-Impact Probabilities for Select Markets,
Top to bottom: Equities SIC6, Bonds SIC6, WTI futures, VIX® futures
Sources: CRSP, Bloomberg, Mergent, WRDS, FINRA, OFR Analysis

index responded strongly and immediately to the run on the repo, but WTI futures did not. Through the crisis, VIX® liquidity remained relatively, but in contrast the liquidity of the equity market largely recovered with the exception of the financial and service sectors. Consistent with the increased uncertainty about the stocks in the financial sector, throughout late 2007 and 2008, remained depressed as we would expect, prior to the crisis liquidity in financial stocks remained depressed. Two other key insights from this analysis are that the liquidity implications of the Lehman Brothers failure were felt broadly for an extended period and that hints of illiquidity foreshocks existed in some markets, including financial stocks (SIC 6) and certain bond sectors, that may ultimately help in crafting liquidity forecasts.

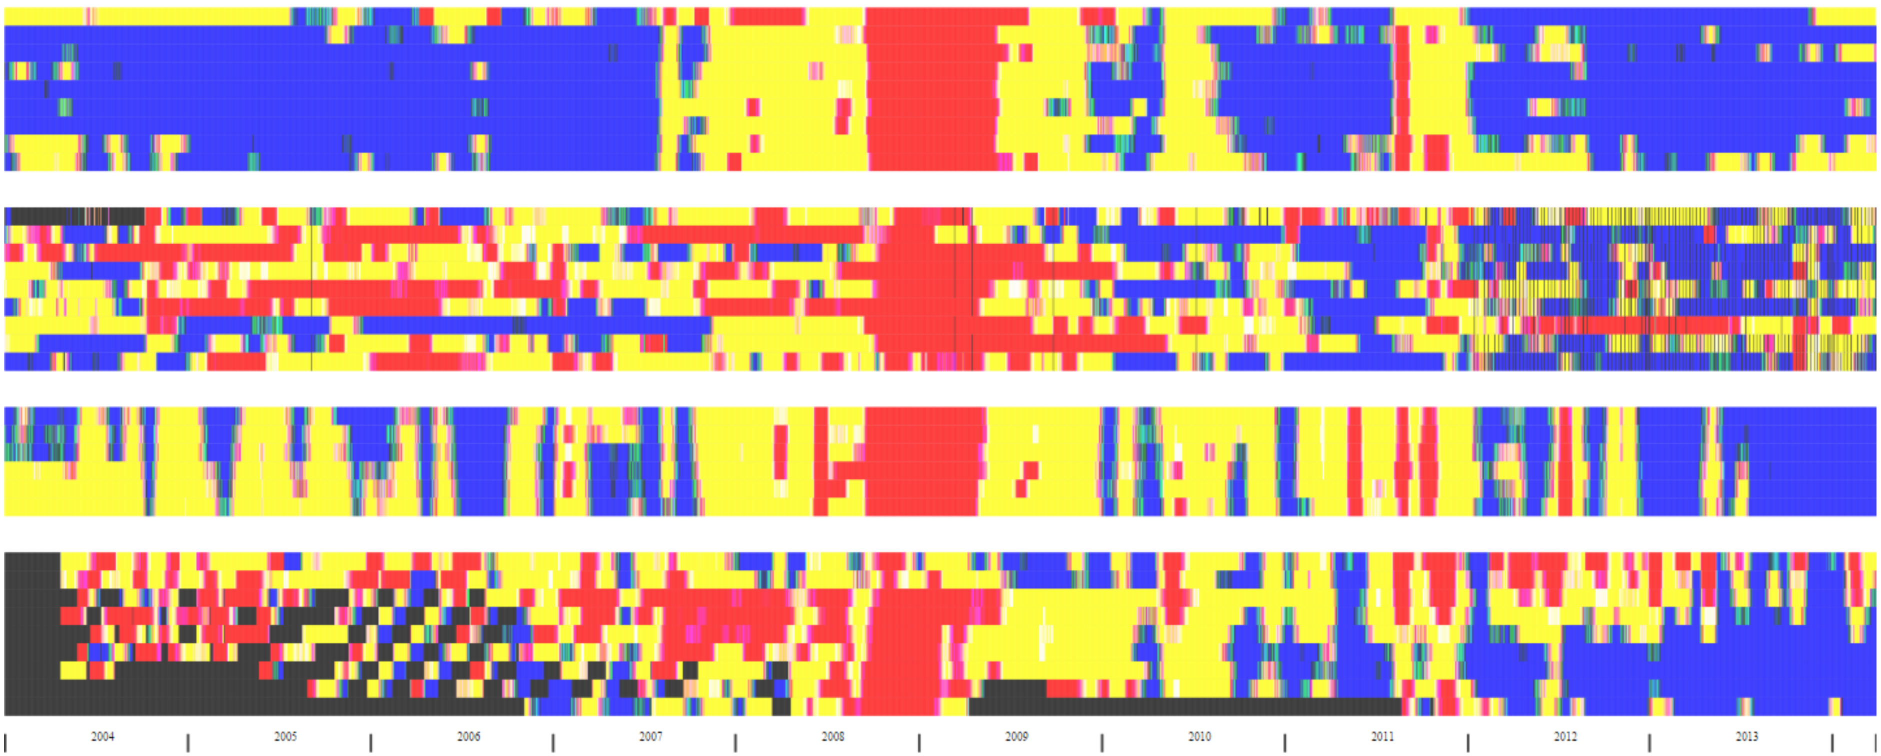

Figure 10: Daily Price-Impact Probabilities across all 33 Markets,
Top to bottom: Equities SIC0-8, Bonds SIC0-8, WTI futures, VIX® futures
Sources: CRSP, Bloomberg, Mergent, WRDS, FINRA, OFR Analysis

4.2 Explaining Liquidity Regimes

There appear to be strong relationships between changes in the level of liquidity and a number of macro variables, and while it is helpful to explore these relationships graphically the hierarchical model allows us to determine whether these relationships are statistically significant, particularly in the presence of other competing macro variables. We restrict our analysis of liquidity dynamics across multiple markets to the US Equity markets. We did this in part because of data consistency issues (there was no missing price-impact data for the U.S. Equity markets over the period of interest) and because these markets exhibit somewhat consistent behavior, see Figure 10. After visually exploring a range of potential macro variables, we selected the 11 macro variables describe in Table 1.

[Insert Table 1 about here]

We test the ability of the macro variables to recover the liquidity dynamics across these markets in two ways. First we calculate a hit rate, which is the proportion of the time that the Probit model, based solely on the macro variables, accurately predicts the state identified by each of the underlying univariate models (or we count the proportion of time that we accurately predict the state of D_{it} for each i and t using the current estimate of β, Σ and the macro data x_{it}). The naive hit rate is 33%, assuming random guessing, and the posterior average of Probit model's hit rate was 66% indicating that the macro variables are explaining a substantial portion of the liquidity dynamics. Second, we plotted the predicted probability of being in each state for each time point, using the Probit model, against the average probability of being in each state for each time point, see Figure 11. The way the predicted probabilities closely tracks the average probabilities confirm, again, the ability of the macro variables to explain the liquidity dynamics.

We standardized the macro variables (mean centered and divided by the standard deviation), to compare the parameter estimates from the Probit portion of the hierarchical

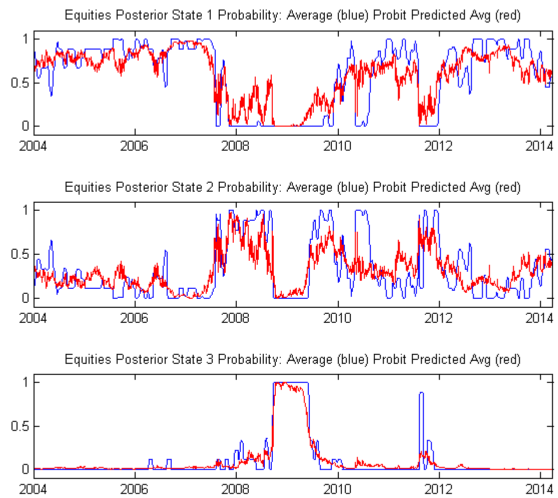

Figure 11: Average State Probabilities vs. Probit Predicted Probabilities
Sources: CRSP, WRDS, Bloomberg, OFR Analysis

model, seen in Table 2, directly with respect to their size. Because we force the latent value for state 1 to always be zero, to address the additive identification restriction, we get only get parameter estimates for states 2 and 3 (which are really the difference between the unrestricted parameters of each of these states relative to state 1). The negative intercepts indicate that state 1, the low price-impact or high-liquidity state, is the most prevalent state and the fact that the intercept for state 3 is more negative than for state 2 indicates that state 3, the low-liquidity state, is the least likely state.

[Insert Table 2 about here]

All but one of the macro variables are statistically significant (only the U.S. 5-year Break-even Inflation is not significant in distinguishing the high-liquidity state from the mid-liquidity state). Within these results, there are some interesting patterns to note. First, there is a natural grouping among the macro variable with regards to the pattern of the signs for the state 2 and 3 parameter estimates. As might be expected, VIX® has a positive-positive pattern indicating that higher levels of VIX® are associated with a higher probability of entering states with low liquidity. The persistence of high levels of VIX® after the crisis makes VIX® more strongly related to the middle liquidity state as opposed to the crisis state. Another group of five macro variables (WTI, 3m Repo Rate, S&P 500 P/B Ratio, TED Spread and Yield Curve (10y–2y)) exhibit a positive-negative pattern indicating that elevated levels of these macro variables lead to a high probability of being in the middle liquidity state. Although this pattern may seem counter intuitive, we observe that market movements, government actions and actions by central banks distorts and lowers these macro variables during times of crisis, which corresponds to times of low liquidity (*i.e.* state 3). The next set of macro variables (Dow Jones Real Estate Index, Moody's Baa Index, and LIBOR-OIS Spread) exhibit a negative-positive pattern which identifies them as measurements which have persistently high levels during times of low liquidity and then "bounce back" sufficiently during the moderate times of low liquidity to have the extremes associated with the crisis. Finally, the DXY Dollar Index shows a significant negative-negative pattern, indicating that higher levels of the dollar are associated with higher probability of entering a high-liquidity state. This is consistent with a flight to quality, in which capital flows into the U.S. during episodes of stress, such as the crisis, simultaneously pushing up the value of the dollar and flooding the domestic market with liquidity.

To assist in our understanding of these parameter estimates, we can compare the time-series plot for individual macro variables versus the Probit-predicted probability for each state. For example, Figure 12 presents this comparison for the TED spread. The spread remains low until mid-2007 and returns to persistent low levels in 2010. The early episode corresponds to consistently high probabilities for the high-liquidity state (i.e., state 1), consistent with the TED spread's role as a bellwether for funding liquidity. Between August 2007 and September 2008, when the TED spread begins to widen, the probability of state 2 jumps, supporting the positive coefficient in Table 2. After September 2008,

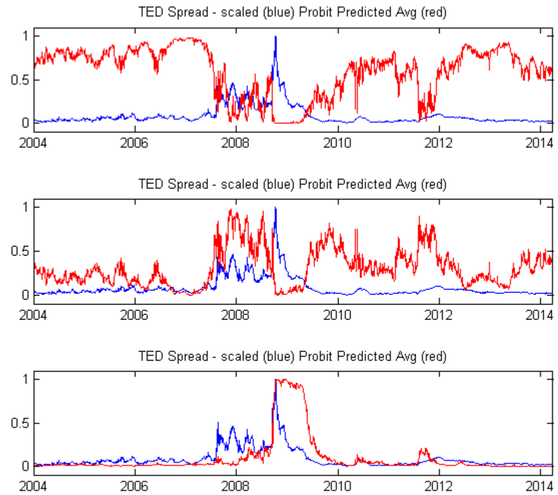

Figure 12: TED Spread and Average Probit Probabilities
Sources: CRSP, WRDS, Bloomberg, OFR Analysis

the TED spread recedes relatively quickly from its peak, compared with the probability of being in state 3, which remains elevated for the next year. This deviation is consistent with the negative coefficient on state 3 in Table 2. In contrast, the VIX® index has positive coefficients on both states 2 and 3 in Table 2. The VIX® is more persistently high after the 2008 shock, consistent with the positive coefficient on state 3 in Table 2. Moreover, it remains moderately elevated for much of the post-crisis period after 2009, when the Probit-predicted probability for state 2 is also raised.

The next four (Moody's Baa Index, VIX®, LIBOR and WTI) have a negative, positive pattern and are clear predictors of periods of low-liquidity. The VIX® has the strongest parameter estimate in absolute value and when the VIX® is high, the probability of being in state 2 drops and the probability of being in state 3 rises dramatically as indicated by plotting VIX® against the predicted probabilities in Figure 13.

The final macro variable (S&P 500 P/B Ratio) has a positive, positive pattern (although only the parameter for state 3 is statistically significant). This gives us the insight that as the price of equities becomes large relative to the underlying book value of the firms,

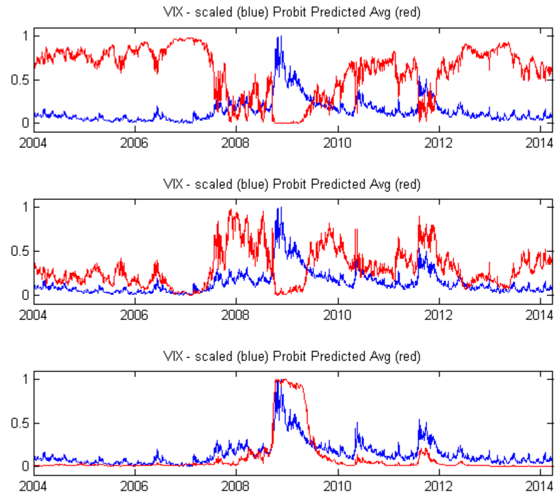

Figure 13: VIX® and Average Probit Probabilities
Sources: CRSP, WRDS, Bloomberg, OFR Analysis

we tend to be in state 3, the low-liquidity state. One possible explanation for this is that a high price to book ratio reflects a potential asset bubble, which could lead investors to engage in herd behavior (e.g., piling into different individual stocks and driving up returns then pulling out suddenly causing large price drops), which could cause not only increased volatility but also larger price-impacts.

Clearly the hierarchical model allows us to obtain interesting insights into how macro variables relate to liquidity dynamics and offers a valuable tool for further investigating and understanding the drivers of liquidity across a wide range of markets.

5 Conclusion

Liquidity is an elusive, yet essential component of the modern financial system. It is elusive because conceptually it is hard to define, and empirically it is hard to measure and predict. More specifically, we attribute the challenges in liquidity measurement to three fundamental aspects of the phenomenon. Liquidity is latent, in the sense that the episodes of illiquidity we seek to understand are rare, and often emerge with little apparent warning. Liquidity is non-linear, in the sense that price impact does not respond proportionately to additional order flow, making it difficult to extrapolate from "ordinary" markets to the behavior of those markets under stress. Liquidty is endogenous, in the sense that it often emerges as a positive externality in very active markets, making those busy venues attractive to others who seek the assurance that counterparties will be available when needed.

We address the challenges of latency, non-linearity and endogeneity statistically with a Bayesian estimation of a hidden Markov chain individually for 33 separate time series covering the CRSP and TRACE universes of U.S. equities and corporate bonds, plus multiple expiries of two key futures contracts, the VIX® volatility contract and the WTI oil contract. Three latent states (high, medium, and low price impact) are adequate to capture the observed liquidity structure of all 33 univariate series.

We also look for cross-sectional structure in the data by estimating a hierarchical Bayesian model, and testing the ability of several macroeconomic time series to recover the estimated aggregate liquidity dynamics. This exercise also permits an attribution of those estimated aggregate dynamics to meaningful economic interpretations. For reasons of data consistency, we have limited our initial efforts in this area to the U.S. equities markets.

Our results at this stage are very preliminary, but also very promising. In addition to testing for robustness and sensitivity, we see several immediate avenues for future research, including expanding the cross section of asset markets in the scope of analysis, comparing in more detail the liquidity behavior of wholesale funding markets, and experimenting wth alternative portfolio formation rules.

References

References

T. Adrian and H. S. Shin. Liquidity and leverage. *Journal of Financial Intermediation*, 19(3):418–437, July 2010. URL http://www.sciencedirect.com/science/article/pii/S1042957308000764.

T. Adrian, P. Colla, and H. S. Shin. Which financial frictions? parsing the evidence from the financial crisis of 2007-9. *NBER Macroeconomics Annual*, 27:159–214, May 2013a. URL http://www.princeton.edu/~hsshin/www/whichfinancialfrictions.pdf.

T. Adrian, D. Covitz, and N. Liang. Financial stability monitoring. Finance and Economics Discussion Series 2013-21, Board of Governors of the Federal Reserve, 2013b. URL http://www.federalreserve.gov/pubs/feds/2013/201321/201321pap.pdf.

Y. Amihud. Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets*, 5:31–56, 2002. URL https://umdrive.memphis.edu/pjain/Microstructure/Papers/Illiquidity%20and%20stock%20returns_Amihud.pdf.

A. Ashcraft, M. L. Bech, and W. S. Frame. The federal home loan bank system: The lender of next-to-last resort? *Journal of Money, Credit and Banking*, 42(4):551–583, June 2010. doi: 10.1111/j.1538-4616.2010.00299.x. URL http://onlinelibrary.wiley.com/doi/10.1111/j.1538-4616.2010.00299.x/abstract.

W. Bagehot. *Lombard Street: A Description of the Money Market*. H.S. King, 1873.

Bank of England. Instruments of macroprudential policy. Discussion paper, Bank of England, 2011. URL http://www.bankofengland.co.uk/publications/other/financialstability/discussionpaper111220.pdf.

M. J. Barclay, T. Hendershott, and K. Kotz. Automation versus intermediation: Evidence from treasuries going off the run. *Journal of Finance*, 61(5):2395–2414, October 2006. ISSN 00221082. URL http://www.jstor.org/stable/3874714.

Basel Committee on Banking Supervision. Basel iii: International framework for liquidity risk measurement, standards and monitoring. Technical Report 188, BIS, December 2010. URL http://www.bis.org/publ/bcbs188.htm.

Basel Committee on Banking Supervision. Basel iii: The liquidity coverage ratio and liquidity risk monitoring tools. Technical report, BIS, 2013. URL http://www.bis.org/publ/bcbs238.htm.

H. Bessembinder, W. Maxwell, and K. Venkataraman. Market transparency, liquidity externalities, and institutional trading costs in corporate bonds. *Journal of Financial Economics*, 82(2):251–288, November 2006. URL http://www.sciencedirect.com/science/article/pii/S0304405X06000699.

J.-P. Bouchaud, J. D. Farmer, and F. Lillo. How markets slowly digest changes in supply and demand. In T. Hens and K. Schenk-Hoppe, editors, *Handbook of Financial Markets: Dynamics and Evolution*, page 57156. Academic Press, 2008. URL http://arxiv.org/pdf/0809.0822.pdf.

K. Boudt, E. C. Paulus, and D. W. Rosenthal. Funding liquidity, market liquidity and ted spread: A two-regime model. Working Paper 244, National Bank of Belgium, November 2013. URL http://tigger.uic.edu/~daler/funding-market-liquidity-poster-ec2.pdf.

N. M. Boyson, C. W. Stahel, and R. M. Stulz. Hedge fund contagion and liquidity shocks. *Journal of Finance*, 65(5):1789–1816, October 2010. URL http://www.cob.ohio-state.edu/fin/faculty/stulz/publishedpapers/jofi1594.pdf.

M. Brunnermeier and L. Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238, June 2009. URL http://rfs.oxfordjournals.org/content/22/6/2201.abstract.

M. K. Brunnermeier. Deciphering the liquidity and credit crunch 2007-2008. *Journal of Economic Perspectives*, 23(1):77–100, Winter 2009. URL http://ideas.repec.org/a/aea/jecper/v23y2009i1p77-100.html.

M. K. Brunnermeier and Y. Sannikov. Macroeconomics with a financial sector. Working Paper 236, National Bank of Belgium, October 2012. URL file:///C:/Users/floodm/Downloads/SSRN-id2160894.pdf.

C. Calomiris and G. Gorton. The origins of banking panics: Models, facts, and bank regulation. In R. G. Hubbard, editor, *Financial Markets and Financial Crises*, chapter 4, pages 109–174. NBER and University of Chicago, 1991. URL http://www.nber.org/chapters/c11484.pdf.

H. Chen, V. Cúrdia, and A. Ferrero. The macroeconomic effects of large-scale asset purchase programs. *Economic Journal*, 122(564):289–315, November 2012a. URL http://www.newyorkfed.org/research/staff_reports/sr527.html.

S. Chen, P. Liu, A. Maechler, C. Marsh, S. Saksonovs, and H. S. Shin. Exploring the dynamics of global liquidity. Technical report, IMF, 2012b. URL http://www.imf.org/external/pubs/ft/wp/2012/wp12246.pdf.

T. Chordia, R. Roll, and A. Subrahmanyam. Commonality in liquidity. *Journal of Financial Economics*, 56(1):3–28, 2000. URL http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.139.1722&rep=rep1&type=pdf.

M. M. Cornett, J. J. McNutt, P. E. Strahan, and H. Tehranian. Liquidity risk management and credit supply in the financial crisis. *Journal of Financial Economics*, 101(2):297–312, August 2011. URL http://www.sciencedirect.com/science/article/pii/S0304405X11000663.

D. Cumming, S. Johan, and D. Li. Exchange trading rules and stock market liquidity. *Journal of Financial Economics*, 99(3):651–671, March 2011. URL http://www.sciencedirect.com/science/article/pii/S0304405X10002370.

S. D'Amico and T. B. King. Flow and stock effects of large-scale treasury purchases: evidence on the importance of local supply. *Journal of Financial Economics*, 108(2):425–448, May 2012. URL http://www.sciencedirect.com/science/article/pii/S0304405X12002425.

S. D'Amico, W. English, D. López-Salido, and E. Nelson. The federal reserve's large-scale asset purchase programs: Rationale and effects. *Economic Journal*, 122(564):415–46, November 2012.

T. V. Dang, G. Gorton, and B. Holmström. Financial crises and the optimality of debt for liquidity provision. Technical report, Yale U., May 2010. URL http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.183.7960&rep=rep1&type=pdf.

T. V. Dang, G. Gorton, and B. Holmström. Ignorance, debt and financial crises. Technical report, Yale U., February 2012. URL http://www.econ.cuhk.edu.hk/dept/chinese/seminar/12-13/1st-term/Paper_Ignorance.pdf.

J. Dick-Nielsen. Liquidity biases in trace. *Journal of Fixed Income*, 19(2):43–55, Fall 2009. doi: DOI: 10.3905/jfi.2009.19.2.043. URL http://www.iijournals.com/doi/abs/10.3905/jfi.2009.19.2.043.

D. Duffie. Presidential address: Asset price dynamics with slow-moving capital. *Journal of Finance*, 65(4):1237–1267, 2010. URL http://www.darrellduffie.com/uploads/pubs/DuffieAFAPresidentialAddress2010.pdf.

N. Economides. The economics of networks. *International journal of industrial organization*, 14(6):673–699, October 1996. URL http://www.sciencedirect.com/science/article/pii/0167718796010156.

D. J. Elliott, G. Feldberg, and A. Lehnert. The history of cyclical macroprudential policy in the united states. Working Paper 0008, OFR, May 2013. URL http://www.treasury.gov/initiatives/ofr/research/Documents/OFRwp0008_ElliottFeldbergLehnert_AmericanCyclicalMacroprudentialPolicy.pdf.

A. J. Entwistle and J. H. Beemer. Approaches to asset recovery for pension fund subprime exposure. Technical report, Entwistle & Cappucci, February 2008. URL http://www.entwistle-law.com/news/publications/000041.

E. Farhi and J. Tirole. Leverage and the central banker's put. *The American Economic Review*, 99(2): 589–593, May 2009. ISSN 00028282. URL http://www.jstor.org/stable/25592463. Papers and Proceedings of the One Hundred Twenty-First Meeting of the American Economic Association.

FINRA. Historic trace data: Enhanced historic time and sales trade record file layout. Technical report, FINRA, 2009. URL http://www.finra.org/web/groups/industry/@ip/@comp/@mt/documents/appsupportdocs/p120989.pdf.

X. Gabaix, P. Gopikrishnan, V. Plerou, and H. E. Stanley. Institutional investors and stock market volatility. *Quarterly Journal of Economics*, 121(2):461–504, 2006. URL http://qje.oxfordjournals.org/content/121/2/461.abstract.

A. Gabrielsen, M. Marzo, and P. Zagaglia. Measuring market liquidity: An introductory survey. Working Paper DSE 802, U. di Bologna, 2011. URL http://arxiv.org/pdf/1112.6169.pdf.

J. Geanakoplos. Liquidity, default, and crashes: Endogenous contracts in general equilibrium. In M. Dewatripont, L. P. Hansen, and S. J. Turnovsky, editors, *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress, Volume II*, 2003.

D. Gefang, G. Koop, and S. M. Potter. Understanding liquidity and credit risks in the financial crisis ? *Journal of Empirical Finance*, 18(5):903–914, December 2011. URL http://www.sciencedirect.com/science/article/pii/S0927539811000508.

C. A. Goodhart. The background to the 2007 financial crisis. *International Economics and Economic Policy*, 4(4):331–346, February 2008. URL http://link.springer.com/article/10.1007%2Fs10368-007-0098-0?LI=true.

A. Hameed, W. Kang, and S. Viswanathan. Stock market declines and liquidity. *Journal of Finance*, 65(1): 257–293, 2010. URL http://onlinelibrary.wiley.com/doi/10.1111/j.1540-6261.2009.01529.x/full.

J. Hasbrouck and D. J. Seppi. Common factors in prices, order flows, and liquidity. *Journal of financial Economics*, 59(3):383–411, 2001. URL http://www.sciencedirect.com/science/article/pii/S0304405X0000091X.

B. Holmström. The nature of liquidity provision: When ignorance is bliss. Presidential address, Econometric Society, 2012. URL http://economics.mit.edu/files/7500.

B. Holmström and J. Tirole. *Inside and Outside Liquidity*. MIT Press, 2013. URL http://www.amazon.com/dp/0262518538/.

G. A. Karolyi, K.-H. Lee, and M. A. van Dijk. Understanding commonality in liquidity around the world. *Journal of Financial Economics*, 105:82–112, 2012. URL http://www.sciencedirect.com/science/article/pii/S0304405X11002844.

C. P. Kindleberger. *A Financial History of Western Europe*. Oxford University Press, 2 edition, 1993.

A. S. Kyle and A. A. Obizhaeva. Market microstructure invariants: Empirical evidence from portfolio transitions. working paper, University of Maryland, December 2011a. URL http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1978943.

A. S. Kyle and A. A. Obizhaeva. Market microstructure invariants: Theory and implications of calibration. working paper, University of Maryland, December 2011b. URL http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1978932.

A. S. Kyle and A. A. Obizhaeva. Large bets and stock market crashes. working paper, University of Maryland, May 2012. URL http://dl.dropbox.com/u/2798811/paper_ko_crashes_bets.pdf.

A. S. Kyle and A. A. Obizhaeva. Market microstructure invariants: Theory and empirical tests. Working paper, University of Maryland, October 2014. URL http://www.rhsmith.umd.edu/faculty/obizhaeva/Kyle_Obizhaeva_20130607.pdf.

M. Marra. Illiquidity commonality across equity and credit markets. working paper, University of Reading, April 2013. URL http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2019512.

P. Martin. Analysis of the impact of competitive rates on the liquidity of nyse stocks and technical memorandum: a nonstationary stochastic process model for aggregate stock price changes. Economic Staff Paper 75-3, Securities and Exchange Commission, July 1975.

R. McCauley. Benchmark tipping in the money and bond markets. *BIS Quarterly Review*, pages 39–45, March 2001. URL http://www.bis.org/publ/r_qt0103.htm.

P. Mehrling. *The New Lombard Street: How the Fed Became the Dealer of Last Resort*. Princeton University Press, November 2010. URL http://www.amazon.com/dp/0691143986.

F.-L. Michaud and C. Upper. What drives interbank rates? evidence from the libor panel. *BIS Quarterly Review*, pages 47–58, March 2008. URL http://www.bis.org/publ/qtrpdf/r_qt0803f.htm.

S. Morris and H. S. Shin. Liquidity black holes. *Review of Finance*, 8(1):1–18, 2004. URL http://cowles.econ.yale.edu/P/cp/p11a/p1114.pdf.

H. G. Moulton. Commercial banking and capital formation. *Journal of Political Economy*, 26(5, 6, 7, 9): 484–508, 638–663, 705–731, 849–881, 1918.

M. Pagano. Trading volume and asset liquidity. *The Quarterly Journal of Economics*, 104(2):255–274, May 1989. URL http://qje.oxfordjournals.org/content/104/2/255.short.

C. M. Reinhart and K. Rogoff. *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press, September 2009. URL http://press.princeton.edu/titles/8973.html.

M. Schularick and A. M. Taylor. Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870–2008. *The American Economic Review*, 102(2):1029–1061, April 2012. ISSN 00028282. URL http://www.jstor.org/stable/23245443.

S. Sundaresan and Z. Wang. Y2k options and the liquidity premium in treasury bond markets. Staff Report 266, Federal Reserve Bank of New York, November 2006. URL http://www.fednewyork.org/research/staff_reports/sr266.pdf.

J. B. Taylor and J. C. Williams. A black swan in the money market. *American Economic Journal: Macroeconomics*, 1:58–83, 2009.

J. Tirole. Illiquidity and all its friends. *Journal of Economic Literature*, 49(2):287–325, 2011. URL http://idei.fr/doc/by/tirole/illiquidity-may12_11.pdf.

B. Toth, Y. Lemperiere, C. Deremble, J. De Lataillade, J. Kockelkoren, and J.-P. Bouchaud. Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X*, 1(2):021006/1–11, 2011. URL http://prx.aps.org/pdf/PRX/v1/i2/e021006.

Tables

Table 1: Macro Variables for Hierarchical (Multiple Market) Model

Macro Variable	Description
Three month Repo Rate	ICAP General Collateral Treasury 3-month repurchase agreement rate
Yield Curve	Yield on the constant maturity 10-year U.S. Treasury bond minus the yield on the constant maturity 2-year U.S. Treasury note
TED Spread	3-month LIBOR rate minus the 3-month U.S. Treasury bill yield
Moody's Baa Corporate Bond Index	Yield on the Moody's investment grade long-term corporate bond index
VIX® Index	Reflects the market estimate of future (30-day) volatility of the S&P 500
Dow Jones U.S. Real Estate Index	Index representing real estate investment trusts (REITs) and other companies investing directly or indirectly in real estate through development, management, or ownership
S&P 500 Price-to-book Ratio	Ratio of equity market value to book value per share of the S&P 500
Three-month LIBOR-OIS Spread	Difference between the 3-month LIBOR and the 3-month USD overnight index swap (OIS) rate
5-year U.S. Breakeven Inflation Rate	Calculated by subtracting the real yield of the 5-year inflation-linked maturity curve from the yield of the closest 5-year nominal Treasury maturity. The result is the market-implied inflation expectation over the next 5 years
WTI Front-Month Price	Futures price for the near-dated expiry of the WTI oil contract
DXY Dollar Index	Indicates the general international value of the USD, by averaging exchange of the USD against other major currencies

Table 2: Posterior Parameter Estimates Probit Portion of Hierarchical Model (Sources: CRSP, WRDS, Bloomberg, OFR Analysis)

	Posterior Mean		**Posterior StDev**		**t-Stat**	
Macro Variable	**State 2**	**State 3**	**State 2**	**State 3**	**State 2**	**State 3**
Intercept	-0.64**	-1.01**	0.02	0.01	-40.27	-102.11
VIX®	0.62**	0.26**	0.03	0.02	23.77	15.67
WTI	0.83**	-0.23**	0.03	0.01	30.65	-16.53
3m Repo Rate	0.68**	-0.41**	0.02	0.01	38.28	-49.12
TED Spread	0.49**	-0.09**	0.03	0.01	18.53	-10.98
Yield Curve (10y–2y)	0.19**	-0.38**	0.02	0.04	9.66	-56.07
S&P 500 P/B Ratio	0.68**	-0.13**	0.03	0.01	23.85	-9.86
Dow Jones Real Estate Index	-1.17**	0.13**	0.02	0.01	-64.48	11.69
Moody's Baa Index	-0.67**	0.47**	0.02	0.01	-43.57	60.83
LIBOR–OIS Spread	-0.64**	0.13**	0.05	0.02	-13.07	7.38
DXY Dollar Index	-0.39**	-0.37**	0.04	0.01	-10.67	-41.05
U.S. 5y Breakeven Inflation	-0.03	0.00	0.02	0.01	-1.43	0.36

** Significant at a 99% confidence level