Liquidity Supply and Demand in the Corporate Bond Market

Jonathan Goldberg* Yoshio Nozawa†

Federal Reserve Board HKUST

July 20, 2018

Abstract

We identify shocks to liquidity supply and demand in the corporate bond market by jointly analyzing dealer capital commitment and bond prices. Liquidity supply falls after the introduction of recent banking regulations. Supply risk is priced in the cross-section of corporate bond returns, implying that investors seek to hedge against low liquidity supply. Dealer capital commitment predicts aggregate bond returns negatively when fluctuations in capital commitment are mostly supply-driven but not when they are demand-driven, underscoring the importance of disentangling liquidity supply and demand.

JEL Classification: G12, G24

Keywords: Corporate bond liquidity and returns, dealer intermediation, return predictability

---

*Federal Reserve Board.

†HKUST. We thank Patrick Augustin (discussant), Jack Bao (discussant), Hui Chen, Valentin Haddad, Yesol Huh, Xiaoxia Lou (discussant), Maureen O'Hara, Or Shachar, Quan Wen (discussant), Kairong Xiao, Alex Zhou and seminar participants at Asian FA Conference, Baltimore Area Finance Conference, CICF, Federal Reserve Board, HEC-McGill Winter Finance Conference, ITAM Finance Conference, Penn State University and the University of California at San Diego 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: ynchoir@gmail.com

1 Introduction

Since the financial crisis in 2008, the drivers of market liquidity and their relation with asset prices have garnered new attention. The literature often highlights financial constraints, regulatory costs, and funding risks that limit liquidity suppliers' ability or willingness to provide liquidity as the key driver. However, variation in common liquidity measures such as dealer capital commitment, turnover, or bid-ask spreads do not only reflect changes in liquidity supply due to such constraints but also liquidity demand by investors. For example, on October 15, 2014, when the US Treasury market experienced a flash rally, record-high trading volume pointed not to a high level of liquidity supply but rather a liquidity shortage. Moreover, the relation between asset prices and common liquidity measures, such as dealer leverage, may be unstable. In the corporate bond market, Adrian, Boyarchenko and Shachar (2017) find that the sign of the relationship between dealer leverage and market liquidity was different before and after the financial crisis.

In this paper, we develop a new method to measure shifts in liquidity supply and demand in the corporate bond market and use this method to study the asset pricing implications of liquidity supply risk and liquidity demand risk. Our method of disentangling supply and demand is based on models of intermediation in segmented markets (Gromb and Vayanos (2017)) and, more generally, models of market illiquidity due to inventory frictions (Grossman and Miller (1988) and Brunnermeier and Pedersen (2009)). We focus on shifts in the compensation that dealers charge to absorb idiosyncratic demand imbalances for bonds with similar maturities and with the same issuer. This type of intermediation does not require dealers to bear the issuer's credit risk. That is, to absorb such imbalances, dealers do not necessarily take a net long or net short position in an issuer's bonds, but dealers do take gross positions (gross long and gross short positions).

The ability or willingness of dealers to absorb demand imbalances, even for bonds from the same issuer and with similar maturity, is limited for many complementary reasons: agency frictions and asymmetric information between dealers and outside investors, regulatory costs, dealer risk aversion, liquidity risk related to dealer funding, and limited capital. As a result, demand imbalances give rise to noise in each issuer's yield curve – deviations of yields on individual bonds from a smooth yield curve. This noise reflects temporary deviations in prices from fundamental values. Therefore, we identify a liquidity supply shift as a shock that leads to a decline in noise and an increase in gross dealer corporate bond positions. In contrast, a demand shock leads to a rise in both noise and gross dealer positions. In this context, it is natural to view noise as a price of liquidity and gross dealer positions as a quantity of liquidity.

An advantage of the corporate bond market in studying liquidity drivers is that it is relatively easy to identify the supplier of liquidity. In particular, we focus on primary dealers. Corporate bonds trade in over-the-counter markets and traditionally these dealers have used their balance sheets to absorb temporary order imbalances 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 has hampered the provision of liquidity by these dealers, illustrating why quantifying liquidity supply by these intermediaries is important for policymakers and researchers.

As a first step to identify liquidity supply, we use micro-level data to construct the price and quantity of corporate bond liquidity suggested by theories of segmented markets. 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 in each bond to measure gross inventory. The quantity of liquidity is the sum of gross positions across bonds and dealers.

Next, we estimate a structural vector autoregression (VAR) model using the price and quantity of liquidity to identify supply and demand shifts. Identification is based solely on the signs of contemporaneous responses of noise and gross dealer positions to liquidity supply and demand shocks. This approach takes advantage of the development of sign-restriction methods in the macroeconomics literature that disentangle supply and demand by jointly analyzing price and quantity (e.g. Uhlig (2005) for money supply shocks and Kilian and Murphy (2014) for oil supply shocks). Advantages of this approach are that we can quantify the magnitude of supply and demand shocks, and that we do not rely on ad-hoc proxies for liquidity supply shocks such as leverage and dealer CDS spreads.

The liquidity supply shocks inferred from the model capture notable episodes of stress in the corporate bond market, with liquidity supply declining following the suspension of redemptions at Bear Stearns hedge funds (which marked the start of the financial crisis), the failure of Lehman Brothers, the European fiscal crisis in 2011, and the taper tantrum in 2013. Liquidity demand, however, does not follow a consistent pattern around stress events: it declines following several stress events in the 2007-2008 financial crisis and the European fiscal crisis, but it spikes following the collapse of Lehman Brothers. Liquidity demand generally rises 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 low.

We also examine the cumulative sum of supply and demand shocks separately for investment grade (IG) and high-yield (HY) bonds over different subsamples. We find that throughout the "Dodd-Frank" period from July 2011 to March 2014 and the "Volcker" period after April 2014, liquidity supply for IG bonds fell notably. In contrast, liquidity supply for HY bonds did not fall as significantly. For the HY market, liquidity demand increased during the Dodd-Frank period but fell notably in the Volcker period. Though it is difficult to pin down the event dates for the introduction of recent banking regulations, our results are broadly consistent with the literature on the effects on market liquidity of post-crisis banking regulations (e.g., Bao, O'Hara and Zhou (2017) and Bessembinder et al. (2017)).

Our methodology enables us to quantify the magnitude of supply and demand shocks and to evaluate their contribution to changes in endogenous variables. We find that both supply and demand shocks are important for explaining variation in the price and quantity of liquidity. In particular, on impact, supply and demand shocks each explain about half of the variance of unexpected changes in dealer capital commitment, while supply shocks explain about 60 percent of the variance in unexpected changes in noise. Because neither supply nor demand drives 100 percent of the variation in dealer gross positions or noise, it is crucial to separate the supply effect from the demand effect to understand changes in dealer inventory.

To determine whether the identified shocks indeed capture liquidity supply and demand shifts, we investigate the relationship between the shocks and other commonly cited proxies for liquidity supply and demand. We find that liquidity supply shocks are strongly correlated with dealers' funding cost and the capital ratio of primary dealer parent companies (He, Kelly and Manela (2017)). Furthermore, we find that demand shocks are positively related to absolute values of corporate bond mutual fund flows. However, there is considerable variation in the liquidity supply and demand shocks that is not explained by standard liquidity supply and demand proxies; kitchen-sink regressions using all of the proxies as explanatory variables yield adjusted R-squared of 0.06 for supply shocks and 0.01 for demand shocks. Therefore, even though the supply and demand shocks identified by our framework are indeed associated with different sets of drivers, the proxies proposed in the literature explain only a small fraction of total liquidity supply and demand shifts. These results highlight an advantage of our approach to quantifying the size of liquidity supply and demand shocks; the results also point to the difficulty of finding proxies that can capture a significant fraction of the variation in liquidity supply and demand.

To demonstrate the importance of the supply-demand decomposition, we show that liquidity supply and demand shocks are priced differently in the corporate bond market. Using the cross-section of corporate bond returns, we estimate betas of corporate bond returns with respect to liquidity supply shocks and liquidity demand shocks. We find that bonds with high betas with respect to liquidity supply shocks earn higher average returns than bonds with low betas. In contrast, bonds with high liquidity demand betas earn lower average returns than those with low betas. Importantly, shocks to aggregate dealer capital commitment are not priced in the cross-section of corporate bonds as capital commitment hides two underlying drivers of liquidity with offsetting effects on expected returns. With the decomposition, we uncover that liquidity supply shocks are priced significantly.

Our estimates of the risk premiums for liquidity supply and demand risk are economically significant and larger in magnitude for liquidity supply risk. For example, for a portfolio that is long the quintile of bonds with the highest liquidity supply betas and short the quintile with the lowest liquidity supply betas, the average excess return is 51 basis points per month. An alternative portfolio long bonds with high liquidity demand betas and short those with low liquidity demand betas earns an average excess return of -22 basis points. The risk premium associated with exposure to liquidity supply risk is significant after controlling for betas with respect to the Amihud or Pastor-Stambaugh measures and for bond characteristics such as the Roll measure, credit ratings, maturity, and issue size, while the risk premium associated with liquidity demand does not always remain statistically significant when including these controls.

Our results support the dealer inventory channel through which liquidity supply risk is priced in securities, as suggested in Kondor and Vayanos (2018). The value of a bond with a high liquidity supply beta tends to fall when dealers are less able or less willing to supply liquidity, and investors demand higher average returns to hold such bonds. Importantly, the noise measure is a deviation of yields from the issuer-specific yield curve and thus our liquidity supply and demand shocks are in principle not affected by asymmetric information regarding the issuer, an alternative channel connecting liquidity and asset prices.

Finally, we examine whether dealer gross positions explain time variation in expected returns for the aggregate corporate bond market. We find that dealer gross positions predict returns on the corporate bond market portfolio negatively, but only when supply shocks have recently dominated demand shocks. When dealer capital commitment falls in response to lower liquidity supply, bond returns in the future are on average higher. However, this relationship disappears or possibly reverses when the fall in dealer positions reflects lower liquidity demand from investors: when demand shocks are greater in magnitude than supply shocks, lower dealer capital commitment predicts lower market returns but this association is not statistically significant. The evidence from the time-series analysis further strengthens our finding that the liquidity supply and demand shocks carry distinct information with respect to risk premiums.

The liquidity measure we use in this paper builds on to the literature on noise in Treasury yield curve (Fontaine and Garcia (2012), Hu, Pan and Wang (2013) and Malkhozov et al. (2016)). 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 Treasury noise 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, Feldhutter and Lando (2012), Edwards, Harris and Piwowar (2007), Feldhutter (2012), Friewald, Jankowitsch and Subrahmanyam (2012), Longstaff, Mithal and Neis (2005),Chernenko and Sunderam (2018)). 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), Goldstein and Hotchkiss (2017) and Pan and Zeng (2018), 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 corporate bond illiquidity using a 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. In the same spirit, Chen, Joslin and Ni (2017) identify the supply of crash insurance in options markets using an approach based on the correlation of recent changes in price and quantity, while Baranova, Chen and Vause (2015) study liquidity using a structural VAR focusing on credit spreads. In contrast, this paper studies the price of liquidity in the corporate bond market motivated by theories of segmented markets.

Lastly, Friewald and Nagler (2016) study variation in dealer inventories across corporate bonds and the connection to expected returns; comparing across bonds in a given week, Friewald and Nagler (2016) show that bonds in which dealers have higher inventory earn higher expected returns. Our study is complementary, as we focus on shocks to liquidity supply and demand at the aggregate level.

The remainder of the paper is organized as follows. Section 2 describes the data and the empirical approach. Section 3 presents the results of the VAR model that identifies liquidity supply and demand shocks. Section 4 examines how the liquidity supply and demand shocks are priced in the corporate bond market. Section 5 concludes.

2 Data and Empirical Approach

In this section, we introduce a simple model of segmented markets in order to motivate a choice of price and quantity measures for corporate bond liquidity. Next, we describe the data used 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 supply and demand shocks.

2.1 Model of Financial Intermediation in a Segmented Market

Building onto Gromb and Vayanos (2002) and Gromb and Vayanos (2017), among others, 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 access only one market and face 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 the risk of a liquidity shock that would force her to unwind her positions before realizing the profits from her long-short portfolio. Therefore, the liquidity provided by the intermediary is generally not sufficient to eliminate the difference in price between the two securities.

In making markets, the dealer focuses on managing liquidity risk. If the securities are in zero net supply, the dealer has no net exposure to the (identical) cash flows of the two securities; that is, the dealer has no net position and does not make a bet on the 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. With securities in zero net supply, the magnitude of the long and short positions are the same. However, as shown in Appendix A, when the securities are not in zero net supply and the dealer chooses to have some net exposure to the underlying payoff of the two securities, dealer gross positions continue to be a measure of market-making activity that are increasing in liquidity supply and liquidity demand. Moreover, the gap in price between the two securities can be thought of as a liquidity price, as this is a deviation in price from the fundamental value due to the need to transact the particular security. Without market segmentation, the liquidity price is zero.

We summarize two key takeaways from the model. First, we show that if investors' desire to trade in their particular securities market increases, then the financial intermediary's gross position increases while the gap between the two securities widens. This shock increases both liquidity price (the gap in security prices) and quantity (dealer 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. In the model, we connect unobservable drivers of liquidity supply and demand 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

Based on the simple model described in Section 2.1 and theories of segmented markets such as Gromb and Vayanos (2017), we use financial intermediaries' gross positions in corporate bonds as a measure of liquidity quantity. In this paper, we focus on primary dealers, the designated trading counterparties of the Federal Reserve Bank of New York listed in Appendix C. We limit our sample to primary dealers since we are interested in market participants who have traditionally been at the center of market-making in the corporate bond market; primary dealers as a group maintains the volume share of around 70% in our sample throughout the sample period. He, Kelly and Manela (2017) also focus on primary dealers, but emphasize their role as marginal investors in bearing aggregate risk.

In order to construct dealers' gross positions in corporate bonds, we use regulatory TRACE data from 2002 to 2016 obtained from FINRA. Regulatory TRACE data is similar to the publicly available TRACE except that the volume for large trades is not truncated and the dealer's name is revealed. 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 an estimate for each dealer of the weekly capital commitment for each bond. 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 trades reversed within a four-week rolling window because a trade that takes longer than four weeks to reverse is likely motivated by a purpose other than market-making (e.g. proprietary trading). Table 1 illustrates the construction of our inventory measure using artificial data for a hypothetical bond over a 5-week period.

Since TRACE was fully implemented in the first quarter of 2005, we start our analysis in April 2005. Some primary dealers have multiple dealer subsidiaries, and thus we aggregate to the primary dealer holding-company level based on the name of each dealer subsidiary. In theory, this cumulative flow should provide a reasonable estimate for the stock of inventory each week. To mitigate the issue of possible mismeasurement, we remove weekly flows that are greater than a third of the amount outstanding of the bond and also remove all flows for bonds with age less than one month. We remove young bonds since a part of sales by dealers likely represents sales of the leftover of bonds underwritten in the primary market. We also make sure that there is no inventory of bonds that have matured.

Furthermore, in order to maintain consistency in pricing among bonds, we merge TRACE data with data on bond characteristics from Mergent FISD; we 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 the TRACE data, which we use 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 of $Q_{d,j,k,t}$:

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

Equation (1) allows the possibility that the inventory $Q_{d,j,k,t}$ for a given bond is negative. Asquith et al. (2013) show that the cost of borrowing corporate bonds to short is comparable to the cost for stocks. Data from the Federal Reserve's Weekly Report of Dealer Positions (known as the FR 2004) also shows that primary dealers have substantial amount of short positions on corporate bonds.[1] Thus, we do not remove observations if $Q_{d,j,k,t}$ is negative. Finally, we scale the aggregate quantity measure Q_t by the consumer price index excluding food and energy to express it in 2005 dollars and use its logarithm $q_t = \log(Q_t/CPI_t)$ in the analysis below.

The top panel of Figure 1 shows the time-series of aggregate dealer gross positions, Q_t, with and without seasonal adjustment. We find that gross positions exhibit strong seasonality, falling at the end of calendar quarters. Thus, we seasonally adjust the series using the ratio to moving average method and use the seasonally-adjusted series in the analysis below.[2] At the beginning of the sample, gross positions remained near \$25 billion for a few years, before starting to decline in late 2007 as the financial crisis began. Following the collapse of Lehman Brothers, gross positions fell to about \$15 billion. Gross positions recovered from 2009 to 2011, and have declined gradually since then. The maturity of gross positions is stable over time, with average of 8.8 years and a standard deviation of 0.28 years.

[1] The correlation between our measure of gross dealer positions, Q_t, and gross dealer 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 imperfect partly because the FR-2004 data includes commercial paper and bonds with embedded options, and inventories are reported on the fair-value basis rather than at book value. Unfortunately, prior to April 2013, the corporate bond reporting category "corporate bonds" did not separate out corporate bonds from other large fixed-income classes with credit risk such as non-agency mortgage-backed securities.

[2] Let $Q_{t,MA}$be the moving average over the past 52 weeks, $Q_{t,MA} = \frac{1}{52}\sum_{m=1}^{52} Q_{t-m}$. We compute the ratio to moving average, $RMA_t = Q_t/Q_{t,MA}$, and compute the arithmetic average of the ratio for each week across years. Then the seasonally-adjusted quantity for m-th week in a year is given by $Q_m^{s.a.} = Q_m/R\bar{M}A_m$.

The ratio of IG bond positions to all positions is on average 58% with a standard deviation of 4%. Moreover, positions in IG bonds and HY bonds are highly correlated (correlation coefficient of 0.76). Thus, the risk profiles of the gross dealer 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 with similar cash flow risks can be regarded as a price of liquidity. In this paper, we assume that preferred habitat investors have idiosyncratic demand for corporate bonds with particular maturity, leading to such discrepancies when interacted with dealer inventory frictions as in Culbertson (1957), Modigliani and Sutch (1966), Malkhozov et al. (2016) and Gromb and Vayanos (2017), among others. Section 1 of Internet Appendix provides empirical evidence for this type of market segmentation based on data on investors' holdings. This assumption leads us to follow Fontaine and Garcia (2012) and Hu, Pan and Wang (2013) and use the deviations of the yield-to-maturity of individual corporate bonds from a fitted yield curve as the basis of our liquidity price measure. In particular, for each week, we estimate a smooth curve for each issuer and measure pricing errors against the curve. The use of the deviation from the issuer-specific yield curve is motivated by our view that one key role of dealers is providing liquidity across bonds 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. For our purposes, the Merrill Lynch data has an advantage over TRACE: it provides end-of-day quotes for all bonds on the daily basis under the same conditions, regardless of whether a transaction occurred for the bond and the characteristics of any trades (such as volume) in each bond. 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 requires bonds to have amount outstanding greater than $100 million 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 during our sample period, the strike price of the option is typically set to vary in a way that options are never in-the-money; Elsaify and Roussanov (2016) argues that the make-whole call options are designed in this way to facilitate cash flow management, rather than to allocate interest rate or default risk between the bondholder and the issuer. 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 their role in cash flow management, 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 at least seven 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 Appendix B and argue that the difference in the two samples does not affect our empirical analysis.

The Nelson-Siegel-Svensson model of Svensson (1994) assumes that the n-period instantaneous forward rate is:

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

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.

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

where ω_k is modified duration for bond k. To estimate the parameters reliably, we set β_3 to zero when there are less than 15 bonds, effectively estimating the Nelson-Siegel curve for those issuers.

The pricing error of a bond is the difference in yield-to-maturity calculated from observed prices and the model-implied 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 each week. The noise spikes during the financial crisis and comes down afterwards. It also goes up somewhat in 2012 reflecting the turmoil in European financial markets and the Greek debt crisis, as well as in late 2015 amid concerns related to China and a selloff in oil markets.[3] The figure also compares p_t with other measures

[3] To address the potential concern about the bias in fitting the Nelson-Siegel-Svensson curve, we show

of illiquidity such as the imputed round-trip cost (IRC) measure of Feldhutter (2012), the Amihud (2002) illiquidity measure, and the negative of CDS-bond basis obtained from JP Morgan. The IRC and Amihud measures are computed each week using 3-month rolling windows and we report the median across all bonds in TRACE data.

The noise measure has two advantages over other alternative measures of transaction costs. First, our measures do not depend on transaction prices and thus are not affected by changing nature of transactions that are not observable to researchers. In Figure 1, IRC and the Amihud measures have downward trend since the crisis while noise does not have such trend. The downward trend in transaction price-based illiquidity measures partly reflects the fact that the average trade size becomes smaller in the post crisis period (Bessembinder et al. (2017)), not necessarily improving market liquidity.

Second, typical measures of illiquidity such as the Amihud price impact proxy are affected both by inventory frictions and information asymmetry between issuers and investors. High price impact does not necessarily imply a high degree of inventory frictions; it could instead reflect a high degree of adverse selection. Since we fit the curve at the issuer level, the noise measure in principle reflects only the inventory friction component of illiquidity, not information asymmetry regarding issuers. We take advantage of the difference between noise and typical illiquidity measures, such as the Amihud and Pastor-Stambaugh measures, and study the source of liquidity risk premium later in Section 4.

An alternative measure of mispricing is difference in credit spreads between corporate bonds and credit default swap (CDS) spreads. As Bai and Collin-Dufresne (2013) show, however, CDS and corporate bond liquidity are driven by different factors.[4] The focus of

examples of fitted curves in the crisis and normal periods, and examine the behavior of fitting errors in Section 1 of Internet Appendix. In particular, we show that the fitted yield curve is flexible enough to capture the de facto seniority of short-term bonds shown by Bao and Hou (2017), such that there is no bias in fitting errors across maturity.

[4]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 liquidity conditions. 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

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.

2.4 Other Data

In studying the cross-section of corporate bonds, we use TRACE data for month-end bond prices and FISD data for bond characteristics. Following Bessembinder et al. (2009), we use the volume-weighted average price for institutional transactions with volume more than \$100,000. We treat the bond price as a month-end observation if the trade occurs in the last five business day of the month. If there are more than one daily price observation in the last five business days, we use the last observation as the month-end price. As before, we focus on US corporate bonds with no embedded options other than make-whole call options.

Finally, we use Merrill Lynch total bond return index for the returns on the bond market portfolio.

2.5 Summary Statistics

Table 2 shows the summary statistics for the liquidity price and quantity variables, p_t and q_t. The unconditional correlation between p_t and q_t is -0.57, suggesting the importance of supply shocks. The table also shows the correlation between q_t and median Amihud illiquidity, imputed round-trip costs and the CDS-bond basis, which are estimated at -0.59, -0.51 and 0.54, 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 is 0.57, the correlation between p_t and imputed round-trip costs is 0.61, and the correlation between p_t and the CDS-bond basis is -0.86. Therefore, our measure of the price of liquidity is correlated with traditional measures of

spreads between CDS and corporate bonds, or the CDS-bond basis, is partly driven by the factors specific to the CDS market.

illiquidity, especially with the CDS-bond basis.

2.6 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:

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

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 $v_t = A^{-1}\xi_t$.[5]

We estimate the model using the pure sign restrictions approach of Uhlig (2005). In

[5]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.

estimating the reduced-form VAR, we use a weak Normal-Wishart prior with $L = 6$ lags of the endogenous variables which is chosen based on the Akaike Information Criterion.

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 from the posterior over B and Σ, 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 the 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 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 it. We repeat the draws for A_m 100 times for each draw of reduced form parameters, $B_1, \ldots, B_L, \Sigma$, which are also drawn 100 times from the posterior distribution. As a result, we use 10,000 draws of parameters to 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 of the structural shocks, 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. Supply declines sharply after the failure of Lehman Brothers in 2008. Liquidity supply recovers subsequently, but decreases sharply during the European fiscal crisis in late 2011 and falls somewhat 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 sum of the liquidity demand shocks, $\sum_{j=1}^{t} \upsilon_j^d$. Immediately after the collapse of Lehman Brothers, liquidity demand spiked up, contributing to the rise in noise during the financial crisis. Liquidity demand stabilized somewhat after the crisis, but generally maintained an upward trend despite some negative shocks during the European fiscal crisis in 2010 and 2011 and the taper tantrum in 2013.

Figure 2 reports changes in liquidity supply and demand over the five subperiods in our sample. The definition of subperiods generally follows Bao, O'Hara and Zhou (2017): i) the pre-crisis period up to June 2007, ii) the crisis period from July 2007 to April 2009, iii) the post-crisis period from May 2009 to June 2010, iv) Dodd-Frank regulation periods from July 2011 to March 2014, and v) the Volcker rule period after April 2014. Despite the names of these subperiods, it is difficult to attribute shocks to changes in banking regulation after the financial crisis, as we do not know the exact event dates on which regulatory changes were important movers of liquidity supply. However, it is still informative to see how liquidity supply and demand change on net in each subsample.

Over the crisis period, both liquidity supply and demand fell on net, before partly recovering during the post-crisis period. During the Dodd-Frank period, liquidity supply decreased notably while liquidity demand increased. The fall in liquidity supply occurred mostly during the period in which the debt crisis in Greece deepens, but it is also possible that prospects for tighter regulations on dealer market activity exacerbated the decrease.

During the Volcker period, both demand shocks and supply shocks were negative, on net. The decrease in liquidity supply occurred around the time when a bond mutual fund (Third Avenue) was forced to halt redemption. The negative demand shocks more recently are consistent with the findings of Anand, Jotikasthira and Venkataraman (2017) and Choi and Huh (2017), who document that non-dealers started to provide liquidity in the corporate bond market in the recent period. As we view the liquidity supply and demand from primary dealers' perspective, an increased supply from non-dealers is interpreted as a decline in liquidity demand.

Overall, our decomposition results show that liquidity supply falls during the period in which a series of banking regulations is introduced, and the effect is more pronounced at the time of market stress, consistent with Bao, O'Hara and Zhou (2017) and Bessembinder et al. (2017). Moreover, the advantage of our approach is that we can quantify the size of shocks precisely and evaluate their importance as a driver for market liquidity. From (3), we can express reduced form shocks as,

$$
\begin{aligned}
\xi_t^p &= a_{11} v_t^s + a_{12} v_t^d \\
\xi_t^q &= a_{21} v_t^s + a_{22} v_t^d
\end{aligned}
$$

where a is an element of matrix A. Using the posterior mean estimate for A, we compute the variance ratio, $a^2/\sigma^2(\xi_t)$, for each pair of reduced-form and structural shocks. By construction, we can fully assign the price shocks and quantity shocks to the two structural shocks, since we have $a_{11}^2 + a_{12}^2 = \sigma^2(\xi_t^p)$ and $a_{21}^2 + a_{22}^2 = \sigma^2(\xi_t^q)$.

For price shocks, we find that $a_{11}^2/\sigma^2(\xi_t^p) = 0.57$ and $a_{12}^2/\sigma^2(\xi_t^p) = 0.43$. For quantity shocks, we estimate the ratio at $a_{21}^2/\sigma^2(\xi_t^q) = 0.51$ and $a_{22}^2/\sigma^2(\xi_t^q) = 0.49$. These estimates imply that, to understand the dynamics of liquidity measures and dealer balance sheets, we have to take both supply and demand shocks into account. While it is natural to think that dealer balance sheets are driven not only by dealer risk aversion but also by investor demand, it poses a challenge to the literature which aims to measure the degree to which financial intermediaries are constrained by examining their balance sheets or leverage. To the extent supply and demand shocks contain different information about the state of economy, using dealer balance sheets as a proxy for financial constraints can lead to incorrect inferences about the relationship between intermediary risk aversion and asset prices. We elaborate further on this issue in Section 4.

3.2 Drivers of Supply and Demand Shocks

Though an advantage of the identification strategy based on sign restrictions is that we do not need ad-hoc proxies for supply and demand, it is natural to ask how the shocks identified by sign restrictions relate to commonly used proxies for liquidity shocks. To examine the link between sign-identified structural shocks and other proxies, we run a regression,

$$v_t = b_1\epsilon_t^{VIX} + b_2|MFFLOW_t| + b_3\Delta ISSUE_t + b_4 HYSHARE_t + b_5\epsilon_t^{LEV} + b_6\epsilon_t^{TED} + b_7 R_{t-1} + u_t \tag{4}$$

where ϵ_t^{VIX} is a shock to VIX, $|MFFLOW_t|$ is the absolute value of a weekly mutual fund flow to US domestic investment-grade corporate bond funds, $\Delta ISSUE_t$ is the log growth rate of the total face value of corporate bonds issued in week t, $HYSHARE_t$ is the share of high-yield bond issues among total corporate bond issues in week t, ϵ_t^{LEV} is a shock to intermediary capital ratio of He, Kelly and Manela (2017), ϵ_t^{TED} is a shock to TED spread, and R_{t-1} is the lagged return on the Merrill Lynch corporate bond index. All variables are standardized to have mean zero and standard deviation equal to one. Shocks to the VIX, the intermediary capital ratio and the TED spread are extracted as a residual of univariate autoregressions on each variable.

We use mutual fund flows, $|MFFLOW_t|$, as one proxy, motivated by Coval and Stafford (2007)'s evidence that equity mutual funds are forced to buy and sell stocks when there are large fund flows. Their evidence suggests that large fund flows to corporate bond mutual funds, whether positive or negative, may force them to trade more and thus lead to an increase in liquidity demand.

To measure the conditions in the primary market, we use the growth rate in new bond issues, $\Delta ISSUE_t$, in week t. Furthermore, we use the share of high-yield bond issues to total corporate bond issues, $HYSHARE_t$, to measure market sentiment. Greenwood and Hanson (2013) show that the share of high-yield debt issuance measures market sentiment and predicts returns on corporate bonds, which may also affect investors' desire to trade and dealers' propensity to intermediate.

To measure dealer financial constraints, He, Kelly and Manela (2017) use the capital ratio of publicly traded dealer parent companies, where the capital ratio is defined as aggregate market equity as a share of the sum of aggregate market equity and the aggregate book value of debt. Thus, this ratio is an inverse of leverage and a higher ratio arguably implies that dealers are less financially constrained. TED spreads measure the funding cost of banks and could affect dealers' willingness to expand inventory holdings.

The top panel of Table 3 reports estimated slope coefficients and regression R-squared from estimating (4). The first seven rows show the results for univariate regressions, while the last row reports the multivariate regression including all explanatory variables at the same time. We find that supply increases when shocks to VIX are negative, new bond issues increase, the share of HY bonds are lower, dealers' capital increases, and the TED spread declines. Economically, dealer capital has the largest significance: a one standard deviation increase in dealer capital corresponds to a 0.16 standard deviation increase in liquidity supply. When dealers are less capital constrained (a positive shock to the capital ratio), they are more willing to supply liquidity in the corporate bond market. On the other hand, when dealers' funding cost is high (a positive shock to TED spread), dealers supply less market liquidity.

For demand shocks, fund flows to corporate bond mutual funds are the only significant driver. When there is a large flow to bond mutual funds, those funds are more likely to trade, which leads to higher liquidity demand. In untabulated results, we also test if the positive flows and negative mutual fund flows affect liquidity demand differently but we do not find an evidence for asymmetry in the estimated slope coefficients.

Comparing supply and demand shocks, we make two observations. First, liquidity demand is not significantly affected by the key drivers of supply shocks, such as shocks to the capital ratio or TED spreads. In our identification strategy, we assume that supply and demand shocks are independent; different loading of shocks to proxies is consistent with this assumption. Second, liquidity supply shocks load on many proxies that are likely to capture shocks to the willingness of the marginal investor to bear risk, while it is less obvious that demand shocks are important state variables of concern for asset pricing. We test the difference in risk premiums for each shock in Section 4.

Later on, we use structural shocks at the monthly frequency to run asset pricing tests in order to be consistent with the practice in literature. Table 4 presents the estimates for regression (4) at the monthly frequency. The results are qualitatively consistent with the weekly frequency, except that now VIX affects demand shocks positively while lagged market returns are positively associated with supply shocks. Thus, at the monthly frequency, the liquidity demand shocks become more systematic than at the weekly frequency.

Overall, the link between the sign-identified structural shocks and observable proxies provides comfort to our methodology of disentangling the roles of liquidity supply and demand in driving dealer gross positions and noise. However, R-squared for the regressions is generally low: the multivariate regression with all proxies yields adjusted R-squared of 0.06 when the supply shock is the dependent variable and only 0.01 when the demand shock is the dependent variable. Low R-squared suggests that, although these proxies do drive a part of liquidity supply or demand, the magnitude of liquidity supply and demand shocks can be only measured through a joint analysis of price and quantities. For example, mutual fund flows may be a valid proxy for liquidity demand but such a shock seems small relative to total demand shocks, perhaps because bond mutual funds hold only a limited fraction of the corporate bond market in the US. [6]

3.3 Discussion and Extension of the Model

3.3.1 What If Investors Wait Longer These Days?

A possible critique of the time-series analysis of liquidity measures during the pre- and post-crisis periods is that dealers' business model and the nature of transactions appearing in the data change dramatically between these two periods. In particular, Goldstein and Hotchkiss (2017) argue that dealers' business model after the crisis changes, with dealers shifting from using their balance sheets to provide liquidity to playing the role of brokers (i.e., prearranging transactions among investors without committing risk capital). Indeed, Bessembinder et al. (2017) report that average trade size becomes smaller in the post-crisis period than the pre-crisis period, suggesting that investors are waiting longer for trades to be executed or splitting large trades into smaller pieces.

These problems arise because liquidity measures used in the literature are based on realized transaction data, and 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, Feldhutter and Lando (2012) and Dick-Nielsen and Rossi (2017)). Our approach addresses this problem by calculating noise using quote prices. Quotes are provided by the same dealer (Merrill Lynch) for a trade with a standard size with no considerations for the characteristics of a particular order. In addition, our database consists of bonds chosen based on objective criteria such as face value and maturity, and those criteria are unchanged through our sample. In contrast, in transactions data, we would observe more transactions

[6] According to the Federal Reserve's flow of funds, mutual funds hold 15% of corporate and foreign bonds issued in the U.S. at the end of 2016

for liquid bonds, and such bias becomes arguably more pronounced over time as dealers shift their focus to more liquid bonds.

Another concern is the impact of technological development in trade execution. Over the sample period, some corporate bonds became available to trade on electronic platforms for large investors in addition to traditional venues based on telephone. With the aid of new technology, dealers may be better able to match customer trades, and thus they can provide liquidity without committing capital. In our framework, we would interpret such changes as negative demand shocks because better technology is likely to reduce capital commitment and noise. We argue that technological development does not invalidate our identification of supply and demand shocks - demand shocks by definition contain everything that moves price and quantity in the same direction. However, the concern about technological development suggests that there are multiple drivers of liquidity demand shocks in addition to the proxies used in the previous exercise.

3.3.2 Investment-Grade and High-Yield Bonds

Our main results above are based on aggregate measures of dealer capital commitment and noise that arises due to idiosyncratic demand imbalances for specific bonds. However, Ellul, Jotikasthira and Lundblad (2011) show that investors for IG bonds are quite different from those for HY bonds. In this article, we do not use the segmentation across credit ratings as a key tool for identification because it is difficult to control for the difference in IG and HY issuer fundamentals. However, it is still interesting to run our structural VAR analysis separately for IG bonds and HY bonds and see if the demand and supply shocks are different in these two segments of the corporate bond market.

To this end, we estimate the VAR in (2) separately for IG bonds and HY bonds, and extract rating-specific supply and demand shocks using sign restrictions. Figure 3 shows the cumulative sum of supply and demand shocks in each market. Though both markets react in the same direction as the aggregate market to notable events of market stress, the magnitude of shocks are quite different. For IG bonds, liquidity supply falls notably during the Volcker period, especially toward the end of 2015 when Third Avenue mutual funds faced liquidity problems. In contrast, the liquidity supply for HY bonds are on net unchanged during the Volcker period.

Figure 3 also shows that there is a notable difference in liquidity demand between IG and HY bonds in the post-crisis period. In particular, HY bonds experience an increase in liquidity demand during the Dodd-Frank period which retraces during the Volcker period.

In addition, our methodology provides an ideal setup to study related research questions, such as whether there is liquidity contagion across markets. 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 the interaction of liquidity in the IG and HY bond markets. Gromb and Vayanos (2002) predict that, if a dealer loses money in one market, it may reduce its liquidity provision in another market due to the collateral constraint. We test this hypothesis by estimating trivariate VARs; the first such VAR includes the IG liquidity price, IG liquidity quantity and HY liquidity price; the second VAR includes the HY liquidity price, HY liquidity quantity and IG liquidity price. Detailed results are provided in the Internet Appendix. In summary, we find that an unexpected rise in noise in the HY market leads to an modest increase in noise in the 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.

3.4 Analysis of Identification using Monte Carlo Simulation

In this section, we verify if our strategy can recover the underlying structural shocks using simulated data from a Monte Carlo simulation. Specifically, we generate 1,000 draws of the history of structural shocks, $\{v_t\}$, using an independent standard normal distribution. We then compute the corresponding endogenous variables Y_t using the the mean estimate of the structural VAR parameters from the benchmark model in Section 2.6. For m-th draw of simulated structural shocks, we use the resulting endogenous variables Y_t to estimate the structural VAR with sign restrictions and generate the posterior distribution of structural shocks. We compute the mean of structural shocks in the posterior distribution, $\hat{v}_t$, and compute "R-squared" for m-th draw of the history of structural shocks as

$$R_m^2 = 1 - \frac{\Sigma_t(\hat{v}_{m,t} - v_{m,t})^2}{\Sigma_t v_{m,t}^2}.$$

If the posterior mean estimates for structural shocks provide an accurate estimate for the true structural shocks, the value of R-squared should be close to one.

Table 5 shows the R-squared from 1,000 sets of simulated data using Monte Carlo simulation. The mean and median R-squared is 0.97 for both supply and demand shocks, implying that the structural shocks identified by imposing sign restrictions on the rotation matrix indeed closely track the true structural shocks. Overall, the simulation results provide additional support for our identification strategy for supply and demand shocks in our sample.

4 Cross-Section and Time-Series of Corporate Bond Returns

4.1 Cross-Section of Corporate Bond Returns: Portfolio Sort

In this section, we study whether the supply and demand shocks are priced differently from each other using the cross-section of corporate bonds as a testing tool. Section 3.2 showed that the liquidity supply shocks are positively correlated with changes in common proxies for systematic risk. The demand shocks, however, are not as strongly associated with these proxies and load instead on a different factor, corporate bond mutual fund flows. Thus, it is natural to conjecture that the price of risk for these two shocks are different from each other. The analysis on the price of risk does not validate or invalidate the supply-demand decomposition. However, it provides another example of the importance of understanding the different drivers of dealer balance sheets.

To estimate the liquidity risk exposure of bond k, we run a factor-model regression of the form,

$$R^e_{k,t} = b_0 + \beta_{k,MKT} R^e_{MKT,t} + \beta_{k,s} v^s_t + \beta_{k,d} v^d_t + \epsilon_{k,t} \tag{5}$$

where $R^e_{k,t}$ is the monthly excess return on bond k and $R^e_{MKT,t}$ is the monthly excess return on the value-weighted corporate bond market portfolio. In this exercise, we aggregate structural shocks within a month to construct monthly series of shocks. We estimate the regression in (5) using three-year rolling windows (with the minimum number of observations of 24 months) to allow the slope coefficients to vary over time.

Once we estimate time-varying betas, we sort corporate bonds into quintiles every month based on their betas and compute the value-weighted average return on each portfolio. The top panels of Table 6 report the average returns on the portfolios in excess of risk-free rate. Corporate bonds with the lowest liquidity supply beta earn a monthly average excess return of 0.49% while the bonds with the highest betas earn 1.01%; the difference is statistically significant. The results suggest that bonds that have higher returns when liquidity supply falls act as a hedge against the liquidity supply contraction and hence earn a lower risk premium. In contrast, bonds with the lowest demand betas have a monthly excess return of 0.85%, which is 0.22% higher than bonds with the highest demand betas. These results suggest that aggregate liquidity demand tends to rise when investors' marginal utility of consumption is high.

The bottom two panels of Table 6 present characteristics of bonds averaged within portfolios. The face values and ages of bonds are similar across portfolios. The Roll measure of illiquidity and imputed round-trip costs have a U-shaped pattern, as bonds in the lowest and highest quintiles are more illiquid than the bonds in the middle quintiles. However, there is a notable difference in the distribution of credit ratings; bonds with low supply betas and high demand betas are predominantly investment-grade bonds while bonds with high supply betas and low demand have a greater tendency to be high-yield bonds.

Since variation in credit ratings leads to differences in average returns (Chordia et al. (2017)), we control for differential risk exposures, including exposure to credit risk, by running return-based factor regressions,

$$R^e_{p,t} = \alpha_p + \Sigma^J_{j=1}\beta_{p,j} f_{j,t} + \epsilon_t, \tag{6}$$

where $R^e_{p,t}$ is the return on portfolio p, and $f_{j,t}$, $j = 1, \ldots, J$ are return-based factors proposed in the literature. Specifically, we use (i) the five-factor model of Fama and French (2015) supplemented by the TERM (difference in returns between long-term Treasuries and T-bills) and DEF factors (difference in corporate bond returns and long-term Treasuries) from Fama and French (1993), (ii) the four-factor model of corporate bond returns from Bai, Bali and Wen (2018), and (iii) the two-factor model of intermediary asset pricing from He, Kelly and Manela (2017). The model of Bai, Bali and Wen (2018) uses four factors including credit risk (returns on bonds sorted on credit rating), downside risk (sorted on the second worst return in the past three years), liquidity (sorted on the Roll measure) and monthly reversals (sorted on returns in the previous month). As this four-factor model includes credit risk and liquidity factors, the model is particularly suitable in this context to adjust for risk exposures. Finally, He, Kelly and Manela (2017) use the market risk factor augmented by shocks to bank holding companies' capital as an additional factor and report that the model prices the cross-section of corporate bonds. Since our shocks are partly extracted from shocks to dealers' capital commitment, it is interesting to examine if their factor model can explain the liquidity supply and demand risk premiums.

The second to fourth panels of Table 6 show the intercept from the regression (6). For the liquidity supply betas, the alpha for the hedge portfolio is 0.51% even when taking into account the 5+2 factors of Fama and French (1993) and Fama and French (2015), 0.54% when taking into account the four factors of Bai, Bali and Wen (2018), and 0.34% when accounting for the two factors of He, Kelly and Manela (2017). For the liquidity demand betas, the alpha for the hedge portfolio is -0.32% against the 5+2 factor models of Fama and French (1993) and Fama and French (2015), -0.22% against the four factor model of Bai, Bali and Wen (2018), and -0.13% against the two factor model of He, Kelly and Manela (2017). These alpha estimates for liquidity supply beta-sorted portfolios are all statistically significantly different from zero, suggesting that the liquidity supply shocks identified in our VAR-based decomposition carry distinct information about the relationship between liquidity and investors' pricing kernel. In contrast, the hedge portfolio sorted on liquidity demand betas no longer yields significant alphas once we control for the risk exposures of the Bai, Bali and Wen (2018) or He, Kelly and Manela (2017) models.

The larger magnitude of the risk premium associated with liquidity supply risk is consistent with models such as Kondor and Vayanos (2018) in which liquidity suppliers' role in overcoming market segmentation leads to liquidity supply being a priced risk factor. The less significant risk premium on the liquidity demand risk makes sense in our setup in which investor's desire to trade a bond of specific maturity gives rise to idiosyncratic demand imbalances. In such a model, we expect that demand shocks are less systematic than supply shocks.

4.2 Cross-Section of Corporate Bond Returns: Fama-MacBeth Regressions

Next, we show that the sensitivity with respect to supply shocks is an important determinant of cross-section of bond returns among various control variables by running monthly cross-sectional regressions of Fama and MacBeth (1973),

$$R_{k,t} = \gamma_{0t} + \gamma_{1t}\beta_{k,t} + \gamma_{3t}X_{k,t} + \epsilon_{k,t},$$

where X_{it} is a vector of control variables, including the logarithm of the Roll measure, maturity, face value of bonds, and dummy variables for credit ratings (Aa+, A, Baa and high-yield). In running the cross-sectional regressions, all variables but rating dummies are standardized to ease the comparison of economic significance.

To start with, Panel A of Table 7 reports the average values of γ_t for betas with respect to shocks to the aggregate capital commitment, ξ_t^q. Since the betas on the right-hand side variables are generated regressors, we correct standard errors following Shanken (1992). The estimated price of risk is weakly positive and generally statistically insignificant.

Panel B of Table 7 instead uses the liquidity supply betas, $\beta_{k,s,t}$, to estimate the price of risk. When we run a univariate regression on the betas with respect to liquidity supply shocks, the average slope coefficient is estimated at 0.35, implying that a one standard deviation increase in beta leads to an increase in average returns of 0.35% per month. When we add bond characteristics as control variables, the estimated slope decreases to 0.23 with a t-statistic of 2.89. Thus, unlike betas with respect to capital commitment, the liquidity supply shock is priced in the cross-section of corporate bonds after accounting for the difference in credit rating, maturity and the level of liquidity.

In Panel B of Table 7, the estimated slope coefficients for the liquidity demand betas are negative and marginally significantly different from zero, consistent with the results in Table 6. The point estimate varies between -0.17 to -0.13 depending on the regression specification; the point estimate is generally smaller in magnitude compared with the coefficient estimate for the liquidity supply betas, but the sign is opposite to the sign of the liquidity supply betas.

The results of the cross-sectional regression highlights the significance of the supply-demand decomposition. Without decomposition, dealer's capital commitment is driven both by liquidity demand and supply shocks that carry risk premiums with opposite signs, leading to statistically insignificant results. The regression results suggest that it is essential to disentangle the drivers of observed liquidity proxies to evaluate the magnitude of liquidity risk premiums.

Previous work shows that a security that comoves positively with marketwide liquidity shocks should earn higher returns than those with negative covariance (Acharya and Pedersen (2005)). Empirically, Pástor and Stambaugh (2003) show that liquidity risk is priced in cross-section of stocks while Lin, Wang and Wu (2011) show that there are positive liquidity risk premiums in the corporate bond market. Thus, we study the asset pricing implications of liquidity supply and demand risk when also taking into account risks as measured by exposure to the Amihud (2002) and Pástor and Stambaugh (2003) measures studied in Lin, Wang and Wu (2011). Though these illiquidity measures are closely related to our liquidity supply and demand shocks, there may be an important difference because the noise measure is not affected by changing characteristics of transactions and our shocks combine the information from noise and dealers' capital commitment.

To show liquidity supply and demand betas carry distinct information from the liquidity risk premiums in Lin, Wang and Wu (2011), we run monthly cross-sectional regressions of Fama and MacBeth (1973) using issue-level monthly returns,

$$R_{k,t} = \gamma_{0t} + \gamma_{1t}\beta_{k,s,t} + \gamma_{2t}\beta_{k,t}^{LIQ} + \gamma_{3t}X_{k,t} + \epsilon_{k,t},$$
$$R_{k,t} = \gamma_{0t} + \gamma_{1t}\beta_{k,d,t} + \gamma_{2t}\beta_{k,t}^{LIQ} + \gamma_{3t}X_{k,t} + \epsilon_{k,t},$$

where $\beta_{k,t}^{LIQ}$ is either betas of bond returns with respect to the Amihud or the Pastor-Stambaugh liquidity factor estimated over a 3-year rolling windows.[7]

[7] We follow Lin, Wang and Wu (2011) in constructing aggregate Amihud (2002) and Pástor and Stambaugh (2003) measures for corporate bonds. For the Amihud (2002) measure, we use the monthly average of the ratio of the absolute value of each day's bond return to the transaction volume on that day; we then average

Panel C of Table 7 reports the average values of γ_t when we use the Amihud measure to estimate $\beta_{k,t}^{LIQ}$. The average slope coefficient for the liquidity supply beta is 0.33 with t-statistic of 3.02, almost unchanged from the results in Panel B. In contrast, the slope coefficient on the Amihud beta is negative, implying that bonds that comove positively with the Amihud illiquidity measure work as a hedge against liquidity risk. Though the sign of the estimated slope coefficient is consistent with the theory in Acharya and Pedersen (2005), the point estimates are economically smaller in magnitude compared with the liquidity supply betas, and statistically insignificant.

In Panel D of Table 7, we use the Pastor-Stambaugh measure to compute the illiquidity betas, $\beta_{k,t}^{LIQ}$, and run the Fama-MacBeth regressions. The estimated slope coefficients for the liquidity supply betas and liquidity demand betas are about the same magnitude as Panel B, and are all statistically significant regardless of the control variables.

The results in this section suggest that the liquidity supply and demand shocks estimated from the corporate bond noise measure and dealer capital commitment carry risk premiums with opposite sign in the cross-section of corporate bonds. These results not only underscore the importance of decomposing supply and demand shocks, but also provides an insight as to why liquidity risk is priced in securities. First, the liquidity risk premium arises mainly due to liquidity supply risk. A bond with high liquidity risk tends to fall in value when dealers become more risk averse (a reduction in liquidity supply) or when investors require more immediacy. Second, since the noise measure is driven by inventory frictions rather than adverse selection regarding issuer default risk, the results in this paper provide support to Kondor and Vayanos (2018)'s model of liquidity premia driven by dealer inventory frictions rather than information asymmetry about the securities being traded.

across bonds to calculate an aggregate price impact measure. Finally, we extract shocks to price impact using a univariate AR(1) regression of the aggregate price impact on its lagged value. For the Pástor and Stambaugh (2003) measure, we regress the daily bond return at the security level on the lagged bond return and lagged signed volume. The Pástor and Stambaugh (2003) measure is the the slope coefficient on signed volume, averaged across bonds. As with the Amihud measure, we extract shocks using a univariate AR(1) regression.

4.3 Time-Varying Risk Premiums in the Bond Market

Given the evidence from the cross-section of corporate bonds, it is natural to conjecture that liquidity supply and demand have different implications for time variation in risk premiums. However, as Vuolteenaho (2002) shows, the drivers of risk premiums across individual securities may be different from the aggregate dynamics of risk premiums. Moreover, Chen, Joslin and Ni (2017) uses dealer option inventory and presents an analytical framework to tie aggregate risk premiums to dealer inventory. Specifically, Chen, Joslin and Ni (2017) shows that dealers' capital commitment in option markets predicts returns better when the capital commitment is driven primarily by supply shocks. Since we have our own measures of capital commitment and liquidity supply and demand shocks, we examine if our quantity measure predicts corporate bond market returns when capital commitment is mainly driven by liquidity supply shocks.

To identify the period when the liquidity demand shocks are the primary drivers, we compare the absolute value of the cumulative sum of shocks over the past 13 weeks and set the dummy to one if the demand shocks are greater in magnitude. Denote

$$D_t = \begin{cases} 1 & \text{if } |\sum_{m=1}^{13} v^d_{t-13+m}| > |\sum_{m=1}^{13} v^s_{t-13+m}|, \\ 0 & \text{otherwise.} \end{cases}$$

Figure 4 shows the capital commitment measure and lagged one-year returns on the Merrill Lynch corporate bond market index. The gray shaded area shows the period in which the supply shocks are more important (i.e., $D_t = 0$). Although the periods in which supply shocks dominate demand shocks are fairly spread out over the sample period, it is notable that the periods right after the Lehman Brothers bankruptcy in 2008 through the end of 2009 are mainly driven by supply shocks. In contrast, demand shocks are more important in 2010 and early 2011 when dealer capital commitment was steadily recovering. If dealers' risk aversion is key in determining risk premiums in the aggregate bond market, changes in dealer capital commitment should predict returns when liquidity supply shocks dominate the variation in quantity. To test this hypothesis, we run the following regressions,

$$
\begin{aligned}
R_{t+h} &= b_0 + b_1 q_t + cX_t + \epsilon_{t+h}, \\
R_{t+h} &= b_0 + b_1 q_t D_t + b_2 q_t (1 - D_t) + cX_t + \epsilon_{t+h},
\end{aligned}
$$

where R_{t+h} is the h-period return on the aggregate corporate bond market and X_t is a set of control variables, including the term spread, the dividend-price ratio, the variance risk premium of Bollerslev, Tauchen and Zhou (2009), and an option-based skewness measure. Skewness is the difference in implied volatility of at-the-money 30-day S&P500 index options and out-of-the-money options (i.e., moneyness = 0.9) with the same maturity.

Table 8 reports the estimated slope coefficients and regression R-squared. In Panel A, we run regressions of the aggregate bond market return on q_t with no control variables. The slope coefficients are weakly positive for prediction horizons of 1 week and 4 weeks, but negative for longer horizons. The slope coefficients are all statistically insignificant, which is expected as dealer inventory is driven both by supply and demand shocks. These two shocks carry different (in fact, opposite) information about risk premiums that cancels out with each other.

Panel B of Table 8 reports the coefficients when we include dummy variables for demand-driven periods. Now, the slope coefficients b_1 and b_2 have opposite signs: a rise in capital commitment driven by positive demand shocks predicts subsequent returns positively (albeit insignificantly) while an increase in capital commitment driven by supply shocks is associated with lower returns in the future. At the one-year horizon, adding dummy variable increases the adjusted R-squared to 0.18 from 0.04 in the univariate regression. In Panel C, we add a set of control variables, and show that dealer capital commitment captures different information about risk premiums than other predictors. The regression results are robust to changes in the time window that defines capital commitment as supply-driven or demand-driven from

4 weeks to 52 weeks.

Despite the short time-series, we find evidence that supply-driven and demand-driven changes in capital commitment by dealers have very different implications for time-variation in risk premiums. Higher dealer gross positions predict lower returns only when changes in positions are driven by variation in the ability or willingness of dealers to provide liquidity, and not when changes in capital commitment are mostly driven by fluctuations in demand to trade specific bonds from the same issuer. This result highlights that the asset-pricing implications of changes in dealer balance sheets depend on *why* observed capital commitment varies. Our findings regarding time-varying risk premiums are consistent with the results of the cross-sectional analysis, in that supply and demand shocks carry risk premiums with opposite signs, and the variation in risk premiums associated with supply shocks are robustly significant while the association between liquidity demand and returns is smaller in magnitude and less robust.

5 Conclusion

In this paper, we provide an analytical framework for extracting shocks to liquidity supply and demand by jointly studying price and quantity measures of liquidity. We focus on dealers' use of their balance sheets to accommodate imbalances in the demand for similar-maturity securities from the same issuer. We present a simple model connecting liquidity supply and demand to noise in issuer yield curves and dealers' gross inventory of bonds. In particular, by imposing reasonable sign restrictions on the initial response of noise and dealer gross positions, we extract shocks to liquidity supply and demand in the corporate bond market. Namely, a positive supply shock is one that increases gross dealer positions in corporate bonds and decreases noise in issuer yield curves.

The decomposition of supply and demand shows that liquidity supply in the corporate bond market declined after 2010, especially for IG bonds. After the financial crisis, liquidity demand increased steadily up to 2014 but it has declined more recently. Overall, our findings are consistent with the view that post-crisis banking regulations have made brokers less willing to provide liquidity. The decomposition also shows that the liquidity supply shock explains about half of the variance of unexpected changes in dealer gross positions, suggesting that supply and demand shocks are both important drivers of dealer balance sheets.

Even though this paper focuses on primary dealers' capital commitment to corporate bonds, our emphasis on extracting liquidity supply using observable liquidity proxies is general and applicable to other measures of liquidity quantity such as trading volume as well as to other financial markets. As a possible extension of our framework, we can use even richer cross-sections of bonds, such as Treasury securities and mortgage-backed securities, to understand the connections across markets with respect to liquidity supply, liquidity demand, and returns. This paper also shows that liquidity supply and demand shocks in the corporate bond market have different sets of drivers, with liquidity supply associated with the capital of dealer parent companies and dealer funding costs, for example. However, we have not fully exhausted the list of potential drivers of shifts in liquidity supply and demand. Our findings are consistent with the argument in the literature that the Volcker rule and other financial regulations lead to a deterioration in liquidity, but there may be other factors that help explain the dynamics of supply and demand. Demonstrating a causal link between changing regulation and liquidity supply in the bond market is also left for future research.

Our approach allows us to estimate the risk premiums associated with liquidity supply risk and liquidity demand risk. We find that liquidity supply risk is significantly positively priced in the cross-section of corporate bonds, while liquidity demand risk is not always priced significantly. Moreover, time-variation in the expected return of the aggregate corporate bond portfolio can be explained by dealer gross positions in a way that depends on whether recent liquidity changes were supply- or demand-driven. In particular, expected returns are decreasing in dealer capital commitment when supply shocks have recently been dominant; in contrast, expected returns are unrelated to or possibly increasing in dealer capital commitment when demand shocks have been dominant. Our results provide empirical support for theories in which liquidity supply is a priced risk factor, such as Kondor and Vayanos (2018).

References

Acharya, Viral V, and Lasse Heje Pedersen. 2005. "Asset Pricing with Liquidity Risk." *Journal of Financial Economics*, 77(2): 375–410.

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.*

Bai, Jennie, Turan G. Bali, and Quan Wen. 2018. "Common Risk Factors in the Cross-Section of Corporate Bond Returns." *Journal of Financial Economics*, forthcoming.

Bao, Jack, and Kewei Hou. 2017. "De Facto Seniority, Credit Risk, and Corporate Bond Prices." *Review of Financial Studies*, 30(11): 4038–4080.

Bao, Jack, Jun Pan, and Jiang Wang. 2011. "The Illiquidity of Corporate Bonds." *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." *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.*

Bollerslev, Tim, George Tauchen, and Hao Zhou. 2009. "Expected Stock Returns and Variance Risk Premia." *Review of Financial Studies*, 22(11): 4463–4492.

Brunnermeier, Markus K, and Lasse Heje Pedersen. 2009. "Market liquidity and funding liquidity." *Review of Financial studies*, 22(6): 2201–2238.

Chen, Hui, Scott Joslin, and Sophie Ni. 2017. "Demand for crash insurance, intermediary constraints, and risk premia in financial markets." *Review of Financial Studies*, forthcoming.

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

Chernenko, Sergey, and Adi Sunderam. 2018. "Corporate Bond Liquidity: A Revealed Preference Approach." Purdue University and Harvard Business School, mimeo.

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

Chordia, Tarun, Amit Goyal, Yoshio Nozawa, Avanidhar Subrahmanyam, and Qing Tong. 2017. "Are Capital Market Anomalies Common to Equity and Corporate Bond Markets? An Empirical Investigation." *Journal of Financial and Quantitative Analysis*, 52(4): 1301–1342.

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

Dick-Nielsen, Jens, and Marco Rossi. 2017. "The cost of immediacy for corporate bonds." Mimeo.

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

Edwards, Amy K., Lawrence E. Harris, and Michael S. Piwowar. 2007. "Corporate Bond Market Transaction Costs and Transparency." *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.*

Fama, Eugene F, and James D MacBeth. 1973. "Risk, Return, and Equilibrium: Empirical Tests." *Journal of Political Economy*, 81(3): 607–636.

Fama, Eugene F., and Kenneth R. French. 1993. "Common risk factors in the returns on stocks and bonds." *Journal of Financial Economics*, 33(1): 3 – 56.

Fama, Eugene F., and Kenneth R. French. 2015. "A five-factor asset pricing model." *Journal of Financial Economics*, 116(1): 1 – 22.

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

Fontaine, Jean-Sébastien, and René Garcia. 2012. "Bond Liquidity Premia." *Review of Financial Studies*, 25(4): 1207–1254.

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.*

Greenwood, Robin, and Samuel G. Hanson. 2013. "Issuer Quality and Corporate Bond Returns." *Review of Financial Studies*, 26(6): 1483â1525.

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

Gromb, Denis, and Dmitri Vayanos. 2017. "The dynamics of financially constrainted arbitrage." *The Journal of Finance*, forthcoming.

Grossman, Sanford J, and Merton H Miller. 1988. "Liquidity and market structure." *The Journal of Finance*, 43(3): 617–633.

He, Zhiguo, Bryan Kelly, and Asaf Manela. 2015. "Intermediary Asset Pricing: New Evidence from Many Asset Classes." Mimeo, University of Chicago.

He, Zhiguo, Bryan Kelly, and Asaf Manela. 2017. "Intermediary asset pricing: New evidence from many asset classes." *Journal of Financial Economics*, 126(1): 1 – 35.

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

Kilian, Lutz, and Daniel P Murphy. 2014. "The role of inventories and speculative trading in the global market for crude oil." *Journal of Applied Econometrics*, 29(3): 454–478.

Kondor, Peter, and Dimitri Vayanos. 2018. “Liquidity Risk and the Dynamics of Arbitrage Capital.” *Journal of Finance*, forthcoming.

Lin, Hai, Junbo Wang, and Chunchi Wu. 2011. “Liquidity Risk and Expected Corporate Bond Returns.” *Journal of Financial Economics*, 99(3): 628–650.

Longstaff, Francis A., Sanjay Mithal, and Eric Neis. 2005. “Corporate Yield Spreads: Default Risk or Liquidity? New Evidence from the Credit Default Swap Market.” *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.” *American Economic Review*, 56(1/2): 178–197.

Pan, Kevin, and Yao Zeng. 2018. “ETF Arbitrage under Liquidity Mismatch.” *Working Paper.*

Pástor, Luboš, and Robert F Stambaugh. 2003. “Liquidity Risk and Expected Stock Returns.” *Journal of Political Economy*, 111(3): 642–685.

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

Shanken, Jay. 1992. “On the Estimation of Beta-Pricing Models.” *Review of Financial Studies*, 5(1): 1–33.

Svensson, Lars EO. 1994. “Estimating and Interpreting Forward Interest Rates: Sweden 1992-1994.” 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.

Vuolteenaho, Tuomo. 2002. “What Drives Firm-Level Stock Returns?” *Journal of Finance*, 57(1): 233–264.

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

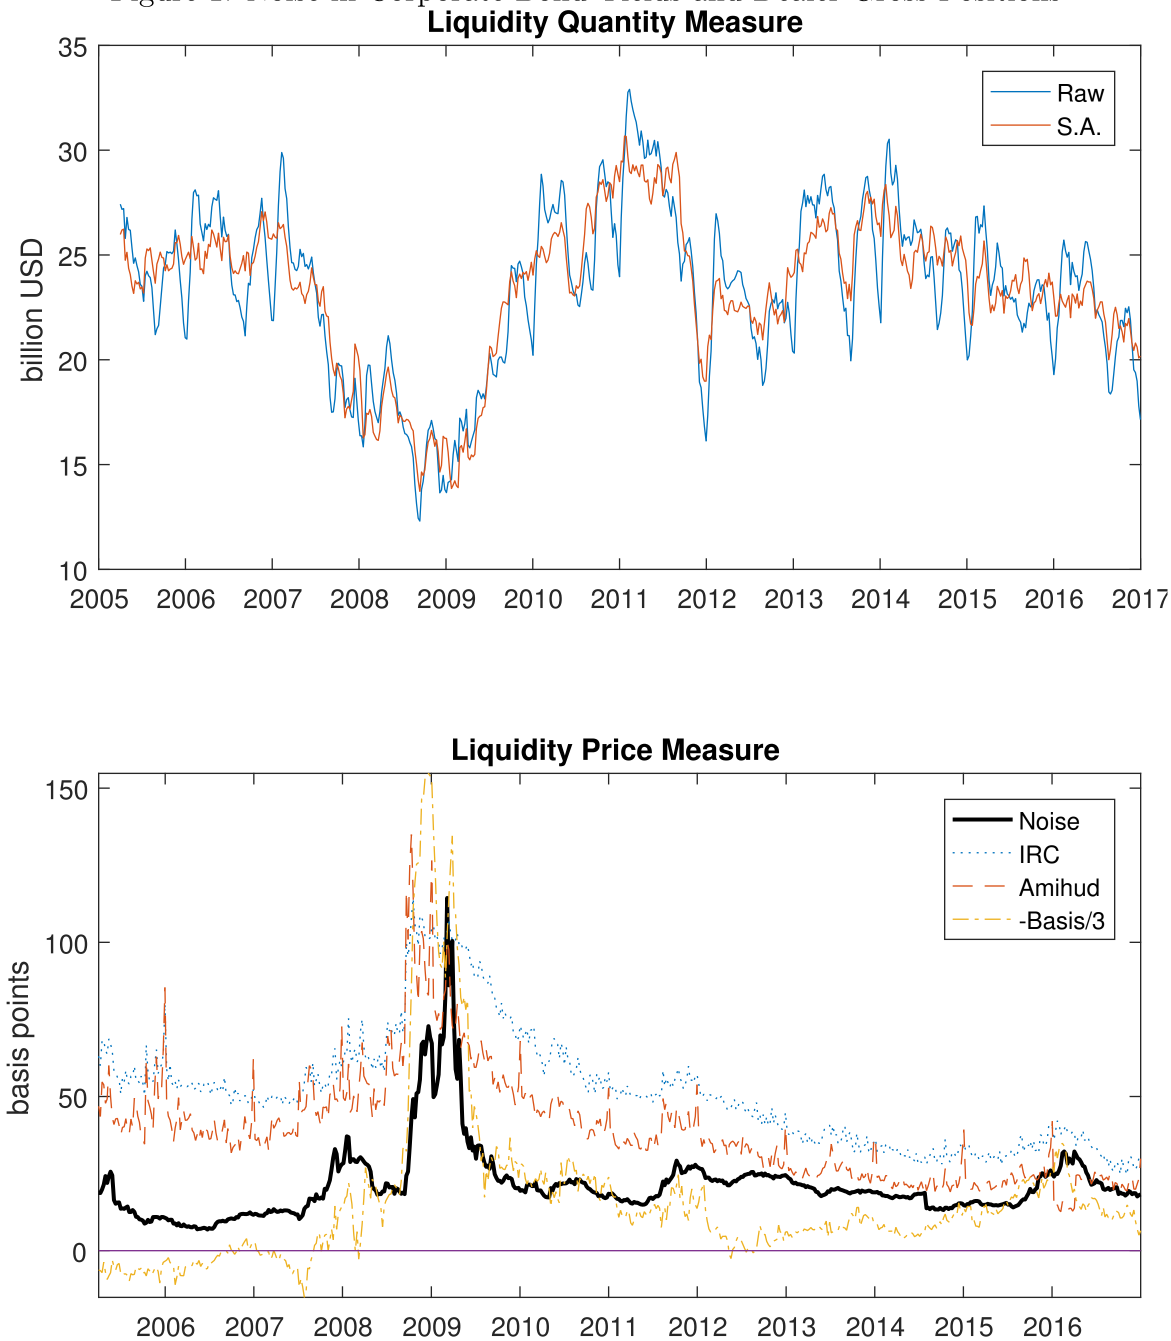

Note: The top panel shows the gross positions in corporate bonds aggregated across primary dealers, scaled by the consumer price index and expressed in 2005 dollars. The bottom panel shows the aggregate measure of noise. Noise is the root mean-squared error of the Nelson-Siegel-Svensson model in basis points, averaged across issuers. IRC is the imputed round-trip cost from Feldhutter (2012), Amihud is Amihud (2002) illiquidity measure, and Basis is the CDS-Bond basis.

Figure 2: Cumulative Sum of Liquidity Shocks

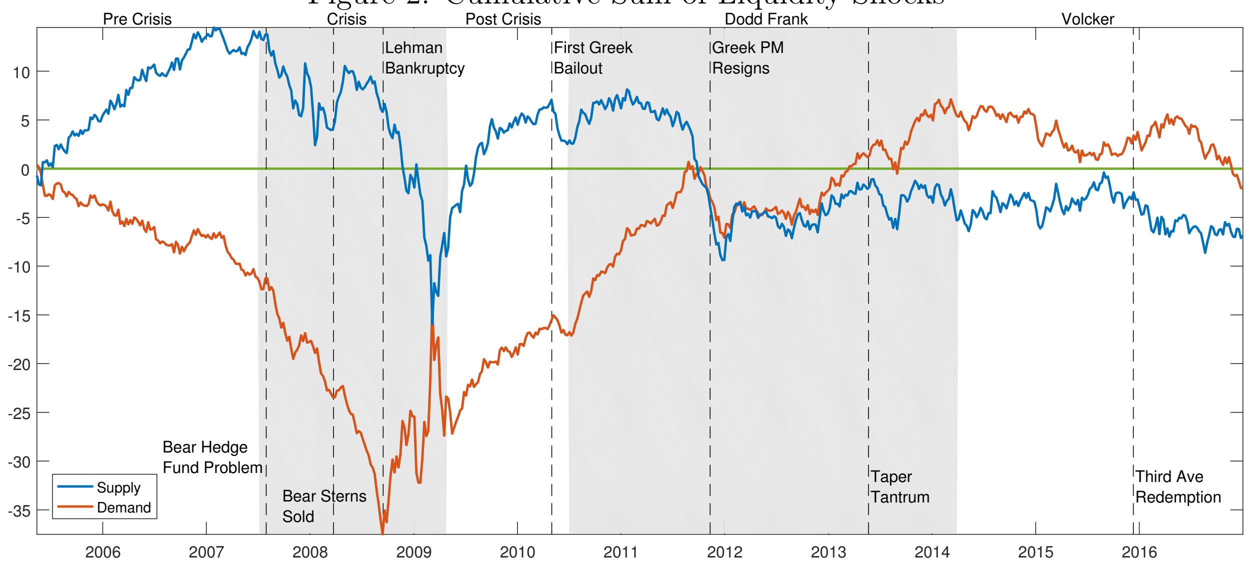

Note: This figure shows the cumulative sums of shocks to liquidity demand and supply. A positive shock reflects an increase in liquidity supply. The data is weekly from 2005 to 2016. The shaded and unshaded area shows five subsample periods: i) the pre-crisis period up to June 2007, ii) the crisis period from July 2007 to April 2009, iii) the post-crisis period from May 2009 to June 2010, iv) the Dodd-Frank regulation period from July 2011 to March 2014, and v) the Volcker rule period after April 2014.

Figure 3: Cumulative Sum of Liquidity Shocks By Ratings

Panel A: Investment-Grade Bonds

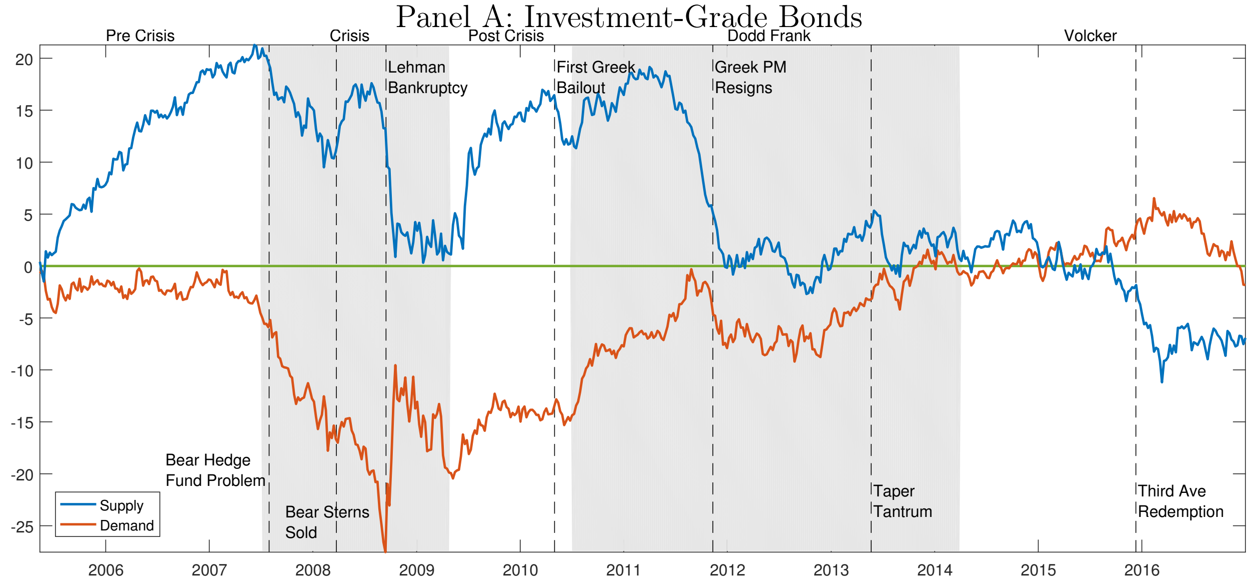

Panel B: High-Yield Bonds

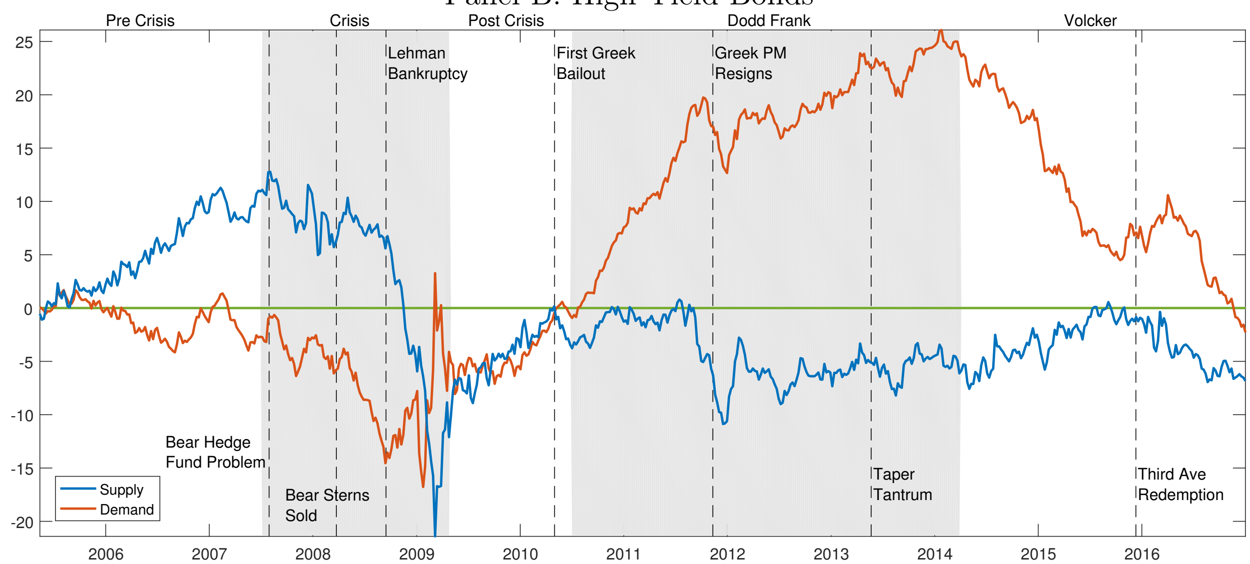

Note: This figure shows the cumulative sums of shocks to liquidity demand and supply. The top panel shows results for the subsample consisting of investment-grade (IG) bonds; the bottom panel uses the subsample of high-yield (HY) bonds. A positive shock reflects an increase in liquidity supply. The data is weekly from 2005 to 2016. The shaded and unshaded area shows five subsample periods: i) the pre-crisis period up to June 2007, ii) the crisis period from July 2007 to April 2009, iii) the post-crisis period from May 2009 to June 2010, iv) Dodd-Frank regulation periods from July 2011 to March 2014, and v) the Volcker rule period after April 2014.

Figure 4: Return and Capital Commitment

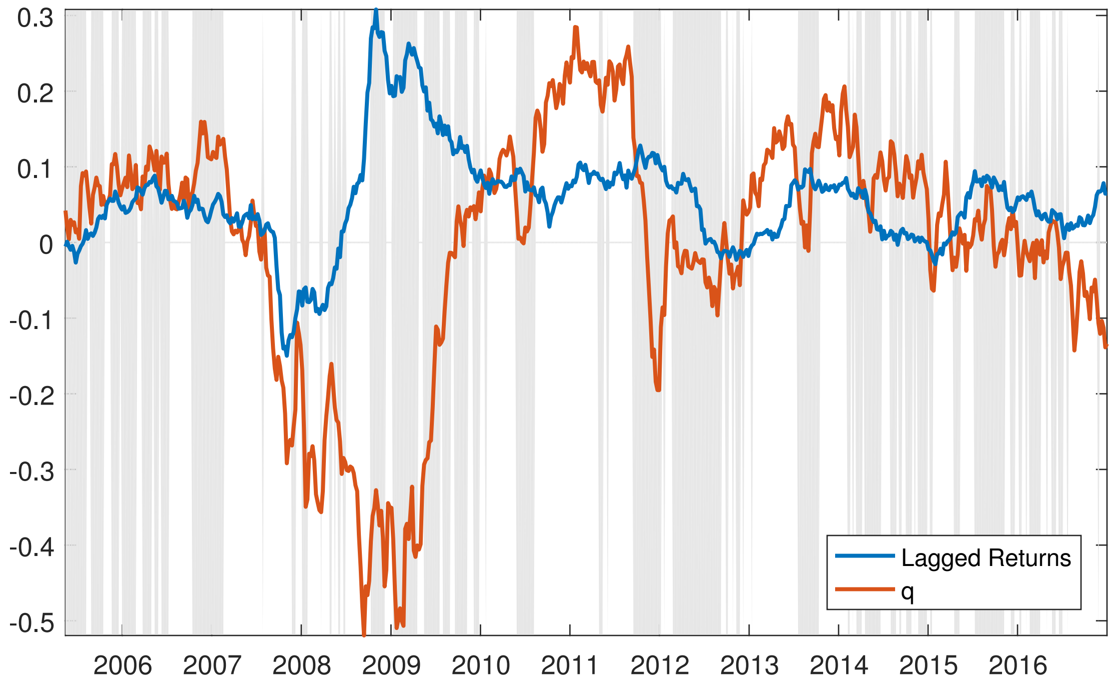

Note: The blue line is the lagged one-year return on the Merrill Lynch corporate bond market index; the red line is seasonally-adjusted dealer capital commitment, q_t. The shaded area shows the period in which liquidity supply shocks are greater than liquidity demand shocks in magnitude, while the unshaded area is the period in which liquidity demand shocks dominate. Specifically, capital commitment in week t is denoted as driven by liquidity demand if $|\sum_{m=1}^{13} v^d_{t-13+m}| > |\sum_{m=1}^{13} v^s_{t-13+m}|$, otherwise it is driven by liquidity supply.

Table 1: Example of the Inventory Construction Method

		Volume	Amount Outstanding for each Transaction ID					End-of-Week
Transaction 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 and a dealer. Zero in bold font emphasizes the fact that we set the inventory balance to zero if the position is not reversed in four weeks. See Section 2.2 for additional details.

Table 2: Summary Statistics: 2005-2016

Panel A: Summary Statistics

	Mean	Std	AR1	AR12
q_t	16.95	0.18	0.98	0.67
p_t	21.45	12.24	0.97	0.68

Panel B: Correlation Matrix

	p_t	Amihud	IRC	Basis
q_t	-0.57	-0.59	-0.51	0.54
p_t		0.57	0.61	-0.86
Amihud			0.93	-0.62
IRC				-0.66

Note: The price variable p_t is our measure of noise in corporate bond yields: the root mean squared error from fitting an issuer-specific Nelson-Siegel-Svensson curve. The RMSE is averaged across issuers each week to obtain the aggregate noise measure. The quantity variable q_t is the log of aggregate gross positions of primary dealers in corporate bonds. Amihud is the median Amihud (2002) measure, IRC is the median imputed round-trip cost measure of Feldhutter (2012), Basis is the CDS-Bond basis. AR1 and AR12 are the autocorrelation at lag 1 and 12, respectively. The frequency is weekly.

Table 3: Structural Shocks and Proxies, Weekly from May 2005 to Dec 2016

	ϵ_t^{VIX}	$\|MFFLOW\|$	$\Delta ISSUE$	$HYSHARE$	ϵ_t^{LEV}	ϵ_t^{TED}	r_{t-1}^{MKT}	$\bar{R}^2$
v_t^s	-0.12							0.02
	(-2.79)							
		-0.01						0.00
		(-0.36)						
			0.12					0.02
			(2.77)					
				-0.08				0.01
				(-2.19)				
					0.16			0.03
					(3.23)			
						-0.12		0.02
						(-3.21)		
							0.01	0.00
							(0.15)	
	0.01	-0.02	0.11	-0.04	0.16	-0.11	0.01	0.06
	(0.25)	(-0.62)	(2.93)	(-1.06)	(2.72)	(-3.16)	(0.28)	
v_t^d	0.10							0.01
	(1.61)							
		0.05						0.00
		(1.90)						
			0.04					0.00
			(1.19)					
				0.00				0.00
				(0.07)				
					-0.07			0.01
					(-1.22)			
						0.07		0.01
						(1.67)		
							-0.04	0.00
							(-0.83)	
	0.07	0.06	0.05	0.00	-0.03	0.06	-0.03	0.01
	(1.17)	(2.02)	(1.26)	(0.08)	(-0.50)	(1.51)	(-0.68)	

Note: The table reports the coefficient estimates for a regression:

$$v_t = b_1\epsilon_t^{VIX} + b_2|MFFLOW_t| + b_3\Delta ISSUE_t + b_4 HYSHARE_t + b_5\epsilon_t^{LEV} + b_6\epsilon_t^{TED} + b_7 r_{t-1}^{MKT} + u_t$$

where ϵ_t^{VIX} is a shock to VIX, $|MFFLOW_t|$ is the absolute value of the weekly mutual fund flow for US domestic investment-grade corporate bond funds, $\Delta ISSUE_t$ is the growth rate in corporate bonds issues in week t, $HYSHARE_t$ is the share of high-yield bond issues among total corporate bond issues in week t, ϵ_t^{LEV} is a shock to the intermediary capital ratio of He, Kelly and Manela (2015), ϵ_t^{TED} is a shock to TED spread, and r_{t-1}^{MKT} is a lagged return on corporate bond market portfolio. All variables are standardized to have mean zero and standard deviation of one. Shocks to VIX, intermediary capital, and the TED spread are a residual of univariate autoregressions on each variable. $\bar{R}^2$ is adjusted R-squared.

Table 4: Structural Shocks and Proxies, Monthly from May 2005 to Dec 2016

	ϵ_t^{VIX}	$\lvert MFFLOW \rvert$	$\Delta ISSUE$	$HYSHARE$	ϵ_t^{LEV}	ϵ_t^{TED}	r_{t-1}^{MKT}	$\bar{R}^2$
v_t^s	-0.13							0.09
	(-3.77)							
		0.01						0.00
		(0.21)						
			0.04					0.01
			(1.25)					
				0.06				0.02
				(1.26)				
					0.16			0.14
					(4.51)			
						-0.06		0.02
						(-1.97)		
							0.10	0.05
							(2.75)	
	-0.03	-0.01	0.03	0.04	0.13	-0.04	0.06	0.16
	(-0.46)	(-0.19)	(1.04)	(0.94)	(2.71)	(-1.19)	(2.19)	
v_t^d	0.08							0.03
	(2.67)							
		0.09						0.04
		(2.17)						
			0.05					0.02
			(1.75)					
				0.01				0.00
				(0.12)				
					-0.03			0.01
					(-0.68)			
						-0.04		0.01
						(-1.42)		
							-0.05	0.01
							(-0.74)	
	0.09	0.06	0.06	0.01	0.03	-0.05	-0.06	0.07
	(3.24)	(1.84)	(1.88)	(0.23)	(0.63)	(-1.10)	(-1.30)	

Note: The table reports the coefficient estimates for a regression:

$$v_t = b_1\epsilon_t^{VIX} + b_2|MFFLOW_t| + b_3\Delta ISSUE_t + b_4 HYSHARE_t + b_5\epsilon_t^{LEV} + b_6\epsilon_t^{TED} + b_7 R_{t-1}^{MKT} + u_t$$

where ϵ_t^{VIX} is a shock to VIX, $|MFFLOW_t|$ is the absolute value of a weekly mutual fund flow to US domestic investment-grade corporate bond funds, $\Delta ISSUE_t$ is the growth rate in corporate bonds issues in week t, $HYSHARE_t$ is the share of high-yield bond issues among total corporate bond issues in week t, ϵ_t^{LEV} is a shock to intermediary capital ratio of He, Kelly and Manela (2015), ϵ_t^{TED} is a shock to TED spread, and r_{t-1}^{MKT} is a lagged return on corporate bond market portfolio. All variables are standardized to have mean zero and standard deviation of one, and shocks are extracted as a residual of univariate autoregressions on each variable. $\bar{R}^2$ is adjusted R-squared.

Table 5: Monte-Carlo Comparison of Sign-Identified Structural Shocks with True Values

	Mean	Std.	Percentiles				
			2.5	16	50	84	97.5
$R^2_{m,s}$	0.97	0.02	0.93	0.96	0.97	0.98	0.99
$R^2_{m,d}$	0.97	0.02	0.93	0.96	0.97	0.98	0.99

Note: The table shows the R^2 measure defined by:

$$R^2_m = 1 - \frac{\Sigma_t(\hat{v}_{m,t} - v_{m,t})^2}{\Sigma_t v^2_{m,t}},$$

where $v_{m,t}$ is the true structural shocks generated by m-th run of Monte Carlo simulation and $\hat{v}_{m,t}$ is the posterior mean structural shocks estimated by applying the structural VAR with sign restrictions on the simulated data. Mean, Std., and Percentiles show the summary statistics of the distribution of R-squared across 1,000 simulations. $R^2_{m,s}$ is R-squared for supply shocks, and $R^2_{m,d}$ is R-squared for demand shocks.

Table 6: Corporate Bond Portfolios Sorted on Supply and Demand Betas, 2007-2016

	Liquidity Supply Betas						Liquidity Demand Betas					
	Low	2	3	4	High	H-L	Low	2	3	4	High	H-L
						Average Excess Returns						
$E[R^e]$	0.49	0.39	0.41	0.54	1.01	0.51	0.85	0.47	0.37	0.48	0.63	-0.22
$t(E[R^e])$	(3.35)	(4.68)	(5.82)	(5.89)	(5.67)	(3.56)	(5.58)	(5.95)	(5.34)	(5.19)	(3.96)	(-2.39)
						Fama-French 5 Factor Model + TERM + DEF						
α	-0.10	0.06	0.13	0.19	0.41	0.51	0.30	0.14	0.10	0.13	-0.02	-0.32
$t(\alpha)$	(-0.86)	(1.10)	(2.64)	(2.73)	(2.93)	(3.75)	(2.33)	(2.39)	(2.02)	(2.11)	(-0.20)	(-3.18)
						Bai, Bali and Wen (2018) 4 Factor Model						
α	-0.18	0.03	0.09	0.12	0.35	0.54	0.18	0.10	0.05	0.07	-0.04	-0.22
$t(\alpha)$	(-2.89)	(0.57)	(1.57)	(1.43)	(3.40)	(3.84)	(2.95)	(1.69)	(1.12)	(1.21)	(-0.39)	(-1.74)
						He, Kelly and Manela (2017) 2 Factor Model						
α	0.36	0.32	0.34	0.42	0.70	0.34	0.56	0.36	0.30	0.42	0.43	-0.13
$t(\alpha)$	(1.89)	(3.12)	(3.90)	(3.97)	(3.88)	(2.29)	(3.53)	(3.88)	(3.54)	(3.59)	(2.18)	(-1.09)
						Average characteristics of bonds:						
β_s, β_d	-2.92	-0.98	-0.26	0.65	4.43		-4.06	-0.78	0.04	0.80	3.20	
Maturity (yr)	13.8	7.3	5.8	6.7	8.2		8.2	6.5	5.8	8.0	13.4	
Size (mil)	821.0	852.1	871.8	808.5	768.4		688.8	798.7	885.1	878.5	872.6	
Age (yr)	6.07	5.76	5.93	6.27	6.84		6.74	6.22	5.93	5.81	6.17	
Roll (%)	1.01	0.63	0.58	0.75	1.20		1.14	0.66	0.56	0.68	1.11	
IRC (%)	0.72	0.53	0.51	0.60	0.87		0.81	0.56	0.49	0.56	0.79	
						Fraction of Credit Ratings						
Aa+	10%	11%	11%	6%	2%		2%	7%	11%	12%	9%	
A	38%	42%	38%	30%	20%		17%	33%	40%	40%	37%	
Baa	31%	34%	37%	40%	31%		32%	38%	37%	36%	30%	
HY	20%	13%	13%	22%	45%		46%	21%	11%	12%	22%	

Note: The table shows monthly returns for portfolios of corporate bonds sorted on betas with respect to liquidity supply and demand shocks. Betas are estimated using 3-year rolling windows. Each month, we sort corporate bonds into five bins based on their betas and compute value-weighted average returns for each bin. α is an intercept from the factor regression:

$$R_t^e = \alpha + \Sigma_{j=1}^J \beta_j f_{j,t} + \epsilon_t,$$

where $f_{j,t}$ is j-th factor of factor pricing models, including Fama and French (2015) augumented with TERM and DEF factors, Bai, Bali and Wen (2018) and He, Kelly and Manela (2017). The data starts in 2005 and ends in 2016; portfolio returns start in 2007 due to an initial 2-year estimation window for betas.

Table 7: Fama-MacBeth (1973) Regressions with Amihud and Pastor-Stambaugh Measures

β^q	β^s	β^d	β^{LIQ}	$\log Roll$	$\log Mat$	$\log Size$	D^A	D^{Baa}	D^{HY}
Panel A: Shocks to Capital Commitment									
0.16									
(1.72)									
0.09				0.28	0.02	0.07	-0.28	-0.23	-0.09
(1.20)				(2.90)	(0.36)	(1.91)	(-3.45)	(-2.69)	(-0.35)
Panel B: Supply and Demand Risk Premiums									
	0.35								
	(3.06)								
	0.23			0.27	0.11	0.06	-0.21	-0.18	0.08
	(2.89)			(3.48)	(1.81)	(1.70)	(-2.91)	(-2.27)	(0.54)
		-0.17							
		(-2.00)							
		-0.15		0.27	0.09	0.07	-0.20	-0.17	0.13
		(-2.26)		(3.48)	(1.55)	(1.86)	(-2.87)	(-2.05)	(0.84)
Panel C: With Amihud (2002) Measure									
	0.33		-0.18						
	(3.02)		(-1.25)						
	0.23		-0.13	0.24	0.13	0.07	-0.20	-0.20	0.02
	(2.93)		(-1.09)	(3.59)	(2.32)	(2.09)	(-2.93)	(-2.56)	(0.14)
		-0.13	-0.08						
		(-1.81)	(-0.58)						
		-0.15	-0.08	0.26	0.11	0.07	-0.20	-0.19	0.06
		(-2.29)	(-0.65)	(3.65)	(1.90)	(2.22)	(-2.99)	(-2.45)	(0.45)
Panel D: With Pastor-Stambaugh (2003) Measure									
	0.40		0.04						
	(3.07)		(0.36)						
	0.27		0.02	0.25	0.12	0.07	-0.21	-0.19	0.03
	(2.97)		(0.17)	(3.43)	(2.15)	(1.88)	(-2.96)	(-2.49)	(0.18)
		-0.19	-0.02						
		(-2.12)	(-0.17)						
		-0.17	0.00	0.28	0.10	0.07	-0.21	-0.18	0.09
		(-2.34)	(-0.01)	(3.53)	(1.62)	(2.06)	(-2.90)	(-2.30)	(0.58)

Note: The table reports average slope coefficients from monthly cross-sectional regressions of corporate bond returns, $R_{k,t} = \gamma_{0t} + \gamma_{1t}\beta^s_{k,t} + \gamma_{2t}\beta^{LIQ}_{k,t} + \gamma_{3t}X_{k,t} + \epsilon_{k,t}$, or $R_{k,t} = \gamma_{0t} + \gamma_{1t}\beta^d_{k,t} + \gamma_{2t}\beta^{LIQ}_{k,t} + \gamma_{3t}X_{k,t} + \epsilon_{k,t}$, where $\beta^s_{k,t}$ is the beta of bond returns with respect liquidity supply shocks, $\beta^d_{k,t}$ is the liquidity demand beta, $\beta^{LIQ}_{k,t}$ is the beta of bond returns with respect to the Amihud or Pastor-Stambaugh measure, and $X_{k,t}$ is a vector of bond-specific control variables including the Roll measure, maturity, face value, and dummy variables for credit ratings. Explanatory variables except for rating dummies are standardized. t-statistics based on Shanken (1992) standard errors are reported in brackets. Betas are estimated using a rolling window of 3 years; since two years are used for the initial estimation window, monthly cross-sectional return regressions begin in 2007.

Table 8: Bond Market Aggregate Return Forecasting Regressions

Horizon (weeks)	1	4	13	26	52
Panel A: Unconditional Forecasting Regressions					
q_t	0.18	0.59	-0.52	-2.98	-8.48
	(0.49)	(0.31)	(-0.11)	(-0.30)	(-0.69)
$\bar{R}^2$	0.00	0.00	0.00	0.01	0.04
Panel B: Conditional Forecasting Regressions					
$q_t D_t$	0.44	2.57	6.17	7.29	5.31
	(0.72)	(0.94)	(1.33)	(1.21)	(0.55)
$q_t(1 - D_t)$	-0.04	-1.36	-7.24	-14.18	-23.48
	(-0.11)	(-0.73)	(-2.18)	(-1.50)	(-2.40)
$\bar{R}^2$	0.01	0.07	0.19	0.17	0.18
Panel C: Conditional Forecasting Regressions with Controls					
$q_t D_t$	0.24	1.93	3.54	2.93	-2.87
	(0.38)	(0.70)	(0.80)	(0.61)	(-0.29)
$q_t(1 - D_t)$	0.07	-0.99	-6.43	-12.20	-19.33
	(0.20)	(-0.48)	(-2.90)	(-1.86)	(-2.78)
VRP_t	0.02	0.04	0.01	0.11	0.02
	(3.97)	(2.45)	(0.30)	(3.48)	(0.24)
DP_t	0.22	0.94	2.76	6.26	10.04
	(1.26)	(1.11)	(0.91)	(1.10)	(1.26)
$TERM_t$	0.02	0.05	0.39	0.37	1.41
	(0.39)	(0.27)	(0.82)	(0.44)	(1.07)
$SKEW_t$	0.02	0.07	0.29	0.51	0.82
	(0.87)	(0.94)	(1.44)	(1.69)	(1.50)
$\bar{R}^2$	0.04	0.10	0.24	0.28	0.34

Note: The table reports the results for estimating regressions of the form,

$$\begin{aligned} R_{t+h} &= b_0 + b_1 q_t + cX_t + \epsilon_{t+h} \\ R_{t+h} &= b_0 + b_1 q_t D_t + b_2 q_t(1 - D_t) + cX_t + \epsilon_{t+h} \end{aligned}$$

where R_{t+h} is the return between weeks t and $t+h$ on the aggregate corporate bond market portfolio, q_t is logarithm of real gross dealer capital commitment and D_t is a dummy variable equal to 1 if demand shocks are greater in magnitude than supply shocks over the past 13 weeks; see Section 4.3 for details. The sample is weekly from May 2005 to December 2016. Standard errors are corrected for overlapping observations following Hansen and Hodrick (1980).

A A Simple Model of Noise in Bond Prices

This appendix develops a highly stylized model of security prices and dealer positions. The main purpose of this model is to illustrate the comparative statics of noise and dealer gross positions in a very simple setting.

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.[8]

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. 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.[9] 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.

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

[8] 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 $v = \frac{1}{R}$.

[9] The model can be modified to include non-dealer intermediaries that, like dealers, are able to trade in both securities markets; the proposition below still holds.

maximize $E[w_i] - \frac{1}{2\gamma} Var[w_i]$, where γ is i-investors' risk tolerance.[10] 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:

$$\psi_i = \mu - p_{i,1} \tag{7}$$

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.

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

$$\Omega = \begin{bmatrix} \sigma + \lambda\sigma_\epsilon & \sigma \\ \sigma & \sigma + \lambda\sigma_\epsilon \end{bmatrix} \tag{8}$$

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

$$x = \Omega^{-1}\gamma_D\psi \tag{9}$$

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

$$y = \frac{1}{\sigma}\left(\gamma\psi - u\begin{bmatrix} 1 \\ -1 \end{bmatrix}\right). \tag{10}$$

Market clearing requires that

$$x + y = g \tag{11}$$

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

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 σ_ϵ) 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

B.1 Difference in sample between price and quantity measures

In estimating p_t, we focus on relatively large issuers while we use inventory for all corporate bonds to compute the quantity measure, q_t. Therefore, it is important to understand if and how the selection bias affects our measure of noise. Table 9 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 9 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.

Table 9: Comparison of Samples Used to Compute Fitting Errors and TRACE Sample

		# obs	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 Feldhutter (2012), Amihud is median Amihud measure (2002), and Volume is median dollar trading volume in millions.

B.2 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 10 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 10 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.

Table 10: 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.

C List of Primary Dealers

Table 11 shows the list of firms designated as primary dealers by the Federal Reserve Bank of New York.

Table 11: 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	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

Internet Appendix to "Liquidity Supply and Demand in the Corporate Bond Market"

Jonathan Goldberg
Federal Reserve Board

Yoshio Nozawa
HKUST

July 24, 2018

1 Behavior of Corporate Bond Noise

1.1 Market Segmentation in Corporate Bond Market

In this section, we examine holding of new corporate bond issues by institutional investors. eMaxx provides the quarterly holding data of corporate bonds for mutual funds, insurance firms and other large institutions. We use all corporate bonds issued between 1998 and 2014, and gather the quarter-end holding information for the quarter when the bond is issued. We classify bonds into issue size-weighted quintiles based on their maturity, such that if an investor buy all new issues proportional to the issue size, the share of each quintile for the investor is 20%. If the investor deviates from the issue size-weighted portfolio, then the share of a quintile for his portfolio may become more or less than 20%.

For each investor in eMaxx, we compute the portfolio share among each maturity bin. For example, if an investor purchases 100 million dollars of 5-year bond and 100 million dollars of 7-year bond over this sample period, then his portfolio share is 50% in the first quintile (which includes bonds with maturity $<=$ 5 years), and 50% in the second quintile (which includes bonds with maturity $>$5 years and $<=$8.2 years). We denote the portfolio share of investor ι on quintile m as $w_{\iota,m}$ such that $\sum_{m=1}^{5} w_{\iota,m} = 1$ holds.

Table 1 presents summary statistics of the weights on the largest quintile, $max_{\iota} w_{\iota,m}$, as well as the Herfindahl index, $\sum_{m=1}^{5} w_{\iota,m}^2$. For any type of investors, the share of the largest maturity quintile is significantly higher than the benchmark of 20%. The average investor in sample puts 55% of funds into his preferred maturity bucket, while median investor puts 48%. About 15% of investors in eMaxx invests only one maturity bucket (not reported in the table). This maturity concentration is prevalent through various investor classes, though mutual funds puts slightly higher weights to their favorite maturity bucket (60% on average) than other investor types. The Herfindahl index also shows significant maturity concentration of corporate bond purchases across investor types, with both average and median investor

Table 1: Shares of Bond Investment in Maturity Buckets: 1998-2014

	Mean	Percentiles				
		5%	10%	50%	90%	95%
Share of Top Quintile						
All investors	0.55	0.27	0.29	0.48	1	1
Mutual funds	0.60	0.27	0.30	0.53	1	1
Life insurers	0.50	0.26	0.29	0.43	0.96	1
Property insurers	0.52	0.28	0.29	0.45	1	1
Others	0.52	0.26	0.28	0.45	1	1
Herfindahl Index						
All investors	0.46	0.21	0.22	0.36	1	1
Mutual funds	0.51	0.21	0.23	0.40	1	1
Life insurers	0.41	0.21	0.22	0.31	0.92	1
Property insurers	0.43	0.22	0.23	0.34	1	1
Others	0.43	0.21	0.22	0.33	1	1

has the Herfindahl index higher than the benchmark case of 20%.

1.2 Fitting Errors and De Facto Seniority of Short-Term Bonds

In this subsection, we examine if there is a systematic bias in fitting the Nelson-Siegel-Svensson curve which may drive our estimates for noise. Figure 1 presents examples of the fitted curves. We plot yield to maturity for the three largest issuers on October 24, 2008 (crisis times) and on December 23, 2016 (normal times) as well as the fitted curve. The Nelson-Siegel-Svensson curve fits the cross-section of yields reasonably well in normal times, when yields are distributed smoothly across maturity. In contrast, during the financial crisis, yield to maturity does not line up well with each other, causing a problem for a smooth curve to fit. Thus, the fit of the curve deteriorates during the crisis, leading to greater noise.

In particular, Bao and Hou (2017) show that short-term bonds are de facto senior to long-term bonds, which may undermine our assumption that the Nelson-Siegel-Svensson curve is flexible enough to capture issuer's default risk, and that the deviation from the fitted curve must come from dealer's inventory frictions. For example, if the fitted curve fails to capture de facto seniority of short-term bonds, the fitting error may be on average negative for short-term bonds and positive for long-term bonds.

To address the concern of the bias in fitting the curve, we examine the summary statistics of fitting errors. Table 2 reports the mean, median and standard deviation of the fitting errors, where the statistics are computed across issuers and time within maturity and credit rating bins. Specifically, we sort sample of fitting errors into five equal-sized maturity buckets and then into four bins based on credit rating.

Mean and median fitting errors are generally small for all maturity and credit rating bins. For investment-grade bonds, the mean fitting errors are -0.4 to -0.8 basis points for bonds with shortest maturity (less than 3.4 years), which are only slighly larger than the mean error for the long-term bonds with average errors of -4.3 to -6.4 basis points. For HY bonds, the mean fitting error is 10.2 basis points for shor-term bonds while the error is -2.4 basis points for long-term bonds.

Overall, we find little evidence for the bias in fitting the yield curve across maturity. Both mean and median errors are economically small compared with their standard deviation. In particular, the fitting errors for short-term bonds are on average not higher than the errors for long-term bonds, which implies that the Nelson-Siegel-Svensson curve is flexible enough to capture the de facto seniority effect of Bao and Hou (2017).

2 Effects of Liquidity Supply and Demand Shocks

Figures 2 and 3 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 2 and 3 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

Figure 1: Examples of the Fit of the Nelson-Siegel-Svensson Curves

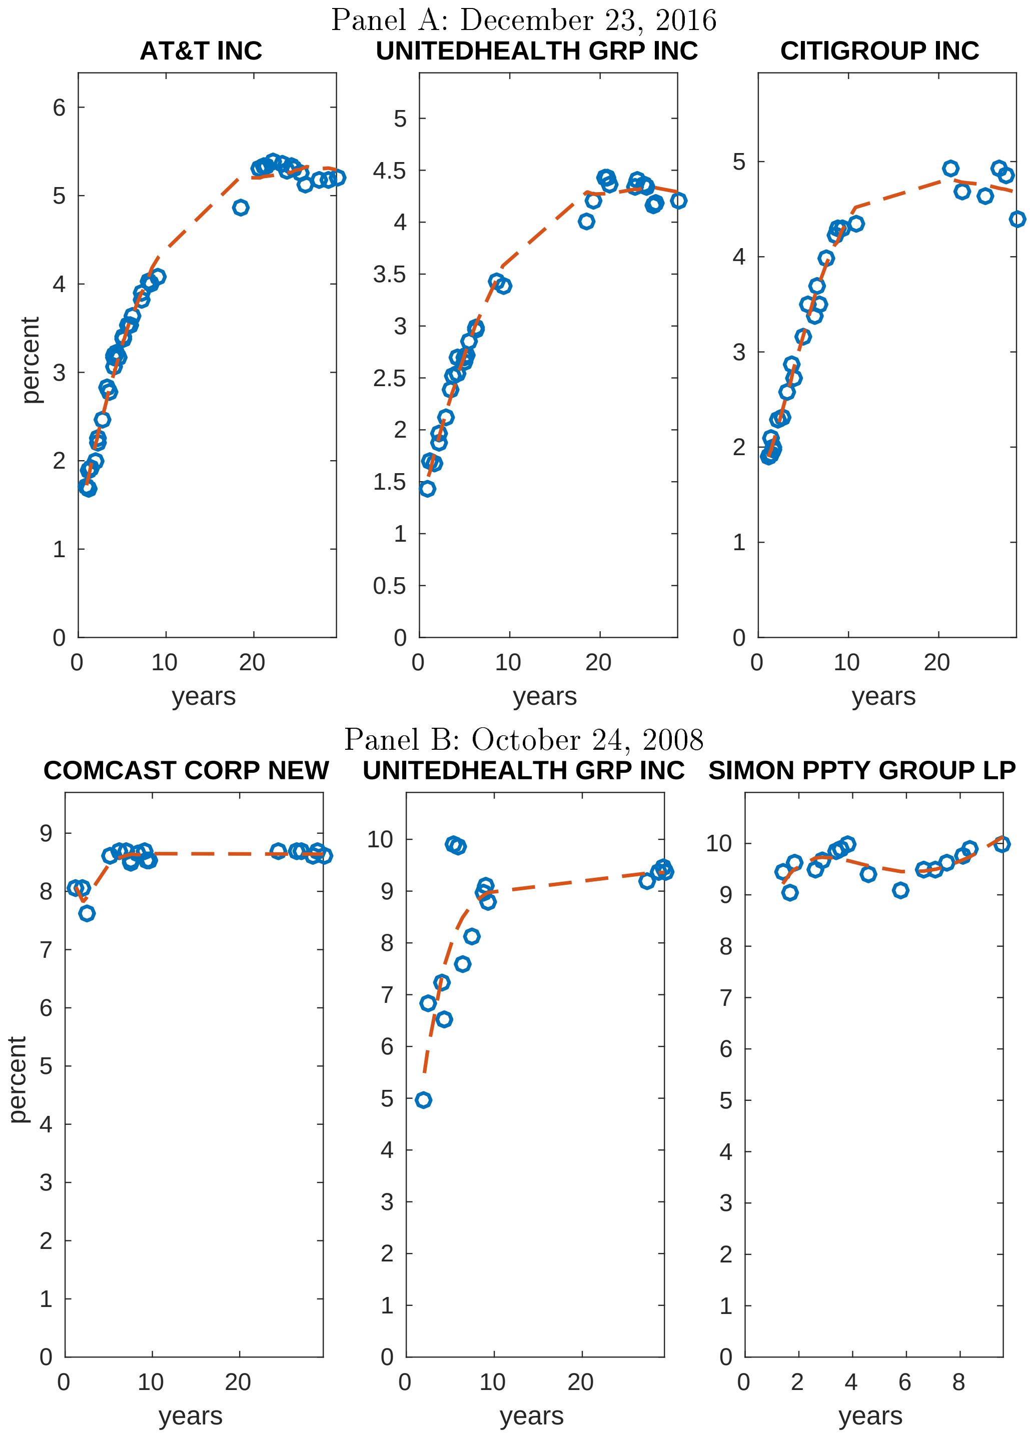

Note: The figures show the observed yields for the three largest issuers in sample on October 24, 2008 (crisis period) and on December 23, 2016 (normal period) and the fitted Nelson-Siegel-Svensson curve. The x-axis shows the time to maturity in years, and the y-axis is yield to maturity in percent.

Table 2: Summary Statistics for Fitting Errors of the Nelson-Siegel-Svensson Curve

	Maturity				
	Short	2	3	4	Long
	Panel A: Mean fitting errors				
Aa+	-0.4	-0.4	-0.6	-0.5	-4.3
A	-0.8	-1.2	0.8	-3.0	-5.2
Baa	-0.8	-1.2	0.7	-3.3	-6.4
HY	10.2	5.1	-5.7	-8.1	-2.4
	Panel B: Median fitting errors				
Aa+	-0.4	-0.5	-0.3	-0.6	-3.6
A	-0.5	-0.4	0.4	-2.9	-3.4
Baa	-0.6	-0.5	0.5	-2.9	-3.7
HY	-0.9	-0.7	-1.1	-2.4	1.9
	Panel C: Standard deviation				
Aa+	7.5	10.1	9.3	10.1	10.8
A	11.0	12.6	13.8	15.4	12.6
Baa	17.4	23.0	16.5	19.6	18.5
HY	277.5	249.1	103.1	78.4	42.3

Note: Table presents the mean, median and standard deviation of the fitting error of the Nelson-Siegel-Svensson model, where statistics are computed across issuers and time. The errors are reported in basis points. Bonds are sorted into five equal-sized maturity buckets, and the cutoff values for the quntiles are 3.4, 5.6, 8.3, 23.8 years to maturity.

Figure 2: Impulse Responses to a Liquidity Supply Shock

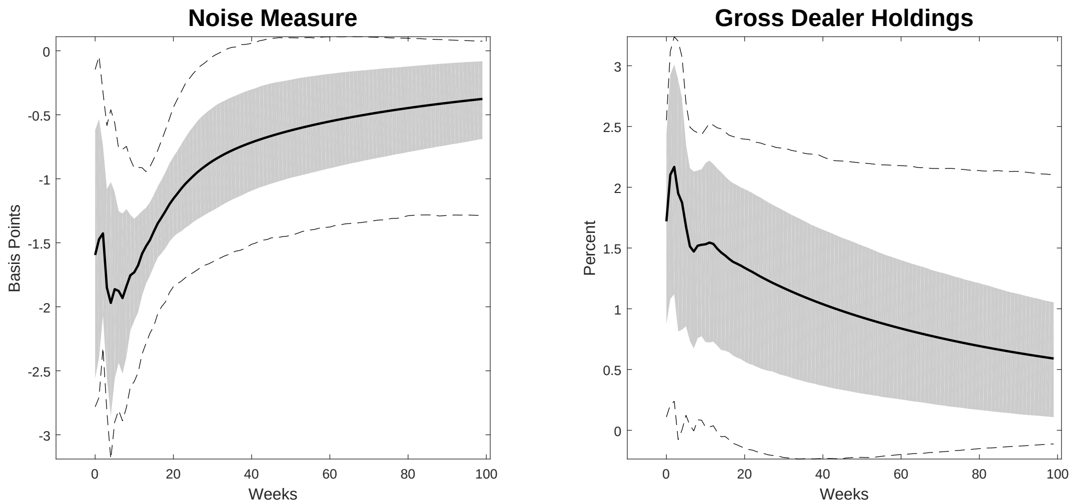

Note: The mean impulse response is shown in black. The shaded area marks a pointwise 68-percent credible interval around the median. The dashed lines mark a pointwise 95-percent credible interval around the median.

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 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 4 shows the share of forecast error variance due to

Figure 3: Impulse Responses to a Liquidity Demand Shock

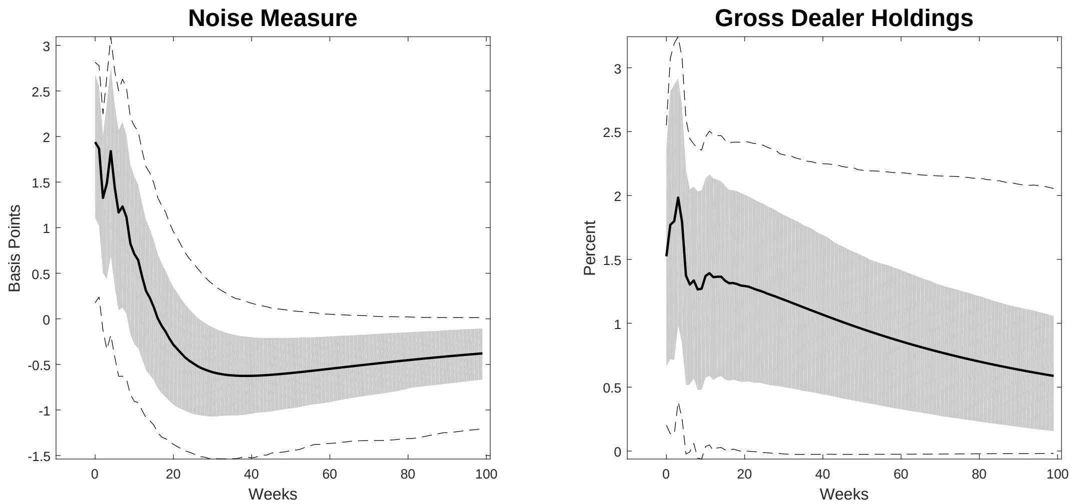

Note: The mean impulse response is shown in black. The shaded area marks a pointwise 68-percent credible interval around the median. The dashed lines mark a pointwise 95-percent credible interval around the median.

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.

3 Liquidity Contagion

We 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 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

Figure 4: Forecast Error Variance Decomposition

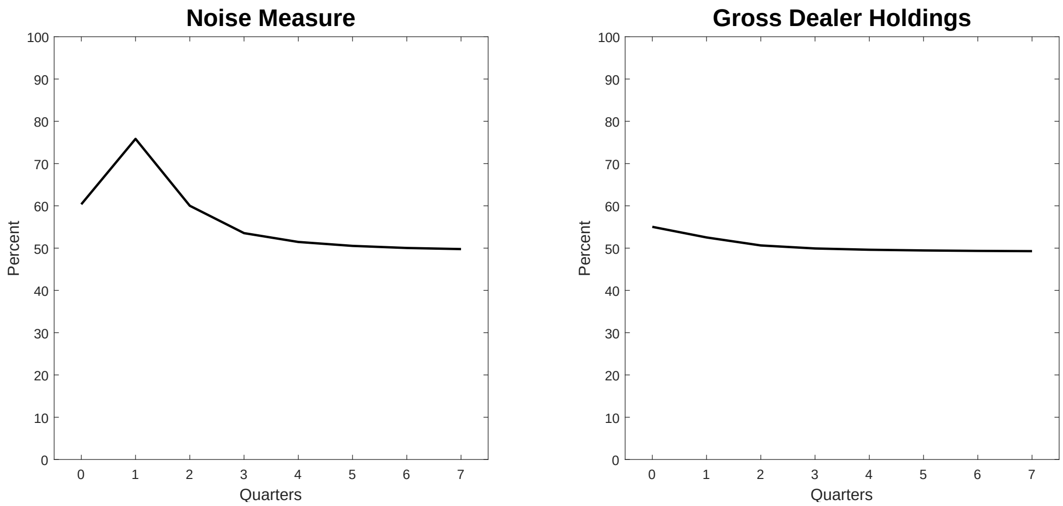

Note: Each plot shows the share of forecast error variance for a given variable due to the liquidity supply shock. The forecast error variance decomposition is calculated by: drawing the parameters of the structural VAR from the posterior distribution; calculating the forecast error variance decomposition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

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 VAR as in the main text,

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

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 5 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 6 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 5: Response of Noise (IG), inventory (IG) and Noise (HY) to a HY shock

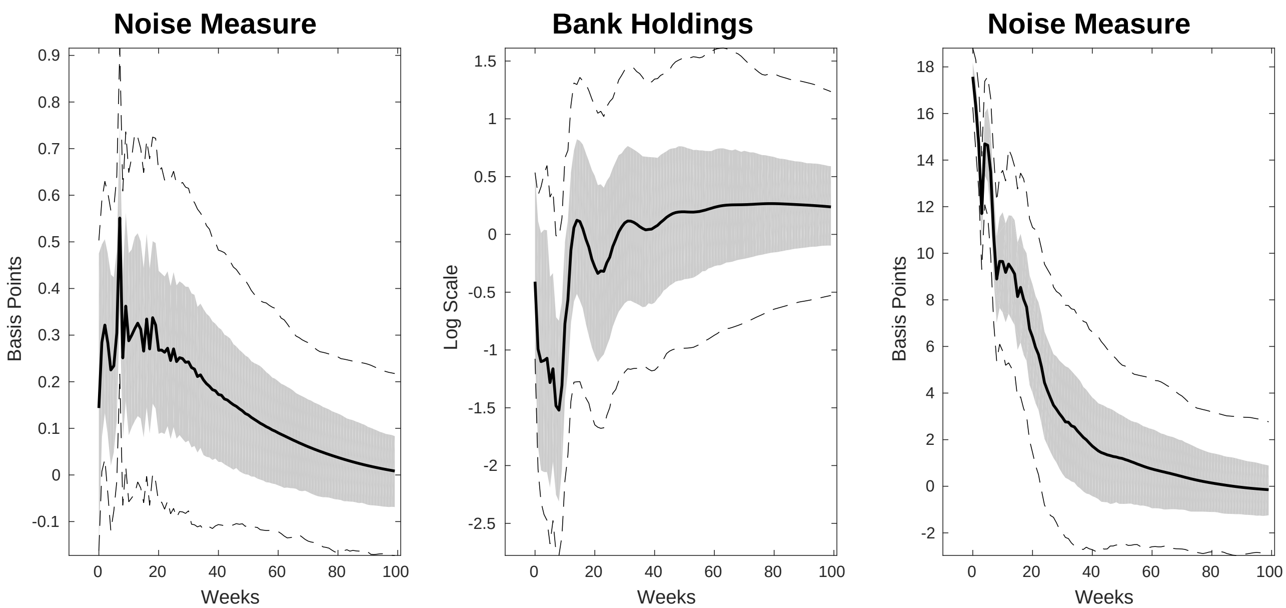

Note: The figures plot the response of endogenous variables to an idiosyncratic shock to the noise in HY bonds. The mean impulse response is shown in black. The shaded area marks a pointwise 68-percent credible interval around the median. The dashed lines mark a pointwise 95-percent credible interval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

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

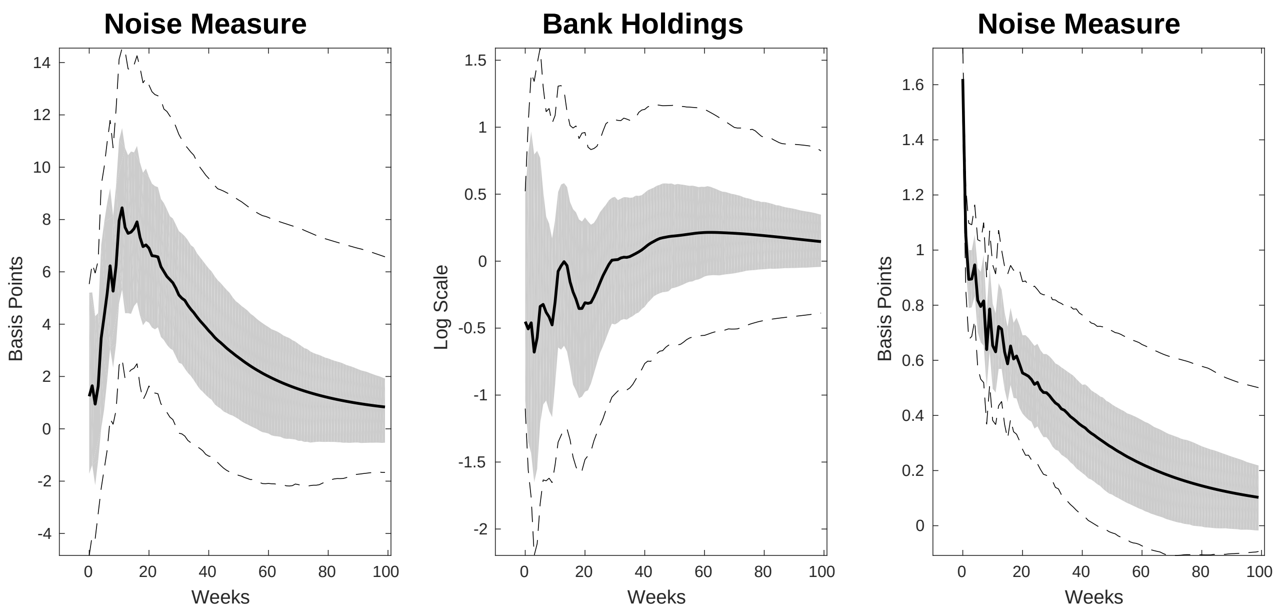

Note: The figures plot the response of endogenous variables to an idiosyncratic shock to the noise in IG bonds. The mean impulse response is shown in black. The shaded area marks a pointwise 68-percent credible interval around the median. The dashed lines mark a pointwise 95-percent credible interval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