Liquidity Supply and Demand in the Corporate Bond Market

Jonathan Goldberg and Yoshio Nozawa*

Federal Reserve Board

January 24, 2018

Abstract

We propose a novel framework to measure liquidity supply in the corporate bond market by jointly analyzing the price and quantity of liquidity. Motivated by theories of segmented markets, we construct the "noise" in corporate bonds as a liquidity price measure, and gross dealer positions on corporate bonds as our quantity measure. Liquidity supply and demand are identified using structural VARs based on price and quantities, together with simple sign restrictions. We find that the elasticity of liquidity supply fell very sharply during the financial crisis and remains below 50 percent of its pre-crisis level, suggesting that the corporate bond liquidity has become more fragile. Our supply-demand decomposition shows that liquidity demand shocks explain slightly less than half of time-variation in aggregate noise and dealer positions. The estimated liquidity supply shocks are negatively associated with uncertainty shocks, whereas liquidity demand shocks are associated with mutual fund flows. These results underscore the importance of separating liquidity supply effects from demand effects when assessing the liquidity condition.

JEL Classification: G12, G24

Keywords: Corporate bond liquidity, dealer intermediation, liquidity supply, post-crisis liquidity

*Federal Reserve Board. We thank Yesol Huh, Alex Zhou and seminar participants at Federal Reserve Board and Penn State University for their comments. The views expressed here are those of the authors and need not represent the views of the Federal Reserve Board or its staff. Contact: yoshio.nozawa@frb.gov

1 Introduction

How has the supply of liquidity in the corporate bond market evolved since the pre-crisis period? And have there been changes in the fragility of the corporate bond market (i.e., how adversely is liquidity affected by a spike in liquidity demand)? Answering these questions is difficult because it is hard to measure the supply of liquidity. For example, since the pre-crisis period, measures of turnover, average trade size, block trade frequency and dealer capital commitment have declined (Bessembinder et al. (2017)). On the other hand, measures of trading costs have declined as well (e.g., Trebbi and Xiao (2015)). Because lower transaction costs may reflect higher supply or lower demand for liquidity in the market, this evidence does not readily permit one to infer whether the supply of corporate bond liquidity has declined, increased or remained the same.

A key contribution of this paper is to develop and implement a new method to measure liquidity supply in the corporate bond market by jointly analyzing the price and quantity of liquidity. The empirical measurement and identification assumptions are based on a simple model of intermediation in segmented markets based on Gromb and Vayanos (2002), which have been applied in particular to yield-curve models with preferred-habitat investors (Vayanos and Vila (2009) and Malkhozov et al. (2016)). In this framework, we focus on how dealers absorb idiosyncratic demand imbalances for bonds with certain maturity. This type of intermediation does not require dealers to take issuers' credit risk. Thus, dealers do not necessarily take a net long or net short position, but they do take gross corporate bond positions (gross long and gross short positions). We view the bond market as being segmented across maturity, and thus idiosyncratic demand imbalances for a bond lead to noise in corporate bond yield curves. The dealers are risk averse and face liquidity risk, and thus in general they do not eliminate noise completely. We show that when dealers are more risk tolerant, their gross positions increase and noise in corporate bond yields decreases. Therefore, we identify changes in liquidity supply as a shock that leads to a decline in noise (i.e. an decrease in price of liquidity) and an increase in gross dealer corporate bond positions (i.e. an increase in quantity of liquidity).

We focus on primary dealers as suppliers of liquidity in the corporate bond market.

Corporate bonds are traded in over-the-counter market and traditionally these dealers have used their balance sheets to absorb temporary order imbalances in the market by purchasing (at a discount) when an investor wants to sell a particular security and selling (at a premium) when an investor wants to buy a particular security. However, the recent introduction of new banking regulations reportedly hampered the provision of liquidity by these dealers, and thus quantifying liquidity supply by these intermediaries is important for policy makers and researchers.

As a first step to identify liquidity supply, we construct the price and quantity of corporate bond liquidity as suggested by theories of segmented markets. We estimate noise in corporate-bond yield curves and calculate dealer gross positions using the micro-level data. To measure the noise in corporate bond yields, we fit the Nelson-Siegel-Svensson (1994) model at the issuer-level for large bond issuers, and compute the root mean-squared error for each issuer. Every week, we average the errors across issuers to obtain an aggregate time series of noise. To measure gross dealer positions in corporate bonds for primary dealers, we use the regulatory version of TRACE made available to us by FINRA. We then accumulate the transactions for each primary dealer for each bond to measure gross inventory of each dealer in each bond. The quantity of liquidity is the sum of gross positions across bonds and dealers.

Next, we estimate a structural VAR model using liquidity price and quantity to identify liquidity supply and demand. The identification assumption is that changes in liquidity supply result in contemporaneous opposite-signed changes in dealer gross positions and noise, while changes in liquidity demand result in changes with the same sign for noise and dealer positions. Our estimation technique takes advantage of the development in the literature in macroeconomics which decomposes supply and demand effects by jointly analyzing the price and quantity (e.g. Uhlig (2005) for money supply shocks, and Kilian and Murphy (2012) for oil supply shocks). An advantage of our approach is that we do not rely on potentially noisy proxies for liquidity supply shocks, such as leverage and CDS spreads of dealers. Instead, our identification relies solely on an economic theory of financial intermediation which provides intuitive sign restrictions on the price and quantity measures.

The liquidity supply shocks inferred from the model capture notable episodes of stress in the corporate bond market, with liquidity supply declining following the failures of Bear Stearns and Lehman Brothers as well as during the European fiscal crisis in 2010 and the taper tantrum in 2013. Liquidity demand increased sharply immediately after major financial crisis stress events, and generally increased in the years between the financial crisis and 2014. However, in recent years, liquidity demand shocks have been negative, which partly explains why estimated transaction costs in recent years are so low. Using the structural VAR, we run a counterfactual experiment in which there are no demand shocks since July 2010, when a series of banking regulations started to be introduced. In the experiment, we show that the noise in corporate bond yields would have been about 10 basis points higher if there had been no demand shocks during this period.

Our methodology enables us to decompose the variance of liquidity measures into the supply and demand components. We find that both supply and demand shocks are important for explaining variation in the price and quantity of liquidity; in particular, at short horizons, supply shocks explain about two-thirds of the time variation in noise, with demand shocks explaining the other one-third. For longer horizons, both supply and demand shocks explain about half of the variation in noise and the size of the dealer balance sheet. For example, an increase in the price of liquidity in 2011, amid the European fiscal crisis, was largely driven by a fall in liquidity supply. These results highlight the importance of taking into account fluctuations in demand for liquidity when assessing how the supply of liquidity has changed over a given period.

Using our estimates of the structural VAR model, we show that the elasticity of liquidity supply after the financial crisis is less than half of the elasticity during the pre-crisis period. This suggests that, for a shock that shifts the demand curve, we would see a change in price of liquidity twice as large in the post crisis period as in the pre-crisis period. Therefore, even though the level of transaction costs have been low in recent years, liquidity risk has increased as dealers' propensity to enlarge their balance sheets to absorb unexpected liquidity demand has declined.

To verify the identified liquidity supply and demand shocks, we investigate the relationship between the shocks and other commonly cited measures of liquidity. We find that liquidity supply shocks are strongly negatively correlated with shocks to swaption implied volatility, a measure for uncertainty in the fixed income markets. Furthermore, we find that demand shocks are positively related to absolute values of corporate bond mutual fund flows, though correlation is low. Therefore, the supply and demand shocks identified by our framework are indeed associated with commonly cited drivers of bond liquidity, though the proper identification requires the information about liquidity price and quantity.

Our methodology provides an ideal setup to test other interesting research questions, such as liquidity contagion. To understand how a shock to liquidity demand in one part of the market affects the price of liquidity in other parts of the market, we study two subsamples of corporate bonds: one for investment grade (IG) bonds and the other for high yield (HY) bonds. Gromb and Vayanos (2002) predict that, if a dealer loses money in one market, it may reduce its liquidity provision in the other market due to the collateral constraint. We test this hypothesis by estimating trivariate VARs including the one with IG liquidity price, IG liquidity quantity and HY liquidity price, and the other with HY liquidity price, HY liquidity quantity and IG liquidity price. We find that an unexpected rise in noise in the HY market leads to an modest increase in noise in IG market, while an unexpected rise in noise in the IG market leads to a significant rise in noise for HY bonds. Thus, a shock to a part of the bond market affects the liquidity in the other part of the market through dealers' balance sheet.

As another extension of the framework, we study the mechanism through which aggregate fluctuations in corporate bond illiquidity affect aggregate credit spreads. We estimate a structural VAR using noise, dealer gross positions and corporate credit spreads, to study whether a demand-driven increase in illiquidity affects credit spreads differently than a supply-driven one. We find that an unexpected decline in liquidity supply leads to a significant increase in credit spreads over the 4 to 6 months horizon that we study. In contrast, a rise in noise driven by increased demand does not predict changes in credit spreads. Thus, the liquidity premium in corporate credit spreads comes primarily from the liquidity supply channel rather than the demand.

Our empirical analysis is motivated by theories of segmented markets in which dealers are able to undo the effects of segmentation – but only partially – by providing liquidity. These papers include, Shleifer and Vishny (1998), Gromb and Vayanos (2002, 2017) and Vayanos and Villa (2009). The liquidity measure we use in this paper builds on to the literature on noise in the Treasury yield curve (e.g. Hu, Pan and Wang (2013)) and liquidity in financial markets. Our contribution is to document fluctuations in noise in corporate bond yields and to decompose the role of liquidity supply and demand in determining noise. Goldberg (2017) uses a similar approach, but focuses on the Treasury market and real activity.

The liquidity of corporate bonds and its effect on bond prices have long been a focus of research, and numerous papers offer alternative measures of illiquidity (Bao, Pan and Wang (2011), Chen, Lesmond and Wei (2007), Dick-Nielsen, Feldhütter and Lando (2012), Edwards, Harris and Piwowar (2007), Feldhütter (2012), Friewald, Jankowitsch and Subrahmanyam (2012), Longstaff, Mithal and Neis (2005)). In addition, there is a recent strand of literature which studies corporate bond liquidity using transactions data, often with a focus on the impact of regulation (e.g. Adrian, Boyarchenko and Shachar (2017), Anand, Jotikasthira and Venkataraman (2017), Bao, O'Hara and Zhou (2017), Bessembinder et al. (2017), Choi and Huh (2017) and Goldstein and Hotchkiss (2017), among others.) This paper is different from others as we study price and quantity jointly to isolate the effect of liquidity supply from the demand effect based on simple sign restrictions.

Though we are among the first to analyze the corporate bond illiquidity using the structural VAR, a related approach is taken in Cohen, Diether and Malloy (2007)'s study of short selling. Unlike our paper, Cohen, Diether and Malloy (2007) classify a supply shock as having occurred if the price and quantity of shorting move in opposite directions. Baranova, Chen and Vause (2015) studies the liquidity using a structural VAR focusing on credit spreads. In contrast, this paper studies the price of liquidity motivated by the preferred habitat investors.

The remainder of the paper is organized as follows. Section 2 provides the summary of the theory of segmented markets, and describes the data and the empirical approach. Section 3 presents the empirical results of the VAR model that identifies liquidity supply and demand shocks. Section 4 extends the framework with a joint analysis on the IG and HY markets, a study on liquidity premiums in credit spreads and an investigation into the heterogeneity of dealers. Section 5 concludes.

2 Data and Empirical Approach

In this section, we introduce a simple model of segmented markets based on Gromb and Vayanos (2002) and Goldberg (2017) in order to motivate a choice of price and quantity measures for corporate bond liquidity. Next, we describe our data sets to construct measures of aggregate noise and dealers' gross positions in the corporate bond market. We then explain how we use these price and quantity measures in the vector autoregression (VAR) framework to separate the supply and demand shocks.

2.1 Model of Financial Intermediation in a Segmented Market

Building onto Gromb and Vayanos (2002), we construct a model in which a risk-averse financial intermediary provides liquidity in two segmented financial markets. The full specification of the model is provided in Appendix A, and here we provide a summary of the results which motivate our choice of price and quantity measures.

In the model, there are two securities with identical cash flows traded in two segmented markets, each of which is populated by investors who can only access to one market and faces idiosyncratic endowment shocks. The risk-averse financial intermediary can access both markets, and provides liquidity by buying the undervalued security and selling the overvalued security. However, she faces liquidity shocks which force her to unwind the positions before realizing the profits from the long-short positions. Therefore, the liquidity provided by the intermediary is generally not sufficient to remove the difference in price between two securities.

In making markets, the dealer focuses on managing liquidity risk and does not make any bets on securities' fundamental value. As a result, the quantity of liquidity she provides can be measured by her *gross* positions, or the sum of the long and short positions. Moreover, the gap in price between the two securities can be thought of as a liquidity price, as this is a deviation in prices from the fundamental value due to the need to transact the particular security. Without the market segmentation, the liquidity price is zero.

We summarize the two key takeaways from the model. First, we show that if investors' idiosyncratic desire to trade the security increases, then the financial intermediary's gross position increases, while the gap between two securities widens. This shock increases both liquidity price (the gap in security prices) and quantity (dealer's gross positions), and thus we call it a demand shock. Second, if the financial intermediary becomes less risk-averse, then she provides more liquidity in the market, leading to larger gross positions and a narrower gap between the two securities. This second shock decreases the liquidity price but increases the liquidity quantity, and we call it a supply shock. As a result, we connect unobservable motive to request for or provide liquidity to observable quantities such as the gap in prices among claims on identical cash flows and dealer's positions on those claims. In the next section, we describe how we construct the liquidity price and quantity measures for empirical analysis.

2.2 Data: Quantities

The theory of a segmented market implies that we use financial intermediaries' gross positions on corporate bonds as a measure of liquidity quantity. In this paper, we focus on primary dealers who are designated as trading counterparties by the Federal Reserve Bank of New York. We limit our sample to primary dealers since we are interested in market participants who are considered as traditional market makers rather than new participants who do not commit their balance sheets for market making activities. The primary dealers as a group maintains the volume share of around 70% in our sample throughout the sample period.

In order to construct dealers' gross positions on corporate bonds as a measure of liquidity quantity, we use the regulatory TRACE data from 2002 to 2016 obtained from FINRA. The regulatory TRACE data is similar to the TRACE data available for the public, except that the volume for large trades is not truncated, and that the dealer's name is revealed. With the data, we identify which dealer is involved in each transaction. Using this identifier, we aggregate trades for each dealer for each bond within a week to construct the weekly flow. Then we accumulate weekly flows to obtain the estimate for the weekly inventory data of bonds for each dealer. We discard the positions if they are not closed within the four week windows using the last-in first-out method. Goldstein and Hotchkiss (2017) report that nearly 60% of paired round-trip trades are completed within a day, and for those which take more than a day, the weighted average holding period of a bond for dealers is 21 days. Thus, we focus on the trades reversed within the four-week rolling window because a trade that

Table 1: Example of the Inventory Construction Method

		Volume	Amount Outstanding for each ID					End-of-Week
ID	Week		1	2	3	4	5	Inventory
1	1	1000	1000					1000
2	2	200	1000	200				1200
3	3	-300	900	0	0			900
4	4	-500	400	0	0	0		400
5	5	100	**0**	0	0	0	100	100

Note: This table shows an example of an artificial data using the method in constructing the estimate for an inventory for a hypothetical bond.

takes longer than four weeks to reverse are likely to be done for the purpose other than market making activities (e.g. proprietary trading).

Table 1 shows an artificial data to highlight our methodology in constructing the sample for a hypothetical bond over a 5-week period. In the first week, the dealer end the week by accumulating net long positions of 1,000 millions. This trade is carried through up to week four, and thus the weekly volume of 300 and 500 millions in sale are net out, leading to the end of the week inventory of 400 millions. In week 5, the unreversed position of the initial trade (with ID=1) is set to zero, so that the inventory estimate for week 5 is 100 millions rather than 500 millions.

Since TRACE is fully implemented in February 2005, we start our analysis from April 2005. Some primary dealers have multiple dealer ids, and thus we aggregate those ids based on the name associated with them. In theory, this cumulative flow should provide a reasonable estimate for the stock every week. Still, a preliminary analysis shows that inventory measures are subject to flows that are likely to reflect mismeasurement in the data. To mitigate the issue, we remove weekly flows that are greater than a third of the amount outstanding of the bond, and remove all flows for the bonds with age less than one month. We remove young bonds since a part of the sales by dealers likely represents the sales of the leftover of the bonds underwritten in the primary market. We also make sure that there is no inventory of bonds that matured.

Furthermore, in order to maintain the consistency in pricing among bonds, we merge TRACE data with the characteristic data from Mergent FISD, and limit our sample to dollar-denominated publicly offered bonds with fixed coupons and no embedded options other than make-whole call provisions. After the filters, we have 18,986 bonds issued by 4,466 issuers over 614 weeks in TRACE data to construct the gross positions.

Let $Q_{d,j,k,t}$ denote the inventory of bond k held by dealer d in week t. Then our quantity measure is computed by aggregating the absolute values $Q_{d,j,k,t}$ by

$$Q_t = \sum_j \sum_k \sum_d |Q_{d,j,k,t}| \,. \tag{1}$$

Equation (1) allows the possibility that the inventory for each bond $Q_{d,j,k,t}$ is negative. Asquith et al. (2013) show that the cost of borrowing corporate bonds to short is comparable to the cost for stocks. The FR-2004 data published by the Federal Reserve Bank of New York also shows that primary dealers have substantial amount of short positions on corporate bonds.[1] Thus, we do not remove the observations if $Q_{d,j,k,t}$ is negative. Finally, we scale the aggregate quantity measure Q_t by consumer price index excluding food and energy to express it in 2005 dollars, and use its logarithm $q_t = \log Q_t$ in the analysis below.

The top panel of Figure 1 shows the time-series of the aggregate dealer gross positions, Q_t. At the beginning of the sample, the gross positions move around 25 billion dollars, and start to decline in late 2007 until the financial crisis in 2008 below 15 billion dollars. The quantity recovers from 2009 to 2011, and declines gradually since then. The maturity of the gross positions is stable over time, with average of 8.8 years and standard deviation of 0.28 years. The ratio of IG bond positions to the all positions is on average 58% with standard deviation of 4%. Moreover, the positions on IG bonds and HY bonds are highly correlated with correlation coefficient of 0.76. Thus, the risk profiles of the dealer gross positions are relatively stable over time, and not likely to be the key driver for the empirical results below.

2.3 Data: Prices

The theory of segmented markets suggests that the discrepancy in market prices among securities that share similar cash flow risks can be regarded as a price of liquidity. In this

[1] The correlation between Q_t and gross outright positions reported in FR-2004 is 0.58 using the data from April 2013, when the FR-2004 data begins for corporate bonds. The correlation is lowered because the FR-2004 data includes commercial papers and bonds with optionalities, and inventories are reported on the fair-value basis rather than the book value.

Figure 1: Noise in Corporate Bond Yields and Dealer Gross Positions

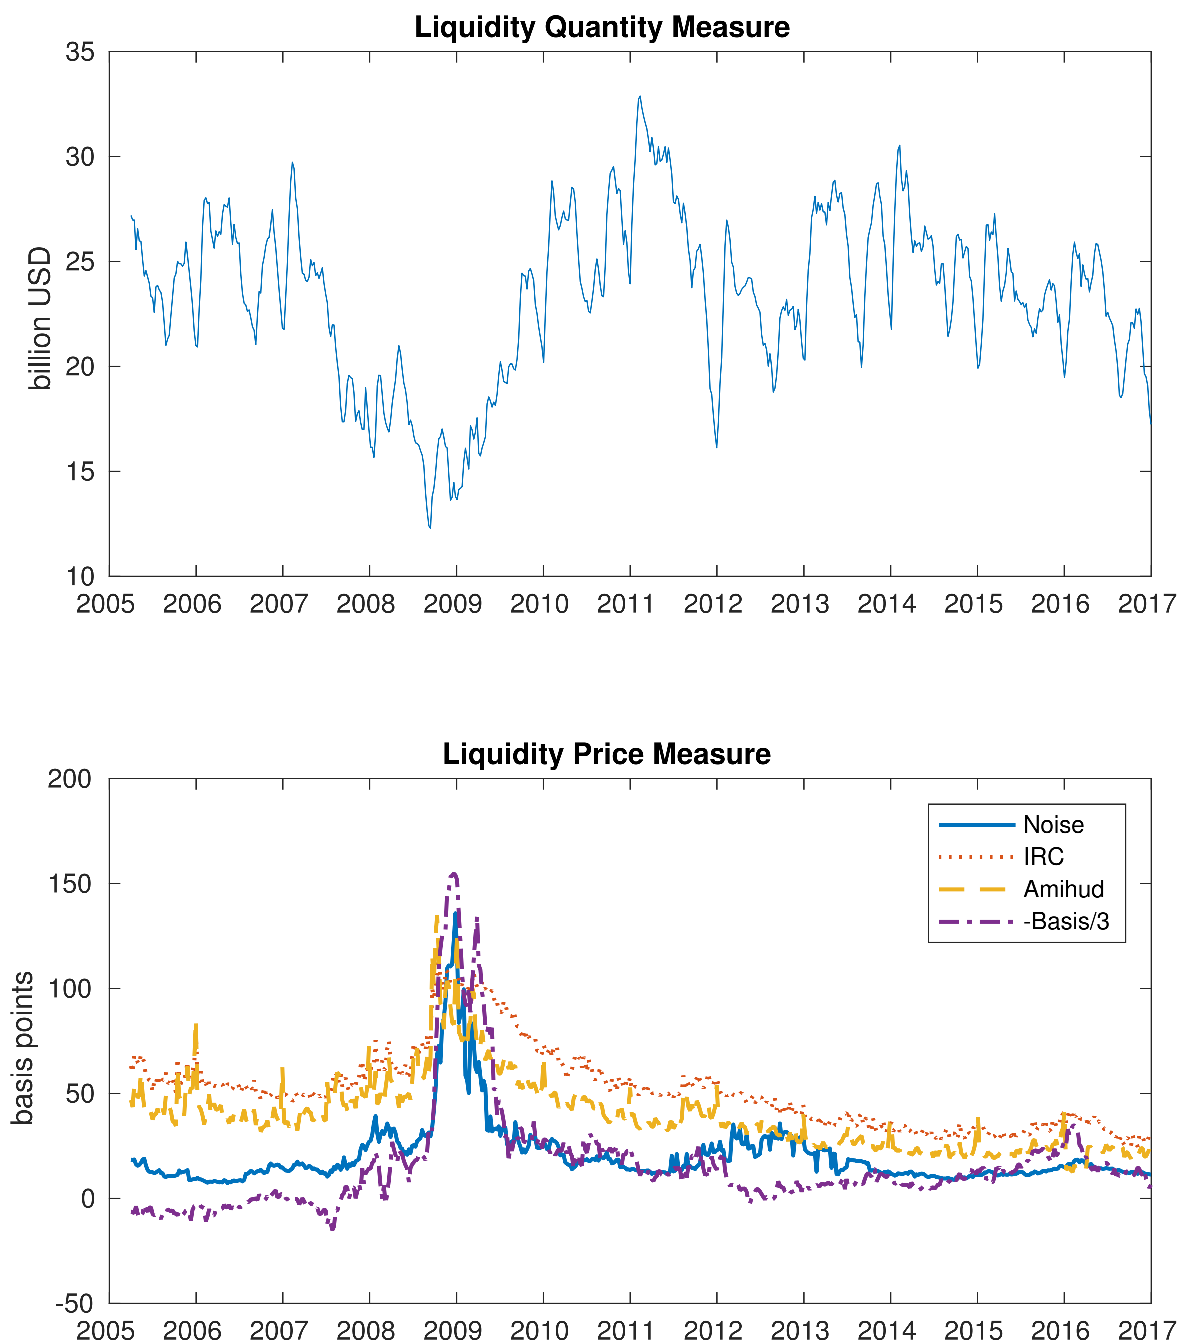

Note: The top panel shows the gross positions on corporate bonds aggregated among primary dealers from 2005 to 2016, scaled by CPI and expressed in 2005 dollars. The bottom panel shows the aggregate measure of noise from 2005 to 2016. It is the root mean-squared error of the Nelson-Siegel-Svensson model in basis points, averaged across issuers.

paper, we view preferred habitat theories of interest rates (Culbertson (1957), Modigliani and Sutch (1966) and Vayanos and Vila (2009)) as the key driver for market segmentation, and assume that the market for corporate bonds is segmented by maturity. This assumption leads us to follow Hu, Pan and Wang (2013) and use deviation of yield to maturity of a corporate bond from a smooth curve as a measure of liquidity price. In particular, we estimate a smooth curve issuer by issuer every week, and measure pricing errors against the curve. The use of the deviation from the issuer-specific yield curve is motivated by our view on the role of dealers, which is to provide liquidity across maturity buckets without making a bet on the default risk of the issuer.

To compute the pricing errors, we obtain weekly bond price data from Merrill Lynch from 2005 to 2016. Though the Merrill Lynch data is based on quotes, in Appendix B, we compare prices between Merrill Lynch and TRACE using overlapping observations, and find that noise measure is unlikely to be affected by our choice of the data set.

Our sample consists of bonds that are in the Merrill Lynch U.S. Corporate Master database, which require bonds to have amount outstanding greater than 100 million dollars and remaining time to maturity greater than one year. As we do for the quantity measure, we limit our sample to dollar-denominated publicly offered bonds with fixed coupons and no embedded options other than make-whole call provisions. Make-whole call options allow an issuer to call the bond before maturity, but the strike price of the option varies in the way that options are never in-the-money. Therefore, the payoff from the options is in principle zero, and so is the value of the option. Since a large fraction of bonds in our sample have this make-whole call provision, we use these bonds in computing pricing errors. Even if these options have non-zero values due to the benefit of cash flow management (Elsaify and Roussanov (2016)), the use of make-whole call bonds does not affect our results so far as the value of options does not vary substantially across maturity.

In order to estimate the curve reliably, we focus on issuers that have more than ten issues outstanding in a given week. Because we focus on large issuers in estimating noise, our sample in constructing noise is smaller than the TRACE sample used to construct the quantity measure. Over the sample period, we have 3,040 bonds issued by 169 unique issuers. We examine potential bias due to the mismatch between the two samples in the next section, and argue that the difference in two samples does not affect our empirical analysis.

The Nelson-Siegel-Svensson model of Svensson (1994) assumes that n-period instantaneous forward rates follow:

$$f(n) = \beta_0 + \beta_1 \exp(-n/\tau_1) + \beta_2(n/\tau_1)\exp(-n/\tau_1) + \beta_3(n/\tau_2)\exp(-n/\tau_2)$$

Since corporate bonds in our sample pay coupons, we compute the model-implied price of corporate bonds as a sum of present values of the cash flows, $p(\theta)$, where $\theta = \{\beta_0, \beta_1, \beta_2, \beta_3, \tau_1, \tau_2, \tau_3\}$. To find parameters, we minimize the sum of squared pricing errors.

$$\hat{\theta} = argmin_{\theta} \sum_k \left(\frac{p_k - p_k(\theta)}{\tau_k} \right)^2$$

where τ_k is modified duration for bond k.

We compute yield-to-maturity for each bond, using observed prices and model-implied prices. The pricing error of a bond is the difference in yield to maturity,

$$\epsilon_k = ytm(p_k) - ytm\left(p_k\left(\hat{\theta}\right)\right)$$

We then compute the root mean-squared error for firm j to obtain the noise in week t.

$$rmse_{j,t} = \sqrt{\frac{1}{K}\sum\nolimits_k \epsilon^2_{k,j,t}}$$

Finally, our measure of aggregate noise is the average across issuers:

$$p_t = \frac{1}{J}\sum_j rmse_{j,t}$$

The bottom panel of Figure 1 shows the time-series of the average noise, in which the average is computed across issuers every week. The noise spikes up during the financial crisis, and comes down afterwards. It also goes up somewhat in 2012 reflecting the turmoils in European financial markets and the debt crisis of Greece, but remains low since 2013. The figure also compares p_t with other measures of illiquidity, such as imputed round-trip costs (IRC) of Feldhütter (2012), the Amihud (2002) illiquidity measure and the negative of

CDS-bond basis obtained from JP Morgan. The IRC and Amihud measures are computed every week using the 3-month rolling windows, and we report the median across all bonds in TRACE data. This pattern is broadly similar to the time-series behavior for other conventional measures of illiquidity, as they increase during the crisis and remain low after the crisis.

Among a variety of alternatives for measures of mispricing, we focus on investors' preferred habitat in terms of maturity, and choose noise as our measure for the deviation in prices among claims on the same cash flows. As in the reduced-form model of Duffie and Singleton (1999), it is natural to conjecture that there are a set of a few state variables (such as default risks) which drive all corporate bonds issued by the same issuer. Then, there is a cross-sectional restriction on yields of those bonds which could be succinctly captured by no-arbitrage conditions, and thus a deviation from the model can be thought of as a pricing error due to a liquidity shock idiosyncratic to a specific maturity. In this paper, we are not interested in identifying those factors that drive issuers' fundamentals, and thus we do not model the term structure of corporate bonds explicitly. Rather, we use the Nelson-Siegel-Svensson curve which these term structure models often try to match as a benchmark, and directly measure pricing errors against it.

An alternative measure of mispricing is a difference in credit spreads between corporate bonds and credit default swap (CDS) spreads. As Bai and Collin-Dufresne (2013) show, however, that the CDS and corporate bond liquidities are driven by different set of factors. On the one hand, corporate bond market liquidity is primarily driven by search frictions in the over-the-counter market and thus the dealers' willingness to hold inventories plays a central role in determining the liquidity condition. On the other hand, CDS liquidity is affected also by other market frictions including counterparty risk, repo market functioning and variation in haircuts on collateral. Thus, the difference in spreads between CDS and corporate bonds, or the CDS-bond basis, is likely to be driven by the same set of liquidity factors. The focus of this paper is to study the drivers of corporate bond liquidity, and thus we choose noise as our price measure rather than the CDS-bond basis.

A common problem in the literature on corporate bond liquidity is that not all the characteristics of a given trade are observable to the researcher; if there is a shift in the composition of trades along unobserved dimensions such as immediacy, liquidity measures based on realized trades may be misleading (Dick-Nielsen, Feldhütter and Lando (2012) and Dick-Nielsen and Rossi (2017)). One version of this story is that dealers have become reluctant to hold inventory since the crisis, but trading costs calculated from realized trades nonetheless are little changed relative to pre-crisis levels because dealers are acting increasingly as agents (simply connecting customers) rather than principals (providing immediacy with their balance sheets). Our approach addresses this problem directly by exploiting joint variation in dealer gross positions and the noise in bond prices. Moreover, the bonds included in the calculation of our positions and noise measures do not shift based on realized trades. In particular, by using quotes, we reduce contamination of our noise measure from shifts in the composition of realized trade volume across a given set of bonds. (Using quotes data for bonds that also have recorded transactions in the TRACE data, we also provide evidence that our results are unlikely to be affected by using quotes rather than prices from realized trades.)

2.4 Summary Statistics

Table 2 shows the summary statistics for the state variables, p_t and q_t. The unconditional correlation between p_t and q_t are -0.62, suggesting the importance of a supply shock. The table also shows the correlation between q_t and median Amihud illiquidity measure, imputed round-trip costs and CDS-bond basis which are estimated at -0.60, -0.52 and 0.53, respectively. As in the bottom panel of Figure 1, p_t is positively correlated with other conventional measures of corporate bond illiquidity. The correlation between p_t and median Amihud illiquidity measure, and the correlation between p_t and imputed round-trip costs are both estimated at 0.71, while the correlation between p_t and CDS-bond basis is -0.87. Therefore, our measure of the price of liquidity is highly correlated with traditional measures of illiquidity.

In order to estimate the yield curve reliably, we use issuers who have at least 10 issues outstanding to estimate the Nelson-Siegel-Svensson curves and noise, p_t, while we use inventory for all corporate bonds to compute the quantity measure, q_t. As a result, we use subsample of relatively large issuers of corporate bonds to construct noise, while no such selection bias exists in constructing the quantity. Therefore, it is important to understand if and how the selection bias affects our measure of noise. Table 3 reports the conventional illiquidity measures between two subsamples: the one matched to Merrill Lynch data sets and used to compute the noise measure, and the other not used to compute noise, either because the bond does not show up in Merrill Lynch database or because the issuer has less than 10 bonds outstanding. Specifically, we compute the average illiquidity measures every week separately for each subsample, and compute the average and correlation between the two series. Panel A reports the correlation coefficients. Using bonds with all credit ratings, the correlation between two groups is 0.96 using imputed round-trip costs, 0.94 using the Amihud measure, and 0.89 using weekly trading volume. The correlation using investment-grade bonds are 0.95, 0.94 and 0.85 using imputed round-trip costs, the Amihud measure and trading volume, respectively. The correlation using high yield bonds are somewhat lower at 0.72, 0.78 and 0.73, respectively.

Panel B of Table 3 reports the average values for each illiquidity measures for each subsample of bonds. The subsample used to compute noise consists of more liquid bonds than the subsample not used to compute noise. On average, the bonds used to compute noise have lower imputed round-trip costs, lower Amihud measure and higher transaction volume. The difference in average illiquidity is natural, as one would expect the bonds issued by large issuers are more liquid than those issued by small issuers. However, in the structural VAR we use, the difference in level of variables does not affect any of our analysis. The bias in the level of a variable changes the intercept of the regression without affecting the estimates for the slope and the variance-covariance matrix of shocks, which are the key objects of our interest.

2.5 Empirical Approach

We identify liquidity supply and demand shocks using a standard VAR model of supply and demand with sign restrictions (e.g., Uhlig (2015)). Denote the vector of noise and dealer gross positions by $Y_t = \begin{pmatrix} p_t & q_t \end{pmatrix}'$. The VAR takes the following form:

$$Y_t = b + B_1 Y_{t-1} + B_2 Y_{t-2} + ... + B_L Y_{t-L} + \xi_t \tag{2}$$

Table 2: Summary Statistics: Weekly 2005-2016

Panel A: Summary Statistics

	Mean	Std	AR1	AR12
q_t	16.95	0.18	0.98	0.67
p_t	20.53	16.43	0.97	0.68

Panel B: Correlation Matrix

	p_t	Amihud	IRC	Basis
q_t	-0.62	-0.60	-0.52	0.53
p_t		0.71	0.71	-0.87
Amihud			0.93	-0.62
IRC				-0.66

Note: Variable p is the root mean squared error of the Nelson-Siegel-Svensson curve averaged across issuers every week. Variable q is log of aggregate gross positions of primary dealers on corporate bonds. Amihud is the median Amihud measures, IRC is median imputed round-trip cost, Basis is the CS-Bond basis.

Table 3: Comparison of Samples Used to Compute Fitting Errors and TRACE Sample

		NObs	IRC	Amihud	Volume
Panel A: Correlation Between Matched and Unmatched Bonds					
All			0.96	0.94	0.89
IG			0.95	0.94	0.85
HY			0.72	0.78	0.73
Panel B: Average values and number of observations					
All	Matched	376,171	0.55	0.44	10320
	Unmatched	1,495,208	1.89	0.62	7420
IG	Matched	351,562	0.52	0.42	10482
	Unmatched	925,402	0.65	0.57	7867
HY	Matched	24,609	0.82	0.64	8014
	Unmatched	569,806	3.79	0.68	6695

Note: Matched bonds are the sample of bonds in TRACE that have more than 10 issues outstanding in the week and are matched to Merrill Lynch database to compute noise. Unmatched bonds are the sample of bonds in TRACE that are used to construct the dealers' inventory measure, but not to compute noise. IRC is median imputed round trip costs of Feldhütter (2012), Amihud is median Amihud measure (2002), and Volume is median dollar trading volume in millions.

where ξ_t is the reduced-form residual, with $E\left[\xi_t \xi_t'\right] = \Sigma$.

Denote the mapping from orthonormal fundamental shocks v_t to the residual ξ_t by the matrix A, with $\xi_t = A v_t$. Consistent with the literature on segmented markets, we assume that a positive liquidity supply shock leads to a decrease in noise and a rise in dealer gross positions. We also assume that a positive liquidity demand shock leads to a rise in both dealer gross positions and noise. Therefore, to identify the structural shock, we impose the following sign restrictions on the rotation matrix A,

$$\begin{pmatrix} \xi_t^p \\ \xi_t^q \end{pmatrix} = \underbrace{\begin{pmatrix} - & + \\ + & + \end{pmatrix}}_{A} \begin{pmatrix} v_t^s \\ v_t^d \end{pmatrix} \tag{3}$$

so that the first entry of v_t corresponds to a supply shock and the second entry corresponds to a demand shock. Using our identification assumptions, we can identify A and thus uncover the supply and demand shocks from $v_t = A^{-1}\xi_t$.[2]

We estimate the model using the pure sign restrictions approach of Uhlig (2005). In estimating the reduced-form VAR, we use a weak Normal-Wishart prior with $L = 13$ lags of the endogenous variables.[3] In Appendix C, we show the posterior mean estimates for the reduced-form parameters B and Σ, and compare them with the OLS estimates. The two estimates are very similar to each other, confirming that our prior to the reduced-form parameters are not driving the results.

To identify structural shocks, we construct the rotation matrix A using the QR factorization, following Uhlig (2005) and Rubio-Ramirez, Waggoner and Zha (2010). For each draw of parameters in the reduced-form VAR from the posterior distributions, we draw entries for a 2-by-2 matrix W from a standard normal distribution. We then apply the QR decomposition to W to obtain the orthogonal matrix Z_W with positive diagonal elements. Based on lower triangular matrix C from the Cholesky decomposition of a draw of reduced-form covariance matrix Σ, we use a product $A_m = CZ_W$ as a candidate rotation matrix from the m-th draw,

[2] In the literature on VARs with sign restrictions, the sign restrictions are typically required to hold not only on impact, but also for a sustained period of time afterward. Here, we assume only that the sign restrictions hold on impact, but later show that the impulse-response functions obey the same sign restrictions for periods up to 16 weeks.

[3] The results are very similar if a time trend is included in the VAR or if the lag length L is doubled.

and check if the candidate matrix obeys the sign restriction in (3). If it does, then we keep the draw and if it does not, we discard the draw. We repeat the draws for A_m 100 times for each draw of reduced form parameters, $B_1, \dots, B_L, \Sigma$, which are also drawn 100 times from the posterior distribution. As a result, we have 10,000 draws of parameters that characterize the dynamics of the structural VAR, including the rotation matrix, history of shocks, and impulse response functions. In the following analysis, we report the mean and the 68 and 95 percent credible sets for the key parameters of interest, computed across these draws from the posterior distribution that satisfy the sign restrictions.

3 Liquidity Supply and Demand

This section presents our estimates for liquidity supply and demand shocks and analyzes their effects on noise and dealer gross positions.

3.1 Time Series of Liquidity Supply and Demand Shocks

In this section, we describe the historical behavior of the supply and demand shocks identified by our methodology. Figure 2 shows the pointwise mean of the cumulative sum of the liquidity supply shocks, $\sum_{j=1}^{t} v_j^s$, where the mean is computed across posterior draws. The supply shocks capture episodes of stress in the corporate bond market as well as more general episodes of financial stress. For example, liquidity supply declines following the suspension of redemption from Bear Stearns hedge funds in the summer of 2007, an event that marked the beginning of the financial crisis. The supply declines sharply after the failure of Lehman Brothers in 2008. The liquidity supply recovers subsequently, but declined somewhat during the European fiscal crisis in 2010 and after the "taper tantrum" in 2013, when interest rates rose as the Federal Reserve considered when to reduce purchasing long-term assets.

Figure 2 also shows the cumulative liquidity demand shocks, $\sum_{j=1}^{t} v_j^d$. Immediately after the collapse of Lehman Brothers, the liquidity demand spikes up, consistent with the rise in noise during the financial crisis. The liquidity demand stabilizes somewhat after the crisis, but generally maintains the upward trend with little changes during and after the European fiscal crisis in 2010 and the taper tantrum in 2013. The increasing trend in liquidity demand

Figure 2: Cumulative Sum of Liquidity Shocks

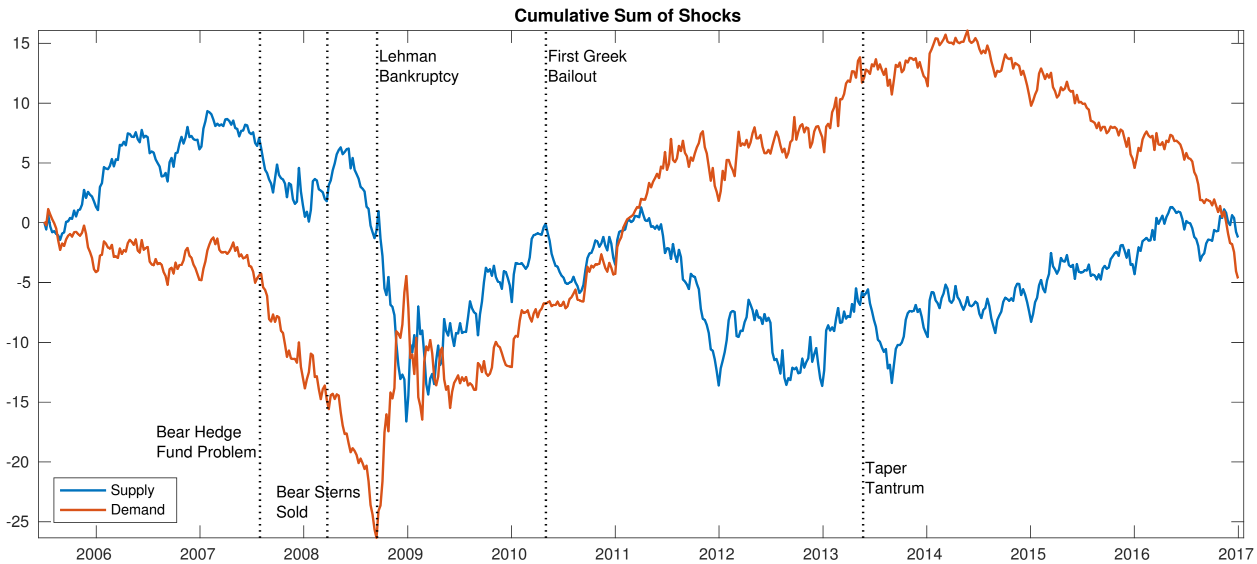

Note: This figure shows the mean of the posterior distribution of the cumulative sum of shocks to liquidity demand and supply. A positive shock reflects an increase in liquidity supply. The data is weekly from 2005 to 2017.

between 2010 and 2014 is consistent with the expanding corporate bond market when firms issue large amount of corporate bonds in response to the lower cost of debt. As the market expands, investors' demand for liquidity naturally increases. However, the demand shocks are generally negative after 2015. The negative demand shocks in the recent period is consistent with the findings of Anand, Jotikasthira and Venkataraman (2017) and Choi and Huh (2017), who document that non-dealers start to provide liquidity in the corporate bond market in the recent period.

The behavior of structural shocks are better understood by comparing those with reduced-form shocks. As an example, we focus on 2008 when the financial crisis occurs. The top panel of Figure 3 plots the cumulative structural shocks $\sum_{j=1}^{t} v_j$ in 2008, while the bottom panel plots the cumulative reduced-form shocks $\sum_{j=1}^{t} \xi_j$ over the same period. Since March 24, 2008, when JP Morgan bailed out Bear Sterns, the dealers' inventory was hit with negative shocks, while the noise was also hit with somewhat smaller negative shocks. The identification strategy based on the sign restrictions interprets these reduced-form shocks as both supply and demand falling, and the demand fell slightly more than the supply, leading to the lower quantity and lower price. Since the bankruptcy of Lehman Brothers on

Figure 3: Structural and Reduced-Form Shocks During the Financial Crisis

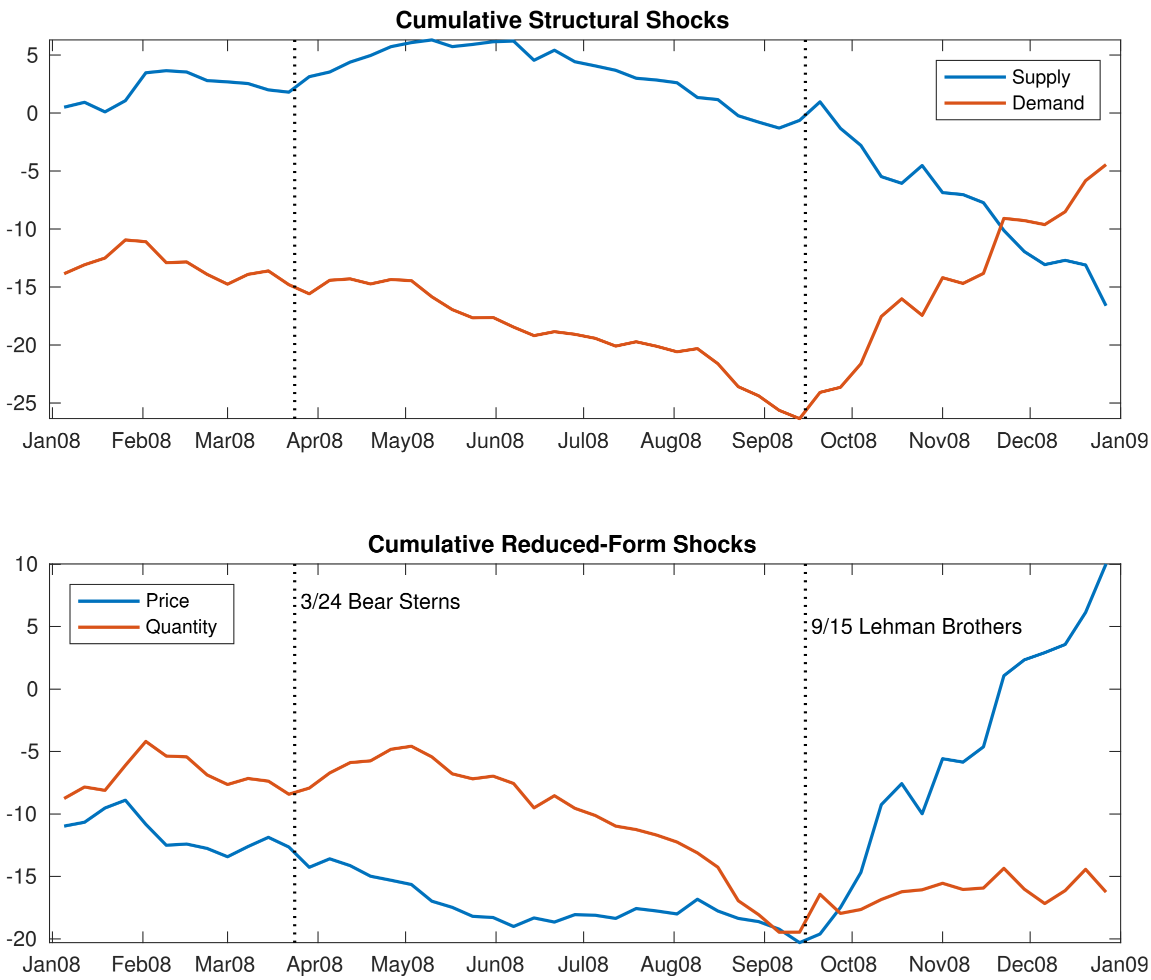

Note: The top panel shows the mean of the posterior distribution of the cumulative sum of shocks to liquidity demand and supply in 2008. The bottom panel shows the cumulative sum of shocks to noise and the aggregate dealer inventory in 2008.

September 15, 2008, the noise measure receives large positive shocks while dealers' inventory is about unchanged. The identification strategy interprets these behavior as significant positive demand shocks and negative supply shocks, which leads to higher price and the similar level of quantity.

Using the estimated structural shocks, we can ask what happens to endogenous variables, if there were no demand shocks. The idea is that we can evaluate the pure effect of supply shocks on corporate bond illiquidity by turning off the demand shock that contaminate the movements of endogenous variables. Bao, O'Hara and Zhou (2017) and Bessembinder et al. (2017) use the period since July 2010 as the period in which a series of banking regulations

Figure 4: Counterfactuals: No Demand Shocks Since July 2010

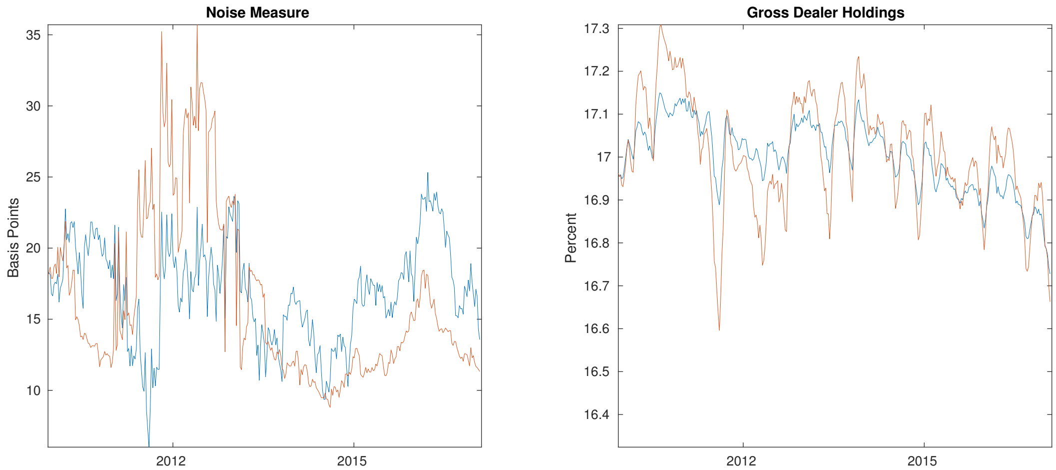

Note: The orange line shows the cumulative shocks in the historical data. The blue line shows the results of the counterfactual experiment in which regulation shocks are set to zero after July 2010.

are introduced. We follow them in defining the post-regulation period, and turn off the demand shocks after July 2010.

Figure 4 shows the variation in endogenous variables in history compared with the counterfactuals without any demand shocks. Overall, without demand shocks, noise measures would have been higher after 2014. This observation is consistent with the pattern in Figure 2, in which demand shocks are generally negative after 2014. As of the end of 2016, the noise measure would have been about 10 basis points higher if the demand had not decreased over this period. This result is consistent with Bao, O'Hara and Zhou (2017), who emphasize the importance of the Volker rule more than the other post-crisis regulations, as the Volker rule becomes effective in 2014. In contrast, gross dealer holdings would have been more stable without demand shocks. Overall, there is some evidence that the liquidity supply is worsened after July 2010, but the magnitude of the effect on noise does not indicate a catastrophic decline in liquidity provision. Appendix D shows the impulse response functions of endogenous variables to supply and demand shocks, as well as the forecast error variance decomposition of liquidity price and quantities. The results suggest that both supply and demand shocks are important drivers for noise, and capital flows slowly to the corporate bond markets so that dealer inventory mean-reverts slowly, especially after a supply shock.

3.2 Liquidity Risk and Elasticity Estimates

The estimated rotation matrix A provides the estimate for the elasticity of supply and demand. Let a_{ij} be an entry for i-th row and j-th column of A. Since the first entry of structural shock vector is a supply shock and the second is a demand shock, we calculate elasticities of supply and demand as,

$$\begin{aligned} \varepsilon^s &= \frac{a_{22}}{a_{12}} \\ \varepsilon^d &= \frac{a_{21}}{a_{11}} \end{aligned} \tag{4}$$

for each draw of matrix A. The idea behind the elasticity estimates is that when a demand (supply) shock hits the market, the ratio of the quantity reaction to the price reaction reveals the local slope of the supply (demand) curve, which is precisely what the elasticity of supply (demand) is. The elasticity of liquidity supply is an interesting object for policy makers, as it shows the tendency of dealers to expand their balance sheet in response to a rise in noise. If the elasticity of supply is high, an increase in noise leads to greater expansion of dealer balance sheet.

Following Bao, O'Hara and Zhou (2017) and Bessembinder et al. (2017), we compare elasticities across different subperiods to evaluate the impact of the new regulations introduced after the financial crisis on dealers' willingness to provide liquidity. However, unlike Bao, O'Hara and Zhou (2017) and Bessembinder et al. (2017) who document the level of liquidity supply, we focus on the sensitivity of liquidity supply to a shock in order to understand what happens to market liquidity if investors' demand for liquidity surges.

Table 4 presents the median and the 68-percent credible interval for estimated elasticities using the full sample as well as three subsamples: i) pre-crisis period (until June 2007), ii) crisis period (July 2007-April 2009), and iii) regulation period (July 2010-). The median estimate for the supply elasticity using the pre-crisis sample is 2.43, implying that a dealers' balance sheet expands 2.43% in response to a basis point rise in noise. During the financial crisis, the supply elasticity falls to 0.43 and increases to 1.07 after the crisis. Even though the supply elasticity recovers after the crisis, the median estimate is less than half of the

Table 4: Elasticities of Liquidity Supply and Demand

		Median	LB	UB	Median	LB	UB
		Elasticity of supply			Elasticity of demand		
Pre-Crisis	2005.4-2007.6	2.43	1.09	6.78	-4.18	-12.97	-0.97
Crisis	2007.7-2009.4	0.43	0.17	1.24	-0.73	-2.25	-0.18
Regulations	2010.7-2016.12	1.07	0.34	4.83	-1.96	-5.69	-0.52
Full Sample	2005.4-2016.12	0.63	0.22	2.09	-1.24	-3.68	-0.32

Note: Using each subsample, we estimate the VAR in (2) and draw 2,500 candidate matrix of A in $\xi_t = Av_t$ from the posterior distribution and retain the candidate that satisfy the sign restrictions. For each draw, we compute elasticities of demand and supply using (4), and report the median across draws, as well as the lower bound (LB) and upper bound (UB) of the 68% credible sets. The elasticity is multiplied by 100 such that, for example, the value of 3 means a 3 percent rise in dealer gross positions in response to a 1 basis point change in noise.

elasticity before the financial crisis. This result reinforces the findings in the literature, in which dealers are less willing to provide capital to the market making activities in the corporate bond market after the financial crisis. Furthermore, our estimates predict that given a demand shock, the median response of noise in the corporate bond market is more than twice as large as the pre-crisis level. Therefore, even if the level of illiquidity might be low during the post-crisis period, *liquidity risk* is high.

Though striking, our results need to be treated with a caveat. A recent work by Baumeister and Hamilton (2016) show that uniform prior over the space of orthogonal matrix Z_W does not necessarily imply uniform prior for estimated elasticities, as they are nonlinear functions of Z_W. However, we emphasize that we use the same prior for all subsamples, and thus the difference in median draw across different subsample cannot be solely due to our choice of prior. Instead, the difference in point estimate must come from the data.

3.3 Drivers of Supply and Demand Shocks

The analysis in the previous section begs the question about the driver of liquidity demand and supply shocks in the corporate bond market. The liquidity demand will be high when investors need to trade specific securities in a large amount immediately. Coval and Stafford (2007) show that equity mutual funds are forced to buy and sell stocks when there is a large fund flow to these funds. Their evidence suggests that fund flows to corporate bond mutual fund may force them to trade more, and thus lead to an increase in liquidity demand. To confirm this idea in the data, we run the following regression

$$v_t^d = b_0 + b_1|flow_t| + \epsilon_t$$
$$v_t^d = b_0 + b_1 flow_t^2 + \epsilon_t$$

where $flow_t$ is the aggregate corporate bond mutual fund flow in week t, obtained separately for investment-grade funds and high yield funds from Investment Company Institute and Morningstar since 2007. We use absolute values or squared flows to capture the idea that large flows, whether positive or negative, will force mutual funds to trade more.

Table 5 reports the estimated slope coefficients and regression R-squared. The slope coefficients are positive and statistically significant for both IG and HY mutual fund flows, regardless of regression specifications. The results imply that large mutual fund flows, regardless of inflows or outflows, are associated with a positive demand shock we estimated using noise and dealer balance sheet data. In untabulated results, we also test if the positive flows and negative flows deferentially affect liquidity demand, but we do not find an evidence for an asymmetry in the estimated slope coefficients. Overall, the demand shocks estimated in the previous sections are indeed tied to the observable investment flows, though R-squared is generally low.

The low R-squared indicates that there are investors other than corporate bond mutual funds that have their own liquidity shocks that are important in understanding the aggregate demand shocks in the corporate bond market. Indeed, the Federal Reserve's flow of funds shows that mutual funds holds only 15% of corporate and foreign bonds issued in the U.S. at the end of 2016. Nonetheless, the evidence that links our demand shock estimates and mutual fund flows provides some comfort to our identification strategy for structural shocks.

Next, we examine what drives the liquidity supply shocks. When uncertainty of the value of the securities is high, a risk-averse financial intermediary is likely to reduce her positions, given the limitation on the amount of risk capital she can commit to market making activities. Thus, we expect that an unexpected rise in uncertainty affects liquidity supply negatively. To examine this hypothesis, we obtain swaption-implied volatility from JP Morgan with

Table 5: Mutual Fund Flows and Liquidity Demand Shocks

RHV	b_1	$t(b_1)$	R^2
IG $flow_t^2$	0.01	(4.66)	0.01
IG $\|flow_t\|$	0.17	(4.43)	0.02
HY $flow_t^2$	0.03	(3.76)	0.02
HY $\|flow_t\|$	0.21	(3.92)	0.03

Note: The table reports the estimated slope coefficient and R-squared of the regression:

$$v_t^d = b_0 + b_1|flow_t| + \epsilon_t$$
$$v_t^d = b_0 + b_1 flow_t^2 + \epsilon_t$$

where $flow_t$ is mutual fund flow in week t. The sample is weekly from January 2007 to December 2016.

Table 6: Correlation Between Volatility, Supply and Demand Shocks

	Expiry / Maturity for Rates			
	1 week / 5-year	1 month / 5-year	1 week, 10 year	1 month / 10-year
3-month rolling window				
Supply shocks	-0.27	-0.27	-0.37	-0.36
Demand shocks	0.11	0.12	0.22	0.24
6-month rolling window				
Supply shocks	-0.30	-0.30	-0.51	-0.51
Demand shocks	-0.12	-0.12	0.07	0.09

Note: Table reports the correlation coefficients between volatility shocks and cumulative supply and demand shocks over the 3-month or 6-month rolling windows. The volatility shocks are extracted by fitting a univariate VAR(13) model to the time-series data of swaption-implied volatility.

1-week and 1-month expiries on 5-year and 10-year rates. We fit a univariate AR(13) model to each time-series and extract volatility shocks.

Table 6 reports the estimated correlation between moving average of volatility shocks for each expiry and maturity, and moving average of the liquidity demand and supply shocks. The correlations between volatility and supply shocks are negative, ranging from -0.27 for 5-year rates using the 3-month window to -0.51 for 10-year rates using the 6-month windows. In contrast, the correlations between volatility and demand shocks have changing signs, and are generally closer to zero. The estimated correlation shows that the liquidity supply shock estimated based on structural VAR are indeed tied to changing uncertainty in the bond markets.

Figure 5 shows the time-series of volatility shocks to 5-year and 10-year rates with 1-month expiries with the time-series of 6-month moving-window average supply shocks. The negative correlation between the two is driven by the pre-crisis period when supply shocks are generally positive and volatility shocks are negative. During the crisis, volatility receives large positive shocks while the liquidity supply is hit by negative shocks. Also, during the "taper tantrum" of May 2013, the rate volatility spikes up while liquidity supply falls. Overall, the negative correlation between the liquidity supply shock and volatility shocks are evident in most of the sample period, and the negative correlation estimates are not the artifact of a few outliers.

However, at the beginning and the end of the sample, volatility shocks are generally close to zero while supply shocks are volatile, lowering the magnitude of correlation between these time-series. Therefore, though the rate volatility may be one of the important drivers of liquidity supply, it is not the only drivers of liquidity supply and it will not be justified to use volatility shocks as a proxy for liquidity supply shocks. Instead, the proper identification of liquidity supply shocks must come from the joint analysis of price and quantities.

4 Extension

The structural VAR analysis on the corporate bond liquidity is useful not only to diagnose the drivers of liquidity in the market, but also provides an analytical framework to answer other key policy questions. As the previous exercise is done at the aggregate level, in this section, we study the IG and HY corporate bond markets separately, and examine their interactions through dealers' balance sheet constraints. Furthermore, we extend the framework to include corporate credit spreads in the analysis, and shed light on why illiquidity is seemingly priced in corporate bond market. Finally, we show the heterogeneity across dealers have different effects on corporate bond liquidity by studying domestic and foreign dealers.

4.1 Liquidity Contagion

We now investigate how liquidity in different markets interact with each other. In Gromb and Vayanos (2002, 2017), when dealers lose money in one market, they may be forced

Figure 5: Cumulative Shocks to Swaption Implied Volatility and Supply Shocks

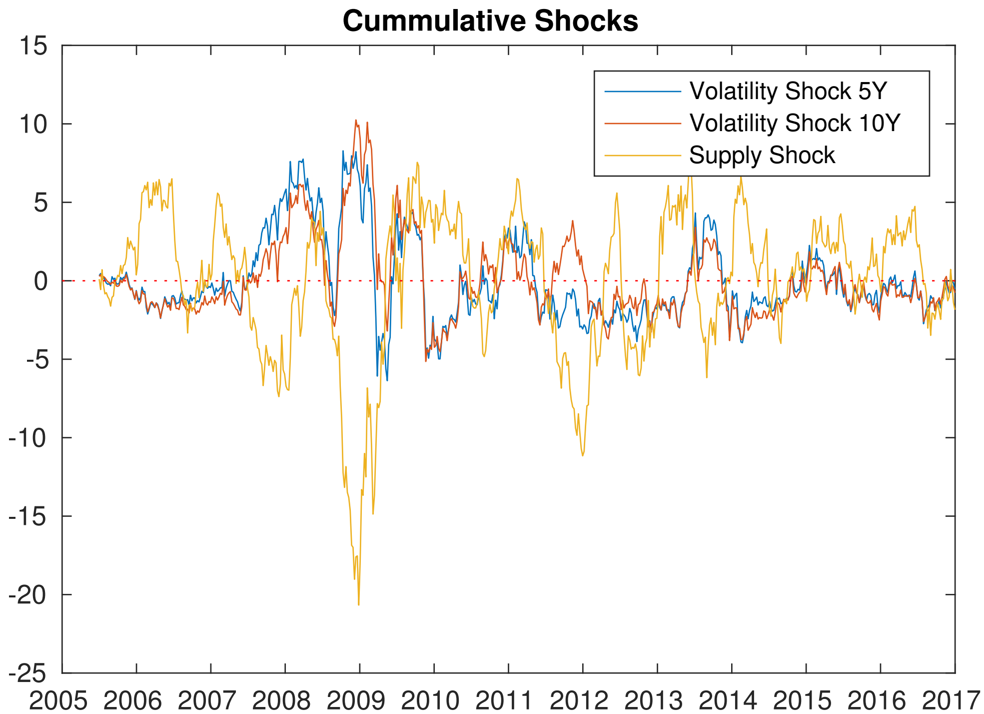

Note: The figure shows the cumulative sum of volatility shocks for swaptions with 1-month expiry on 5-year and 10-year swap rates. Supply shock is the cumulative sum of liquidity supply shocks identified by the sign restriction in (3).

to reduce liquidity supply in the other market due to collateral constraints. Since dealers provide liquidity by buying an undervalued security and selling an overvalued security, an increase in noise leads to a loss for the dealers, shrinking their wealth. Thus, we study whether a rise in noise in one market leads to higher liquidity price and lower quantity in the other market. In particular, we focus on the two segments of corporate bond market defined by credit ratings: investment-grade (IG) and high yield (HY). As Ellul, Jotikasthira and Lundblad (2011) show, these two markets are segmented due to regulations and different investor base. Thus, we construct our liquidity price measure by taking average root mean squared errors across issuers with IG or HY ratings separately. We also aggregate primary dealers' inventory for IG and HY bonds separately to create the quantity measure specific to each rating category.

We run the same VARs as in (2) but with the new vectors of state variables

$$
\begin{aligned}
Y_t^{HY\to IG} &= \begin{pmatrix} p_t^{IG} & q_t^{IG} & p_t^{HY} \end{pmatrix}' \\
Y_t^{IG\to HY} &= \begin{pmatrix} p_t^{HY} & q_t^{HY} & p_t^{IG} \end{pmatrix}'
\end{aligned}
$$

where superscripts IG and HY denote the investment-grade and high yield for each variable, respectively.

To identify structural shocks, we impose the following sign restrictions on the rotation matrix,

$$
\begin{pmatrix} \xi_t^{p,IG} \\ \xi_t^{q,IG} \\ \xi_t^{p,HY} \end{pmatrix} = \underbrace{\begin{pmatrix} - & + & 0 \\ + & + & 0 \\ ? & ? & + \end{pmatrix}}_{A} \begin{pmatrix} v_t^s \\ v_t^d \\ v_t^{HY} \end{pmatrix},
$$

and

$$
\begin{pmatrix} \xi_t^{p,HY} \\ \xi_t^{q,HY} \\ \xi_t^{p,IG} \end{pmatrix} = \underbrace{\begin{pmatrix} - & + & 0 \\ + & + & 0 \\ ? & ? & + \end{pmatrix}}_{A} \begin{pmatrix} v_t^s \\ v_t^d \\ v_t^{IG} \end{pmatrix},
$$

where $\xi_t = Av_t$. The restriction implies that we identify an idiosyncratic shock to HY (IG) liquidity price assuming that on impact, the effect of the shock to the IG (HY) market is close to zero. Specifically, we restrict the ratio of the forecast error variance explained by the HY (IG) shock to the total forecast error variance of the endogenous variables in the IG (HY) market to be less than 0.1. Thus, on impact, the correlation in liquidity between two markets is small by construction. We also expect that in the long run, a price shock in one market has insignificant impact on the liquidity in the other market, as capital eventually flows in to exploit better investment opportunities. Therefore, the hypothesis of liquidity contagion that we test is whether, over the medium term, the liquidity price and quantity in the IG (HY) market respond significantly to an idiosyncratic shock to the noise in the HY (IG) market or not.

In this exercise, we include two liquidity price variables and one quantity variable, as we are not interested in identifying why noise in the other market increases. Rather, we take an increase in noise in the other market as given, and observe how the shock leads to a change in liquidity supply in the market of our interest.

Figure 6 provides the impulse-response functions of IG noise and dealer positions with respect to a shock to HY noise. For the first few months after the shock, the noise in IG bonds increase modestly up to 0.5 basis points while dealer gross positions decline up to 1.5%. Though the magnitude of the reaction of noise does not seem striking, the effect is not negligible compared with the standard deviation of the reduced-form shock to IG noise, which is estimated at 1.7 basis points. Thus, we find an evidence for a modest contagion from a shock to HY bonds to IG bonds through dealers balance sheet.

Figure 7 shows the impulse-response functions of HY noise and dealer positions in response to a shock to IG noise. In this case, the HY noise increases up to 8 basis points over the medium term, while dealer gross positions shrink up to 0.5%. Thus, we find an evidence for contagion across two markets, though the effect is stronger for the contagion from the IG market to HY bond liquidity. The difference in magnitude between IG shocks and HY shocks may be due to the difference in market size. Since the IG bond market is larger than the HY market, a shock to the noise in the IG market may have a greater impact on dealer balance sheets than a shock in the HY noise.

Figure 6: Response of Noise (IG), inventory (IG) and Noise (HY) to a HY shock

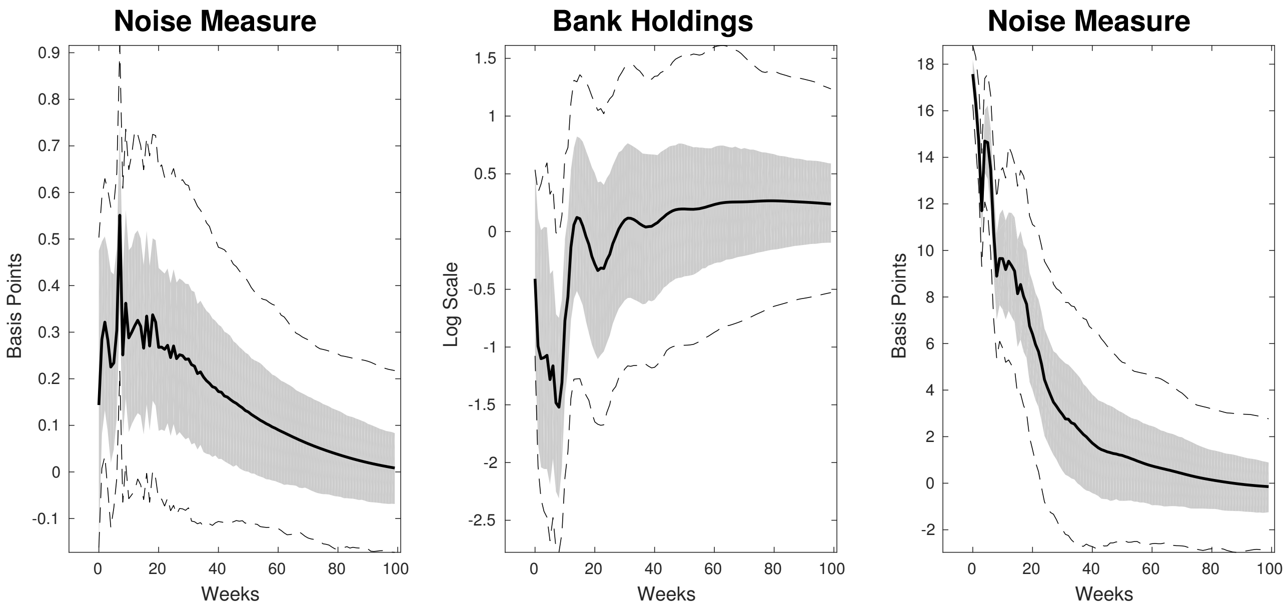

Note: The figures plot the response of endogenous variables to an idiosyncratic shock to the noise in HY bonds. The mean impulse response is shown in black. The shaded area marks a pointwise 68-percent credible interval around the median. The dashed lines mark a pointwise 95-percent credible interval around the median.

Figure 7: Response of Noise (HY), inventory (HY) and Noise (IG) to an IG shock

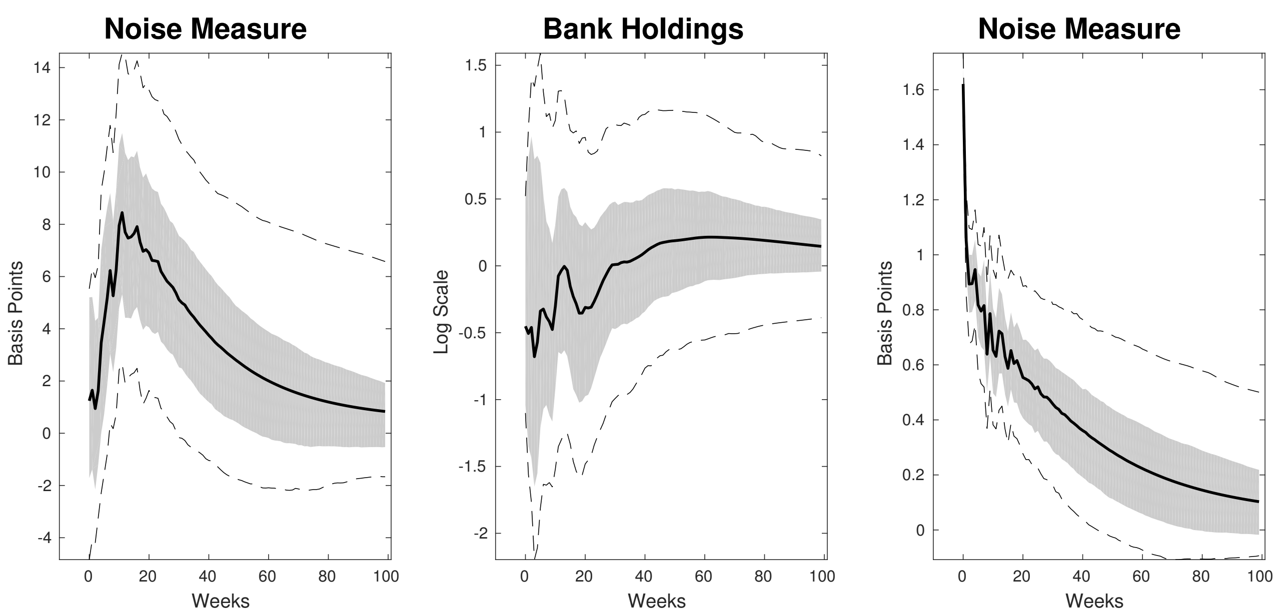

Note: The figures plot the response of endogenous variables to an idiosyncratic shock to the noise in IG bonds. The mean impulse response is shown in black. The shaded area marks a pointwise 68-percent credible interval around the median. The dashed lines mark a pointwise 95-percent credible interval around the median.

4.2 Effect of Liquidity Supply and Demand on Credit Spreads

Numerous papers (e.g. Bao, Pan and Wang (2011) and Feldhütter (2012)) document that liquidity is priced in the corporate bond market. Generally, the previous literature finds that when a measure for corporate bond illiquidity is high, credit spreads for the bond tend to be high as well. With our structural VAR framework, we can dissect deeper into the mechanism through which changing illiquidity affects credit spreads of corporate bonds. In particular, higher credit spreads may correspond to lower liquidity supply or higher liquidity demand. If financial intermediaries are the marginal investor in this market, then a change in their risk aversion could affect the liquidity supply and price of corporate bonds simultaneously. On the other hand, if investors' demand to trade occurs in response to changing credit spreads, liquidity demand may affect credit spreads.

To address this question, we run the same VARs as in (2) but with another vector of state variables

$$Y_t = \begin{pmatrix} p_t^{IG} & q_t^{IG} & cs_t \end{pmatrix}'$$

where cs_t is the Merrill Lynch Baa-rated average credit spreads over Treasury securities (BAMLC0A4CBBB series in Federal Reserve Economic Data). The results using high-yield credit spreads are qualitatively similar with greater effects in magnitude, and thus omitted.

To identify the structural shock, we impose the following sign restrictions on the rotation matrix,

$$\begin{pmatrix} \xi_t^p \\ \xi_t^q \\ \xi_t^{cs} \end{pmatrix} = \underbrace{\begin{pmatrix} - & + & ? \\ + & + & ? \\ ? & ? & + \end{pmatrix}}_{A} \begin{pmatrix} v_t^s \\ v_t^d \\ v_t^{cs} \end{pmatrix}.$$

In this identification scheme, we do not take a stand on how liquidity supply and demand shocks affect credit spreads on impact, as security price may or may not react to liquidity shocks immediately. However, this set of sign restrictions does not provide enough identification assumptions, as a shock that affects all endogenous variables positively can be either v_t^d or v_t^{cs}. Thus, an additional assumption is in order. Rather than restricting the response

Figure 8: Impulse Responses to a Liquidity Supply Shocks

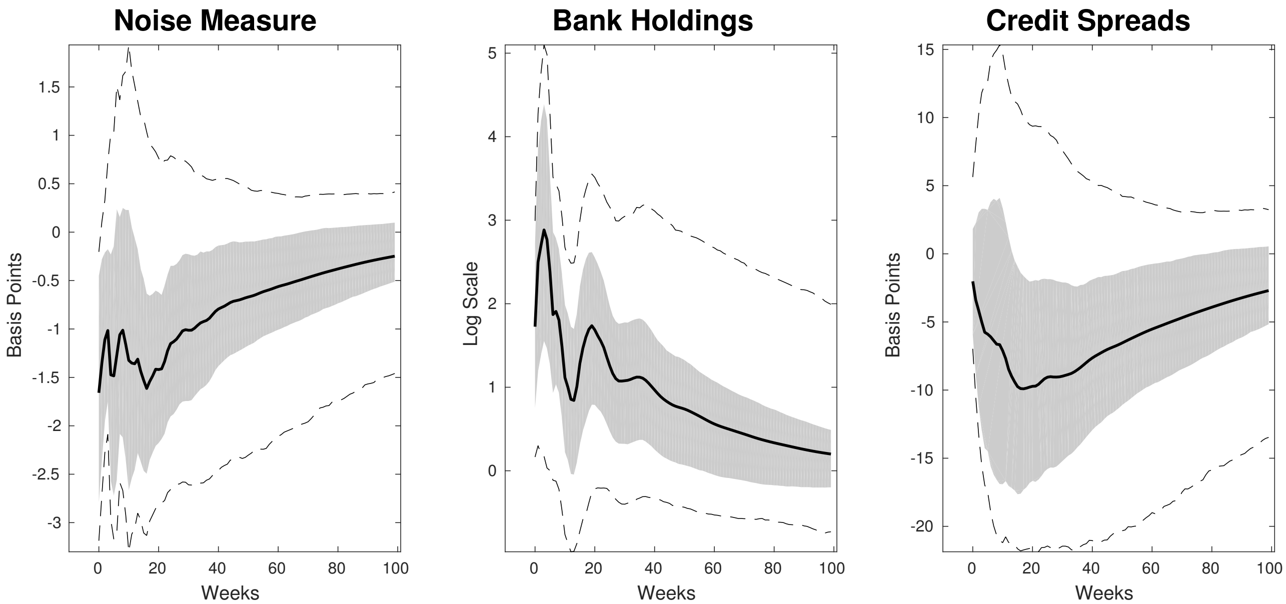

Note: The figures plot the response of endogenous variables to a liquidity supply shocks. The mean impulse response is shown in black. The shaded area marks a pointwise 68-percent credible interval around the median. The dashed lines mark a pointwise 95-percent credible interval around the median.

on impact, we impose a restriction in the long-run response to identify structural shocks. In particular, we assume that liquidity shocks affect bond prices temporarily, as an arbitrage capital eventually flows in and erases the effect of liquidity shocks in the long run. Thus, we impose a restriction that the ratio of the variance of credit spreads explained by v_t^s and v_t^d to the total forecast error variance for credit spreads cannot exceed 25% for any time 52 weeks after the shock. This assumption potentially biases our estimated response of credit spreads to liquidity shocks downward over the medium-term (which is the horizon of our interest), so our estimates provide the lower bound for the response of credit spreads.

Figure 8 shows the impulse-response functions to a liquidity supply shock. As before, the liquidity supply shock lowers the noise and increases the dealer inventory temporarily. In addition, the liquidity supply shock has a significant negative effect on credit spreads over the medium term, and the credit spreads decrease up to 10 basis points in 4 months after the shock. These results suggest that, when dealers become risk averse, not only liquidity supply declines but also credit spreads for the corporate bond rises.

Figure 8 shows the impulse-response functions to a liquidity demand shock. In contrast

Figure 9: Impulse Responses to a Liquidity Demand Shocks

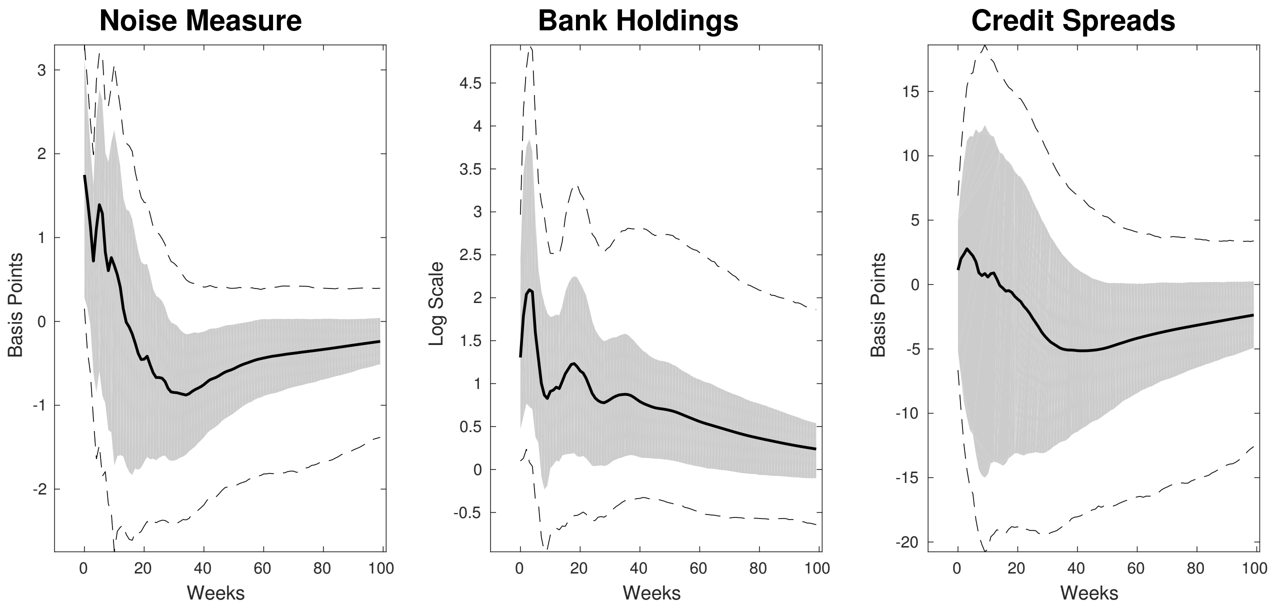

Note: The figures plot the response of endogenous variables to a liquidity demand shocks. The mean impulse response is shown in black. The shaded area marks a pointwise 68-percent credible interval around the median. The dashed lines mark a pointwise 95-percent credible interval around the median.

to the responses to a liquidity supply shock, the liquidity demand shock does not affect credit spreads significantly over the medium term. Therefore, we conclude that the observed relationship between illiquidity measures and credit spreads mainly comes from the supply channel, and risk-aversion of financial intermediary is important in understanding the dynamics of credit spreads over the medium-term horizon.

4.3 Domestic versus Foreign Dealers

As Adrian, Boyarchenko and Shachar (2017) and Friewald and Nagler (2016) show, the heterogeneity across dealers can affect bond markets differently. In particular, the difference between domestic and foreign dealers may shed light on how the heterogeneity in risk aversion or financial constraints across dealers affect the liquidity in the US corporate bond market. Foreign dealers are likely to have different market coverage and exposures than domestic dealers, and thus hit by shocks differently. To investigate the heterogeneity across dealers, we aggregate the bond holdings by domestic and foreign dealers separately. Dealers are classified into domestic or foreign based on their headquarter location. Appendix E shows the classification of domestic and foreign dealers.

We run the same VARs as in (2) but with the new vectors of state variables

$$Y_t = \begin{pmatrix} p_t & q_t^D & q_t^F \end{pmatrix}'$$

where superscript D and F denote the domestic and foreign for each variable, respectively.

To identify structural shocks, we impose the following sign restrictions on the rotation matrix,

$$\begin{pmatrix} \xi_t^p \\ \xi_t^{q,D} \\ \xi_t^{q,F} \end{pmatrix} = \underbrace{\begin{pmatrix} - & + & - \\ + & + & 0 \\ 0 & + & + \end{pmatrix}}_{A} \begin{pmatrix} v_t^{s,D} \\ v_t^d \\ v_t^{s,F} \end{pmatrix},$$

where $\xi_t = Av_t$. The restriction implies that we identify an idiosyncratic shock to domestic supply shock assuming that on impact, the effect of the domestic supply shock to the foreign dealers balance sheet is close to zero. As in the previous analysis on liquidity contagion, we restrict the ratio of the forecast error variance explained by the domestic (foreign) shock to the total forecast error variance of the gross positions for foreign (domestic) dealers to be less than 0.1. Thus, on impact, the correlation in liquidity supply shocks between two types of dealers is small by construction.

Figure 10 shows the cumulative supply shocks for domestic dealers, $\Sigma_{j=1}^t v_j^{s,D}$, and for foreign dealers, $\Sigma_{j=1}^t v_j^{s,F}$. The two supply shocks share broadly a similar pattern over time until 2013, though there is some interesting difference in response to notable economic events. For example, right after the bankruptcy of Lehman Brothers, the liquidity supply of domestic dealers increases slightly while the foreign supply falls sharply. Moreover, when the Greek government agreed the first bailout plan with the IMF amid the violent protests in May 2010, foreign dealers are hit by negative supply shocks more sharply than domestic counterparts. In contrast, the taper tantrum in 2013 led to a series of sharp negative supply shocks for domestic dealers, while foreign dealers are not affected as much. Interestingly, after 2014, domestic dealers a series of positive supply shocks and maintain the level of inventory in corporate bonds, while foreign dealers reduce their inventory, which we interpret as receiving

Figure 10: Cumulative Supply Shocks for Domestic and Foreign Banks

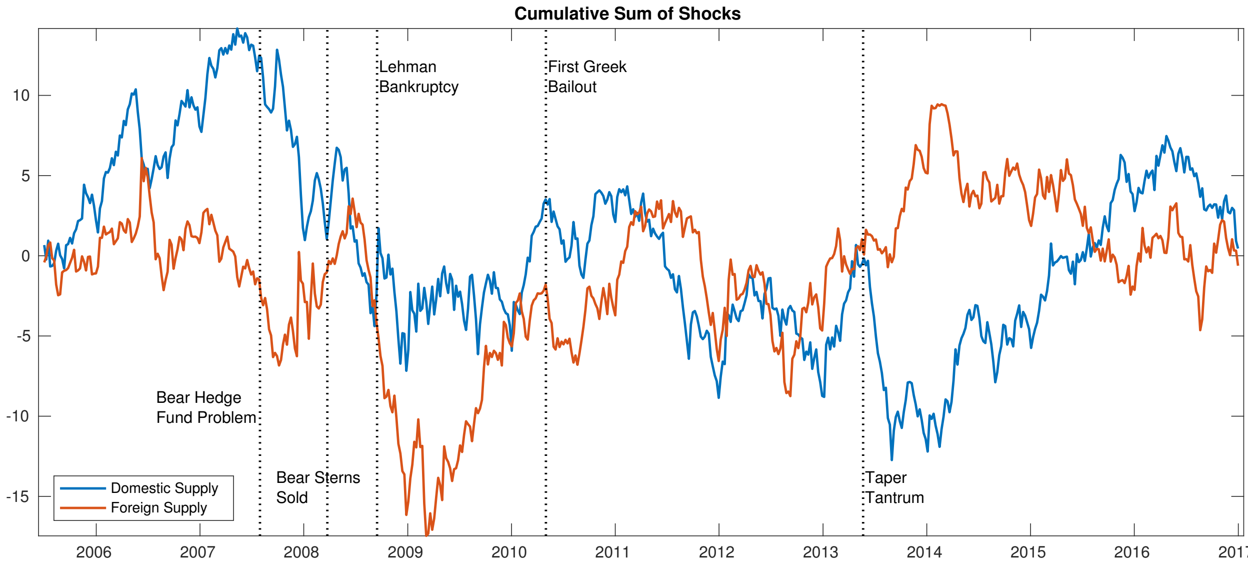

Note: This figure shows the mean of the posterior distribution of the cumulative sum of shocks to domestic and foreign liquidity supplies. A positive shock reflects an increase in liquidity supply. The data is weekly from 2005 to 2017.

negative supply shocks. Indeed, much of the decline in primary dealers' balance sheet for corporate bonds since 2014 seems to come from foreign dealers rather than domestic ones.

5 Conclusion

In this paper, we provide an analytical framework to extract the supply shocks to liquidity by jointly studying the price and quantity measures of liquidity. As dealers use their balance sheet to provide liquidity and thereby correct the deviation of bond prices from the reference yield curve, the noise in yield curve and dealers' gross inventory of bonds are naturally connected by supply and demand shocks. In particular, by imposing reasonable sign restrictions on the impulse-response functions of a shock, we extract the liquidity supply shocks to the bond market. Namely, a positive supply shock must increase dealers' gross position on corporate bonds while decreasing the noise in yield curves.

The decomposition of supply and demand shows that the liquidity supply in the corporate bond market declines in the financial crisis in 2008 and recovers partly afterwards. The liquidity demand spikes up during the crisis, but decreases over the recent period. The decline in liquidity demand partly explains improving traditional liquidity measures (e.g., the narrowing bid-ask spreads), and we estimate that absent demand shocks, the noise in the corporate bond market after 2014 would have been higher by around 10 basis points. In addition, we find an evidence that elasticity of liquidity supply declines after the financial crisis relative to the pre-crisis level, implying that liquidity risk has risen despite the stable level of illiquidity. Overall, our findings are consistent with the narrative that newly introduced banking regulations post crisis make brokers less willing to provide liquidity. The decomposition also shows that about half of the residual variance of noise comes from the liquidity supply shock. In addition, we show our estimates of the supply shocks are correlated with swaption-implied volatilities while the demand shocks are associated with corporate bond mutual fund flows.

Even though this paper focuses on the corporate bond market, our emphasis on disentangling the supply effects from observable liquidity proxies is general and applicable to other markets. As a possible extension of our framework, we can use even richer cross-section of bonds, such as Treasury securities and MBS, to understand the contagion effect across markets. In addition, we do not fully exhaust the list of the drivers behind the changing supply and demand of corporate bond liquidity. Our findings are consistent with the argument in the literature that the Volker rule and other financial regulations lead to a deterioration in liquidity, but there may be other factors that moved over the same period. Establishing the causation between changing regulation and liquidity in the bond market is also left for future research.

References

Adrian, Tobias, Nina Boyarchenko, and Or Shachar. 2017. "Dealer Balance Sheets and Bond Liquidity Provision." *Journal of Monetary Economics, forthcoming.*

Amihud, Yakov. 2002. "Illiquidity and stock returns: cross-section and time-series effects." *Journal of financial markets*, 5(1): 31–56.

Anand, Ambar, Chotibhak Jotikasthira, and Kumar Venkataraman. 2017. "Do Buy-Side Institutions Supply Liquidity in Bond Markets? Evidence from Mutual Funds." *Working Paper.*

Asquith, Paul, Andrea S. Au, Thomas R. Covert, and Parag A. Pathak. 2013. "The market for borrowing corporate bonds." *Journal of Financial Economics,*, 107(1): 155–182.

Bai, Jennie, and Pierre Collin-Dufresne. 2013. "The Bond-CDS Basis." *Working Paper.*

Bao, Jack, Jun Pan, and Jiang Wang. 2011. "The Illiquidity of Corporate Bonds." *The Journal of Finance*, 66(3): 911–946.

Bao, Jack, Maureen O'Hara, and Xing Zhou. 2017. "The Volcker rule and corporate bond market-making in times of stress." *Journal of Financial Economics, forthcoming.*

Baranova, Yuliya, Louisa Chen, and Nicholas Vause. 2015. "Has corporate bond makret liquidity fallen?" *Bank Underground, Bank of England.*

Bessembinder, Hendrik, Kathleen M. Kahle, William F. Maxwell, and Danielle Xu. 2009. "Measuring Abnormal Bond Performance." *The Review of Financial Studies*, 22(10): 4219–4258.

Bessembinder, Hendrik, Stacey E. Jacobsen, William F. Maxwell, and Kumar Venkataraman. 2017. "Capital Commitment and Illiquidity in Corporate Bonds." *Working Paper.*

Chen, Long, David A. Lesmond, and Jason Wei. 2007. "Corporate Yield Spreads and Bond Liquidity." *The Journal of Finance*, 62(1): 119–149.

Choi, Jaewon, and Yesol Huh. 2017. "Customer Liquidity Provision: Implications for Corporate Bond Transaction Costs." *Working Paper.*

Cohen, Lauren, Karl B Diether, and Christopher J Malloy. 2007. "Supply and demand shifts in the shorting market." *The Journal of Finance*, 62(5): 2061–2096.

Coval, Joshua D., and Erik Stafford. 2007. "Asset Fire Sales (and Purchases) in Equity Markets." 86: 479–512.

Culbertson, J. M. 1957. "The Term Structure of Interest Rates." *The Quarterly Journal of Economics*, 71(4): 485–517.

Dick-Nielsen, Jens, Peter Feldhütter, and David Lando. 2012. "Corporate bond liquidity before and after the onset of the subprime crisis." *Journal of Financial Economics*, 103(3): 471 – 492.

Duffie, Darrell, and Kenneth Singleton. 1999. "Modeling Term Structures of Defaultable Bonds." *Review of Financial Studies*, 12(4): 687–720.

Edwards, Amy K., Lawrence E. Harris, and Michael S. Piwowar. 2007. "Corporate Bond Market Transaction Costs and Transparency." *The Journal of Finance*, 62(3): 1421–1451.

Ellul, Andrew, Chotibhak Jotikasthira, and Christian Lundblad. 2011. "Regulatory pressure and fire sales in the corporate bond market." *Journal of Financial Economics*, 101(3): 596–620.

Elsaify, Amora, and Nikolai Roussanov. 2016. "Why Do Firms Issue Callable Bonds?" *Working Paper.*

Feldhütter, Peter. 2012. "The Same Bond at Different Prices: Identifying Search Frictions and Selling Pressures." *Review of Financial Studies*, 25: 1155–1206.

Friewald, Nils, and Florian Nagler. 2016. "Dealer Inventory and Cross-Section of Corporate Bond Returns." *Working Paper.*

Friewald, Nils, Rainer Jankowitsch, and Marti G. Subrahmanyam. 2012. "Illiquidity or credit deterioration: A study of liquidity in the US corporate bond market during financial crises." *Journal of Financial Economics*, 105(1): 18–36.

Goldberg, Jonathan. 2017. "The Supply of Liquidity and Real Economic Activity." *Working Paper.*

Goldstein, Michael, and Edith Hotchkiss. 2017. "Dealer Behavior in Illiquid Markets." *Working Paper.*

Gromb, Denis, and Dimitri Vayanos. 2002. "Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs." *Journal of Financial Economics*, 66(2): 361–407.

Hu, Grace Xing, Jun Pan, and Jiang Wang. 2013. "Noise as Information for Illiquidity." *The Journal of Finance*, 68(6): 2341–2382.

Longstaff, Francis A., Sanjay Mithal, and Eric Neis. 2005. "Corporate Yield Spreads: Default Risk or Liquidity? New Evidence from the Credit Default Swap Market." *The Journal of Finance*, 60(5): 2213–2253.

Malkhozov, Aytek, Philippe Mueller, Andrea Vedolin, and Gyuri Venter. 2016. "International Illiquidity." Working paper.

Modigliani, Franco, and Richard Sutch. 1966. "Innovations in Interest Rate Policy." *The American Economic Review*, 56(1/2): 178–197.

Pedersen, Lasse Heje. 2015. *Efficiently Inefficient.* Princeton, NJ:Princeton University Press.

Rubio-Ramirez, Juan F., Daniel F. Waggoner, and Tao Zha. 2010. "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference." *The Review of Economic Studies*, 77(2): 665–696.

Trebbi, Francesco, and Kairong Xiao. 2015. "Regulation and market liquidity." National Bureau of Economic Research.

Uhlig, Harald. 2005. "What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure." *Journal of Monetary Economics*, 52(2): 381–419.

Uhlig, Harald. 2015. "Shocks, sign restrictions and identification." University of Chicago, mimeo.

Vayanos, Dimitri, and Jean-Luc Vila. 2009. "A preferred-habitat model of the term structure of interest rates." National Bureau of Economic Research.

A A Simple Model of Noise in Bond Prices

This appendix, based on Goldberg (2017), develops a highly stylized model of security prices and dealer positions. In the model, there are two securities that represent claims to the same long-term cash flow. However, the securities trade in segmented markets, potentially at different prices; dealers hold long and short positions in the securities to partially overcome market segmentation. The noise in the security prices compensates dealers for making markets.

There are: two investors, A investors and B investors; two securities, A securities and B securities; and three periods, t=1, 2, 3. A and B securities represent a claim to an uncertain cash flow v in period 3.[4]

In period 1, A and B investors have complementary trading needs: A and B investors receive endowment shocks in period 3 that are equal in magnitude but opposite in sign and these endowment shocks are correlated with the cash flow v. However, the markets are segmented: investor A is only able to trade A securities and investor B is only able to trade B securities.[5] Hence, gains from trade between the investors can only be realized by trading through a dealer. Market making by dealers involves risk: in period 2, intermediaries may be forced to liquidate their positions at uncertain prices. As a result, unless dealers are risk neutral, the securities will trade at different prices.

i-investors, with $i \in \{A, B\}$, can trade only in the i-bond and money. Dealers can trade in both markets and money.[6] Financial markets are competitive. At $t = 1$, investors and dealers trade in the i-markets. The period-t price of the i-security is $p_{i,t}$. The gross interest rate is normalized to one. The A and B securities each have net supply g.

[4]It is straightforward to slightly modify the model to allow the securities to be interpreted as Treasury bonds. Specifically, add a fourth period to the model, in which the securities mature with known value equal to one. Assume there is a perfectly elastic supply of central bank reserves at the exogenous interest rate and that the interest rate between periods 3 and 4 is a random variable R revealed in period 3. Then $v = \frac{1}{R}$.

[5]Examples of investors that are willing to trade only a particular security include, per Pedersen (2015), "price-insensitive insurance companies who need [a given bond] for a specific reason." A related example is an investor who owns a particular security and no longer wants the associated risk may not be willing to shed that risk by going short a different security with the same payoff; she simply wants to sell, even if doing so is more expensive. Alternatively, consider a sophisticated investor who has shorted a particular security by obtaining the security on loan and then selling it; when the investor wants to close out the trade, she needs to buy back that particular security in order to return it to the lender of the security.

[6]In an appendix (available upon request), I modify the model to include non-dealer intermediaries that, like dealers, are able to trade in both securities markets; I show that the main results of the model still hold.

The mean of the cash flow v, conditional on period 1 information, is denoted by μ. The variance is denoted by σ. That is,

$$E[v] = \mu$$

and

$$Var[v] = \sigma.$$

The cash flow v is revealed in period 2. Also, with probability λ, dealers are forced to liquidate their positions at uncertain prices: $p_{i,2} = v + \epsilon_i$, where ϵ_A and ϵ_B have variance σ_ϵ.

i-investors have mean-variance preferences over period-3 wealth w_i. That is, i-investors maximize $E[w_i] - \frac{1}{2\gamma} Var[w_i]$, where γ is i-investors' risk tolerance.[7] Dealers also have mean-variance preferences. The risk tolerance of dealers is denoted by γ_D. This risk tolerance γ_D is a proxy for liquidity supply.

The i-investors have a motivation to hedge. In particular, $e_A = -e_B$ and $Cov(v, e_A) = u > 0$.

I denote the period-1 position of dealers in the i-security by x_i; the period-1 position of the i-investor in the i-security is y_i. I denote the period-1 risk premia by ψ, where the i-th element of ψ is:

$$\psi_i = \mu - p_{i,1} \tag{5}$$

At $t = 3$, i-investors receive endowment e_i.

The cash flow v, the liquidation price shocks ϵ_A and ϵ_B, and the realization of the liquidation event are mutually independent. Also, the liquidation price shocks ϵ_A and ϵ_B, the realization of the liquidation event and the endowment e_A are mutually independent.

Define

$$g^* = \left(1 + \frac{2\gamma\sigma}{\gamma_D\sigma + \gamma\lambda\sigma_\epsilon}\right)\frac{u}{\sigma} > 0.$$

I assume that $|g| < g^*$. This assumption guarantees that, in equilibrium, the dealer has a strictly positive position in security B and a strictly negative position in security A, consistent with the role of a market maker.

[7]Without loss of generality, i-investors and dealers have zero initial wealth.

Equilibrium. For dealers, the variance-covariance matrix of the payoffs associated with the A and B securities is given by:

$$\Omega = \begin{bmatrix} \sigma + \lambda\sigma_\epsilon & \sigma \\ \sigma & \sigma + \lambda\sigma_\epsilon \end{bmatrix} \tag{6}$$

The vector of dealers' demand, $x = [x_A \; x_B]'$, is given by:

$$x = \Omega^{-1}\gamma_D\psi \tag{7}$$

and the vector of investors' demand, $y = [y_A \; y_B]'$, is:

$$y = \frac{1}{\sigma}\left(\gamma\psi - u\begin{bmatrix} 1 \\ -1 \end{bmatrix}\right). \tag{8}$$

Market clearing requires that

$$x + y = g \tag{9}$$

There is a unique equilibrium. Define price dispersion as $|p_{B,1} - p_{A,1}|$ and dealer gross positions as $|x_A| + |x_B|$. Then,

$$\frac{|p_{B,1} - p_{A,1}|}{2} = \frac{1}{\gamma_D \frac{1}{\lambda}\frac{\sigma}{\sigma_\epsilon} + \gamma}u$$

and

$$\frac{|x_A| + |x_B|}{2} = \frac{1}{\sigma + \frac{\gamma}{\gamma_D}\lambda\sigma_\epsilon}u$$

Proposition 1. *An increase in dealer risk tolerance γ_D leads to lower price dispersion and higher dealer gross positions. That is, $\frac{d\left[|p_{B,1}-p_{A,1}|\right]}{d\gamma_D} < 0$ and $\frac{d[|x_A|+|x_B|]}{d\gamma_D} > 0$. An increase in investor risk tolerance γ or a decrease in investor trading needs u also leads to lower price dispersion; however, dealer gross positions decrease.*

Proposition 1 reflects the intuition behind the sign restrictions discussed in the main text. Changes in dealer risk tolerance or liquidation risks (the probability of liquidation λ and the riskiness of liquidation prices κ) lead to opposite-signed changes in the dispersion of bond prices and dealer gross positions. Changes in liquidity demand u lead to same-sign changes.

B Comparing Merrill Lynch and TRACE Data

In this section, we compare the price in Merrill Lynch data and TRACE data using the month-end overlapping observations from 2005 to 2016. By merging the two data sets, we have 229,228 bond-month observations. For a price in TRACE, we follow Bessembinder et al. (2009) to compute volume-weighted average prices using transactions with volume above 100,000 dollars. For each observation, we compute yield to maturity based on two prices.

Panel A of Table 7 reports the average yield to maturity from the two data sets for each credit rating and maturity. The average yields are very similar between two data sets, though Merrill Lynch data has slightly higher yields for short- and medium-term bonds. Such a difference can be easily accounted for by either level or slope of yield curves in the Nelson-Siegel-Svensson model.

Since the focus of this article is on the liquidity effect and noise in corporate bond prices, it is also important to compare the two samples under illiquid market conditions. To this end, we compute the average for the subsample using year-end price observations. Panel B of Table 7 reports the average yield to maturity for the year-end observations. The yields at the end of years are also fairly close to each other, and thus the noise constructed using Merrill Lynch data is likely to capture the mispricing due to illiquidity as well as TRACE data does.

C Reduced-Form VAR Estimates

In this section, we report the mean and credible sets for the posterior distributions for the reduced-form parameters in (2). Columns 2 to 3 of Table 8show the OLS estimates for the slope and error covariance matrix of the VAR, while Columns 4 and 5 show the mean posterior draws from the Bayesian estimates. Both estimates are close to each other,

Table 7: Comparing Yield to Maturity Between Merrill Lynch and TRACE Data

	Merrill Lynch				TRACE			
	-4yr	4-7yr	7-12yr	12yr-	-4yr	4-7yr	7-12yr	12yr-
Panel A. Average, Full Sample								
AAA	3.39	4.03	4.40	4.97	3.31	3.99	4.35	4.96
AA	2.99	3.86	4.55	5.12	2.92	3.81	4.51	5.10
A	2.88	3.76	4.51	5.31	2.82	3.72	4.48	5.29
BBB	3.34	4.28	4.92	5.87	3.28	4.24	4.89	5.85
HY	11.18	9.57	8.12	9.36	11.01	9.44	8.06	9.25
Panel B. End of Year Only								
AAA	3.24	4.24	4.54	4.83	3.08	4.07	4.44	4.78
AA	3.03	3.82	4.55	5.08	2.91	3.72	4.46	5.04
A	2.85	3.83	4.70	5.39	2.74	3.71	4.61	5.33
BBB	4.00	4.50	5.24	6.14	3.86	4.37	5.16	6.11
HY	16.34	11.81	8.86	12.95	16.05	11.70	8.78	12.46

Note: The table reports the average yield to maturity for each credit rating and maturity using the overlapping month-end observations between Merrill Lynch and TRACE from 2005 to 2016. To obtain prices in TRACE, we follow Bessembinder et al. (2009) and compute volume-weighted average price on the day, using transactions above 100,000 dollars.

confirming that the prior we use in estimating the VAR is indeed weak and do not affect our statistical inferences.

D Effects of Liquidity Supply and Demand Shocks

Figures 11 and 12 show the impulse responses to liquidity supply and demand shocks. By assumption, on impact, dealer gross positions weakly increase after both types of shocks, but noise weakly decreases only following a liquidity supply shock. Figures 11 and 12 show that on impact, the magnitude of the shocks are similar to each other. Inspecting the subsequent reactions of endogenous variables, the reactions of the noise measure to a supply shock is persistent, and the effect slowly decays over the two-year period. In contrast, the response to a demand shock is less persistent and reverts to zero in about 20 weeks. The response of gross dealer holdings has a pronounced reactions over the first ten weeks, followed by small but persistent effects.

The persistent response of dealer holdings in response to a negative supply shock suggests the existence of slow-moving capital in the corporate bond market. When a financial

Table 8: Parameter Estimates for the Reduced-Form VAR

	OLS Estimates		Mean Posterior Draws		Credible Sets (16%)		Credible Sets (84%)	
	p_t	q_t	p_t	q_t	p_t	q_t	p_t	q_t
b	50.734	0.181	51.281	0.171	38.060	0.069	67.122	0.278
B_1	0.906	0.000	0.915	-0.001	0.860	-0.001	0.962	0.000
	1.792	1.419	1.869	1.430	-2.618	1.390	6.436	1.470
B_2	0.003	0.001	-0.011	0.001	-0.079	0.000	0.061	0.001
	-7.543	-0.421	-8.900	-0.432	-19.863	-0.497	-1.357	-0.368
B_3	-0.049	-0.001	-0.054	-0.001	-0.108	-0.001	0.005	0.000
	-7.957	-0.020	-6.246	-0.025	-14.385	-0.102	0.597	0.043
B_4	0.323	0.001	0.341	0.001	0.286	0.000	0.395	0.001
	26.002	-0.074	24.924	-0.070	17.025	-0.153	32.541	0.014
B_5	-0.037	-0.001	-0.055	-0.001	-0.106	-0.001	0.003	0.000
	-6.290	-0.062	-5.919	-0.065	-12.482	-0.153	1.650	0.010
B_6	-0.138	0.001	-0.134	0.001	-0.198	0.000	-0.069	0.001
	0.382	0.218	1.611	0.227	-6.707	0.140	9.581	0.342
B_7	-0.078	-0.001	-0.082	-0.001	-0.147	-0.002	-0.020	-0.001
	-24.974	-0.034	-26.150	-0.036	-36.603	-0.125	-17.604	0.041
B_8	-0.065	0.001	-0.047	0.001	-0.111	0.001	0.020	0.002
	19.202	-0.152	20.416	-0.148	13.851	-0.223	27.021	-0.060
B_9	0.096	0.000	0.082	0.000	0.025	0.000	0.143	0.001
	6.275	0.029	5.552	0.020	-2.051	-0.040	14.234	0.099
B_{10}	-0.068	0.000	-0.049	0.000	-0.104	0.000	0.000	0.001
	-12.501	0.032	-13.601	0.034	-21.332	-0.028	-6.111	0.099
B_{11}	0.105	0.000	0.100	0.000	0.031	-0.001	0.177	0.000
	3.463	0.125	5.268	0.130	-3.935	0.066	13.323	0.189
B_{12}	0.036	0.000	0.026	0.000	-0.035	0.000	0.098	0.001
	-5.437	-0.079	-7.078	-0.084	-16.283	-0.157	3.115	0.003
B_{13}	-0.094	0.000	-0.091	0.000	-0.145	-0.001	-0.042	0.000
	4.701	0.009	5.340	0.009	0.531	-0.038	9.941	0.050
Σ	12.375	0.005	12.314	0.005	11.591	0.000	13.081	0.010
	0.005	0.001	0.005	0.001	0.000	0.001	0.010	0.001

Figure 11: Impulse Responses to a Liquidity Supply Shock

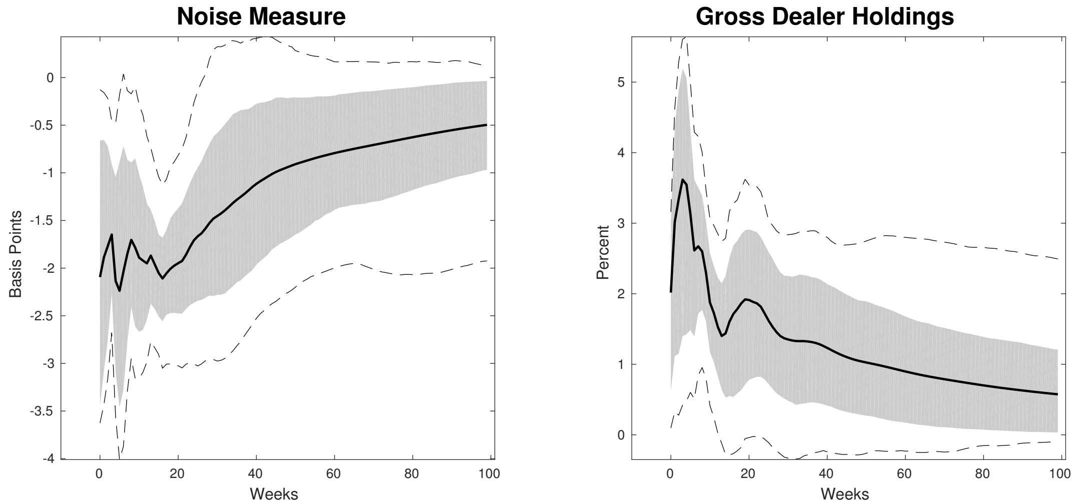

Note: The mean impulse response is shown in black. The shaded area marks a pointwise 68-percent credible interval around the median. The dashed lines mark a pointwise 95-percent credible interval around the median.

intermediary becomes more risk averse or more financially constrained, market liquidity deteriorates, making the market making activities more profitable on average. As a result, a new capital flows in to take advantage of the increased noise, and the gross dealer holding reverts to the original level. However, the recovery of the dealer holding is slow, and it takes more than a year to fully mean revert.

With the structural VAR framework, we measure the economic significance of each structural shock by quantifying the contribution of each shock to the total forecast error variance of endogenous variables, $\sigma^2(\xi)$. Figure 13 shows the share of forecast error variance due to the liquidity supply shock. The share of forecast error variance explained by liquidity supply shocks is more than half for both noise and dealer gross positions both at short and long horizons. In particular, in the first few quarters, the supply shock accounts for about two-thirds of the forecast error variance for the noise. Since neither supply nor demand dominates in explaining the shocks to endogenous variables, the result highlights the importance of isolating supply effects from demand effects when assessing changes in liquidity.

Figure 12: Impulse Responses to a Liquidity Demand Shock

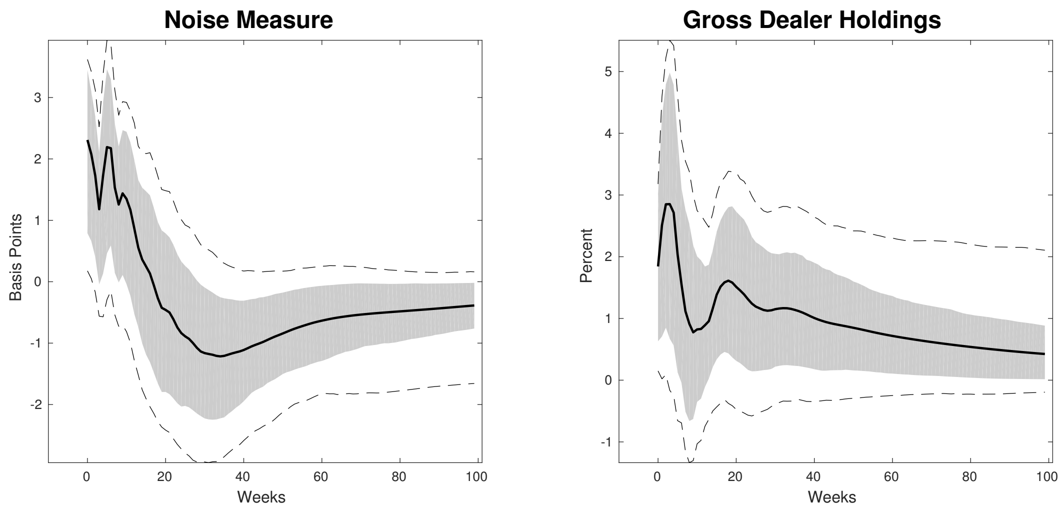

Note: The mean impulse response is shown in black. The shaded area marks a pointwise 68-percent credible interval around the median. The dashed lines mark a pointwise 95-percent credible interval around the median.

Figure 13: Forecast Error Variance Decomposition

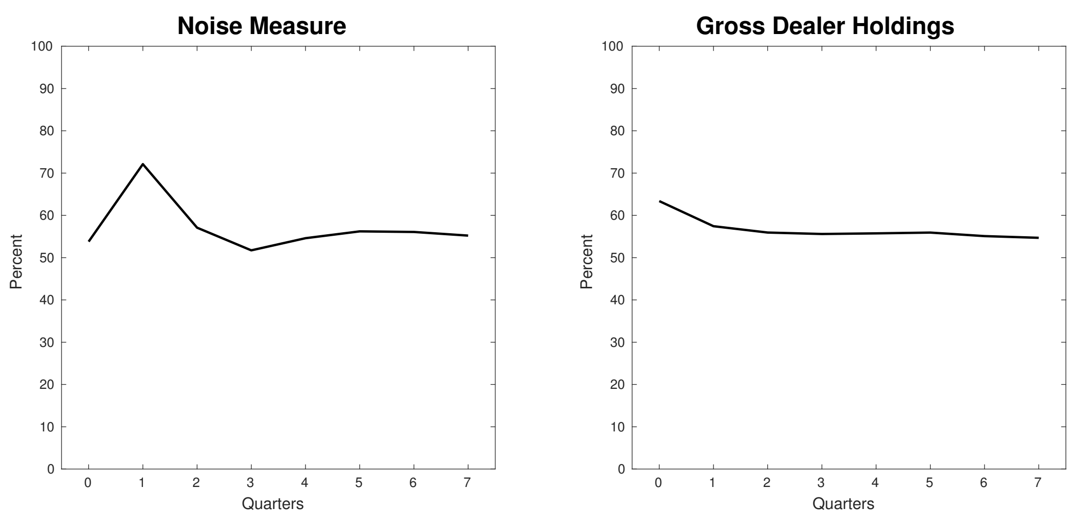

Note: Each plot shows the share of forecast error variance for a given variable due to the liquidity supply shock. The forecast error variance decomposition is calculated by: drawing the parameters of the structural VAR from the posterior distribution; calculating the forecast error variance decomposition at different time horizons for each draw; and then taking the mean across draws and across weeks within a given quarter.

Table 9: List of Primary Dealers

Primary Dealer	Start Date	End Date	Foreign Dealer
BANC OF AMERICA SECURITIES LLC	19990517	20101101	NO
BANK OF NOVA SCOTIA	20111004	N.A.	YES
BARCLAYS CAPITAL INC.	19980401	N.A.	YES
BMO CAPITAL MARKETS CORP.	20111004	N.A.	YES
BNP PARIBAS SECURITIES CORP.	20000915	N.A.	YES
CANTOR FITZGERALD & CO.	20060801	N.A.	NO
CIBC WORLD MARKETS CORP.	19990503	20070208	YES
CITIGROUP GLOBAL MARKETS INC.	20030407	N.A.	NO
CREDIT SUISSE SECURITIES (USA) LLC	20030117	N.A.	YES
DAIWA CAPITAL MARKETS AMERICA INC.	19861211	N.A.	YES
DEUTSCHE BANK SECURITIES INC.	20020330	N.A.	YES
GOLDMAN, SACHS & CO.	19741204	N.A.	NO
HSBC SECURITIES (USA) INC.	19990601	N.A.	YES
J.P. MORGAN SECURITIES LLC	20010501	N.A.	NO
JEFFERIES LLC	20090618	N.A.	NO
MERRILL LYNCH GOVERNMENT SEC. INC.	19600519	20090211	YES
MERRILL LYNCH, PIERCE, FENNER & SMITH INCORPORATED	19600519	N.A.	NO
MF GLOBAL	20110202	20111031	YES
MIZUHO SECURITIES USA INC.	20020401	N.A.	YES
MORGAN STANLEY & CO. LLC	19780201	N.A.	NO
NOMURA SECURITIES INTERNATIONAL,INC	19861211	N.A.	YES
RBC CAPITAL MARKETS, LLC	20090708	N.A.	YES
RBS SECURITIES INC.	20090401	N.A.	YES
SG AMERICAS SECURITIES, LLC	20110202	N.A.	YES
TD SECURITIES (USA) LLC	20140211	N.A.	YES
UBS SECURITIES LLC.	20030609	N.A.	YES

E List of Primary Dealers

Table 9 shows the list of firms designated as primary dealers by the Federal Reserve Bank of New York.