DISCUSSION PAPER

2016/50

Heterogeneous Liquidity Effects in Corporate
Bond Spreads

Hafner, C. and F. Walders

Heterogeneous Liquidity Effects in Corporate Bond Spreads

Christian M. Hafner^a and Fabian Walders^b

^aUniversité catholique de Louvain, ISBA & CORE, B-1348 Louvain-la-Neuve, Belgium.

Christian.hafner@uclouvain.be

^bUniversity of Bonn, BGSE & Institute for Financial Economics and Statistics, Bonn, Germany.

fwalders@uni-bonn.de

October 20, 2016

Abstract

Among the many sources of risk explaining corporate bond spreads, the role of liquidity is the least well understood. This paper investigates the impact of liquidity risk of unknown functional form on the yield spread over time. Heterogeneity is introduced via a latent group structure explaining differences in nonlinear liquidity effects across groups. A key feature of the model is that it can be estimated from highly unbalanced longitudinal data, allowing us to work with data at minimum levels of temporal aggregation. In an extensive empirical study we apply the suggested method to a large panel of trade data for US corporate bonds. Our procedure identifies nonlinear liquidity effects for a large fraction of the securities. Nonlinearities become more pronounced as bond-idiosyncratic illiquidity increases. The classification clearly distinguishes groups differing e.g. in bond characteristics such as spread levels and trading activity. While groups share similar dynamics of liquidity effects, their magnitudes as well as the interplay between idiosyncratic and market illiquidity are different across groups.

Keywords: Yield Spread, Liquidity Risk, Classification of Regression Functions, Semi-parametric Panel Data Model

JEL Classification: C58, C14, G01, G12

1 Introduction

Yield spreads of corporate bonds reflect multiple sources of risks. While credit risk is typically considered to be the most dominant driver, recent works point at the fundamental importance of liquidity risk.¹ The precise individual contributions of risk factors to the yield spread is of particular interest for investors, bond issuers as well regulatory authorities and central banks. Such information allows to improve risk management for bond portfolios, the timing and design of new bond issues as well as the regulatory assessment of asset riskiness and implementation of monetary policy.

According to [Engle and Lange \(2001\)](#), liquidity of an asset is *the ability to transact at low cost*. In their seminal work, [Amihud and Mendelson \(1986\)](#) introduced the idea that investors demand a premium for holding an illiquid asset. For bonds, this premium is to be viewed as one constituent of the yield spread. Liquidity premiums are typically difficult to extract from data as liquidity is latent. Furthermore, the way how a premium enters the yield spread is unknown and is likely to differ across assets.

The preceding remarks ultimately motivate the research questions addressed in this paper: (i) in which way does liquidity risk affect bond spreads, (ii) how do these effects vary over time, and (iii) what governs the differences of liquidity risk across assets? Using three different liquidity proxies reflecting idiosyncratic and market liquidity, we find that their contribution to the yield spread is in many cases highly nonlinear. This nonlinearity is particularly pronounced in times of low liquidity as it was the case, for example, during the 2008 financial crisis. However, the magnitude and range of the premium varies considerably among classes of bonds: in our data we identify two large classes with moderate effects and three smaller classes with large effects. These classes are distinguished by the spread level, the financial strength of the issuer and the trading activity. Interestingly, dynamic features of liquidity effects in the yield spreads seem to be quite similar.

A number of recent studies has addressed the link between corporate bond liquidity and its contribution to yield spreads, e.g. [Friewald et al. \(2012\)](#), [Dick-Nielsen et al. \(2012\)](#), [Chen et al. \(2007\)](#) and [Houweling et al. \(2005\)](#). The vast majority of these findings are gathered from trade data, which typically are temporally aggregated in order to construct a balanced panel, from which sub-samples are taken of bonds with similar time to maturity and rating class. The coefficients of the liquidity proxies in linear panel data regressions are then estimated for each sub-sample separately. [Dick-Nielsen et al. \(2012\)](#) estimate linear panel regressions for a priori chosen subsamples employing a linear combination of liquidity measures as a new proxy. In particular, their results confirm the presence of liquidity effects on corporate bond spreads, varying over time, time to maturity, and rating classes. Based on similar data, [Friewald et al. \(2012\)](#) take a slightly different approach and estimate a

¹See for example [Longstaff et al. \(2005\)](#)

linear panel model as well as Fama-MacBeth cross sectional regressions using a battery of classical liquidity measures. Effects of rating classes as well as changes in marginal effects of the liquidity measures in times of financial stress are covered using dummy variables. Their main finding confirms the presence of liquidity effects, which are particularly pronounced in times of financial turmoil.

While Friewald et al. (2012) and Dick-Nielsen et al. (2012) aggregate all data to weekly, monthly and quarterly observations, Chen et al. (2007) employ even yearly data in a very similar framework. All studies basically find positive relations between measures of illiquidity and corporate bond spreads, while the magnitude of identified effects varies among different levels of temporal aggregation.

Contributing to this strand of literature, the three novelties of our work can be summarized as follows.

1. Our framework is capable to deal with nonlinear effects of liquidity risk on yield spreads, which are motivated by the varying role of liquidity. For example, a deterioration of a single bond's liquidity will affect the yield spread differently in situations of high or low liquidity in the fixed income market. Moreover, different liquidity proxies reflect different facets of illiquidity, and mixed effects between idiosyncratic and market liquidity of unknown form constitute a potential source of nonlinearity. Nonlinearities may also arise in a feasible regression even if the true liquidity-spread relation is linear, depending on the relation between the true but unobserved liquidity variable and its approximation.
2. We estimate liquidity effects based on *event time models*, thus avoiding high levels of temporal aggregation and loss of information. Despite the high degree of non-synchronous trading, common components of the regression functions can still be estimated using time series *and* cross sectional variation.
3. Our approach allows for substantial heterogeneity of liquidity effects across bonds. We avoid subjective divisions into subsamples by endowing the model with a latent group structure, which can be estimated from the data using standard cluster algorithms. In particular, there is no classification of time varying quantities such as time to maturity or credit ratings, avoiding inconsistencies when classifying bonds based on data over a large time horizon.

The remainder of this paper is structured as follows. Section two introduces the formal model, presents estimation procedures, and addresses inferential questions and corresponding statistical tests. In the third section, liquidity measures and their adaptation to the model framework are discussed. The fourth section offers an application of the outlined procedure to trade data from the US corporate bond market between 2004 and 2012. A last section concludes.

2 An Econometric Assessment of Liquidity Risk

In the following we present our additive panel data model with a latent group structure in event time. For a formal presentation of the framework, index the available bonds by $i \in \{1, \dots, N\} =: \mathbf{N}$. Suppose that for each bond there are observations over time on a bond-specific time-scale. Precisely, for an $i \in \mathbf{N}$ denote the ordered set of corresponding time stamps as $\mathcal{T}_i$, with cardinality $T_i := |\mathcal{T}_i|$. These sets might well vary over the index i as is most likely the case with event times. For notational simplicity we use a generic time index $t \in \mathcal{T}_i$ for a given $i \in \mathbf{N}$. Note that the time index t is to be interpreted as i -specific.

Given these time scales, suppose that for each bond $i \in \mathbf{N}$ a sample of time series observations $\{y_{it}, \mathbf{x}_{it}, z_{it} : t \in \mathcal{T}_i\}$ is available, where y_{it} denotes the spread at time t , z_{it} is an L -vector of control variables, and the vector $\mathbf{x}_{it} \in \mathcal{I} \subset \mathbb{R}^m$ collects m liquidity measures, with values in some set $\mathcal{I}$, independent of the index i .

2.1 Model Formulation and Estimation

We suggest incorporating nonlinear liquidity effects in a regression model according to

$$y_{it} = M_i(\mathbf{x}_{it}) + b_i(t) + \gamma_i^T z_{it} + \epsilon_{it}, \quad (1)$$

where $M_i : \mathcal{I} \rightarrow \mathbb{R}$ and $b_i : \mathcal{T}_i \rightarrow \mathbb{R}$ denote unknown smooth functions. The term $M_i(\mathbf{x}_{it})$ represents the effects of liquidity $\mathbf{x}_{it}$, while $b_i(t)$ captures idiosyncratic local time trends. The latter is introduced to the regression in order to capture (potentially unobserved) risk factors beyond the control variables z_{it} . Regression (1) thus models liquidity risk as contributing to the deviation of the yield spread from a local aggregate risk variable. The parameter vector γ_i contains bond specific marginal effects of the control variables.

For the above model we rely on the following set of standard assumptions. Suppose that for all $i = 1, \dots, N$, $\{(\epsilon_{it}, \mathbf{x}_{it}, z_{it})\}_{t \in \mathbb{Z}}$ forms a strictly stationary and ergodic sequence of random variables. For any i this sequence is as usual indexed over the set of integers for simplicity. The error terms are assumed to satisfy for all $i \in \mathbf{N}$, $\epsilon_{it} \perp\!\!\!\perp (\mathbf{x}_{is}, z_{is})$ for any $t, s \in \mathbb{Z}$. Further we assume in line with [Vogt and Linton \(2015\)](#) that errors are independent across time and bonds, i.e. $\epsilon_{it} \perp\!\!\!\perp \epsilon_{is}$ as well as $\epsilon_{ik} \perp\!\!\!\perp \epsilon_{jl}$ for all $i \neq j$, $i, j \in \mathbf{N}$ and any $t, s, k, l \in \mathbb{Z}$ with $t \neq s$. We further assume that for every i and t , the ϵ_{it} are normally distributed with mean 0 and bond specific variance σ_i^2 . Suppose further that for both $\mathbf{x}_{it}$ and z_{it} moments of sufficiently high order exist and that correlation over time and individuals is sufficiently weak. Note that under this set of assumptions, constants in $M_i(\mathbf{x}_{it})$, $b_i(t)$ and $\gamma_i^T z_{it}$ are certainly not separately identifiable.

As a central feature of (1), we postulate a latent group structure determining heterogeneity of liquidity effects. In the fashion of [Vogt and Linton \(2015\)](#) and [Su et al. \(2014\)](#), we assume that cross sectional units are organized in K groups, denoted $G_1, \dots, G_K$, which form an unknown partition of the index set $\{1, \dots, N\}$. Within a group, the nonparametric regression function is identical among assets, i.e. $\forall k = 1, \dots, K$ it holds that

$$M_i = g_k \quad \forall i \in G_k. \quad (2)$$

The functions $g_k : \mathcal{I} \rightarrow \mathbb{R}$ are the group specific liquidity effects - the quantity of main economic interest. Note that the parameters in γ_i as well as the local trend $b_i(t)$ might vary across bonds, also within a group. Hence, the *only* (a priori) common feature of two assets $i, j \in G_k$ is the liquidity effect g_k . Model (1) can be interpreted not only as a generalization of the classical linear panel regression model with fixed effects, but also of the longitudinal data models presented in [Vogt and Linton \(2015\)](#) and [Su et al. \(2014\)](#). It should be noted that the fully parametric linear models estimated in [Friewald et al. \(2012\)](#) and [Dick-Nielsen et al. \(2012\)](#) are comparable to (1) in the sense that they also possess a group structure, which is, however, determined by the rating class and time to maturity and is hence assumed to be known a priori.

Another important difference between the aforementioned models and (1) is the time dimension. As outlined before, we understand the time index t as labeling event time. The precise meaning of this is introduced in Section 3. Reading (1) as a typical panel data model, it is likely that estimation will suffer from unbalanced observations in practice, because trading is typically highly non-synchronous even at a daily level.

Associated with the above model, there are three estimation problems. In a meaningful order, the first problem is to estimate the unknown components in $E[y_{it}|x_{it}, z_{it}]$ for each $i \in \mathbf{N}$. A second problem is to classify the regression curves, and thus the bonds, into groups according to (2). Based on estimates for the individual conditional mean as well as the groups, the estimation of the group specific liquidity effects g_k constitutes the third problem.

The separation of the first and the third estimation problem allows to estimate a common structure from temporally unbalanced panel data without additional aggregation. The implementation might proceed along the lines of existing work: adapting the idea of [Vogt and Linton \(2015\)](#) to our setup, we suggest the following three step procedure.

First Stage

At the first stage, event time model (1) is estimated for each bond $i \in \mathbf{N}$ separately from corresponding time series data $\{(y_{it}, x_{it}, z_{it}) : t \in \mathcal{T}_i\}$. This is implemented running N times a penalized maximum likelihood approach. The smooth components of the regression

are modeled using splines due to their numerical convenience and flexibility. While the local trend b_i can be represented as a univariate spline, a multivariate spline is required to represent M_i . We decide to approximate the splines by a cubic regression spline for b_i and a thinplate regression spline (see Section 4.1.5 in [Wood \(2006\)](#)) of bond specific order for M_i . The penalty order of the spline approximation for M_i is set to two for all bonds, ensuring that all functions M_i are estimated under the same notion of smoothness. Regarding the choice of the spline basis for M_i , note that multiplicative mixed effects between different measures of illiquidity are allowed explicitly.

Model (1) can be restated as a multiple regression model under an identification constraint for the smooth terms². Corresponding coefficients of the truncated basis expansions of M_i and b_i are collected in the vectors $\lambda_{1,i}$ and $\lambda_{2,i}$.

Roughness penalties to be employed in the objective function are the well known scaled quadratic forms $\nu_{j,i}\lambda_{j,i}^T S_{j,i} \lambda_{j,i}$ for $j = 1, 2$. Here $S_{j,i}$ are square matrices constituted by (known) second derivatives of elements in the corresponding spline basis³ and $\nu_{j,i} > 0$ are tuning parameters. The latter are estimated using generalized cross validation. For notational efficiency collect all but the tuning parameters in the vector $\theta_i := [\lambda_{1,i}^T, \lambda_{2,i}^T, \gamma_i^T, \sigma_i^2]^T$ and denote the unpenalized Gaussian log-likelihood as $\mathcal{L}_i(\theta_i)$. The final objective function is the penalized likelihood given by

$$\mathcal{L}_{i,p}(\theta_i) = \mathcal{L}_i(\theta_i) - 1/2 \sum_{j=1}^2 \nu_{j,i} \lambda_{j,i}^T S_{j,i} \lambda_{j,i}. \quad (3)$$

The estimator $\hat{\theta}_i$ maximizes (3) respecting the imposed identification constraints. Then, the estimator $\hat{M}_i(x)$, $x \in \mathcal{I}$, is simply a linear combination of the basis functions evaluated at x . The weights in this linear combination are the coordinates in $\hat{\theta}_i$ estimating $\lambda_{1,i}$. Denoting the T_i -vector of residuals $\hat{\epsilon}_i$ and the influence matrix P_i , the estimator of the error variance is obtained as $\hat{\sigma}_i^2 = \frac{\hat{\epsilon}_i^T \hat{\epsilon}_i}{T_i - \text{tr}(P_i)}$.

Second Stage

Along the lines of [Vogt and Linton \(2015\)](#), we suggest to employ the estimates $\hat{M}_i$ from the previous stage to classify bonds into groups G_k . For this purpose, a hierarchical cluster algorithm appears particularly suited as it (i) offers a natural notion of a distance between groups, and (ii) is computationally efficient even for a large number of bonds. For any cluster algorithm we need a suitable metric to measure the distance between two regression functions. Since these are smooth, the L^2 metric is an appropriate choice, which for two

²In particular, as in [Wood \(2006\)](#) pp. 163-164 we impose the constraints $\sum_{t \in \mathcal{T}_i} M_i(\mathbf{x}_{it}) = \sum_{t \in \mathcal{T}_i} b_i(t) = 0$.

³See [Wood \(2003\)](#) or the aforementioned section in [Wood \(2006\)](#) for details.

unknown population functions M_i and M_j is defined by

$$\Delta_{i,j} := \left(\int_{\mathcal{I}} (M_i(\mathbf{x}) - M_j(\mathbf{x}))^2 d\mathbf{x} \right)^{1/2}. \quad (4)$$

In the same way, we define $\hat{\Delta}_{i,j}$ to be the L^2 distance between $\hat{M}_i$ and $\hat{M}_j$. Intuitively, given $i, j \in G_k$ these functions should be close to each other and so should be $\hat{\Delta}_{i,j}$ and $\Delta_{i,j} = 0$. Based on this metric, a hierarchical cluster algorithm yields a set of clusters, for a given number of K groups, including distances between clusters.

The statistics literature offers a number of cost criteria allowing to select the number of groups K , see e.g. [Charrad et al. \(2014\)](#) for an overview. We suggest to employ one such cost criterion in order to obtain an estimate $\hat{K}$. Given this estimate, ultimate outcomes of the stage are estimated groups $\hat{G}_1, \dots, \hat{G}_{\hat{K}}$.

Regarding implementation we note that, depending on the number of bonds N and the number of liquidity measures m , this stage might be computationally demanding, because it involves $O(N^2)$ computations of m dimensional integrals.

Third Stage

In a third step the group specific liquidity effects g_k are estimated. Following [Vogt and Linton \(2015\)](#) this is implemented by averaging over estimated curves within groups. Unweighted averaging however might lead to inefficient estimation of g_k , because error variances σ_i^2 are allowed to vary over bonds i , also within groups G_k . We hence suggest to obtain an estimate $\hat{g}_k$ of the group function g_k at $\mathbf{x} \in \mathcal{I}$ by weighted averaging according to

$$\hat{g}_k(\mathbf{x}) := \sum_{i \in \hat{G}_k} w_{ik} \hat{M}_i(\mathbf{x})$$

where weights are defined as

$$w_{ik} := \frac{\hat{\sigma}_i^{-1}}{\sum_{j \in \hat{G}_k} \hat{\sigma}_j^{-1}}.$$

This weighting scheme ideally leads to an efficient estimator of $\hat{g}_k$, given the assumption of uncorrelated errors across i . Note that controlling for potential correlation is infeasible in our case as it would require temporally balanced data.

2.2 Inference

A large fraction of existing studies find significant linear liquidity effects from temporally aggregated data, but the results are specific to the level of aggregation. In our event time study, the type of liquidity effect is *a priori* unclear, which motivates to test the form of the

effect from the data. To do so, we suggest a (conditional) two-step testing procedure. It is designed to conclude whether there is (i) no effect, (ii) a linear or (iii) a nonlinear effect for each bond separately after the first stage regressions. In particular, we test in a first step whether liquidity effects are significant at all. Subsequently it is tested for which bonds the liquidity effect, conditional on being significant, deviates from linearity. Note that the first-mentioned test also allows a *pre-classification*: bonds that do not show any significant liquidity effect will be classified in a separate group, say $\hat{G}_0$. All regression functions $\hat{M}_i$ with $i \in G_0$ are not used in the second and third stages.

Approaching the aforementioned procedure more formally, we consider the following two testing problems, **P1** and **P2**:

P1

$$H_0^{(1)} : M_i(x) = 0 \quad \forall x \in \mathcal{I}$$

versus

$$H_1^{(1)} : \exists x \in \mathcal{I} : M_i(x) \neq 0.$$

For the second testing problem, define for some $\beta \in \mathbb{R}^m$ the event $\mathcal{C}$ via all outcomes satisfying

$$E[y_{it} | \mathbf{x}_{it}, z_{it}] = \mathbf{x}_{it}^T \beta + z_{it}^T \gamma + b_i(t).$$

Testing such a fully parametric specification against a semiparametric specification can formally be obtained according to⁴

P2

$$H_0^{(2)} : \exists \beta \in \mathbb{R}^m : P(\mathcal{C}) = 1$$

versus

$$H_1^{(2)} : P(\mathcal{C}) < 1 \quad \forall \beta \in \mathbb{R}^m.$$

Hypothesis $H_0^{(1)}$ is tested separately for all $i \in \mathbf{N}$. Analogously, the testing problem **P2** is considered separately for different bonds $i \in \mathbf{N} \setminus \hat{G}_0$. For both hypotheses, corresponding F-tests are readily available, because the spline approximation of M_i nests a zero function as well as a linear model.

⁴The null is formulated along the lines of [Fan and Li \(1996\)](#) and [Hsiao et al. \(2007\)](#), while it is certainly not identical. It is a formal variant of the corresponding hypothesis stated in [Hayfield et al. \(2008\)](#).

Denoting the (unscaled) deviances⁵ of the null and alternative models as D_0 and D_1 , the test statistic obtains as

$$F = \frac{(D_0 - D_1)/(df_1 - df_0)}{D_1/(T_i - df_1)}. \quad (5)$$

The quantities df_0 and df_1 are the model’s (effective) degrees of freedom, defined as the trace of the corresponding influence matrix. The asymptotic distribution of F under the null is an $F_{T_i - df_1}^{df_1 - df_0}$ -distribution. In the penalized ML approach that we use to obtain point estimates, the effective degrees of freedom are likely to be non-integer quantities, such that only an approximate F-distribution and corresponding p-values would be available. Thus, we follow the advice in [Wood \(2006\)](#) and ultimately employ estimates obtained under unpenalized ML for testing **P1** and **P2**.

3 Corporate Bond Liquidity

In the previous section we described a modeling approach allowing to identify the effect of an m -vector of liquidity proxies, $\mathbf{x}_{it}$, on the yield spread given a bond specific event time scale. In this section we make the proxies in $\mathbf{x}_{it}$ as well as the time scale precise. In what follows, suppose that transaction prices are available without explicit information on bid and ask quotes.

The literature offers a battery of liquidity proxies, designed to cover a range of different facets of illiquidity. Most common proxies in empirical studies are the bid-ask spread, the number of zero-returns, the [Amihud \(2002\)](#) measure of price impact, and measures of price dispersion as suggested e.g. by [Han and Zhou \(2006\)](#). These quantities are all designed to measure idiosyncratic bond liquidity over a fixed, pre-specified, time horizon.

We focus on three different measures in order to explore the effects of, and the interplay between, idiosyncratic and market liquidity. In particular, we consider the indirect measure of the bid ask spread proposed by [Han and Zhou \(2006\)](#), a market counterpart of it, and the duration between active trading periods. In order to introduce these measures more formally, let us first discuss the considered time scale in more detail, recalling that one of our aims is to estimate liquidity effects at low levels of temporal aggregation. As the liquidity measure of [Han and Zhou \(2006\)](#) is calculated over fixed time intervals, temporal aggregation cannot be avoided completely. We consider intervals of one day for this purpose, which is a plausible compromise between a low level of aggregation (yielding larger T_i s) and exclusion of very thinly traded bonds (yielding smaller N). Thus, we finally define as an *event* (in the time scale sense) a calendar day at which for the considered bond there are sufficiently many transactions available to calculate the liquidity measure.

⁵For the definition, see [Wood \(2006\)](#), p.69 in the case of GLMs and p. 191 in the case of GAMs.

We now describe the Han and Zhou (2006) measure. Given that there are no quotes available, Han and Zhou (2006) suggest an indirect proxy for the bid ask spread based on the price dispersion, arguing that in the absence of news about fundamentals, price dispersion within a day is mainly due to the bid ask spread. At some active trading day t for a bond i , collect the time stamps of transactions in a set $\mathbf{T}_{it}$. Further denote the price reported on transaction $\tau \in \mathbf{T}_{it}$ as $p_{i\tau}$ and the r -quantile of the empirical distribution of $\{p_{i\tau} : \tau \in \mathbf{T}_{it}\}$ as $\mathcal{Q}_{it}(r)$. Then, the proposed quantity is a robust measure of intraday price dispersion, defined as

$$x_{1,it} := \frac{\mathcal{Q}_{it}(0.75) - \mathcal{Q}_{it}(0.25)}{\bar{p}_{i,t}}.$$

In the denominator, $\bar{p}_{i,t}$ denotes the arithmetic mean of transaction prices on day t . We consider t to be an *active* trading day of bond i if $|\mathbf{T}_{it}| \geq 2$.

The above measure approximates idiosyncratic illiquidity. Bonds may also be exposed to *market illiquidity*, and we suggest to proxy market illiquidity by the arithmetic mean of $x_{1,it}$ over all bonds which are traded on this particular day. Collecting these bonds on day t in the index set $\mathbf{N}_t$, our market liquidity proxy is defined as

$$x_{2,it} := \frac{1}{|\mathbf{N}_t|} \sum_{j \in \mathbf{N}_t} x_{1,jt},$$

where the index i is added because in event time the index t has an idiosyncratic interpretation.

In practice, the above illiquidity measure will also be affected by news-based trading. In order to control for such effects, we also consider the duration between days with trading activity. Ideally this quantity allows to identify price dispersion triggered by news based trading, which is typically associated with higher trading activity over larger time horizons. To be precise, $x_{3,it}$ is defined as the number of calendar days between trading days t and $t - 1$ for a given bond i .

Model (1) enables us to discover rich mixed effects between the different liquidity measures. Besides theoretical moment conditions, it was assumed that $\mathbf{x}_{it}$ takes values in $\mathcal{I}$, a subset of $\mathbb{R}^m$, where in our case $m = 3$. This assumption seems to be weak, but for practical purposes it deserves closer attention: A reliable estimation of M_i over all $\mathcal{I}$ in a first stage regression will require the data to cover also the boundaries of the interval sufficiently well. Simply defining $\mathcal{I}$ as set spanned by minima and maxima of x_1, x_2 and x_3 over all bonds and active trading days will usually not lead to such a situation.

We therefore suggest to standardize the liquidity measures. As one can expect many values to be close to zero in applications, we first rescale the measures to obtain a broader distribution,

and second, a standardization to the interval $[0, 1]$ such that $\mathcal{I} = [0, 1]^3$. Concerning the rescaling we define for each bond $i \in \mathbf{N}$ and liquidity measures $j = 1, 2, 3$

$$\tilde{x}_{j,it} := \log \left(\exp(1) + \frac{x_{it}}{\hat{s}_{j,i}^2} \right)$$

where $\hat{s}_{j,i}^2$ is the empirical variance of $x_{j,i1}, \dots, x_{j,i,T_i}$. The standardization of these rescaled measures to the interval $[0, 1]$ is then obtained as

$$\mathbf{x}_{j,it} := \frac{\tilde{x}_{j,it} - \min_{1 \leq t \leq T_i} \tilde{x}_{j,it}}{\max_{1 \leq t \leq T_i} \tilde{x}_{j,it} - \min_{1 \leq t \leq T_i} \tilde{x}_{j,it}}. \quad (6)$$

Hence, every bond also covers boundary regions of $\mathcal{I}$, while the rescaling delivers a more equally spread distribution on this interval.

In what follows, we collect the three liquidity proxies in a vector $\mathbf{x}_{it} = [\mathbf{x}_{1,it}, \mathbf{x}_{2,it}, \mathbf{x}_{3,it}]^T$, serving as argument of M_i .

4 Empirical Evidence: The US Corporate Bond Market

We confront the approach outlined in the previous sections with data for the US corporate bond market. The primary reason is the availability of rich data on corporate bond transactions. As imposed by the US Financial Regulatory Authority, almost all important brokers and dealers in the US are obliged to report every transaction. This data has been used recently e.g. by [Dick-Nielsen et al. \(2012\)](#) or [Friewald et al. \(2012\)](#). In this section, after a description of the data and its processing, we discuss the empirical estimation results for our approach.

4.1 Data

We employ trade data for a large number of US corporate bonds gathered from TRACE (Trade Reporting and Compliance Engine). Furthermore, we employ CUSIP data to obtain bond and issuer characteristics and COMPUSTAT data to obtain firm fundamentals. Spreads are constructed using reference yields from the daily sovereign yield curve, interpolating linearly the yield curve to obtain values at all times to maturity. Information about this yield curve is obtained from the data set provided by [Gurkaynak et al. \(2006\)](#). Data on the Dow Jones Industrial Average index are gathered from Yahoo Finance.

The considered sample covers the time between 2004-10-01 and 2012-12-31. As in [Dick-Nielsen et al. \(2012\)](#), we discard earlier observations because between 2002-07-01 and 2004-09-30 TRACE was not fully implemented and only covered dominantly large investment grade bonds.

In order to deal with errors in the data, we first employ the filter suggested in [Dick-Nielsen \(2009\)](#). Then, inspired by the median filter (see [Friewald et al. \(2012\)](#) and references therein), we cope with potential outliers deleting all trades for which the reported price deviates by more than 10% from its daily median.

We calculate the duration between trading days including also days at which there is only one trade. The duration at some day is defined as the number of calendar days since the last day at which a trade took place. After having calculated this quantity, we delete for each bond days at which the security was traded only once. Moreover, we also discard bonds for which no information was available on issuance, issuer and firm fundamentals.

At any point in time, only bonds with time to maturity between one and thirty years are considered. Furthermore, we remove securities with less than 50 observations. Data points associated with a duration of more than 100 days from the last active trading day are deleted as well.

From the remaining observations, we calculate for each bond a daily spread by averaging over the calculated intraday spreads. The liquidity measures are computed from the intraday observations as outlined in Section 3. Finally, our sample comprises 5755 bonds of 1499 issuers. Thus, the number of assets is smaller than in other studies, but the sample still covers the full range of bond characteristics.

Interest Rate and Equity Market Controls

We employ three control variables capturing level, slope and curvature effects of the risk free yield curve. Each one is calculated from the daily US sovereign yield curve reported in [Gurkaynak et al. \(2006\)](#). Shifts in the level are reflected by the mean yield over all available maturities, while the slope is calculated as the spread between yields at 29 years and 1 year time to maturity. A curvature effect is captured by the so called butterfly spread, defined as the difference between the yield at time to maturity 15 years and the mean from yields corresponding to 29 years and 1 year time to maturity.

As corporate bond yields typically exhibit a dependence on equity dynamics, we also consider daytime returns of the Dow Jones Industrial Average index.

Bond Characteristics and Issuer Information

The bond characteristics used in the analysis are the time to maturity and dummies signaling embedded call options and fixed rates. The considered information on the issuing firms comprises operating income/total assets, long term debt/ total assets, and the S&P quality rank. A summary of the considered variables is given in Table (7) in the appendix.

4.2 Estimating Liquidity Effects

We follow the three-step procedure as described in Section 2 using model (1) in levels. In several related studies, panel regressions are estimated in first (time) differences to account for nonstationarity, while our approach uses local trends that naturally incorporates potential nonstationarity.

In Section 3 we introduced a distribution stabilizing transformation for the liquidity measures, the effect of which can be inferred from Figure 6 in the appendix. Despite the important number of zeros for x_1 , the distributions of x_1 and x_2 are well spread over the interval $[0, 1]$, hence providing sufficient information on situations of illiquidity. For the durations between active trading days, x_3 , the dominance of zeros is much more present, which should be taken into account in the interpretation of estimated liquidity components.

First Stage: Bond Specific Event Time Regressions

For the first stage regressions, we select interest rate and equity market control variables, as well as the bond's time to maturity and three autoregressive terms. As indicated before, we employ a bond specific complexity of the spline approximation for M_i , the rank of which is chosen to depend on the sample size according to $\lfloor 10 \cdot T_i^{2/9} \rfloor$.⁶ For the local trend $b_i(\cdot)$, cubic regression splines are used.

To reduce the impact of outliers, we employ a bond specific definition of an outlier and remove highly influential points identified by a Cook's distance larger than $4/T_i$. After deleting outliers, we relabel event time (and lags) accordingly. Furthermore, the liquidity measures are again transformed according to (6) in order to ensure that minimum and maximum are at zero and one, respectively.

A battery of specification tests is applied to the fitted regressions. First of all, we run Augmented Dickey Fuller (ADF) tests on the residual series. Using a specific lag length of $\lfloor (T_i - 1)^{1/3} \rfloor$ and no trend or constant, the presence of a unit root is rejected for 94 per cent of the assets at the 5 per cent level. Then, the results of residual autocorrelation and normality tests are summarized in Table 1.

	$p > 0.10$	$p > 0.05$	$p > 0.01$
Ljung-Box	0.63	0.70	0.79
Box-Pierce	0.65	0.71	0.80
Jarque-Bera	0.28	0.31	0.35

Table 1: **Residual Diagnostics.** *Share of estimated models for which the tests deliver a p-value as indicated. Lag length for the Box tests is set to ten.*

The Box tests indicate that for a majority of the regressions there is negligible autocorrelation in the residuals. For example, at the 1 percent level this holds for about 80 per cent of the

⁶See the discussion and corresponding reference in Wood (2006) p. 157.

residual series. The results of the Jarque-Bera test indicate deviations from normality for around two thirds of the regressions. Nevertheless the estimators can still be interpreted as pseudo-ML estimators as long as the true distributions are sufficiently close to normality.

	Q(0.10)	Q(0.25)	Q(0.50)	Q(0.75)	Q(0.90)
Adjusted R^2	0.68	0.86	0.94	0.97	0.98
Explained Deviance	0.74	0.89	0.95	0.97	0.98

Table 2: **Adjusted R^2 s and Explained Deviances of first Stage Regressions.** *Table reports quantiles of the empirical distribution of adjusted R^2 s and explained deviances over different bonds.*

Table 2 reports classical goodness-of-fit measures. As 90 per cent of the regressions yield an adjusted R^2 higher than 68 per cent, we regard the quality of the corresponding fits as globally satisfactory, a view supported by the explained deviances. To more clearly isolate the importance of the liquidity effect, we run the tests shown in Section 2 for problems **P1** and **P2**. Table 3 summarizes the results for **P1**.

	$p < 0.10$	$p < 0.05$	$p < 0.01$	$p < 0.001$
F-Test	0.67	0.59	0.46	0.34

Table 3: **Test Results for Problem P1.** *Quantiles of the empirical distributions of p -values across bonds for the F-Test shown in (5) in Section 2.*

For about two thirds of the bonds, the inclusion of a liquidity effect is significant at the 10 per cent level, while for almost half of the bonds it is significant at the 1 per cent level. Whether this effect is linear is tested using the F-test of Section 2. The corresponding results are provided in Table 4.

	$p < 0.10$	$p < 0.05$	$p < 0.01$	$p < 0.001$
F-Test (Full Sample)	0.67	0.60	0.49	0.39
F-Test (Reduced Sample, $p < 0.05$)	0.74	0.68	0.57	0.47
F-Test (Reduced Sample, $p < 0.01$)	0.77	0.72	0.62	0.52
F-Test (Reduced Sample, $p < 0.001$)	0.81	0.76	0.67	0.58

Table 4: **Test Results for Problem P2.** *Quantiles of the empirical distributions of p -values across bonds for the deviance based F-Test in (5) with parametric liquidity contribution (additive, no mixed effects) in the null model. The reduced samples consist of all bonds for which there is a significant liquidity effect at the level indicated in parentheses.*

The test results reported in Table 4 show a striking tendency: for the vast majority of those bonds for which there is a significant liquidity effect, this effect is significantly nonlinear. We also find that with increasing significance of the liquidity effect, we also get more evidence for nonlinearities. Note however that this finding does not imply that $\hat{M}_i$ is nonlinear in all directions as the tested hypothesis is that of *joint linearity*.

It may also be the case that, depending on the values of the remaining components in $\mathbf{x}_i$, one component of liquidity has linear effects in some situations and non-linear ones in others.

Figure 1: **Liquidity Effects along the Main Diagonal.** τ -quantiles, Q_τ , of the empirical distribution of function values among bonds. The functions are evaluated at the main diagonal, e.g. a value of 0.25 on the horizontal axis mean that the function is evaluated at $\mathbf{x}^T = [0.25, 0.25, 0.25]$. Only bonds are considered that have significant liquidity effects at the 1 per cent level according to the F-Test for problem **P1** in the first stage.

This is of particular economic interest: for example, is the impact of idiosyncratic liquidity higher in times of high or low market liquidity? The suggested model (1) allows to answer such questions explicitly from data.

Figure 1 as well as Figures 7 and 8 in the appendix illustrate the contributions of the measures $\mathbf{x}_1$, $\mathbf{x}_2$ and $\mathbf{x}_3$ to the yield spreads in different situations. First, Figure 1 shows the effect along the *main diagonal*, i.e. the evaluation of $\hat{M}_i$ at liquidity measures taking all the same values. Second, 7 shows the effects of idiosyncratic measures $\mathbf{x}_1$ and $\mathbf{x}_3$ in cases of high or low market liquidity. Finally, Figure 8 shows the effect of duration for situations of high or low intraday price dispersion.

The cross-section distributions of the main diagonal values of $\hat{M}_i$ (Figure 1) reveal substantial differences of liquidity effects below and above a threshold of about 0.5. Beyond this value, the distribution is more dispersed and probability mass shifts to the right, meaning that when all liquidity measures jointly increase towards 1, the marginal impact on bond spreads increases as well. Moreover, the dispersion of the distribution is increasing, so that heterogeneity among bonds is stronger in situations of little liquidity. As Figures 7 and 8 suggest, this result is less obvious but still confirmed in situations when only a subset of the measures exceed the threshold, while others remain at fixed values.

Second Stage: Classification of Regression Curves

In the first stage, a number of 2554 bonds were found to have significant liquidity effects at the 1 per cent level. As indicated in Section 2, all other bonds are *pre-classified* to be members of the group $\hat{G}_0$. In order to avoid obtaining groups having only one or two members, we further discard bonds for which $\hat{M}_i$ takes *extreme* values. To be precise, we calculate global minima and maxima of $\hat{M}_i$ and delete those assets having minima below the 0.05 quantile of the minimum distribution and/or maxima above the 0.95 quantile of the maximum distribution. Focusing on the resulting set of 2351 bonds, estimated curves $\hat{M}_i$ are first clustered using Ward's algorithm, based on the dissimilarity measure in (4).⁷

Figure 2: **Radial Dendrogram from Clustering Liquidity Effects.** *The dendrogram displays group and distance structures for estimated liquidity effects $\hat{M}_i$. Classification results are reflected by the color scheme.*

To estimate the number of groups, we employ the *McClain* index as described e.g. in [Charrad et al. \(2014\)](#). We set the lower and upper bounds to be 5 and 15, avoiding the appearance of very small or very large groups. For comparison, note that in [Dick-Nielsen et al. \(2012\)](#) the number of classes resulting from liquidity/time to maturity buckets is determined to be 15 (both, in pre- and post-subprime crisis samples). According to the McClain criterion, we find $\hat{K} = 7$, which appears to be quite robust to the choice of the cost function: The *Silhouette* and *Dunn* indexes as described in [Charrad et al. \(2014\)](#) suggest $\hat{K} = 7$ and $\hat{K} = 6$. The seven clusters constructed by Ward's algorithm are depicted in Figure 2. The algorithm identifies two large classes accounting together for more than half of the assets ($k = 5, 6$).

⁷Regarding implementation, we used the algorithm outlined in [Murtagh and Legendre \(2014\)](#).

The remaining bonds are found to constitute two more groups of medium size ($k = 3, 4$) and three smaller groups ($k=1,2,7$).

Table 8 in the appendix summarizes the bond and firm characteristics within the different groups. Interestingly, the level of spreads seems to be somewhat smaller for the two large groups. This is contrasted by the three small groups revealing in tendency higher levels: in particular the seventh group contributes to this result. The similarity of the two big groups as well as their distinction from the smaller groups also transfers to the time to maturity, the quality rank as well as the trading activity. Bonds in those two groups have larger time to maturity, their firms receive better quality ranks and they are traded more actively. On the other hand, smaller groups seem to collect bonds which are less actively traded and issued by firms with less solid quality rankings. As can be seen from Figure 2, the cluster algorithm seems to confirm this result also with respect to dissimilarity between the different regression functions.

Furthermore, Table (9) reports the distribution of bonds with respect to their respective industry sectors within the corresponding groups. Most groups show a similar distribution: dominant sectors are manufacturing as well as finance and insurance and for most groups, except for the seventh, also utilities.

Third Stage: Estimation of Group specific Liquidity Effects

Given the classification of the second stage, we obtain group specific curves $\hat{g}_k$ as the weighted average of estimated regression functions $\hat{M}_i$ belonging to group k . The shape of the resulting estimates $\hat{g}_1, \dots, \hat{g}_8$ is depicted in Figure (3). For the sake of depicting a three-dimensional nonparametric regression function, we evaluate it at arguments (x_1, x_2, x_3) taking the same value x , say, which is given at the horizontal axis. Small (high) values of x represent small (high) illiquidity.

Except for the second group, the result from the second stage is confirmed: the two large groups share similar shapes. The small groups one and seven reveal much higher levels and higher degrees of nonlinearity for larger illiquidity while remaining relatively similar. It can be concluded that the distinction between groups is mainly due to situations of high illiquidity.

The estimates $\hat{g}_k$ now also allow us to assess differences in liquidity effects between groups on the whole domain $\mathcal{I} = [0, 1]^3$. Central questions of economic relevance are twofold: first, do the effects vary equally strong for different degrees of illiquidity among groups? And second: which combination of liquidity measures maximizes the distinction between groups? The first question can simply be tackled considering the difference between ranges of function values, i.e.

Figure 3: **Estimated Liquidity Effects at the Main Diagonal for different Groups.** *Liquidity effects in different groups evaluated at $x_1 = x_2 = x_3$, which is the value on the horizontal axis. Numbers in parentheses indicate group sizes.*

$$R_k := \max_{x \in \mathcal{I}} \hat{g}_k(x) - \min_{x \in \mathcal{I}} \hat{g}_k(x) \quad (7)$$

for $k = 1, \dots, \hat{K}$.

We consider the second question from two different perspectives: a level and a slope perspective. Regarding the former we suggest to compare functions $\hat{g}_k$ and $\hat{g}_l$ according to the location of their maximum distance, i.e.

$$D_{k,l}^0 := \arg \max_{x \in \mathcal{I}} |\hat{g}_k(x) - \hat{g}_l(x)|. \quad (8)$$

For the slope we can use a similar procedure but now applied to first derivatives. Indicating the gradient with ∇ we suggest to compare slopes of liquidity effects according to

$$D_{k,l}^1 := \arg \max_{x \in \mathcal{I}} \|\nabla \hat{g}_k(x) - \nabla \hat{g}_l(x)\|. \quad (9)$$

The results for R_k as well as locations of maxima and minima are reported in Table 5, the locations of level and slope differences in Table 6.

	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6	Group 7
R_k	3.73	2.46	1.04	1.68	0.46	0.15	2.61
max	(1, 1, 1)	(0, 1, 1)	(1, 1, 1)	(1, 1, 1)	(1, 1, 0)	(1, 1, 0)	(1, 0, 1)
min	(.4, 0, .8)	(1, 0, 1)	(.3, 0, 0)	(.2, .1, 0)	(.7, .2, 1)	(.7, 0, 0)	(.2, .3, 0)

Table 5: **Range of Function Values, Maximizer and Minimizer of $\hat{g}_k$.** *Table reports rounded values.*

The most striking observation is that group functions have, in most cases, optima as well as maximum differences at the boundaries. Moreover, all but function $\hat{g}_2$ have a maximum at which idiosyncratic illiquidity ($\mathbf{x}_1$) is at one. For all but the seventh group, market illiquidity $\mathbf{x}_2$ is close to one at the maximum and close to zero at the minimum. This behavior matches the intuition: large values of $\mathbf{x}_1$ and $\mathbf{x}_2$ mean high illiquidity, which tends to imply large liquidity premia. However, the magnitude of these premia seems to be fundamentally different between groups. The range of liquidity effects over $\mathcal{I}$ is comparably small for the large groups. The opposite is true for the three small groups, in particular for the first group. Between groups 6 and 1, for example, the difference can be quantified with a scaling factor larger than 20. For those bonds which tend to have higher spreads and are less actively traded, the effects of liquidity vary much stronger.

k	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6	Group 7
1		(1, 1, 1)	(1, 1, 1)	(1, 1, 1)	(1, 1, 1)	(1, 1, 1)	(1, 0, 1)
2	(1, 1, 1)		(1, 0, 1)	(1, 0, 1)	(1, 0, 1)	(1, 1, 1)	(1, 0, 1)
3	(.7, 1, 1)	(1, 0, 1)		(1, 1, 1)	(1, 1, 1)	(1, 1, 1)	(1, 0, 1)
4	(1, 1, 1)	(1, 0, 1)	(.2, 1, 1)		(1, 1, 1)	(1, 1, 1)	(1, 0, 1)
5	(1, 1, 1)	(1, 0, 1)	(.1, 1, .9)	(1, 1, 1)		(1, 1, .5)	(1, 0, 1)
6	(1, 1, 1)	(1, 0, 1)	(.2, .2, .7)	(1, 1, 1)	(0, 1, 1)		(1, 0, 1)
7	(1, 0, 1)	(1, 0, 1)	(1, 0, 1)	(1, 0, 1)	(1, 0, 1)	(1, 0, 1)	

Table 6: **Locations of maximum Level Differences $D_{k,l}^0$ (lower Triangular) and Slope Differences $D_{k,l}^1$ (upper Triangular).** *Table reports rounded locations of level and slope differences as in (8) and (9).*

The exact constellations of liquidity measures at the maximum and minimum effects differ considerably between groups. In particular, the role of the trading activity variable seems to be ambiguous. For groups 1, 2 and 5, the minimum effect is associated with values of the trading activity variable close to one. For the remaining groups this variable is close to zero at the minimum. Finally, groups 3 and 4 show an intuitive pattern: at the maximum, all liquidity measures are close to one, at the minimum all are close to zero.

Differences in levels and slopes seem to confirm the intuition gathered from Figure 3: For many groups, differences to other groups are most pronounced for large values of the liquidity proxies. For both, levels and slopes, it is remarkable that groups seem to differ particularly often in cases of high idiosyncratic illiquidity. The slope and magnitude of liquidity contribu-

tions sharply discriminates groups in situations of high intraday price dispersion combined with a large duration between active trading days.

4.3 Time Variation of Liquidity Effects

Our estimation procedure revealed substantial differences of magnitudes, shapes and interpretations of the functions $\hat{g}_k$ between groups. Whether this translates into different function values over time certainly depends on the evolution of liquidity proxies over time.

Figure 4: **Norms of Average Daily Liquidity Measures by Group.** *Time series of $\|\bar{\mathbf{x}}_{kt}\|$ for each group with $\bar{\mathbf{x}}_{kt}$ as in (10), smoothed by the Hodrick-Prescott filter. Vertical dotted red lines are at 2007-04-02 (New Century Financial bankruptcy) and 2008-09-15 (Lehman Brothers bankruptcy).*

We aim to assess these two aspects as follows. For each group $k \in \{1, \dots, \hat{K}\}$ we first calculate an average liquidity measure,

$$\bar{\mathbf{x}}_{kt} := \frac{1}{|\mathbf{N}_{kt}|} \sum_{i \in \mathbf{N}_{kt}} \mathbf{x}_{it} \quad (10)$$

for every day. The set $\mathbf{N}_{kt} \subset \hat{G}_k$ comprises all members of group k for which there is an observation on day t . Figure 4 visualizes the group specific time series of Euclidean norms of $\bar{\mathbf{x}}_{kt}$, smoothed using the Hodrick-Prescott filter. Next, we evaluate the estimated group

functions $\hat{g}_k$ at the corresponding average illiquidity time series. This ultimately leads to a time series $\{\hat{g}_k(\bar{\mathbf{x}}_{kt})\}_t$. All series capture the time span from 2004-10-06 to 2012-12-31.

Figure 5: **Estimated Liquidity Effects.** *Time series of function values $\hat{g}_k(\bar{\mathbf{x}}_{kt})$ for each group, smoothed by the Hodrick-Prescott filter. Vertical dotted red lines are the same as in figure (4).*

Figure (5) allows to visually compare the time series of liquidity effects. The most striking observation is that all series tend to have a very similar behavior up to a scaling factor. This impression is particularly supported by the 2008 financial crisis: for all groups, the liquidity contribution increases quite exactly at Lehman Brothers' bankruptcy. The magnitude of this increase differs among groups and, in particular, is more than a percentage point for the seventh group while it is only around 10 basis points for the fifth group.

Albeit the absolute levels differ among groups, timing and direction of changes are strikingly similar between all groups. This is underlined by the empirical correlation between function values among groups. As can be seen from Figure 9 in the appendix, correlations are between 80 and 99 per cent. In light of this observation it seems as if (i) liquidity effects are similar among bonds up to scale, and (ii) further differ substantially only in times of financial turmoil.

5 Conclusion

In this work we introduced a novel way of estimating liquidity effects in corporate bond yield spreads from an unbalanced panel of securities. The suggested semiparametric longitudinal

data model enables us to avoid high degrees of temporal aggregation as typically found in the related literature. We allow for a latent group structure in the cross section dimension to govern similarities and dissimilarities in the effects of illiquidity. The group structure can be estimated from the data. The actual liquidity effects are allowed to reflect a potentially nonlinear interplay between bond idiosyncratic and market liquidity.

In our empirical study, we estimated liquidity effects as well as groups of assets sharing similar effects from transaction data of about 5700 bonds traded in the US. For the majority of assets, we found significant and significantly nonlinear effects. Based on the similarity of estimated liquidity effects in the L^2 metric, we classified assets into seven groups. Two large groups collected actively traded bonds with moderate spread levels from high quality issuers. The effects of illiquidity on the yield spread were found to be quite small, even in the 2008 financial crisis. This was contrasted by three smaller groups. Such groups contained less actively traded high-yield bonds. At the beginning of the crisis, liquidity contributions jumped of about more than a percentage point.

An analysis of group specific effects showed that, in particular for extreme values of the liquidity proxies, functions differed among groups. The groups to distinguish in that respect were particularly the two large and the three small groups. Although the effects of illiquidity differed considerably in their functional form, function values were highly correlated over time.

References

- Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of financial markets*, 5(1):31–56.
- Amihud, Y. and Mendelson, H. (1986). Asset pricing and the bid-ask spread. *Journal of financial Economics*, 17(2):223–249.
- Charrad, M., Ghazzali, N., Boiteau, V., Niknafs, A., and Charrad, M. M. (2014). Package ‘nbclust’. *Journal of statistical Soft*, 61:1–36.
- Chen, L., Lesmond, D. A., and Wei, J. (2007). Corporate yield spreads and bond liquidity. *The Journal of Finance*, 62(1):119–149.
- Dick-Nielsen, J. (2009). Liquidity biases in trace. *Journal of Fixed Income*, 19(2).
- Dick-Nielsen, J., Feldhütter, P., and Lando, D. (2012). Corporate bond liquidity before and after the onset of the subprime crisis. *Journal of Financial Economics*, 103(3):471–492.
- Engle, R. F. and Lange, J. (2001). Predicting vnet: A model of the dynamics of market depth. *Journal of financial markets*, 4(2):113–142.
- Fan, Y. and Li, Q. (1996). Consistent model specification tests: omitted variables and semiparametric functional forms. *Econometrica: Journal of the econometric society*, pages 865–890.
- Friewald, N., Jankowitsch, R., and Subrahmanyam, M. G. (2012). Illiquidity or credit deterioration: A study of liquidity in the us corporate bond market during financial crises. *Journal of Financial Economics*, 105(1):18–36.
- Gurkaynak, R., Sack, B., and Wrigth, J. (2006). The u.s. treasury yield curve: 1961 to the present. Finance and economics discussion series, 2006-28, Federal Reserve Board.
- Han, S. and Zhou, H. (2006). Nondefault bond spread and market trading liquidity. Technical report, Federal Reserve Board.
- Hayfield, T., Racine, J. S., et al. (2008). Nonparametric econometrics: The np package. *Journal of statistical software*, 27(5):1–32.
- Houweling, P., Mentink, A., and Vorst, T. (2005). Comparing possible proxies of corporate bond liquidity. *Journal of Banking & Finance*, 29(6):1331–1358.
- Hsiao, C., Li, Q., and Racine, J. S. (2007). A consistent model specification test with mixed discrete and continuous data. *Journal of Econometrics*, 140(2):802–826.
- Longstaff, F. A., Mithal, S., and Neis, E. (2005). Corporate yield spreads: Default risk or liquidity? new evidence from the credit default swap market. *The Journal of Finance*, 60(5):2213–2253.
- Murtagh, F. and Legendre, P. (2014). Ward’s hierarchical agglomerative clustering method: Which algorithms implement ward’s criterion? *Journal of Classification*, 31(3):274–295.
- Su, L., Shi, Z., and Phillips, P. (2014). Identifying latent structures in panel data. *Working paper, Dept. of Economics, Yale University*.
- Vogt, M. and Linton, O. B. (2015). Classification of nonparametric regression functions in heterogeneous panels. *Forthcoming in Journal of the Royal Statistical Society*.
- Wood, S. (2006). *Generalized additive models: an introduction with R*. CRC press.
- Wood, S. N. (2003). Thin plate regression splines. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(1):95–114.

Appendix

	Quantiles	Q(0.05)	Q(0.25)	Q(0.50)	Q(0.75)	Q(0.95)
Observations		129	351	640	990	1552
Spread (%)		0.25	0.83	1.64	3.34	9.00
Trades/Day		2.00	3.00	6.00	13.00	41.00
Time to Maturity (Years)		1.63	3.64	6.20	9.60	27.80
Mean Rate Yield Curve (%)		2.06	3.28	3.96	4.57	5.03
Slope Yield Curve (%)		-0.23	0.55	2.62	3.78	4.46
Butterfly Spread Yield Curve (%)		0.07	0.42	1.09	1.70	2.07
DJIA Returns (%)		-1.93	-0.45	0.06	0.56	1.77
Quality Rank of Issuer		1.00	3.00	4.00	5.00	7.00
Op. Inc. / Tot. Ass.		-0.00	0.07	0.10	0.15	0.25
L.T. Debt / Tot. Ass.		0.05	0.17	0.28	0.40	0.75

Table 7: **Summary Statistics.** *Quantiles of the empirical distributions from the sample after cleaning as described in paragraph 4.1. The quality rank was translated to a numeric integer scale: the worst rating is decoded as 1, the best as 8.*

Figure 6: **Distribution of liquidity measures.** *Distributions are calculated from pooled data. For x_2 only values for unique days are shown.*

Figure 7: **Effects of idiosyncratic liquidity in different market situations.** τ -quantiles, Q_τ , of the empirical distribution of function values among bonds. For the solid lines market liquidity was at $\mathbf{x}_2 = 0.9$ (low liquidity) for the dotted lines at $\mathbf{x}_2 = 0.1$ (high liquidity). Only bonds with significant liquidity effects at 0.01 per cent according to the F -Test in the first stage were considered.

Figure 8: **Effects of trading activity with high/low price dispersion.** τ -quantiles, Q_τ , of the empirical distribution of function values among bonds. For the solid lines functions were evaluated at $\mathbf{x}_1 = \mathbf{x}_2 = 0.9$ for the dotted lines at $\mathbf{x}_1 = \mathbf{x}_2 = 0.1$. Only bonds with significant liquidity effects at 0.01 per cent according to the F -Test in the first stage were considered.

	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6	Group 7
Bonds	130	60	299	365	714	726	57
Firms	113	56	216	247	371	372	54
Spread	3.15	3.11	1.46	2.37	1.61	1.1	4.47
Time to Mat.	5.54	5.82	4.76	6.1	7.27	7.33	6.16
Quality Rank	3 (0.81)	3 (0.8)	5 (0.79)	4 (0.74)	5 (0.79)	5 (0.82)	3 (0.86)
Trades/Day	5	5	7	5	7	8	5
Op. Inc./Tot. Ass.	0.09	0.09	0.1	0.1	0.11	0.13	0.08
L.T. Debt/Tots Ass.	0.31	0.34	0.24	0.28	0.23	0.21	0.38
Callable	0.58	0.72	0.66	0.66	0.73	0.76	0.58
Fixed Rate	0.88	0.95	0.93	0.92	0.95	0.97	0.91

Table 8: **Characteristics of Estimated Groups.** *Summary statistics for bonds / firms within the estimated groups. The quality rank was translated to an numeric integer scale as in table (7). The number in parentheses is the fraction of bonds for which the quality rank of the firm was available. Rows 3 to 8 report medians.*

	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6	Group 7
Finance and Insurance	0.21	0.13	0.21	0.17	0.19	0.19	0.11
Manufacturing	0.22	0.33	0.27	0.31	0.29	0.35	0.33
Utilities	0.13	0.13	0.14	0.16	0.13	0.10	0.04
Transportation and Warehousing	0.06	0.03	0.06	0.07	0.05	0.03	0.05
Retail Trade	0.05	0.03	0.05	0.05	0.06	0.08	0.05
Mining- Quarrying- and Oil and Gas Extraction	0.07	0.08	0.06	0.08	0.09	0.06	0.04
Information	0.05	0.03	0.07	0.07	0.09	0.10	0.18
Other	0.21	0.22	0.13	0.09	0.10	0.09	0.21

Table 9: **Industries in estimated groups.** *Distribution of bonds over different industries within a given group. Note that double counting issuers is allowed. For industry classification we use the NAICS key.*

Figure 9: **Correlation of time series of function values $\hat{g}_k(\bar{\mathbf{x}}_{kt})$.** Visualization of the correlation of $\hat{g}_k(\bar{\mathbf{x}}_{kt})$ over time. Entries below the main diagonal are rounded correlations.