An asset pricing approach to liquidity effects in corporate bond markets[*]

Dion Bongaerts, Frank de Jong and Joost Driessen

September 2011

Abstract

We use an asset pricing approach to compare the effects of expected liquidity and liquidity risk on expected U.S. corporate bond returns. Liquidity measures are constructed for bond portfolios using a Bayesian approach to estimate Roll's measure. The results show that expected bond liquidity and exposure to equity market liquidity risk affect expected bond returns, and that these liquidity effects explain a substantial part of the credit spread puzzle. In contrast, we find robust evidence that exposure to corporate bond liquidity shocks carries an economically negligible risk premium. We develop a simple theoretical model to explain this result.

Keywords: Liquidity premium, liquidity risk, corporate bonds, credit spread puzzle

JEL: C51, G12, G13

1 Introduction

Illiquidity plays a major role in corporate bond markets. While some corporate bonds are traded on a daily basis, many other bonds trade less frequently. The corporate bond market is therefore very well suited to study the price effects of liquidity. Several studies have recently examined whether illiquidity affects corporate bond prices. Most of these studies regress a panel of credit spreads on liquidity measures, thus using liquidity as a bond characteristic. A few recent articles analyze whether there is a premium associated with exposure of corporate bond returns to systematic liquidity shocks in the corporate bond market or equity market (see Section 2).

The first contribution of this paper is that we integrate these two approaches. We perform a detailed comparison of the effects of liquidity as a bond characteristic (liquidity level) and various forms of liquidity risk (both equity market liquidity risk and corporate bond liquidity risk). We do this using a formal asset pricing approach. Given that liquidity level and liquidity risk exposures are typically highly correlated, neglecting either the liquidity level or liquidity risk may lead to misleading conclusions on the effects of these different liquidity measures (Acharya and Pedersen (2005) illustrate this for the equity market). Determining which liquidity channel is most important is relevant for several reasons. First, most theoretical models that generate price effects of liquidity focus on the liquidity level, and not on liquidity risk (see, for example, Vayanos (2004) and Vayanos and Wang (2009)). Second, the extent to which optimal financial portfolios are affected by illiquidity also depends on whether liquidity risk or the liquidity level is priced. Finally, disentangling these liquidity effects is important for the valuation of illiquid assets (Longstaff (2010)). Our results show that both the liquidity level and exposure to equity market liquidity risk have a strong and robust effect on corporate bond prices, while the effect of systematic corporate bond liquidity risk is mostly insignificant

1

and always economically negligible.

Why is corporate bond liquidity risk not priced, while exposure to equity market liquidity risk does carry a risk premium? Our second contribution is to provide a simple theoretical model that explains these empirical findings. In our model, investors prefer to trade and rebalance their portfolio using liquid assets such as equities, and avoid trading in the relatively illiquid corporate bonds as much as possible.[1] Our model then predicts that the liquidity risk associated with relatively liquid assets such as equities is important and carries a risk premium, since this liquidity risk captures the extent to which trading costs increase in bad times. In contrast, exposure to liquidity risk of illiquid assets (corporate bonds) will not be priced in equilibrium since these shocks are less relevant for the investor as he avoids trading these illiquid assets when transaction costs are high. To provide empirical evidence for this hypothesis, we analyze turnover patterns in the equity and corporate bond market. We find that average turnover in the corporate bond market is much lower than turnover in the equity market. In addition, we find that corporate bond turnover goes down in bad times (when prices decline and liquidity costs go up). In contrast, for equity markets turnover actually increases in bad times, in line with the notion that in bad times, investors need to trade more and choose to use the most liquid assets to do so.

Our third contribution is to show that our liquidity-based asset pricing model sheds light on the "credit spread puzzle". This puzzle states that credit spreads and expected returns on corporate bonds are much higher than what can be justified by expected losses and exposure to market risk factors (see Elton, Gruber, Agrawal and Mann (2001) and Huang and Huang (2003)). We show that liquidity effects play an important role in explaining this credit spread puzzle. Especially for high-rated bonds, a considerable

---

[1]Indeed, several articles study optimal rebalancing of assets in case of transaction costs, and derive no-trade ranges that are higher when transaction costs are higher (see Constantinides (1986) for a seminal contribution).

part of the expected return can be explained by the illiquidity of these bonds.

This paper is related to existing work on corporate bonds and liquidity. As discussed in Section 2 in more detail, our paper thus contributes to this literature by (i) studying both expected liquidity and liquidity risk effects (using a formal asset pricing model with forward-looking expected returns), (ii) developing a simple theoretical model to explain why corporate bond liquidity risk is not priced, and (iii) studying the implications for the credit-spread puzzle.

Our analysis uses data from *TRACE* (Trade Reporting and Compliance Engine) for a 2005 to 2008 sample period, which thus includes the 2007-2008 crisis period. Since 2005 essentially all U.S. corporate bond transactions have been recorded in *TRACE*. We have data at the transaction level but do not know who initiated the trade. We also do not have price quotes hence we cannot use the Lee and Ready (1991) method to assess the trade directions. In this context, Hasbrouck (2009) proposes a Bayesian approach to estimate the Roll (1984) measure of effective transaction costs. We extend his approach to a portfolio setting and adapt it to fit the bond market. Using the Gibbs sampler, this approach provides us with time series of returns and liquidity estimates at the portfolio level. For the equity market liquidity, we use both Amihud's (2002) ILLIQ measure and the Pastor-Stambaugh (2003) measure.

A critical issue in any asset pricing test is the measurement of expected returns. This is particularly true for corporate bonds. Average returns on corporate bonds critically depend on the number of defaults over the sample period, and given the rare occurrence of default events this implies that average returns are noisy estimates of expected returns. In addition, transaction data for corporate bonds are only available for short sample periods. Also, using average returns in the presence of microstructure noise may bias towards finding liquidity effects, see Asparouhova, Bessembinder, and Kalcheva (2010).

Therefore, we follow Campello, Chen and Zhang (2008), de Jong and Driessen (2006) and Bongaerts, de Jong and Driessen (2011) and construct forward-looking estimates of expected returns. We do this by correcting the credit spread, which captures the return of holding corporate bonds to maturity in excess over the government bond return, for the expected default losses. This expected loss is calculated using default probability estimates from *Moody's-KMV* and assumptions on the loss rate in case of default.

We then construct various double-sorted corporate bond portfolios, sorting first on credit quality (credit rating, estimated default probabilities) and then on liquidity proxies (trading volume, bond age, amount issued, liquidity betas). In a first step, we estimate exposures of these portfolio returns to equity market risk, volatility risk, corporate bond liquidity risk and equity market liquidity risk. Corporate bond liquidity risk is captured by innovations in the aggregate Roll measure. In a second step, we regress the cross-section of forward-looking expected returns on the portfolio liquidity levels, market betas and the various liquidity betas.

The first-step results show that corporate bonds have significant exposures to equity market returns, volatility risk, corporate bond market returns, and systematic liquidity risk measures for the equity and corporate bond markets. Equity market returns, volatility risk and liquidity risk together explain about 65% of the time-series variation in corporate bond returns.

The second-step cross-sectional regressions generate several key findings. First, the liquidity level (expected liquidity) substantially affects expected returns, leading to higher expected returns for portfolios with lower expected liquidity, even when controlling for equity market, liquidity and volatility factors. This expected liquidity premium is both economically and statistically significant. Second, we find that exposure to equity market liquidity risk is also priced, irrespective of whether we use the Amihud or

Pastor-Stambaugh measure. Third, the corporate bond liquidity risk premium is economically negligible in all specifications. Finally, we also find significant and robust premia for equity market risk and volatility risk.

The finding that corporate bond liquidity risk is not priced is surprising, especially given existing work (which we discuss in Section 2). We therefore perform several robustness checks to validate this result. First, we find similar results when using a Fama-MacBeth approach where we incorporate time-variation in expected returns, betas and liquidity levels. Second, using a pre-crisis subsample also generates very similar results. Third, we construct portfolios that are directly sorted on corporate bond liquidity betas and find that even in the cross-section of these portfolios corporate bond liquidity risk is not priced. Fourth, instead of using Roll's liquidity measure, we use the market average of the imputed roundtrip cost measure of Feldhütter (2011) and Dick-Nielsen, Feldhütter and Lando (2011) to measure corporate bond liquidity risk. This does not affect the results. In fact, this measure is highly correlated with the aggregate Roll measure of corporate bond liquidity risk. Fifth, it may be that liquidity and credit risk are correlated. We therefore include the *Moody's-KMV* default probability estimates as a control variable, and find that the results do not change substantially. Finally, we use an alternative liquidity pricing model, following Acharya and Pedersen (2005) and Bongaerts, de Jong and Driessen (2011). In these models various liquidity covariances can affect expected returns, but our results show that the premia related to these corporate bond liquidity risk measures have a negligible effect on expected returns.

Another concern could be that estimation error in the corporate bond liquidity betas makes it hard to find a substantial risk premium, thus making the comparison with liquidity level unfair. We deal with this in several ways. First, we note that in the first-step time-series regressions corporate bond liquidity betas are estimated quite precisely, with an average t-statistic of -8.6 in univariate regressions and -4.0 in multivariate

regressions. Second, we stress that the liquidity level is also estimated with error. In fact, the average t-statistic of the liquidity level estimate is equal to 6.2, hence in the same range as the liquidity beta t-statistics. Third, the time-series regressions show that corporate bond liquidity shocks alone explain 26.8% of the time-series variation in corporate bond returns, which shows that our corporate bond liquidity measure does not simply reflect noise. Fourth, a final concern could be that the significance of the liquidity beta estimates is driven only by the large liquidity shocks in the Fall of 2008. We therefore use a subsample up to August 2008 and find average t-statistics of liquidity betas of -7.5 (univariate) and -5.7 (multivariate).

In sum, we show that an asset pricing model with expected liquidity and premia to equity market liquidity risk, equity market risk and volatility risk provides a very good fit of expected bond returns, with a cross-sectional R^2 of about 70%. Across all portfolios, the average expected excess bond return equals about 1.9% per year, of which about 1% is due to expected liquidity, while equity market liquidity risk, equity market risk and volatility risk each contribute about 0.3% to the expected excess return. This model fits both expected returns on high-rated and low-rated bonds very well, and thus goes a long way in explaining the credit spread puzzle. Including expected liquidity is particularly important for explaining the high returns on high-rated bonds.

The remainder of this paper is organized as follows. In Section 2 we discuss the related literature. Section 3 introduces the asset pricing models that we estimate. Section 4 describes the data and the Bayesian approach to estimate Roll's model. Section 5 contains the empirical results. Section 6 presents various robustness checks. Section 7 concludes.

2 Comparison with existing literature

Our paper is related to two streams in the literature on corporate bonds and liquidity. The first stream uses liquidity as a bond characteristic, and analyzes, typically in a panel setting, the relation between the credit spread on a corporate bond and its liquidity. This stream includes Houweling, Mentink and Vorst (2005), Covitz and Downing (2006), Nashikkar and Subrahmanyam (2006), Chen, Lesmond and Wei (2007), Bao, Pan and Wang (2010), and Friewald, Jankowitsch and Subrahmanyam (2010). Our paper differs from this stream in two important ways. First, instead of analyzing credit spreads in a panel setting, we estimate a formal asset pricing model, where we explain (in two steps) the time-series of returns and the cross-section of expected returns. Second, we include both liquidity level (a bond characteristic) and several liquidity risk exposures in the asset pricing model. The advantage of an asset pricing model is that it puts structure on the model specification and allows for a direct interpretation of the coefficients in terms of risk exposures and risk premia.

The second, smaller, stream in this literature analyzes the effect of liquidity risk on corporate bonds. De Jong and Driessen (2006) show that equity market liquidity risk is priced in a cross-section of corporate bond portfolios, while Acharya, Amihud and Bharath (2010) show that corporate bonds are exposed to liquidity shocks in equity and treasury markets. Both articles do not investigate corporate bond liquidity risk, nor do they incorporate the liquidity level.

Four recent articles study the pricing of corporate bond liquidity risk. Dick-Nielsen, Feldhütter and Lando (2011) mainly focus on liquidity levels to explain credit spread levels (while we analyze expected returns), but do find some effect of liquidity betas on credit spread levels as well. However, their focus is on explaining the panel of individual credit spreads within each rating category, while our focus is to explain variation across

portfolios sorted on credit and liquidity proxies. They do not estimate an asset pricing model. For example, the coefficient on liquidity betas (the liquidity risk premium) is estimated separately for each rating category, which results in very different coefficient estimates. Chacko (2005), Downing, Underwood and Xing (2005), and Lin, Wang and Wu (2010) construct various corporate bond liquidity risk measures, and show these are priced in a cross-section of corporate bond returns. There are two important differences between these three studies and our work. First, we include both expected liquidity and liquidity risk. As discussed in the introduction, given that liquidity level and liquidity risk exposure are correlated, omitting one of the two may affect the results. Indeed, if we only include corporate bond liquidity risk exposure in our regressions (without liquidity level or market risk exposures), we do find a significant corporate bond liquidity risk premium, although the effect is economically small.[2] Second, while the existing studies use realized corporate bond returns to estimate expected returns, our work complements these studies by using a forward-looking measure of expected returns. Given the short sample period available for corporate bonds, and given the skewed nature of corporate bond returns (depending on the number of defaults in the sample period), we believe that it is worthwhile to explore the effects of liquidity on forward-looking expected returns. A further concern is survivorship bias. In actual returns of defaulted bonds, the returns at default often do not show up leading to upward biased average returns. In our forward looking measure, we account properly for the possibility of default events. An additional argument for using forward-looking expected returns is that using average returns in the presence of microstructure noise may bias towards finding liquidity effects (see Asparouhova, Bessembinder, and Kalcheva (2010)).

---

[2]Lin, Wang, and Wu (2010) provide a robustness check where they control for the liquidity level, by multiplying Amihud's ILLIQ measure with the turnover rate of corporate bonds and subtracting this from the average bond returns. This assumes that the ILLIQ level itself equals the transaction costs of trading, which is not necessarily the case as the scale and the trend in these measures are quite different, see Acharya and Pedersen (2005).

Our paper is also related to the broader literature investigating liquidity effects in financial markets. In particular, and in line with out findings, several articles have found that equity market liquidity risk is priced outside the cross-section of equities (see for example Franzoni, Novak, and Phalippou (2011) for the private equity market and Sadka (2009) for hedge funds). Also related is recent work of Lou and Sadka (2010) that compares the role of liquidity level and liquidity risk in the equity market during the recent financial crisis, and finds that stocks with high liquidity risk underperformed during the crisis relative to stocks with low liquidity risk, while there is less effect of liquidity level on returns during the crisis.

Finally, our liquidity-based asset pricing model helps to explain the "credit spread puzzle". In addition to the seminal work of Elton et al. (2001) and Huang and Huang (2003), previous work on this puzzle includes Cremers, Driessen and Maenhout (2005), David (2008) and Chen, Collin-Dufresne and Goldstein (2009). None of these articles incorporates liquidity effects.

3 Asset pricing model

In the benchmark analysis we use a standard risk factor approach to formalize the impact of liquidity on corporate bond prices, following Pastor and Stambaugh (2003) who use this approach to study liquidity risk effects in equity markets. We regress the time series of corporate bond excess returns r_{it} on a set of risk factor innovations F_t (not necessarily returns)

$$r_{it} = \beta_{0i} + \beta_i' F_t + \epsilon_{it}. \tag{1}$$

Our forward-looking estimate of the expected excess returns $\widehat{E}(r_{it})$, as constructed from credit spreads corrected for expected default losses (see Section 4.4), is then regressed

9

on the betas and the expected transaction costs

$$\widehat{E}(r_{it}) = \lambda' \beta_i + \zeta E(c_{it}) + \alpha_i, \tag{2}$$

where c_{it} denotes the transaction costs (relative to the asset price) and α_i denotes the error term of the cross-sectional regression, which can be interpreted as the pricing error of asset i. The theory predicts that the intercept in this regression is zero since we focus on excess returns. The coefficients λ measure the market prices of factor risk, and ζ measures the impact of transaction costs and can, under some assumptions, be interpreted as the turnover rate of the asset (Amihud and Mendelson (1986)). The risk factors we include are the equity market return, innovations in corporate bond market liquidity and equity market liquidity, and innovations in the VIX index. As a robustness check, we apply the liquidity asset pricing approach of Acharya and Pedersen (2005) in Section 6.7.

Note that this risk-factor approach can be used to study the credit spread puzzle, as long as we do not use the corporate bond market return as risk factor to avoid that the puzzle is present on both the left-hand side and right-hand side of the equation. We therefore explain the expected corporate bond returns from equity market and volatility risk exposure and liquidity effects.

4 Measuring bond returns and liquidity

4.1 Portfolio selection

For our analysis we use individual bond transaction data from the *TRACE* database.[3]
From July 2002 onwards the NASD discloses all corporate bond trades that all its
affiliated traders are required to report. Initially only trades in a limited number of
bonds were disclosed, but gradually disclosure expanded to reach full disclosure from
October 2004 on. We thus download all trade data from *TRACE* from October 2004 up
to end of December 2008 so that we have a sample with homogeneous coverage. After
applying several data filters (see Appendix A for details) we end up with approximately
4.4 million bond trades. For each bond we calculate a yield and a credit spread by
comparing the bond yield with a duration-based weighted average of the yield on two
treasuries with bracketing duration.

As is usual in the asset pricing literature, we fit the model to different test portfolios
rather than to individual assets. To this end, we form portfolios which are sorted first on
credit quality and thereafter on liquidity. To increase the number of test assets, we sort in
each dimension using different variables. To conduct the credit quality sorts, we use the
S&P credit rating at the end of the previous quarter (AAA, AA, A, BBB, BB, B, CCC)
or the cumulative default probability over the life of the bond estimated by *Moody's-
KMV* EDFs (quintile portfolios). For the liquidity dimension, we sort by amount issued,
bond age, and number of trades in the previous quarter. Amount issued and age have
been shown to be good proxies for liquidity by Houweling, Mentink and Vorst (2005),
while typically the number of trades will be higher for more liquid securities. In the
liquidity dimension, we categorize a bond as either liquid or illiquid. The cutoff point

---

[3]A good description of the TRACE data can be found in Lin, Wang and Wu (2010).

for amount issued and age is the median, whereas for the trade count it is the 70%
percentile. This proportion is required to ensure that there are enough trades in the low
activity portfolio. The AAA and CCC rated portfolios contain too few observations to
conduct a double sort, but are included as rating portfolios. This yields 62 portfolios
consisting of almost 15,000 different bonds. These portfolios form the basis of our tests.

4.2 The Roll model for bond returns

Estimating returns and transaction costs from the $TRACE$ data is not trivial. The data
contain a record of transaction prices and trade volume, but no quote or bid-ask spread
information. The data also do not indicate whether the transaction was a buy or a sell.
The data are also irregularly spaced: some bonds trade several times a day, but many
bonds trade very infrequently. To deal with these issues, we use the basic Roll (1984)
model suggested by Hasbrouck (2009) as the basis of our analysis, and adapt it to a
setting where we form portfolios of bonds. We start by modeling the credit spread of
bond i at time t, denoted CS_{it} as

$$CS_{it} = m_{it} + c_{it}q_{it}, \tag{3}$$

where m_{it} is the efficient credit spread level and q_{it} is an i.i.d. trade indicator that can
take values $+1$ and -1 with equal probability. The coefficient c_{it} is the effective bid-ask
half-spread in yield terms (effective transaction costs). We focus on credit spreads rather
than prices to take out most of the effects of interest-rate risk and implicit weighting
induced by maturity differences.

Following Hasbrouck (2009), we write this model in first difference form

$$CS_{it} - CS_{i,t-1} = \Delta m_{it} + c_{it}q_{it} - c_{i,t-1}q_{i,t-1}, \tag{4}$$

where Δm_{it} is the innovation in the efficient credit spread. We model the change in this efficient credit spread as the sum of an element common to the portfolio to which bond i is allocated, and an idiosyncratic component

$$\Delta m_{it} = z_{it}\Delta M_t + u_{it}v_t, \tag{5}$$

where $\Delta M_t \sim N(0, \sigma_M^2)$ represents the portfolio-level spread change, $u_{it} \sim N(0, \sigma_u^2)$ the idiosyncratic shock with v_t an observable scale factor that captures heteroskedasticity. This is important as the volatility of idiosyncratic shocks may change over time. Empirically we use the level of the VIX index for v_t. We let the loading on the common factor ΔM be dependent on the bond duration with

$$z_{i,t_{ik}} = 1 + \gamma(Duration_{ik} - \overline{Duration}), \tag{6}$$

where γ is estimated in a first step, $Duration_{ik}$ is the duration of bond i at trade k, and $\overline{Duration}$ is the average duration of all bonds in the portfolio.[4] This factor z_{it} captures patterns in the term structure of volatilities. For example, if long-term credit spreads are less volatile than short-term credit spreads, one would expect a negative γ. The latent components ΔM and u are independent. Furthermore, we assume that the transaction costs are the same for all bonds in the same portfolio, $c_{it} = c_t$.

In our analysis, we use hourly time intervals, but not every bond trades each hour and we therefore use a repeat sales methodology (see, for example, Case and Shiller (1987)). Let t_{ik} denote the time of the k'th trade in bond i. Taking differences with respect to the previous trade of bond i, these assumptions lead to the complete model

---

[4]Specifically, γ is estimated by using a repeat sales methodology to estimate a restricted version of equation (7) with the transaction costs c set to zero.

for all data in the same portfolio,

$$CS_{i,t_{ik}} - CS_{i,t_{i,k-1}} = \sum_{s=t_{i,k-1}+1}^{t_{ik}} z_{is}\Delta M_s + c_{t_{ik}}q_{i,t_{ik}} - c_{t_{i,k-1}}q_{i,t_{i,k-1}} + e_{it}, \qquad (7)$$

where $e_{it} = \sum_{s=t_{i,k-1}+1}^{t_{ik}} u_{is}v_s$. We estimate the components of equation (7) using a Bayesian approach and the Gibbs sampler, following Hasbrouck (2009). For each portfolio, this approach gives us posterior distributions for the time series of the common credit-spread factor ΔM, the transaction costs c, and the posterior probabilities of the trade indicators q. The identification assumptions of the model imply that the values of the transaction costs c are always positive. Appendix B details this estimation method in full detail.

In the actual estimation, we assume the transaction costs to be constant within every week, and estimate credit spread changes ΔM for every hour. We transform these credit spread changes to returns r by multiplying these changes with (minus) the duration of the bond portfolio. This gives (to first order) the return on the corporate bond portfolio in excess of the government bond return. Similarly, the transaction costs in terms of yields are transformed to price-based transaction costs by multiplying the costs c by the bond duration.[5].

These returns are then aggregated to weekly returns, so finally the Roll model produces a time series of weekly portfolio excess returns and transaction costs. In the equations above we suppressed the subscript of each portfolio j. In the remainder of the paper, the subscript j refers to portfolio j.

---

[5]See Bongaerts, de Jong and Driessen (2011) for a derivation of the relation between yield-based and price-based transaction costs.

4.3 Validation of the liquidity estimates

As of November 2008, the $TRACE$ data do contain the trade indicators q_{it}: for each transaction it is recorded whether this was buyer-initiated or seller-initiated. This allows us to do a strong check on the estimation of the Roll model describe above. We thus estimate the transaction costs in equation (7) in two ways. First, we use the Gibbs sampler, where we do not use information on the trade indicators ("indirect" approach). Second, we use the observed trade indicators in which case (7) can directly be estimated using a repeat-sales regression approach ("direct" approach).

We perform this analysis on the portfolios where we first sorts on rating or EDF, and then on amount issued. We calculate the correlation between the weekly series of transaction costs, estimated using either the direct or indirect approach, using data until end of 2009. We find that the average of these time-series correlations equals 78%, and the correlations range between 63% and 97% across portfolios (except for one portfolio which has correlation of -4.5%, but this portfolio has relatively few bond issues). Also, the average level of the direct and indirect transaction costs is very similar: 1.35% (direct) versus 1.32% (indirect) on average. This shows that, even though we do not observe the trade indicators in 2005 to 2008, it is possible to reliably estimate transaction costs on corporate bonds using the Gibbs sampling method.[6]

4.4 Time series model for liquidity

The betas in the asset pricing model are defined as the ratio of conditional covariances and variances, that is, the (co)variances of the innovations in returns and costs. We assume that returns have no serial correlation and we take the residuals of an autore-

---

[6]We do not use 2009 data for our asset pricing tests since we do not have EDF data for this period.

gressive AR(2) model as the liquidity innovations[7]

$$c_{j,t} = b_{0j} + b_{1j}c_{j,t-1} + b_{2j}c_{j,t-2} + \varepsilon_{j,t}, \tag{8}$$

where $c_{j,t}$ now denotes the portfolio-level transaction costs. We also estimate the market-wide transaction costs by averaging costs across portfolios. We use the innovations in these market-wide costs as liquidity risk factor in our asset pricing model.

We also analyze the effects of exposure to equity market liquidity and the VIX volatility index. Following Acharya and Pedersen (2005), we construct this measure by taking AR(2)-innovations to the equally-weighted mean of Amihud's (2002) ILLIQ measure across all stocks in *CRSP*. Innovations in the VIX are also constructed using an AR(2) model.

4.5 Expected return estimates

To estimate the expected excess return $E(r_j)$ on a corporate bond portfolio j, we take the observed credit spread and correct it for expected default losses. This procedure follows de Jong and Driessen (2006), Campello, Chen and Zhang (2008), and Bongaerts, de Jong and Driessen (2011), who show that it yields much more accurate estimates of expected returns than simple averaging of historical excess returns.

The method works as follows. Consider bond i, and denote the cumulative default probability over the entire maturity of the bond π_{it}, the loss given default L, the yield on the bond y_{it} and the corresponding government bond yield y_{gt}. We approximate the coupon-paying bond by a zero-coupon bond with maturity equal to the duration of the coupon-paying bond, T_{it}. Assuming that default losses are incurred at maturity, the

---

[7]We have checked that residuals from the AR(2) series do not contain significant autocorrelation.

expected return of holding the bond to maturity equals $(1 + y_{it})^{T_{it}}(1 - L \cdot \pi_{it})$. We then annualize this number and subtract the annualized expected return on the corresponding government bond to obtain our expected excess return estimate

$$E_t(r_i) = (1 + y_{it})(1 - L \cdot \pi_{it})^{1/T_{it}} - (1 + y_{gt}). \tag{9}$$

Note that this gives an estimate of the bond-level expected return at each point in time t. The expected excess portfolio return is then constructed each week by averaging the expected excess bond returns over all trades in the portfolio in that week that have an EDF available. The unconditional expected return for a portfolio $\widehat{E}(r_{it})$ is given by the time-series average across weeks.

The loss given default L is assumed to be 60%. Default probability estimates π_{it}, needed to construct these expected excess returns, are obtained from *Moody's-KMV* EDF database. We have data on the average 1-year and 5-year annualized expected default frequencies (EDFs), which capture the conditional default rate in the first and fifth year, respectively. We construct the conditional expected default frequency for every bond as the duration weighted average of the one-year and five-year EDFs. For durations longer than 5 years, we assume that the conditional default rate is flat beyond 5 years. From these bond-specific EDFs we obtain the expected cumulative default probabilities over the entire maturity of the bond (π_{it}). We prefer using *Moody's-KMV* EDFs over rating-based default probability estimates because we observe a strong increase in the EDFs in the last two years of the sample (2007 and 2008). It is not obvious how to adjust for these new market circumstances when using rating-based default probabilities.

5 Empirical results

5.1 Correlations, expected returns, transaction costs and betas

We first check the time-series correlations between the four factors in our benchmark model: corporate bond liquidity shocks (CBLIQ), equity market liquidity shocks (EQLIQ), equity market (S&P 500) returns (EQ), and VIX innovations (VIX). Panel A of Table 1 shows that only the VIX and equity market returns exhibit very strong (negative) correlation (the so-called "leverage" effect). Liquidity shocks across bond and equity markets have a correlation of 36%. Hence, there is some commonality in liquidity across markets, but most of the liquidity variation is market-specific.

Panel B of Table 1 presents averages of expected returns, costs, betas, and associated t-stats across portfolios and over the full 2005-2008 sample period. The first key result in Panel B is that the estimated one-way transaction costs are substantial, on average 0.83% across portfolios and over time. These numbers are very similar to those of Bao, Pan and Wang (2010) who use a different method to estimate Roll's model for corporate bonds. They report a median bid-ask spread of 1.50%, implying one-way transaction costs of 0.75%, close to our estimates. As noted by Bao, Pan and Wang (2010), these estimated costs are higher than quoted bid-ask spreads as found in Bloomberg, and they argue that the Roll model thus captures liquidity effects that go beyond the quoted bid-ask spread. The second result in Panel B of Table 1 is that we find large positive expected returns (in excess of government bonds), around 1.9% per year on average, in line with earlier evidence on the credit spread puzzle. Note that this is an average of the weekly expected return estimate since we construct the forward-looking expected return measure each week. The Newey-West corrected t-statistics on these average expected return estimates are high (average t-stat of 4.8), which shows the usefulness of estimating

expected returns from credit spread levels.

The time series of market-wide average transaction costs is shown in Figure 1 (along with an alternative liquidity series based on Feldhütter (2011) and Dick-Nielsen, Feldhütter and Lando (2011), to be discussed in Section 6). Clearly, liquidity peaks during events in the credit crisis, such as the March 2008 Bear Sterns failure and the September 2008 Lehman collapse. Bao, Pan and Wang (2010) report similar illiquidity spikes in 2008.

Next we turn to the betas. Following equation (1), Panel B of Table 1 also reports results of univariate and multivariate regressions of bond returns on the four factors (CBLIQ, EQLIQ, EQ, and VIX). We scale all factors such that they have the same standard deviation as equity returns, so that one can easily compare the betas across factors. Also, we use overlapping four-weekly returns and innovations to estimate betas.

We see that corporate bonds have significant equity market exposure, which by itself explains on average 50% of the time-series variation. Equity market liquidity risk and volatility risk have similar explanatory power. We also see that corporate bond returns have significant negative exposure to systematic bond liquidity shocks, measured by innovations in the market-wide level of corporate bond transaction costs. This exposure explains alone on average about 26% of the time-series return variation. When we look at the multivariate betas, we see that the equity and volatility betas both become substantially smaller, which is due to the strong negative correlation between equity returns and volatility shocks (the "leverage effect"). The average time-series R^2 for this multivariate regression is 65%.

It is important to note that all these betas are estimated with quite high precision. In particular, the corporate bond liquidity betas have an average t-stat of -8.6 (univariate) and -4.0 (multivariate). To put this in perspective, the average t-stat of the estimate

for expected liquidity, $E(c_j)$, is equal to 6.2. Hence, estimation error is of similar size for expected liquidity and liquidity betas.

Finally, Panel B (Table 1) reports the returns, costs and betas of high-liquidity and low-liquidity portfolios. Recall that our portfolios are first sorted on rating or EDF, and then on one of the three liquidity proxies, bond age, amount issued and volume. For each rating level or EDF quintile, we thus have a high-liquidity and low-liquidity portfolio for each liquidity proxy (except for the AAA and CCC ratings). We then report averages across all rating-based and EDF-based portfolios and across the three liquidity proxies. We see that low-liquidity portfolios have higher expected returns and higher estimated transaction costs, suggesting an effect of transaction costs on expected returns. In contrast, there is little difference in equity or volatility betas, which shows that, once we sort on rating or EDF, the liquidity sort is indeed capturing liquidity effects and not differences in market or volatility risk exposure. We also see that the liquidity betas of the low-liquidity portfolios are closer to zero than those of the high-liquidity portfolios. There is substantial variation in corporate bond liquidity betas across portfolios however: the bond market liquidity exposures range from -0.56 to about -0.11 (or -0.50 to 0.11 for the multivariate betas) across portfolios (non-tabulated). Keeping in mind that we scale the liquidity factor to have the same volatility as the equity return, this variation in liquidity betas is substantial and is expected to be informative about the presence of a liquidity risk premium.

5.2   Benchmark asset pricing results

In this subsection we focus on the cross-sectional asset pricing results with the betas of the four factors (CBLIQ, EQLIQ, EQ, and VIX) as determinants of expected returns, in addition to the average transaction costs as portfolio characteristic (equations (1)

and (2)). Specifically, we run a cross-sectional regression of the average expected excess returns on the estimated risk factor exposures and the average transaction costs. The averages and the betas are estimated over the full sample period 2005-2008. As the betas and the expected costs contain estimation noise, the standard errors of the regression are calculated using an extension of the method by Shanken (1992).[8] Notice that the regressions do not contain an intercept, which is consistent with the model in equation (2).

We first check whether there are any multicollinearity issues for the cross-sectional regression. Panel C of Table 1 presents cross-sectional correlations of the various betas that enter the cross-sectional regression, along with the expected liquidity of each portfolio. Not surprisingly, the highest correlation is between equity and volatility betas (67%). Interestingly, the equity liquidity and corporate bond liquidity betas are negatively correlated. This shows that equity liquidity and corporate bond liquidity generate quite different beta patterns.

5.2.1 Cross-sectional pricing regressions: Univariate

Table 2 presents the main results from the cross-sectional asset pricing regressions. We first present univariate regressions in specifications (1) to (5). We see that the expected cost has a positive and significant coefficient (specification (1)), and explains about 17% of the cross-sectional variation. The CBLIQ beta has a negative coefficient (specification (2)). Given the negative liquidity betas, the product of the liquidity beta and liquidity coefficient is positive. Hence we find a positive corporate bond liquidity premium when we do not include any other variables, although it has low cross-sectional explanatory power as it has a negative cross-sectional R^2. This implies that the model with priced

---

[8]The Internet Appendix of Bongaerts, de Jong and Driessen (2011) provides more details on the procedure.

corporate bond liquidity risk explains less of the expected return variation than a model with a constant term only.

Specifications (3) to (5) show that equity liquidity betas, equity betas, and volatility betas are all important determinants of expected corporate bond returns. Recall that we scaled all factors to have the same standard deviation. This allows for a straightforward comparison of the betas. Since Table 1 shows that these betas are quite comparable in size, the risk premium coefficients λ can also be compared directly. Table 2 then shows that the premia for equity risk, equity liquidity risk, and volatility risk are of similar size.

Note that most coefficients in the cross-sectional regression have quite high t-statistics. This is a direct consequence of the use of forward-looking expected returns as left-hand-side variables. These are estimated much more precisely than realized average returns.

5.2.2 Cross-sectional pricing regressions: Multivariate

Next we turn to the multivariate regressions to see whether corporate bond expected returns are affected by both expected transaction costs and liquidity risk premia (Table 2, specifications (6) to (11)). The first key result is that the expected liquidity effect remains significant and positive across all specifications. In contrast, as soon as we both include expected liquidity and corporate bond liquidity risk, the corporate bond liquidity risk coefficient becomes positive and mostly insignificant. This implies a counter-intuitive negative liquidity risk premium. The economic size of this corporate bond liquidity risk premium is negligible however. For example, when we add corporate bond liquidity risk to the model with equity risk only, the cross-sectional R^2 increases by 0.1% (47.2% versus 47.3%).

In contrast, the effect of expected liquidity is economically large even when we control

for corporate bond liquidity risk and equity risk (specification (7)): adding expected liquidity increases the R^2 from 47.3% to 65.1%. The coefficient on expected liquidity can be related to the trading frequency of bonds under some assumptions (Amihud and Mendelson (1986)). The coefficient of 1.172 in specification (7) corresponds to a turnover frequency of about 10 months.[9]

While there is no evidence for a substantial corporate bond liquidity risk premium, we do find robust evidence that equity market illiquidity shocks are priced (see the multivariate regression specifications (8)–(11)). This finding adds to the evidence that equity market liquidity risk is priced in various markets, such as private equity (Franzoni, Novak, and Phalippou (2011)) and hedge funds (Sadka (2009)). We also find that equity volatility shocks are priced. Note that the firm-value approach of Merton (1974) directly implies that volatility exposure should affect corporate bond expected returns if there is a volatility risk premium. Such a volatility risk premium has been found in the cross-section of stock returns (see Ang, Hodrick, Xing and Zhang (2006)) and for index options (see e.g. Bollerslev, Tauchen and Zhou (2009)).

The effect of expected liquidity remains positive and significant across these specifications. In contrast, the bond liquidity risk premium continues to have the "wrong" sign and is economically small. Finally, note that the estimated equity premium is always significantly positive and reasonable in size (between 2% and 5.3% per year).

In Section 6 we present a wide range of robustness checks on the asset pricing results presented here. Across all these robustness checks, we continue to find that both expected liquidity and equity market liquidity risk affect corporate bond expected returns, while exposure to corporate bond liquidity risk has an economically negligible

---

[9]Note that this should not be interpreted as the equally-weighted average turnover across bonds. Since we use transaction data, bonds that trade more often have a higher weight in our sample. Hence, the expected liquidity coefficient captures a trade-frequency weighted average of turnover across bonds.

effect.

5.2.3 Fit of expected returns and the credit-spread puzzle

Figure 2 graphs the fitted values of the risk premium according to specification (11), for the portfolios sorted on rating and liquidity proxies (Panel A) and for the sorts on EDFs and liquidity (Panel B). The graphs present the average across the three liquidity proxies per rating/EDF category and show that the expected liquidity premium, together with risk premia for equity, volatility and equity liquidity risk, explains most of the observed credit spreads, with a negligible effect of corporate bond liquidity risk.

These results shed light on the credit spread puzzle. Huang and Huang (2003) show that structural models of default risk generate credit spreads well below observed credit spread levels. We find similar results using our asset pricing approach. Equity market and volatility risk exposure explain only a part of the level of expected bond returns. In particular, equity and volatility betas of high-rated bonds are very low, so that only with extremely high equity and volatility risk premia it would be possible to explain the relatively high expected returns on these bonds. However, such high market risk premia would (i) be inconsistent with risk premia observed in for example equity markets, and (ii) imply too high expected returns on lower-rated bonds, given that these bonds have high exposure to market and volatility risk. Incorporating liquidity effects, mainly expected liquidity and equity market liquidity risk, resolves this puzzle. As shown in Figure 2, a substantial part of the expected return of high-rated bonds is due to these liquidity effects. The model provides a very good fit of expected bond returns across all portfolios, and does not underestimate the expected return on high-rated bonds. In fact, for high-rated bonds the model predicts expected returns that are slightly higher than the observed average returns. Figure 2 also shows that the economic size of the

24

corporate bond liquidity risk premium is very small across all portfolios.

5.3 Why is corporate bond liquidity risk not priced?

Our results show that expected liquidity and equity market liquidity risk have a strong effect on corporate bond prices, while corporate bond liquidity risk exposure does not. To understand these empirical findings, first note that U.S. equities are typically much more liquid than U.S. corporate bonds.[10] Then consider an investor holding both liquid assets (U.S. equities) and illiquid assets (U.S. corporate bonds). There are various reasons why investors trade these assets over time, such as rebalancing, risk-shifting, satisfying regulatory capital requirements, exogenous liquidity needs, etc. In the presence of transaction costs, the investor then faces a trade-off between having an optimally diversified portfolio and minimizing trading costs, see for example Constantinides (1986).

In Appendix C we formalize such a setup using a simple two-period asset pricing model. In this model, a mean-variance investor can invest in two assets, a low-cost asset and high-cost asset. After one period, the investor is forced to liquidate part of his portfolio due to an exogenous liquidity shock. If the difference between transaction costs of the two assets is sufficiently high, it is optimal for the investor to absorb the liquidity shock by selling only low-cost assets and avoid trading in the high-cost asset. In equilibrium, this implies that the risk of shocks to the liquidity of the high-cost asset is not priced. Only exposure to the liquidity risk of the low-cost asset is priced. Applying this result to our setting, it would imply that corporate bond liquidity risk is not priced, which is in line with our empirical findings, while exposure to equity market liquidity

---

[10]Hasbrouck (2009) has estimated transaction costs for U.S. equities using Roll's measure. Using his data, we find that for the S&P 500 stocks the average cost is 0.15% in 2005-2007, and 0.32% in 2008. In contrast, Figure 1 reports market-wide corporate bond transaction costs between 0.5% and 1% in 2005-2007 and around 2.5% in 2008. Even for the most liquid corporate bond portfolio costs are around 0.4% before the crisis and between 1% and 3% during the crisis.

risk should be priced. This is also in line with our empirical findings: in Table 2 we find that exposure to equity market illiquidity has a significant risk premium with the expected sign. Alternatively, if the difference between transaction costs of the two assets is sufficiently small, our simple model predicts that both assets are used to absorb the liquidity shock, although the low-cost asset is still used relatively more. In this case, shocks to transaction costs of both assets are priced.

This simple model can be used to generate some additional testable hypotheses. First, the model predicts higher turnover for the low-cost asset. Second, the pricing of liquidity and the turnover levels may be different in normal times versus crisis periods. To see this, consider the case where, in normal times, transaction costs on both assets are sufficiently small so that both assets will be used for trading. In this case, liquidity risk on both the low-cost asset (equity) and high-cost asset (corporate bonds in our case) is priced. However, in normal times the variation in transaction costs is quite small (see Figure 1) and hence the effect of liquidity risk on prices may be negligible. Next, consider a crisis period during which transaction costs on the high-cost asset increase substantially so that in equilibrium this asset will not be used to absorb liquidity shocks. The model then predicts that trading in the liquid asset actually increases when moving from normal times to crisis periods, while trading in the illiquid asset decreases.

To support these claims, we study turnover patterns in equity and corporate bond markets. We test two implications. First, turnover in liquid markets should be higher than turnover in illiquid markets. Second, as markets become less liquid and/or prices go down, turnover in the illiquid market decreases, while turnover in the liquid market increases.

For our 2005-2008 sample period we calculate monthly turnover for each stock in the *CRSP* database and each bond in our *TRACE* sample, by dividing monthly dollar

26

trading volume by the dollar value of the market capitalization (or amount issued in bond markets). We first focus on the value-weighted average turnover levels in both markets. We indeed find that equity market turnover is much higher than corporate bond turnover. For stocks, the value-weighted average turnover level equals 68.3% per quarter, while for corporate bonds the corresponding number is 6.6% per quarter.

Subsequently, we study the dynamics of turnover. Figure 3 graphs the time series of the value-weighted mean and the median turnover across stocks and bonds, respectively. We see that, as the crisis unfolds, equity turnover increases while corporate bond turnover goes down. To analyze these time series in more detail, we perform two sets of regressions. First, we regress monthly changes in the value-weighted average equity turnover on the equity market return and the change in Amihud's (2002) ILLIQ measure. This ILLIQ measure is calculated for each stock, and we then use the change in the value-weighted average ILLIQ across stocks as regressor. Results in Table 3 show that a decrease in liquidity or decrease in the stock market index level imply an increase in equity market turnover. Both effects are significant in these univariate regressions. Given that the change in ILLIQ and the equity market return are strongly negatively correlated, a multivariate regression suffers from multicollinearity issues, as shown in Table 3. These results show that equity market turnover increases in bad, illiquid times. As discussed above, this may be because investors are essentially forced to trade despite the higher transaction costs.

The second regression focuses on corporate bond turnover. In this case, we perform a panel regression of the monthly change in individual bond turnover on the change in the market-wide level of corporate bond transaction costs, as constructed in this paper, and the change in the average credit spread across all bonds. The results in Table 4 show that corporate bond turnover decreases when the market becomes less liquid and when credit spreads increase. We also investigate a flight to quality effect. We analyze

27

this by interacting the liquidity and credit spread variables with a dummy that equals one when a bond is rated AAA or AA. The results of the turnover regression in Table 4 show that the turnover decrease in bad times is indeed larger for low-rated (and hence less liquid) corporate bonds.

In sum, these results shows that stocks have much higher turnover than corporate bonds, and that this difference is larger when prices go down and when markets are overall less liquid. Hence, in line with our hypothesis, when transaction costs increase investors prefer to trade the relatively more liquid assets and avoid trading in less liquid assets.

6 Robustness checks

In this section we present a wide range of robustness checks on the empirical findings presented in Section 5.

6.1 Sorting on liquidity betas

The results above indicate that the effect of corporate bond liquidity risk on corporate bond prices is economically small. However, this finding may be caused by a lack of cross-sectional variation in liquidity betas making estimation of a risk premium difficult. Therefore, similar to Pastor and Stambaugh (2003) we now construct test portfolios that are also sorted on corporate bond liquidity betas. This requires liquidity beta estimates at the individual bond level. The challenge here is that many bonds do not trade very frequently and estimation of individual betas for these assets is problematic. Moreover, beta estimates for individual instruments can be rather unstable and sensitive to outliers.

To deal with these issues, we use a Bayesian approach. Our liquidity beta for each bond is calculated as a weighted average of the direct regression estimate (obtained by regressing individual bond returns on the liquidity factor) and a portfolio-based beta. This portfolio beta is obtained by using the liquidity beta of the portfolios to which the bond was assigned in the analysis above. The liquidity beta of these portfolios is our "best guess" (or, in Bayesian terms, "prior") of the true liquidity beta of the bond in case insufficient trading data for this bond is available. The more precisely we can estimate the individual bond liquidity beta from transaction data, the less weight we want to give to these portfolio betas. This is exactly what our Bayesian solution achieves. Appendix D provides details of the procedure.

Table 5 shows the cross-sectional results when we add these liquidity-beta portfolios to our cross-section of portfolios. Comparing these results with the benchmark results, we see that adding liquidity-beta portfolios hardly affects the estimates for the risk premia and expected liquidity. In particular, the liquidity risk premium still has the "wrong" sign and its economic impact remains very small. Even when we only use the liquidity-beta portfolios for the cross-sectional estimation, we find a small coefficient with the "wrong" sign for the liquidity risk premium (non-tabulated).

6.2   Other liquidity measures

Liquidity has many dimensions, and therefore we check whether our results our robust to the use of different liquidity measures. First, we consider an alternative measure of corporate bond market liquidity risk. Feldhütter (2011) and Dick-Nielsen, Feldhütter, and Lando (2011) propose to measure the bid-ask spread using imputed roundtrip trades (IRT). The idea of this measure is to find trades of the same size in a given bond within a given time interval. In many cases, these trades represent a dealer acting as market

maker, who buys from and sells to end users and collects a bid-ask spread as fee. For each bond, we thus take the difference between the highest price and the lowest price of trades with the same trade size on a given day as an estimate of the bid-ask spread. We obtain an aggregate measure by averaging this bid-ask spread across all bonds, and in Figure 1 we compare this IRT measure with the Roll measure used in the benchmark analysis. Clearly, both measures are highly correlated. Hence, two very different liquidity measures generate very similar liquidity factors.

We then estimate the asset pricing model using the IRT-based liquidity factor, using an AR(2) model to model shocks to this liquidity measure. Table 5 shows that the results are extremely close to the benchmark results.

We also use an alternative equity market liquidity measure, replacing Amihud's ILLIQ measure by the measure used by Pastor and Stambaugh (PS, 2003).[11] Table 5 shows that this hardly affects the results. In particular, the corporate bond liquidity risk premium continues to have the "wrong" sign, and its economic significance remains small.

6.3 Credit risk as a portfolio characteristic

Our analysis focuses on explaining expected excess corporate bond returns. As discussed above, these expected returns are constructed by correcting credit spreads for expected default losses. Still, even though we control for expected default losses, it is possible that the liquidity effects pick up some missing credit risk effects.[12] To analyze this, we add the *Moody's-KMV* default probability estimates as a portfolio characteristic to

---

[11]We construct the PS measure for each stock on each day in our sample using a rolling window of 22 trading days, and take an equally weighted average. We then estimate an AR(2) model for this average at the weekly frequency to obtain liquidity shocks.

[12]He and Xiong (2010) propose a theoretical model in which rollover risk generates an endogenous relation between liquidity and credit risk.

the second-step cross-sectional regression. This variable is constructed by averaging the default probability estimates over time and across all firms in a given portfolio. Table 5 shows that, even though the default probability variable is significant, its presence does not change the results on liquidity and liquidity risk.

6.4 Pre-crisis sample

Market liquidity deteriorates quickly after the Lehman default in September 2008. Hence, it is interesting to see to what extent our results are driven by the large liquidity shocks during this crisis period. We therefore rerun the asset pricing regressions using a pre-crisis sample, using data up to August 2008. We first check whether the first-step liquidity betas are still significant, since it may be that the strong significance of these betas in the full sample (as reported in Table 1) was purely driven by the crisis period. We find that this is not the case: using the pre-crisis sample, the corporate bond liquidity betas have average t-statistics of -7.5 (univariate) and -5.7 (multivariate), very similar to the full-sample results. Table 5 shows the results of the cross-sectional regression for the pre-crisis sample. Again, we obtain very similar results. In this case, the corporate bond liquidity risk premium does have the "correct" sign, but it is both statistically and economically insignificant.

6.5 Intercept in the pricing model

So far, we have not included a constant term in second-step regression for the asset pricing model, since the model predicts a zero constant term. Hence, a simple specification test is to add an intercept to the second-step regression and see whether it is significant. Table 5 shows that this is not the case. The intercept is equal to -0.012% (in annual terms), and highly insignificant with a t-stat of -0.05. Also, the coefficients for expected

31

liquidity and the risk premia are essentially unchanged when adding an intercept.

6.6 Time variation in betas and premia

Next we analyze whether there is any variation in the betas and risk premia over our sample period. Our forward-looking measures of expected returns and transaction cost estimates are available on a weekly basis. However, the right-hand-side betas cannot be calculated with only one week of data. Therefore, we use a two-sided kernel to estimate the beta at each point in time. The kernel function gives higher weight to observations close to the measurement date (see Ang and Kristensen (2009) for more details) and is two-sided, thus using observations before and after the measurement date. For simplicity, we use a "tent shaped" kernel function, with linearly decreasing weights for observations further away from the measurement date, and zero weights for observations that are more than 52 weeks away.[13] The second-step equations are then estimated using four-week rolling windows. Of course, our sample period is short and hence it is difficult to precisely assess the time variation in betas and risk premia. Therefore, this analysis should be considered as explorative.

Figure 4 graphs the results for the model including the equity, equity liquidity and corporate bond liquidity betas, and expected costs.[14] The left-hand panels of Figure 4 show the estimated betas and transaction costs. For each week, the graphs show the cross-sectional average over all portfolios. We see that equity betas increase from the start of the financial crisis in mid-2007. The transaction costs also increase from that period onwards, and increase very substantially around the Lehman collapse in September 2008. The middle panels show the estimated coefficients from the weekly

---

[13]For dates close to the start or end date of the sample a truncated tent shape is used.

[14]To keep the graphs readable, we exclude the volatility premium because it is highly correlated with the equity premium. Results are very similar when we include volatility risk.

second step regressions. The coefficient for expected liquidity is remarkably stable over time, while the liquidity premia are less stable. The right-hand panels show the implied risk premia (top-right panel), and the expected liquidity premium (bottom-right panel), obtained by multiplying the betas with the estimated coefficients from the second step regressions. These graphs clearly show an increase in the equity risk premium from mid-2007. The corporate bond liquidity risk premium is close to zero and unstable, though. The expected liquidity premium increases from about 60 basis points to around 2.5 percent for the average portfolio.

The average of the weekly estimates are reported in Table 5 (now including the volatility risk factor).[15] The average estimated equity premium is around 5.4 percent, while the equity liquidity risk premium is a bit smaller than in the full sample, perhaps due to larger estimation error for the kernel-based betas. The corporate bond liquidity risk premium has the correct sign, but its effect is again economically small. The estimated coefficient for expected liquidity is in the same order of magnitude as in the full sample estimates. In sum, this explorative time-varying analysis supports our main finding that expected liquidity has a strong effect on bond prices, while the effect of corporate bond liquidity risk is small.

6.7 Liquidity-CAPM approach

In the analysis above we focused on one liquidity risk exposure, the covariance between portfolio returns and market-wide liquidity shocks. However, the liquidity CAPM of Acharya and Pedersen (AP, 2005) suggests that other liquidity risk covariances may also matter. We therefore now focus on the implications of the AP model, which implies

---

[15]Calculating standard errors that incorporate the Shanken (1996) correction is not trivial in this setting since we use four-week windows for expected returns and costs, and 52-week windows to get betas.

that

$$E(r_{it}) = \zeta E(c_{it}) + \phi \frac{Cov(r_{it} - c_{it}, r_{mt} - c_{mt})}{Var(r_{mt} - c_{mt})}, \tag{10}$$

where $r_{mt} - c_{mt}$ is the average (value weighted) net return on the corporate bond market.

One possible approach would be to use this AP model for the corporate bond market, in isolation from other markets. This is not very realistic as corporate bond returns are known to be strongly correlated with equity returns and also with volatility changes. Bongaerts, de Jong and Driessen (BDD, 2011) build a formal model of liquidity and liquidity risk pricing in markets with hedging pressure and (potentially) short selling, where asset returns can be partly hedged by other (so-called benchmark) assets. In the absence of short-selling, the equilibrium pricing equation is similar to the AP model, but the returns and costs are orthogonalized for their covariance with a set of benchmark assets. Formally, we have

$$E(\widehat{r}_{it}) = \zeta E(\widehat{c}_{it}) + \phi \frac{Cov(\widehat{r}_{it} - \widehat{c}_{it}, \widehat{r}_{mt} - \widehat{c}_{mt})}{Var(\widehat{r}_{mt} - \widehat{c}_{mt})}, \tag{11}$$

with

$$\widehat{r}_{it} = r_{it} - \beta_i^{r\prime} r_{b,t}, \quad \beta_i^r = Var(r_{b,t})^{-1} Cov(r_{it}, r_{b,t}),$$

and

$$\widehat{c}_{it} = c_{it} - E_{t-1}(c_{it}) - \beta_i^{c\prime} r_{b,t}, \quad \beta_i^c = Var(r_{b,t})^{-1} Cov(c_{it} - E_{t-1}(c_{it}), r_{b,t}),$$

and with r_{bt} the vector with returns on the benchmark assets.[16] Similar to AP, we incorporate that transaction costs are persistent over time by focusing on innovations $c_{it} - E_{t-1}(c_{it})$. The 'market' return and cost factors $\widehat{r}_m$ and $\widehat{c}_m$ are value weighted

---

[16]This model is simpler than the original BDD model: it assumes there are no non-traded risk factors that correlate with corporate bond returns.

averages of the individual returns and costs.[17] Empirically, we allow for two benchmark assets, the equity market index and the VIX index.

The empirical model proceeds in two steps. In the first step, corporate bond excess returns and corporate bond transaction cost innovations are regressed on a set of benchmark assets. This produces estimates of the exposure coefficients β_i^r and β_i^c. We also calculate the elements of the last term in equation (11)

$$
\begin{aligned}
\beta_i^{rr} &= \frac{Cov(\widehat{r}_{it}, \widehat{r}_{mt})}{Var(\widehat{r}_{mt} - \widehat{c}_{mt})} \\
\beta_i^{rc} &= \frac{Cov(\widehat{r}_{it}, \widehat{c}_{mt})}{Var(\widehat{r}_{mt} - \widehat{c}_{mt})} \\
\beta_i^{cr} &= \frac{Cov(\widehat{c}_{it}, \widehat{r}_{mt})}{Var(\widehat{r}_{mt} - \widehat{c}_{mt})} \\
\beta_i^{cc} &= \frac{Cov(\widehat{c}_{it}, \widehat{c}_{mt})}{Var(\widehat{r}_{mt} - \widehat{c}_{mt})}.
\end{aligned}
\tag{12}
$$

The expected returns then follow from

$$
E(r_{it}) - E(r_{bt})'\beta_i^r = \zeta(E(c_{it}) - E(r_{bt})'\beta_i^c) + \phi(\beta_i^{rr} - \beta_i^{rc} - \beta_i^{cr} + \beta_i^{cc}).
\tag{13}
$$

In practice, we do not know the expected return on the benchmark assets $E(r_{bt})$ and treat it as a parameter vector (λ) to be estimated. This model is linear in all parameters except ζ. However, if we take a preliminary estimate $\zeta = \zeta_0$ to construct $\beta_i = \beta_i^r - \zeta_0\beta_i^c$, the model is linear and can be estimated by OLS. In the empirical work, we set $\zeta_0 = 1.2$ (implying a turnover rate of about 10 months), which roughly corresponds to the average estimate across the specifications in Table 2. Following AP, we also allow each component of the final covariance term to have a separate impact the expected return. The final

---

[17]Another possible approach would be to apply the Acharya-Pedersen model to both the entire equity market and the corporate bond market. This would assume perfect integration of the two markets, and require a liquidity factor that combines equity and corporate bond market liquidity.

model thus is

$$\widehat{E}(r_{it}) = \lambda'(\beta_i^r - \zeta_0\beta_i^c) + \zeta E(c_{it}) + \phi_1\beta_i^{rr} + \phi_2\beta_i^{rc} + \phi_3\beta_i^{cr} + \phi_4\beta_i^{cc} + \alpha_i. \qquad (14)$$

The residuals α_i can again be interpreted as pricing errors.

Table 6 presents summary statistics on the different betas in this model. Most notably, we see that the different liquidity betas have the expected sign: β^{rc} and β^{cr} are negative on average: low returns coincide with higher transaction costs. Also, the average β^{cc} is positive, suggesting the presence of commonality in liquidity.

Turning to the cross-sectional regressions in Table 7, we estimate a variety of specifications for the BDD model in equation (14). All specifications include the equity beta, volatility beta and the expected transaction costs. In addition, either the 'net' beta or all or some of its components (β^{rr}, β^{rc}, β^{cr}, and β^{cc}) are in the regression. The equity beta and the transaction cost have positive and significant coefficients in every specification, and the volatility risk premium is significantly negative in all specifications, in line with the results in Section 5. The magnitudes of the equity risk premium and the transaction cost premium are fairly stable across specifications. Together, expected liquidity and the equity and volatility risk premia explain 73% of the cross-sectional variation in expected returns.

Without additional variables, the estimated equity risk premium is 3.9% per year, and the expected liquidity coefficient is 0.932. Adding the corporate bond market risk premium (orthogonalized for the equity risk) in Table 7, specification (2) gives similar results with exposure to β^{rr} not significant. Hence, once we control for equity market and volatility risk, exposure to the market-wide corporate bond return does not have explanatory power.

Next we add various liquidity risk factors. In specifications (3) and (4), the exposure of bond returns to corporate bond market transaction costs β^{rc} is added. This factor is significant and has a counter-intuitive positive sign, implying a negative liquidity risk premium (as β^{rc} is negative for every portfolio), but the economic effect is small as the cross-sectional R^2 increases only marginally. Even when we add all components of liquidity risk with separate coefficients in specification (5) the cross-sectional R^2 does not increase substantially, and multicollinearity across the different liquidity betas leads to "wrongly-signed" coefficients on some of the liquidity betas. When we impose the restriction that all coefficients on the liquidity betas are the same ($\beta^{other} = -\beta_i^{rc} - \beta_i^{cr} + \beta_i^{cc}$, specification (6)), we again find the "wrong" sign for this liquidity risk premium, and again the effect is economically small. Finally, when we include the total 'net' beta as the regressor in equation (14), $\beta_i^{net} = \beta_i^{rr} - \beta_i^{rc} - \beta_i^{cr} + \beta_i^{cc}$, this net beta has a negative but insignificant coefficient (specification (7)). In sum, even when we allow for various forms of liquidity risk exposure, we do not find that liquidity risk is priced in the cross-section of corporate bond portfolios. The effect of expected liquidity is remarkably constant over all specifications, though, with a coefficient around one.

7   Conclusion

This paper explores the asset pricing implications of expected liquidity and liquidity risk for expected corporate bond returns. We measure corporate bond liquidity using a Bayesian estimation of Roll's effective cost model. We then construct liquidity levels and liquidity innovations for a set of corporate bond portfolios. Several asset pricing models, including Acharya and Pedersen's liquidity CAPM, are then estimated using the cross-section of corporate bond portfolios. Overall, we find a strong effect of expected liquidity and equity market liquidity risk on expected corporate bond returns, while there is little

evidence that corporate bond liquidity risk covariances explain expected corporate bond returns, even during the recent financial crisis. We show that incorporating liquidity effects goes a long way in explaining the high returns on high-rated corporate bonds (the "credit spread puzzle"). We also find that equity risk and volatility risk (exposure to VIX shocks) are priced in the cross-section of corporate bonds. Our findings are consistent with a setting in which investors strategically choose to use the assets with low liquidity costs to rebalance their portfolio and satisfy their exogenous liquidity needs.

A   Data filters

We apply several filters to our dataset to remove bonds with special features and to remove erroneous entries. Our filters are very similar to those employed in Bongaerts, Cremers and Goetzmann (2010). We remove all trades that include commission, that have a settlement period of more than 5 days, and all trades that are canceled or reversed. Trade volumes are truncated by the system and we replace truncated trade volumes by their respective truncation barrier ($5 million for Investment Grade and $1 million for High Yield). We remove all trades for which we have a negative reported yield, since these will be mainly driven by implicit option premia in the yield. We use *Bloomberg* to match the trades to bond characteristics and S&P ratings using CUSIPs. We discard all bonds with convertibility options, that are putable, that have a non-fixed coupon, that are subordinated, secured or guaranteed. We keep callable bonds because they comprise a large part of the sample. Removing these callable bonds would substantially reduce the sample size and the precision of our transaction costs estimates. We discard all zero-coupon bonds. We also remove trades with a settlement date later than or equal to the maturity date. Furthermore, we found several duplicate records, resulting from both parties involved in a trade reporting to the system. We filter out these trades by consecutively sorting on bond, date and volume and removing identical consecutive records. Moreover some of the yield changes are unrealistically high. Therefore, we remove trades with yield *changes* of more than 1000 basis points (about 0.15% of our trades). We also exclude all trades with a volume lower than $10,000 since very small trades may have substantially higher transaction costs. As Ford and GM together were responsible for more than 10% of all corporate bond trading, we excluded these issuers to avoid portfolios to be completely driven by individual companies.

B   Gibbs sampler for the Roll model

Estimation of the coefficients of the Roll model is done by means of the Gibbs sampling method developed by Hasbrouck (2009), combined with the repeat sales methodology. In the Gibbs sampler, the parameters c and σ_u^2 and the latent series ΔM_t and r are simulated step-by-step from their Bayesian posterior distributions. In every step, one set of parameters or latent variables is simulated, conditional on the values of the other parameters and latent variables from the previous simulation round. Each step then is a relatively simple application of Bayesian regression.

Simulating q

The first step in each iteration of the Gibbs sampler is the simulation of the trade indicators q. In Hasbrouck's model, these can take only two values, $+1$ and -1. The prior is equal probabilities, i.e. $Pr[q_{i,t_{ik}} = 1] = 1/2$. After observing p, the posterior odds are

$$\frac{Pr[q_{i,t_{ik}} = 1]}{Pr[q_{i,t_{ik}} = -1]} = \frac{f(e_{t_{ik}}|q_{i,t_{ik}} = 1)f(e_{t_{i,k+1}}|q_{i,t_{ik}} = 1)}{f(e_{t_{ik}}|q_{i,t_{ik}} = -1)f(e_{t_{i,k+1}}|q_{i,t_{ik}} = -1)} \tag{15}$$

where

$$f(e_{t_{ik}}|q_{i,t_{ik}} = q) = \phi\left(\frac{CS_{i,t_{ik}} - CS_{i,t_{i,k-1}} - \sum_{s=t_{i,k-1}+1}^{t_{ik}} z_i \Delta M_s - c_{t_{ik}} q + c_{t_{i,k-1}} q_{i,t_{i,k-1}}}{\sigma_u^2 \sum_{s=t_{i,k-1}+1}^{t_{i,k}} v_s^2}\right) \tag{16}$$

and

$$f(e_{t_{i,k+1}}|q_{i,t_{ik}} = q) = \phi\left(\frac{CS_{i,t_{i,k+1}} - CS_{i,t_{ik}} - \sum_{s=t_{ik}+1}^{t_{i,k+1}} z_i \Delta M_s - c_{t_{i,k+1}} q_{i,t_{i,k+1}} + c_{t_{ik}} q}{\sigma_u^2 \sum_{s=t_{i,k-1}+1}^{t_{i,k}} v_s^2}\right) \tag{17}$$

From the posterior odds ratio, the posterior probabilities for $q = \{1, -1\}$ are easily calculated.

Simulating c

The liquidity cost of a particular week $w = t_{i,k}$ realized in a particular trade k shows up in two credit-spread equations:

$$CS_{i,t_{ik}} - CS_{i,t_{i,k-1}} - \sum_{s=t_{i,k-1}+1}^{t_{ik}} z_i \Delta M_s = c_{t_{ik}} q_{i,t_{ik}} - c_{t_{i,k-1}} q_{i,t_{i,k-1}} + e_{it_{ik}}, \qquad (18)$$

$$CS_{i,t_{i,k+1}} - CS_{i,t_{i,k}} - \sum_{s=t_{i,k}+1}^{t_{ik+1}} z_i \Delta M_s = c_{t_{ik+1}} q_{i,t_{ik+1}} - c_{t_{i,k}} q_{i,t_{i,k}} + e_{it_{ik+1}} \qquad (19)$$

The posterior mean for c_w is found from a linear regression of the two return equations stacked on top of each other.

Let us first work out equation (18). If both t_{ik} and $t_{i,k-1}$ fall in the same week w_{ik}, the equation is

$$CS_{i,t_{ik}} - CS_{i,t_{i,k-1}} - \sum_{s=t_{i,k-1}+1}^{t_{ik}} z_i \Delta M_s = c_{w_{ik}} (q_{i,t_{ik}} - q_{i,t_{i,k-1}}) + e_{it_{ik}} \qquad (20)$$

If $t_{i,k-1}$ happens to be in an earlier week, we write

$$CS_{i,t_{ik}} - CS_{i,t_{i,k-1}} - \sum_{s=t_{i,k-1}+1}^{t_{ik}} z_i \Delta M_s + \widehat{c}_{w_{i,k-1}} q_{i,t_{i,k-1}} = c_{w_{ik}} q_{i,t_{ik}} + e_{it_{ik}} \qquad (21)$$

where $\widehat{c}_{w_{i,k-1}}$ is the most recent simulation of the earlier week's transaction cost.

Working out equation (19), we get again that if both t_{ik} and $t_{i,k+1}$ fall in the same

week w_{ik}, the equation is

$$CS_{i,t_{ik+1}} - CS_{i,t_{i,k}} - \sum_{s=t_{i,k}+1}^{t_{ik+1}} z_i \Delta M_s = c_{w_{ik}}(q_{i,t_{ik+1}} - q_{i,t_{i,k}}) + e_{it_{ik+1}} \qquad (22)$$

If $t_{i,k+1}$ happens to be in a later week, we write

$$CS_{i,t_{ik+1}} - CS_{i,t_{i,k}} - \sum_{s=t_{i,k}+1}^{t_{ik+1}} z_i \Delta M_s - \widehat{c}_{w_{i,k+1}} q_{i,t_{i,k+1}} = -c_{w_{ik}} q_{i,t_{ik}} + e_{it_{ik+1}} \qquad (23)$$

where $\widehat{c}_{w_{i,k+1}}$ is the simulation of the subsequent week's transaction cost from the previous iteration. Estimation of the posterior mean of c_w is then done by stacking these equations. Formally, we estimate $y = Xc_w + e$ with $y = (y^{cont}, y^{fut})'$ and

$$y_{ik}^{cont} = CS_{i,t_{ik}} - CS_{i,t_{i,k-1}} - \sum_{s=t_{i,k-1}+1}^{t_{ik}} z_i \Delta M_s + (1 - I_{w_{ik}=w_{i,k-1}})\widehat{c}_{w_{i,k-1}} q_{i,t_{i,k-1}}) \qquad (24)$$

$$y_{ik}^{fut} = CS_{i,t_{ik+1}} - CS_{i,t_{i,k}} - \sum_{s=t_{i,k}+1}^{t_{ik+1}} z_i \Delta M_s - (1 - I_{w_{ik}=w_{i,k+1}})\widehat{c}_{w_{i,k+1}} q_{i,t_{i,k+1}}) \qquad (25)$$

and $x = (x^{cont}, x^{fut})'$ with

$$x_{ik}^{cont} = q_{i,t_{ik}} - I_{w_{ik}=w_{i,k-1}} q_{i,t_{i,k-1}} \qquad (26)$$

$$x_{ik}^{fut} = -(q_{i,t_{ik}} - I_{w_{ik}=w_{i,k-1}} q_{i,t_{i,k+1}}) \qquad (27)$$

for all $w_{ik} = w$ and is estimated using all data in that week. Notice that the error term e_{it} is a sum $u_{it}v_t$ for $t = t_{i,k-1}$ to $t = t_{i,k}$ and therefore heteroskedastic.

In the model, the transaction costs c_w are identified up to their sign, see Hasbrouck (2009), and we therefore impose $c_w > 0$. Hence, the posterior distribution of

c_w is a truncated normal

$$c_w \sim N^+((X'\Sigma_e^{-1}X)^{-1}X'\Sigma_e^{-1}y, \ (X'\Sigma_e^{-1}X)^{-1}) \tag{28}$$

with Σ_e a diagonal matrix with elements $\sigma_u^2 \sum_{s=t_{i,k-1}+1}^{t_{i,k}} v_s^2$.

Simulating ΔM

The most complex step is the simulation of the latent portfolio-level changes in credit spreads ΔM_t. This step is absent in Hasbrouck's model but necessary here as Δm consists of two components (simulating u is not necessary as it follows immediately from the observed values of CS and the simulated values of q, c and ΔM). We draw ΔM from a normal distribution with mean $\widehat{\Delta M}$ and variance $\widehat{V}$, where $\widehat{\Delta M}$ is the OLS estimate of a repeat sales regression

$$y = X\Delta M + e \tag{29}$$

with the matrixes y and X have rows

$$y_{ik} = CS_{i,t_{ik}} - CS_{i,t_{i,k-1}} - c_{t_{ik}} q_{i,t_{ik}} + c_{t_{i,k-1}} q_{i,t_{i,k-1}} \tag{30}$$

and

$$x_{ik} = (0'..z_{ik}\iota'..0') \tag{31}$$

for $k = 1, .., K(i)$ and $i = 1, .., N$ stacked, where $K(i)$ denotes the total number of transactions for bond i and N is the number of bonds allocated to the portfolio. ι is a vector of ones with length $t_{ik} - t_{i,k-1}$. The OLS estimator then is $\widehat{\Delta M} = (X'X)^{-1}X'y$ with variance $\widehat{V} = (X'X)^{-1}X'\Sigma_e X(X'X)^{-1}$. We neglect any serial correlation in credit spread changes, and thus take the diagonal of $\widehat{V}$ to draw ΔM. This procedure occasionally has

'gaps' i.e. periods with no or too few transactions. In such case, adjoining periods are clustered and the procedure estimates the cumulative return over the clustered periods.

C Model

In this appendix we derive a simple asset pricing model that helps to explain why corporate bond liquidity risk is not priced, while equity market liquidity risk is priced. The model has two assets, a liquid asset a (equity) with low transaction costs c_a and an illiquid asset b (corporate bond) with high costs c_b. We assume two investment periods, starting at time $t = 0$ and $t = 1$ and ending at $t = 2$. At $t = 0$ and $t = 2$ the investor can trade both assets without costs and borrow at the risk-free rate, while at $t = 1$ trading involves costs c_a and c_b, respectively, and borrowing is not possible. Hence, $t = 1$ represents a liquidity crisis state.

We first study the decision problem at time $t = 1$. Let N_a and N_b be the positions in assets a and b at time $t = 0$ (in terms of number of assets). The prices of the assets are normalized to 1 a time 1. We assume that the investor faces a liquidity problem at $t = 1$ and is forced to liquidate an amount D of his risky asset holdings. This can be done by selling either asset a, asset b or both (it is assumed that the amount D to be liquidated is in excess of any cash holdings). The numbers of shares sold are denoted by $\Delta N_a \geq 0$ and $\Delta N_b \geq 0$, where the value of the amounts sold, after transaction costs, must add up to D, so that $\Delta N_a(1 - c_a) + \Delta N_b(1 - c_b) = D$. The wealth at time $t = 2$ then is given by

$$W_2 = (N_a - \Delta N_a)R_a + (N_b - \Delta N_b)R_b, \tag{32}$$

where R_a and R_b are the gross returns on the assets in the second period. Substituting

the restriction $\Delta N_a(1 - c_a) + \Delta N_b(1 - c_b) = D$ we can write this as

$$W_2 = N_a R_a + N_b R_b + \Delta N_a \left(\frac{1 - c_a}{1 - c_b} R_b - R_a \right) - \frac{D}{1 - c_b} R_b. \qquad (33)$$

We now linearize this equation around $c_a = 0$, $c_b = 0$, and $R_b = 1$ to obtain

$$W_2 = N_a R_a + N_b R_b + \Delta N_a (R_b - R_a + c_b - c_a) - D R_b - D c_b. \qquad (34)$$

Applying a mean-variance optimization at time $t = 1$ with respect to ΔN_a, taking c_a and c_b as known, we find the optimal amount to be sold for asset a. If there is an interior solution with $\Delta N_a > 0$ and $\Delta N_b > 0$, the first order condition is

$$(\mu_b - \mu_a + c_b - c_a) - \alpha Var(R_b - R_a)\Delta N_a - \alpha Cov(N_a R_a + (N_b - D)R_b, R_b - R_a) = 0, \quad (35)$$

which gives the optimal amount ΔN_a to be sold

$$\Delta N_a = \frac{\mu_b - \mu_a + c_b - c_a}{\alpha Var(R_b - R_a)} + \frac{Cov(N_a R_a + (N_b - D)R_b, R_b - R_a)}{Var(R_b - R_a)}. \qquad (36)$$

This amount is increasing in the difference in transaction costs between asset b and asset a. However, if the cost difference becomes very large, at some point the optimal ΔN_a will exceed the gross amount to be liquidated $D/(1 - c_a)$. In that case, a corner solution $\Delta N_a = D/(1 - c_a)$ and $\Delta N_b = 0$ will be optimal.

Now we turn to the optimization problem at time $t = 0$. Suppose it is known in advance to the investor that in the case of a forced liquidation of assets, there will be a corner solution with $\Delta N_a = D/(1 - c_a)$. Then we can find the $t = 2$ wealth as seen

from $t = 0$ as (linearized around $c_a = 0$ and $R_a = 1$)

$$W_2 = \widetilde{N}_a \widetilde{R}_a + \widetilde{N}_b \widetilde{R}_b - D - Dc_a, \tag{37}$$

where $\widetilde{R}_a$ and $\widetilde{R}_b$ denote the two-period returns, i.e. the return from time 0 to time 2, and $\widetilde{N}_a$ and $\widetilde{N}_b$ denote amount invested in the assets at time 0, i.e. $\widetilde{N}_a = N_a P_a$ with P_a the price of asset a at $t = 0$. A simple mean-variance problem, without leverage constraints, produces the optimal time 0 investments

$$\widetilde{N} = \alpha^{-1} \Sigma^{-1} (\mu - r_f) + D\Sigma^{-1} Cov(\widetilde{R}, c_a), \tag{38}$$

with $\widetilde{N} = (\widetilde{N}_a, \widetilde{N}_b)'$, $\widetilde{R} = (\widetilde{R}_a, \widetilde{R}_b)'$ and $\mu = E(\widetilde{R})$ and $\Sigma = Var(\widetilde{R})$.

Now consider the equilibrium implications in a setting with several investors, which may differ in terms of the size of the liquidity shock D. The equilibrium depends on which traders are forced to liquidate assets and how many end up at the corner solution $\Delta N_a = D$. In order to provide intuition for the outcome, assume that all traders are at this corner solution, where each investor has to liquidate an amount D_i. In equilibrium, then, with common α, μ and Σ we find

$$\overline{\widetilde{N}} = \alpha^{-1} \Sigma^{-1} (\mu - r_f) + \bar{D}\Sigma^{-1} Cov(\widetilde{R}, c_a), \tag{39}$$

where $\bar{D}$ is the average of the individual amounts to be liquidated. The equilibrium risk premia then are

$$\mu - r_f \iota = \alpha \Sigma \overline{\widetilde{N}} - \alpha \bar{D} Cov(\widetilde{R}, c_a). \tag{40}$$

The risk premia are the sum of a usual market risk premium, plus a liquidity premium for the covariance between the asset returns and the transaction cost on the (more liquid)

asset c_a. We see that there is no risk premium for correlation with the transaction cost of the less liquid asset, c_b. Notice that μ_b may contain a liquidity risk premium for the covariance between the illiquid asset return $\widetilde{R}_b$ and the transaction costs on the liquid asset c_a, but there is no liquidity risk premium for the covariance between the illiquid asset return $\widetilde{R}_b$ and its own transaction cost c_b.

D Liquidity beta sorts

In this appendix we provide details about the liquidity beta portfolio sorts. We estimate the portfolio liquidity betas and their standard errors from a regression of bond portfolio excess returns on market liquidity innovations for all double sorted portfolios across quality (rating and EDF) and liquidity (age, issue size, trading volume) as well as the AAA and CCC rated portfolio. We do this on a weekly basis using a one-year rolling window. For each portfolio we then create quarterly betas and standard errors by averaging the betas over all weeks in the quarter and calculating the appropriate covariances. Next we average for each bond-quarter the betas and squared standard errors[18] across all portfolios in which that bond was contained.

For each bond, we also estimate the direct liquidity beta. To this end, we estimate a beta and standard error from the univariate regression of the individual bond excess returns on corporate bond market liquidity innovations. We again do this on a weekly basis using the last trade available every week on a one-year rolling window, where we require at least 25 observations and where the smallest and largest observation are winsorized.

---

[18]This implicitly implies a correlation of one between beta estimates of different portfolios; betas are typically highly correlated across liquidity sorts. If anything, this would put too little weight on the portfolio betas.

When constructing the portfolios for a given quarter, say Q2 2006, the one-year rolling window used to estimate these betas includes this quarter, hence we use data from Q3-Q4 2005 and Q1-Q2 2006 in this example. We thus essentially use a mixture of the pre-ranking and post-ranking betas to form portfolios. This has the advantage that it generates more variation in the liquidity betas used to estimate the asset pricing model. The disadvantage is that these portfolios are not ex-ante tradable portfolios, but this is not an issue for estimating and testing asset pricing models.[19]

Having obtained the portfolio beta and the direct beta we can now use the standard Bayesian formula to calculate our posterior beta

$$\beta_{liq}^{post} = \frac{var(\hat{\beta}_{liq}^{port})^{-1}\hat{\beta}_{liq}^{port} + var(\hat{\beta}_{liq}^{direct})^{-1}\hat{\beta}_{liq}^{direct}}{var(\hat{\beta}_{liq}^{port})^{-1} + var(\hat{\beta}_{liq}^{direct})^{-1}}. \tag{41}$$

Portfolio double sorts are then conducted again as before using a sequential sort, first on credit quality (rating or EDF) and then on liquidity beta.

---

[19]When we use pre-ranking betas, we find very similar results (not reported).

References

Acharya, Viral, and Lasse Pedersen, 2005, Asset pricing with liquidity risk, *Journal of Financial Economics* 77, 375–410.

Acharya, Viral, Yakov Amihud, and Sreedhar Bharath, 2010, Liquidity risk of corporate bond returns, Working paper, NYU Stern.

Amihud, Yakov and Haim Mendelson, 1986, Asset pricing and the bid-ask spread, *Journal of Financial Economics* 17, 223-249.

Amihud, Yakov, Haim Mendelson and Lasse Pedersen, 2005, Liquidity and asset prices, Foundations and Trends in Finance, Vol. 1, No. 4.

Ang, Andrew, and Dennis Kristensen, 2011, Testing Conditional Factor Models, Working paper, Columbia University.

Ang, Andrew, Robert Hodrick, Yuhang Xing, and Xiaoyong Zhang,.2006, The cross-section of volatility and expected Returns. *Journal of Finance*, 61: 259–299.

Asparouhova, Elena, Hendrik Bessembinder, and Ivalina Kalcheva, 2010, Liquidity biases in asset pricing tests, *Journal of Financial Economics*, 96: 215-237.

Bao, Jack, Jun Pan, and Jiang Wang, 2010, The illiquidity of corporate bonds, *Journal of Finance*, forthcoming.

Bollerslev, Tim, George Tauchen, and Hao Zhou, 2009, Expected stock returns and variance risk premia, *Review of Financial Studies*, 22: 4463-4492.

Bongaerts, Dion, Martijn Cremers, and William Goetzmann, 2010, Tiebreaker: certification and multiple credit ratings, *Journal of Finance*, forthcoming.

Bongaerts, Dion, Frank de Jong and Joost Driessen, 2011, Derivative pricing with liquidity risk: Theory and evidence from the credit default swap market, *Journal of Finance*, 66: 203–240.

Campbell, John, and Glen Taksler, 2003, Equity volatility and corporate bond yields. *Journal of Finance*, 58: 2321–2350.

Campello, Murillo, Long Chen, and Lu Zhang, 2008, Expected returns, yield spreads, and asset pricing tests, *Review of Financial Studies* 21, 1297–1338.

Case, Karl, and Robert Shiller, 1987, Prices of single-family homes since 1970: New indexes for four cities, *New England Economic Review*, 45–56.

Chacko, George, 2005, Liquidity risk in the corporate bond markets, Working paper, Harvard Business School.

Constantinides, George, 1986, Capital market equilibrium with transaction costs, *Journal of Political Economy,* 94(4): 842-862.

Chen, Long, David Lesmond and Jason Wei, 2007, Corporate yield spreads and bond liquidity, *Journal of Finance* 62, 119-149.

Chen, Long, Pierre Collin-Dufresne, and Robert Goldstein, 2009, On the relation between the credit spread puzzle and the equity premium puzzle, *Review of Financial Studies*, 22(9): 3367-3409.

Covitz, Dan and Chris Downing, 2007, Liquidity or credit risk? The determinants of very short-term corporate yield spreads, *Journal of Finance*, 62: 2303–2328.

Cremers, Martijn, Joost Driessen, and Pascal Maenhout, 2008, Explaining the level of credit spreads: Option-implied jump risk premia in a firm value model, *Review of Financial Studies*, 21(5): 2209-2242.

David, Alexander, 2008, Inflation uncertainty, asset valuations, and the credit spreads puzzle, *Review of Financial Studies*, 21(6): 2487-2534.

de Jong, Frank, and Joost Driessen, 2006, Liquidity risk premia in corporate bond markets, Working paper, University of Amsterdam.

Dick-Nielsen, Jens, Peter Feldhütter and David Lando, 2011, Corporate bond liquidity before and after theonset of the subprime crisis, *Journal of Financial Economics*, forthcoming.

Downing, Chris, Shane Underwood, and Yuhang Xing, 2005, Is liquidity risk priced in the corporate bond market?, Working paper, New York University.

Elton, Edward, Martin Gruber, Deepak Agrawal, and Christopher Mann, 2001, Explaining the rate spread on corporate bonds, *Journal of Finance*, 56, 247–277.

Feldhütter, Peter, 2011, The same bond at different prices: Identifying search frictions and selling pressures, *Review of Financial Studies*, forthcoming.

Friewald, Nils, Rainer Jankowitsch, and Marti Subrahmanyam, 2010, Illiquidity or credit deterioration: A study of liquidity in the U.S. corporate bond market during financial crises, *Journal of Financial Economics*, forthcoming.

Hasbrouck Joel, 2009, Trading costs and returns for U.S. equities: Estimating effective costs from daily data, *Journal of Finance* 64, 1445–1477.

He, Zhiguo, and Wei Xiong, 2010, Rollover risk and credit risk, Working paper, University of Chicago.

Houweling, Patrick, Albert Mentink, and Ton Vorst, 2005, Comparing possible proxies of corporate bond liquidity, *Journal of Banking & Finance* 29, 1331-1358.

Huang, Jing-Zhi and Ming Huang, 2003, How much of corporate-treasury yield spread is due to credit risk?: A new calibration approach, Working paper Cornell University.

Lee, Charles and Mark Ready, 1991, Inferring trade direction from intraday data,

Journal of Finance,46: 733-746.

Lin, Hai, Junbo Wang, and Chunchi Wu, 2010, Liquidity risk and the cross-section of expected corporate bond returns, *Journal of Financial Economics*, forthcoming.

Longstaff, Francis, 2010, Asset pricing in markets with illiquid assets, *American Economic Review*, forthcoming.

Lou, Xiaoxia, and Ronnie Sadka, 2010, Liquidity level or liquidity risk? Evidence from the financial crisis, *Financial Analysts Journal*, forthcoming.

Merton, Robert, 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance* 29, 449-470.

Nashikkar, Amrut, and Marti Subrahmanyam, 2006, The determinants of liquidity in the corporate bond markets: An application of latent liquidity, Working paper, New York University.

Pastor, Lubos and Stambaugh, Robert, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy*, 111, 642-85.

Roll Richard, 1984, A simple implicit measure of the bid-ask spread in an efficient market, *Journal of Finance*, Vol. 39, No. 4, pp. 1127-1139.

Shanken, Jay, 1992, On the estimation of beta-pricing models, *Review of Financial Studies,* 5 (1): pp. 1-55.

Vayanos, Dimitri, 2004, Flight to quality, flight to liquidity, and the pricing of risk, Working paper, London School of Economics.

Vayanos, Dimitri, and Jiang Wang, 2009, Liquidity and asset prices: A unified framework, Working paper, London School of Economics.

Table 1: **Summary statistics on returns, factors, costs, and betas**

Panel A presents time-series correlations of the four factors: corporate bond liquidity shocks (CBLIQ), equity market liquidity shocks (EQLIQ), equity returns (EQ), and shocks to the VIX (VIX). The sample period is 2005 to 2008 with weekly observations, and we use overlapping four-weekly changes for the factors. Panel B presents statistics for the return and liquidity data, across 62 double-sorted portfolios sorted first on rating or EDF and then on the basis of trading activity, average bond age or issue size. The first three columns present averages across all portfolios. The second and third column present average t-stats and average R^2 of the first step regressions. The final two columns show average values for low liquidity and the high liquidity portfolios. Annualized expected returns (in excess of government bond returns) are denoted by $E(r)$ and average transaction costs by $E(c)$, both in percentages. Expected costs are calculated for each portfolio as the average transaction costs over the full sample, estimated using Roll's approach. The betas capture exposure of corporate bond returns to the CBLIQ, EQLIQ, EQ and VIX factors. All factors are scaled to have the same standard deviation as the equity market return. Panel C provides cross-sectional correlations of the multivariate factor betas and expected liquidity $E(c)$ across all portfolios.

Panel A: Time-series correlations of factors

	CBLIQ	EQLIQ	EQ	VIX
CBLIQ	100%	36.3%	-34.6%	36.2%
EQLIQ		100%	-70.2%	71.3%
EQ			100%	-91.6%
VIX				100%

Panel B: Expected returns, costs, and betas

rating	average	t-stat	R^2	low liq	high liq
$E(r)(\%)$	1.868	[4.80]		1.938	1.778
$E(c)(\%)$	0.833	[6.15]		0.933	0.712
univariate					
β^{CBLIQ}	-0.264	[-8.63]	26.8%	-0.240	-0.248
β^{EQLIQ}	-0.435	[-13.95]	47.2%	-0.406	-0.463
β^{EQ}	0.380	[14.62]	49.8%	0.358	0.399
β^{VIX}	-0.360	[-15.43]	52.4%	-0.338	-0.382
multivariate					
β^{CBLIQ}	-0.101	[-4.04]		-0.087	-0.111
β^{EQLIQ}	-0.158	[-3.65]		-0.140	-0.175
β^{EQ}	0.076	[1.10]		0.088	0.052
β^{VIX}	-0.163	[-3.48]	64.9%	-0.147	-0.191

Panel C: Cross-sectional correlations of betas and expected liquidity

	$E(c)$	β^{CBLIQ}	β^{EQLIQ}	β^{EQ}	β^{VIX}
$E(c)$	100%	-36.9%	4.7%	46.4%	26.7%
β^{CBLIQ}	-36.9%	100%	-58.5%	-11.3%	-24.5%
β^{EQLIQ}	4.7%	-58.5%	100%	-33.4%	6.3%
β^{EQ}	46.4%	-11.3%	-33.4%	100%	66.6%
β^{VIX}	26.7%	-24.5%	6.3%	66.6%	100%

Table 2: **Risk-factor model: Cross-sectional regression estimates**

This table presents estimates of the risk factor model in equation (2), from a cross-sectional regression of expected excess corporate bond returns on corporate bond liquidity betas (CBLIQ), equity liquidity betas (EQLIQ), equity market returns (EQ), VIX innovations (VIX), and expected transaction costs ($E(c)$). The sample period is 2005 to 2008 and the portfolios are based on various sequential sorts; first, the portfolios are sorted on rating or EDF, then each rating/EDF portfolio (except the AAA and CCC rating portfolios) is sorted on the basis of trading activity, firm size or issue size. In total, there are 62 portfolios. Shanken (1992) t-statistics are given in square brackets. The betas that go into the cross-sectional regressions are calculated with all factors scaled to have the same standard deviation as the equity market return (except for the expected liquidity $E(c)$).

model	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
$\zeta^{E(c)}$	2.194						1.172		0.923	1.296	1.263
	[4.05]						[4.00]		[4.25]	[4.61]	[4.50]
λ^{CBLIQ}		-6.367				-1.935	0.633			0.357	0.508
		[-5.20]				[-3.98]	[2.30]			[1.15]	[1.70]
λ^{EQLIQ}			-4.072					-2.629	-2.637	-2.297	-2.083
			[-5.32]					[-7.91]	[-7.56]	[-7.01]	[-6.9]
λ^{EQ}				4.743		4.713	3.089	5.273	2.231	2.038	3.246
				[5.29]		[5.55]	[5.81]	[4.12]	[3.06]	[2.64]	[4.63]
λ^{VIX}					-5.032						-1.798
					[-5.26]						[-2.49]
R^2	16.8%	-25.5%	42.9%	47.2%	43.1%	47.3%	65.1%	47.8%	60.8%	66.1%	68.7%

Table 3: **Equity turnover regressions**

The table contains estimates of a linear regression of monthly changes in the value-weighted mean of log turnover across stocks on the value-weighted equity market excess return and changes in the value-weighted ILLIQ across stocks, corrected for inflation. Data runs from Jan 2005 to Dec 2008. t-statistics are in brackets.

model	(1)	(2)	(3)
Δ Equity market illiquidity	0.140		-0.004
	[1.93]		[-0.03]
Equity market return		-1.725	-1.755
		[-2.51]	[-1.53]
Constant	0.001	0.001	0.001
	[0.32]	[0.33]	[0.34]
N	47	47	47
R^2	7.6%	12.3%	12.3%

Table 4: **Corporate bond turnover regressions**

The table contains estimates of a linear panel regression of quarterly changes in individual bond log turnover on changes of the Roll-based market-wide transaction cost measure for the corporate bond market and changes in the market-wide credit spreads (average credit spread across firms). We also include interactions between a high quality dummy (equal to one for bonds rated AAA or AA) and the change in cost index or market-wide credit spread, respectively. Bond-specific fixed effects are included and standard errors are clustered by quarter. The sample runs from 2005Q1 to 2008Q4 and contains data on the same sample of bonds as is used in the previous tables. Bond turnover is winsorized at 5% top and 5% bottom. t-statistics are in brackets.

model	(1)	(2)	(3)	(4)	(5)
Δ Corporate bond market illiquidity	-0.149 [-7.29]		-0.118 [-7.62]	-0.144 [-6.63]	-0.119 [-7.88]
Δ Credit spread		-7.844 [-2.45]	-3.440 [-2.31]	-3.526 [-2.36]	-5.327 [-2.67]
Δ Corporate bond market illiquidity * HQ dummy				0.0841 [2.71]	
Δ Credit spread * HQ dummy					5.965 [2.46]
HQ dummy (AAA/AA)				0.0173 [0.31]	0.00632 [0.11]
Constant	-0.0371 [-1.93]	-0.0279 [-1.45]	-0.0275 [-1.42]	-0.0326 [-1.32]	-0.0295 [-1.20]
N	113776	113776	113776	113776	113776
R^2	4.9%	4.8%	5.0%	5.0%	5.0%
# clusters	15	15	15	15	15

Table 5: **Risk-factor model: Robustness checks**

This table presents various robustness checks on estimates of the risk factor model in Table 2 (equation (2)): (i) adding portfolios sorted on corporate bond liquidity betas, (ii) using the imputed roundtrip cost measure to construct the corporate bond liquidity factor, (iii) using the Pastor-Stambaugh measure to construct the equity market liquidity factor, (iv) using $Moody's - KMV$ default probability estimates as credit risk portfolio characteristic, (v) using a subsample up to August 2008, (vi) adding a constant term to the model, and (vii) using a Fama-MacBeth approach with time-varying betas, expected returns and costs. The Fama-MacBeth estimates are based on four-week rolling regressions with betas estimated on a 52-week rolling window.

	Benchmark	CBLIQ beta sort	IRT measure	PS measure	Credit control	Pre-crisis	Intercept	FB avg.
$\zeta^{E(c)}$	1.263	1.478	1.230	1.178	1.233	0.960	1.272	0.917
	[4.50]	[3.72]	[4.38]	[3.94]	[4.37]	[4.04]	[7.54]	
λ^{CBLIQ}	0.508	0.489	0.959	1.182	1.320	-0.129	0.504	-0.772
	[1.70]	[1.478]	[2.88]	[4.34]	[4.63]	[-0.66]	[1.49]	
λ^{EQLIQ}	-2.083	-2.454	-2.122	-4.338	-1.391	-4.238	-2.087	-1.324
	[-6.90]	[-8.65]	[-7.01]	[-3.92]	[-4.5]	[-8.07]	[-7.44]	
λ^{EQ}	3.246	2.890	3.375	3.229	2.924	5.755	3.247	5.452
	[4.63]	[5.05]	[4.79]	[5.56]	[3.94]	[6.81]	[4.70]	
λ^{VIX}	-1.798	-1.604	-1.747	-1.954	-2.198	-4.562	-1.812	-4.123
	[-2.49]	[-2.83]	[-2.50]	[-3.37]	[-3.43]	[-6.70]	[-3.71]	
Credit control					0.253			
					[2.67]			
Intercept							-0.012	
							[-0.05]	
R^2	68.7%	69.6%	68.0%	70.3%	72.0%	55.6%	68.7%	

56

Table 6: **Market and liquidity betas across corporate bond portfolios**

The table presents average betas (across all portfolios) of the Bongaerts, de Jong and Driessen (2011) model in equation (14). We have $\beta^{eq} = \beta^{r,eq} - \zeta_0 \beta^{c,eq}$ and $\beta^{vix} = \beta^{r,vix} - \zeta_0 \beta^{c,vix}$, with $\zeta_0 = 1.2$, $\beta_{other} = -\beta^{rc} - \beta^{cr} + \beta^{cc}$ and $\beta_{net} = \beta^{rr} + \beta_{other}$. The other betas are defined in equations (11) and (12).

$\beta^{r,eq}$	0.158
$\beta^{r,vix}$	-0.165
$\beta^{c,eq}$	-0.013
$\beta^{c,vix}$	0.018
β^{eq}	0.174
β^{vix}	-0.187
β^{rr}	0.760
β^{rc}	-0.098
β^{cr}	-0.127
β^{cc}	0.044
β^{other}	0.268
β^{net}	1.029

Table 7: **BDD model: Cross-sectional regression estimates**

This table present estimates of the Bongaerts, de Jong and Driessen (2011) model in equation (14). The sample runs from 2005 to 2008. t-statistics are given in square brackets.

model	(1)	(2)	(3)	(4)	(5)	(6)	(7)
β^{eq}	3.902	3.844	3.844	2.969	2.354	3.286	3.955
	[6.95]	[7.03]	[7.03]	[3.78]	[3.17]	[3.95]	[5.94]
β^{vix}	-2.149	-2.154	-2.195	-1.500	-1.718	-1.943	-2.220
	[-3.18]	[-2.51]	[-3.08]	[-1.69]	[-1.93]	[-2.13]	[-2.67]
$E(c)$	0.932	0.933	1.043	1.067	1.112	1.156	0.961
	[4.36]	[3.62]	[3.95]	[4.12]	[4.85]	[4.90]	[3.58]
β^{rr}		-0.004		0.598	0.569	0.427	
		[-0.02]		[3.12]	[2.59]	[1.84]	
β^{rc}			0.906	2.927	0.471		
			[1.55]	[5.27]	[0.92]		
β^{cr}					-5.169		
					[-3.37]		
β^{cc}					-19.02		
					[-4.92]		
β^{other}						-1.380	
						[-4.07]	
β^{net}							-0.044
							[-0.38]
R^2	72.8%	72.8%	73.3%	74.5%	77.7%	73.9%	72.8%

Figure 1: **Time series of corporate bond transaction costs: Roll measure and imputed roundtrip costs**
The figure shows weekly time series of the transaction costs obtained by using the Gibbs sampler to estimate the Roll model (solid line), averaged across all portfolios, and the imputed roundtrip cost measure of Feldhütter (2011) and Dick-Nielsen, Feldhütter, and Lando (2011), averaged across all bonds (dashed line). The sample period is 2005 to 2008.
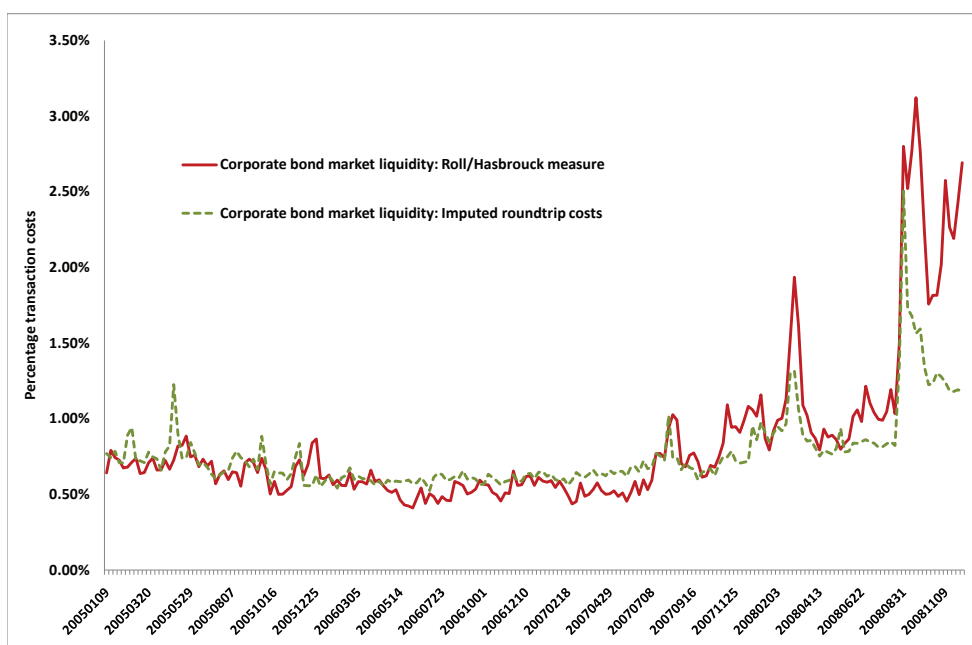
Figure 2: **Fit of expected returns across portfolios**

The figure shows the fitted values of the expected bond returns, obtained by multiplying the estimated coefficients in Table 2, specification (11) with the estimated expected cost and the estimated betas. The equity risk premium and volatility risk premium are presented together. Alpha is the pricing error as defined in equation (2). The fit is presented for portfolios across rating categories / EDFs and liquidity proxies, averaged across the three liquidity proxy sorts (amount issued, age, and activity). For example, "AA-hi" refers to the high-liquidity AA portfolios, while "AA-lo" refers to the low-liquidity AA portfolios.
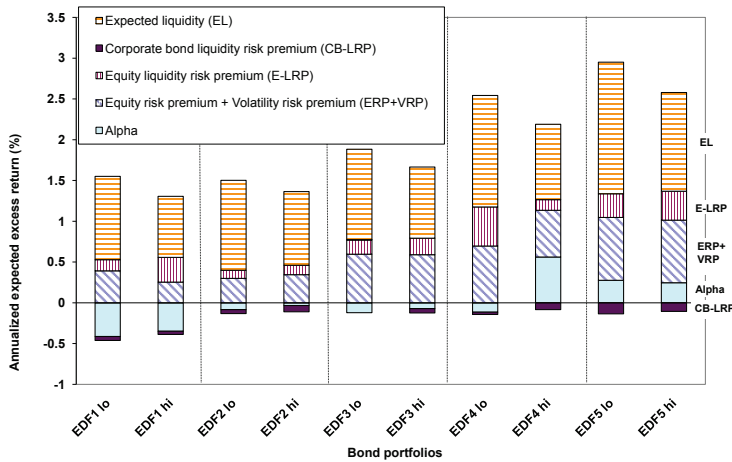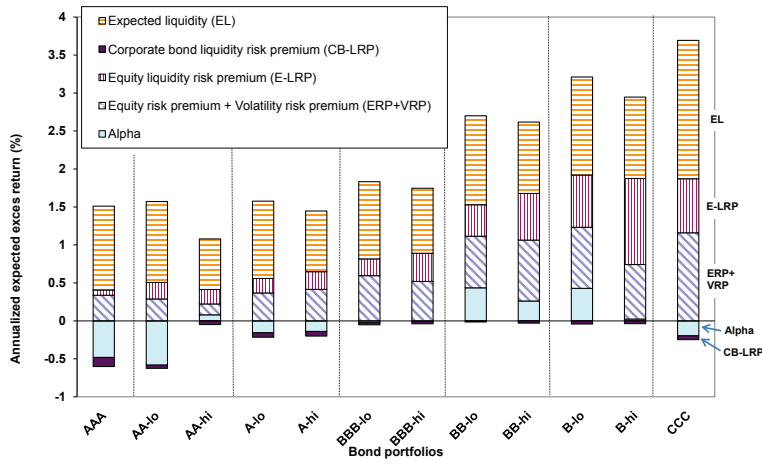
Figure 3: **Turnover in equity and corporate bond markets**
The figure shows quarterly turnover in the equity and corporate bond market. For the equity market, the time series of both the value-weighted mean and median of quarterly turnover across all CRSP stocks is shown. Similarly, for corporate bonds the time series of the value-weighted mean and median of quarterly bond turnover across all TRACE bonds is shown.
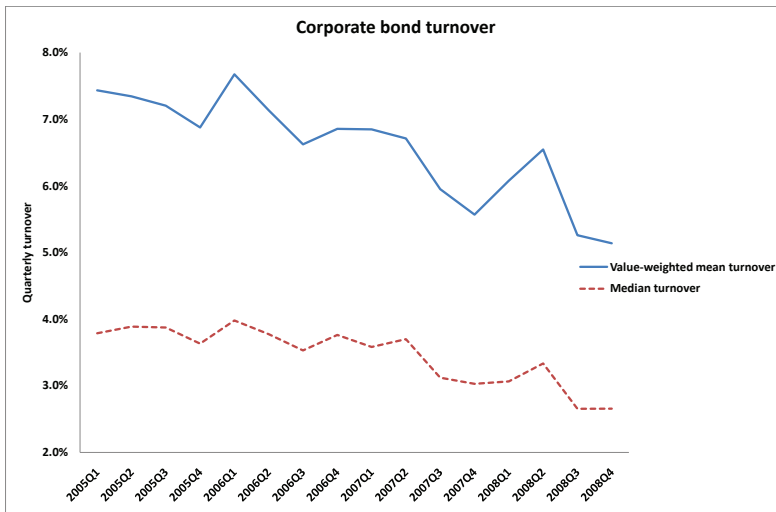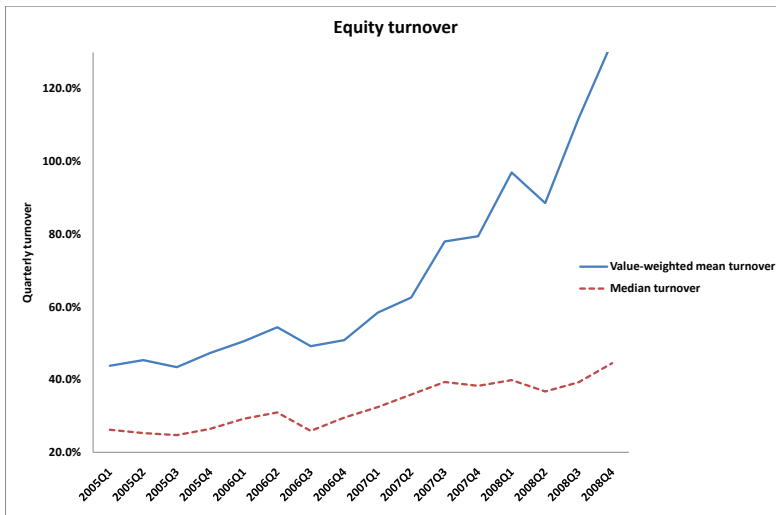
Figure 4: **Time-varying estimates**

The top panels of this figure show the estimated betas, second step regression coefficients and implied risk premiums for the equity return (EQ, solid line), equity market liquidity innovations (EQLIQ, dotted line) and corporate bond liquidity innovations (CBLIQ, dashed line). The bottom panels show the same items for the expected liquidity. The estimates are obtained from four-week rolling windows of expected return and cost data, with a two-sided kernel estimate of the betas.