A Conditional Analysis of Liquidity on Quality in the U.S. Corporate Bond Universe

Arwed Ebner *
Stockholm School of Economics

Paul Schempp *
Stockholm School of Economics

May 19, 2014

Abstract

Why are high-yield bonds more severely hit by large liquidity dry outs than investment grade bonds? This study investigates the effects of liquidity shocks on returns in the U.S. corporate bond market during times of heightened liquidity stress, using a comprehensive data set of 13,500 bonds between October 2004 and September 2013. Applying a Markov regime-switching model, we identify liquidity stress periods in which illiquid bonds underperform by as much as 21.7% relative to their liquid counterparts in the high-yield segment, while the same return differential amounts to only 5.4% for investment grade bonds. We show that classical explanatory approaches fail to describe these differing effects on returns: neither the pre-crisis liquidity levels nor the liquidity shocks during the stress periods show an asymmetric distribution across ratings. Thus, a puzzling aspect of investors' behavior can be inferred: during times of distress, investors punish the same unit of illiquidity differently across credit quality. In order to explain this phenomenon, we develop a model in non-formal reasoning. It is grounded on the idea that investors perceive liquidity dry outs as transitory and therefore only penalize assets that are likely to be sold in the short-term. Since investors are more risk-averse in times of distress, high-yield bonds are the first of the investors' assets in line to be liquidated, and thus investment grade bonds only show marginal return effects with respect to liquidity shocks, perceived to be of temporary nature.

*This paper was submitted as Master thesis at the Department of Finance of the Stockholm School of Economics. We would like to express our gratitude to our supervisor Jungsuk Han for his guidance and support throughout each stage of our thesis. We also want to thank Markus Ibert, Paul Schultz, Michel Bauer, and one anonymous referee for helpful comments. All errors remain our own.
E-mail addresses: 40396@student.hhs.se (Arwed Ebner), 40379@student.hhs.se (Paul Schempp)

Contents

List of Figures

List of Tables

1 Introduction

During both the fall of Long-Term Capital Management and the peak of the subprime crisis, corporate bond spreads soared to far greater levels than the traditional fundamentals could explain. Practitioners quickly saw the reason for this unexpected blow to prices in a simple concept: the absence of buyers and hence of market liquidity. Thereupon, academics have tracked down the striking characteristic of liquidity: its feature of being more time-varying than most other risk sources. While playing a non-trivial but only moderate role during normal periods, liquidity has the power to exert a major influence in times of heightened market uncertainty. One specific effect during liquidity stress periods raises our attention, which is the empirical observation that liquidity disruptions affect low-rated bonds to a much higher extent than high-rated bonds.

Our sample of U.S. corporate bonds also exhibits this effect: During the liquidity stress periods as identified in our study (comprising 10 to 14 months, mainly during the subprime crisis) the prices of illiquid bonds underperform their liquid counterparts by as much as 21.7% in the high-yield segment, while the same return differential amounts to only 5.4% for investment grade bonds. Hence, the purpose of this study is to explore the reason behind the *difference in return differentials* between high-yield and investment grade bonds since no research has disentangled the underlying effects so far. The concept of "differences in return differentials" may appear elusive at first, can however figuratively be clarified by thinking in terms of four bond portfolios: One illiquid and one liquid portfolio for each high-yield and investment grade rating class. It is obvious that the prices of the illiquid portfolios must lie below the prices of their liquid counterparts within each rating category since low liquidity is undesirable. The goal of our work is to understand why this price difference widens by far more for high-yield bonds than for investment grade bonds (21.7% compared to 5.4% in our sample) during liquidity stress periods. Interestingly, we draw a puzzling conclusion by falsifying the two possible rational explanatory approaches: neither the pre-crisis liquidity levels nor the liquidity shocks during the stress periods show an asymmetric distribution across ratings.

First, the two high-yield portfolios might generally exhibit a lower portfolio-specific liquidity level than the investment grade portfolios. Consequently, a rising liquidity premium across all four portfolios would punish the lower liquidity levels of high-yield bonds by a higher absolute extent and thus result in a more pronounced price differential between the bonds in the high-yield than between the bonds in the investment grade class. Interestingly, we challenge the prevailing opinion about the lock step of liquidity and quality by finding that bonds show the *same* level of liquidity across credit ratings. Hence, a rise in the overall liquidity premium would affect the bonds across ratings to the same extent.

The second explanation would be that a withdrawal of market-wide liquidity hits the four portfolios in an asymmetric manner, e.g. that predominantly the illiquid portfolio of the high-yield segment drops in its liquidity. As a result, the difference in the underlying characteristic of liquidity among the high-yield segment would have risen in comparison to the investment grade portfolios. If the extent of the liquidity increases between two portfolios, their price difference must increase as well. Surprisingly, we show that liquidity dry outs are equally spread across credit ratings and that the dispersion of the liquidity shock within each rating category is the same, i.e. illiquid bonds of all ratings experience the same strong liquidity shock so that the liquidity dispersion[1] across ratings stays the same.

As a result, we infer the puzzling conclusion that investors punish the same unit of illiquidity differently across credit ratings in times of distress. During those times, a severe decrease in the bond-specific liquidity levels across all ratings hits high-yield bonds by a large extent, while bonds with ratings of BBB and above are only marginally affected by the same magnitude of shock. Investors appear to disregard liquidity shocks to investment grade bonds during times of distress. Besides not only arguing that investors punish the same magnitude of liquidity shock differently across ratings, we show that the liquidity premium must exclusively increase for high-yield bonds during stress times in order to explain the observed price differentials.

To sum up, analyzing the reason for the differences in returns allows us to eventually draw inferences about the absolute effects of liquidity. We conclude that liquidity dry outs affect all rating classes to the same nominal amount, but that investors punish the withdrawal of liquidity differently according to the quality of the bond.

In order to explain the discriminating treatment of the same unit of liquidity, we develop a model in non-formal reasoning. Our explanation for the asymmetric liquidity effects also sheds light on another phenomenon visible in our sample: We observe a *flight-from-junk* effect in the form that high-yield bonds drop by about 20%, while investment grade bonds suffer only 2% over the course of the identified liquidity stress period. The rationale for the model developed is grounded on the work of Vayanos (2004) and Brunnermeier and Pedersen (2009) that both predict rising market volatility to increase the threat of unwinding positions in order to meet withdrawals or funding margins. We argue that this rising probability to liquidate positions increases selling pressure and thus decreases prices for the assets likely to be sold. Assuming that high-risk assets are sold off first in stress periods (due to an increase in risk-aversion), the selling pressure predominantly hits the high-yield segment, leading to the flight-from-junk effect. Furthermore, the increasing probability for earlier than initially expected liquidation of high-yield bonds decreases their expected holding period which in turn increases the effective transaction costs (i.e. the transaction costs per holding period). Consequently, the same nominal transaction costs (i.e. liquidity) affect the investor to a higher extent, leading to an increasing liquidity premium. We then show that if investors perceive the accompanied liquidity dry out as transitory and thus non-persistent, they only penalize liquidity shocks for assets that are likely to be sold during the stress times. Since high-yield bonds are first in line to be liquidated, and thus better-rated bonds more likely to be unwound later when markets have returned to normality, investment grade bonds only show marginal return effects to these temporary illiquidity shocks.

Turning towards the technical side of the paper, an analysis of liquidity effects during liquidity stress periods naturally demands for measuring liquidity and identifying the corresponding stress periods. Both tasks might appear simple, but are actually far from obvious since on the one hand liquidity is not directly measurable and on the other hand liquidity stress periods arise very suddenly and are short in nature.

First, liquidity, in the form of ease of finding a counterparty to trade with, is a nebulous and multi-faceted phenomenon: One knows it when one sees it, but it is hard to define. Therefore, a

[1] The difference in liquidity between illiquid and liquid bonds.

plethora of liquidity measures has been developed over the years to proxy for different dimensions of liquidity. We employ four well-known liquidity measures (initially developed for the stock market, but successfully applied to corporate bonds by many studies) on a daily basis, namely the Amihud (2002) price impact measure, the Pastor-Stambaugh (2003) price reversal measure, the Roll (1984) bid-ask spread estimator, and the Corwin-Schultz (2012) bid-ask spread estimator. While the former three have been widely used for many years, the relatively new Corwin-Schultz estimator is currently attracting a lot of attention due to its simple implementation and high accuracy. In order to apply the idea of the Corwin-Schultz measure on infrequently traded bonds, we algebraically derive a more flexible, enhanced version and further show that this extension provides an equal or more accurate estimation of the bid-ask spread in c.85% of the transactions in our sample compared to the original measure. In order to examine whether our liquidity measures accurately measure liquidity, we estimate unconditional risk premia of up to 1% for the liquidity level and up to 0.5% for the liquidity risk (the bond's return co-movement with market-wide liquidity) whose magnitudes are in line with those found in the academic literature. Therefore, we perform the Fama-MacBeth procedure on five standard risk factors together with the bond-specific liquidity level and liquidity risk on a monthly return level. The results hold regardless of liquidity measure employed.

Second, since the severe effects of liquidity arise suddenly and only within a short time frame, an accurate identification of the liquidity stress periods is of utmost importance in order to precisely capture the effects of the *typically sleeping, but at times rampaging giant* called liquidity. Therefore, we use Markov regime-switching models since Watanabe and Watanabe (2008) and Acharya et al. (2013) show the models' power in identifying periods with severe liquidity effects. These models are based on the idea that the stochastic process of a variable (e.g. the return) can abruptly change between two underlying processes with different statistical properties. We not only refine their approach with the Baum-Welch algorithm, but further apply the Viterbi algorithm in order to determine an actual series of liquidity stress periods rather than mere probabilities for being in liquidity distress for each point in time. The identified liquidity stress regime comprises 10 to 14 months, mainly during the subprime crisis (depending on the model specification) and shows, despite being endogenously generated by a purely statistical procedure, a high economic relevance.

Every analysis is only as good as its data and in order for liquidity effects to become more visible, it is important for the sample to consist of a high dispersion in liquidity across bonds. Therefore, we base our analysis on the unique data of the Trade Reporting and Compliance Engine (TRACE) which contains approximately 90% of all corporate bond transactions in the U.S. from October 2004 on. This allows us to not only capture the bulk of illiquid bonds, but also to directly work with the most detailed primary data for corporate bonds since bond dealers have to directly report every single transaction to TRACE. We merge the transaction data from TRACE from October 2004 to September 2013 (from the moment of the full implementation of TRACE to the most recent date for which data is available to us) with bond characteristics (especially ratings) from the Fixed Income Securities Database (FISD). Extensively preparing and filtering the data set, we obtain a unique sample covering almost 330,000 bond months, more than 13,500 bonds and nearly $15tn in trading volume, representing one of the most comprehensive data sets for analyzing liquidity in the corporate bond market. While other studies tend to apply rough filters which primarily exclude illiquid bonds, we only omit observations that are specifically unfeasible for the constructed liquidity measures in order to minimize the omission of the illiquid parts of the market.

The remainder of the paper is organized as follows: we demarcate us from the relevant literature in Section 2. In Section 3, we explain the composition of our data set as well as the filters and matching procedures and show summary statistics. Section 4 describes our methodology employed to measure liquidity, to perform asset pricing tests of liquidity, and to determine liquidity stress periods. Section 5 presents our results. Section 6 discusses theoretical explanations for our findings. Section 7 summarizes our main findings and concludes the paper.

2 Literature Review

There are two main strands of literature that are in close relation to our paper.

The first explores the relevance of liquidity on asset prices and more specifically on corporate bond returns. Due to the difficulties in identifying and measuring liquidity, the academic literature focusing on the effect of liquidity on asset prices developed late, but is by now vast and growing further at high speed[2]. Liquidity in the context of asset pricing can be roughly divided into two subsections: *liquidity level* and *liquidity risk*. Before demarcating our work from the existing literature, we want briefly survey the most important studies of liquidity in asset pricing, first for stocks then for bonds.

The studies of Amihud and Mendelson (1986), Brennan and Subrahmanyam (1996), and Brennan et al. (1998), are commonly considered the seminal works on the effect of liquidity level on asset returns and find that illiquidity increases the required return on a stock.

Due to the finding of commonality of liquidity – the phenomenon that an asset's individual liquidity co-moves with market-wide liquidity – in the studies of Chordia et al. (2000), Hasbrouck and Seppi (2001), and Huberman and Halka (2001), the attention of researchers moved to the question of whether there is a systematic and undiversifiable liquidity risk factor (i.e. co-movement of a bond's return with the market-wide liquidity) relevant in asset pricing. Several studies find that liquidity risk is priced in stocks, such as Pastor and Stambaugh (2003), Acharya and Pedersen (2005), Liu (2006), Bekaertet al. (2007), Korajczyk and Sadka (2008), and Lee (2011)[3].

Although the relevance of liquidity for corporate bonds has never been doubted, the first studies that investigate this issue are related to the so-called "credit spread puzzle", which is connected to the fact that bond yield spreads are larger than what can be explained by the mere default risk component of the bond (see Huang and Huang (2012) and Eom et al. (2004)). These studies (e.g. Elton et al. (2001) and Collin-Dufresne et al.(2001)) focus on the default risk component of corporate bond returns and partly attribute the unexplained portion of the spread to illiquidity. The first studies that directly analyze the price impact of liquidity on corporate bond returns all found evidence for the importance of both liquidity level and liquidity risk in bond returns and spreads (such as Houweling et al. (2005),

[2]See, for example, Amihud et al. (2005, 2013) and Vayanos and Wang (2013) for encompassing surveys about this literature.

[3]There is also evidence that liquidity risk can explain asset pricing anomalies (Sadka(2006) and Asness (2013)) and can explain parts of investment fund performance (Dong et al. (2014), Sadka(2010, 2012), and Franzoni et al.(2012)).

Downing et al. (2005), Chacko (2006), Covitz and Downing (2007), Chen et al. (2007), Bao et al. (2011), Lin et al. (2011), Bongaerts et al. (2012), and De Jong and Driessen (2012)).

The second strand of literature explores the relevance and magnitude of corporate bond liquidity on prices during distressed times. In contrast to the first strand, the number of these studies is moderate. There are three main papers that are of utmost relevance for our study, namely Dick-Nielsen et al. (2012), Friewald et al. (2012), and Acharya et al. (2013). They all show the exceptional feature of liquidity which is its large difference in importance across time. Dick-Nielsen et al. (2012) and Friewald et al. (2012) uncover the strong, but transitory importance of liquidity during crises also for the corporate bond market. For example, Dick-Nielsen et al. (2012) find that the spread between very liquid and average liquid bonds increases rapidly during the Lehman Brothers default in the fall of 2008 and returns to pre-crisis levels already in the summer of 2009. The crucial question for understanding the time-varying nature of liquidity is whether this spread is caused by decreasing liquidity, increasing liquidity premia or non-textbook aspects. Since Dick-Nielsen et al. (2012) show that illiquidity rises significantly during distressed times, it is unclear whether investors merely punish this increasing illiquidity or are additionally becoming more averse to carry each unit of illiquidity. Friewald et al. (2012) show that investors actually increase their unwillingness to hold illiquidity since the same unit of bond-specific illiquidity shocks doublse in its effect on the bond's returns during the subprime crisis.

However, our work aims at clarifying a higher dimension, i.e. the *cross-sectional* effects of liquidity in crises periods. Both Dick-Nielsen et al. (2012) and Friewald et al. (2012) illustrate that liquidity does not only matter strongly in crises, but also affects speculative grade bonds to a much larger extent. For example, Dick-Nielsen et al. (2012) uncover that the liquidity spread component of AAA rated bonds merely increase from 1bp to 5bp, while those of BBB bonds rise from 4bp to 93bp and those of the high-yield universe from 58bp to staggering 197bp during the subprime crisis. Acharya et al. (2013) underline this pattern by displaying an even more pronounced picture. During stress periods in their sample, high-yield bonds show highly negative return reactions to drops in market-wide liquidity, while the returns of investment grade bonds actually slightly increase with respect to the market's liquidity deterioration. This high dispersion of severity in which liquidity crises hit differently rated bonds appears puzzling. The simplest answer would be that liquidity only deteriorates for the lower end of credit ratings during stress periods. Acharya et al. (2013) sketch another explanation and assume an increasing demand for liquidity (i.e. a higher premium for illiquidity) during these periods. They consider speculative bonds less liquid than investment grade bonds so that investors would exhibit a flight-to-liquidity in the sense of re-allocating capital from illiquid high-yield towards liquid investment grade bonds. However, the sample of Friewald et al. (2012) which is based on a more recent and larger data set does not exhibit a co-movement of lower credit quality and illiquidity, so that the effect in Acharya et al. (2013) might also stem from a flight-to-quality instead of a flight-from-liquidity.

To sum up, we are interested in why speculative bonds are hit harder by liquidity during liquidity stress times. Specifically, no study so far has examined potential reasons such as whether low-rated bonds show a generally lower liquidity level that gets amplified in stress periods, drop more strongly in liquidity during stress periods, or suffer from increasing demand for liquidity in case of adverse economic situations. We aim at filling the gap of research by forming portfolios in both dimensions of liquidity and quality in order to isolate the effects. This is especially important liquidity stress periods are closely related to general quality stress periods as shown by the theoretical models of Vayanos (2004) and Brunnermeier and Pedersen (2009)[4].

3 Data

3.1 Data Description

In general, corporate bonds represent an ideal laboratory to scrutinize liquidity effects for four reasons. First, the infrequent nature of bond trading accentuates the role of liquidity. Second, the major risk component of bonds, the default risk, is directly observable via credit ratings so as to allow for isolating liquidity concerns to a high extent. Third, the yield-to-maturity of bonds uniquely reveals the expected returns of investors giving researchers the opportunity to increase accuracy in uncovering risk premia. Fourth, through the Trade Reporting and Compliance Engine (TRACE), detailed primary data on transaction level is available for almost the entire U.S. corporate bond market.

Our dataset consists of information from two sources; we obtain transaction data from the TRACE and bond characteristics from the FISD (Fixed Income Securities Database) databases[5]. Under pressure from the U.S. Congress, buy-side traders and the SEC, the Financial Industry Regulatory Agency (FINRA) implemented the TRACE database in July 2002 in order to increase transparency, to provide greater regulatory insight, and to decrease transaction costs for retail and small institutional investors in the over-the-counter U.S. corporate bond market (Bushman et al., 2010). In the beginning, the reporting was only limited to investment grade issues as there existed conflicting views about the effect of improving transparency. However since October 1, 2004, essentially all secondary over-the-counter U.S. corporate bond transactions are reported within 15 minutes to TRACE. As a result, the TRACE database comprises a very unique dataset for an over-the-counter market due to its level of detail and scope. As less than 10% of all U.S. corporate bond trades are executed via public exchanges and c.99% of all trades of the over-the-counter market is captured by the TRACE reporting (Edwards et al., 2007), our dataset embraces almost the entire universe of corporate bond transactions during our sample period. Besides its unique scope, TRACE also allows us to observe prices with the utmost degree of detail since it comprises information on intra-day transaction level such as actual transaction prices, trading volumes and the yield-to-maturity.

We span our sample period from the moment of the full implementation of TRACE to the most recent date for which data is available to us. This leaves us with a sample period from October 1, 2004, to September 30, 2013. We merge these transactions with bond issue-specific information from the Mergent FISD database to obtain additional bond characteristics such as maturity, coupon and credit

[4]Beber et al. (2009) show that liquidity stress periods partly differ from distressed quality times for the European sovereign bond market. They can fall apart since an increasing sovereign bond issue enhances market liquidity, but decreases the issuer's quality due to excessive leverage. For corporate bonds, liquidity appears to be much more related to funding capital for speculators which is why Beber et al. (2009) is rather not relevant for our scope.

[5]We access both databases via the Wharton Research Data Services (WRDS) webpage.

rating. Several papers have used the TRACE database to analyze liquidity effects in the U.S. corporate bond market[6]. Given TRACE's comprehensive scope and extraordinary detail for corporate bonds, the number of authors using TRACE is still relatively moderate which can probably be explained by the high challenges arising in the course of compiling this large set of primary data (which is by now more than 10 gigabytes large), but also by the relatively short period of time from when the data is available.

We undertake an extensive data preparation and filtering process to arrive at our final data sample. Table 1 and 2 display the effect of our filtering steps onto the data set. Appendix A describes our entire data preparation process in detail and might serve as a comprehensive guideline for future academic research since a thorough description of editing the TRACE's data set does not yet exist to the best of our knowledge. Our initial sample of TRACE's corporate bond trades consists of 77,255,887 transactions from 161,896 bonds issued by 16,665 firms over a number of 11,862,357 total bond trading days with a total dollar volume of \$26.3tn. We then apply the Dick-Nielsen (2009) error filters in order to correct for three types of cancelled and corrected trade reports resulting in dropping 4.8% of the transactions. We further apply additional price filters that shall detect wrong transaction reports which have not been corrected or cancelled. Those comprise mistakes such as typos or, by mistake, reporting yields as prices and vice versa. Consequently, the filters aim at identifying unusual price jumps among transactions. We construct the filters according to Edwards et al. (2007), Han and Zhou (2011) and Friewald et al. (2012) and detect 0.6pp of erroneous price entries. As we perform our liquidity measures on daily prices, we aggregate these intra-day transactions for each bond into a single summary observation that consists of the dollar volume-weighted price, the high and low prices and total trading volume of each day. In order to precisely capture the total returns sensitivity to liquidity changes, we transform the clean prices of TRACE into dirty prices according to the following formula:

$$Price_{dirty} = Price_{clean} + Accrued\,Interest_t + \sum_{t=0}^{T} Coupon_t, \tag{1}$$

Manually examining the dataset, we identify that erroneous price records are likely to appear in clusters (e.g. the broker submits the yield instead of the price for many consecutive transactions before correcting his reporting behavior). Consequently, we extent the price filters of the transactions above also to daily levels in order to capture these clusters of false records and identify 3.5pp of bond-specific trading days as invalid. We then merge the TRACE data set with credit rating and bond characteristics from FISD. We translate the official rating categories into a numeric scale from 10 to 0 with 10 being the highest rating (10=AAA, 9.5=Aa1/AA+, 9=Aa2/AA ... 1.5=Caa2/CCC, 1=Caa3/CCC-, 0.5=Ca/CC, 0=C/D) and assign the most recently published credit rating to each trading day of a specific bond. We then exclude convertible, exchangeable, puttable and perpetual bonds, as well as bonds denoted in foreign currencies. However, we retain callable bonds as this feature is very standard for bonds and applies to 56.9% of the bonds in our sample. In order to investigate the

[6]See, for example, Edwards et al. (2007), Goldstein et al. (2007), Ronen and Zhou (2009), Nashikkar et al. (2011), Lin et al. (2011), Jankowitsch et al. (2011), Feldhütter (2012), Friewald et al. (2012), Bao et al. (2011), Dick-Nielsen et al. (2012), and Das et al. (2014).

effects of liquidity changes on bond returns, we chose a monthly time horizon. Therefore, we exclude all months of a bond with less than 10 traded days since its liquidity measures which are constructed on a daily level might exhibit large biases. We define a month's return by the relative change from the preceding month's latest dirty price to the current month's latest dirty price. Furthermore, we require that these two prices stem from transactions in the last weeks of the respective months since we assume that such a return calculation is more accurate than one, which uses interpolation. Every month of a bond which does not meet this criterion is disregarded. The entire filtering and aggregation process results in a final sample of 329,252 monthly time-series observations from 13,842 bonds issued by 2,089 firms with a trading volume of $14.9tn.

Table 1: Data Prepration and Filtering Process I

This table shows all the steps of the preparation and filtering process for our data, which are descibed in more detail in Appendix A. *Transactions*, *Trading Days*, and *Trading Months* are reported in total numbers. The percentage numbers next to each variable show how much of the inital (non-filtered) data set our own data set after the particular data preparation step represents.

Data Preparation Steps	**Transactions**		**Trading Days**		**Trading Months**	
Transaction level						
Initial data set	77,255,887	(100.0%)	11,862,357	(100.0%)	1,579,321	(100.0%)
Dick-Nielsen filter 1	77,253,140	(100.0%)	11,835,978	(99.8%)	1,574,457	(99.7%)
Dick-Nielsen filter 2	76,580,843	(99.1%)	11,822,564	(99.7%)	1,571,103	(99.5%)
Dick-Nielsen filter 3	73,539,606	(95.2%)	11,809,741	(99.6%)	1,570,222	(99.4%)
Absolute price filter	73,477,694	(95.1%)	11,798,451	(99.5%)	1,568,104	(99.3%)
Intra-day median filter	73,455,481	(95.1%)	11,754,951	(99.1%)	1,565,791	(99.1%)
Preceding-transactions median filter	73,080,660	(94.6%)	11,728,350	(98.9%)	1,563,114	(99.0%)
Daily level						
Aggregated data set			11,728,350	(98.9%)	1,563,114	(99.0%)
Trading day median filter			11,320,426	(95.4%)	1,553,984	(98.4%)
Merging with FISD			10,222,333	(86.2%)	1,354,652	(85.8%)
Filtering according to bond characteristics			9,785,026	(82.5%)	1,196,737	(75.8%)
Filtering bond months with more than 10 trading days			6,811,664	(57.4%)	427,136	(27.0%)
Monthly level						
Aggregated data set					427,136	(27.0%)
Final data set (filtered by prices in month's last week)					329,252	(20.8%)

Table 2: Data Prepration and Filtering Process II

This table shows all the steps of the preparation and filtering process for our data, which are descibed in more detail in Appendix A. *Volume* is reported in $bn, *#Bonds* and *#Companies* are reported in total numbers. The percentage numbers next to each variable show how much of the inital (non-filtered) data set our own data set after the particular data preparation step represents.

Data Preparation Steps	Volume ($bn)		#Bonds		#Companies	
Transaction level						
Initial data set	26,306	(100.0%)	161,896	(100.0%)	16,665	(100.0%)
Dick-Nielsen filter 1	26,305	(100.0%)	161,893	(100.0%)	16,665	(100.0%)
Dick-Nielsen filter 2	26,017	(98.9%)	158,759	(98.1%)	16,009	(96.1%)
Dick-Nielsen filter 3	23,817	(90.5%)	147,173	(90.9%)	15,598	(93.6%)
Absolute price filter	23,804	(90.5%)	147,700	(91.2%)	15,465	(92.8%)
Intra-day median filter	23,802	(90.5%)	146,327	(90.4%)	15,298	(91.8%)
Preceding-transactions median filter	23,655	(89.9%)	145,259	(89.7%)	15,265	(91.6%)
Daily level						
Aggregated data set	23,655	(89.9%)	145,259	(89.7%)	15,265	(91.6%)
Trading day median filter	22,751	(86.5%)	139,392	(86.1%)	14,682	(88.1%)
Merging with FISD	21,271	(80.9%)	86,242	(53.3%)	5,576	(33.5%)
Filtering according to bond characteristics	19,525	(74.2%)	47,088	(29.1%)	3,283	(19.7%)
Filtering bond months with more than 10 trading days	17,040	(64.8%)	18,935	(11.7%)	2,452	(14.7%)
Monthly level						
Aggregated data set	17,040	(64.8%)	18,935	(11.7%)	2,452	(14.7%)
Final data set (filtered by prices in month's last week)	14,910	(56.7%)	13,842	(8.5%)	2,089	(12.5%)

In our work, we aim at capturing the default risk component of a bond's return by its credit rating, although Credit Default Swaps (CDS) may appear more precise in quantifying default risks. Becker and Milbourn (2011) and Bolton et al. (2012) provide theoretical and empirical evidence of competition and conflicts of interest within the credit rating industry that may result in inefficient ratings with decreasing ability to predict default, while Bar-Isaac and Shapiro (2013) detect countercyclical credit rating to the extent that credit rating agencies are more likely to issue less accurate ratings in boom times than during recessionary periods which is supported by Baghai et al. (2014) and Xia (2014), who find that rating agencies become more conservative and adopt stricter rating standards over time.

On the other hand, CDSs have seen an enormous rise in popularity during the last decade and Ericsson et al. (2009) even show that the CDS spread measure of a firm's default risk more precisely than credit ratings. Furthermore, Flannery et al. (2010) evaluate the viability of CDS spreads as substitutes for credit ratings and support using CDS for regulatory purposes. Although, a CDS spread does not only reflect the default risk, but also contains a liquidity premium itself (Bongaerts et al. (2011), Junge and Trolle (2013), and Tang and Yan (2013)) as well as counterparty risk (Arora et al., 2012), we consider CDSs as a more accurate measure of default risk than credit ratings. However, we rely on credit ratings due to a very practical reason; CDSs are only traded for a very small set of bonds so that we would not be able to match every bond of our final data set with a respective CDS spread. Furthermore, as our work aims at examining the differing behavior of bonds with very different credit quality, we believe that credit ratings are sufficiently accurate in order to correctly categorize the bonds according to their credit risk.

3.2 Summary Statistics

As displayed in Table 2, the number of bonds in our final data set amounts to only 8.5% of the initial scope retrieved from TRACE; however these bonds comprise 56.7% of the dollar volume of the initial sample. Hence, the scope of the liquidity analyses in our work is limited to a relatively small set of bonds that accounts for approximately half of the markets turnover. The major part of bonds is filtered out due to lacking information in the FISD database (32.8% of bonds), non-standard embedded options (24.2% of bonds) and too infrequent trading activity for applying our set of liquidity measures (17.4% of bonds)[7]. Especially given the latter selection criteria of frequent trading days, the bonds in our sample are likely to be more liquid than a typical bond. Table 3 displays the average bond's characteristics for our sample in comparison to the initial cleaned sample (after the Dick-Nielsen (2009) and various price filters) over time. The average monthly dollar volume in our sample is $36.7m compared to $14.3m in the initial cleaned sample and the average number of trading days per year amount to 143 days in our sample compared to 55 days in the initial cleaned sample. Thus, the bonds in our sample are also relatively more liquid. Given that our focus is to study the significance of illiquidity for corporate bonds, such a bias in our sample toward more liquid bonds, although not ideal, will only help to strengthen our results if they already show up for the more liquid bonds.

[7]The share of bonds filtered out by these criteria changes according to the sequence in which the filters are applied, e.g., bonds not listed by the FISD exhibit a higher proportion of bonds with infrequent trading days.

However, we find that our data set's market coverage is at a considerably higher level compared to similar work since we do only exclude observations that are specifically unfeasible for our liquidity measures and subsequent analyses. Other research in the field tends to apply rough filters that facilitate the data preparation process, but drop actually useful observations. Representing the papers most comparable to our investigation with respect to data, Bao et al. (2011), Lin et al. (2011), and Dick-Nielsen et al. (2012) base their analyses on a sample of bonds that amounts to only 7.5%, 84.7%[8], and 38.8% of the number of bonds in our sample, respectively. To the best of our knowledge, there is only one study, which comprises a larger dataset than ours, namely that of Friewald et al. (2012) with more than 20,000 bonds. However, they use a different approach to get the inputs for the regression analysis. More specifically, they simply average each variable for every week and run regression on those averages regardless how many observation this week has. In contrast, we compute everything on a monthly basis (aggregating the daily observation), but require at least 10 trading days in a month; otherwise the bond month is filtered out. The problem associated with the approach of Friewald et al. (2012) in using weekly averages is that it might lead to very noisy estimates, since a weekly average might be inferred from only one trade. Another fact puts the large number of the bonds of Friewald et al. (2012) into perspective: the major difference of us is applying the filter of eliminating bond months with less than 10 trading days. This reduces the number of bonds by 60%, but the overall trading volume only shrinks by 13%. On the contrary, our approach reduces the risk of a few observations, but considerably distorting observations. Thus, when comparing our sample size with that of Friewald et al. (2012), not our final number of bonds after all data preparation steps would be of relevance, but the number of bonds before eliminating months with less than 10 trading days. Before this step, we have a total of 47,000 bonds and hence more than twice as many as Friewald et al. (2012). Concluding, we can say that we have one of the most comprehensive data sets for the U.S. corporate bond market in the academic literature and are consequently able to capture a very high fraction of the illiquid bonds.

[8]In light of the longer sample period of Lin et al. (2011) amounting to c.15 years compared to our sample of c.10 years, their number of bonds totaling 85% of the number of our sample appears to represent an even smaller share of the entire bond universe. Lin et al. (2011) are able to create such a long-span sample as they consolidate TRACE prior to its full implementation with the NAIC data set that is unfortunately not available to us. Besides, Mahanti et al. (2008) and Bao et al. (2011) argue that some issues with the construction of reliable illiquidity measures exist in the NAIC data.

Table 3: Summary Statistics over Time

This table reports summary statistics for our sample of bonds and for bonds in TRACE (after adjusting for the Dick-Nielsen (2009) and various price error filters (as described in Appendix A) over our sample period from 11/04 – 09/13. *#Bonds* is the number of bonds. *Rating* is a numerical translation of the credit rating agency's rating: 10=AAA, 9.5=Aa1/AA+, 9=Aa2/AA … 1.5=Caa2/CCC, 1=Caa3/CCC-, 0.5=Ca/CC, 0=C/D. *Original maturity* is the bond's total maturity in years. *Current maturity* is the bond's time left until the maturity date in years. *Coupon*, reported only for fixed coupon bonds, is the bond's annualized coupon payment in percent. *Volume* is the bond's monthly dollar trading volume in $m. *#Trading days* is the bond's total number of trading in the respective period. For each bond, we also calculate the time-series mean and standard deviation of its monthly log returns, whose cross-sectional mean is reported under *Return* and *Volatility* (both in %).

	Q4/2004	**2005**	**2006**	**2007**	**2008**	**2009**	**2010**	**2011**	**2012**	**2013**	**Average**
Panel A: All bonds reported in TRACE after the Dick-Nielsen (2009) and various price error filters											
#Bonds	10,552	19,649	19,041	20,965	19,647	19,251	18,124	21,107	17,712	15,017	18,107
Rating	7.1	6.9	6.9	7.2	7.1	6.5	6.5	7.0	6.4	6.5	
Original maturity	11.2	10.8	11.1	11.5	11.6	12.0	12.6	11.9	13.1	13.4	
Current maturity	8.3	8.0	7.9	8.2	7.9	8.0	8.2	8.2	8.7	8.7	
Coupon	5.71	5.68	5.63	5.54	5.32	5.35	5.43	4.96	5.18	5.16	
Volume ($m)	17.2	12.0	12.1	11.2	10.1	14.3	15.6	15.1	17.4	18.3	14.3
#Trading days	19	54	55	48	45	60	70	62	71	68	55.3
Return (in %)	0.58	-0.01	0.68	0.21	-0.82	1.70	0.58	0.33	0.54	0.01	
Volatility (in %)	2.48	3.26	2.49	2.54	6.89	5.36	2.07	2.10	1.59	1.70	
Panel B: Bonds in our sample											
#Bonds	2,582	4,325	4,266	3,883	3,450	5,235	5,803	5,864	6,051	5,840	4,730
Rating	6.6	6.1	6.1	6.3	6.3	5.8	6.0	6.2	6.1	6.2	
Original maturity	11.1	10.6	11.3	12.3	12.2	12.6	13.4	13.1	13.1	12.8	
Current maturity	7.7	7.4	7.7	8.5	8.2	8.6	9.2	9.0	9.0	8.7	
Coupon	6.38	6.34	6.34	6.33	6.30	6.26	6.21	6.00	5.66	5.37	
Volume ($m)	48.2	38.2	35.8	37.7	38.2	41.0	39.0	39.2	39.7	39.7	39.7
#Trading days	44	155	152	150	156	155	162	163	158	134	142.9

4 Methodology

4.1 Liquidity Measures

The dimension of liquidity that we focus on in our work represents the ease of finding counterparties to trade with. For a given bond, this ease of trading differs with respect to type of order, e.g., for a given bond, a block trade may be easier to carry out, while executing orders in a short period of time may be more difficult. As a consequence, liquidity comprises three dimensions: tightness represents the magnitude of the bid-ask spread and measures the implicit "transaction costs" charged by market makers, depth is the degree to which one can execute orders immediately with little market impact, and resilience describes how fast prices recover to "equilibrium" levels after trading shocks. Since the latter two are very closely related, they are often subsumed under the term of market impact. In over-the-counter markets, data about the bid-ask spreads is often not available for researcher and market impact is per se not directly observable; therefore numerous studies have proposed diverse measures to capture and proxy the different dimensions of liquidity. Due to the sheer number of different liquidity measures, there are only a few studies aiming at answering the question which liquidity measure performs best. Examples for such studies are Goyenko et al. (2009) for exchange traded stocks, Marshall et al. (2012) for commodity futures, and – more importantly for us – Schestag et al. (2013) for corporate bonds. Importantly, Schestag et al. (2013), besides other authors, find that all liquidity measures show a high level of co-movement, irrespective of the dimension they measure, and that low-frequency proxies (which are based on daily price levels) show similar results to more sophisticated high-frequency measures (which are based on intra-day data).

One caveat about applying liquidity measures on corporate bonds needs to be mentioned since almost all liquidity measures have initially been developed for the stock market. The vital difference of bonds to stocks is its lower trading frequency, which makes it rather impossible to determine high-frequency measures. Therefore, we focus on liquidity proxies on daily price levels, which are generally not inferior to their more sophisticated high-frequency counterparts as described above. We select our specific measures among the wide pool of proxies according to either high estimation power in comparative studies or the ability to detect statistically and economically significant liquidity premia in papers directly investigating corporate bonds. In order to obtain a comprehensive set of measures, we chose two proxies for both the market impact and the bid-ask spread each. The market impact shall be captured by the Amihud (2002) as well as the Pastor-Stambaugh (2003), and the bid-ask spread by the Roll (1984) and a proprietarily extended version of the Corwin-Schultz (2012) measure[9].

The Amihud measure is ranked among the best proxies in studies of Goyenko et al. (2009) and Fong et al. (2013) concerning stocks as well as in the study of Schestag et al. (2013) for corporate bonds and is successfully used, among others, in the work of Dick-Nielsen et al. (2012) and Friewald et al. (2012). The Pastor-Stambaugh measure convinces by capturing significant liquidity effects in the article of Lin et al. (2011), which is highly comparable to our work with respect to data and the level of time horizon. Schestag et al. (2013) also states the Roll and Corwin-Schultz measures as very accurate liquidity estimators. While the Roll measure is also effective in the analysis of Friewald et al.

[9] For ease of reading, we forego in the following to name the year in which the measure was developed.

(2012), the Corwin-Schultz measure additionally convinces by its simplicity of reasoning in deriving an easily programmable, closed-form solution for the bid-ask spread[10].

For creating liquidity measures and analyzing their effects on asset prices, it is important to pay close attention to the sign and time horizon in which the measure is defined. We specify the measures as *il*liquidity measures rather than liquidity measures: They are monotone positive and increase in illiquidity, not in liquidity, i.e. the higher a measure the more illiquid the bond is. Consequently, we describe our measures as "illiquidity measures" in the following. The second result-determining factor is the time dimension. We focus on monthly price changes in bonds and, importantly, ascribe for every bond the illiquidity measure of a month to the return of the exact same month. In comparison, other studies assign, for example, illiquidity measures of the three previous months to a month's return. The difference appears at first to be small, but has major effects. Our definition leads to an analysis of actual returns, i.e. how returns change according to the illiquidity characteristics. For example should we expect that returns *decrease* when the illiquidity of a bond increases. The second definition aims at expected returns, i.e. how returns compensate for an expected illiquidity characteristic that co-moces strongly with its historical magnitude. For example should we then assume that returns are *higher* when the illiquidity of a bond has been high during the preceding months.

The lack of frequent trading of corporate bonds raises another issue. The Amihud, Pastor-Stambaugh, and Roll measures are based on the daily return of a security which is typically defined as the return between two consecutive trading days. For ease of reading, we define the most recent preceding day with any trading activity of a specific bond as its *previous day of trading*, so, for example, the *previous day of trading* for a day can be actually five days ago if there does not occur any trading during the four days in between. Due to the infrequent nature of corporate bond trading, we define the return on a day as the relative price change of the *previous day of trading* to the current day. Consequently, the return is not necessarily defined over consecutive trading days. Recalling that a bond's clean return for a trading day also includes the bond's drift towards its principal over time (in case that the bond is not at par), a return defined over t trading days also includes t times the daily drift. So if our daily return is, for instance, actually defined over several trading days, this return includes price drifts of more than one day. As a result, the illiquidity measure of a discount bond (which exhibits a positive drift) would be calculated on the basis of a higher return in case of trading days without any trading than in case of trading activity every day, all else being equal. In order to limit this bias in our liquidity measures and to have enough illiquidity estimates per month to derive a monthly illiquidity measure, we impose the requirement of more than ten trading days per month for including a bond's month in our sample. As, on the other hand, the Corwin-Schultz measure is specifically based upon the idea of the variance of a return being proportional to its return period, the approximation of treating the *previous day of trading* as the preceding trading day would distort the measure in its fundamental idea. Consequently, we derive an extended closed-end formula of the Corwin-Schultz estimator in order to specifically account for the number of days that lie between two days of actual trading activity.

[10]Except for the liquidity measure comparison study of Schestag et al. (2013), the Corwin-Schultz measure has not yet been applied to the corporate bond market, but found application in liquidity studies for options (Deuskar et al., 2011) and for stocks (Karstanje et al. (2013) and Kim and Lee (2014)).

4.1.1 Amihud Price Impact Measure

The illiquidity measure developed by Amihud (2002) is one of the most widely used illiquidity measures and was originally developed for the equity market and is conceptually based on Kyle's (1985) λ [11]. Illiquidity for a security is high if a large volume can only be traded with a large price impact, and vice versa. The intuition behind the measure is that a security is more illiquid if there are less market participants willing to absorb any order flow so they demand a higher price change in their favor in order to trade. Consequently, this security shows a higher absolute price change for every U.S. dollar of volume traded so that the daily Amihud measure is the ratio of the absolute return of a security to its volume traded on a given day. We then average the daily measures up to the monthly Amihud estimator of

$$AM_{i,m} = \frac{1}{D_m} \sum_{d=1}^{D} \frac{|r_{i,d,m}|}{V_{i,d,m}}, \tag{2}$$

where D_m is the number of days for which security i is traded in month m , $r_{i,d,m}$ is the return on security i on day d of month m , and $V_{i,d,m}$ is the non-zero dollar trading volume of security i on day d of month m.

4.1.2 Pastor-Stambaugh Price Reversal Measure

Pastor and Stambaugh (2003) develop an illiquidity measure which captures temporary price changes associated with order flow. It aims at capturing the extent to which a price recovers to its "equilibrium" level after a trading induced shock. The rationale is that the more an initial price change reverses, the more it is just a temporary price distortion caused by order flow with too few absorbing counterparties, rather than a fair adjustment to the fundamental value of the security. Hence, the reversion is stronger the more illiquid the stock is. That is, the bigger the liquidity motivated price impact on the day before.[12] Alike the idea underlying the Amihud measure in which the return is proportional to trading volume for a given level of illiquidity, the price reversal is assumed to be proportional to the volume of the day before for securities with the same illiquidity. Only if the price reversal is stronger in proportion to its causing trading volume, the security is considered more illiquid. The monthly Pastor-Stambaugh measure is defined by the coefficient γ in the following regression

$$r^e_{i,d+1,m} = \theta_{i,m} + \phi_{i,m} r_{i,d,m} + \gamma_{i,m} sign\left(r^e_{i,d,m}\right) V_{i,d,m} + \varepsilon_{i,d+1,m} \qquad d = 1, 2, ..., D, \tag{3}$$

where $r^e_{i,d+1,m}$ is the bond's return in excess of the risk free rate of day $d+1$ in month m , $r_{i,d,m}$ is the return of bond i on the *previous day of trading* of $d+1$, $sign\left(r^e_{i,d,m}\right)$ is an indicator whether $r^e_{i,d,m}$ is positive or negative, and $V_{i,d,m}$ is the bond's dollar volume on the *previous day of trading* of

[11] Hasbrouck (2009) shows that the Amihud measure and Kyle's λ have a correlation of 0.82.

[12] The insidious thing about logic itself is that the mere absence of illogicality does not necessarily lead to the truth. In this example, one could argue that the more a price reverses the next day, the *more* liquid the security is because then an initial price shock would be absorbed to a higher extent identifying a higher number of trading parties. However, this argumentation bears the hidden assumption of the initial price change being independent of the liquidity of the security. The Pastor-Stambaugh measure assumes that the initial price distortion is related to liquidity, but its reversal during the next day is not.

$d+1$. The sign indicator of the previous day's return is necessary to define the sign of price reversal. If γ is negative it captures a price reversal, otherwise it captures a momentum effect. In the latter case we set γ to zero for our illiquidity measure, otherwise we defined the monthly Pastor-Stambaugh illiquidity measure as the negative of γ:

$$PS_{i,m} = max\left(-\gamma_{i,m}, 0\right). \tag{4}$$

4.1.3 Roll Bid-Ask Spread Estimator

While the two measures just described above focus on the price impact, the Roll measure serves as a proxy for the bid-ask spread, i.e. the direct costs associated with a transaction. A long time before liquidity actually became again a hot topic in finance during the 2000s, Roll (1984) bases his spread measure on the well-known statement that, if markets are efficient and trading costs are zero, prices fluctuate randomly and consequently show no serial dependence in successive price changes (aside from that resulting from the serial dependence of expected returns). If, however, a bid-ask spread exists prices will jump back and forth between the ask and bid quota according to whether the trade is buyer or seller initiated. This bid-ask bounce results in negative serial correlation among transactions and also across trading days. The larger the negative serial correlation is, the larger the bid-ask spread of the security is. Consequently, the monthly Roll measure is

$$RL_m = 2\sqrt{-Cov\left(r_{i,d,m}, r_{i,d-1,m}\right)}, \tag{5}$$

where $r_{i,d,m}$ is the return of security i on day d in month m. The monthly covariance is determined on basis of all traded days within the respective month. If the covariance is positive, we set the Roll measure to zero.

4.1.4 Extended Corwin-Schultz Bid-Ask Spread Estimator

Corwin and Schultz (2012) just recently developed a new bid-ask proxy based on the simple idea that the variance increases proportional in time. Considering the high-low spread, i.e. the difference in the highest to the lowest price within a given trading day, they assume that high prices are a result of a buy order while low prices correspond to sell orders. Consequently, the spread between the high and low prices during a day represents the daily variance plus the bid-ask spread. Since the variance is proportional to the return interval, the high-low spread over two consecutive days (i.e. the difference in the highest to the lowest price over both days), represents twice the daily variance, but just once the bid-ask spread. This simple, yet brilliant idea allows solving for both the spread and the variance by deriving two equations, the first as a function of the high-low ratios on 2 consecutive single days and the second as a function of the high-low ratio from a single 2-day period. This results in a closed-end formula for the bid-ask spread.

It has to be noted that the Corwin-Schultz measure was actually created for stocks and assumes a geometric Brownian motion. The bond's characteristic of pull-to-par (i.e. the bond's price approaching its par value over time) contradicts the random walk, however we believe that for short time periods of several days, the pull-to-par effect can be neglected, so that we can also assume the geometric Brownian motion for daily bond returns. Since the pivotal idea of the Corwin-Schultz measure is that the variance of a return doubles over two consecutive days, the measure is only applicable to two consecutive days with trading. Due to infrequent trading nature of corporate bonds, this restriction would limit our claim on scope. Applying the original formula simply on the closest days with any trading activity (i.e. disregarding the trading days without any trading activity in between) would distort its central idea. Paul Schultz, contacted by us in March 2014, confirms our remark: "It is hard to estimate liquidity with corporate bonds because they trade infrequently. Hence I think our estimator would yield very noisy estimates". Consequently, it is crucial to account for the fact of trading days without any trading activity. Therefore, we derive an extended closed-end formula that is flexible to the number of trading days that lie between two days with trading activity. If, for example, trading only occurs on a Monday and Friday in a given week, the high-low spread over the days of Monday and Friday represents once the bid-ask spread and five times the daily variance instead of only twice the daily variance as in the original Corwin-Schultz measure. Appendix B shows the detailed derivation of our extension, which leads to the Extended Corwin-Schultz measure

$$CS_i = \frac{2\left(e^{\alpha}-1\right)}{1+e^{\alpha}}, \tag{6}$$

where

$$\alpha = \frac{\sqrt{\beta}-\sqrt{\gamma}}{\sqrt{T}-1} \tag{7}$$

$$\beta = \sum_{j=0}^{T-1}\left[ln\left(\frac{H^0_{t+j}}{L^0_{t+j}}\right)\right]^2 \tag{8}$$

$$\gamma = \left[ln\left(\frac{H^0_{t,t+T-1}}{L^0_{t,t+T-1}}\right)\right]^2, \tag{9}$$

Where $H_{t;t+T-1}$ is the highest price over the T days of $t; t+1; \ldots; t+T-1$ and $L_{t;t+T-1}$ is the lowest price over the T days of $t; t+1; \ldots; t+T-1$[13]. In case of just a single trade per day (leading to a high-low ratio of 1), we set the bid-ask proxy to missing. We define the monthly Extended Corwin-Schultz measures as the average of the daily estimates.

We take the trading volume-weighted average of every bond-specific illiquidity measure in order to derive market-wide illiquidity measures. Since the scales of the illiquidity measures are arbitrary, we standardize every (bond-specific and market-wide) measure by its mean to allow for comparability.

[13]We use observed high-low ratios to calculate the spread proxy, while the estimator is actually derived using expectations. As the variance and bid-ask spread are non-linear functions of the high-low spread, the average spread estimates are biased. Corwin and Schultz (2012) show that this does not affect the outcome of their estimator to a large extent.

4.2 Fama-MacBeth Procedure

4.2.1 Risk Source Identification

Liquidity affects asset prices in various ways. As already outlined in the literature review, many studies have examined the effect of a bond's specific liquidity (liquidity level), while recent studies have shifted their attention to whether there is a premium associated with the co-movement of a bond's return with market wide liquidity shocks (liquidity risk).

As Bongaerts et al. (2012) show, both liquidity level and liquidity risk are typically highly related so that omitting one of the two could lead to distorted results. Consequently, we include both types of liquidity of relevance for asset pricing, but want to point out that liquidity level is a characteristic, while liquidity risk describes a factor. This difference has the implication that a priced characteristic is compensated directly, by the mere fact of featuring it, while a priced factor leads only to a premium if the asset co-moves with the factor. Hence, not the level of the factor is compensated, but the specific exposure of an asset's return to it. Applied to our case, we try to analyze whether, on the one hand, the specific illiquidity (liquidity level) and, on the other hand, the bond's sensitivity to market-wide liquidity shocks (liquidity risk) lead to a higher compensation.

4.2.2 Risk Factor Sensitivities

In order to isolate the effect of liquidity, we follow Elton et al. (2001) and Lin et al. (2011) by adjusting for five common risk factors, namely the term (TERM), default (DEF), market (MKT), size (SMB) and book-to-market (HML) factor. In their seminal paper of 1993, Fama and French show that those five risk factors seem to explain average returns on both stocks and bonds, while the first two appear to be intuitively connected to bonds (Gebhardt et al., 2005). TERM describes the risk from unexpected changes in the term structure of interest rates, while DEF arises from changes in the overall default probability of the bond market in response to changing economic conditions. The latter three represent the well-known Fama-French three factors of the standard stock pricing model which is why they do not need further explanation here. However, they also affect bond returns since both bond and stocks are the investor's claims on the same underlying assets of the firm so that common stock market factors such as market, size and book-to-market should have spill-over effects to bonds.

In order to investigate whether liquidity is important in bond asset pricing, we perform the Fama-MacBeth (1973) procedure, which estimates the premium rewarded for every risk factor exposure. The procedure itself consists of two steps, a bond-specific time series regression to determine every bond's sensitivities to the risk factors and a cross-section regression of the factor betas and characteristic on expected returns to derive risk premia (we use the term "risk source" to describe both of the aforementioned characteristics and factors). First, we estimate betas of the identified risk factors in a time-series regression of a linear factor model for every individual bond

$$\begin{aligned} r_{it} - r_{ft} = \alpha_i + \beta_{iTERM} TERM_t + \beta_{iDEF} DEF_t + \beta_{iMKT} MKT_t + \beta_{iSMB} SMB_t \\ + \beta_{iHML} HML_t + \beta_{iLIQ_{mkt}} LIQ_{mkt,t} + \varepsilon_{it}, \end{aligned} \tag{10}$$

where $TERM$ is the difference in the monthly 30-year T-bond return and one month T-bill returns, DEF is the difference between the monthly return on a value-weighted portfolio of our entire bond sample and the average return on government bonds, i.e. 30-year T-bond minus one month T-bill (all government bond returns are obtained from Datastream). The Fama-French three factors MKT, SMB, and HML are obtained from Kenneth French's data library, and therefore are defined as described in Fama and French (1993). The liquidity factor, LIQ_{mkt}, is the negative of the residual from an AR(1) process on the market-wide illiquidity measures as described in section 4.1. We will run the regression for each of our four illiquidity measures separately. Following many studies (such as Pastor and Stambaugh (2003), Acharya and Pedersen (2005); and Bao et al.(2011)), we aim at eliminating the auto-regressive moments of the market-wide illiquidity that could be introduced by serial correlation of the returns on which they are defined. The Bayesian information criteria estimates an AR(1) process as most adequate for all of the four illiquidity measures. We take the negative of the illiquidity measure in order to define the factor in the dimension of liquidity (and not *il*liquidity), since co-movements with the market-wide liquidity introduce cyclicality so as to create a positive (and not negative) risk premium.

4.2.3 Risk Premia

The second step of the Fama-MacBeth procedure runs a cross-section regression of the bonds' risk sources (factor betas and characteristics) on its expected returns for every month separately

$$\begin{aligned} E_t(r_i^e) = \alpha_i + \lambda_{t,TERM}\beta_{i,TERM} + +\lambda_{t,DEF}\beta_{i,DEF} + \lambda_{t,MKT}\beta_{i,MKT} + \lambda_{t,SMB}\beta_{i,SMB} \\ + \lambda_{t,LIQ_{mkt}}\beta_{i,LIQ_{mkt}} + \lambda_{t,LIQ_{lvl}}\beta_{i,LIQ_{lvl}} + \varepsilon_{it}, \end{aligned} \tag{11}$$

where λ represents the coefficient of each risk source in the cross-section. The rationale is that, assuming the risk source is priced, a higher factor exposure, i.e. beta/higher characteristic should also be compensated by a higher expected return. However, if the risk source is not priced by the market participants, a high beta of the factor/high characteristic should not lead to a higher expected return so as to result in an insignificant or zero lambda. As investors should expect a certain return conditional on the beta up to that date, one should actually run the cross-section regression of expected returns on the beta of some period before that month. Fama and MacBeth (1973) use a five-year rolling window, however, we simplify by using the beta of the entire sample period for the regression in every month because the accurate procedure would trim the time period of the regression and as Cochrane (2005) states that the simplified version does not distorts the results by a significant extent. As mentioned above we also include liquidity level, LIQ_{lvl}, as a risk source in the regression here[14]. The regression results in one coefficient (risk premium) for every risk source and for every month. We then average the coefficients for every risk factor over time to derive an average risk premium per factor for our sample.

[14]Since adjusting every company-specific liquidity level for serial auto-correlation would result in an extreme computational demand, we neglect this adjustment here.

4.2.4 Expected Returns

Whenever a risk factor is priced, it should become visible in the expected returns. However, most tests of asset pricing models and factors use actual returns as proxies for expected returns. It implicitly assumes rational expectations (Campbell et al., 1998), i.e. that the average realization is a good proxy for the expectation. This practice is justified on the grounds that in sufficiently long time periods, average actual returns will "catch up and match" the expected returns. However, this assumption has significant limitations. Lundblad (2005) and Pastor et al. (2008) show via simulations that actual returns do not necessarily converge to expected returns, except for very long time periods. Furthermore, Elton (1999) observes that there have been time periods of more than 10 years in which realized stock returns are lower than the risk-free rate (1973-1984) and periods of 50 years in which high-yield bonds underperform the risk-free rate (1927-1981).

Especially for our sample of corporate bonds, we expect that historical returns are extremely noisy because of the skewed nature of bond returns due to the rare occurrence of the very severe event of default. Furthermore, our sample period of 9 years can be considered relatively short. Another source of noise is created by the subprime crisis in our sample. During an unexpected recession, bond returns are typically negative, while increasing risk premia (or increasing co-movements to risk factors) typically *increase* expected returns. Hence, actual and expected returns move in opposite directions to an economically significant extent.

For all those reasons, we rely on the approach to determine expected returns from yield spreads which was first formulated by Elton et al. (2001) and then applied by only a small group of researchers such as Campello et al. (2008), Huang and Huang (2012), and De Jong and Driessen (2012), who show that it yields much more accurate estimates than historical averages[15]. To the best of our knowledge, Bongaerts et al. (2012) is the only paper using this approach to identify the effects of liquidity. The approach derives the expected excess return $E(r^e_{it})$ from taking the observed credit spread (yield-to-maturity) directly submitted to TRACE and correct it not only for the corresponding government yield, but also for the major driver of bond returns, the expected default loss. Hence this expected excess returns is closely related to the "credit spread puzzle" since it describes the part of the expected return that is not explained by the time value of money or credit risk caused by the bond issuer. This expected excess return is then defined as

$$E_t(r^e_i) = (1 + y_{it})(1 - L\pi_{it})^{1/T_{it}} - (1 + r_{ft}), \tag{12}$$

where y_{it} is the observed yield-to-maturity, L is the loss given default, π_{it} is the cumulative default probability over the entire maturity and T_{it} is the duration of the bond i, while r_{ft} represents the yield of the government bond equal to the maturity of the bond. The bond's duration is used as the bond is approximated by a zero-coupon paying bond with the same duration in order to circumvent the tedious

[15] If expected returns reduce the noisiness of realized returns, it might appear reasonable to also use them in the first step of the Fama-MacBeth (1973) procedure in order to increase the accuracy in estimating factor betas. This is, however, a fallacious conclusion. Consider, for example, that a bond has a constant beta over time to only one risk factor. Hence, its actual returns co-move with that factor to always the same extent. This should result in the expected return of the bond staying constant over time as the inherent risk exposure (the factor beta) does not change. Consequently, the realized return shows a beta loading, while the expected return does not.

issue of accounting for coupon payments. Assuming that default losses are incurred only at maturity, the expected return from holding the security until maturity is$(1 + y_{it})^{T_{it}(1-L/pi_{it})}$. Annualizing this and subtracting the government bond equal to the maturity of the bond leads to formula (10). We obtain the cumulative default probabilities from Standard and Poor's, which are based on the time period of 1981 to 2005 and can be found in Table 9 in the Appendix[16]. The cumulative default probabilities incorporate both the probability of directly defaulting from the current credit rating, but also the probability to first deteriorate in rating and subsequently default. They typically increase with maturity since a longer time period makes it generally more probable for the bond to default. As the information of default probabilities for investors should stem from a time period *before* the actual determination of expected returns, we use those default probabilities prior to our sample period (there is however an overlap for the years 2004 and 2005 as we could not access summary default probabilities on the basis of information before 2004). Our use of default probabilities introduce a certain imprecision since we do not update the historical rates on a rolling basis (e.g. an investor in year 2008 also incorporates the default rates of 2006 and 2007) and as more detailed forward-looking estimates of default probabilities exist (e.g. Moody's-KMV database to which we unfortunately do not have access). As in Bongaerts et al. (2012), we assume a constant loss given default of 60%[17].

4.3 Identification of Liquidity Stress Periods

To further deepen our analysis of liquidity in distressed times, we want to investigate whether there are flight-to-liquidity phenomena in the U.S. corporate bond market during our sample period. A flight-to-liquidity describes the sharply decreased willingness of investors to hold illiquid assets, which is why they "fly" from such illiquid assets towards very liquid assets that can easily be sold when funding is needed.

One of the most delicate parts during a study of a flight-to-liquidity – besides its exact definition – is the identification of periods where a flight-to-liquidity occurs. Most studies that study liquidity in distressed times, put liquidity stress periods on a level with the time periods of financial crises (see, e.g., Chordia et al., 2005; Dick-Nielsen et al., 2012; and Friewald et al., 2012).This approach, however, may not be suitable for an investigation of flight-to-liquidity periods, since these usually are of relatively short duration and create the problem of determining the exact end of the stress period.

Given these weaknesses, we make use of the burgeoning applications of Markov regime-switching models in empirical finance[18] and endogenously identify liquidity stress periods with the help of Markov regime-switching models, first proposed by Hamilton (1989). The implicit assumption of Markov regime-switching models is that the data is driven by a process that undergoes abrupt changes,

[16]Surprisingly, the Canadian Gas Transmission appears to be very interested in default characteristics and offers the report for download online under http://www.gastransmissionnw.com/rate_case_filings/documents/SPGTN18.pdf.

[17]Elton et al. (2001) and De Jong and Driessen (2012) use rating-specific recovery rates from Altman et al. (2000). However, we think that this exerts some flaws. Consider a recovery rate of e.g. 60% for AAA rated bonds. This recovery rate is only suitable if the bond defaults while having the AAA rating. If the bond deteriorates to BBB and then defaults, the recovery rate of BBB ratings is adequate. Since the cumulative default probability does not only incorporate the probability of directly defaulting from the current credit rating, but also the probability to first deteriorate in rating and subsequently default, it is not possible to specifically ascribe a bond to a rating-specific recovery rate. For that reason, we assume a constant recovery rate of 40% (i.e. loss given default of 60%) for all ratings.

[18]See Guidolin (2011a, 2011b) for a survey of this literature and Hamilton (1994, 2008) for concise introductions to Markov regime-switching models.

induced for example by exceptional events. This means that we ask the data directly to tell us when the system is likely to be in a liquidity stress regime and do not derive it from other indirect (e.g. macroeconomic or financial) indicators. Consequently, we derive the regimes for the liquidity stress periods in purely statistical terms without any apparent underlying economic intuition, what also creates a problem related to endogeneity. Nonetheless, this approach is to our belief superior to a more naïve one for identifying stress times[19]. In addition, we will check the plausibility of our identified stress periods using economic reasoning.

We follow Acharya et al. (2013) in estimating a number of Markov regime-switching regressions. The intuitive idea behind this approach is that in a regression model the slope and intercept coefficients may follow a Markov regime-switching dynamics in times of normal and distressed times. The regime-switching regression results in probabilities of being in a particular regime at every point in time of the sample period. Unlike most other studies, which leave the modeling at that, we use the result of these regressions to determine the concrete most probable regime path (opposed to abstract probabilities) over our time horizon by using two powerful algorithms, namely the Baum-Welch algorithm and the Viterbi algorithm. This path returns for every months in our sample the predicted regime by the Markov regime-switch model on which we can further analyze possible flights-to-liquidity. We describe the entire procedure in detail in the following subsections.

4.3.1 Markov Regime-Switching Regression

We estimate a Markov regime-switching model for corporate bond betas, allowing the liquidity risk beta coefficient of our corporate bond pricing model to vary between two regimes, which should represent normal and stress liquidity periods. Moreover, we allow the variance of the normally distributed error term to change between the two regimes, which represents a higher uncertainty regarding the predictive power of the model in each state of the world. Concluding, this means that we solely allow the liquidity coefficient of the time-series regression (8) to switch between two regimes, while the other coefficients remain constant. This is reasonable as investors should be more sensitive towards market-wide liquidity during flight-to-liquidity periods. Formally, the Markov regime-switching regression can be represented in the following form:

$$\begin{aligned} r_{it} - r_{ft} = \alpha_i &+ \beta_{iTERM} TERM_t + \beta_{iDEF} DEF_t + \beta_{iMKT} MKT_t + \beta_{iSMB} SMB_t \\ &+ \beta_{iHML} HML_t + \beta^{s_t}_{iLIQ_{mkt}} LIQ_{mkt,t} + \varepsilon^{s_t}_{it} \qquad with \;\; \varepsilon_{it} \sim N\left(0, \sigma^2_{s_t}\right), \end{aligned} \tag{13}$$

where all coefficients and factors are defined as in (8) and the superscript index s_t indicates whether the coefficient is allowed to switch between regime $s_t \epsilon \, (1, 2)$. The transition between regimes is stochastic in Markov regime-switching models, so that there is no clear answer to whether a flight from one to the other regime actually occurs or not. It is, however, assumed that s_t follows a homogeneous first order Markov chain[20] with constant transition probabilities, which can be represented in the following

[19] See, for example, Barrell et al. (2010) for a discussion about the problems associated with exogenous crisis identification.

[20] A first order Markov chain is a special kind of time-discrete stochastic process that undergoes transitions from one state to another (in a finite state space). Furthermore, a first order Markov chain is characterized by the fact that the probability of a certain observation x_t at time t depends only on the most recent observation x_{t-1} at time $t-1$ and hence is independent from the sequence of observation that preceded it, which is the so-called Markov property.

transition matrix

$$\mathbf{P} = \begin{bmatrix} \mathbb{P}(s_t = 1|s_{t-1} = 1) & \mathbb{P}(s_t = 2|s_{t-1} = 1) \\ \mathbb{P}(s_t = 1|s_{t-1} = 2) & \mathbb{P}(s_t = 2|s_{t-1} = 2) \end{bmatrix} = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix}, \tag{14}$$

where p_ij $(i, j = 1, 2)$ denote the transition probabilities of $s_t = j$ given that $s_{t-1} = i$. This means that the probability of a switch from state 1 to state 2 between time $t-1$ and t will be given byp_{12}[21]. Likewise the probability of staying in state 1 is given by p_{11}. The second row of P can be interpreted analogously. Given this, the transition probabilities of each row of the transition matrix have to sum up to 1, since each row represents the full probabilities of the process for all states. For an intuitively accessible illustration of these relations the reader is referred to Figure 1, which displays a state diagram for a two-state Markov model.

Figure 1: State Diagram for a Two-State Markov Model

This figure shows the state diagram for a two-state Markov model. The circles in the diagram represent the two possible states of the process and the arrows represent the transitions between the states. The label on each arrow represents the probability of that particular transition.

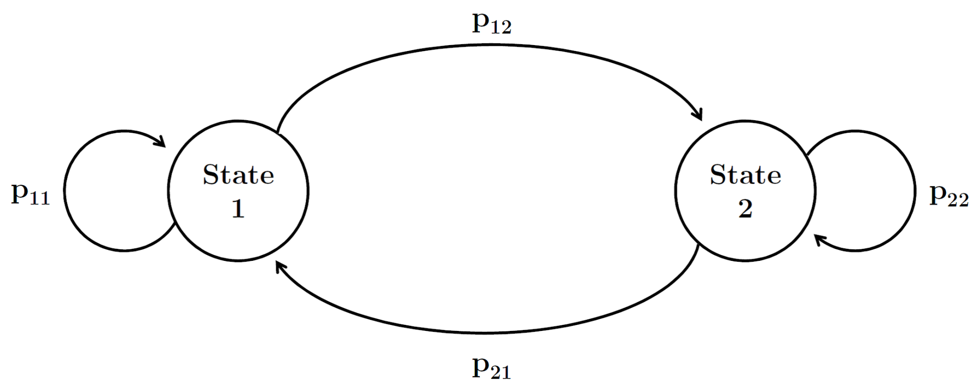

The parameters of Markov regime-switching models are generally estimated via MLE and since analytical solutions due to the non-linearity of the problem do not exist, solutions are typically derived via direct numerical maximization or via algorithms. The parameters of our model are estimated by the MATLAB Markov regime-switching routine of Perlin (2012), which directly maximizes the log likelihood function (see Perlin (2012) for further details on this procedure) and is based on the Hamilton filter – the most popular method of regime-switching calibration in economics and finance. Since the parameter estimation via algorithms seems to be superior to the direct numerical maximization by employing the Hamilton filter (see Mitra and Date, 2010) and since we are more interested in a binary series of regimes identified by the regime-switching model, we undertake a few more steps using the abovementioned algorithms.

[21] In line with the academic literature, we will use the terms regime and state interchangeably.

4.3.2 Determination of Regime Path

In order to obtain such a binary regime series, we have to apply the so-called Viterbi algorithm (Viterbi, 1967), for which we have to introduce hidden Markov models (HMMs)[22], which usually find less application in economics or finance and more in speech recognition (e.g. Rabiner, 1989), gene finding (e.g. Burge and Karlin, 1997), protein secondary structure prediction (e.g. Krogh et al., 2001), and volcano-seismic signal detection (e.g. Ibáñez et al., 2009), among others.

Although HMMs and Markov regime-switching models describe the same kind of models in the sense that one observes a set of variables (in our case regression coefficients) and aims for determining the unobservable (hidden) states (in our case liquidity stress period or normal period), we have to make some adaptions in order to use HMMs in our context. As already mentioned, our regime-switching routine maximizes the log likelihood function directly – with all its weaknesses – which is why we refine the estimation of our model parameters by using the Baum-Welch algorithm (Baum et al., 1970) – a special form of the Expectation Maximization (EM) algorithm for HMMs. Before going into detail about the algorithm itself, we first have to introduce HMM terminology (for the two-state case), which we will use afterwards.

The variables relevant for our application of HMMs are the state sequence, which is the (hidden) sequence of states (liquidity stress period or normal period) over the sample period:

$$s = \{s_1, s_2, ..., \; s_{T-1}, \; s_T\} \qquad s_t \,\epsilon\, (1, 2), \tag{15}$$

the observation sequence, which is defined as the sequence of observations (represented as integers since not the manifestation of the observation itself is of interest but the fact whether the observations are different in value), over the whole sample period:

$$x = \{x_1, x_2, ..., \; x_{T-1}, \; x_T\} \qquad x_t \,\epsilon\, (1, ..., N), \tag{16}$$

the transition probabilities, defined as before as the probabilities of going from one state to the other or staying in one state:

$$\mathbf{P} = \begin{bmatrix} \mathbb{P}(s_t = 1|s_{t-1} = 1) & \mathbb{P}(s_t = 2|s_{t-1} = 1) \\ \mathbb{P}(s_t = 1|s_{t-1} = 2) & \mathbb{P}(s_t = 2|s_{t-1} = 2) \end{bmatrix} = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix}, \tag{17}$$

and the emission probabilities, defined as the likelihood of a certain observation x_t, if the model is in state s_t, which in matrix notation is equal to a $s \times N$ matrix[23]:

$$\mathbf{E} = \begin{bmatrix} \mathbb{P}(x_t = 1|s_t = 1) & ... & \mathbb{P}(x_t = N|s_t = 1) \\ \mathbb{P}(x_t = 1|s_t = 2) & ... & \mathbb{P}(x_t = N|s_t = 2) \end{bmatrix} = \begin{bmatrix} e_{11} & ... & e_{1N} \\ e_{21} & ... & e_{2N} \end{bmatrix}. \tag{18}$$

In order to illustrate all these terms and formulas more descriptively, Figure 2 shows the general

[22] Strictly speaking, hidden Markov models and Markov regime-switching models describe the same kind of models with the same underlying idea, but with different fields of application and slightly different methodologies due to different needs.

[23] In fact, in our case the observations would be continuously valued (i.e. $x_t \epsilon \, \mathbb{R}$), which is why we would get probability density functions over the observation space for the system being in state s_t. Such emission probability density functions would usually need to be parameterized, e.g, by Gaussians mixture models. However, for simplicity reasons and negligible loss of generality, we assume the values of the observations to be discrete.

structure of HMMs graphically.

Figure 2: Basic Structure of a Hidden Markov Model

This figure shows the basic structure of a hidden Markov model. A HMM is firstly defined by the transition probabilities represented by a Markov chain, which determines the (hidden) state s_t at time t and is secondly defined by the emission probabilities, which determine the observation x_t at time t. The light circles in the graph represent the hidden states and the darker circles represent the observation process, where each observation depends *only* on the present (hidden) state of the Markov model.

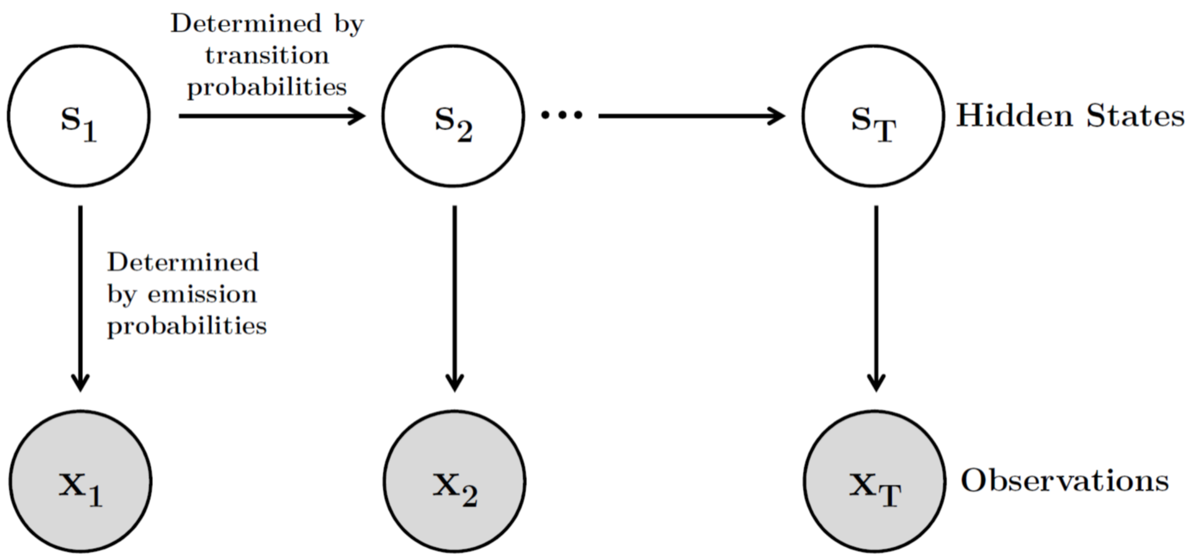

Baum-Welch Algorithm

The Baum-Welch algorithm is basically a method that estimates the parameters of a HMM when the state sequence is unknown, what means that it tries to set the parameters (i.e. transition and emission probabilities) of the HMM in such way that they model a phenomenon best possible. Consequently, the Baum-Welch algorithm itself is capable of finding a solution to a HMM just like the Hamilton filter. However, we use the Baum-Welch algorithm more as a training for our HMM, what means that we let it improve our estimates from the Markov regime-switching regression (adapted to HMM terminology) before we go on to the next step, which consists of the determination of the regime path by applying the Viterbi algorithm.

The functioning of the algorithm is as follows. In a first step, the probabilities of realization of an arbitrarily model with some initial transition and emission probabilities are calculated. During the calculation, the algorithm records how often transitions and particular emissions were used. In a second step, the parameters of the HMM are again calculated, but with the difference that transitions and emissions used more often get assigned a higher probability while transitions and emissions used less often get assigned a lower probability. These two steps are iterated until the parameters of the model do not change substantially or if the maximum number of predefined iterations is reached.

Hence, the Baum-Welch algorithm requires initial guesses for the transition and emission probability matrices given an observation sequence. We use the outputs of our Markov regime-switching regression (13) as initial guesses and not some random guesses, e.g. drawn from a simulation. The rationale behind this approach is twofold. First, the Baum-Welch is proved to converge to local maxima (e.g. Dempster et al., 1977), which is why good initialization is crucial. Therefore, using the transition and emission probability matrices implied by the Markov regime-switching regression minimizes the risk of ending up in a bad local maximum since these probabilities should also be sufficiently good without the use of the Baum-Welch algorithm. Second, the transition and emission probability matrices implied by the Markov regime-switching regression are already probabilities and hence this poses not the problem of translating identified regime periods into probabilities. Ergo, our transition matrix is given by the estimates of the Markov regime-switching regression. Our guesses for emission probabilities are as follows. From our Markov regime-switching regression, we get filtered probabilities as output, which describe the probabilities predicted by the regime-switching model to be in state 1 or state 2. To transform these probabilities into emission probabilities, we sum all probabilities for each state up and divide the single probabilities by the whole sum of them:

$$e_{s_t x_t} = \frac{\mathbb{P}_{filtered,t_{s_t}}}{\sum_{t=1}^{T} \mathbb{P}_{filtered,t_{s_t}}} \qquad s_t \,\epsilon\, (1,2) \tag{19}$$

Given these initial guesses for the transition and emission probabilities, the Baum-Welch algorithm iteratively improves the transition and emission probabilities in order to maximize the probability of the observation sequence. We use 10,000 iterations until we allow the Baum-Welch algorithm to halt. Since the purpose of the Baum-Welch algorithm in our context is just for improving our HMM parameter estimation, we just sketched the functioning of the algorithm in a non-formal manner, wherefore the mathematical versed reader curious about the algorithm's precise procedure is referred to Baum et al. (1970), Juang and Rabiner (1991), and Bilmes (1998), where the Baum-Welch algorithm is explained more detailed.

Viterbi Algorithm

These improved transition and emission probabilities are then used for the Viterbi algorithm, an efficient dynamic programming algorithm, which was originally proposed by Viterbi (1967) as a decoding method for convolutional codes in noisy digital communication applications. The purpose of the Viterbi algorithm is that it translates the parameters of a HMM into a state sequence – called the Viterbi path – that has the maximum likelihood with respect to the given HMM. Or in mathematical terms, the Viterbi path describes the state sequence that maximizes the posterior probability, i.e. the probability of observing the state sequence s given the observation sequence x:

$$s^* = \underset{s}{argmax}\; \mathbb{P}\,(x, s)\,, \tag{20}$$

where s^* describes the most likely state path over the observation time and x and s are the observation and state sequence, respectively. Hence, the Viterbi algorithm uses the observation sequence as well as the transition and emission probabilities as input and gives one the hidden state path s^* that maximizes equation (20) as output. The key idea behind the Viterbi algorithm is to find the most probable path for every intermediate and for the final state in our model, which means that it calculates at each time t the most likely path until t and uses only this path for time$t + 1$ (i.e. the probabilities for all previous steps/iterations can be discarded), and so forth until the final period T is reached. By recursively maximizing the joint probability for each possible new state, the Viterbi algorithm maximizes the posterior probability of the entire sequence of states (formula (20)), which is optimal in the maximum likelihood sense.

The precise procedure of the Viterbi algorithm is nontrivial and the gentle reader curious about this procedure is therefore referred to Viterbi (1967), Forney (1973), and Lou (1995), where the exact functioning of the Viterbi algorithm is elaborated in more detail.

5 Results

5.1 Comparison of Extended and Original Corwin-Schultz Illiquidity Measure

In the following subsection, we test the performance of our extended version of the high-low spread estimator by Corwin and Schultz (2012) against its original counterpart. By using the bid-ask spread reported in TRACE as liquidity benchmark, we run a horse race of monthly estimates of our extended illiquidity measure against the original one from 11/08 – 09/13.

Corwin and Schultz (2012) originally developed their illiquidity measure for the usually very liquid stock market and therefore derive their measure only for trading days that are preceded by another trading day. Although their measure can be applied par for par to the bond market, the fact that the original Corwin-Schultz measure neglects the possibility of trading days without any trading activity, was our rationale behind extending it.

While the fundamental rationale of the Corwin-Schultz measure – the high–low price ratio reflects both the security's variance and its bid-ask spread – is not altered by our extension, we relax its assumption of two consecutive days of trading. Consequently, we apply the original concept on the closest days with any trading activity by adjusting for the fact that there are trading days without any trading activity in between.

This means that the main input which differentiates our measure from the original one is the number of days between two consecutive trades, which is why we compare our measure with the original one for different timespans of non-trading days between two consecutive trades.

In order to compare the two measures with each other, we use the bid-ask spread reported in TRACE as liquidity benchmark. This bid-ask spread was introduced and reported in TRACE from November 2008 onwards, wherefore our horse race is performed in monthly terms during the time-period from 11/08 – 09/13. Next to a graphical comparison, we use the mean absolute error (MAE) as statistical tool to measure how close our liquidity measures are to the benchmark. The MAE takes the absolute value[24] of the difference of our liquidity measures and the bid-ask spread reported in TRACE and averages them over the entirety of the comparison period. In mathematical terms the MAE is defined as follows:

[24]Taking an absolute value of a number disregards whether the number is negative or positive and, thus avoids the positives and negatives canceling each other out.

$$MAE = \frac{1}{n}\sum_{i=1}^{n} |LIQ_i - BM_i| \,, \tag{21}$$

where LIQ_i is either our extended version of the Corwin-Schultz measure or the original one and BM_i is benchmark liquidity measure, which is the bid-ask spread reported in TRACE.

Figure 3 illustrates the run of the bid-ask spread reported in TRACE, the Extended Corwin-Schultz, and the original Corwin Schultz measure for a zero, one, two, and three no-trading day gap between two consecutive trades[25]. In Panel A, one clearly sees that the Extended Corwin-Schultz measure is identical to the original one for two consecutive trading days without any non-trading day gap, which is an additional proof that we derived our measure algebraically correctly. However, both measures are quite far off from the liquidity benchmark. Moving to a non-trading day gap of one day duration between two consecutive trades (Panel B) changes the picture significantly, since the Extended Corwin-Schultz measure now displays the run of the bid-ask benchmark correctly, while the original Corwin-Schultz measure is relatively far away from the benchmark, which is supported in statistical terms by a MAE of 0.002 of our measure compared to a 0.008 for the original one. This superiority of our measure continues for a two-day gap between two consecutive trades (Panel C), but only slightly since now our measure underestimates the bid-ask spread while the original one overestimates it, which results in a MAE of 0.006 and 0.007, respectively. The pattern that our measure starts to underestimate the bid-ask spread with every additional day in the gap perpetuates when we look at Panel D, where the gap is three days long. Here, the underestimation of our extended model is more severe than the overestimation of the original model (MAE of 0.008 and 0.006 for the extended and original version, respectively). Although we do not display it graphically, our version of the Corwin-Schultz measure, compared to the original one, results in less accurate the longer the no-trading time gaps between two consecutive trading days.

Table 4: MAE Extended and Original Corwin Schultz Illiquidity Measure against Benchmark

This table shows the comparison of the monthly Extended Corwin-Schultz bid-ask spread estimator and the original one against the bid-ask spread of TRACE, describing the liquidity benchmark, from 11/08 – 09/13. Panel A, Panel B, Panel C, and Panel D report the Mean Absolute Error (MAE), which describes the average in absolute values of the difference of the correpsonding liquidity measures and the bid-ask spread reported in TRACE.

Panel A: Zero-Day Gap	
Mean Absolute Error	
Extended	Original
0.0141	0.0141

Panel B: One-Day Gap	
Mean Absolute Error	
Extended	Original
0.0021	0.0079

Panel C: Two-Day Gap	
Mean Absolute Error	
Extended	Original
0.0064	0.0069

Panel D: Three-Day Gap	
Mean Absolute Error	
Extended	Original
0.0083	0.0065

In a nutshell, our Extended Corwin-Schultz illiquidity measure is more accurate than the original one for no-trading day gaps of one and two days between two consecutive trades, while by definition

[25]For reasons of simplification, we limit our analysis to a few universal examples.

being equally accurate for two consecutive trading days. Given the trading characteristics of our entire sample, our illiquidity measure provides in 84.6% of all transactions an equal or more precise estimate of the bid-ask spread than the original Corwin-Schultz measure.

The main challenge of both Corwin-Schultz illiquidity measure versions can be found in the measures' underlying assumption of a geometric Brownian motion of bond prices. Although this assumption is not optimal since geometric Brownian motions are more suitable for the courses of stock prices, it is a reasonable one for bond prices over a relatively short time period. However, it becomes less reasonable for longer time periods due to the pull-to-par effect of bond prices, which describes the phenomenon that bond prices converges to par value for bonds close to maturity. Still, due to the noisiness that is existent in all bid-ask spread estimators, such illiquidity measures are not really applicable for very infrequently traded securities. Therefore, we infer that our extension of the Corwin-Schultz measure seems to be more suitable for the bond market than the original one since it provides a more accurate estimation of the bid-ask spread for the relevant time periods, where the assumption of a bond price following a geometric Brownian motion is still reasonable (up to two days of no trading between two consecutive trades).

Nonetheless, there remains the question how adjusting for overnight price changes or for negative daily values affects the accuracy of our measure, which we, however, leave for further research.

Figure 3: Comparison Extended and Original Corwin Schultz Illiquidity Measure

This figure shows the comparison of the monthly Extended Corwin-Schultz bid-ask spread estimator and the original one against the bid-ask spread of TRACE, describing the liquidity benchmark, from 11/08 – 09/13. Panel A, Panel B, Panel C, and Panel D display the runs of the illiquidity measures and the liquidity benchmark for no-trading day gaps between two consecutive trades of zero, one, two, and three duration, respectively.

Panel A: Zero-Day Gap

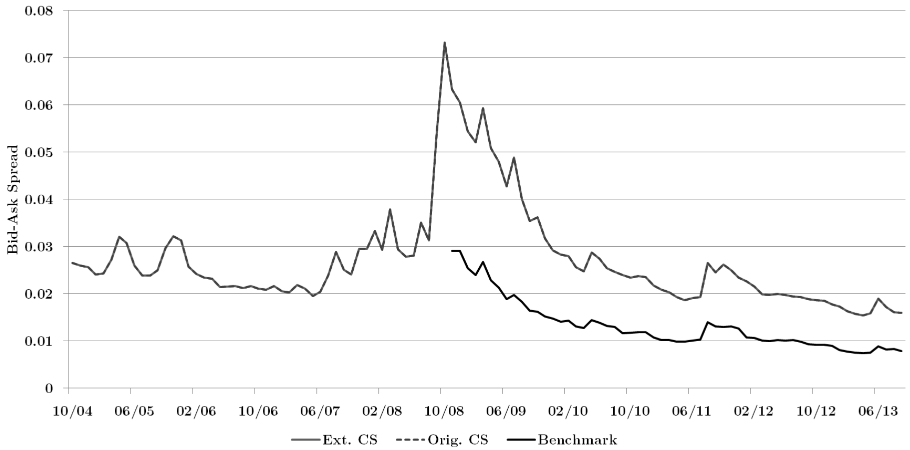

Panel B: One-Day Gap

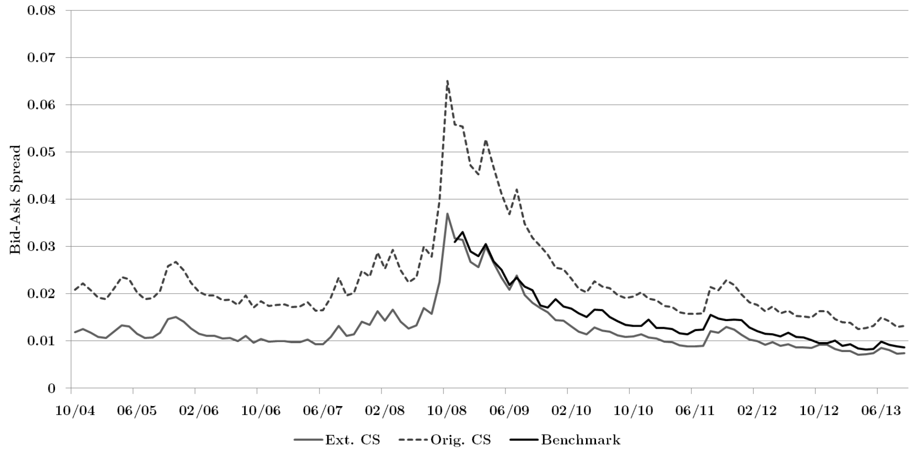

Panel C: Two-Day Gap

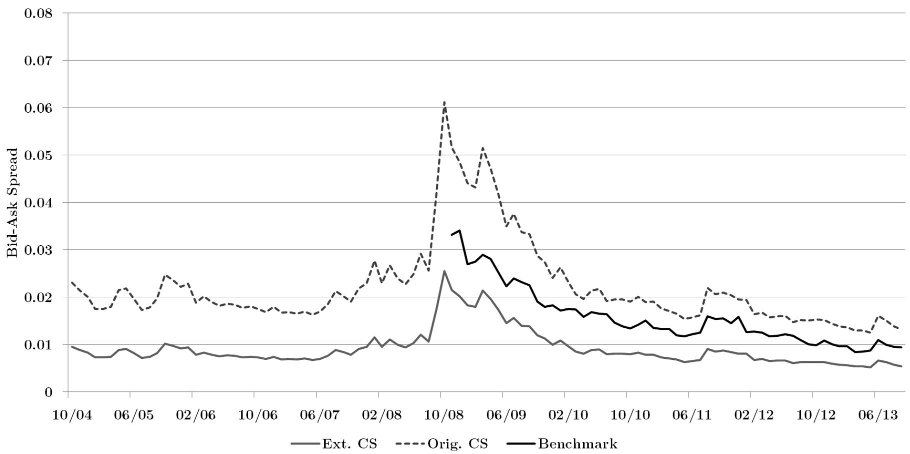

Panel B: Three-Day Gap

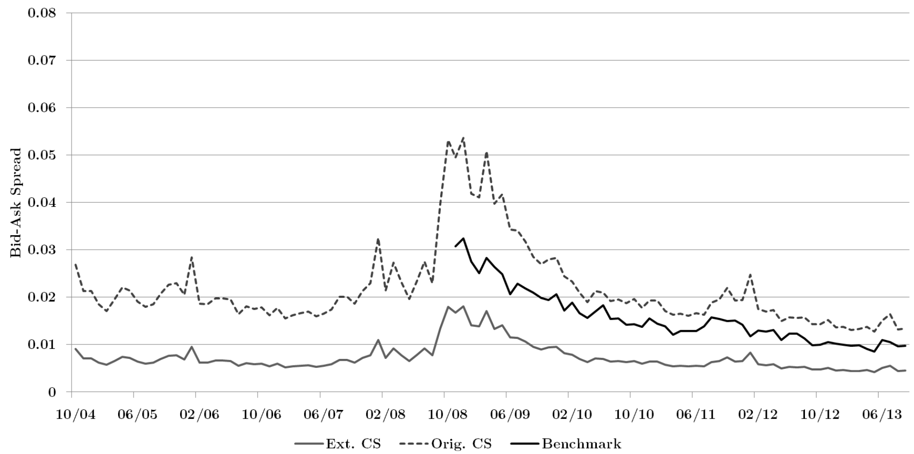

5.2 Unconditional Analysis

In this subsection, we test to what extent liquidity level and liquidity risk are priced in the U.S. corporate bond market during our sample period. Before performing the two-step Fama-MacBeth (1973) procedure to determine factor loadings and corresponding risk premia, we show descriptive statistics of our sample with respect to liquidity.

5.2.1 Liquidity Characteristics

Liquidity Across Rating

Since the crucial point of our analysis is the interplay of credit quality with liquidity, Table 5 displays the average monthly illiquidity measures (standardized through dividing by its mean) across credit ratings for our sample and the sample of Chen et al. (2007), who use information from Datastream on c.4,000 bonds. Chen et al. (2007) find that bonds with lower ratings exhibit significantly lower liquidity, while our sample shows almost no major differences in illiquidity measures across ratings. Panel B and C of Table 5 display the three illiquidity measures employed by Chen et al. (2007): the simple percentage of zero returns, a modified version of Lesmond et al. (1999) model's effective spread estimator that is based on the assumption of informed trading on non-zero-return days and the absence of informed trading on zero-return days, and the relative bid-ask spread. All three illiquidity measures increase significantly with deteriorating credit rating; for short maturities they are 5.93, 7.88, and 24.51 for AAA bonds and 46.31, 933.06, and 77.00 for CCC-D bonds, respectively. In our study, both price impact measures are in fact higher for top rated than for junk bonds with an Amihud measure of 1.66 for AAA bonds and 1.56 for CCC-D bonds and a Pastor-Stambaugh measure of 1.97 for AAA bonds and 1.89 for CCC-D bonds. With respect to our bid-ask proxies, the Roll and Extended Corwin-Schultz measures are about half as high for AAA bonds than for CCC-D bonds, however they do not share the large difference found by Chen et al. (2007). Therefore, our results challenge their findings to a very high extent. As we filter out many illiquid bonds in our data preparation (especially due to the filter of omitting bonds with less than ten trading days per month), we could incur a sample bias if those neglected bonds are primarily low-rated bonds. However, we find that the omitted bonds even show a higher credit rating on average. Consequently, the differences in results may arise from differences in the sample periods (2004 to 2013 in our sample compared to 1995 to 2003 in Chen et al., 2007) or from differences in the specific samples (c.13,850 bonds from TRACE in our sample compared to c.4,000 bonds from Datastream in Chen et al., 2007). Friewald et al. (2012) confirm our methodological approach since they show a very similar illiquidity pattern to our results with respect to investment grade and high yield bonds for both the Amihud and Roll measures. As they also base their analysis on the TRACE database and the time period of 2004 to 2008, it stands to reason whether Datastream omits a large part of lowly rated, but highly liquid bonds and whether high yield bonds have notably increased in liquidity from the 1990s to the 2000s.

Assuming the correctness of our liquidity dispersion across ratings, the conclusion of Acharya et al. (2013) about a flight-to-liquidity within the U.S. corporate bond market is called into question (at

Table 5: Illiquidity Summary Statistics by Rating Category

This table reports illiquidity summary statistics for our sample of bonds and the sample of Chen et al. (2007). Each rating category comprises all of its gradations and its equivalent Standard & Poor's and Fitch rating, e.g., column AA comprises AA+, AA and AA-, as well as Moody's Aa1, Aa2 and Aa3. *#Bonds* is the number of bonds. Panel A comprises our sample which is based on the period from October 2004 until September 2013. We display the average time-series mean and standard deviation of the bond's monthly log returns, whose cross-sectional mean is reported under *Return* and *Volatility* (both in %). *#Trading days* is the bond's total number of trading in the respective period. Volume is the bond's average monthly dollar trading volume in $m. *AM* is the Amihud, *PS* is the Pastor-Stambaugh, *RL* is the Roll and *CS* is the Extended Corwin-Schultz illiquidity measure, all normalized to a mean of zero and a standard deviation of one. Panel B and Panel C comprise the bond sample of Chen et al. (2007) for maturities of 1-7 and 7-15 years, respectively. Their sample period ranges from January 1995 to December 2003. *% Zeros* is the percentage of zero returns for a given year adjusted for missing prices. *LOT* refers to the modified Lesmond et al. (1999) model's liquidity estimate. The *bid–ask* is the proportional spread derived from quarterly quotes from Bloomberg. The bid-ask measure is calculated on a smaller sub-sample of bonds.

	Investment Grade				High Yield		
Rating	AAA	AA	A	BBB	BB	B	CCC-D
Panel A: Our Sample, 2004 - 2013							
#Bonds	495	2,368	5,578	5,391	2,914	2,262	1,996
Return (in %)	0.16	0.24	0.33	0.35	0.47	0.44	1.62
Volatility (in %)	2.44	1.96	2.32	2.61	3.11	3.32	6.51
#Trading days	146	153	147	147	149	146	148
Volume ($m)	46.9	48.0	42.7	45.4	26.1	26.5	27.3
AM	1.66	1.13	1.02	1.17	1.44	1.04	1.56
PS	1.97	1.23	1.02	0.94	1.53	1.00	1.89
RL	1.05	0.86	0.90	1.09	1.28	1.14	2.02
CS	1.01	0.93	0.92	1.03	1.16	1.06	1.91
Panel B: Chen et al. (2007) with Short Maturity (1-7 years), 1995 - 2003							
#Bonds	87	336	1,162	1,234	333	167	119
Zeros (%)	5.93	4.10	3.88	8.43	40.63	44.71	46.31
LOT (bp)	7.88	9.63	10.51	34.99	201.45	458.86	933.06
Bid-ask (bp)	24.51	26.02	25.82	91.01	54.26	58.76	77.00
Panel C: Chen et al. (2007) with Medium Maturity (7-15 years), 1995 - 2003							
#Bonds	49	120	539	730	152	78	44
Zeros (%)	9.79	12.59	10.61	11.94	36.99	38.71	34.96
LOT (bp)	24.28	47.26	57.74	70.29	259.34	342.5	941.84
Bid-ask (bp)	49.52	36.57	38.20	44.02	54.65	60.44	180.35

least for the recent years). They argue that higher liquidity betas of high yield compared to investment grade bonds depict a flight-to-liquidity due to the fact that high yield bonds are less liquid than their investment grade counterparts. Relaxing the assumption of liquidity dispersion across ratings, the difference in liquidity betas would then rather indicate a flight-to-quality (conditional on liquidity shocks in stress periods).

Liquidity across both Rating and Liquidity

It is worthwhile mentioning that our sample's characteristic of similar liquidity levels across rating categories forms a major feature of our ability to identify the specific effect of credit rating on a bond's return. Otherwise, in case of positively correlated credit risk and illiquidity, it would be difficult to empirically isolate the effect of credit rating. In order to also visualize the specific effect of liquidity, we divide our total sample not only in credit categories, but also among the dimension of liquidity. Therefore, we sort all bonds in five portfolios according to those two dimensions, named quality (credit rating) and liquidity, so that we obtain a total of 25 portfolios which we will use throughout most of our analyses. The bonds are ascribed towards quality portfolios according to their credit rating and assigned to liquidity portfolios according to the average liquidity of the Amihud and Roll measures. We sort so that every portfolio contains the same number of total months across its bonds. Consequently, the number of bonds in the portfolios differs, but not the number of time-series observations. Every portfolio is given a number from 5 to 1 for both dimensions with a high number representing high credit rating / high liquidity (=low illiquidity), and vice versa. Consequently, the best rated and most liquid bonds are in the portfolio with a quality and a liquidity number of both 5. A bond with the highest rating, but lowest liquidity is ascribed to the portfolio with a quality number of 5 and a liquidity number of 1. The sorting is based on the average manifestation of both characteristics for each bond over the whole sample period.

Figure 4 and Table 10 (see Appendix) displays the average illiquidity according to the four measures[26]. As described above, both price impact measures Amihud and Pastor-Stambaugh do not show any consistent in- or decrease across quality, while both bid-ask proxies Roll and Extended Corwin-Schultz slightly increase with lower credit quality. Furthermore, the price impact measures show a higher cross-sectional dispersion than the bid-ask spread estimators, i.e. illiquid bonds have very high price impact while only having a moderately high bid-ask spread. Liquid bonds, on the other hand, show almost no price impact, but still a considerable wide bid-ask spread. Hence, the liquidity's aspect of price impact is the characteristic that varies the most between liquid and illiquid bonds; the bid-ask spread also widens when one moves from liquid to illiquid bonds, but at a considerably lower pace. Furthermore, we shift our attention towards the Extended Corwin-Schultz measure because of its novelty of being used in a modified version for bonds. We conclude that its pattern is very similar to that of the Roll measure, except for the fact that the dispersion of illiquidity is less pronounced across ratings.

Returns across both Liquidity and Rating

The analysis of returns across both dimensions gives a first impression on how credit rating and liquidity are priced, however the Fama-MacBeth regression later will be more precise as it adjusts for additional risk sources that could vary among the portfolios. Panel A of Figure 5 displays the average realized return over the sample period which is between 2% and 6% p.a. In general, there exists a pattern that returns increase in lower quality and lower liquidity, however returns fall dramatically for

[26]Throughout the following, we will strongly rely on three-dimensional graphs to display our findings. With justification, the academic literature of asset pricing shies away from graphical representation which allows for biasing numerical results through specifically chosen scales, shadings or perspectives, while hard numbers do not leave any room for distortion. However, we believe that a graphical display of our portfolio characteristics is adequate as it increases readability by a large extent. All results can be found as tables in the appendix.

Figure 4: Portfolio-Sorted Illiquidity Measures

This figure shows the average bond portfolio illiquidity for 25 portfolios sorted by liquidity and quality during our sample period from 11/04 – 09/13. Illiquid and liquid are defined as the bond portfolios with the highest and lowest illiquidity, respectively and junk and quality are defined as the bond portfolios with the lowest and highest rating, respectively. Panel A, Panel B, Panel C, and Panel D display the level of illiquidity based on the Amihud measure, the Pastor-Stambaugh measure, the Roll measure, and the Extended Corwin-Schultz measure, respectively.

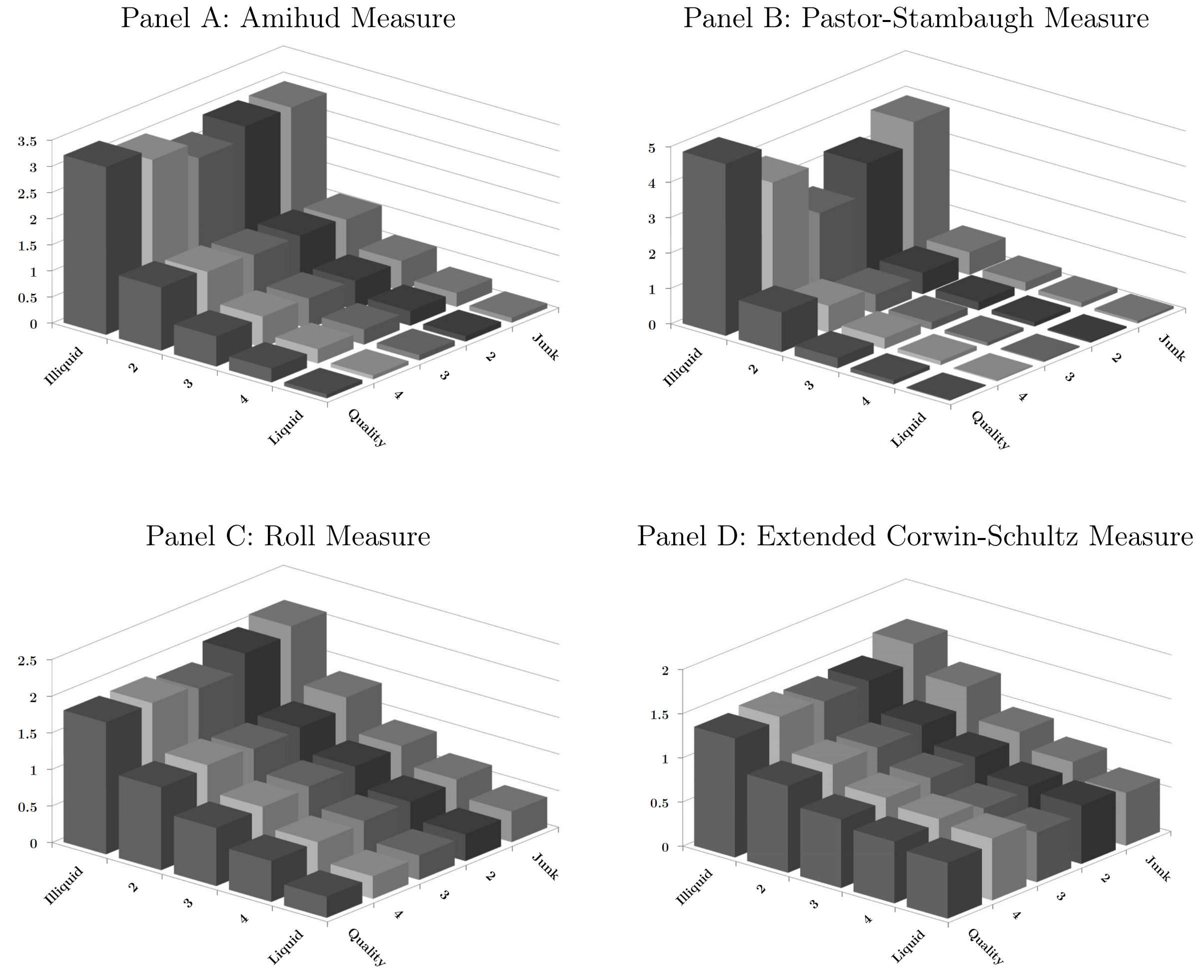

most of the lowest rated bonds. Besides, Panel B shows the expected excess returns over the risk-free rate derived from the quoted yield (before adjusting for the default component). Expected returns are far higher than their realized counterparts and lie between 3% and 10% p.a. and show a consistent pattern across quality and liquidity. The average expected return amounts to 5.5% p.a. while the actual returns only average 3.2%. We can conclude that actual returns stay behind expectations over our sample period which underlines the importance of using expected returns for deriving risk premia. One explanation for the difference may be that the negative effect of the financial crisis in 2008/2009 represents a fat tail event that exceeds what investors had expected over the 9 year period.

All five junk portfolios in Panel B show a significant jump in expected returns compared to the second lowest rating portfolios. This return gap accounts for the fact that the junk portfolios also have a significantly lower credit rating than the other portfolios as shown in Panel D (please note that the axes are reversed in Panel D in order to facilitate readability) and that the lower ratings exhibit

Figure 5: Portfolio-Sorted Returns and Rating

This figure shows the average monthly bond portfolio returns in excess of the one-month T-bill and the rating of 25 portfolios sorted by liquidity and quality during our sample period from 11/04 – 09/13. Illiquid and liquid are defined as the bond portfolios with the highest and lowest illiquidity, respectively and junk and quality are defined as the bond portfolios with the lowest and highest rating, respectively. Panel A, Panel B, Panel C, and Panel D display the average actual returns, the average yield-to-maturity, the average expected returns, and the average rating of the portfolios, respectively.

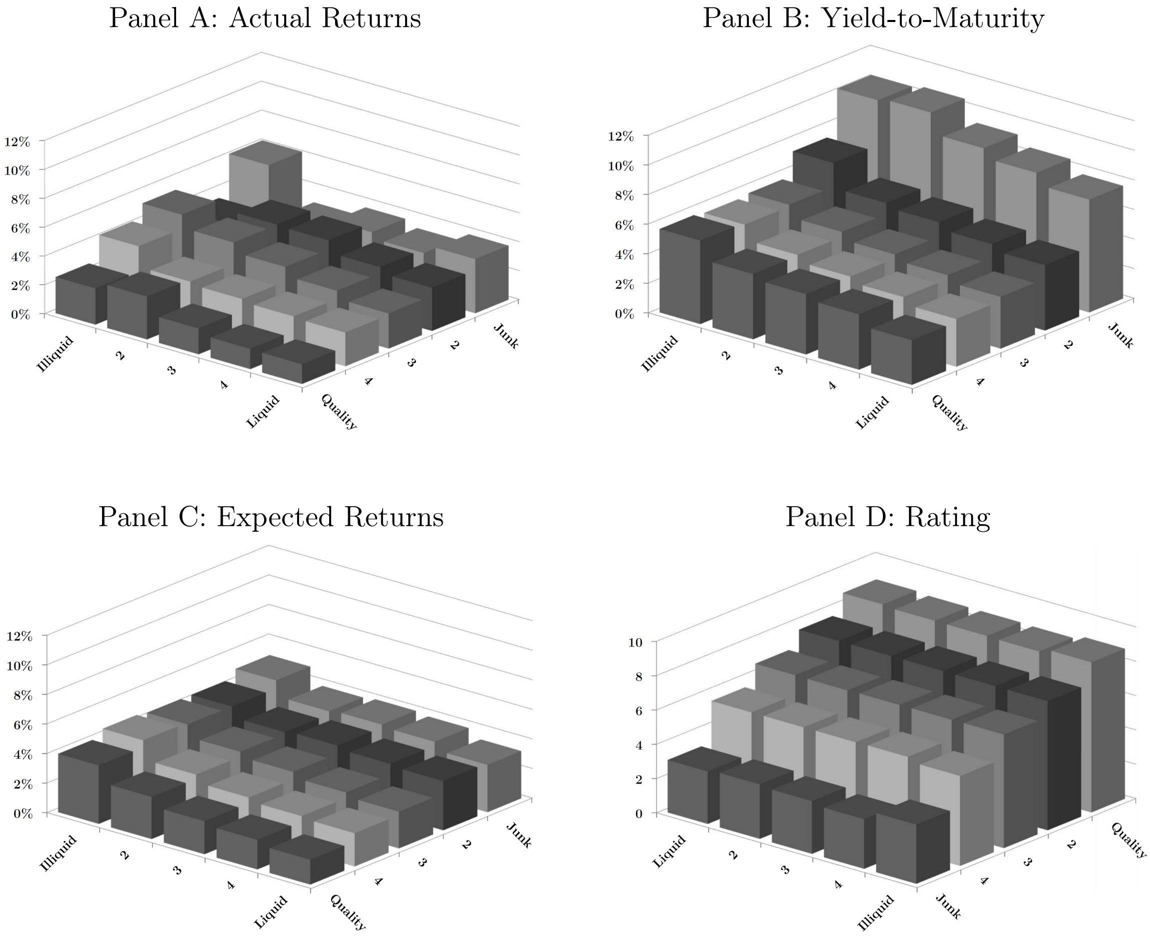

a disproportionally high default probabilities (as can be seen in Table 11 in the Appendix). Panel C accounts for these differences and displays expected excess returns after adjusting for the expected default loss. The remaining expected returns are between 1.5% and 5% p.a. and still average to 3.3%. They consequently represent what is commonly known as the "credit spread puzzle", i.e. the fact that bond expected and realized returns are consistently higher than predicted according to the time value of money (risk-free rate) and credit risk (expected default loss). In the following section, we set out to analyze whether the five Fama-French (1993) factors and especially liquidity risk and liquidity level can explain these expected excess returns.

Market-Wide Liquidity

We first check the time-series correlation of our factors which include the market-wide illiquidity measures derived as described in section 4.1. The factors comprise the term (TERM), default (DEF), market (MKT), size (SMB) and book-to-market (HML), as well as our market-wide liquidity measures (the market-wide versions of the liquidity measures are described with the subscript of “mkt“). The market-wide liquidity measures in Table 6 are the negative of the volume-weighted average of all company-specific illiquidity proxies for every month across our sample of bonds. The illiquidity measures are multiplied with -1 so that a positive (negative) shock to the measure represents an increase (decrease) in market-wide liquidity.

Table 6: Correlations among Factors

The table displays the time-series correlation of five common risk factors and four market-wide liquidity measures. *TERM* is the difference between the monthly 30-year T-bond return and the one month T-bill returns, *DEF* is the difference between the monthly return on a value-weighted portfolio of our entire bond sample and the average return on government bonds (i.e. 30-year T-bond minus one month T-bill), the Fama-French three factors *MKT*, *SMB*, and *HML* are obtained from Kenneth French's data library, and therefore are defined as described in Fama and French (1993). AM_{mkt}, PS_{mkt},RL_{mkt}, and$CS_{mkt}$are the volume-weighted market-wide liquidity levels according to the Amihud, Pastor-Stambaugh, Roll and the Extended Corwin-Schultz illiquidity measures. The illiquidiy measures are multiplied with -1 so that a positive (negative) shock to the measure represents an increase (decrease) in market-wide liquidity. The market-wide liquidity measures are divided by their mean in order to create comparability.

	TERM	*DEF*	*MKT*	*SMB*	*HML*	AM_{mkt}	PS_{mkt}	RL_{mkt}	CS_{mkt}
TERM	1.00								
DEF	0.10	1.00							
MKT	0.02	0.60	1.00						
SMB	0.13	0.20	0.47	1.00					
HML	-0.02	0.07	0.36	0.21	1.00				
AM_{mkt}	0.31	0.37	0.35	0.07	0.07	1.00			
PS_{mkt}	0.00	0.34	0.41	0.06	0.11	0.82	1.00		
RL_{mkt}	0.03	0.24	0.42	0.10	0.17	0.85	0.86	1.00	
CS_{mkt}	-0.08	0.14	0.24	0.01	0.14	0.55	0.50	0.61	1.00

Table 6 shows that the liquidity measures have a moderately correlation to the five common factors (between -0.08 and 0.42) so that the inclusion of liquidity as a potential additional factor is justified. However, the liquidity factors themselves show a high correlation between 0.55 and 0.86. Hence, price impact and the bid-ask spread strongly co-move and the two illiquidity measures within both liquidity dimensions appear to be accurate measures. Due to the high correlation of the illiquidity measures and thus the possible problem of multicollinearity, we only include one dimension at a time for our pricing analysis. Figure 6 plots the four market-wide illiquidity measures AM_{mkt}, PS_{mkt}, RL_{mkt}, and CS_{mkt} over time. AM_{mkt}, PS_{mkt}, and RL_{mkt} show a very high co-movement, while the CS_{mkt} stands out due to its very high peaks.

5.2.2 Pricing of Liquidity

Time-Series Regression of Factor Loadings

We now run the first step of the of the Fama-MacBeth (1973) procedure which estimates the time-series sensitivities of every bond's actual returns to the identified factors of TERM, DEF, MKT, SMB, HML and one market-wide liquidity risk factor (generically named LIQ_{mkt} in the following). We do not form portfolios, but run the regression on every bond individually. The number of bonds with

Figure 6: Illiquidity Measures over Time

This figure shows the run of our four illiquidity measures over our sample period from from 11/04 – 09/13. Panel A shows the marketwide illiquidity for the Amihud, Pastor-Stambaugh, Roll, and Extended Corwin-Schultz illiquidity measure.

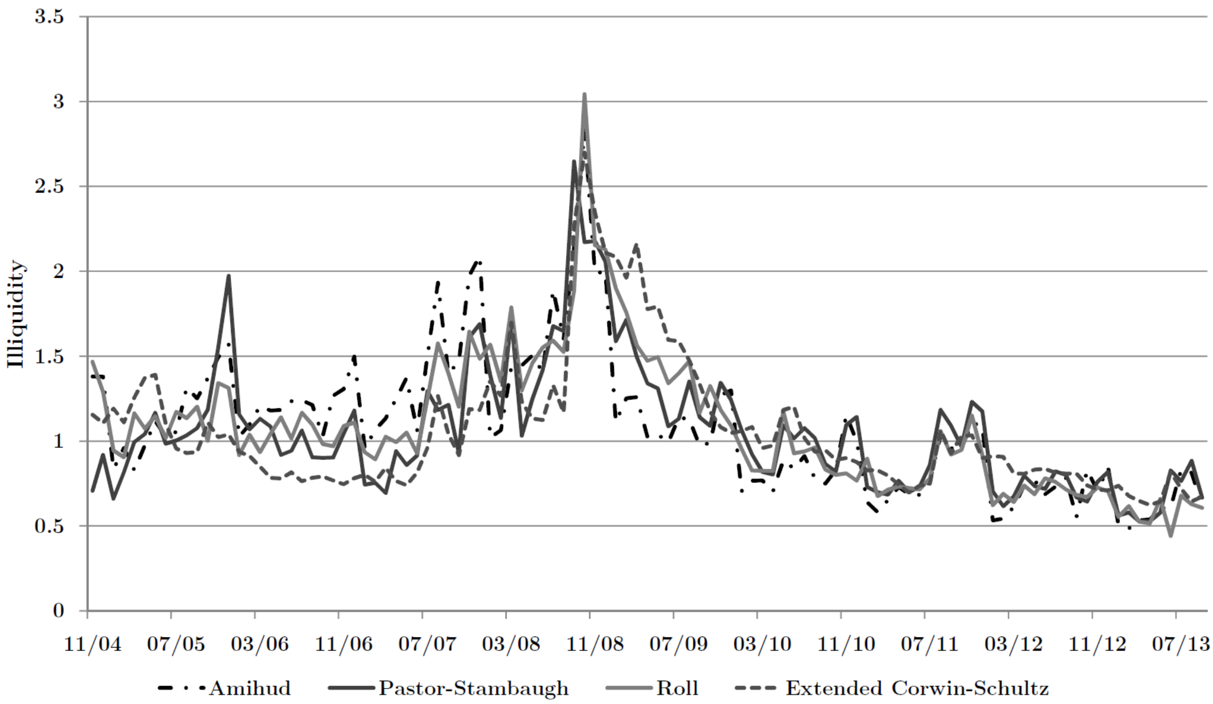

enough time-series observations in order to determine betas accounts for 7,716.

Panel A in Table 7 shows the average beta loading across all the bonds in our sample. It comprises five regressions: BASIC does only incorporate the Fama-French five factors, while the other four include their respective market-wide liquidity factor. TERM and DEF have beta loadings of 0.54 to 0.67 and 0.76 to 0.78, respectively, while MKT, SMB, and HML are at around 0.01 to 0.02. These results hold regardless of employed illiquidity measure. DEF and MKT are statistically significant while TERM, SMB, and HML are not. Due to our standardization of the liquidity measures to a mean of one, their coefficients are small. However, they show the right sign and the liquidity factors according to RL and ECS are significant at the 10% level while those of the PS measure are significant at the 5% level. The AM measure is slightly below the 10% significance level. In general we can conclude that bonds co-move statistically significantly with market-wide liquidity. Whether investors also require a compensation for that co-movement needs to be examined in the second step of the Fama-MacBeth procedure.

Our coefficients are in line with Lin et al. (2011), however they show a low level of statistical significance which can probably be accredited to the noisy nature of our realized returns. Since the insignificant TERM factor shows almost the same coefficient as its significant counterpart in the work of Lin et al. (2011) and the SMB and HML factor have been proven to drive bond returns in several other studies (such as Fama and French (1993) and Elton et al. (2001)) we include them also in the cross-section regression to derive risk premia. Since the identified risk factors are scaled differently,

Table 7: Six-Factor Regression for Monthly Excess Returns

This table displays the results of a linear time series regression of a six-factor model on monthly bond return series from 11/04 – 09/13 (first step of Fama-MacBeth procedure). The dependent variable is the return of bond i in excess of the one-month T-bill return and the explanatory variables are *TERM,* defined as the difference between the monthly 30-year T-bond return and the one month T-bill returns, *DEF* defined as the difference between the monthly return on a value-weighted portfolio of our entire bond sample and the average return on government bonds (i.e. 30-year T-bond minus one month T-bill), MKT, defined as excess return on the market minus the one-month T-bill rate, *SMB*, defined as the average return on small size portfolios minus the average return on big size portfolios and *HML*, defined as the average return on value portfolios minus the average return on growth portfolios, and LIQ_{mkt} describes the market-wide liquidity for either the Amihud (AM), Pastor-Stambaugh (*PS*), Roll (*RL*), and Extended Corwin-Schultz (*CS*) illiquidity measure. Panel A shows the average regression coefficents, where the β regression coefficients measures the bond i's risk loading on the particular factors and the α measures the part of the return that cannot be explained by the loadings on the different factors. The complete regression in compact matrix form is the following:

$$R = F\beta_n + \varepsilon_n,$$

where R is the vector of excess returns, F is the factor matrix with the factors defined as above, β_n is the vector of factor loadings and all first elements are the intercept α, and ε_n is the vector of error terms. One and two asterisks indicate that the regression coefficient is statistically significant at the 10% and 5% significance level, respectively. The t-statistic for each regression coefficient is displayed in parentheses below it and is calculated from Newey-West (1987) standard errors, which are corrected for heteroskedasticity and serial correlation. Panel B presents the ratio of the standard deviation of returns to standard deviation of factors.

Panel A: Average Coefficients from Time-Series Regression (Factor Loadings)

	α	$TERM$	DEF	MKT	SMB	HML	LIQ_{mkt}	Adj. R^2
BASIC	0.0001	0.64	0.78**	0.02**	0.02	0.01		19.3%
	(1.38)	(1.32)	(3.04)	(1.97)	(1.46)	(1.57)		
AM	0.0003	0.54	0.76**	0.02*	0.01	0.02	0.0010	33.7%
	(1.42)	(1.33)	(2.96)	(1.93)	(1.48)	(1.56)	(1.60)	
PS	0.0000	0.70	0.78**	0.02*	0.01	0.01	0.0005**	34.0%
	(1.39)	(1.34)	(3.07)	(1.94)	(1.47)	(1.56)	(2.11)	
RL	0.0002	0.59	0.75**	0.02**	0.01	0.01	0.0012*	26.8%
	(1.40)	(1.35)	(2.97)	(1.99)	(1.44)	(1.57)	(1.69)	
CS	0.0002	0.62	0.76**	0.02*	0.02	0.02	0.0006*	33.8%
	(1.39)	(1.34)	(2.75)	(1.92)	(1.48)	(1.56)	(1.65)	

Panel B: Ratio of Standard Deviation of Returns to Standard Deviation of Factors Above

	α	$TERM$	DEF	MKT	SMB	HML	LIQ_{mkt}
BASIC	0.84%	6.22%	93.61%	7.57%	2.64%	2.65%	
AM	2.68%	5.25%	91.40%	7.64%	2.16%	2.99%	2.05%
PS	-0.31%	6.79%	93.49%	7.34%	1.66%	0.97%	1.00%
RL	1.57%	5.73%	90.65%	8.16%	1.41%	2.65%	2.46%
ECS	1.51%	6.03%	91.58%	7.17%	3.69%	3.45%	1.21%

Panel B of Table 7 displays the coefficients as the ratio of standard deviation of the returns to the standard deviation of the factor. This represents how many standard deviations the return changes in a change of one standard deviation of the factor. Thus, it gives the economic significance of the underlying factors, i.e. it makes the sensitivities to the factors comparable. DEF leads to the highest fluctuation in returns with a ratio of c.92% (so one standard deviation shock to DEF increases returns by 0.92 of the return's standard deviation), followed by MKT and TERM with c.5% - 8%. Market-wide liquidity has a minor effect with a ratio of 1.00% to 2.46%. Whether every co-movement does actually result in a risk premium required by investors will be examined in the next section.

Cross-Sectional Regression to Determine Risk Premia

To analyze the compensation for factor loadings and the liquidity characteristic we perform the cross-section regression of factor betas from Table 7 together with the liquidity level on expected returns, which is generally known as the second step of the Fama-MacBeth procedure. Thereby, we first normalize each right-hand side variable by its monthly standard deviation and annualize it so that its coefficient is readily interpretable as the annualized premium per unit of standard deviation. The second step of the Fama-MacBeth regression performs the cross-section regression first on every month individually and then averages the resulting coefficients over time.

Table 8 displays the average of the monthly coefficients of the cross-section regressions. $BASIC_r$ represents the regression on realized, while $BASIC_{E(r)}$ is based on the expected returns after adjusting for the default probability. The negative risk premia for TERM (-0.87%) and MKT (-0.42%) of $BASIC_r$ once again underline the importance of using expected returns in the area of corporate bonds, especially in cases of small sample periods and fat tail events. The benchmark regression, $BASIC_{E(r)}$, shows an only slightly negative TERM premium and relatively small SMB and HML effects, while DEF and MKT with risk premia of 0.90% and 0.76%, respectively, are the most important risk factors. This is also supported by the t-statistics, which show especially for DEF and MKT highly significant results, while not showing a consistent statistical significance pattern for the other risk factors. Before comparing the illiquidity measures with each other, we first show the importance of including the liquidity level when analyzing (market-wide) liquidity risk. Table 8 shows that the regressions omitting liquidity level (AM_{mkt}, PS_{mkt}, RL_{mkt}, and CS_{mkt}) results market-wide liquidity risk premium, LIQ_{mkt}, of 0.41% to 0.62% for the four illiquidity measures. The liquidity level-only regressions (AM_{lvl}, PS_{lvl}, RL_{lvl} and CS_{lvl}) display a liquidity level premium of 0.38% to 0.92%. In the regressions including risk sources, liquidity level and liquidity risk, the liquidity risk premium falls by up to 50% while the liquidity level premium falls by up to 25%. We can conclude that each liquidity dimension captures part of the other dimension, while liquidity risk appears to be affected more strongly.

In terms of statistical significance, all risk premia for all three specifications, namely only liquidity level, only liquidity risk, and both liquidity level and liquidity risk, are highly significant and hence we can infer that there is a risk premium investors require for liquidity/illiquidity. Consequently, studies such as Downing et al. (2005), Chacko (2006), Lin et al. (2011), and De Jong and Driessen (2012) which do not include liquidity level in their analysis of liquidity risk may overstate the effect of market-wide bond liquidity variations on returns, but also the large literature focusing on liquidity level run the risk of slightly overrating the liquidity level effects. The Extended Corwin-Schultz measure appears to be able to grasp liquidity effects since it shows high risk premia for the liquidity effects and the highest adjusted R^2 of 30.91%. Our findings show about the same size of liquidity effects as other studies. De Jong and Driessen (2012) and Lin et al. (2011) estimate a liquidity risk premium (without accounting for liquidity level) of c.0.6% and 1% p.a., respectively, while our average over the four liquidity measures is 0.42% (based on the models without liquidity level). Bongaerts et al. (2012) who include both liquidity dimensions estimate a liquidity risk premium of c.0.3% and a liquidity level premium of c.1% per standard deviation comparable to 0.4% and 0.5%, respectively, regarding the average of our measures. In general, we can conclude that liquidity has an economically significant effect on corporate bond returns, with liquidity level mattering slightly more than liquidity risk and both together even exceeding the risk premia associated with the default or market risk.

5.3 Effects of Liquidity Stress Periods

To further deepen our analysis of the importance of liquidity in distressed times, we want to investigate whether flight-to-liquidity phenomena, i.e. time periods with a sudden severe rise in the importance of liquidity leading investors to shy away from illiquid and re-allocate towards liquid bonds, exist in the U.S. corporate bond market during our sample period. After identifying specific months in which liquidity increases in importance through the usage of Markov regime-switching models, we form 25 (5x5) bond portfolios sorted by liquidity and quality and examine the excess returns of each portfolio during these stress months in order to scrutinize them for flights. The sorting has the purpose of partially disentangling the potential phenomena of flight-to-liquidity from flight-to-quality (the increasing demand for holding low-risk assets), which might coincide to a large extent. This allows for studying both the effect of credit risk and of liquidity on realized returns in distressed periods.

There are numerous approaches to define the flight effects[27]. In its fundamental concept, a flight represents a capital flow from one to another set of assets. While many studies utilize indirect measures such as correlations and volatility to define flight periods, we opine that solely focusing on the returns of both asset sets is the best proxy for order flow. Furthermore, returns additionally capture price effects of a changing willingness to hold assets (that can be accompanied without unusual capital flows) representing an "implicit" flight which is eventually also of interest for investors. It is important to mention that we do not adjust the returns for other risk factors, hence an identified return pattern across portfolios with different levels of liquidity might be caused by another hidden risk source. The reason is that adjusting returns in distressed times would demand for a conditional asset pricing model whose formulation exhibits considerable difficulties in our case. Probably most severe is the question about the investors' information set. If, for example, a financial crisis (triggering large changes in risk factor characteristics) starts to unfold, it is unclear whether investors' expectations are based on long-term historical, current or anticipated factor exposures and risk premia. Therefore, we assume in the following that return differences along the dimension of liquidity / quality are not caused by any other risk characteristic. Admittedly, the assumption is quite strong; however a large return differential would serve as a strong indication for the impact of the analyzed risk dimensions.

[27] There is no consensus in the literature on how to define flight-to-quality and flight-to-liquidity episodes. In Bekaert et al. (2009) a flight-to-quality is defined as the joint occurrence of higher economic uncertainty with lower equity prices and low real rates. Baur and Lucey (2009), on the other hand, define a flight-to-quality period with a significant decrease in the correlation in a (stock market) crisis period compared to a benchmark period resulting in a negative correlation level. Baele et al. (2013) define flight-to-safety periods as periods with market stress (high equity and perhaps bond return volatility), a simultaneous high bond and low equity return, low (negative) correlation between bond and equity returns, while Goyenko and Sarkissian (2014) define a flight-to-liquidity using illiquidity in short-term U.S. Treasuries.

Table 8: Average Coefficients from Cross-Section Regression (Risk Premia)

This table reports results of cross-sectional regression tests of individual bonds using the Fama-MacBeth procedure over our sample period from 11/04 – 09/13 (second step Fama-MacBeth procedure). The dependent variable is a bond's monthly expected return in excess of the one-month T-bill return. The explanatory variables are the regression coefficents for *TERM, DEF, MKT, SMB, HML*, LIQ_{mkt}, and LIQ_{lvl} from the time-series regression (10), where each one is normalized by its cross-sectional standard deviation every month. The regressions are run either without any liquidity component or only liquidity risk, or only liquidity level or with both liquidity level and liquidity risk. The complete regression in compact matix form is the following:

$$R_t = \hat{\beta}\lambda_t,$$

Where R_t is an vector of expected asset returns, $\hat{\beta}$ is a vector of factor loadings where all elements in the first column are 1, and λ_t is a vector of factor premia where all elements in the first row are the intercept. One asterisks indicates that the regression coefficient is statistically significant at the 5% significance level. The t-statistic for each regression coefficient is displayed in parentheses below it and is calculated from Newey-West (1987) standard errors, which are corrected for heteroskedasticity and serial correlation.

	α	$TERM$	DEF	MKT	SMB	HML	LIQ_{mkt}	LIQ_{lvl}	Adj. R^2
$BASIC_r$	2.31%*	-0.87%*	0.15%*	-0.42%*	0.56%*	-0.16%*			25.2%
	(5.34)	(2.95)	(8.74)	(8.12)	(5.26)	(3.70)			
$BASIC_{E(r)}$	2.01%*	-0.10%	0.90%*	0.76%*	0.20%*	0.22%*			19.3%
	(19.81)	(1.91)	(10.85)	(5.86)	(2.18)	(2.35)			
AM_{mkt}	2.01%*	-0.13%	0.90%*	0.74%*	0.20%*	0.23%*	0.27%*		20.1%
	(19.96)	(1.73)	(10.70)	(5.61)	(2.13)	(2.42)	(2.54)		
AM_{lvl}	1.92%*	-0.13%	0.83%*	0.78%*	0.20%*	0.18%*		0.38%*	21.0%
	(19.15)	(1.71)	(10.08)	(6.03)	(2.14)	(2.06)		(6.75)	
AM	1.87%*	-0.12%	0.86%*	0.74%*	0.19%*	0.21%*	0.23%*	0.35%*	22.05%
	(19.05)	(1.69)	(10.40)	(5.67)	(2.10)	(2.27)	(2.22)	(6.14)	
PS_{mkt}	2.03%*	-0.12%	0.92%*	0.66%*	0.18%	0.18%	0.62%*		24.01%
	(14.32)	(1.57)	(7.88)	(3.97)	(1.86)	(1.77)	(4.44)		
PS_{lvl}	2.02%*	-0.18%	0.80%*	0.64%*	0.12%	0.16%		0.40%*	21.35%
	(13.93)	(1.54)	(7.10)	(3.73)	(1.68)	(1.62)		(5.04)	
PS	1.96%*	-0.12%	0.89%*	0.66%*	0.16%	0.16%	0.56%*	0.30%*	25.42%
	(14.13)	(1.57)	(7.74)	(3.94)	(1.83)	(1.67)	(3.95)	(3.91)	
RL_{mkt}	1.99%*	-0.09%*	0.88%*	0.65%*	0.25%*	0.22%*	0.28%*		20.62%
	(18.90)	(2.26)	(10.10)	(4.96)	(2.14)	(1.96)	(3.11)		
RL_{lvl}	1.48%*	-0.13%*	0.67%*	0.59%*	0.21%	0.16%		0.67%*	21.35%
	(14.97)	(2.00)	(7.78)	(4.71)	(1.78)	(1.50)		(9.88)	
RL	1.45%*	-0.08%	0.69%*	0.59%*	0.23%*	0.17%	0.20%*	0.64%*	26.77%
	(15.04)	(1.93)	(8.14)	(4.73)	(1.99)	(1.62)	(2.47)	(9.18)	
CS_{mkt}	2.02%*	-0.14%*	0.94%*	0.69%*	0.22%*	0.26%*	0.61%*		21.65%
	(20.06)	(2.11)	(10.82)	(5.52)	(2.34)	(2.63)	(5.21)		
CS_{lvl}	1.25%*	-0.18%	0.52%*	0.43%*	0.13%*	0.13%		0.92%*	28.34%
	(10.43)	(1.88)	(6.79)	(3.83)	(2.07)	(1.63)		(9.09)	
CS	1.17%*	-0.13%*	0.67%*	0.51%*	0.17%*	0.17%	0.46%*	0.84%*	30.91%
	(10.23)	(1.98)	(7.96)	(4.41)	(2.34)	(1.83)	(4.16)	(8.30)	

5.3.1 NBER Recession as Benchmark

Before analyzing the effects of liquidity stress periods, we briefly examine the price reactions of our 25 portfolios during the period of the macroeconomic crisis in order to compare the extent in which liquidity stress periods differ later on. We use the crisis period from January 2008 to June 2009 as defined by the National Bureau of Economic Research (NBER) recession indicator, which is a time series of dummy variables that represents periods of expansion and recession based on the data of the National Bureau of Economic Research (NBER) in the U.S. Figure 7 and Table 12 (see Appendix) show the monthly excess returns for the 25 portfolios.

Figure 7: Portfolio-Sorted Excess Returns during NBER Recession

This figure shows the average monthly bond portfolio returns in excess of the one-month T-bill return of 25 portfolios sorted by liquidity and quality during recession periods identified by the NBER recession indicator for our sample period from 11/04 – 09/13. Illiquid and liquid are defined as the bond portfolios with the highest and lowest illiquidity, respectively, and junk and quality are defined as the the bond portfolios with the lowest and highest rating, respectively.

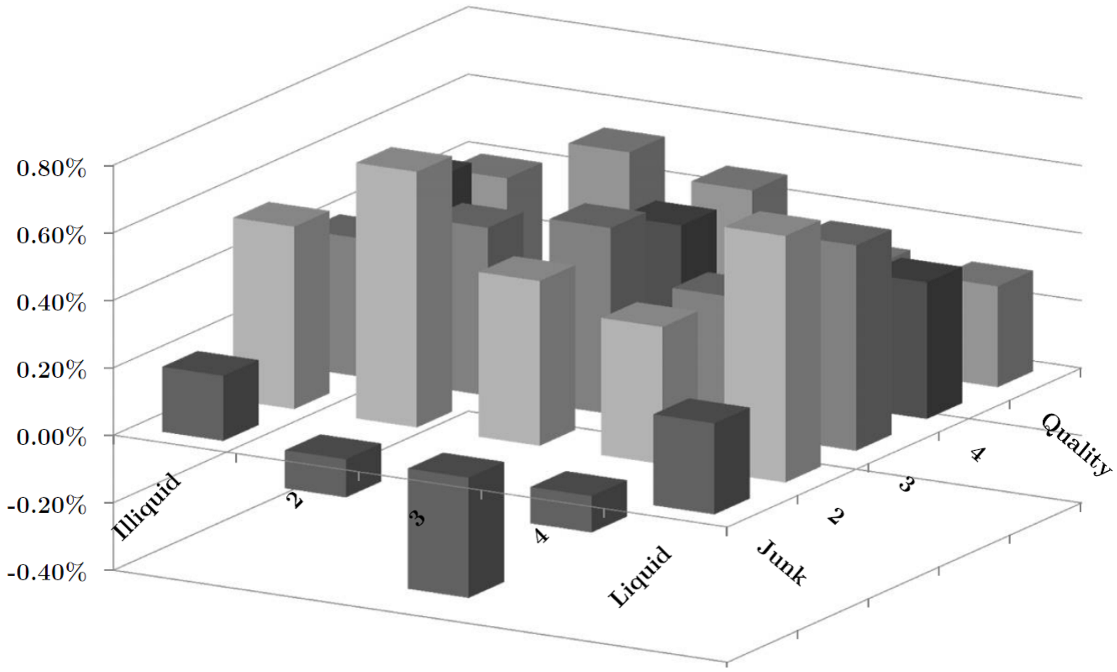

While the lowest quality quintile suffers slightly with an average of -0.02% p.m., all other portfolios show solid positive returns with 0.47% p.m. We can assume that investors shy away from junk bonds, but consider all bonds with non-junk ratings more valuable. Interestingly, there is no gradual price pattern in the dimension of quality, i.e. the four highest quality quintiles show all about the same level of returns. Consequently, the effect appears to be rather a "flight-from-junk" than a "flight-to-quality" in the way that the lowest quintile decreases in value, while there is no apparent return difference among the other four quintiles[28]. This is supported in statistical terms since only for two of the five quality-minus-junk spreads[29] the spread is statistically different from zero. One can hypothesize that

[28]We define a "flight-from" event as an asymmetric effect in the sense that one subset of assets underperforms to all other subsets without the latter showing return differences among each other. We refer to a "flight-to" situation when the returns of the subsets change gradually in the observed characteristic.

[29]Quality-minus-junk is defined as the difference in the return of the highest quality portfolio (regardless of liquidity)

all non-junk bonds with an average monthly return of 0.47% during the recession compared to 0.25% for the entire sample period benefit from capital inflows from the more risky equity markets. Junk bonds however are negatively hit by the adverse economic situation. A strong liquidity effect is not visible as there are no consistent changes in returns for portfolios with different liquidity levels and the liquid-minus-illiquid spread is never statistically different from zero. We can conclude that liquidity appears to not exert strong effects over the total course of the last macroeconomic crisis.

5.3.2 Liquidity Stress Periods

Although liquidity is not apparently visible during the macroeconomic recession in our bond sample, this sub-section shall examine whether specific short time periods exist in which liquidity unfolds severe effects. Therefore, we first have to define liquidity stress periods. Intuitively, high levels of market-wide illiquidity should be able to identify such periods. However, Figure 6 shows that illiquidity is non-stationary, i.e. illiquidity fluctuates around different levels for larger time spans. For example, the average illiquidity before the subprime crisis is notably higher than the illiquidity during recent years. Formulating dynamic criteria for triggering thresholds on illiquidity levels such as moving averages encounter further problems of parameterization. Therefore, we employ a Markov regime-switching model because of its power to identify periods with unusual statistical properties. We specify the model so as to uncover time periods in which the average market-wide bond return exhibits extreme liquidity betas on the basis of both the Amihud and the Extended Corwin-Schultz illiquidity measure[30]. The liquidity beta is chosen as the determining factor since liquidity should matter the most during periods of the highest sensitivity of bond returns to liquidity shocks. Extended with the Baum-Welch algorithm to increase accuracy, the Markov regime-switch results in a probability series over time that depicts the likelihood of a stress period for every point in time as shown in Figure 8. One clearly sees that not only there are abrupt positive shocks for the stress probabilities but also there is a clustering of stress probabilities during the months before, during, and after the Lehman Brothers bankruptcy. More pronounced than in the case of the Amihud measure, the Extended Corwin-Schultz measure leads to probability series with amplitudes of short duration before the subprime crisis in October 2006 (Ford Motor, one of the largest corporate borrowers in U.S. high-yield indexes, announces massive restructuring plans with huge layoff) and in November 2007 to January 2008 (several big banks reports remarkable losses due to poor investments in U.S. subprime mortgages and fears funding shortages). Both liquidity measures lead to similar regime-switching results and show a correlation of 0.71 between their probabilities series.

Since the Baum-Welch algorithm only returns a probabilities for stress periods over time, we apply the Viterbi algorithm to determine the so-called Viterbi path– the most likely binary series of states (either stress or normal period for every time-series observation) given a previously determined probability series – which consists of 10 months on the basis of the Amihud and of 14 months on the

and the return of the lowest quality portfolio (regardless of liquidity). Liquid-minus-illiquid is defined as the difference in the return of the highest liquidity portfolios (regardless of rating) and the return of the lowest liquidity portfolios (regardless of rating).

[30] For reasons of simplification we refrain from undertaking the regime-switch for all four illiquidity measures and instead limit our analysis to the Amihud and the Extended Corwin-Schultz measure to have one price impact measure and one bid-ask spread estimator, which both are superior to their counterparts.

Figure 8: Probability of High Illiquidity Regime from Markov Regime-Switching Model

This figure shows the probabilities of being in the stress regime (regime 2) estimated from the Markov regime-switching model (13) and refined by the Baum-Welch algorithm over our sample period 11/04 – 09/13. The liquidity factor in the Markov regime-switching regression is either based on the Amihud (Panel A) or the Extended Corwin-Schultz (Panel B) illiquidity measure.

Panel A: Regime-Switch Amihud Illiquidity Measure

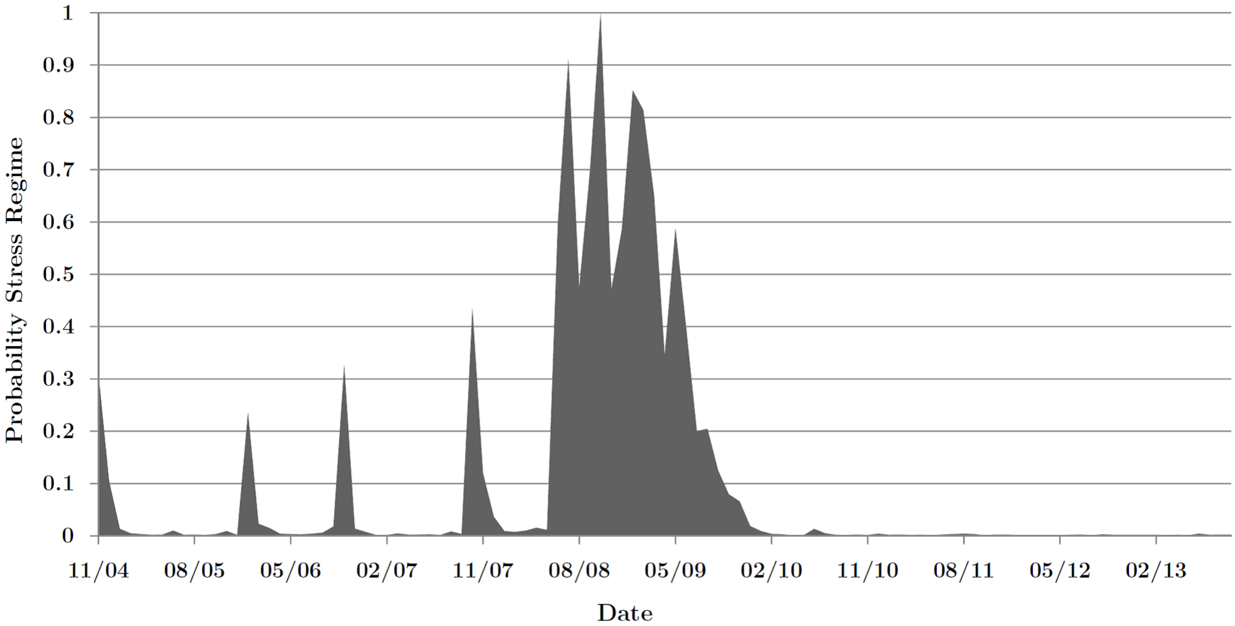

Panel B: Regime-Switch Extended Corwin-Schultz Illiquidity Measure

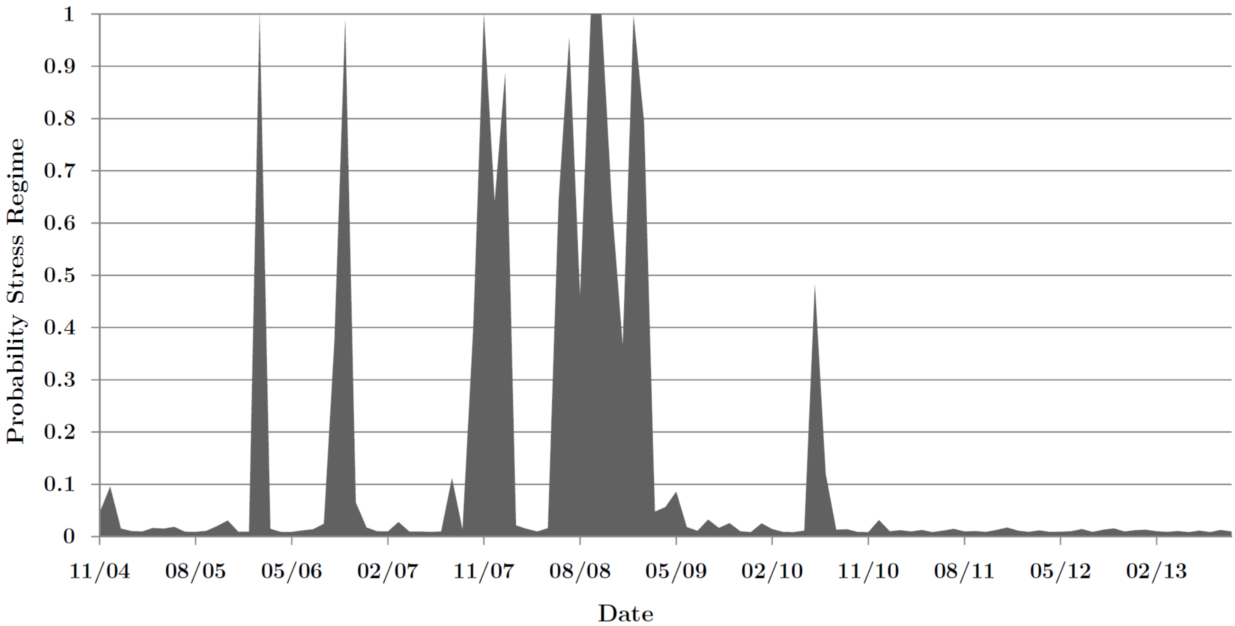

basis of the Extended Corwin-Schultz measure; they are shown in Figure 9. The run of both series illustrates that it is highly coincident with the stress probabilities for both regime-switches. The two regime series also show that the Viterbi algorithm ignores drops in the stress probability if those are girded by months with high stress probabilities (e.g. during 08/08 and 11/08 in both regime-switches) and that it ignores isolated stress probability shocks (e.g. during 05/09 for the regime-switch based on the Amihud measure). The latter two features of the Viterbi algorithm show the difference of undertaking this extra step for the stress period determination compared to a naïve approach that translates smoothed probabilities in regimes by the usage of some probability threshold analysis.

Figure 9: Regime Sequence Predicted by Viterbi Algorithm

This figure shows the most probable sequence of regimes (Viterbi path) given the probabilities of being in a normal or stress regime from Figure 8 and predicted by the Viterbi algorithm over our sample period 11/04 – 09/13. Panel A displays the Viterbi path for the regime-switch based on the Amihud measure and Panel B displays the Viterbi path for the regime-switch based on the Extended Corwin-Schultz measure. The liquidity stress regime is defined as regime 2 and the normal regime is defined as regime 1.

Panel A: Regime-Switch Amihud Illiquidity Measure

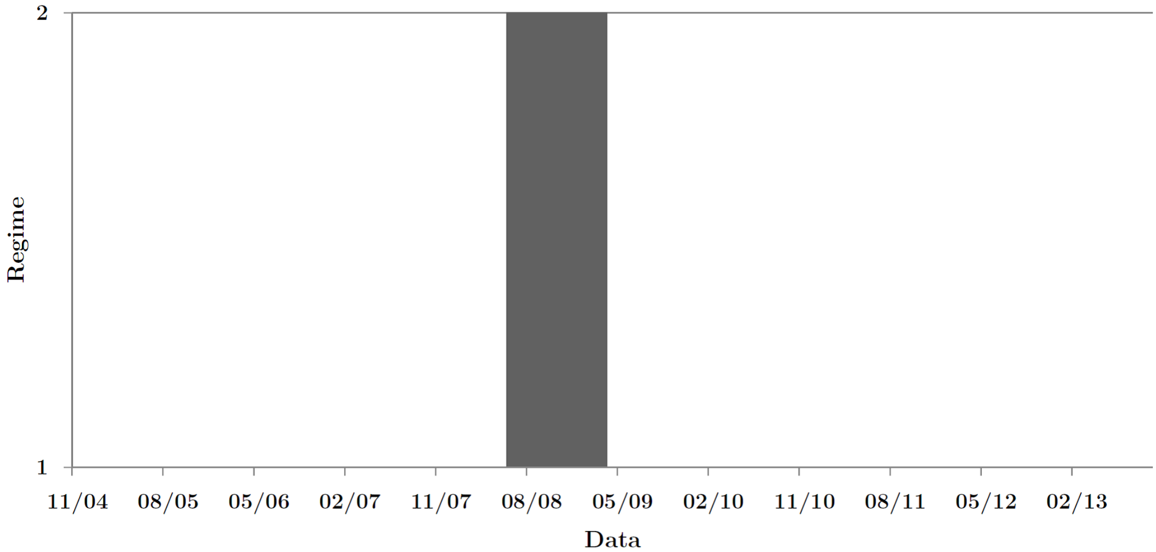

Panel B: Regime-Switch Extended Corwin-Schultz Illiquidity Measure

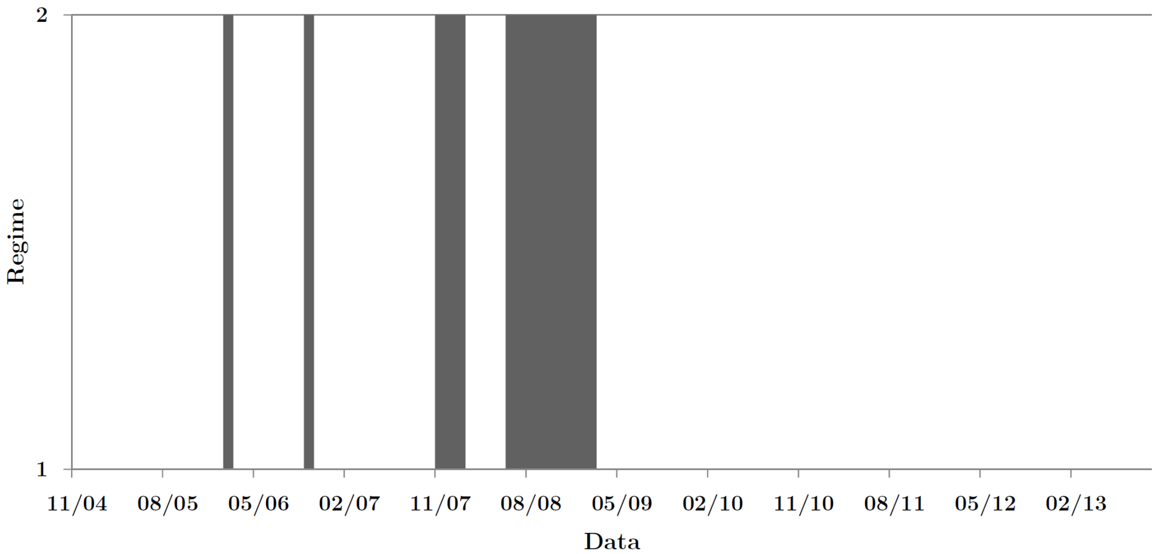

Although endogeneity might be the biggest criticism of using a regime-switching model in crisis identification, we do not undertake statistical robustness checks of our results using exogenously specified financial crisis periods in combination with probit and logit models such as Acharya et al. (2013) and Schuster and Uhrig-Homburg (2013), but intuitively compare our identified stress periods with the market-wide illiquidity levels and NBER recession as shown in Figure 10. One sees that the regime-switch meets its purpose of detecting liquidity stress regimes since it coincides with periods of high market-wide illiquidity relatively to its surrounding illiquidity levels. Moreover, the stress periods identified by the regime-switch coincide with the NBER recession, but seem to be more specific and focused to smaller time intervals.

Figure 10: Robustness Check Viterbi Path

This figure shows the Viterbi path against the NBER recession indicator and the market-wide illiquidity over our sample period 11/04 – 09/13. Panel A displays the situation with the Viterbi path for the regime-switch based on the Amihud measure and the market-wide illiquidity is based on the Amihud measure and Panel B displays the situation with the Viterbi path for the regime-switch based on the Extended Corwin-Schultz measure and the market-wide illiquidity is also based on the Extended Corwin-Schultz measure.

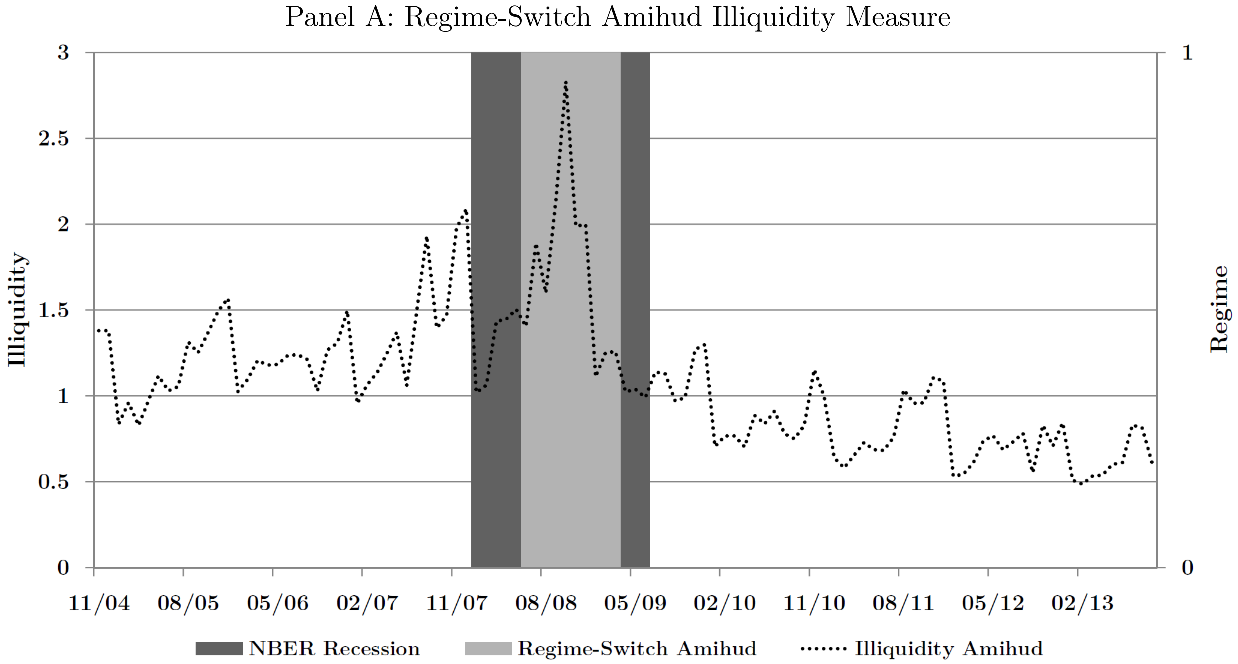

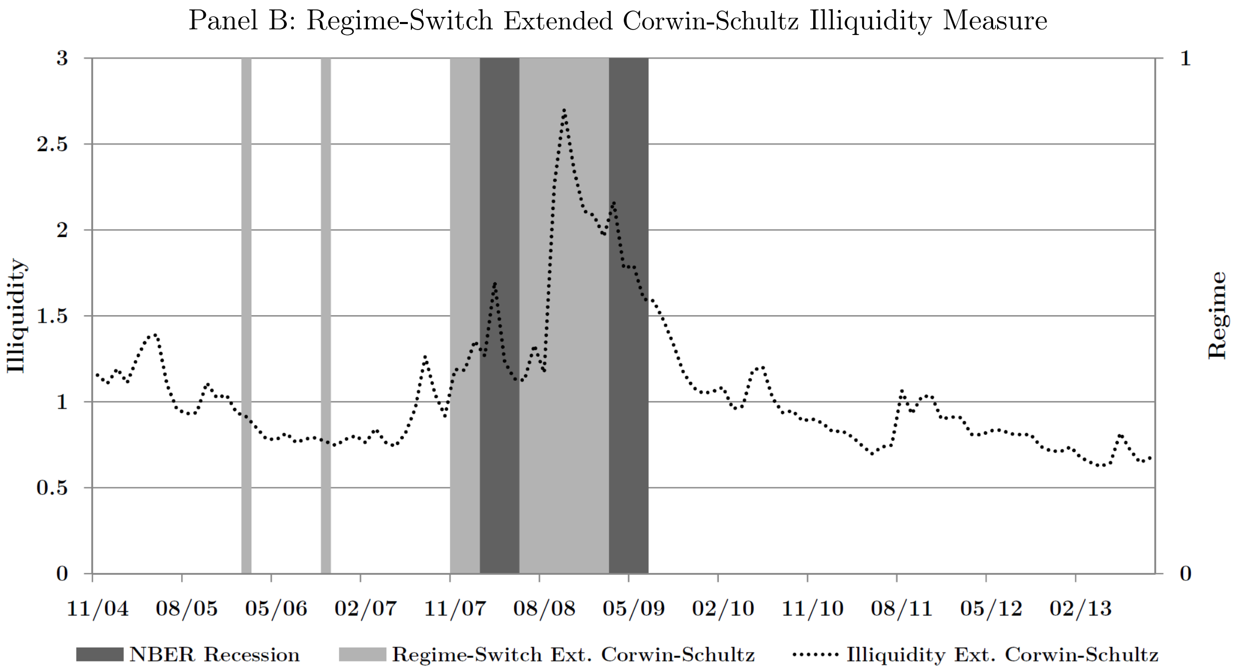

5.3.3 Return Analysis

Figure 11 and Table 13 (see Appendix) show the average monthly excess returns for each of our 25 portfolios during the stress months predicted by our two regime-switches.

Flight-from-Junk

It is eye-catching that for both regime-switches the returns on all lowest-quality (junk) portfolios are highly negative (ranging from -3.1% to -0.9% p.m.), while all other portfolios show just slightly negative returns with an average of -0.1% p.m. Besides the junk portfolios, the other quintiles show a slight increase in returns moving to higher rated bonds for the Extended Corwin-Schultz measure. However, one cannot speak of a strong intra-asset class flight-to-quality since the four highest quality quintiles show no strong gradual change in returns with respect to quality, although the quality-minus-junk spreads all show statistically significant positive spreads, which however mainly stem from the heavy flight-from-junk. Similar to the NBER recession, junk-bonds seem to be affected by a negative shock independent of the other bond portfolios which is why we speak of a flight-from-junk instead of an overarching flight-to-quality for the entire corporate bond market. Interestingly, the flight-to-quality effect is considerably stronger in the liquidity stress periods with an average return of -1.94% p.m. compared to an average return of -0.2% p.m. during the NBER recession, which is a first indication that our regime-switch specification identifies really unfavorable periods for high-yield bonds. Over the total liquidity stress time period, junk bonds fall by a significant extent of 21% on average.

Flight-to-Liquidity

Switching to the liquidity perspective, we can observe that there is a strong flight-to-liquidity within the lowest-rated bond portfolios since the most liquid portfolio outperforms the most illiquid portfolio by a large return of 2.2% p.m. for the Amihud and 1.2% p.m. for the Extended Corwin-Schultz measure. Importantly, the portfolios in between show a gradual return decline from the most liquid to the most illiquid portfolio. Although, all five lowest-rated bond portfolios exhibit negative returns, we still denote the liquidity effect as a "flight-to" situation: There is a general negative return shock to that credit rating, but relative to this shock, it appears as if investors show a gradually rising desire to hold liquid bonds. Among the four highest quality quintiles, we only see a weak flight-to-liquidity which is smooth for the Extended Corwin-Schultz, but quite noisy for the Amihud measure. Although this effect is nontrivial in nature, it could be very likely be driven by other risk factors, so that we do not see it as very strong evidence for a flight-to-liquidity. This is additionally supported by the t-tests on the liquid-minus-illiquid spreads, which show statistically significant results only for a few spreads. However, the flight-to-liquidity within junk bonds is extremely pronounced leading to an approximate 20% return differential between the most liquid and illiquid junk bonds over the entire stress period, which is highly statistically significant. We can conclude that our findings are in line with those of Dick-Nielsen et al. (2012), Friewald et al. (2012), and Acharya et al. (2013) in the sense that liquidity matters much more for speculative bonds during crisis periods. However, we differ with respect to the effect of liquidity on middle-rated bonds. While Friewald et al. (2012) and Acharya et al. (2013) only compare two groups, namely speculative and investment grade bonds, Dick-Nielsen et al. (2012) present the liquidity effects for five different rating categories. Interestingly, they find that AA-, A-, and BBB-rated bonds do also exhibit a large increase in sensitivity to liquidity in stress periods, while AAA-rated bonds are relatively unaffected. Contrary, our analysis shows that liquidity only affects the lowest quality quintile which ranges up to BB-rated bonds; liquidity does not affect ratings of BBB-rated or higher by a large extent. Hence, we find a much more isolated effect of the increasing importance of liquidity.

Since the stress periods are solely defined via liquidity characteristics, but show a strong quality effect, flight-to-liquidity and flight-to-quality appear to be closely related to each other. Comparing these results with the ones found for the NBER recession, we can infer that our regime-switch approach is able to identify the worst time period for corporate bonds in general since investors do not only flee from low-quality bonds to the highest extent but also punish illiquidity enormously. Although quality matters the most in absolute terms, it seems that the importance of liquidity rises by a larger extent than the one of quality. We now turn to the central question why only the speculative bond segment is affected by the heightened influence of liquidity.

5.3.4 Liquidity Shocks across Ratings

There are three hypothetical explanations about the isolated increase in the liquidity return differential for speculative bonds :

- If speculative bonds have generally a higher illiquidity, a rising liquidity premium would affect those bonds by a larger absolute magnitude.
- Liquidity shocks could be asymmetric during distressed times, i.e. primarily hitting illiquid speculative bonds, which would result in deteriorating characteristics relative to their liquid speculative counterparts.
- Assuming same liquidity (and liquidity shock) levels across credit ratings, investors may punish the same unit of illiquidity by a larger extent for speculative than investment grade bonds.

As shown in section 5.2 and in contrast to Acharya et al. (2013), liquidity is uniformly spread across credit ratings. Speculative and investment grade bonds show in general the same liquidity levels resulting in the first hypothesis to be obsolete. Figure 12 and Table 14 (see Appendix) show the difference in illiquidity levels of the stress periods to the non-stress levels and consequently displays the average illiquidity shocks during that time. Panel A displays the results on the basis of the Amihud and Panel B on the basis of the Extended Corwin-Schultz measure.

The crucial point is that illiquidity level shocks are commonly distributed across ratings for both measures. The dry out of market liquidity affects the price impact (proxied by the Amihud measure in Panel A) of illiquid bonds to a much higher absolute extent, while the most liquid bonds only increase slightly in their price impact dimension. Interestingly, the relative increase of the Amihud illiquidity is quite similar across the five liquidity quintiles: The illiquidity of the first (=most illiquid), second, third, fourth, and fifth liquidity quintiles of bonds increases by 37%, 62%, 54%, 39% and 48%, respectively. One the other hand, the bid-ask spread according to the Extended Corwin-Schultz

Figure 11: Portfolio-Sorted Excess Returns during Liquidity Stress Regime

This figure shows the average monthly bond portfolio returns in excess of the one-month T-bill return of 25 portfolios sorted by liquidity and quality during the liquidity stress regimes identified by the Viterbi paths. Illiquid and liquid are defined as the bond portfolios with the highest and lowest illiquidity, respectively, and junk and quality are defined as the the bond portfolios with the lowest and highest rating, respectively. Panel A displays the excess returns for the stress periods predicted by the regime-switch based on the Amihud measure and Panel B displays the excess returns for the stress periods predicted by the regime-switch based on the Extended Corwin-Schultz measure.

Panel A: Regime-Switch Amihud Illiquidity Measure

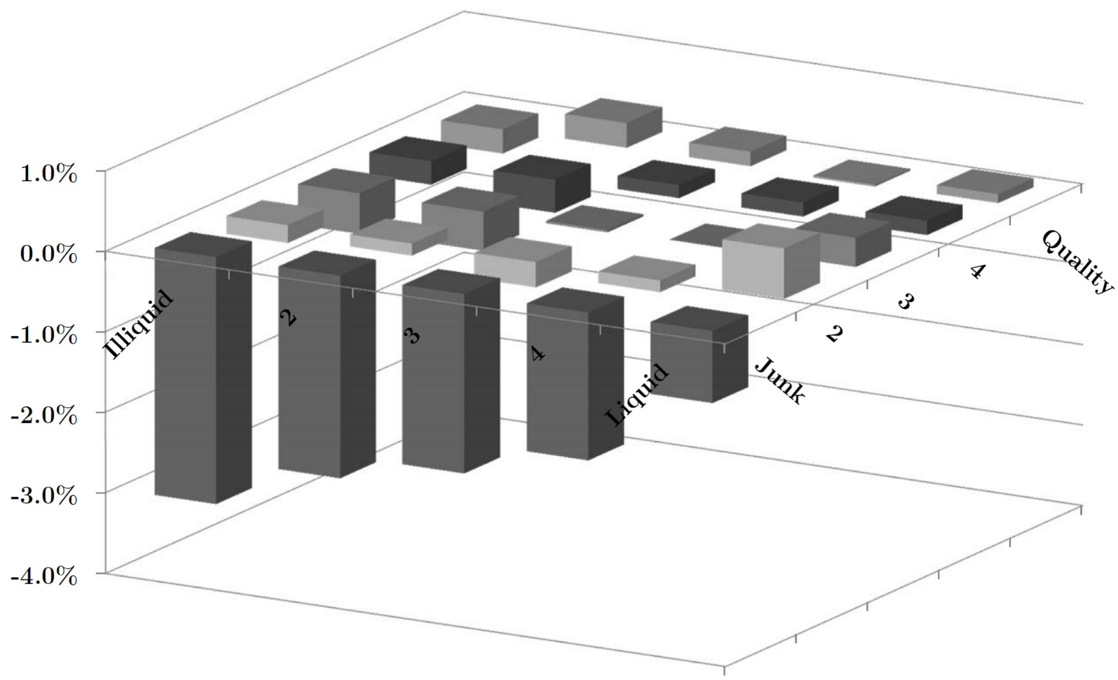

Panel B: Regime-Switch Extended Corwin-Schultz Illiquidity Measure

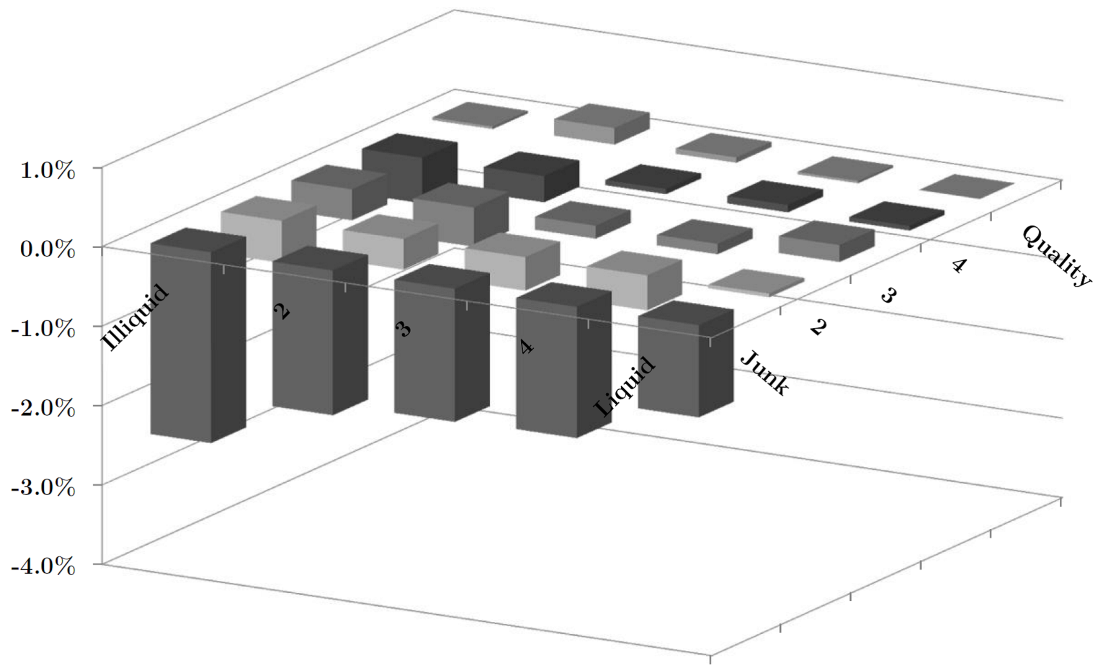

Figure 12: Illiquidity Level Shocks for Liquidity Stress Periods

This figure shows the difference in average liquidity level of 25 portfolios sorted by liquidity and quality during the liquidity stress regimes identified by the Markov regime-switch based on the Amihud illiquidity measure and the normal stress periods. Illiquid and liquid are defined as the bond portfolios with the highest and lowest illiquidity, respectively, and junk and quality are defined as the the bond portfolios with the lowest and highest rating, respectively. Panel A displays the average illiquidity shock for the Amihud measure and Panel B displays the the average illiquidity shock for the Extended Corwin-Schultz measure.

Panel A: Amihud Illiquidity Measure

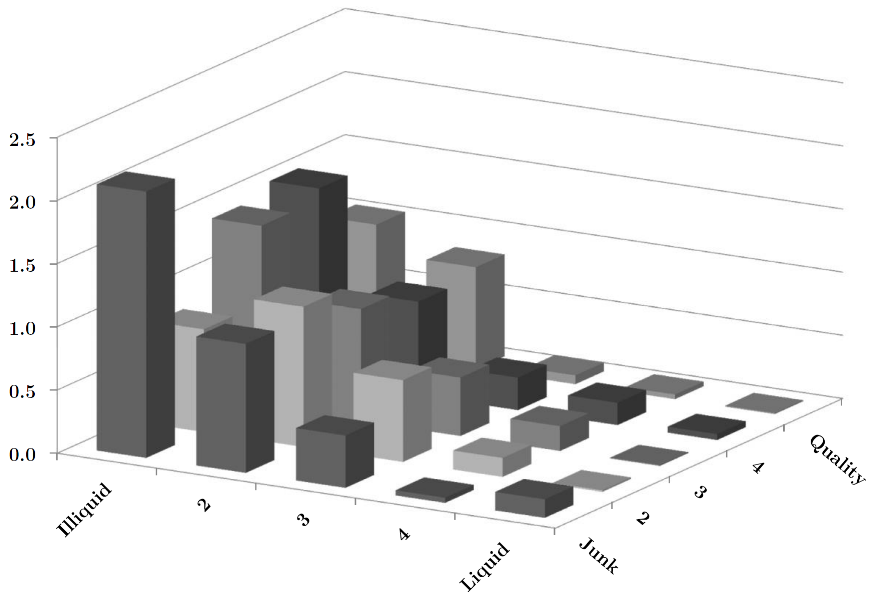

Panel B: Extended Corwin-Schultz Illiquidity Measure

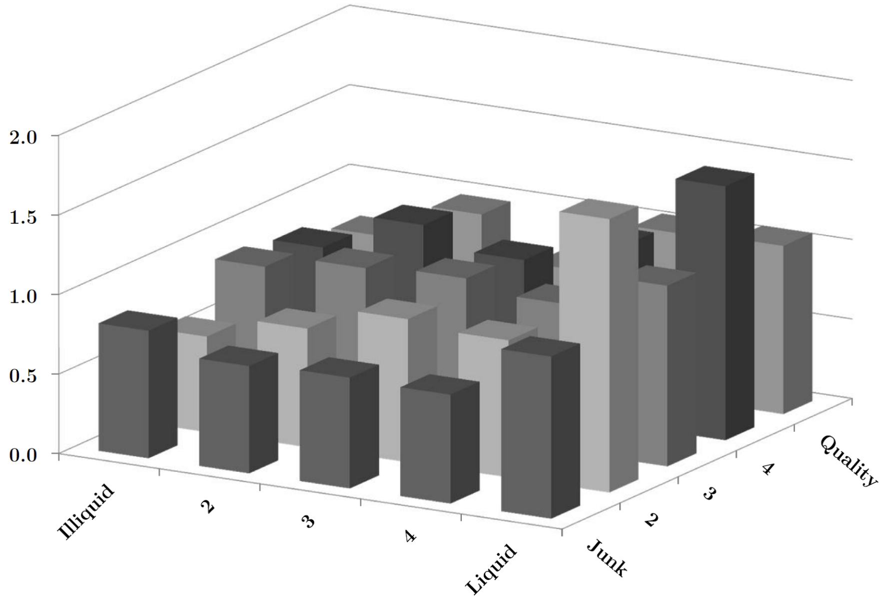

measure in Panel B shows a tendency to increase slightly more for liquid bonds. Both the price impact and bid-ask spread increases contradict the second hypothesis of an asymmetric illiquidity shock to ratings and leads to the *pivotal finding of our work*:

Liquidity stress periods appear to not affect bonds with ratings of BBB or above. During those times, not only the liquidity level premium stays the same, but also severe deterioration in liquidity levels fizzles out for investment grade bonds. Contrary, the most illiquid bonds underperform their liquid counterparts in the high-yield segment by c.20% over the course of the liquidity crisis. It appears that not only the increasing price-impact dispersion, but also a rising liquidity premium drives the liquidity return differential.

To illustrate that not only the larger increase in price-impact of illiquid compared to liquid high-yield bonds, but also a higher liquidity level premium drives the returns, we perform a parsimonious estimate of the relative effects. Under the assumption of a constant liquidity level premium over time, a shock to the illiquidity of a specific bond i should result in the following total price change (return)

$$\triangle P_i = -\triangle ILLIQ_{i,lvl} \cdot \lambda_{ILLIQ_{lvl}} \cdot D_i \tag{22}$$

where $\triangle ILLIQ_{i,lvl}$ is the change in the illiquidity level, $\lambda_{ILLIQ_{lvl}}$ is the illiquidity level risk premium and D_i is the duration of the bond. $\triangle ILLIQ_{i,lvl} \cdot \lambda_{ILLIQ_{lvl}}$ represents the change in expected return (=discount rate for the bond's cash flows) caused by the illiquidity shock. Since the duration represents not only a weighted maturity measure of the bond, but also the relative price change of the bond according to one percentage point change in the yield-to-maturity, the product of above depicts the total price change to the illiquidity shock. We use this formula to determine which constant risk premium would explain the return differential (of 20% between the most illiquid and the most liquid high-yield bonds ($\triangle P_{illiquid,HY}$-$\triangle P_{liquid,HY}$). Therefore, we assume an average duration of 8 years for all high-yield bonds (which roughly matches the maturities in Table 3) and base the analysis on the Amihud measure (since its illiquidity shocks are much more widespread). The Amihud measure of the most illiquid high-yield portfolio increases by approximately two standard deviations compared to its increase of the most liquid high-yield portfolio during the stress period. Plain algebra shows that, on the basis of these parameters, the constant liquidity level premium must amount to 1.25% per annum in order to explain the liquidity return differential. Our unconditional Fama-MacBeth procedure in section 5.2 estimates a risk premium of 0.23% for the Amihud measure, while De Jong and Driessen (2012) and Lin et al. (2011) estimate unconditional premia of up to 1%. Since the actual premium appears to be below 1.25%, the illiquidity shocks to the high-yield segment must lead to a lower spread in returns given the assumption of a constant premium. Therefore, we conclude that the liquidity return differential for speculative bonds cannot be fully explained by the increasing dispersion in illiquidity. It rather appears that the liquidity risk premium increases on top during liquidity crisis, but puzzlingly only for the high-yield segment.

6 Discussion

In the previous section we confirm the findings of Dick-Nielsen et al. (2012) and Friewald et al. (2012) that liquidity stress affects speculative bonds to a higher extent than their investment grade counterparts. The uniqueness of our work is the disentanglement of the underlying drivers. While the liquidity return differential of high-yield bonds appears to be caused by both a further increasing gap in liquidity and a rising liquidity premium, it remains puzzling why the severe shocks to the bond-specific illiquidities of bonds with rating BBB and above (that lie between 0% and 62%) are only marginal in effect. At this point, astrophysicist Sir Arthur Eddington would probably sound a note of caution similar to his reminder from 1934: "It is a good rule not to put overmuch confidence in the observational results that are put forward until they have been confirmed by theory." Therefore, this section will give possible economic explanations for liquidity effects during distressed periods. We first present the two theoretical models of Vayanos (2004) and Brunnermeier and Pedersen (2009) that explain the close connection between liquidity and quality effects visible in our sample. This interrelation does not represent a novelty to the academic literature. However, we build upon their understanding in order to explain the flight-from-junk and the unique finding of our work that investors punish the same unit of illiquidity level differently across credit ratings during times of distress.

6.1 Interrelations of Liquidity and Quality

Liquidity and quality stress periods are widely known to significantly co-move and reinforce each other. While quality matters also during normal time periods, liquidity is rather a *typically sleeping, but at times rampaging giant.* Just as in our analysis, liquidity does usually cause only moderate premia, but can rise rapidly in importance during high levels of uncertainty. This convex nature of liquidity to volatility makes liquidity concerns increase by a larger extent than quality matters during certain time periods. More precisely, the strong impact of liquidity appears to unfold in times of highest market turmoil; our liquidity stress period is accompanied by a loss of 20% of illiquid in relation to liquid bonds in the high-yield sector. Vayanos (2004) argues that the risk-aversion of fund managers increases due to the increasing threat of investors to withdraw capital in light of market stress. Brunnermeier and Pedersen (2009) describe a spiral in which increasing volatility and decreasing market liquidity reinforce each other through the effect of funding margins.

6.1.1 Increasing Risk-Aversion according to Vayanos (2004)

We apply the reasoning of Vayanos (2004) in order to explain an increase in risk-aversion during high-volatility times which amplifies liquidity concerns to a large extent. His model consists of fund managers that are subject to withdrawals for both random reasons and when the fund performance falls below a certain threshold. While Vayanos (2004) assumes that investors withdraw the fund capital all at once, we modify his model by introducing multiple thresholds that each only account for withdrawing a slice of the fund's total capital. As in real life, withdrawals are personally costly for the fund manager as they reduce the management fee and result in frictions such as the need for down-scaling of operations or a decreasing fund reputation. During volatile times or after negative market shocks, the probability that the fund performance falls below the threshold increases, and withdrawals become more likely. As a result, the fund manager becomes more risk-averse (in order to avoid the personal costs caused by withdrawals). Furthermore, the fund manager is faced with an increasing probability of being forced to liquidate positions to meet withdrawals. In case of withdrawals, transaction costs decrease the fund performance further (assuming that the fund value is based on gross or mid-point prices) so that the liquidity premium increases in line with the probability of withdrawals and hence also with the increase in overall risk aversion – thus, default risk as the primary component of a bond's risk and liquidity form part of the same equation. Importantly, the probability of withdrawals increases disproportionally in volatility resulting in the convex nature of risk-aversion and liquidity. The intuition is that when volatility is low, managers are not concerned with withdrawals because the event that performance falls below the threshold requires a movement of several standard deviations (i.e. not very likely). Thus, liquidity premia are very small, and almost insensitive to volatility. When volatility increases, however, the probability of withdrawals starts to increase rapidly, and so do the liquidity premia.

6.1.2 Price-Liquidity Spiral according to Brunnermeier and Pedersen (2009)

Brunnermeier and Pedersen (2009) show in their theoretical model that funding margins together with negative market shocks can cause liquidity dry ups so that market volatility and the level of market liquidity are closely related and reinforce each other. The mechanisms are illustrated in Figure 13. Their model consists of speculators who provide market liquidity (i.e. serve as market makers for customers) and finance their trades through collateralized borrowing from financiers who set margins to control their value-at-risk. Specifically, financiers set the margins large enough to ensure that a certain loss can only be exceeded by a certain probability (hence the underlying principle of a threshold of the capital-giving party is the same as in the model of Vayanos(2004)). Consequently, financiers increase the margins in case of intensified volatility and the speculators' equity falls below the required margins when they suffer losses. In both cases (higher volatility and losses) speculators have to reduce some positions in order to meet the margins. This forced sell-off pushes prices away from fundamentals which in turns leads to higher volatility (and hence higher margins) and further losses on existing positions. As a result, speculators have to continue unwinding assets resulting in a loss-margin spiral. Since speculators provide market liquidity by filling the temporary order imbalance of customers arriving sequentially to the market, a reduction in their positions results in less capital provision for keeping the market liquid. As a consequence, increasing volatility, market losses and market liquidity deterioration go hand in hand and amplify each other.

6.2 Isolated Effect of Liquidity Stress Periods on High-Yield Bonds

Next, we give an explanation for the pronounced flight-from-junk (large losses for high-yield bonds, while investment grade bonds remain relatively stable) and the limitation of liquidity effects on the high-yield segment during the liquidity stress periods.

6.2.1 Rational Explanations for the Flight-from-Junk

We find that high-yield bonds are much more strongly affected by the negative effects during the subprime crisis. In general, it is a stylized fact that high market uncertainty and volatility increase a firm's probability to default so that stress periods trigger overall rising default rates. Elton et al. (2001) show that, for high yield bonds, the marginal probability to default actually decreases over time,

Figure 13: Loss-Margin Spiral

This figure (taken from Brunnermeier and Pedersen (2009), page 2204) shows the origin and reinforcement of a liquidity spiral through a loss spiral and a margin/haircut spiral. A *margin spiral* emerges if margins are increasing in market illiquidity. In this case, a funding shock to the speculators lowers market liquidity, leading to higher margins, which tightens speculators' funding constraint further, and so on. A *loss spiral* arises if speculators hold a large initial position that is negatively correlated with customers' demand shock. In this case, a funding shock increases market illiquidity, leading to speculator losses on their initial position, forcing speculators to sell more, causing a further price drop, and so on. These liquidity spirals reinforce each other, implying a larger total effect than the sum of their separate effects.

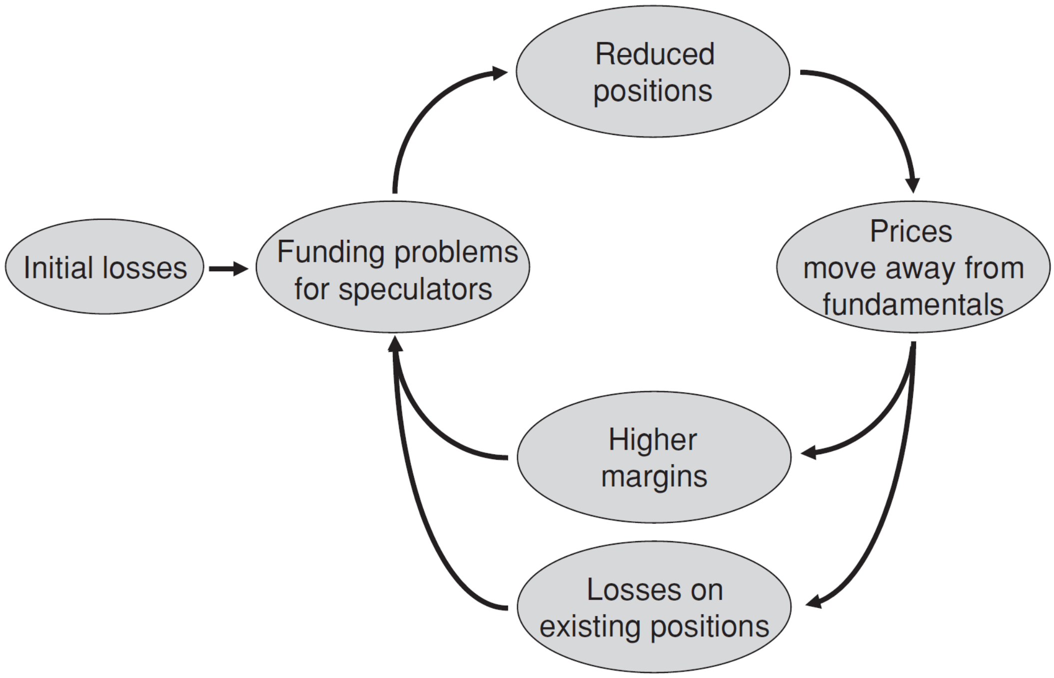

i.e. the bulk of the cumulative default probability concerns the next one to three years. Hence, a larger part of the default risk of high-yield bonds is caused by the ability for near-future financing so that high-yield bonds are more vulnerable to negative market shocks. Historical default rates underline this rationale: According to the Standard & Poor's Default Report 2013, the default rates of investment grade bonds over the course of the subprime crisis only slightly increased in absolute terms (0.00% in 2007, 0.42% in 2008, 0.32% in 2009 and 0.00% in 2010) while those of the high-yield bonds rose by a considerable large amount (0.90% in 2007, 3.65% in 2008, 9.75% in 2009 and 2.94% in 2010). Bernanke et al. (1994) argue that adverse shocks to the economy typically get amplified by worsening credit-market conditions. Thereby, lenders are reluctant to provide capital and are confronted with especially high agency costs concerning low-rated issuers resulting in disproportionally low credit extension for these borrowers. Consequently, the default risk of low-rated bonds is much more sensitive to market slowdowns than those of highly-rated bonds. However, we still find it puzzling that in our case only the lowest rating quintile of bonds suffers large losses during the stress period, while all four other quintiles are relatively unaffected and do not show a gradual price pattern moving from low to high-rated bonds, e.g. that the second lowest rating quintile also underperforms. Therefore, it would be interesting to examine the credit default spreads for different rating classes to examine whether a jump in the sensitivity of default probability to crisis periods exists between the lowest layer of bond ratings and the higher ones.

6.2.2 Non-Formal Model

The models of Vayanos (2004) and Brunnermeier and Pedersen (2009) both predict that rising market volatility increases the threat to unwind positions (in order to meet withdrawals or funding margins). We argue that this rising probability to liquidate positions, increases risk and the liquidity premium only for the assets likely to be sold. Since high-risk assets are typically sold off first in stress periods, the rising liquidity premium hits the high-yield bond segment the most, leading to a flight-from-junk. We then show that if investors perceive the accompanied liquidity dry out as transitory, they only penalize large negative liquidity shocks for assets that are likely to be sold during the stress times. This explains why the illiquidity shocks on investment grade bonds do only show a marginal return effect. Before explaining our rationale, we explicitly state the strongest assumptions of the model

- Confronted with liquidity needs during crises, investors tend to liquidate high-risk assets first.
- Investors care about the *effective* transactions costs, i.e. the transaction costs incurred through a transaction divided by the period for which the investor holds the asset, rather than the nominal transaction costs.
- Investors perceive liquidity dry outs during liquidity stress periods as transitory.

Increasing selling pressure on high-yield bonds to meet increasing withdrawals and funding margins

Just like Vayanos (2004) and Brunnermeier and Pedersen (2009), we assume that heightened market volatility increases the threat to unwind positions (in order to meet withdrawals or funding margins). While investors also hold a minimum tranche of assets as cushion for cash needs during normal times, they have to "put aside" more assets for potential liquidation in stress periods. Thus, we assume that rising volatility leads to an increasing probability for earlier than initially expected liquidation. The crucial point is the type of assets first in line for liquidation in case of cash needs: We argue that the increasing risk-aversion in stress times, makes investors unload high-risk and hence low-rated bonds first which is why the selling pressure increases risk premia, pushes prices down and thus explains the flight-from-junk. This effect gets amplified by a rising liquidity premium for this first sold-off tranche as explained in the following paragraph.

Increasing liquidity premia due to shortening holding period

Inspired by the clientele effect of Amihud and Mendelson (1986), we argue that a decrease in the expected holding period increases the assets' effective transaction costs (i.e. its yearly return deduction caused by transaction costs) so that the same nominal transaction costs (i.e. liquidity) hurt more leading to an increasing liquidity premium. We assume that an investor assesses an investment's return on the basis of the average net return over the holding period. Hence, the gross return is adjusted by the expected transaction costs incurred at the moment of sale and divided among the expected years of holding the asset. If, for example, the investor aims at holding a security for many years, the eventual transaction costs in the far future, only marginally affect the average net return over the holding period. Consequently, the investor cares about the transaction costs per holding period (hereinafter called the *effective* transaction costs). If now an unexpected rise in the probability to sell-off positions (caused by rising withdrawals or funding margins) occurs, the expected holding period shortens so as to increase the *effective* transaction costs. Hence, the same level of transaction costs and thus illiquidity hurts the investor more, resulting in higher liquidity premia.As shown above, a rising volatility leads to an increasing probability for earlier than initially expected liquidation. As high-yield bonds are assumed to be the tranche first in line to be sold (shown above), the effective transaction costs decreasemore for these bonds.

The asymmetric effect of general liquidity dry outs

Besides being amplified by increasing liquidity premia, the net effect of liquidity also increases due to general liquidity dry outs during stress periods. As Brunnermeier and Pedersen (2009) show, high volatility is accompanied by decreasing market liquidity since both trigger and reinforce each other. We now assume that investors perceive those large negative liquidity shocks as only transitory phenomena, i.e. investors believe that liquidity will return to normal levels after a certain period of time. Consequently, the effective transaction due to increasing temporary illiquidity only for assets that are likely to be sold during the stress times. Since high-yield bonds are first in line to be sold, and thus better-rated bonds more likely to be unwound later when markets have returned to normality, investment grade bonds only show marginal return effects to these temporary illiquidity shocks.

7 Conclusion

The purpose of this study is to answer the question why high-yield bonds are more severely hit by large liquidity dry outs than investment grade bonds. Before examining the underlying drivers, we first apply four liquidity measures, namely the Amihud, Pastor-Stambaugh, Roll, and Corwin-Schultz estimators, on our data set in order to proxy for the not directly observable phenomenon of liquidity. In order to apply the idea of the Corwin-Schultz measure on infrequently traded bonds, we algebraically derive a more flexible, enhanced version and further show that this extension provides an equal or more accurate estimation of the bid-ask spread in c.85% of the transactions in our sample compared to the original measure.

In order to examine whether our liquidity measures accurately measure liquidity, we estimate unconditional risk premia of up to 1% for the liquidity level and up to 0.5% for the liquidity risk whose magnitudes are in line with those found in the academic literature. The determination of liquidity premia in our sample serves as a good case study in order to illustrate the potential flaws that realized returns as proxies for expected returns can exhibit. Furthermore, we show that liquidity level and liquidity risk must both be included in liquidity premia analyses since omitting one of them may overstate the other's effect.

We show that during liquidity stress periods as identified by our Markov regime-switching model (comprising 10 to 14 months, mainly during the subprime crisis) illiquid bonds underperform their liquid counterparts by as much as 21.7% in the high-yield segment, while the same return differential amounts to only 5.4% for investment grade bonds. In order to return to the main question of our work, we then examine the potential underlying drivers of this asymmetric effect on return differentials. We show that classical explanatory approaches fail to describe this difference in the effect of liquidity on returns. Neither the pre-crisis liquidity levels nor the liquidity shocks during the stress periods show an asymmetric distribution across credit ratings.

As a result, we infer the puzzling conclusion that investors punish the same unit of illiquidity differently across credit ratings in times of distress. During those times, a severe decrease in the bond-specific liquidity levels across all ratings hits high-yield bonds by a large extent, while ratings of BBB and above are only marginally affected by the same magnitude of shock. Investors appear to disregard liquidity shocks to investment grade bonds during times of distress. Besides not only arguing that investors punish the same magnitude of liquidity shock differently across ratings, we show that the liquidity premium must exclusively increase for high-yield bonds during stress times in order to explain the observed price differentials.

In order to explain the discriminating treatment of the same unit of liquidity, we develop a model in non-formal reasoning. Our explanation for the asymmetric liquidity effects also sheds light on another phenomenon visible in our sample: We observe a *flight-from-junk* effect in the form that high-yield bonds drop by about 20%, while investment grade bonds suffer only 2% over the course of the identified liquidity stress period. It is grounded on the idea that investors perceive liquidity dry outs as transitory and therefore only penalize assets that are likely to be sold in the short-term. Since investors are more risk-averse in times of distress, high-yield bonds are the first in line to be liquidated, and thus investment grade bonds only show marginal return effects with respect to liquidity shocks, perceived to be of temporary nature.

These results can be useful for several practical applications. First, it may serve as a basis to further clarify the extent to which apparent *abnormal returns* of investment strategies simply display liquidity premia or are indeed harvested by superior skills. Second, it may shed light on potential policy reactions to severe liquidity dry outs in the corporate bond market. As predicted by the model of Brunnermeier and Pedersen (2009), decreasing liquidity and increasing volatility can reinforce each other. In order to break the downward spiral, policy makers would be advised to direct potential liquidity provisions (comparable to the current quantitative easing in the U.S.) on the high-yield sector. However, we want to cite Tolstoy from his novel Anna Karenina: “Happy families are all alike; every unhappy family is unhappy in its own way”. As our analysis is predominantly based on the subprime crisis, future liquidity stress periods may exert different characteristics. However, we hope that our analysis serves as a good starting point to understand the behavior of corporate bond returns and thus of investors themselves.

References

Acharya, Viral V., Yakov Amihud, and Sreedhar T. Bharath. 2013. "Liquidity Risk of Corporate Bond Returns: Conditional Approach." *Journal of Financial Economics* 110 (2): 358–86.

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

Altman, Edward I., Naeem Hukkawala, and Vellore Kishore. 2000. "Defaults & Returns on High Yield Bonds: Analysis Through 1999 and Default Outlook for 2000-2002." *Working Paper*, New York University.

Amihud, Yakov. 2002. "Illiquidity and Stock Returns: Cross-Section and Time-Series Effects." *Journal of Financial Markets* 5 (1): 31–56.

Amihud, Yakov, and Haim Mendelson. 1986. "Asset Pricing and the Bid-Ask Spread." *Journal of Financial Economics* 17 (2): 223–49.

Amihud, Yakov, Haim Mendelson, and Lasse Heje Pedersen. 2005. "Liquidity and Asset Prices." *Foundations and Trends in Finance* 1 (4): 269–364.

Amihud, Yakov, Haim Mendelson, and Lasse Heje Pedersen. 2013. *Market Liquidity: Asset Pricing, Risk and Crises.* New York, New York: Cambridge University Press.

Arora, Navneet, Priyank Gandhi, and Francis A. Longstaff. 2012. "Counterparty Credit Risk and the Credit Default Swap Market." *Journal of Financial Economics* 103 (2): 280–93.

Asness, Cliff S., Tobias J. Moskowitz, and Lasse Heje Pedersen. 2013. "Value and Momentum Everywhere." *Journal of Finance* 68 (3): 929–85.

Baele, Lieven, Geert Bekaert, Koen Inghelbrecht, and Min Wei. 2013. "Flights to Safety.“ Netspar Discssion Paper 05/2013-034.

Baghai, Ramin P., Henri Servaes, and Ane, Tamayo. 2014. "Have Rating Agencies Become More Conservative? Implications for Capital Structure and Debt Pricing.“ *Journal of Finance, forthcoming.*

Bao, Jack, Jun Pan, and Jiang Wang. 2011. "The Illiquidity of Corporate Bonds." *Journal of Finance* 66 (3): 911–46.

Bar-Isaac, Heski, and Joel Shapiro. 2013. "Ratings Quality over the Business Cycle." *Journal of Financial Economics* 108 (1): 62–78.

Barrell, Ray, E. Philip Davis, Dilruba Karim, and Iana Liadze. 2010. "Bank Regulation, Property Prices and Early Warning Systems for Banking Crises in OECD Countries." *Journal of Banking and Finance* 34 (9): 2255–64.

Baum, Leonard E., Ted Petrie, George Soules, and Norman Weiss. 1970. "A Maximization Technique Occurring in the Statistical Analysis of Probabilistic Function of Markov Chains." *The Annals of Mathematical Statistics* 41 (1): 164–71.

Baur, Dirk G., and Brian M. Lucey. 2009. "Flights and Contagion – An Empirical Analysis of Stock-Bond Correlations." *Journal of Financial Stability* 5 (4): 339–52.

Beber, Alessandro, Michael W. Brandt, and Kenneth A. Kavajecz. 2009. "Flight-to-Quality or Flight-to-Liquidity? Evidence from the Euro-Area Bond Market." *Review of Financial Studies* 22 (3): 925–57.

Becker, Bo, and Todd Milbourn. 2011. "How Did Increased Competition Affect Credit Ratings?" *Journal of Financial Economics* 101 (3): 493–514.

Bekaert, Geert, Eric Engstrom, and Yuhang Xing. 2009. "Risk, Uncertainty, and Asset Prices." *Journal of Financial Economics* 91 (1): 59–82.

Bekaert, Geert, Campbell R. Harvey, and Christian Lundblad. 2007. "Liquidity and Expected Returns: Lessons from Emerging Markets." *Review of Financial Studies* 20 (6): 1783–831.

Bernanke, Ben, Mark Gertler, and Simon Gilchrist. 1994. "The Financial Accelerator and the Flight to Quality." *Working Paper*, National Bureau of Economic Research.

Bilmes, Jeff A. 1998. "A Gentle Tutorial of the EM Algorithm and its Application to Parameter Estimation for Gaussian Mixture and Hidden Markov Models." Technical Report TR-97-021, International Computer Science Institute Berkeley.

Bolton, Patrick, Xavier Freixas, and Joel Shapiro. 2012. "The Credit Ratings Game." *Journal of Finance* 67 (1): 85–111.

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 (1): 203–40.

Bongaerts, Dion, Frank de Jong, and Joost Driessen. 2012. "An Asset Pricing Approach to Liquidity Effects in Corporate Bond Markets." *Working Paper.*

Brennan, Michael J., and Avanidhar Subrahmanyam. 1996. "Market Microstructure and Asset Pricing: On the Compensation for Illiquidity in Stock Returns." *Journal of Financial Economics* 41 (3): 441–64.

Brennan, Michael J., Tarun Chordia, and Avanidhar Subrahmanyam. 1998. "Alternative Factor Specifications, Security Characteristics, and the Cross-Section of Expected Stock Returns." *Journal of Financial Economics* 49 (3): 345–73.

Brunnermeier, Markus K., and Lasse Heje Pedersen. 2009. "Market Liquidity and Funding Liquidity." *Review of Financial Studies* 22 (6): 2201–38.

Burge, Chris, and Samuel Karlin. 1997. "Prediction of Complete Gene Structures in Human Genomic DNA." *Journal of Molecular Biology* 268 (1): 78–94.

Bushman, Robert M., Anh Le, and Florin P. Vasari. 2010. "Implied Bond Liquidity." *Working Paper*, University of North Carolina at Chapel Hill.

Campbell, John Y., Andrew W. Lo, A. Craig MacKinlay, and Robert F. Whitelaw. 1998. "The Econometrics of Financial Markets." *Macroeconomic Dynamics* 2 (4): 559–62.

Campello, Murillo, Long Chen, and Lu Zhang. 2008. "Expected Returns, Yield Spreads, and Asset Pricing Tests." *Review of Financial Studies* 21 (3): 1297–338.

Chacko, George. 2006. "Liquidity Risk in the Corporate Bond Markets." *Working Paper*, Harvard Business School.

Chen, Long, David A. Lesmond, and Jason Wei. 2007. "Corporate Yield Spreads and Bond Liquidity." *Journal of Finance* 62 (1): 119–49.

Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam. 2000. "Commonality in Liquidity." *Journal of Financial Economics* 56 (1): 3–28.

Chordia, Tarun, Asani Sarkar, and Avanidhar Subrahmanyam. 2005. "An Empirical Analysis of Stock and Bond Market Liquidity." *Review of Financial Studies* 18 (1): 85–129.

Cochrane, John H. 2005. *Asset Pricing Revised Edition.* Princeton, NJ: Princeton University Press.

Collin-Dufresne, Pierre, Robert S. Goldstein, and J. Spencer Martin. 2001. "The Determinants of Spread Changes." *Journal of Finance* 56 (6): 2177–207.

Corwin, Shane A., and Paul H. Schultz. 2012. "A Simple Way to Estimate Bid-Ask Spreads from Daily High and Low Prices." *Journal of Finance* 67 (2): 719–59.

Covitz, Dan, and Chris Downing. 2007. "Liquidity or Credit Risk? The Determinants of Very Short-Term Corporate Yield Spreads." *Journal of Finance* 62 (5): 2303–28.

Das, Sanjiv, Madhu Kalimipalli and Subhankar Nayak. 2014. "Did CDS Trading Improve the Market for Corporate Bonds?" *Journal of Financial Economics* 111 (2): 495–525.

De Jong, Frank, and Joost Driessen. 2012. "Liquidity Risk Premia in Corporate Bond Markets." *Quarterly Journal of Finance* 2 (2).

Dempster, A. P., N. M. Laird, and D. B. Rubin. 1977. "Maximum Likelihood from Incomplete Data via the EM Algorithm." *Journal of the Royal Statistical Society. Series B (Methodological)* 39 (1): 1–38.

Deuskar, Prachi, Anurag Gupta, and Marti G. Subrahmanyam. 2011. "Liquidity Effect in OTC Options Markets: Premium or Discount?" *Journal of Financial Markets* 14 (1): 127–60.

Dick-Nielsen, Jens. 2009. "Liquidity Biases in TRACE." *Journal of Fixed Income* 19 (2): 43–55.

Dick-Nielsen, Jens, Peter Feldhütter, and David Lando. 2012. "Corporate Bond Liquidity Before and After the Onset of the Subprime Crisis." *Journal of Financial Economics* 103 (3): 471–92.

Dong, Xi, Shu Feng, and Ronnie Sadka. 2014. "Does Liquidity Beta Predict Mutual-Fund Alpha?" *Working Paper*, Boston College.

Downing, Chris, Shane Underwood, and Yuhang Xing. 2005. "Is Liquidity Risk Priced in the Corporate Bond Market?" *Working Paper*, Rice University.

Edwards, Amy K., Lawrence E. Harris, and Michael S. Piwowar. 2007. "Corporate Bond Market Transaction Costs and Transparency." *Journal of Finance* 62 (3): 1421–51.

Elton, Edwin J. 1999. "Presidential Address: Expected Return, Realized Return, and Asset Pricing Rests." *Journal of Finance* 54 (4): 1199–220.

Elton, Edwin J., Martin J. Gruber, Deepak Agrawal, and Christopher Mann. 2001. "Explaining the Rate Spread on Corporate Bonds." *Journal of Finance* 56 (1): 247–77.

Eom, Young Ho, Jean Helwege, and Jing-Zhi Huang. 2004. "Structural Models of Corporate Bond Pricing: An Empirical Analysis." *Review of Financial Studies* 17 (2): 499–544.

Ericsson, Jan, Kris Jacobs, and Rodolfo A. Oviedo. 2009. "The Determinants of Credit Default Swap Premia." *Journal of Financial and Quantitative Analysis* 44 (1): 109–32.

Fama, Eugene F., and Kenneth R. French. 1993. "Common Risk Factors in the Returns of Stocks and Bonds." *Journal of Financial Economics* 33 (1): 3–56.

Fama, Eugene F., and James D. MacBeth. 1973. "Risk, Return, and Equilibrium: Empirical Tests." *Journal of Political Economy* 81 (3): 607–36.

Feldhütter, Peter. 2012. "The Same Bond at Different Prices: Identifying Search Frictions and Selling Pressures." *Review of Financial Studies* 25 (4): 1155–206.

Flannery, Mark J., Joel F. Houston, and Frank Partnoy. 2010. "Credit Default Swap Spreads as Viable Substitutes for Credit Ratings." *University of Pennsylvania Law Review* 158 (7): 2085–123.

Fong, Kingsley Y. L., Craig W. Holden, and Charles A. Trzcinka. 2013. "What Are the Best Liquidity Proxies for Global Research?" *Working Paper.*

Forney Jr., G. David. 1973. "The Viterbi Algorithm." *Proceedings of the IEEE* 61 (3): 268–78.

Franzoni, Francesco, Eric Nowak, and Ludovic Phalippou. 2012. "Private Equity Performance and Liquidity Risk." *Journal of Finance* 67 (6): 2341–73.

Friewald, Nils, Rainer Jankowitsch, and Marti G. Subrahmanyam. 2012. "Illiquidity or Credit Deterioration: A Study of Liquidity in the U.S. Corporate Bond Market during Financial Crises." *Journal of Financial Economics* 105 (1): 18–36.

Garman, Mark B., and Michael J. Klass. 1980. "On the Estimation of Security Price Volatilities from Historical Data." *Journal of Business* 53 (1): 67–78.

Gebhardt, William R., Soeren Hvidkjaer, and Bhaskaran Swaminathan. 2005. "The Cross-Section of Expected Corporate Bond Returns: Betas or Characteristics?" *Journal of Financial Economics* 75 (1): 85–114.

Goldstein, Michael A., Edith S. Hotchkiss, and Erik R. Sirri. 2007. "Transparency and Liquidity: A Controlled Experiment on Corporate Bonds." *Review of Financial Studies* 20 (2): 235–73.

Goyenko, Ruslan Y., Craig W. Holden, and Charles A. Trzcinka. 2009. "Do Liquidity Measures Measure Liquidity?" *Journal of Financial Economics* 92 (2): 151–82.

Goyenko, Ruslan Y., and Sergei Sarkissian. 2014. "Treasury Bond Illiquidity and Global Equity Returns." *Journal of Financial and Quantitative Analysis, forthcoming.*

Guidolin, Massimo. 2011a. "Markov Switching Models in Empirical Finance." in David M. Drukker (ed.): *Missing Data Methods: Time-Series Methods and Applications (Advances in Econometrics, Volume 27)*, Emerald Group Publishing Limited: 1–86.

Guidolin, Massimo. 2011b. "Markov Switching in Portfolio Choice and Asset Pricing Models: A Survey." in David M. Drukker (ed.): *Missing Data Methods: Time-Series Methods and Applications*

(Advances in Econometrics, Volume 27), Emerald Group Publishing Limited: 87–178.

Hamilton, James D. 1989. "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle." *Econometrica* 57 (2): 357–84.

Hamilton, James D. 1994. *Time Series Analysis.* Princeton, NJ: Princeton University Press.

Hamilton, James D. 2008. "Regime-Switching Models." In Steven Durlauf and Lawrence Blume (ed.): *New Palgrave Dictionary of Economics* (Palgrave McMillan Ltd., Basingstoke).

Han, Song, and Hao Zhou. 2011. "Effects of Liquidity on the Nondefault Component of Corporate Yield Spreads: Evidence from Intraday Transactions Data." HKIMR Working Paper No.02/2011.

Hasbrouck, Joel. 2009. "Trading Costs and Returns for U.S. Equities: Estimating Effective Costs from Daily Data." *Journal of Finance* 64 (3): 1445–77.

Hasbrouck, Joel, and Duane J. Seppi. 2001. "Common Factors in Prices, Order Flows, and Liquidity." *Journal of Financial Economics* 59 (3): 383–411.

Houweling, Patrick, Albert Mentink, and Ton Vorst. 2005. "Comparing Possible Proxies of Corporate Bond Liquidity." *Journal of Banking and Finance* 29 (6): 1331–58.

Huang, Jing-Zhi, and Ming Huang. 2012. "How Much of the Corporate-Treasury Yield Spread is Due to Credit Risk?" *Review of Asset Pricing Studies* 2 (2): 153–202.

Huberman, Gur, and Dominika Halka. 2001. "Systematic Liquidity." *Journal of Financial Research* 24 (2): 161–78.

Ibáñez, Jesús M., Carmen Benítez, Ligdamis A. Gutiérrez, Guillermo Cortés, Araceli García-Yeguas, and Gerardo Alguacil. 2009. "The Classification of Seismo-Volcanic Signals Using Hidden Markov Models as Applied to the Stromboli and Etna Volcanoes." *Journal of Volcanology and Geothermal Research* 187 (3–4): 218–26.

Jankowitsch, Rainer, Amrut Nashikkar, and Marti G. Subrahmanyam. 2011. "Price Dispersion in OTC Markets: A New Measure of Liquidity." *Journal of Banking and Finance* 35 (2): 343–57.

Juang, Biing-Hwang, and Lawrence R. Rabiner. 1991. "Hidden Markov Models for Speech Recognition." *Technometrics* 33 (3): 251–72.

Junge, Benjamin, and Anders B. Trolle. 2013. "Liquidity Risk in Credit Default Swap Markets." Swiss Finance Institute Research Paper No. 13-65.

Karstanje, Dennis, Elvira Sojli, Wing WahTham, and Michel van der Wel. 2013. "Economic Valuation of Liquidity Timing." *Journal of Banking and Finance* 37 (12): 5073–87.

Kim, Soon-Hu, and Kuan-Hui Lee. 2014. "Pricing of Liquidity Risks: Evidence from Multiple Liquidity Measures." *Journal of Empirical Finance* 25: 112–33.

Korajczyk, Robert, and Ronnie Sadka. 2008. "Pricing the Commonality across Alternative Measures of Liquidity." *Journal of Financial Economics* 87 (1): 45–72.

Krogh, Anders, Björn Larsson, Gunnar von Heijne, and Erik L. L. Sonnhammer. 2001. "Predicting Transmembrane Protein Topology with a Hidden Markov Model: Application to Complete Genomes." *Journal of Molecular Biology* 305 (3): 567–80.

Kyle, Albert S. 1985. "Continuous Auctions and Insider Trading." *Econometrica* 56 (4): 1315–35.

Lee, Kuan-Hui. 2011. "The World Price of Liquidity Risk." *Journal of Financial Economics* 99 (1): 131–61.

Lesmond, David A., Joseph P. Ogden, and Charles A. Trzcinka. 1999. "A New Estimate of Transaction Costs." *Review of Financial Studies* 12 (5): 1113–41.

Lin, Hai, Junbo Wang, and Chunchi Wu. 2011. "Liquidity Risk and Expected Corporate Bond Returns." *Journal of Financial Economics* 99 (3): 628–50.

Liu, Weimin. 2006. "A Liquidity-Augmented Capital Asset Pricing Model." *Journal of Financial Economics* 82 (3): 631–71.

Lou, Hui-Ling. 1995. "Implementing the Viterbi Algorithm." *IEEE Signal Processing Magazine* 12 (5): 42–52.

Lundblad, Christian. 2005. "The Risk Return Tradeoff in the Long Run: 1836–2003." *Journal of Financial Economics* 85 (1): 123–50.

Mahanti, Sriketan, Amrut Nashikkar, Marti G. Subrahmanyam, George Chacko, and Gaurav Mallik. 2008. "Latent Liquidity: A New Measure of Liquidity, with an Application to Corporate Bonds." *Journal of Financial Economics* 88 (2): 272–98.

Marshall, Ben R., Nhut H. Nguyen, and Nuttawat Visaltanachoti. 2012. "Commodity Liquidity Measurement and Transaction Costs." *Review of Financial Studies* 25 (2): 599–638.

Mitra, Sovan, and Paresh Date. 2010. "Regime Switching Volatility Calibration by the Baum–Welch Method." *Journal of Computational and Applied Mathematics* 234 (12): 3243–60.

Nashikkar, Amrut, Marti G. Subrahmanyam, and Sriketan Mahanti. 2011. "Liquidity and Arbitrage in the Market for Credit Risk." *Journal of Financial and Quantitative Analysis* 46 (3): 627–56.

Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55 (3): 703–8.

Parkinson, Michael. 1980. "The Extreme Value Method for Estimating the Variance of the Rate of Return." *Journal of Business* 53 (1): 61–5.

Pastor, Lubos, Meenakshi Sinha, and Bhaskaran Swaminathan. 2008. "Estimating the Intertemporal Risk–Return Tradeoff Using the Implied Cost of Capital." *Journal of Finance* 63 (6): 2859–97.

Pastor, Lubos, and Robert F. Stambaugh. 2003. "Liquidity Risk and Expected Stock Returns." *Journal of Political Economy* 111 (3): 642–85.

Perlin, Marcelo. 2012. "MS_Regress – The Matlab Package for Markov Regime Switching Models." Available at http://ssrn.com/abstract=1714016.

Rabiner, Lawrence R. 1989. "A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition." *Proceedings of the IEEE* 77 (2): 257–86.

Roll, Richard. 1984. "A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market." *Journal of Finance* 39 (4): 1127–39.

Ronen, T., X. Zhou. 2009. "Where Did All the Information Go? Trade in the Corporate Bond Market." *Working Paper*, Rutgers University.

Sadka, Ronnie. 2006. "Momentum and Post-Earnings-Announcement Drift Anomalies: The Role of Liquidity Risk." *Journal of Financial Economics* 80 (2): 309–49.

Sadka, Ronnie. 2010. "Liquidity Risk and the Cross-Section of Hedge-Fund Returns." *Journal of Financial Economics* 98 (1): 54–71.

Sadka, Ronnie. 2012. "Hedge Fund Performance and Liquidity Risk." *Journal of Investment Management* 10 (2): 60–72.

Schestag, Raphael, Philipp Schuster, and Marliese Uhri-Homburg. 2013. "Measuring Liquidity in Bond Markets." *Working Paper*, Karlsruhe Institute of Technology.

Schuster, Philipp, and Marliese Uhri-Homburg. 2013. "The Term Structure of Bond Market Liquidity Conditional on the Economic Environment: An Analysis of Government Guaranteed Bonds."

Working Paper, Karlsruhe Institute of Technology.

Tang, Dragon Yongjun, and Hong Yan. 2012. "Liquidity and Credit Default Swap Spreads." *Working Paper.*

Vayanos, Dimitri. 2004. "Flight to Quality, Flight to Liquidity, and the Pricing of Risk." NBER Working Paper No.10327.

Vayanos, Dimitri, and Jiang Wang. 2013. "Market Liquidity – Theory and Empirical Evidence." In George Constantinides, Milton Harris, and Rene Stulz (ed.): *Handbook of the Economics of Finance* (North Holland, Amsterdam).

Viterbi, Andrew J. 1967. "Error Bounds for Convolutional Codes and Asymptotically Optimum Decoding Algorithm." *IEEE Transactions on Information Theory* 13 (2): 260–9.

Watanabe, Akiko, and Masahiro Watanabe. 2008. "Time-Varying Liquidity Risk and the Cross Section of Stock Returns." *Review of Financial Studies* 21 (6): 2449–86.

Xia, Han. 2014. "Can Investor-Paid Credit Rating Agencies Improve the Information Quality of Issuer-Paid Rating Agencies?" *Journal of Financial Economics* 111 (2): 450–68.

Appendices

Appendix A: Detailed Data Preparation and Filtering Process

In the following, we describe in detail our data preparation and filtering process since all results are very sensitive to small changes in the data preparation process and as a thorough description of how to edit TRACE's transaction data for future academic research does not exist so far to the best of our knowledge. We hope that this overview might serve as a comprehensive guideline about the different steps which are necessary to obtain a useful data set for studying liquidity with the TRACE database so that future researchers may save a considerable amount of time. We are happy to provide our Stata and Matlab codes, as well as our data sets upon request.

Our sample data is obtained from the two databases TRACE and FISD which we both access via the Wharton Research Data Services.

- **Step 1 – Obtaining and filtering transaction data from TRACE:** We start by retrieving all transactions reported to the TRACE database from October 1, 2004 until September 30, 2013. Those transactions are secondary over-the-counter transactions and comprise approximately 90% of the entire trading volume of corporate bonds in the U.S.. We collect transaction information such as trade date, trade price, trade size and the corresponding CU.S.IP identifier and issuing company.

- **Step 2 – Applying the Dick-Nielsen (2009) error filters:** The FINRA regulation requires most of the secondary OTC transactions in the corporate bond market to be recorded within 15 minutes. However, since the reporting is done manually by the involved brokers, Dick-Nielsen (2009) finds that erroneous reporting accounts for approximately 8% of all TRACE transactions and develops three filters in order to detect and correct these deficiencies:

  - Deleting true duplicates: In case of inter-dealer trades, each transaction is transmitted by both dealers and consequently reported twice. The duplicate transactions are identified by having the same intra-day message sequence number which specifically identifies each transaction.
  - Deleting same-day corrections: If the disseminating party realizes an erroneous report within the same day, it can report a corrected transaction notice. Since the original report still shows up in TRACE, we identify them by the message sequence number in the corrected report and delete them. Furthermore, the disseminating party can simply cancel a report by sending a cancellation notice. In that case, both the original and the cancellation record have to be deleted.
  - Deleting reversals: Reversals are corrections that are performed on a later date than the original report. For each reversal, the original report is deleted.

- **Step 3 – Applying price filters on transaction level:** We apply several filters that shall detect wrong reports which have not been corrected or cancelled. Due to the importance of price for our subsequent analyses, we focus on detecting erroneous price entries. Those comprise mistakes such as typos or, by mistake, reporting yields as prices. Consequently, the filters aim at identifying unusual price jumps among transactions. The following filters are based on the filters of Edwards et al. (2007), Han and Zhou (2011) and Friewald et al. (2012).

- Absolute price filter: We exclude all transactions with prices less than $2 and more than $500.
- Intra-day median filter: The filter eliminates any transaction where the price deviates by more than 25% from the daily median.
- Preceding transactions median filter: The filter eliminates any transaction where the price deviates by more than 25% from the median of the five preceding transactions. Those preceding transactions can either fall in the same trading day or in the most recent traded day before.

- **Step 4 – Aggregating transactions to daily data:** As we perform our liquidity measures on daily prices, we aggregate the intra-day transactions for each bond into a single summary observation that consists of the dollar volume-weighted price, the high and low prices and total trading volume of each day. In order to precisely capture the total returns sensitivity to liquidity changes, we transform the clean prices of TRACE into dirty prices according to the following formula:

$$Price_{dirty} = Price_{clean} + Accrued\,Interest_t + \sum_{t=0}^{T} Coupon_t$$

- **Step 5 – Applying price filters on daily level:** Manually examining the dataset, we identify that erroneous price records are likely to appear in clusters (e.g. the broker submits the yield instead of the price for many consecutive transactions before correcting his reporting behavior). Consequently, we redo the median filter on trading day level in order to capture these clusters of false records.

  - Trading day median filter: The filter eliminates any transaction where the price deviates by more than 25% from the median of the five preceding or subsequent days of trading.

- **Step 6 – Merge the TRACE database with credit rating information from FISD:** The Mergent FISD database provides information about credit ratings from Standard & Poor's, Moody's, Fitch Ratings, and Duff and Phelps Rating for almost all U.S. corporate bonds. We assign the credit rating of the most recent date at which a bond's rating was initiated or changed to every trading day of the specific bond, irrespective from which of the four agencies the rating stems from. Furthermore, we translate the official rating categories into a numeric scale from 10 to 0 with 10 being the highest rating (10=AAA, 9.5=Aa1/AA+, 9=Aa2/AA ... 1.5=Caa2/CCC, 1=Caa3/CCC-, 0.5=Ca/CC, 0=C/D).

- **Step 7 – Merge the TRACE database with bond characteristics from FISD:** Besides rating information, the Mergent FISD database also offers a wide range of bond issue-specific characteristics that we merge with our TRACE data set. We exclude convertible, exchangeable, puttable and perpetual bonds, as well as bonds denoted in foreign currencies. However, we retain callable bonds as this feature is very standard for bonds and applies to 56.9% of the bonds in our sample.

- **Step 8 – Construct liquidity measures on daily level:** The filtered and merged data set is used to apply our inter-day liquidity measures on daily price / return information. Detailed information about the specification of the liquidity measures can be found in section 4 (Methodology).

- **Step 9 – Generate monthly returns and liquidity measures:** In order to investigate the effects of liquidity changes on bond returns, we chose a monthly time horizon. Therefore, we exclude all months of a bond with less than 10 traded days since its liquidity measures which are constructed on a daily level might exhibit large biases. We define a month's return by the relative change from the preceding month's latest dirty price to the current month's latest dirty price. Furthermore, we require that these two prices stems from transactions in the last weeks of the respective months. Every month of a bond which does not meet this criterion is disregarded. The daily liquidity measures are aggregated to a monthly basis according to the rationales outlined in section 4 (Methodology).

- **Step 10 – Winsorizing liquidity measures:** An analysis of the distribution of monthly liquidity shows that there are a few very severe outliers. Since all liquidity measures are based on price changes to certain dimensions, corporate events such as bankruptcy, restructuring and takeover announcements result in extremely high illiquidity scores so that we exclude them by winsorizing the 1% of highest illiquidity measures.

Appendix B: Derivation of Extended Corwin-Schultz Liquidity Measure[31]

We assume that the bond price follows a diffusion process and that the spread of $S\%$ is constant over the estimation period of T days. We further assume that the daily high price is buyer-initiated and is grossed up by half of the spread, whereas the daily low price is assumed to be seller-initiated and is discounted by half the spread. Hence, we can write

$$\left[ln\left(\frac{H_t^0}{L_t^0}\right)\right]^2 = \left[ln\left(\frac{H_t^A\,(1+S/2)}{L_t^A\,(1-S/2)}\right)\right]^2, \tag{23}$$

where H_t^0 and L_t^0 denote the observed high and low bond prices for day t, respectively and H_t^A and L_t^A denote the actual high and low bond price for day t, respectively. Rearranging this expression leads to

$$\left[ln\left(\frac{H_t^0}{L_t^0}\right)\right]^2 = \left[ln\left(\frac{H_t^A}{L_t^A}\right)\right]^2 + 2\left[ln\left(\frac{H_t^A}{L_t^A}\right)\right]\left[ln\left(\frac{2+S}{2-S}\right)\right] + \left[ln\left(\frac{2+S}{2-S}\right)\right]^2. \tag{24}$$

Under the assumption that the bond price follows a geometric Brownian motion and is observed continously over a relatively small period of time, Parkinson (1980) and Garman and Klass (1980) show that the expection of the natural logarithm of the the high to low bond ratio is proportional to the variance of the bond:

$$E\left\{\frac{1}{T}\sum_{t=1}^{T}\left[ln\left(\frac{H_t}{L_t}\right)\right]^2\right\} = k_1\sigma_{HL}^2, \tag{25}$$

where H_t and L_t denote the high and low bond prices for day t and $k_1 = 4ln(2)$. In a similar setting, Parkinson (1980) shows that the same general rationale also holds for a non-squared ratio

$$E\left\{\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\left[ln\left(\frac{H_t}{L_t}\right)\right]\right\} = k_2\sigma_{HL}, \tag{26}$$

where H_t and L_t denote the high and low bond prices for day t and $k_2 = \sqrt{\frac{8}{\pi}}$. Taking expectations of (21) and substituting the expressions from (22) and (23) yields to

$$E\left\{\left[ln\left(\frac{H_t^0}{L_t^0}\right)\right]^2\right\} = k_1\sigma_{HL}^2 + 2k_2\sigma_{HL}\left[ln\left(\frac{2+S}{2-S}\right)\right] + \left[ln\left(\frac{2+S}{2-S}\right)\right]^2. \tag{27}$$

The expectation of the sum of $\sum_{j=0}^{T-1} ln\left(\frac{H_{t+j}^0}{L_{t+j}^0}\right)$over T single days is

$$E\left\{\sum_{j=0}^{T-1}\left[ln\left(\frac{H_{t+j}^0}{L_{t+j}^0}\right)\right]^2\right\} = Tk_1\sigma_{HL}^2 + 2Tk_2\sigma_{HL}\left[ln\left(\frac{2+S}{2-S}\right)\right] + \left[ln\left(\frac{2+S}{2-S}\right)\right]^2. \tag{28}$$

For simplification purposes we set

$$\alpha = \left[ln\left(\frac{2+S}{2-S}\right)\right], \tag{29}$$

which allows us to rewrite (25) as

[31] Our derivation follow closely the procedure in Corwin and Schultz (2012).

$$E\left\{\sum_{j=0}^{T-1}\left[ln\left(\frac{H^0_{t+j}}{L^0_{t+j}}\right)\right]^2\right\} = Tk_1\sigma^2_{HL} + 2Tk_2\sigma_{HL}\alpha + T\alpha^2. \tag{30}$$

We further set

$$\beta = E\left\{\sum_{j=0}^{T-1}\left[ln\left(\frac{H^0_{t+j}}{L^0_{t+j}}\right)\right]^2\right\}, \tag{31}$$

and hence we can simplify (27) as follows

$$Tk_1\sigma^2_{HL} + 2Tk_2\sigma_{HL}\alpha + T\alpha^2 - \beta = 0, \tag{32}$$

In order to solve this equation, we need an additional second equation with the same unknowns. For this purpose we square the equation for the high-low ratio from the T-day period

$$\left[ln\left(\frac{H^0_{t,t+T-1}}{L^0_{t,t+T-1}}\right)\right]^2 = \left[ln\left(\frac{H^A_{t,t+T-1}}{L^A_{t,t+T-1}}\right)\right]^2 + 2\left[ln\left(\frac{H^A_{t,t+T-1}}{L^A_{t,t+T-1}}\right)\right]\alpha + \alpha^2, \tag{33}$$

where $H_{t,t+T-1}$ is the high price over the T days $t, t+1, ..., t+T-1$ and $L_{t,t+T-1}$ is the low price over the T days $t, t+1, ..., t+T-1$. Taking expectations and setting

$$\gamma = E\left\{\left[ln\left(\frac{H^0_{t,t+T-1}}{L^0_{t,t+T-1}}\right)\right]^2\right\} \tag{34}$$

and further assuming that

$$E\left[ln\left(\frac{H_{t,t+T-1}}{L_{t,t+T-1}}\right)\right]^2 = E\left[\frac{1}{T}\sum_{t=1}^{T}ln\left(\frac{H_t}{L_t}\right)\right]^2 = k_1\sigma^2_{HL}, \tag{35}$$

and

$$E\left[ln\left(\frac{H_{t,t+T-1}}{L_{t,t+T-1}}\right)\right] = E\left[\frac{1}{\sqrt{T}}\sum_{t=1}^{T}ln\left(\frac{H_t}{L_t}\right)\right] = k_2\sigma_{HL}, \tag{36}$$

we get

$$Tk_1\sigma^2_{HL} + 2\sqrt{T}k_2\sigma_{HL}\alpha + \alpha^2 - \gamma = 0. \tag{37}$$

Solving (29) for α yields to

$$\alpha = -k_2\sigma_{HL} + \sqrt{k_2^2\sigma^2_{HL} - k_1\sigma^2_{HL} + \frac{\beta}{T}}. \tag{38}$$

Rearranging this expression gives us

$$\alpha = -k_2\sigma_{HL} + \sqrt{\sigma^2_{HL}\left(k_2^2 - k_1\right) + \frac{\beta}{T}}. \tag{39}$$

When we further plug (36) into (34), we get

$$Tk_1\sigma^2_{HL} + 2\sqrt{T}k_2\sigma_{HL}\left(-k_2\sigma_{HL} + \sqrt{\sigma^2_{HL}\left(k_2^2 - k_1\right) + \frac{\beta}{T}}\right) + \left(-k_2\sigma_{HL} + \sqrt{\sigma^2_{HL}\left(k_2^2 - k_1\right) + \frac{\beta}{T}}\right)^2 - \gamma = 0. \tag{40}$$

Rearranging this expression yields to

$$\sigma_{HL}^2\left[k_2^2\left(2-2\sqrt{T}\right)+k_1\left(T-1\right)\right]+\sigma_{HL}k_2\left(2\sqrt{T}-2\right)\sqrt{\sigma_{HL}^2\left(k_2^2-k_1\right)+\frac{\beta}{T}}+\frac{\beta}{T}-\gamma=0. \tag{41}$$

This equation can be solved numerically for σ_{HL}. When we further substitute this solution in α, we can calculate the spread S. To simplify the estimation in practice, we ignore Jensen's inequality in (23) and assume the following

$$E\left\{\frac{1}{T}\sum_{t=1}^{T}\left[ln\left(\frac{H_t}{L_t}\right)\right]\right\}=\sqrt{E\left\{\frac{1}{T}\sum_{t=1}^{T}\left[ln\left(\frac{H_t}{L_t}\right)\right]\right\}^2}=\sqrt{k_1\sigma_{HL}^2}=\sqrt{k_1}\sigma_{HL}. \tag{42}$$

This means that $k_2^2=k_1$,which simplifies (36) (now with $\bar{\alpha}$ instead of α) to

$$\bar{\alpha}=-k_2\sigma_{HL}+\sqrt{\frac{\beta}{T}}. \tag{43}$$

Using this simplified $\bar{\alpha}$ in equation (38) leads to

$$\sigma_{HL}^2k_2^2\left(T-2\sqrt{T}+1\right)+\sigma_{HL}k_22\left(\sqrt{T}-1\right)\sqrt{\frac{\beta}{T}}+\frac{\beta}{T}-\gamma=0. \tag{44}$$

We can rewrite this expression as

$$\sigma_{HL}^2k_2^2\left(\sqrt{T}-1\right)^2+\sigma_{HL}k_22\left(\sqrt{T}-1\right)\sqrt{\frac{\beta}{T}}+\frac{\beta}{T}-\gamma=0. \tag{45}$$

Dividing the whole term by $k_2^2\left(\sqrt{T}-1\right)^2$yields to

$$\sigma_{HL}^2+\frac{2\sqrt{\frac{\beta}{T}}}{k_2\left(\sqrt{T}-1\right)}\sigma_{HL}+\frac{\frac{\beta}{T}}{k_2^2\left(\sqrt{T}-1\right)^2}=\frac{\gamma}{k_2^2\left(\sqrt{T}-1\right)^2}. \tag{46}$$

Further simplifiying this expression gives us

$$\left(\sigma_{HL}+\frac{\sqrt{\frac{\beta}{T}}}{k_2\left(\sqrt{T}-1\right)}\right)^2=\frac{\gamma}{k_2^2\left(\sqrt{T}-1\right)^2}. \tag{47}$$

Solving this equation for σ_{HL} yields to

$$\sigma_{HL}=\frac{\sqrt{\gamma}-\sqrt{\frac{\beta}{T}}}{k_2\left(\sqrt{T}-1\right)}. \tag{48}$$

Plugging (44) into (40) gives us

$$\bar{\alpha}=\frac{\sqrt{\frac{\beta}{T}}-\sqrt{\gamma}}{\sqrt{T}-1}+\sqrt{\frac{\beta}{T}}=\frac{\sqrt{\frac{\beta}{T}}-\sqrt{\gamma}+\sqrt{\beta}-\sqrt{\frac{\beta}{T}}}{\sqrt{T}-1}. \tag{49}$$

Further simplifiying yields

$$\bar{\alpha} = \frac{\sqrt{\beta} - \sqrt{\gamma}}{\sqrt{T} - 1}. \tag{50}$$

This $\bar{\alpha}$ can then be used for the generalized high-low spread estimate, which is a transformation of α in (26)

$$S = \frac{2\left(e^{\bar{\alpha}} - 1\right)}{1 + e^{\bar{\alpha}}}. \tag{51}$$

Appendix C: Supplementary Tables and Figures

Table 9: Cumulative Average Default Rates in the U.S. from 1981 – 2005

This table displays the historical cumulative default probability per rating and per time left until maturity. They are based on the time period of 1981 to 2005 for the U.S.. The cumulative default probabilities incorporate both the probability of directly defaulting from the current credit rating, but also the probability to first deteriorate in rating and subsequently default. The information is from the Standard & Poor's Annual 2005 Global Corporate Default Study which can be downloaded under the following link: http://www.gastransmissionnw.com/rate_case_filings/documents/SPGTN18.pdf

	Time horizon						
	1	2	3	4	5	6	7
AAA	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%	0.1%
AA	0.0%	0.0%	0.1%	0.2%	0.3%	0.4%	0.5%
A	0.1%	0.2%	0.3%	0.5%	0.7%	0.9%	1.2%
BBB	0.3%	0.7%	1.2%	1.9%	2.6%	3.4%	4.0%
BB	1.1%	3.2%	5.8%	8.2%	10.3%	12.4%	14.0%
B	5.4%	11.8%	17.3%	21.4%	24.4%	26.7%	28.7%
CCC/C	27.2%	37.0%	42.8%	46.7%	50.1%	51.4%	52.6%
	9	10	11	12	13	14	15
AAA	0.3%	0.4%	0.4%	0.4%	0.4%	0.5%	0.6%
AA	0.7%	0.8%	0.8%	0.9%	1.0%	1.1%	1.1%
A	1.8%	2.1%	2.3%	2.5%	2.7%	2.9%	3.1%
BBB	5.1%	5.7%	6.3%	6.8%	7.2%	7.7%	8.3%
BB	16.8%	17.9%	18.9%	19.6%	20.3%	20.7%	21.3%
B	31.4%	32.7%	33.8%	34.7%	35.8%	36.7%	37.5%
CCC/C	54.8%	55.7%	56.4%	57.2%	57.8%	58.4%	58.4%

Table 10: Portfolio-Sorted Illiquidity Measures

This table shows the average bond portfolio illiquidity for 25 portfolios sorted by liquidity and quality during our sample period from 11/04 – 09/13. Illiquid and liquid are defined as the bond portfolios with the highest and lowest illiquidity, respectively and junk and quality are defined as the the bond portfolios with the lowest and highest rating, respectively. Panel A, Panel B, Panel C, and Panel D display the level of illiquidity based on the Amihud measure, the Pastor-Stambaugh measure, the Roll measure, and the Extended Corwin-Schultz measure, respectively.

Panel A: Amihud Measure

	Illiquid	*2*	*3*	*4*	*Liquid*
Quality	3.20	1.21	0.57	0.26	0.07
4	2.98	1.15	0.59	0.27	0.07
3	2.66	1.11	0.58	0.29	0.10
2	2.90	1.11	0.56	0.28	0.10
Junk	2.90	1.06	0.57	0.26	0.09

Panel B: Pastor-Stambaugh Measure

	Illiquid	*2*	*3*	*4*	*Liquid*
Quality	4.84	1.11	0.29	0.12	0.03
4	3.77	0.76	0.30	0.12	0.03
3	2.37	0.53	0.21	0.09	0.04
2	3.23	0.60	0.23	0.12	0.04
Junk	3.86	0.64	0.24	0.13	0.07

Panel C: Roll Measure

	Illiquid	*2*	*3*	*4*	*Liquid*
Quality	1.80	1.13	0.79	0.56	0.29
4	1.81	1.18	0.82	0.56	0.31
3	1.75	1.14	0.84	0.59	0.33
2	1.97	1.17	0.86	0.59	0.37
Junk	2.08	1.33	0.88	0.64	0.42

Panel D: Extended Corwin-Schultz Measure

	Illiquid	*2*	*3*	*4*	*Liquid*
Quality	1.35	0.98	0.77	0.70	0.63
4	1.38	1.01	0.81	0.75	0.70
3	1.35	1.00	0.83	0.66	0.56
2	1.38	1.05	0.86	0.70	0.66
Junk	1.59	1.28	0.94	0.77	0.60

Table 11: Portfolio-Sorted Return and Rating

This table shows the average monthly bond portfolio returns in excess of the one-month T-bill return on 25 portfolios sorted by liquidity and quality dduring our sample period from 11/04 – 09/13. Illiquid and liquid are defined as the bond portfolios with the highest and lowest illiquidity, respectively and junk and quality are defined as the the bond portfolios with the lowest and highest rating, respectively. Panel A, Panel B, Panel C, and Panel D display the actual returns, the yield-to-maturity, the expected returns, and the rating, respectively.
loss spira

Panel A: Actual Returns

	Illiquid	*2*	*3*	*4*	*Liquid*
Quality	2.52%	2.98%	1.79%	1.41%	1.37%
4	4.21%	2.79%	2.61%	2.43%	2.34%
3	5.17%	4.27%	3.61%	2.99%	2.44%
2	3.39%	4.29%	4.22%	3.40%	2.99%
Junk	6.16%	2.89%	3.56%	3.15%	3.75%

Panel B: Yield-to-Maturity

	Illiquid	*2*	*3*	*4*	*Liquid*
Quality	5.64%	4.40%	4.07%	3.74%	3.01%
4	5.50%	4.47%	4.04%	3.68%	3.22%
3	5.62%	4.81%	4.31%	3.93%	3.49%
2	7.23%	5.54%	5.26%	4.81%	4.41%
Junk	10.25%	10.44%	9.05%	8.41%	7.62%

Panel C: Expected Returns

	Illiquid	*2*	*3*	*4*	*Liquid*
Quality	3.95%	2.85%	2.23%	1.99%	1.68%
4	4.43%	3.13%	2.61%	2.38%	2.25%
3	4.28%	3.45%	3.12%	2.75%	2.48%
2	4.28%	3.57%	3.69%	3.47%	3.33%
Junk	4.83%	3.76%	3.88%	3.74%	3.20%

Panel D: Rating

	Illiquid	*2*	*3*	*4*	*Liquid*
Junk	3.06	3.18	3.07	2.89	3.45
2	5.47	5.47	5.47	5.49	5.25
3	6.60	6.58	6.60	6.61	6.62
4	7.55	7.55	7.55	7.53	7.53
Quality	8.65	8.56	8.55	8.52	8.75

Table 12: Portfolio-Sorted Excess Returns during NBER Recession

This table shows the average monthly bond portfolio returns in excess of the one-month T-bill return of 25 portfolios sorted by liquidity and quality during recession periods identified by the NBER recession indicator for our sample period from 11/04 – 09/13. Illiquid and liquid are defined as the bond portfolios with the highest and lowest illiquidity, respectively, and junk and quality are defined as the the bond portfolios with the lowest and highest rating, respectively. LMI is defined as the difference in the return of the highest liquidity portfolios (regardless of rating) and the return of the lowest liquidity portfolios (regardless of rating). QMJ is defined as the difference in the return of the highest quality portfolio (regardless of liquidity) and the return of the lowest quality portfolio (regardless of liquidity). The statistical significance of the spread is tested via a t-test, where an asterik indicates whether the spread is statistically significant at the 5% significance level.

	Junk	*2*	*3*	*4*	*Quality*	*QMJ*
Illiquid	0.19%	0.54%	0.41%	0.50%	0.40%	0.21%
2	-0.11%	0.76%	0.50%	0.19%	0.53%	0.65%*
3	-0.36%	0.49%	0.55%	0.46%	0.47%	0.83%*
4	-0.11%	0.41%	0.40%	0.40%	0.31%	0.41%
Liquid	0.27%	0.73%	0.61%	0.40%	0.30%	0.03%
LMI	0.08%	0.19%	0.20%	-0.09%	-0.10%	

Table 13: Portfolio-Sorted Excess Returns during Liquidity Stress Regime

This figure shows the average monthly bond portfolio returns in excess of the one-month T-bill return of 25 portfolios sorted by liquidity and quality during the liquidity stress regimes identified by the Viterbi paths. Illiquid and liquid are defined as the bond portfolios with the highest and lowest illiquidity, respectively, and junk and quality are defined as the the bond portfolios with the lowest and highest rating, respectively. LMI is defined as the difference in the return of the highest liquidity portfolios (regardless of rating) and the return of the lowest liquidity portfolios (regardless of rating). QMJ is defined as the difference in the return of the highest quality portfolio (regardless of liquidity) and the return of the lowest quality portfolio (regardless of liquidity). The statistical significance of the spread is tested via a t-test, where an asterik indicates whether the spread is statistically significant at the 5% significance level. Panel A displays the excess returns for the stress periods predicted by the regime-switch based on the Amihud measure and Panel B displays the excess returns for the stress periods predicted by the regime-switch based on the Extended Corwin-Schultz measure.

Panel A: Regime-Switch Amihud Illiquidity Measure

	Junk	*2*	*3*	*4*	*Quality*	*QMJ*
Illiquid	-3.07%	-0.22%	-0.51%	-0.29%	-0.30%	2.77%*
2	-2.52%	-0.15%	-0.49%	-0.41%	0.31%	2.83%*
3	-2.23%	-0.32%	-0.03%	0.16%	0.18%	2.41%*
4	-1.84%	-0.15%	-0.01%	0.18%	-0.02%	1.81%*
Liquid	-0.90%	0.63%	0.36%	0.18%	0.11%	1.01%*
LMI	2.17%*	0.85%*	0.87%*	0.47%*	0.41%	

Panel B: Regime-Switch Extended Corwin-Schultz Illiquidity Measure

	Junk	*2*	*3*	*4*	*Quality*	*QMJ*
Illiquid	-2.40%	-0.53%	-0.39%	-0.57%	-0.03%	2.37%*
2	-1.82%	-0.39%	-0.49%	-0.33%	0.21%	2.03%*
3	-1.68%	-0.42%	-0.16%	0.07%	0.07%	1.74%*
4	-1.65%	-0.44%	-0.13%	0.09%	-0.03%	1.62%*
Liquid	-1.16%	-0.05%	0.21%	0.06%	0.00%	1.17%*
LMI	1.24%*	0.48%	0.60%*	0.63%*	0.04%	

Table 14: Illiquidity Level Shocks for Distressed Times

This table displays the difference in average liquidity level of 25 portfolios sorted by liquidity and quality during the liquidity stress regimes identified by the Markov regime-switch based on the Amihud illiquidity measure and the normal stress periods. Illiquid and liquid are defined as the bond portfolios with the highest and lowest illiquidity, respectively, and junk and quality are defined as the the bond portfolios with the lowest and highest rating, respectively. Panel A displays the average illiquidity shock for the Amihud measure and Panel B displays the the average illiquidity shock for the Extended Corwin-Schultz measure.

Panel A: Amihud Illiquidity Measure

	Junk	*2*	*3*	*4*	*Quality*
Illiquid	2.11	0.81	1.43	1.52	1.03
2	1.02	1.11	0.88	0.74	0.81
3	0.41	0.65	0.46	0.26	0.07
4	0.04	0.15	0.20	0.18	0.04
Liquid	0.15	0.01	0.01	0.05	0.01

Panel B: Extended Corwin-Schultz Illiquidity Measure

	Junk	*2*	*3*	*4*	*Quality*
Illiquid	0.80	0.60	0.88	0.83	0.74
2	0.68	0.75	0.96	1.07	0.97
3	0.70	0.90	0.99	0.94	0.71
4	0.68	0.86	0.92	1.13	1.04
Liquid	1.02	1.72	1.14	1.60	1.06