

Boyarchenko, Nina; Giannone, Domenico; Shachar, Or

Working Paper
Flighty liquidity

Staff Report, No. 870

Provided in Cooperation with:
Federal Reserve Bank of New York

Suggested Citation: Boyarchenko, Nina; Giannone, Domenico; Shachar, Or (2018) : Flighty liquidity, Staff Report, No. 870, Federal Reserve Bank of New York, New York, NY

This Version is available at:
<https://hdl.handle.net/10419/210722>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Federal Reserve Bank of New York
Staff Reports

Flighty Liquidity

Nina Boyarchenko
Domenico Giannone
Or Shachar

Staff Report No. 870
October 2018



This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York, the Federal Reserve System, or the Financial Industry Regulatory Authority. Any errors or omissions are the responsibility of the authors.

Flighty Liquidity

Nina Boyarchenko, Domenico Giannone, and Or Shachar

Federal Reserve Bank of New York Staff Reports, no. 870

October 2018

JEL classification: C22, G12, G17

Abstract

We study the conditional distribution of future liquidity in the secondary market for corporate bonds as a function of current liquidity. Increases in liquidity are persistent for investment-grade bonds and flighty for high-yield bonds. Greater liquidity of high-yield bonds is associated with lower uncertainty about future liquidity of investment-grade bonds, but greater liquidity of investment-grade bonds is associated with greater uncertainty about future liquidity of high-yield bonds. Finally, we show that measures of market-wide volatility and market-maker constraints do not contain information useful for predicting the distribution of future liquidity over and above that contained in the recent history of bid-ask spreads.

Key words: corporate bond liquidity, liquidity uncertainty, quantile regressions

Boyarchenko, Giannone, Shachar: Federal Reserve Bank of New York (emails: nina.boyarchenko@ny.frb.org, domenico.giannone@ny.frb.org, or or.shachar@ny.frb.org). The authors thank Patrick Adams, Daniel Roberts, and Francisco Ruela for excellent research assistance. They also thank FINRA for providing corporate bond data. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York, the Federal Reserve System, or the Financial Industry Regulatory Authority.

To view the authors' disclosure statements, visit https://www.newyorkfed.org/research/staff_reports/sr870.html.

1 Introduction

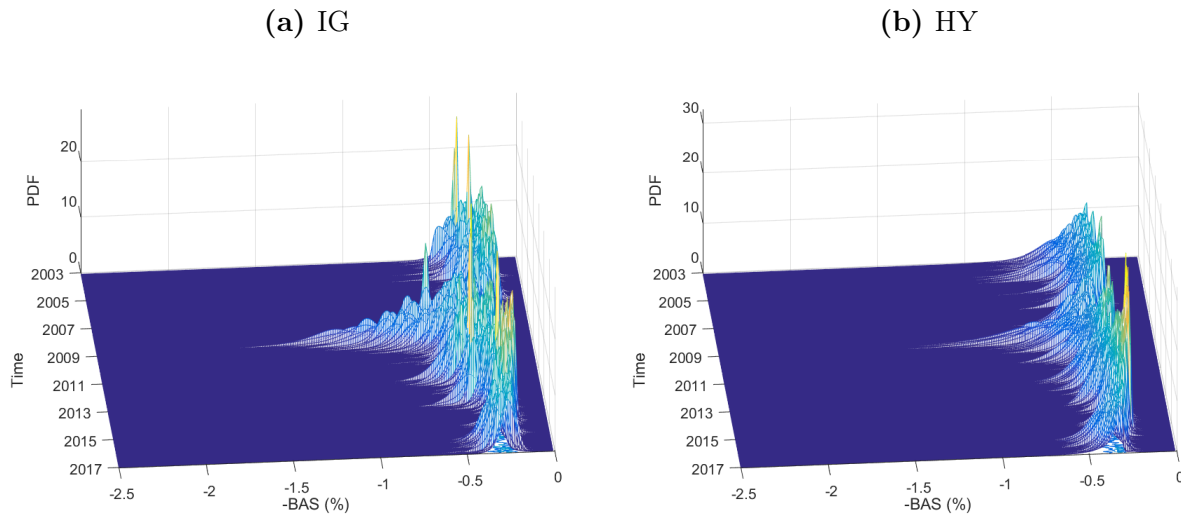
“To date, observed changes in liquidity do not suggest that shifts in liquidity are having a notable effect on the cost of trading. Nonetheless, the potential for liquidity to evaporate in times of stress deserves careful scrutiny—along with broader risks to financial stability associated with changes in markets.”

(Former Vice Chairman Stanley Fischer, November 15, 2016, Is There a Liquidity Problem Post-Crisis?)

Traditional measures of liquidity focus on the cost, either in terms of money or in terms of time, of buying or selling an asset given contemporaneous market conditions. By their nature, such measures capture the contemporaneous state of market liquidity but remain silent on whether market liquidity will change in the future. In this paper, we model empirically the full distribution of future corporate bond market liquidity at the credit rating level as a function of current liquidity and recent history of liquidity.

We estimate the distribution using quantile regressions and smooth the estimated quantile distribution every week by interpolating between the estimated quantiles using the skewed- t distribution. This allows us to transform the empirical quantile distribution into an estimated conditional distribution of credit-rating-level liquidity, plotted in Figure 1. Two features are striking about the estimated distributions. First, the entire distribution evolves over time, with both the mean and the higher moments fluctuating with aggregate market conditions. Second, although on average investment grade bonds are more liquid than high yield bonds, average future liquidity is lower and uncertainty about future liquidity is higher for investment grade than high yield bonds during periods of market stress, such as the recent financial crisis. That is, during the financial crisis, the downside risk to liquidity of investment grade bonds was greater than the downside risk to liquidity of high yield bonds, potentially reflecting market expectation of future downgrades of investment grade bonds.

Figure 1. Distribution of liquidity over time. This figure plots the time series of one-week-ahead predictive distribution of volume-weighted average (negative) bid-ask spread by credit rating category, based on quantile regressions with current bid-ask spreads for both investment grade and high yield bonds as conditioning variables.



We estimate a quantile vector autoregression in the spirit of Koenker and Xiao (2006), allowing the predictive distribution of liquidity for a given credit rating to depend on current and lagged bid-ask spreads for both credit ratings. We find that the liquidity of investment grade bonds is highly persistent across all quantiles while the persistence of high yield bonds is asymmetric. In the left tail, the liquidity of high yield bonds is as persistent as the liquidity of investment grade bonds but becomes increasingly more mean reverting when moving to the right tail of the distribution: downside liquidity risk tends to be persistent while the probability of high liquidity mean reverts quickly.

We argue that these relationships are a robust and stable feature of the data, and, thus, that our approach can be used to monitor potential risks to liquidity in real-time. We begin by showing that out-of-sample estimates of the conditional distributions of the future bid-ask spreads are very similar to the in-sample distributions. We further document our strong out-of-sample performance by analyzing predictive scores and probability integral transforms. We show that the conditional distribution is well-calibrated and performs better

out-of-sample than the unconditional distribution for both investment grade and high yield bonds. This suggests that the recent history of bid-ask spreads for both credit ratings robustly reflects information relevant for the future evolution of liquidity.

We illustrate how our methodology could be used to monitor market responses to unanticipated shocks, such as the Taper Tantrum in June 2013, as well as to evaluate the evolution of the conditional distribution around anticipated shocks, such as the liquidation of the Third Avenue Focused Credit Fund¹ in December 2015. We show that the conditional distribution can be informative about the duration of the market response to these shocks and often interprets realized shocks as being less further in the left tail of the distribution than the unconditional distribution does.

We show that including proxies of demand-side and supply-side funding liquidity pressures in the market do not lead to consistent improvements in the in-sample accuracy of the predicted distribution and, in most cases, lead to decreases in the out-of-sample accuracy of the predictive distribution. That is, market-wide proxies for uncertainty, risk premia, overall financial conditions, and measures of dealer activity in the corporate bond market do not contain information useful for predicting future bond market liquidity above what is contained in the recent history of bid-ask spreads for both credit ratings.

Finally, we summarize the downside and upside risks to the median liquidity forecast using two metrics: (1) the upside and downside entropy of the unconditional distribution of bid-ask spreads relative to the empirical conditional distribution; (2) the five percent expected shortfall and its upper tail counterpart, the five percent expected longrise. While downside relative entropy captures the conditional risks of liquidity deteriorating in excess of the downside risks predicted by the unconditional distribution, expected shortfall measures the average level of liquidity conditional on the bottom five percent tail outcomes realizing. Thus, downside relative entropy and the expected shortfall measure two complementary but

¹Although the precise timing of the liquidation was unknown, market participants had anticipated distress in funds with a focused investment in high yield bonds following distress in the oil production sector earlier in 2015.

distinct features of upside risk to liquidity. Downside relative entropy captures the probability of a negative liquidity shock occurring, relative to the probability predicted by the unconditional distribution, while expected shortfall captures the expected extreme effects of a negative liquidity shock. Similarly, upside relative entropy captures the probability of a positive liquidity shock occurring, relative to the probability predicted by the unconditional distribution; the expected longrise captures the expected extreme effects of a positive liquidity shock.

Our paper contributes to the large literature on liquidity of the corporate bond market. Corporate bonds used to be traded in an opaque environment where transaction prices were not made public. In July 2002, the Transaction Reporting and Compliance Engine (TRACE) was introduced, requiring trades in publicly issued corporate bonds to be reported to the National Association of Security Dealers, which in turn made transaction data available with some delay to the public. This was a major evolution in the corporate bond market. The impact of transparency on liquidity and on dealers' propensity to provide liquidity have been debated, but most academic papers find that the implementation of TRACE benefited clients overall, lowering transaction costs (Bessembinder et al., 2006; Goldstein et al., 2007; Edwards et al., 2007; Asquith et al., 2013). Asquith et al. (2013), however, find that market activity, as measured by trading volume divided by issue size, declined significantly for high yield bonds. Bessembinder and Maxwell (2008) provide an overview of the impact of the increase in transparency on the market.

The 2007-2009 financial crisis highlighted the need to better understand corporate bond market liquidity. Friewald et al. (2012) document that liquidity explains about one third of the variation in the aggregate market corporate yield spread in the time-series, and about half during the crisis. Direct measures of trading activity, such as trade volume, and other commonly-used liquidity measures, do not show significant explanatory power. In the cross-section, they find that the overall liquidity of bonds issued by financial firms is higher on average, than those of industrial firms. Dick-Nielsen et al. (2012) document that liquidity

deteriorated for both investment grade and high yield bonds, but it was slow and persistent for the first and short-lived for the latter. Moreover, they find consistent evidence with flight-to-quality only for AAA-rated bonds. Bao et al. (2011) calculate the *Roll* liquidity measure at the bond-level and then aggregate the liquidity measure across individual bonds. Using the aggregate measure they find that the aggregate illiquidity doubled relative to its pre-crisis average when the credit problem first broke out in August 2007, and subsequently tripled in March 2008 when Bear Stearns collapsed. Their measure peaks in October 2008, after Lehman's default and the bailout of AIG, and slowly declines thereafter. Adrian et al. (2017) show that the relationship between bond-level liquidity and dealer-level constraints changes with the introduction of post-crisis regulation, with bonds traded by more levered institutions and institutions with investment bank like characteristics less liquid after the financial crisis.

Our paper deviates from the prior literature in focusing on predicting the future evolution of liquidity, forecasting both the expected future liquidity and the downside risks to liquidity. In this aspect, our paper is related to prior literature that has investigated the time series properties of liquidity in other markets. Chordia et al. (2004) estimate a vector autoregression for stock and Treasury bond liquidity, and find cross-market dynamics from volatility to liquidity in both markets. More recently, Nagel (2012) argues that market liquidity declines during the financial crisis is partially explained by demands for higher expected returns by liquidity providers. Similarly, Comerton-Forde et al. (2010) show that when the revenues of NYSE specialists are low, liquidity on the NYSE is low as well. Relatedly, Baele et al. (2018) find increases in the VIX and the TED spread are associated with decreases in Treasury bond liquidity.

The relationship between VIX and liquidity is further investigated by Chung and Chuwonganant (2014), who show that VIX has a pervasive impact on liquidity. Karolyi et al. (2012) examine additional proxies for demand-side and supply-side pressures and find that liquidity in several countries varies across time because of demand-side reasons and not with

proxies for funding liquidity. In contrast, Karnaukh et al. (2015) document that FX liquidity declines with both funding constraints and global risk, with stronger comovement of FX liquidity when funding is constrained, global volatility is high, and FX speculators incur losses. Similarly, Mancini et al. (2013) find commonality across liquidity measures for FX, U. S. stock, U. S. Treasury, and U. S. corporate bond markets. While our paper also finds common variation in liquidity across investment grade and high yield bonds, unlike this prior literature, we find that market-wide volatility and proxies for market-maker constraints do not help predict future liquidity once we control for the recent liquidity of both credit rating categories.

From a technical perspective, our paper contributes to the growing literature that has uncovered interesting patterns by analyzing the entire predictive density. We use the methodology from Adrian et al. (Forthcoming), who find that financial conditions are an important driver of macroeconomic vulnerabilities, measured as downside risk of GDP growth. Relatedly, Smith and Vahey (2016) show substantial asymmetries in that the forecast densities of GDP growth and inflation during the great recession. In financial markets, Ghysels (2014) documents that there are substantial and time-varying asymmetries of the predictive distribution of returns. Similarly, Schmidt and Zhu (2016) show that, while the tails of the predictive distribution of stock returns vary over time, the median of the distribution is essentially time invariant. Using the quantile regression approach of Adrian et al. (Forthcoming), Crump et al. (2018) find that current realized volatility of stock returns has strong predictive content for the uncertainty of future returns and, thus, for the overall future distribution of market returns. Our paper is complementary to this prior literature as it studies the entire predictive density in a novel setting, the credit-rating-level liquidity of corporate bonds.

The rest of the paper is organized as follows. Section 2 describes the construction of measures of liquidity studied in this paper. Section 3 lays out the empirical methodology, and documents the basic features of the conditional distribution of illiquidity. We present the out-

of-sample evidence, and illustrate real-time monitoring applications using two event studies in Section 4. Section 5 investigates the information contained in alternative explanatory variables. We construct measures of flightiness of liquidity in Section 6. Section 7 concludes.

2 Data Description and Sample Construction

2.1 Corporate bond market liquidity

We use corporate bond transaction data from a supervisory version of TRACE, which contains the uncapped trade size, price, buyer and seller identities. FINRA members are identified by a designated Market Participant Identifier, MPID, and non-FINRA members are identified either as C (for client), or as A (for a non-member affiliate). Our trades dataset spans from July 2002, when TRACE was introduced, to December 2017. Real-time, public dissemination of trades was staggered, and its full implementation was completed on February 7, 2005, when all U. S. corporate bonds, except the TRACE-eligible Rule 144A bonds, were subject to dissemination. Therefore, we limit our sample to start on January 2005. We address the data issues in TRACE and clean the data as described in Adrian et al. (2017).

Using the traded prices in TRACE, we calculate the weekly effective bid-ask spread at the bond-level. The effective bid-ask is the difference between the dollar weighted average price of the buy trades and the dollar weighted average price of the sell trades (see Hong and Warga 2000 and Chakravarty and Sarkar 2003):

$$\text{BAS}_{b,t} = \sum_{n=1}^N P_n^B W_n^B - \sum_{m=1}^M P_m^S W_m^S.$$

The measure is calculated using only client-dealer trades, and requires at least one client buy trade and one client sell trade each day.

We merge the weekly measure of bond-level liquidity with Mergent FISD to get the characteristics of the bonds. We exclude bonds denominated in foreign currency, which are

agency backed, or issued as private placements, unit deals, perpetual, and preferred. We also drop bonds with a maturity of less than one year, and unrated bonds. We exclude trades of bonds 30 days prior to default, and, if the bond is reinstated, then we exclude the first 30 days after it was reinstated.

Using the credit rating information from Mergent FISD, we construct aggregate liquidity measures for the portfolio of AAA-rated bonds, investment-grade (excluding AAA) rated bonds, and high yield rated bonds as the gross-trading-volume-weighted average of bid-ask spreads for the corporate bonds with the corresponding trading volume. Figure 2a plots the time series of bid-ask spreads for these portfolios. Three features are worth noting about these time series. First, bid-ask spreads increase dramatically during periods of market stress, such as the 2007-2009 financial crisis. Second, during these stress periods, bonds with higher credit ratings have higher bid-ask spreads than bonds with lower credit ratings, suggesting that the market anticipates the eventual downgrade of these bonds. Finally, after August 2011, the bid-ask spread for the AAA category is extremely volatile. This is due to the fact that, after August 2011, very few corporate bonds actually have AAA credit rating. Indeed, Figure 2b shows that, after August 2011, the fraction of gross trading volume accounted for by trades in AAA-rated bonds drops dramatically. Because of this dramatic decrease in trading volume, we exclude AAA bonds from the results reported in the main body of the paper.

Additionally, the bid-ask spread series for high-yield bonds exhibits a year-end seasonality when trading in the corporate bond market is thin. We correct for this seasonality by regressing the bid-ask spread of high-yield bonds on a year-end indicator, and work with seasonality-adjusted bid-ask spreads for the rest of our analysis.

2.2 Market-wide variables

In some empirical specifications, we control also for market-wide proxies of liquidity demand and supply. We proxy for liquidity demand using measures of option-implied equity volatility

(VIX), Treasury volatility (MOVE 1M), and interest rate swap volatility (SMOVE 1M), as well as the Baa-Aaa spread (which proxies for credit risk premia), the Treasury slope (difference between yields on a 10 year and a 3 month Treasury, which proxies for term premia), and the Chicago Fed National Financial Conditions Index (NFCI, which proxies for economy-wide financial conditions). On the supply-side, we use data from FR 2004 on corporate securities transactions and repo market activity by primary dealers, as well as delivery fails into corporate securities borrowing agreements. While the first two measures proxy for funding liquidity in the corporate bond market as they capture the willingness and ability of the traditional market makers to trade in and provide financing against corporate securities, the third measure captures the scarcity of desirable bonds. For VIX, MOVE 1M, SMOVE 1M, Baa-Aaa spread and the Treasury slope, we aggregate the daily market prices into weekly measures by averaging within the week. The rest of the variables are available at a weekly frequency only.

3 Empirical Methodology

In this section, we describe how we apply the methodology in Adrian et al. (Forthcoming) to construct conditional distributions of corporate bond market liquidity (rather than of real GDP growth). We refer the interested reader to Adrian et al. (Forthcoming) for more details on the quantile-regression methodology itself.

3.1 Conditional inverse CDF

We begin by characterizing the relationship between future bid-ask spreads and current bid-ask spreads using quantile regressions. In particular, let $y_{i,t+h}$ be the log bid-ask spread for portfolio i in future week $t + h$, and denote by x_t the vector of conditioning variables, including a constant. In a quantile regression of $y_{i,t+h}$ on x_t , the regression slope $\beta_{i,\tau,h}$ is

chosen to minimize the quantile-weighted absolute value of prediction errors

$$\hat{\beta}_{i,\tau,h} = \operatorname{argmin}_{\beta_{i,\tau,h} \in \mathbb{R}^k} \sum_{t=1}^{T-h} \left(\tau \cdot \mathbb{1}_{(y_{i,t+h} \geq x_t \beta_{i,\tau,h})} + (1 - \tau) \cdot \mathbb{1}_{(y_{i,t+h} < x_t \beta_{i,\tau,h})} \right) |y_{i,t+h} - x_t \beta_{i,\tau,h}|, \quad (1)$$

where $\mathbb{1}_{(\cdot)}$ denotes the indicator function. Unlike ordinary least squares, which predicts the average realization of $y_{i,t+h}$ conditional on x_t , the predicted value from the regression above is the quantile of $y_{i,t+h}$ conditional on x_t

$$\hat{Q}_{y_{i,t+h}|x_t}(\tau|x_t) = x_t \hat{\beta}_{i,t+h,\tau}.$$

To reduce the influence of outliers in bid-ask spreads on the estimated coefficients, we estimate the quantile regression (1) for the natural logarithm of the bid-ask spread for a particular portfolio. We include four lags of bid-ask spreads in our regressions to capture the dependence on the whole pattern of liquidity over the previous month. That is, we parametrize the quantile function of the *negative* log bid-ask spread of portfolio i in week t , $y_{i,t}$, as

$$Q_{y_{i,t+h}|x_t}(\tau|x_t) = \alpha_{i,h,\tau} + \sum_{l=1}^4 \varphi_{i,l,h,\tau} y_{i,t-l+1} + \sum_{l=1}^4 \gamma_{i,l,h,\tau} y_{-i,t-l+1} + \epsilon_{i,h,t,\tau}, \quad (2)$$

where $y_{-i,t}$ is the realization of the negative log bid-ask spread for the other portfolio. We focus on the negative logarithm of the bid-ask spread to have a measure of liquidity: higher bid-ask spreads correspond to higher *illiquidity* of the bond, while higher negative (log) bid-ask spreads correspond to higher *liquidity* of the bond. Including the lagged bid-ask spreads of both portfolios into the specification (2) allows us to study the differential persistence of bid-ask spreads at various quantiles (through the coefficients $\{\varphi_{i,l,h,\tau}\}$), as well as the differential correlation of bid-ask spreads across credit ratings at various quantiles (through the coefficients $\{\gamma_{i,l,h,\tau}\}$).

In the following, we report the *cumulative* effect of a change in either own or other credit

rating log bid-ask spreads on the quantile function. That is, when we report regression coefficients, we are reporting

$$\varphi_{i,h,\tau} \equiv \sum_{l=1}^4 \varphi_{i,l,h,\tau}, \quad \gamma_{i,h,\tau} \equiv \sum_{l=1}^4 \gamma_{i,l,h,\tau},$$

respectively.

Figure 3 shows the scatter plot of one-week-ahead negative log bid-ask spreads for investment grade and high yield bonds against the current realization of negative log bid-ask spreads for investment grade and high yield bonds, as well as the univariate quantile regression lines for the fifth, fiftieth and ninety-fifth quantiles and the OLS regression line. Consider first the relationship between future bid-ask spreads and own current bid-ask spreads. For investment grade bonds (Figure 3a), the slopes of the three quantile regression lines are similar to each other and, moreover, similar to the linear regression slope, suggesting a linear relationship between current and future investment grade bid-ask spreads. Instead, for high yield bonds (Figure 3d), the slope of the ninety-fifth percentile is noticeably different from the slopes of the other two quantile regression lines and the OLS regression line, suggesting that bid-ask spreads on high yield bonds have different persistence across different quantiles. Turning next to the cross-credit-rating relationship between future and current bid-ask spreads, we can see that there is a non-linear relationship between one-week-ahead bid-ask spreads on investment grade bonds and current bid-ask spreads on high yield bonds (Figure 3c), but a potentially linear relationship between one-week-ahead bid-ask spreads on high yield bonds and current bid-ask spreads on investment grade bonds (Figure 3b).

We test formally the marginal effects of including the history of bid-ask spreads for both credit ratings in a multivariate regression setting in Figure 4.² Consider first the

²The confidence bounds plotted in Figure 4 are the 95 percent confidence bounds for the null hypothesis that the true data-generating process is a flexible and general linear model for liquidity. In particular, we estimate a vector autoregression (VAR) with four lags, Gaussian innovations, and a constant using the full-sample evolution of log bid-ask spreads, and bootstrap 1000 samples to compute bounds at different confidence levels for the OLS relationship. Quantile coefficient estimates that fall outside of this confidence bound thus indicate that the relationship between log bid-ask spreads and the predictive variable is non-linear.

estimated coefficients from the quantile regression of one-week-ahead negative log bid-ask spreads for investment grade bonds, plotted in the left column of Figure 4. Bid-ask spreads for investment-grade bonds are extremely persistent, with the estimated autoregressive coefficient of around 0.9. This persistence is mostly flat across quantiles but increases slightly for the right-most quantiles (most liquid) and decreases slightly for the left-most quantiles (least liquid). Turning next to the loading on current bid-ask spreads for the high yield bonds, we see that there is a positive relationship between future bid-ask spreads on investment grade bonds and current high yield bid-ask spreads in the left tail of the bid-ask spread distribution. We also observe a negative relationship between the right tail of the future bid-ask spreads on investment grade bonds and current high yield bid-ask spreads. That is, when high yield bonds are relatively more liquid, both downside and upside risks to liquidity of investment grade bonds are lower and the distribution is more concentrated around the mean.

The right column of Figure 4 plots the estimated coefficients from the quantile regression of one-week-ahead negative log bid-ask spreads for high yield bonds. Liquidity of high yield bonds is much less persistent than the liquidity of investment grade bonds, with the estimated autoregressive coefficient at the median of around 0.7. In addition, persistence increases for the leftmost quantiles and decreases for the rightmost quantiles for high yield rated bonds, with both of these extremes different from the median estimate at the 5 percent confidence level. Taken together, the results for the autoregressive coefficients across quantiles and credit ratings suggest that liquidity is more likely to evaporate for high yield bonds than for investment grade bonds; that the risks of illiquidity of high yield bonds tend to be persistent; and that upside risks to liquidity of investment grade bonds tend to be persistent but reverse quickly for high yield rated bonds. Finally, turning to the loading on current bid-ask spreads for the investment grade bonds, we see that higher liquidity of investment grade bonds is associated with larger upside and downside risks to the liquidity of high yield bonds, so that the future liquidity of high yield bonds is more uncertain when investment grade liquidity

is higher. Figure A.1 in the Appendix shows that these patterns also hold for the estimated coefficients for the four-weeks-ahead distribution.

Turning to the implications of these relationships for the dynamic evolution of risks to liquidity, Figure 5 shows realized liquidity together with the conditional median and the conditional 5th, 25th, 75th, and 95th percentile quantiles of the one-week-ahead and four-weeks-ahead predicted distribution across credit rating categories.³ This figure demonstrates one of the key results of the paper: while the distribution around the median for investment-grade bonds is largely symmetric, there is significant asymmetry between the upper and lower conditional quantiles for high yield bonds. That is, for high yield bonds, the lower quantiles vary significantly over time but the upper quantiles are stable. Figure 6 shows that, across credit ratings, the median and the interquartile range are strongly negatively correlated, i.e., deteriorations in aggregate liquidity are associated with a decrease in median liquidity for both credit rating categories and an increase in the uncertainty around the median. Thus, the left tail of the distribution shifts to the left: the fifth quantile has a negative relationship with the interquartile range. Figure 6 shows these decreases in median liquidity and increases in the uncertainty around the median persist at the four-weeks-ahead horizon.

3.2 Conditional distribution

The quantile regression (1) provides us with estimates of the quantile function, a representation of the inverse cumulative distribution function (ICDF). Prior literature has struggled with inverting the empirical ICDF produced from quantile regressions to obtain a conditional probability distribution function. Instead, we follow Adrian et al. (Forthcoming) and smooth the quantile distribution function using the skewed- t distribution developed by Azzalini and

³We transform the conditional distribution for log bid-ask spread from the quantile regression (2) to the conditional distribution for the negative bid-ask spread using the change of variables formula for distributions.

Capitanio (2003):⁴

$$f(y; \mu, \sigma, \alpha, \nu) = \frac{2}{\sigma} t\left(\frac{y - \mu}{\sigma}; \nu\right) T\left(\alpha \frac{y - \mu}{\sigma} \sqrt{\frac{\nu + 1}{\nu + \left(\frac{y - \mu}{\sigma}\right)^2}}; \nu + 1\right) \quad (3)$$

where $t(\cdot)$ and $T(\cdot)$ respectively denote the PDF and CDF of the Student t -distribution. The four parameters of the distribution pin down the location μ , scale σ , fatness ν , and shape α . Relative to the t -distribution, the skewed t -distribution adds the shape parameter which regulates the skewing effect of the CDF on the PDF. The skewed t -distribution is part of a general class of mixed distributions proposed by Azzalini (1985) and further developed by Azzalini and Dalla Valle (1996). The intuition for the derivation is that a base probability distribution – in this case $t\left(\frac{y - \mu}{\sigma}; \nu\right)$ – gets shaped by its cumulative distribution function, and rescaled by a shape parameter α . The notable special case is the traditional t -distribution when $\alpha = 0$. In the case of both $\alpha = 0$ and $\nu = \infty$, the distribution reduces to a Gaussian with mean μ and standard deviation σ . When $\nu = \infty$ and $\alpha \neq 0$, the distribution is a skewed normal.

Besides its flexibility, an advantage of using the skewed- t distribution is that it has closed-form expressions for both the PDF and the ICDF. This allows us to fit the skewed- t distribution f in week t by minimizing the distance between the estimated quantile function $\hat{Q}_{y_{i,t+h}|x_t}(\tau)$ and the ICDF $F^{-1}(\tau; \mu_{i,t,h}, \sigma_{i,t,h}, \alpha_{i,t,h}, \nu_{i,t,h})$ of the skewed- t distribution. More specifically, for each week and each credit rating, we choose the four parameters $\{\mu_{i,t,h}, \sigma_{i,t,h}, \alpha_{i,t,h}, \nu_{i,t,h}\}$ to match the fifth, twenty-fifth, fiftieth, seventy-fifth and ninety-fifth percent conditional quantiles

$$\{\hat{\mu}_{i,t,h}, \hat{\sigma}_{i,t,h}, \hat{\alpha}_{i,t,h}, \hat{\nu}_{i,t,h}\} = \underset{\mu, \sigma, \alpha, \nu}{\operatorname{argmin}} \sum_{\tau} \left(\hat{Q}_{y_{i,t+h}|x_t}(\tau) - F^{-1}(\tau; \mu_{i,t,h}, \sigma_{i,t,h}, \alpha_{i,t,h}, \nu_{i,t,h}) \right)^2,$$

⁴An alternative approach to smoothing the quantile densities is to interpolate the quantile function using splines. Imposing monotonicity and smoothness requires additional modeling choices, as in for example Schmidt and Zhu (2016).

where $\tau \in \{0.05, 0.25, 0.75, 0.95\}$, $\hat{\mu}_{i,t,h} \in \mathbb{R}$, $\hat{\sigma}_{i,t,h} \in \mathbb{R}^+$, $\hat{\alpha}_{i,t,h} \in \mathbb{R}$, and $\hat{\nu}_{i,t,h} \in \mathbb{N}$.⁵ This can be viewed as an exactly identified nonlinear cross-sectional regression of the predicted quantiles on the quantiles of the skewed- t distribution.⁶

Figure 7 plots the estimated conditional quantile distributions $\hat{Q}_{y_{i,t+h}|x_t}(\tau)$ and two versions of the fitted inverse cumulative skewed- t distribution $F^{-1}(\tau; \mu_{i,t,h}, \sigma_{i,t,h}, \alpha_{i,t,h}, \nu_{i,t,h})$ – the one conditional on bid-ask spreads and the unconditional distribution – for two sample dates at different points in the liquidity cycle: September 19, 2008, the week after the liquidation of Lehman Brothers; and January 13, 2006, which represents normal liquidity conditions. Across both credit rating categories and both dates, the skewed- t distribution is sufficiently flexible to smooth the estimated quantile function while passing through the targets of interest. Figure 7 also shows that the conditional distribution can deviate substantially from the unconditional distribution. The conditional distribution is significantly below the unconditional distribution for both credit rating categories during the height of the financial crisis. Instead, for investment grade bonds, the conditional distribution is noticeably above the unconditional distribution during normal times, and, for high yield bonds, the conditional distribution mostly coincides with the unconditional distribution during normal times. Thus, conditioning information is particularly important during episodes of low liquidity.

Figure 8 then plots the two versions of the density functions of negative bid-ask spreads for the same two dates. Comparing the conditional density across the two dates and credit ratings, we see significant variation in the conditional density across both time and credit ratings. During liquidity dry-ups, the conditional distribution has higher variance, greater negative skewness, and lower mean than the unconditional distribution. These changes are particularly pronounced for investment-grade bonds, suggesting that the risks of further liquidity deterioration are greater for higher rated bonds. This may be due to market

⁵Notice that these parameters are functions of the conditioning variables in week t .

⁶We fit the skewed- t distribution to the quantile function of log bid-ask spreads, and then use change of variables formula for distributions to convert that to the distribution for bid-ask spreads.

expectations of credit rating downgrades for such bonds.

The estimated skewed- t distribution allows us to formally test the in-sample differences in the conditional and unconditional distributions in Table 1. For both credit ratings and both forecast horizons, the conditional model significantly outperforms the unconditional model. Thus, in-sample, lagged bid-ask spreads contain information that is crucial for predicting the distribution of future liquidity outcomes.

4 Out-of-Sample Performance

The previous Section demonstrates that the predictive model which includes lagged bid-ask spreads for both credit rating categories outperforms the unconditional model in-sample. We now turn to evaluating the performance of the conditional and unconditional models out-of-sample, and illustrate the out-of-sample gains of using the conditional model using two event studies.

4.1 Statistical performance

We backtest the model by replicating the analysis that an economist would have done using the proposed methodology in real time. We produce predictive distributions recursively for two horizons (1 week and 4 weeks), starting with the estimation sample that ranges from January 1, 2003 to August 1, 2007. The first out-of-sample estimates are thus for the average liquidity in the week ending on August 8, 2007 (one-week-ahead) and the average liquidity in the week ending on September 1, 2007 (4 weeks ahead). We then iterate the same procedure, expanding the estimation sample one week at a time, until the end of the sample (December 31, 2017). At each iteration, we repeat the estimation steps above, estimating quantile regressions and matching the skewed t -distribution. The outcome of this procedure is a ten year time series of out-of-sample density forecasts for each of the two forecast horizons and each of the two credit ratings.

We perform two types of out-of-sample analyses. First, we study the robustness of our predicted distributions by comparing the in-sample predicted distributions with their real time counterparts. Second, we evaluate the out-of-sample accuracy and calibration of the density forecasts by analyzing the predictive score and the probability integral transform (PIT); that is, the predictive density and cumulative distribution evaluated at the outturn, respectively.

We begin by comparing the in-sample and out-of-sample predicted distribution, presented in Figure 9. The figure illustrates that the in-sample and out-of-sample estimates of the quantiles are virtually indistinguishable for both horizons and both credit ratings. The only case in which the in-sample and out-of-sample quantiles deviate noticeably is for the bottom fifth percentiles of liquidity during the financial crisis, with the out-of-sample more negative than the in-sample estimate. The full sample estimate incorporates the reversion of bond liquidity to more normal levels in the post-crisis period, while the real time procedure estimates a somewhat lower worst case outcomes. The similarities are more striking as the financial crisis of 2007–2009 is a significant tail event that is not in the data when estimating the out-of-sample distributions. The similarity between in-sample and out-of-sample estimates suggests that our methodology can be used to detect liquidity risks in real time.

Next, we assess the reliability of the predictive distribution using the predictive score, computed as the predictive distribution generated by a model (either the conditional or the unconditional model) and evaluated at the realized value of the time series. Higher predictive scores indicate more accurate predictions on average as higher predictive scores indicate that outcomes that the model considers more likely are closer to the ex-post realization. Figures 10a and 10c plot the time series of the scores of the conditional and unconditional one-week-ahead predictive distribution for investment grade and high yield bonds, respectively. For both investment grade and high yield bonds, the predictive score for the conditional distribution is almost always above the predictive score for the unconditional model,

indicating that the conditional model is almost always more accurate than the unconditional model. We test the predictive scores differences formally in Table 2. The conditional distribution outperforms the unconditional distribution across both horizons (one-week-ahead and four-weeks-ahead) and both credit ratings.

We conclude the out-of-sample evaluation by analyzing the calibration of the predictive distributions. We compute the empirical cumulative distribution of the PITs, which measures the percentage of observations that are below any given quantile. A model is said to be better calibrated the closer the empirical cumulative distribution of the PITs is to the 45 degree line. In a perfectly calibrated model, the cumulative distribution of the PITs is exactly the 45 degree line, so that the fraction of realizations falling below any given quantile $Q_{y_{i,t+h}|x_t}(\tau)$ of the predictive distribution is exactly equal to τ . We plot the PITs for the conditional and unconditional one-week-ahead distribution for investment grade and high yield bonds, together with the corresponding confidence bounds,⁷ in Figures 10b and 10d. For investment grade bonds, the empirical distribution of the PITs for the conditional model is well within the confidence bands across all quantiles, while the empirical distribution of the PITs for the unconditional model falls outside the confidence bands for the bottom half of the distribution. For high yield bonds, for both the condition and unconditional distributions, the empirical distribution of the PITs is well within the confidence bands for the lower quantiles, though the empirical distribution falls outside the confidence bands in the center of the distribution.

Overall, the results in Figure 10, Figure A.3 in the Appendix, and Table 2 suggest that the quantile regression approach generates robust predictive distributions, across multiple predictive horizons and across both credit ratings, and is able to capture the downside vulnerability of liquidity particularly well. We turn next to quantifying the amount of upside and downside risks present in the conditional predictive distributions of liquidity.

⁷We follow Rossi and Sekhposyan (2017) in computing the bounds. The confidence bounds should be taken as general guidance since they are derived for forecasts computed using a rolling, rather than expanding, sample. For the one-week-ahead and the four-weeks-ahead, the bands are based on critical values derived under the null of uniformity and independence of the PIT.

4.2 Event studies

The fact that the conditional model outperforms the unconditional model both in- and out-of-sample suggests that the method proposed in this paper can be used to monitor upside and downside risks to liquidity in real time. We now illustrate that the conditional distribution constructed in this paper performs well not just on average but in times of stress for the corporate bond market using two event studies: the “Taper Tantrum” on June 19, 2013 in response to Chair Bernanke’s Congressional testimony, and the liquidation of the Third Avenue Focused Credit Fund on December 11, 2015. The Taper Tantrum episode was characterized by a sell-off of longer maturity Treasuries on fears of faster-than-anticipated tapering of asset purchases by the Federal Reserve, and a resulting decline in liquidity across fixed-income markets, including the U. S. corporate bond market. The period following the liquidation of the Third Avenue Focused Credit Fund was also characterized by a decrease in liquidity of corporate bonds due to concerns about possible distress of other bond mutual funds. For both events, we compute the out-of-sample predicted distribution one- to four-weeks-ahead of the event.

Taper Tantrum. Figure 11 plots the one- and four-weeks-ahead conditional and unconditional distribution for (negative) bid-ask spreads of investment grade and high yield bonds, together with the realized bid-ask spreads the week of June 19, 2013. Two features are worth noting. First, across both forecast horizons and both credit rating categories, the conditional distribution is much more concentrated than the unconditional distribution. Thus, as with the in-sample estimates in Figures 8b and 8d, during periods of low stress for the market, the out-of-sample conditional distribution exhibits lower uncertainty than the unconditional distribution. Second, the conditional distribution assigns a much higher probability to the realized bid-ask spread in the week of June 19, 2013 than the unconditional distribution.

Third Avenue Focused Credit Fund Liquidation. Figure 12 plots the one- to four-weeks-ahead predictive distributions on December 11 and December 18, 2015 for investment grade and high yield bonds. The difference between the distributions at these two dates can be thought of as a density impulse response function in response to the liquidation of the Third Avenue Fund on December 11, 2015. For both investment grade and high yield bonds, the predictive distribution shifts down after the liquidation, reflecting decreased liquidity of the market, and the left tail of the distribution shifts out (becomes more negative), reflecting increased risk of further liquidity deterioration. Nonetheless, the realized path of the bid-ask spread lies well within the uncertainty bands predicted as of December 11, 2015. For investment grade bonds, the realized average bid-ask spread in the week ending on December 18, 2015, falls within the 10-25 percentile of the distribution as of December 11, 2015. The liquidity of investment grade bonds rebound by the week ending on January 1, 2016, falling within the second quartile of the distribution as of December 11 and the third quartile of the distribution as of December 18, 2015.

In contrast, the realized average bid-ask spread for high yield bonds in the week ending on December 18, 2015, falls within the bottom fifth percentile of the distribution as of December 11, 2015, but their liquidity rebounds within one week. The average liquidity of high yields bonds in the week ending on December 25, 2015 falls within the second quartile of the distribution as of December 11 and the third quartile of the distribution as of December 18, 2015.

5 Other Predictors

Prior literature (see e.g. Nagel, 2012; Chung and Chuwonganant, 2014) has shown that measures of market-wide volatility are correlated with measures of market liquidity. In this Section, we examine whether such proxies for demand-side pressures as well as proxies for liquidity supply contain additional information about the future distribution of corporate

bond liquidity. In particular, we augment the quantile regression specification (2) to include observations of market-wide variables z_t as predictors

$$Q_{y_{i,t+h}|x_t}(\tau|x_t) = \alpha_{i,h,\tau} + \sum_{l=1}^4 \beta_{i,l,h,\tau} y_{i,t-l+1} + \sum_{l=1}^4 \gamma_{i,l,h,\tau} y_{-i,t} + \sum_{k=1}^K \eta_{i,h} z_{k,t-1+1} + \epsilon_{i,h,t,\tau}, \quad (4)$$

where $z_{k,t}$ is the observation of the k^{th} market-wide variable in week t . The market-wide variables that we consider here are defined in Section 2.2. We compare the in-sample and out-of-sample performance of the distribution conditional on lagged bid-ask spreads and market-wide variables to the performance of the distribution conditional on lagged bid-ask spreads only by conducting log-likelihood ratio tests for both credit ratings and both horizons.

5.1 In-sample performance

Consider first the results of the in-sample log-likelihood ratio comparisons, reported in Table 3a. Each column (except for the first column, which reports the log-likelihood ratio between the unconditional and the conditional models) in Table 3a corresponds to the log-likelihood ratio between the augmented conditional distribution (4) and the baseline conditional distribution (2) for different market-wide variables, with positive numbers indicating better performance of the augmented model than the baseline model in-sample. Table 3a shows the striking result that, even in-sample, market-wide predictors do not consistently contain information about the future distribution of liquidity over and above the information contained in lagged bid-ask spreads. For investment grade bonds, the augmented model outperforms the baseline conditional model at the one week horizon only if all the dealer condition variables or the first-order principal component of the market-conditions variables (or both) are included; at the four week horizon, the augmented model outperforms the baseline model only if VIX or all of the market-conditions variables are included. For high yield bonds, instead, the one-week-ahead performance is improved if either the VIX, MOVE, Treasury slope or the first-order principal component of the market-conditions variable are

included, and the four-weeks-ahead performance is improved if either all the market conditions variables or the gross volume of primary dealer activity in corporate securities-backed repo (or both) are included. Thus, although proxies for demand-side pressures may help predict future average liquidity, augmenting the conditional model (2) with these variables does not consistently improve the ability of the model to predict the full distribution of future liquidity.

5.2 Out-of-sample performance

Turn now to the out-of-sample performance of the augmented models reported in Table 3b. For investment grade bonds, for both horizons, the augmented models consistently underperform the baseline conditional model. Strikingly, if we include either all the predictors (for both horizons) or the first principal component of the market-conditions variables and the first principal component of the dealer conditions variables (for the one week horizon), the point estimate of the log-likelihood ratio is at least two standard deviations below zero, suggesting that the corresponding augmented models underperform the baseline model in a statistically-meaningful way.

For high yield bonds, the underperformance of the augmented models is not as persistent, with the augmented models outperforming the baseline model in a statistically-meaningful way when any of the dealer conditions (or their first principal component) are included. At the four week horizon, only the model that includes the first principal component of the market-conditions variables and the first principal component of the dealer conditions variables outperforms the baseline model in a statistically-meaningful way; interestingly, the model that includes instead all the predictive variables underperforms the baseline model in a statistically-meaningful way.

Overall, the results of this Section suggest that market-wide measures do not consistently improve the predictive performance of the conditional model in-sample and, for most specifications, either do not improve or even detriment the out-of-sample performance. Thus,

market-wide measures of demand-side and supply-side pressures do not seem to contain information about future credit-rating-level liquidity beyond the information contained in the history of bid-ask spreads.

6 Measuring Liquidity Flightiness

The median of the predicted density provides the modal forecast for liquidity one week or four weeks ahead. However, as illustrated by the quote from former Vice Chairman Fisher in the introduction, policymakers are frequently concerned with downside risks to liquidity or, in other words, how likely is liquidity to evaporate. In this Section, we summarize the risks encoded in the conditional distributions of bond liquidity using two measures proposed by Adrian et al. (Forthcoming): upside and downside relative entropy and expected longrise and shortfall.

6.1 Upside and downside relative entropy

We start with upside and downside relative entropy, which measures the “extra” probability assigned by the conditional model to outcomes above and below the median of the distribution, respectively, relative to the probability assigned to the same outcomes by the unconditional distribution. Put simply, upside relative entropy measures to what extent “good” outcomes are more likely to happen under the conditional distribution than under the unconditional distribution. Similarly, downside relative entropy measures to what extent “bad” outcomes are more likely to happen under the conditional distribution than under the unconditional distribution. Formally, we denote by $\hat{g}_{y_{i,t+h}}$ the unconditional density computed by matching the unconditional empirical distribution of the log bid-ask spread on bonds with credit rating i and by $\hat{f}_{y_{i,t+h}|x_t}(y|x_t) = f(y; \hat{\mu}_{i,t+h}, \hat{\sigma}_{i,t+h}, \hat{\alpha}_{i,t+h}, \hat{\nu}_{i,t+h})$ the estimated skewed t -distribution. Then the upside, $\mathcal{L}_{i,t}^U$, and downside, $\mathcal{L}_{i,t}^D$, entropy of $\hat{g}_{y_{i,t+h}}(y_i)$

relative to $\hat{f}_{y_{i,t+h}|x_t}(y_i|x_t)$ are defined as

$$\mathcal{L}_{i,t}^D \left(\hat{f}_{y_{i,t+h}|x_t}; \hat{g}_{y_{i,t+h}} \right) = - \int_{-\infty}^{\hat{F}_{y_{i,t+h}|x_t}^{-1}(0.5|x_t)} \left(\log \hat{g}_{y_{i,t+h}}(y_i) - \log \hat{f}_{y_{i,t+h}|x_t}(y_i|x_t) \right) \hat{f}_{y_{i,t+h}|x_t}(y_i|x_t) dy_i, \quad (5)$$

$$\mathcal{L}_{i,t}^U \left(\hat{f}_{y_{i,t+h}|x_t}; \hat{g}_{y_{i,t+h}} \right) = - \int_{\hat{F}_{y_{i,t+h}|x_t}^{-1}(0.5|x_t)}^{+\infty} \left(\log \hat{g}_{y_{i,t+h}}(y_i) - \log \hat{f}_{y_{i,t+h}|x_t}(y_i|x_t) \right) \hat{f}_{y_{i,t+h}|x_t}(y_i|x_t) dy_i, \quad (6)$$

where $\hat{F}_{y_{i,t+h}|x_t}(y_i|x_t)$ is the cumulative distribution associated with $\hat{f}_{y_{i,t+h}|x_t}(y_i|x_t)$ and $\hat{F}_{y_{i,t+h}|x_t}^{-1}(0.5|x_t)$ is the conditional median.

Figures 13a and 13b illustrate the downside entropy calculation for the one-week-ahead distribution of liquidity of investment grade bonds on two dates at different points in the liquidity cycle: September 19, 2008, the week after the liquidation of Lehman Brothers; and January 13, 2006, which represents normal liquidity conditions. On September 19, 2008, the conditional distribution is much more pessimistic than the unconditional distribution so that the conditional distribution at the median is above the unconditional distribution, and downside relative entropy is the area between the conditional and unconditional distribution, shaded in grey. Thus, in periods of time when the conditional distribution is more pessimistic than the unconditional distribution, the downside relative entropy is positive. On January 13, 2006, instead, the conditional distribution is somewhat more optimistic than the unconditional distribution, and the median of the conditional distribution is above the median of the unconditional distribution. The downside relative entropy is then the area between the conditional and unconditional distribution below the conditional median and above the unconditional median (in grey) less the area between the conditional and unconditional distribution below the unconditional median (shaded in red). Thus, in periods when the conditional distribution is more optimistic than the unconditional distribution, the downside relative entropy is negative, and the upside relative entropy is positive.

Figures 14a and 14c plot the time series of one-week-ahead upside and downside relative entropy for investment grade and high yield bonds, respectively, and Figures A.4a and A.4c plot the time series of four-weeks-ahead upside and downside relative entropy. On average, upside relative entropy is more volatile than downside relative entropy for investment-grade bonds, indicating that, on average, there is more uncertainty about whether investment grade bonds become more liquid than uncertainty about whether investment-grade bonds become less liquid than predicted by the unconditional distribution. For both investment grade and high yield bonds, upside and downside entropy co-move positively during periods of market stress, so that there is greater overall uncertainty about corporate bond liquidity during market downturns. Since the financial crisis, the upside relative entropy is more volatile than the downside relative entropy for high yield bonds as well. Thus, after the financial crisis, there is more uncertainty about the upside risks to liquidity of high yield bonds than about the downside risks.

6.2 Expected longrise and shortfall

In addition to upside and downside entropy, we also study the expected shortfall encoded in the conditional distribution of liquidity together with its right tail counterpart, the expected longrise. The expected longrise measures how high the average liquidity in the conditional top fifth percentile is, while the expected shortfall captures how low the average liquidity in the bottom fifth percentile of the conditional distribution is. Formally, for a chosen target probability π , the expected shortfall and longrise are defined, respectively, as

$$SF_{i,t+h}(\pi) = \frac{1}{\pi} \int_0^\pi \hat{F}_{y_{i,t+h}|x_t}^{-1}(\tau|x_t) d\tau; \quad LR_{i,t+h}(\pi) = \frac{1}{\pi} \int_{1-\pi}^1 \hat{F}_{y_{i,t+h}|x_t}^{-1}(\tau|x_t) d\tau.$$

We illustrate the expected five percent shortfall calculation for the one-week-ahead distribution of investment grade bonds in Figures 13c and 13d. For both dates, the expected shortfall is measured as the grey area between the zero line and the conditional ICDF for

quantiles between 0 and 0.05.⁸ On September 19, 2008 (13c), that area is substantially larger than the area on January 13, 2006 (13d), reflecting the greater illiquidity of the corporate bond market during the financial crisis. Unlike the relative upside and downside entropy calculation described above, expected longrise and shortfall do not compare the conditional and unconditional distributions. Thus, the fact that the unconditional ICDF is uniformly above the conditional ICDF during periods of market stress is not reflected in the expected shortfall.

Figures 14b and 14d plot the time series of one-week-ahead expected longrise and shortfall for investment grade and high yield bonds, respectively.⁹ For both credit rating categories, expected longrise and shortfall co-move positively during the entire sample, with the expected shortfall more volatile than the expected longrise. Thus, although the probability of liquidity increases is more volatile than the probability of liquidity decreases, how adverse the bottom five percent liquidity outcomes are is more volatile than how beneficial the top five percent liquidity outcomes are.

7 Conclusion

This paper studies the predictability of liquidity and downside risk to liquidity of U. S. investment grade and high yield corporate bonds. We find evidence of liquidity spillovers across credit rating categories: greater current liquidity of high yield bonds is associated with lower uncertainty about future liquidity of investment grade bonds, while greater liquidity of investment grade bonds is associated with greater uncertainty about future liquidity of high yield bonds. We find that augmenting the baseline model with proxies for demand-side and supply-side pressures in the market does not improve the predictive performance

⁸A further measure of downside risk is “liquidity-at-risk”, or, in our notation, $\hat{F}_{y_{i,t+h}|x_t}^{-1}(\tau|x_t)$ for a given level of τ , which corresponds to, e.g., the fifth percentile worst outcome of liquidity. Expected shortfall, instead, averages across all percentiles below (and including) the target quantile and thus provides a more comprehensive metric of the severity of worst-case outcomes.

⁹Figures A.4b and A.4d plot the time series of four-weeks-ahead expected longrise and shortfall.

of the model. This suggests that, although dealer balance sheet constraints affect liquidity at the bond level, the current state of dealer balance sheets does not contain additional information about the future evolution of credit-rating-level liquidity beyond that contained in the current level of bid-ask spreads.

The global financial crisis highlighted the importance of understanding risks to liquidity for both individual institutions and the financial system as a whole. As a result, a number of jurisdictions have introduced liquidity stress tests, arguing that liquidity stress tests generate valuable information on institutions' liquidity profile beyond that captured by standardized liquidity metrics, such as the Liquidity Coverage Ratio and the Net Stable Funding Ratio.¹⁰ In this paper, we show that the predictive model that conditions on recent history of liquidity of investment grade and high yield corporate bonds performs well out-of-sample, both on average and around stress events in the market, suggesting that the model could be used to produce plausible, date- and horizon-dependent liquidity stress scenarios.

¹⁰See the overviews in e.g. BCBS (2013) and Jobst et al. (2017).

References

- ADRIAN, T., N. BOYARCHENKO, AND D. GIANNONE (Forthcoming): “Vulnerable Growth,” *American Economic Review*.
- ADRIAN, T., N. BOYARCHENKO, AND O. SHACHAR (2017): “Dealer Balance Sheets and Bond Liquidity Provision,” *Journal of Monetary Economics, Carnegie-Rochester-NYU Conference Series*, 89, 92–109.
- ASQUITH, P., T. COVERT, AND P. PATHAK (2013): “The Effects of Mandatory Transparency in Financial Market Design: Evidence from the Corporate Bond Market,” Tech. rep., National Bureau of Economic Research.
- AZZALINI, A. (1985): “A class of Distributions which Includes the Normal Ones,” *Scandinavian Journal of Statistics*, 171–178.
- AZZALINI, A. AND A. CAPITANIO (2003): “Distributions Generated by Perturbation of Symmetry with Emphasis on a Multivariate Skew t-distribution,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65, 367–389.
- AZZALINI, A. AND A. DALLA VALLE (1996): “The Multivariate Skew-normal Distribution,” *Biometrika*, 83, 715–726.
- BAELE, L., G. BEKAERT, K. INGHELBRECHT, AND M. WEI (2018): “Flights to Safety,” Working Paper No. 230, National Bank of Belgium.
- BAO, J., J. PAN, AND J. WANG (2011): “The Illiquidity of Corporate Bonds,” *Journal of Finance*, 66, 911–946.
- BCBS (2013): “Liquidity stress testing: a survey of theory, empirics and current industry and supervisory practices,” *Basel Committee on Banking Supervision*, 24.
- BESSEMBINDER, H. AND W. MAXWELL (2008): “Markets transparency and the corporate bond market,” *The Journal of Economic Perspectives*, 22, 217–234.
- BESSEMBINDER, H., W. MAXWELL, AND K. VENKATARAMAN (2006): “Market transparency, liquidity externalities, and institutional trading costs in corporate bonds,” *Journal of Financial Economics*, 82, 251–288.
- CHAKRAVARTY, S. AND A. SARKAR (2003): “Trading costs in three US bond markets,” *The Journal of Fixed Income*, 13, 39–48.
- CHORDIA, T., A. SARKAR, AND A. SUBRAHMANYAM (2004): “An empirical analysis of stock and bond market liquidity,” *The Review of Financial Studies*, 18, 85–129.
- CHUNG, K. H. AND C. CHUWONGANANT (2014): “Uncertainty, market structure, and liquidity,” *Journal of Financial Economics*, 113, 476–499.

- COMERTON-FORDE, C., T. HENDERSHOTT, C. JONES, P. MOULTON, AND M. SEASHOLES (2010): “Time Variation in Liquidity: The Role of Market-Maker Inventories and Revenues,” *Journal of Finance*, 65, 295–331.
- CRUMP, R. K., D. GIANNONE, AND S. HUNDTOFTE (2018): “Changing risk-return profiles,” Staff Report No. 850, Federal Reserve Bank of New York.
- DICK-NIELSEN, J., P. FELDHÜTTER, AND D. LANDO (2012): “Corporate Bond Liquidity before and after the Onset of the Subprime Crisis,” *Journal of Financial Economics*, 103, 471–492.
- EDWARDS, A. K., L. E. HARRIS, AND M. S. PIWOWAR (2007): “Corporate bond market transaction costs and transparency,” *The Journal of Finance*, 62, 1421–1451.
- FRIEWALD, N., R. JANKOWITSCH, AND M. 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, 18–36.
- GHYSELS, E. (2014): “Conditional Skewness with Quantile Regression Models: SoFiE Presidential Address and a Tribute to Hal White,” *Journal of Financial Econometrics*, 12, 620–644.
- GOLDSTEIN, M. A., E. S. HOTCHKISS, AND E. R. SIRRI (2007): “Transparency and liquidity: A controlled experiment on corporate bonds,” *Review of Financial Studies*, 20, 235–273.
- HONG, G. AND A. WARGA (2000): “An empirical study of bond market transactions,” *Financial Analysts Journal*, 56, 32–46.
- JOBST, A. A., L. L. ONG, AND C. SCHMIEDER (2017): “Macroprudential Liquidity Stress Testing in FSAPs for Systemically Important Financial Systems,” Working Paper 17/102, International Monetary Fund.
- KARNAUKH, N., A. RANALDO, AND P. SÖDERLIND (2015): “Understanding FX liquidity,” *The Review of Financial Studies*, 28, 3073–3108.
- KAROLYI, G. A., K.-H. LEE, AND M. A. VAN DIJK (2012): “Understanding commonality in liquidity around the world,” *Journal of Financial Economics*, 105, 82–112.
- KOENKER, R. AND Z. XIAO (2006): “Quantile autoregression,” *Journal of the American Statistical Association*, 101, 980–990.
- MANCINI, L., A. RANALDO, AND J. WRAMPPELMAYER (2013): “Liquidity in the foreign exchange market: Measurement, commonality, and risk premiums,” *The Journal of Finance*, 68, 1805–1841.
- NAGEL, S. (2012): “Evaporating liquidity,” *The Review of Financial Studies*, 25, 2005–2039.
- ROSSI, B. AND T. SEKHPOSYAN (2017): “Alternative tests for correct specification of conditional predictive densities,” Working Paper, Texas A&M University.

- SCHMIDT, L. AND Y. ZHU (2016): “Quantile Spacings: A Simple Method for the Joint Estimation of Multiple Quantiles Without Crossing,” Working Paper, University of Chicago.
- SMITH, M. AND S. VAHEY (2016): “Asymmetric Forecast Densities for US Macroeconomic Variables from a Gaussian Copula Model of Cross-Sectional and Serial Dependence,” *Journal of Business & Economic Statistics*, 34, 416–434.

Table 1: In-sample log-likelihoods. This table reports the in-sample log-likelihoods for the unconditional and conditional models for log bid-ask spreads. The conditional model includes four lags of log bid-ask spreads. Numbers reported below the diagonal correspond to the log-likelihood ratio between the unconditional and conditional model. HAC standard errors reported in parentheses below the estimates. Bolded estimates indicate point estimates at least two standard deviations away from 0.

(a) IG: 1-week ahead

	Uncond	BAS
Uncond	-0.35	
(s.e.)	(0.10)	
BAS	0.81	0.46
(s.e.)	(0.08)	(0.05)

(b) HY: 1-week ahead

	Uncond	BAS
Uncond	0.01	
(s.e.)	(0.08)	
BAS	0.51	0.52
(s.e.)	(0.07)	(0.05)

(c) IG: 4-weeks ahead

	Uncond	BAS
Uncond	-0.35	
(s.e.)	(0.10)	
BAS	0.69	0.34
(s.e.)	(0.08)	(0.06)

(d) HY: 4-weeks ahead

	Uncond	BAS
Uncond	0.01	
(s.e.)	(0.09)	
BAS	0.38	0.39
(s.e.)	(0.06)	(0.06)

Table 2: Out-of-sample log-likelihoods. This table reports the out-of-sample log-likelihoods for the unconditional and conditional models for log bid-ask spreads. The conditional model includes four lags of log bid-ask spreads. Numbers reported below the diagonal correspond to the log-likelihood ratio between the unconditional and conditional model. HAC standard errors reported in parentheses below the estimates. Bolded estimates indicate point estimates at least two standard deviations away from 0.

(a) IG: 1-week ahead

	Uncond	BAS
Uncond	-0.53	
(s.e.)	(0.22)	
BAS	1.08	0.55
(s.e.)	(0.10)	(0.05)

(b) HY: 1-week ahead

	Uncond	BAS
Uncond	-0.07	
(s.e.)	(0.12)	
BAS	0.39	0.33
(s.e.)	(0.13)	(0.11)

(c) IG: 4-weeks ahead

	Uncond	BAS
Uncond	-0.55	
(s.e.)	(0.24)	
BAS	0.91	0.36
(s.e.)	(0.10)	(0.09)

(d) HY: 4-weeks ahead

	Uncond	BAS
Uncond	-0.09	
(s.e.)	(0.13)	
BAS	0.37	0.28
(s.e.)	(0.10)	(0.07)

Table 3: Comparisons to alternative models. This table reports the in-sample and out-of-sample log-likelihood ratios between the conditional model and models that also include additional predictors. All models include four lags of log bid-ask spreads. “Slope” is the difference between the yield on a 10 year Treasury and the yield on a 3 month Treasury. In the “All Predictors” category, the “1st PCs” column includes the first principle component of the market conditions variables and the first principle component of the dealers conditions variables. HAC standard errors reported in parentheses below the estimates. Bolded estimates indicate point estimates at least two standard deviations away from 0.

(a) In-sample log-likelihood ratios

IG	Uncond.	Market conditions predictors										Dealer conditions predictors				All predictors	
		VIX	MOVE IM	SMOVE IM	BAA-AAA	Slope	NFCI	All	1st PC	Fails	Transactions	Repo	All	1st PC	All	1st PCs	
$h = 1$	-0.823 (0.077)	-0.015 (0.035)	0.021 (0.015)	0.017 (0.013)	-0.054 (0.064)	0.004 (0.010)	-0.036 (0.071)	0.053 (0.028)	0.031 (0.014)	0.009 (0.014)	0.019 (0.011)	-0.010 (0.014)	0.025 (0.012)	0.010 (0.009)	0.071 (0.027)	0.006 (0.027)	
$h = 4$	-0.701 (0.079)	0.042 (0.014)	0.001 (0.012)	-0.004 (0.010)	0.017 (0.013)	-0.002 (0.009)	0.035 (0.021)	0.072 (0.027)	0.005 (0.021)	-0.004 (0.011)	0.015 (0.013)	0.001 (0.015)	0.016 (0.014)	0.004 (0.014)	0.086 (0.026)	0.027 (0.016)	

HY	Uncond.	Market conditions predictors										Dealer conditions predictors				All predictors	
		VIX	MOVE IM	SMOVE IM	BAA-AAA	Slope	NFCI	All	1st PC	Fails	Transactions	Repo	All	1st PC	All	1st PCs	
$h = 1$	-1.515 (0.065)	0.030 (0.009)	0.032 (0.012)	0.035 (0.017)	0.037 (0.021)	0.058 (0.025)	0.045 (0.027)	0.060 (0.034)	0.054 (0.022)	0.037 (0.027)	0.033 (0.028)	0.036 (0.027)	0.053 (0.029)	0.052 (0.028)	0.051 (0.039)	0.069 (0.028)	
$h = 4$	-0.380 (0.058)	0.011 (0.012)	0.028 (0.022)	0.040 (0.021)	0.018 (0.011)	0.029 (0.019)	0.094 (0.018)	0.024 (0.024)	0.017 (0.022)	0.016 (0.018)	0.016 (0.021)	0.041 (0.020)	0.041 (0.022)	0.018 (0.025)	0.109 (0.025)	0.070 (0.020)	


(b) Out-of-sample log-likelihood ratios

IG	Uncond.	Market conditions predictors										Dealer conditions predictors				All predictors	
		VIX	MOVE IM	SMOVE IM	BAA-AAA	Slope	NFCI	All	1st PC	Fails	Transactions	Repo	All	1st PC	All	1st PCs	
$h = 1$	-1.082 (0.104)	-0.048 (0.046)	0.004 (0.013)	-0.002 (0.017)	-0.045 (0.027)	-0.016 (0.014)	-0.116 (0.093)	-0.126 (0.049)	-0.012 (0.027)	-0.091 (0.071)	-0.021 (0.016)	-0.060 (0.043)	-0.114 (0.058)	-0.084 (0.068)	-0.168 (0.059)	-0.111 (0.053)	
$h = 4$	-0.908 (0.105)	0.013 (0.019)	0.004 (0.014)	-0.072 (0.062)	-0.154 (0.108)	-0.041 (0.024)	0.038 (0.028)	-0.443 (0.206)	0.028 (0.020)	-0.142 (0.120)	-0.021 (0.025)	-0.120 (0.068)	-0.274 (0.165)	-0.076 (0.049)	-0.658 (0.214)	-0.087 (0.063)	

HY	Uncond.	Market conditions predictors										Dealer conditions predictors				All predictors	
		VIX	MOVE IM	SMOVE IM	BAA-AAA	Slope	NFCI	All	1st PC	Fails	Transactions	Repo	All	1st PC	All	1st PCs	
$h = 1$	-0.391 (0.130)	-0.003 (0.078)	0.073 (0.107)	0.106 (0.084)	0.079 (0.071)	0.158 (0.074)	0.060 (0.094)	0.050 (0.104)	0.099 (0.074)	0.147 (0.073)	0.168 (0.080)	0.171 (0.070)	0.008 (0.132)	0.201 (0.076)	-0.176 (0.143)	0.167 (0.088)	
$h = 4$	-0.374 (0.101)	0.007 (0.032)	-0.034 (0.048)	-0.011 (0.036)	-0.078 (0.071)	0.056 (0.039)	-0.076 (0.090)	-0.063 (0.072)	0.022 (0.032)	-0.032 (0.047)	0.053 (0.049)	0.086 (0.047)	-0.017 (0.065)	0.067 (0.046)	-0.287 (0.107)	0.092 (0.043)	

Figure 2. Illiquidity and Trading Over Time. This figure plots the time series of volume-weighted average bid-ask spread by credit rating category (2a), and the share of total trading volume represented by each credit rating category (2b).

(a) Bid-ask spread



(b) Relative trading volume

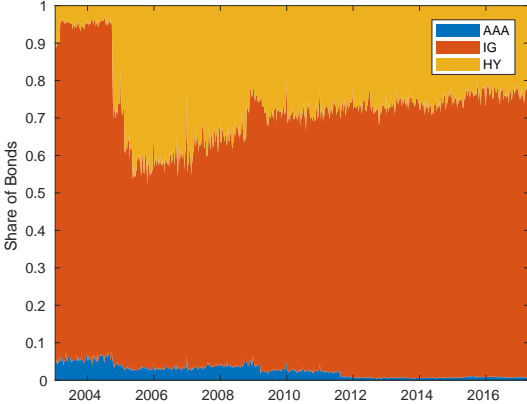
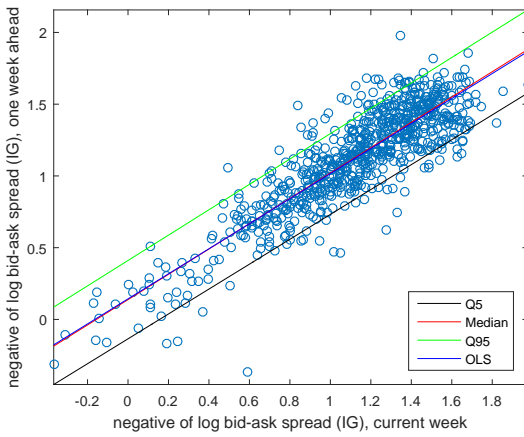



Figure 3. Quantile Regressions. This figure shows the univariate quantile regressions of one-week-ahead bid-ask spreads for investment grade and high yield bonds on current bid-ask spread for investment grade (left column) and high yield bonds (right column).

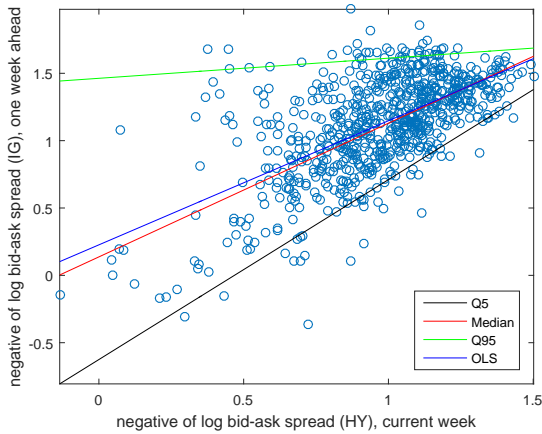
(a) IG: IG



(b) HY: IG



(c) IG: HY



(d) HY: HY




Figure 4. Estimated Quantile Regression Coefficients. This figure shows the estimated coefficients in quantile regressions of one-week-ahead bid-ask spreads for investment grade and high yield bonds on four lags of log bid-ask spreads for investment grade (left column) and high yield bonds (right column). Regression coefficients reported as the sum of the coefficients on the four lag of the respective variable. We report confidence bounds for the null hypothesis that the true data-generating process is a general, flexible linear model for bid-ask spreads (VAR with 4 lags); bounds are computed using 1000 bootstrapped samples.

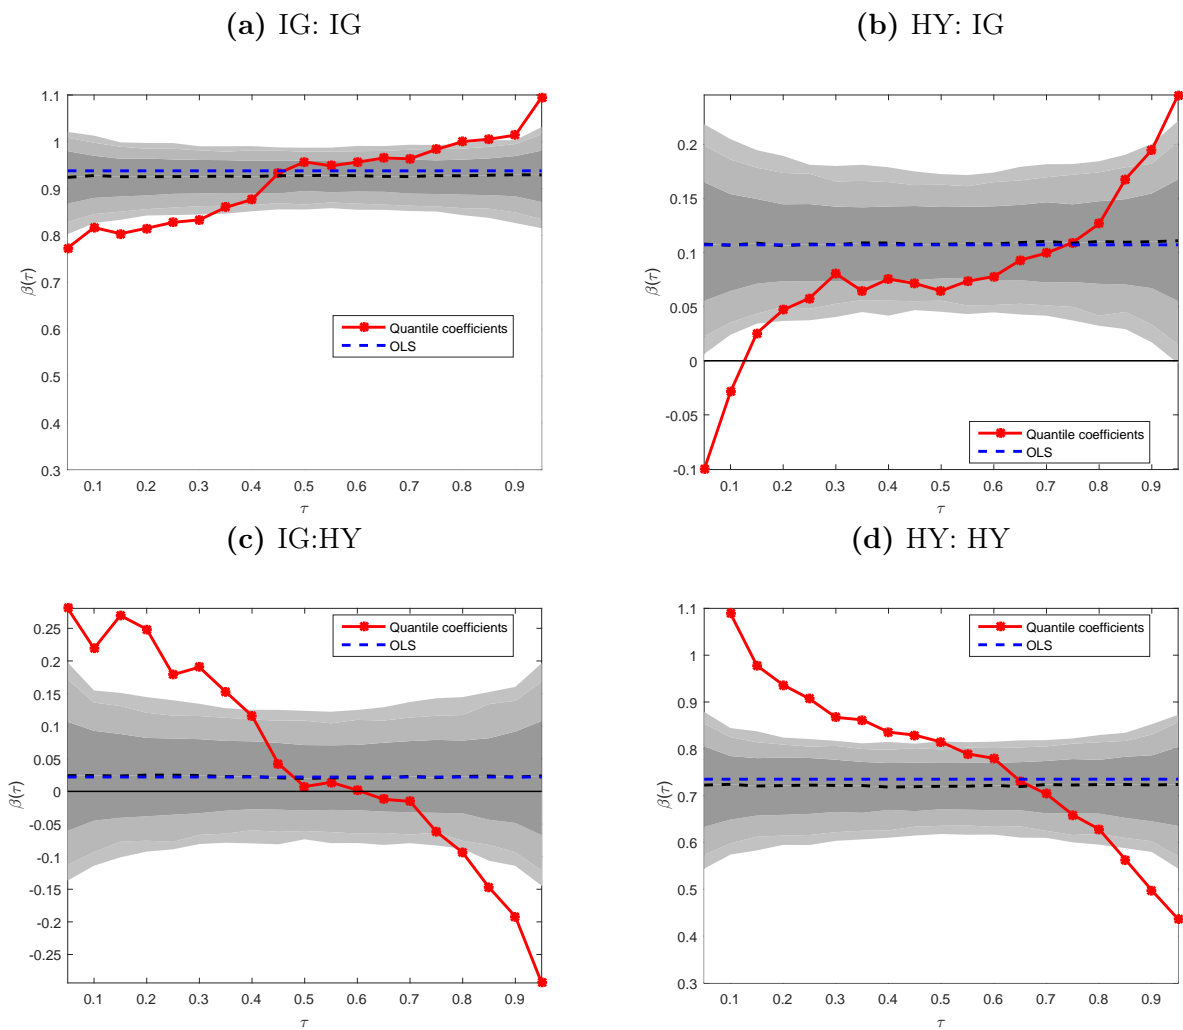




Figure 5. Predicted Distributions. This figure shows the time series evolution of the predicted distribution one-week-ahead (left column) and four weeks ahead (right column) of volume-weighted average bid-ask spread by credit rating category. Shaded areas correspond to the (5%, 95%), (10%, 90%) and (25%, 75%) interquartile ranges, respectively.

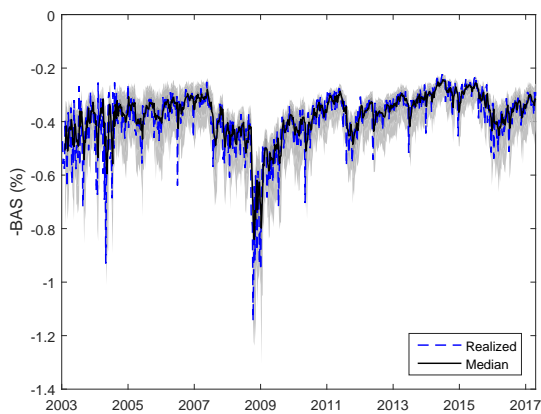
(a) IG: 1 week ahead



(b) IG: 4 weeks ahead



(c) HY: 1 week ahead



(d) HY: 4 weeks ahead




Figure 6. Median, Interquartile Range and Tail Outcomes. This figure shows relationship between the interquartile range and the median (left column) and the interquartile range and the 95th percentile (right column) of the one-week-ahead conditional distribution of volume-weighted average bid-ask spread by credit rating category.

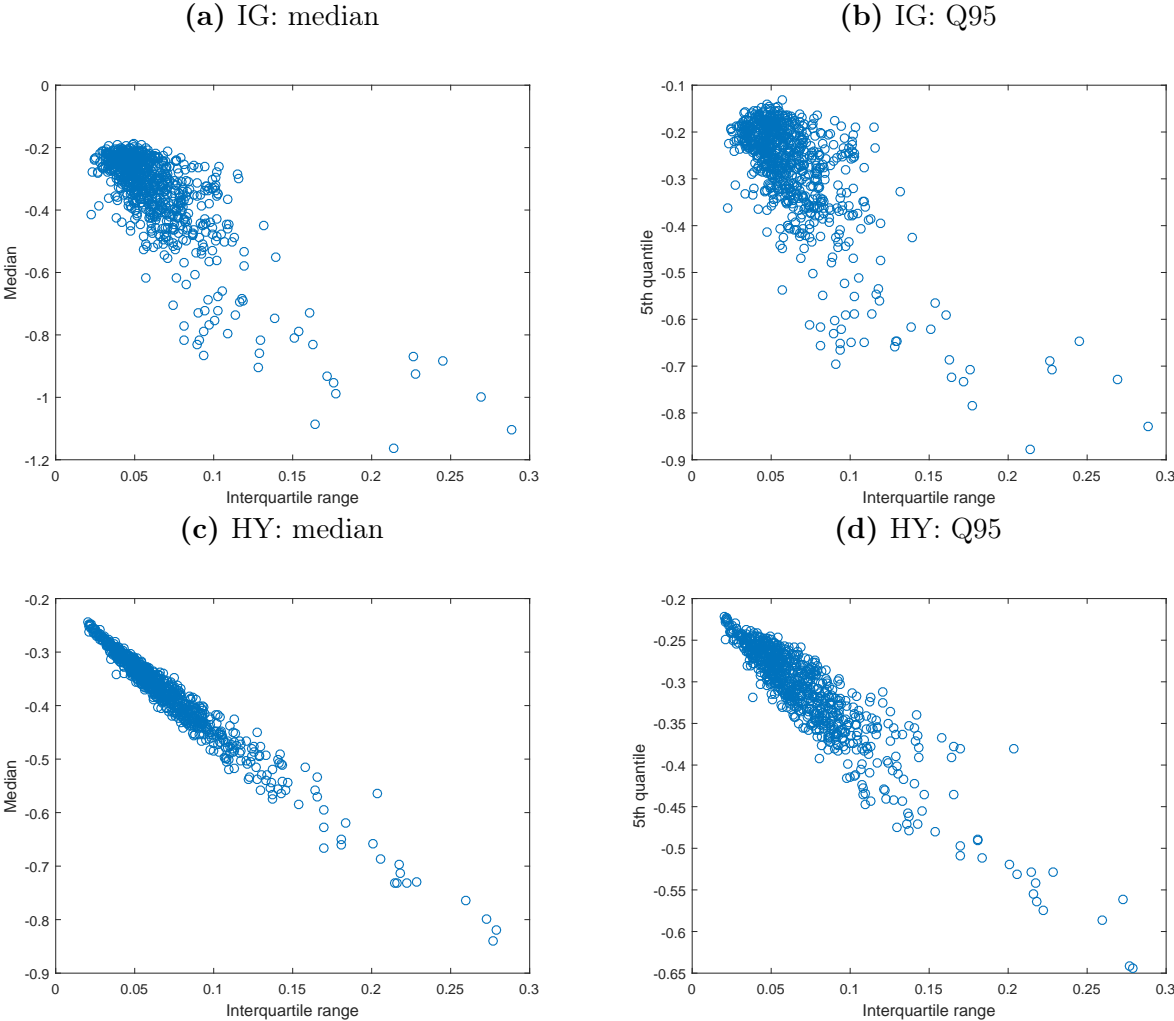




Figure 7. The Conditional Quantiles and the Skewed- t Distribution. This figure shows the conditional quantiles together with the estimated skewed- t inverse cumulative distribution functions for one-week-ahead volume-weighted average bid-ask spread by credit rating category on two dates: September 19, 2008 (“stressed” market; left column) and January 13, 2006 (“calm” market; right column). The plots include the empirical unconditional distribution for comparison.


(a) IG: 9/19/2008



(b) IG: 1/13/2006



(c) HY: 9/19/2008



(d) HY: 1/13/2006






Figure 8. Probability Densities. This figure shows the estimated skewed- t density functions for one-week-ahead volume-weighted average bid-ask spread by credit rating category on two dates: September 19, 2008 (“stressed” market; left column) and January 13, 2006 (“calm” market; right column). The plots include the empirical unconditional density for comparison.


(a) IG: 9/19/2008



(b) IG: 1/13/2006



(c) HY: 9/19/2008



(d) HY: 1/13/2006



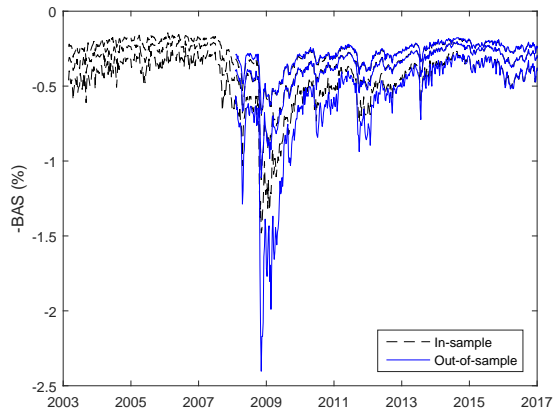


Figure 9. Out-of-Sample Predictions. This figure compares the in-sample and out-of-sample predicted distribution one-week-ahead (left column) and four weeks ahead (right column) of volume-weighted average bid-ask spread by credit rating category. The quantiles plotted are the 5th, 50th and 95th percentile.

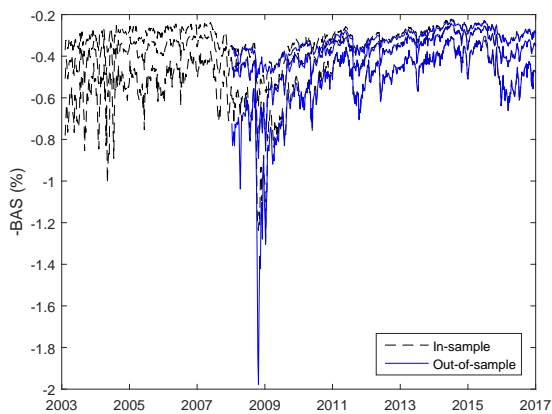
(a) IG: 1 week ahead



(b) IG: 4 weeks ahead



(c) HY: 1 week ahead



(d) HY: 4 weeks ahead




Figure 10. Out-of-Sample Accuracy. This figure reports the predictive scores (left column) and the cumulative distribution of the probability integral transform (PITs) for the one-week-ahead predictive distribution of volume-weighted average bid-ask spread by credit rating category. Predictive distribution conditions on lagged bid-ask spreads for both credit rating categories. Scores and PITs for the unconditional distribution included for comparison. Critical values obtained as in Rossi and Sekhposyan (2017).

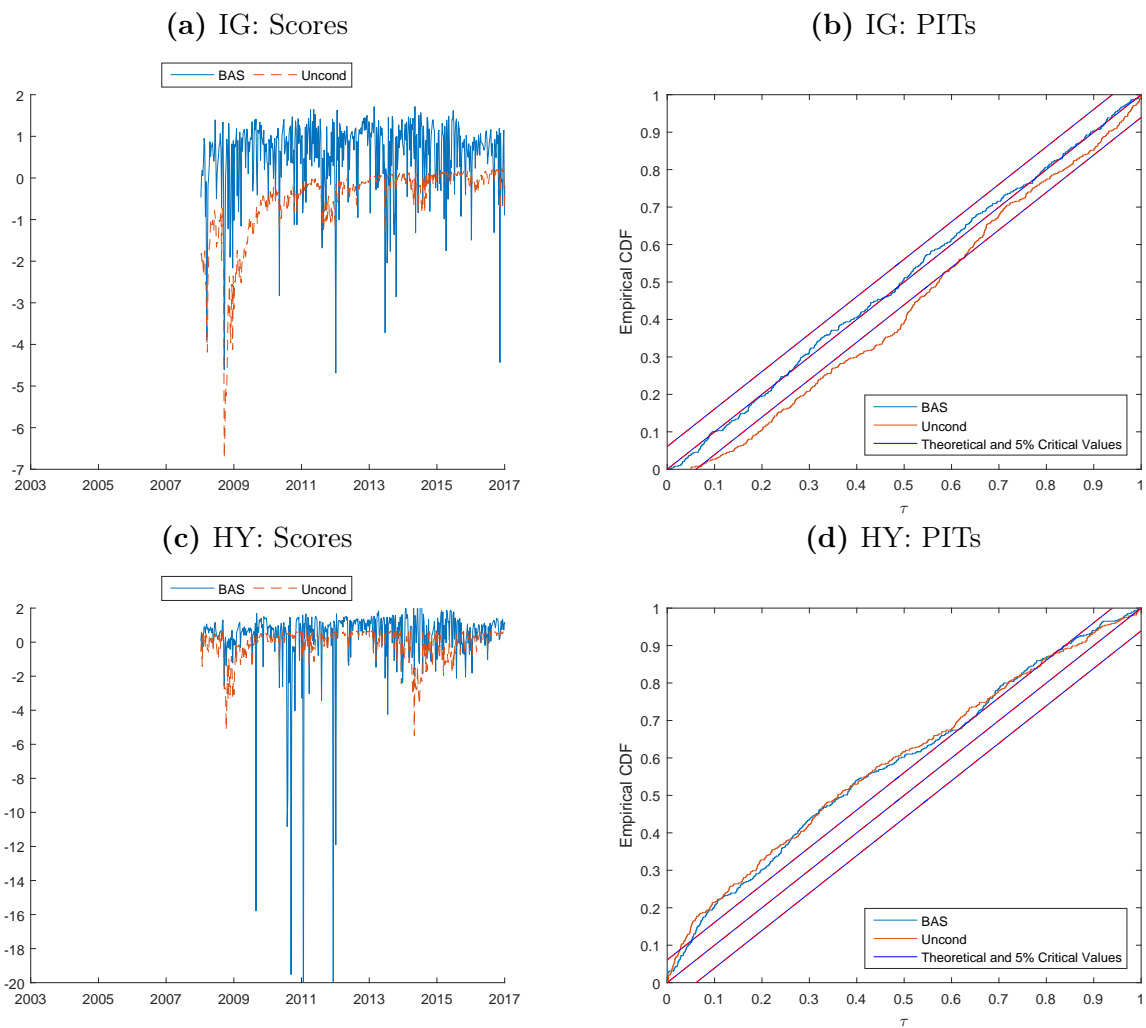


Figure 11. Conditional Distribution after Taper Tantrum. This figure shows the conditional distribution of one-week-ahead (left column) and four weeks ahead (right column) of volume-weighted average bid-ask spread by credit rating category after Taper Tantrum (June 19, 2013). Predictive distribution conditions on lagged bid-ask spreads for both credit rating categories. The unconditional distribution is included for comparison. Both distributions are estimated using bid-ask spread and market volatility data up to June 21, 2013. Reported p -values correspond to the upper tail quantiles.

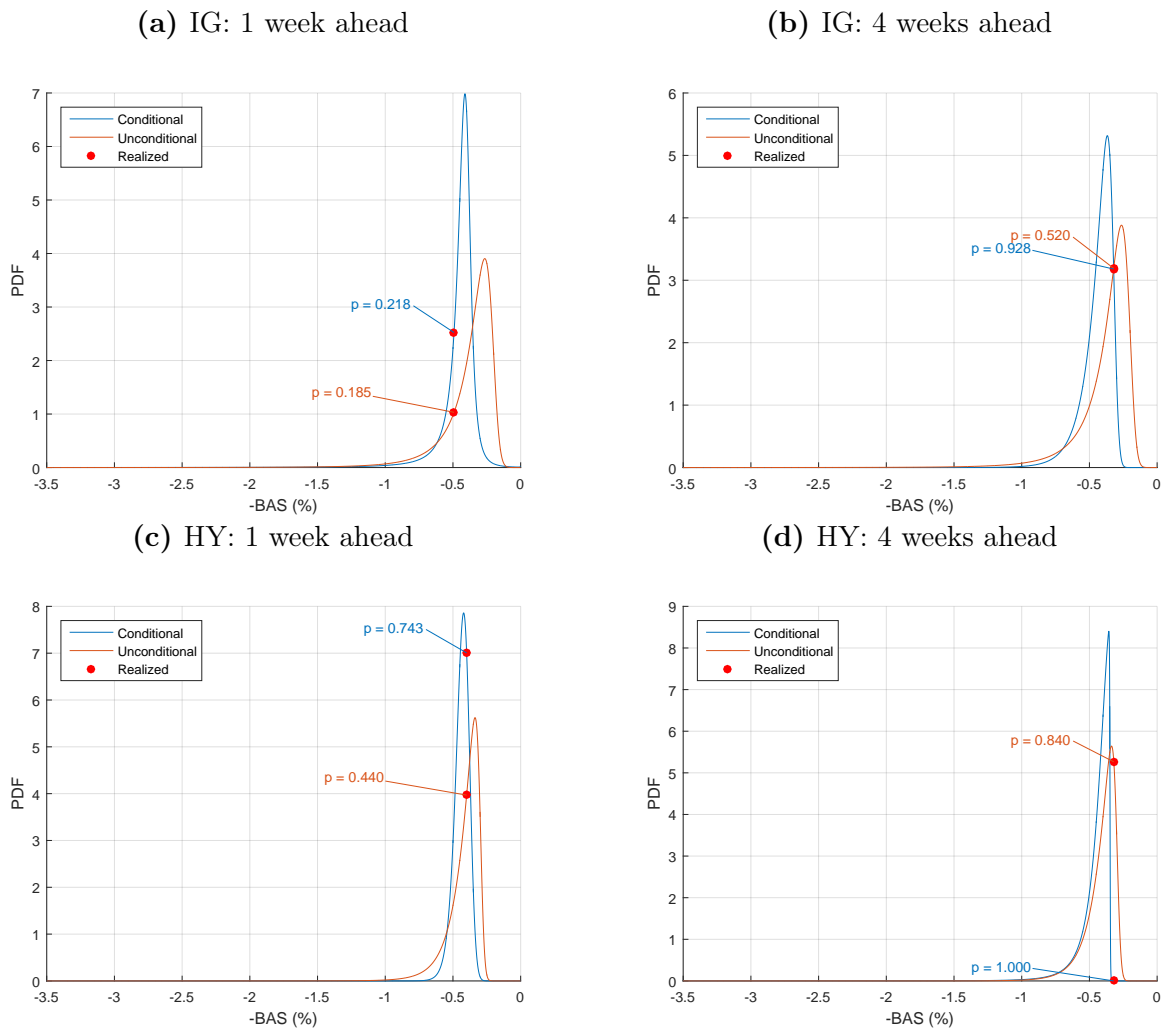
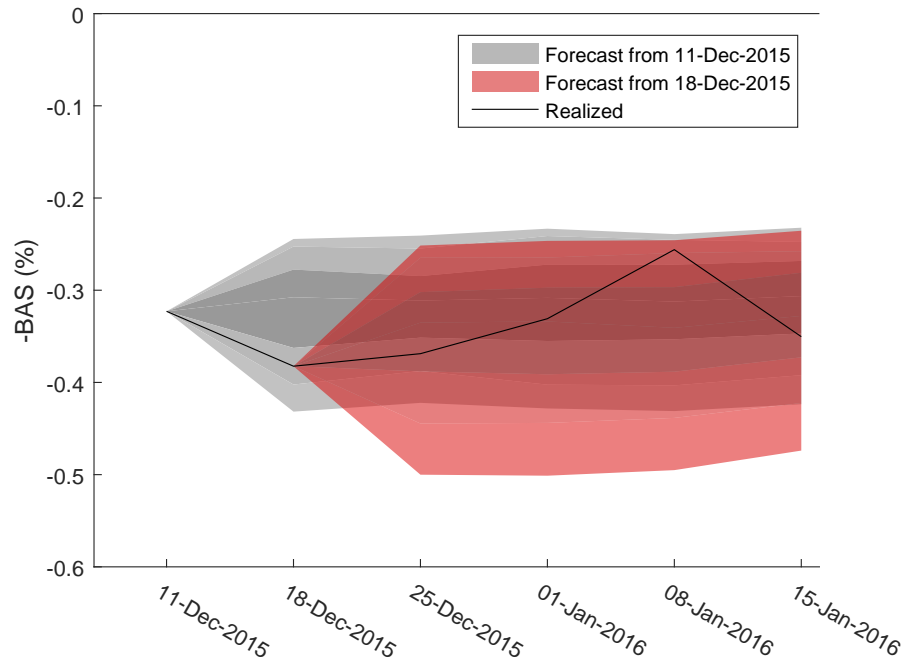


Figure 12. Conditional Distribution around Third Avenue Liquidation. This figure shows the out-of-sample conditional ICDF the week before Third Avenue Fund liquidation on December 15, 2015 and the week of the liquidation of volume-weighted average bid-ask spread by credit rating category around Third Avenue Fund liquidation on December 15, 2015.

(a) IG



(b) HY






Figure 13. Information Captured by Downside Entropy and Expected Shortfall. This figure illustrates the information captured by downside relative entropy and the five percent expected shortfall for the one-week-ahead out-of-sample predicted distribution of investment grade (excluding AAA) bid-ask spread on September 19, 2008 (“stressed market”; left column) and January 13, 2006 (“calm” market; right column). Downside entropy is the area below the conditional median between the conditional and the unconditional distribution. Expected shortfall is the expected bid-ask spread in the worst (bottom) five percent outcomes.


(a) Relative entropy: 9/19/2008



(b) Relative entropy: 1/13/2006



(c) Expected shortfall: 9/19/2008



(d) Expected shortfall: 1/13/2006






Figure 14. Measures of Liquidity Flightiness. This figure plots the time series evolution of the relative downside and upside entropy (left panel) and the five percent expected shortfall and longrise (right panel) for one-week-ahead volume-weighted average bid-ask spread by credit rating category.


(a) IG: Relative entropy



(b) IG: Shortfall/Longrise



(c) HY: Relative entropy



(d) HY: Shortfall/Longrise




Figure A.1. Estimated Quantile Regression Coefficients. This figure shows the estimated coefficients in quantile regressions of four-weeks-ahead bid-ask spreads for investment grade and high yield bonds on four lags of log bid-ask spreads for investment grade (left column) and high yield bonds (right column). Regression coefficients reported as the sum of the coefficients on the four lag of the respective variable. We report confidence bounds for the null hypothesis that the true data-generating process is a general, flexible linear model for bid-ask spreads (VAR with 4 lags); bounds are computed using 1000 bootstrapped samples.

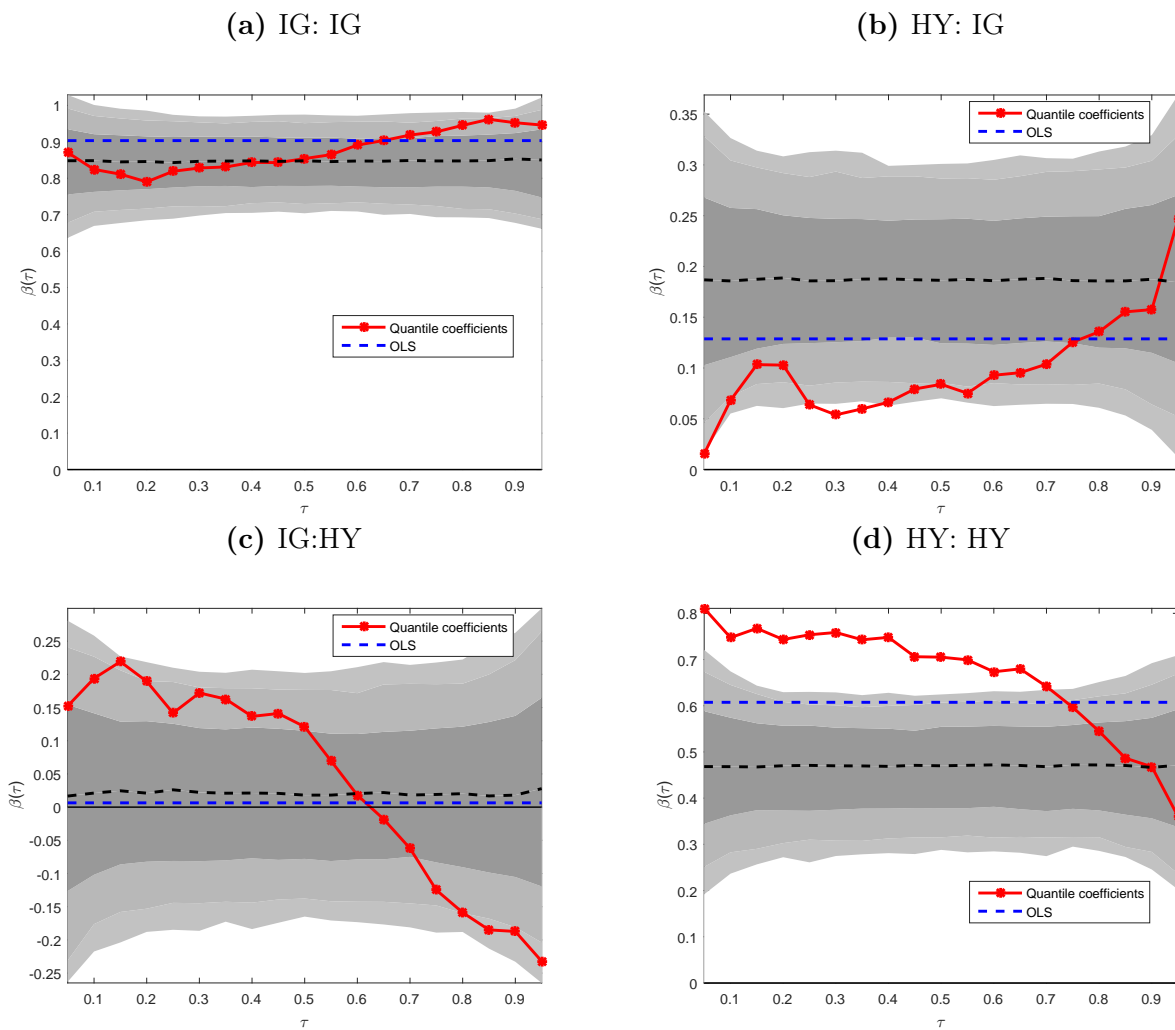




Figure A.2. Median, Interquartile Range and Tail Outcomes. This figure shows relationship between the interquartile range and the median (left column) and the interquartile range and the 95th percentile (right column) of the four-weeks-ahead conditional distribution of volume-weighted average bid-ask spread by credit rating category.


(a) IG: median



(b) IG: Q95



(c) HY: median



(d) HY: Q95




Figure A.3. Out-of-Sample Accuracy. This figure reports the predictive scores (left column) and the cumulative distribution of the probability integral transform (PITs) for the four-weeks-ahead predictive distribution of volume-weighted average bid-ask spread by credit rating category. Predictive distribution conditions on lagged bid-ask spreads for both credit rating categories. Scores and PITs for the unconditional distribution included for comparison. Critical values obtained as in Rossi and Sekhposyan (2017).

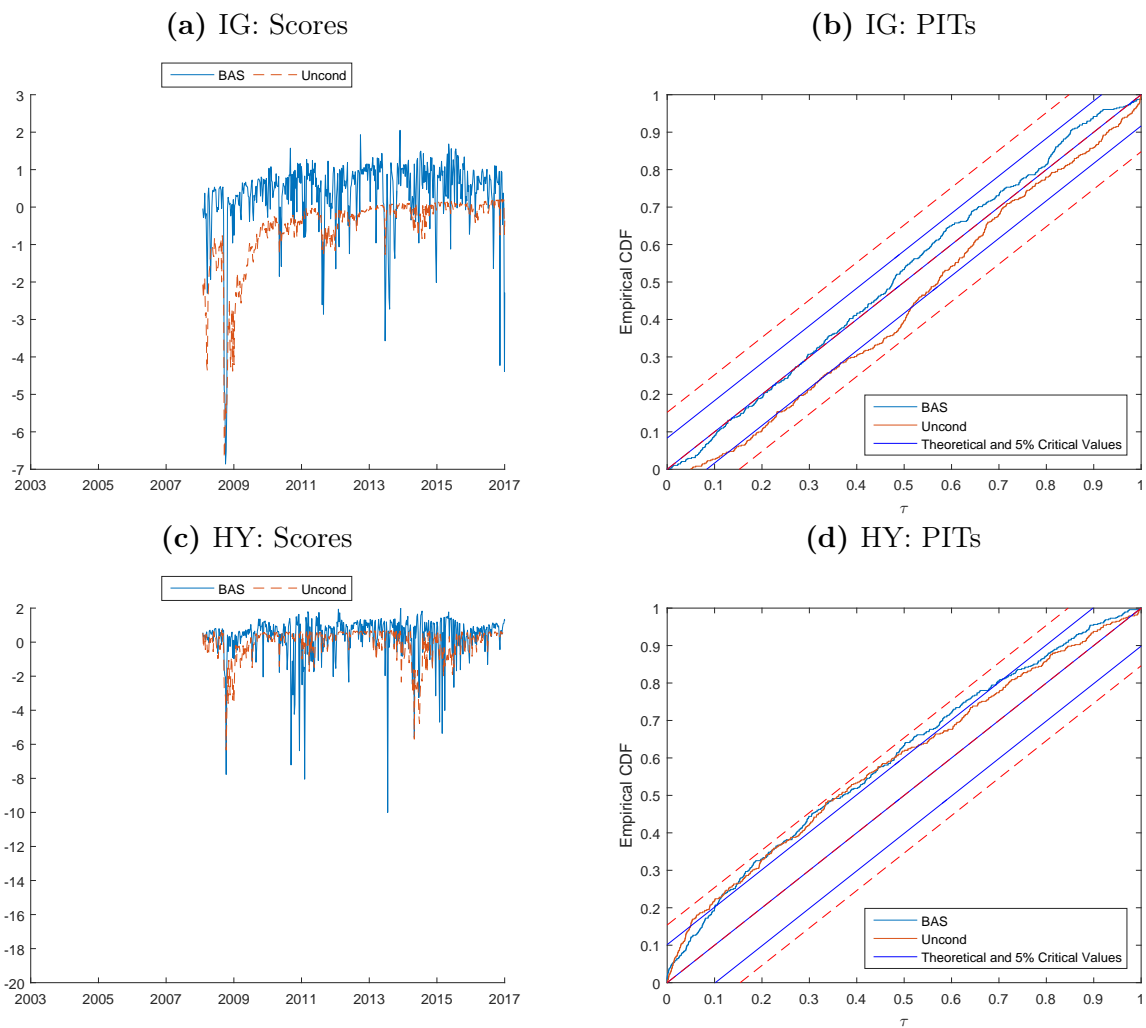




Figure A.4. Measures of Liquidity Flightiness. This figure plots the time series evolution of the relative downside and upside entropy (left panel) and the five percent expected shortfall and longrise (right panel) for four-weeks-ahead volume-weighted average bid-ask spread by credit rating category.


(a) IG: Relative entropy



(b) IG: Shortfall/Longrise



(c) HY: Relative entropy



(d) HY: Shortfall/Longrise

