SCUOLA
NORMALE
SUPERIORE
PISA

Market Microstructure, Price Impact and Liquidity in Fixed Income Markets

Tesi di Perfezionamento
in
Matematica per la Finanza

presentata da
Michael Schneider

Supervisors

Prof. Fabrizio Lillo
Prof. Loriana Pelizzon

Executive summary

The field of market microstructure connects the workings of financial markets on the micro-level, e.g. of single agents or single orders and often at ultra-fast time scales, to emergent phenomena and behaviours at a more macroscopic level and slower time scales. Classical examples include the price formation process, especially the impact of trading on prices, and the estimation and modelling of volatility and liquidity - often with a special focus on fluctuations and extreme events. Corresponding applications are manifold and diverse, ranging from optimal execution problems faced by financial investors to market design and the detection of fraudulent trading behaviours by regulators. Increasing availability of high-frequency and high-quality data, especially in the wake of electronification of trading and regulatory initiatives for higher transparency, have led the field to flourish in recent decades.

Most of these advances have been concentrated on international equities markets and, after the introduction of mandatory post-trade reporting through TRACE, the U.S. fixed-income market. Considerably less attention has been devoted to fixed-income markets in Europe. This imbalance is not so much reflective of the importance of these markets, but rather due to the availability of reliable data sources. As a matter of fact the market for European fixed-income securities and especially sovereign debt is one of the most important worldwide.

Like in most fixed-income markets trading occurs predominantly over-the-counter, with smaller market shares being taken by retail and interdealer electronic platforms. Until recently studies of this market have been few and - especially from a market microstructure point of view - been based mainly on data from the interdealer platform MTS. However the European sovereign debt crisis, the Quantitative Easing program of the ECB and the environment of low and negative bond yields as well as upcoming regulatory changes have drawn new attention to fixed income markets in Europe.

This thesis aims to shed light on the market microstructure in European fixed income markets from a number of different viewpoints, focusing on interconnections among securities, trading structures and markets. One limitation of such an endeavor, as alluded above, is the availability of representative data due to the existing trading structures. In the first part of this thesis we overcome this by using data from electronic platforms, while in the second part we make use of a regulatory dataset that includes over-the-counter transactions. The remainder of this thesis is organized as follows.

Chapter 1 Here we provide an overview of the main contributions in this thesis.

Chapter 2 This chapter introduces the themes and research questions that we address in this thesis for a general audience. A reader already familiar with the concepts of market microstructure could skip this chapter.

Chapter 3 This chapter presents several stylized facts of trading and quoting activity on the electronic retail platform MOT.

Chapter 4 In this chapter we study limits to cross-asset arbitrage in market impact models, extending the framework of Gatheral (2010) to the multi-dimensional case where trading in one asset has a cross-impact on the price of other assets. From the condition of absence of dynamical arbitrage we derive theoretical limits for the size and form of cross-impact and we test these constraints on cross-impact estimated from trades in sovereign bonds on the platform MOT. While we find significant violations of the no-arbitrage conditions, we show that these are not arbitrageable with simple strategies because of the presence of the bid-ask spread.

Chapter 5 Here we model the cross-asset dynamics of liquidity, considering especially significant deteriorations of liquidity conditions. We propose a peak-over-threshold method to identify abrupt liquidity drops from limit order book data and we model the time-series of these illiquidity events across multiple assets as a multivariate Hawkes process. This allows to quantify both the self-excitation of extreme changes of liquidity in the same asset (illiquidity spirals) and the cross-excitation across different assets (illiquidity spillovers). Applying the method to the MTS sovereign bond market, we find significant evidence for both illiquidity spillovers and spirals.

Chapter 6 Dealers in European sovereign bonds have the option to trade either over-the-counter or on an exchange. In this chapter we study the drivers behind their decision for either venue, using a regulatory dataset of transactions in German federal government debt which we match with the limit order book of the interdealer platform MTS. We find that both factors related to cost and immediacy are major drivers of venue choice. Over-the-counter trades on average occur at prices that are favorable with respect to those offered on the exchange and we investigate also the drivers for this OTC discount.

Chapter 7 The final chapter assesses liquidity of German corporate bonds in comparison to the benchmark U.S. market. We describe a market that consists of a majority of infrequently traded bonds and a fraction of few actively traded bonds. For the latter we compute liquidity metrics and determine their driving factors. In a matched sample analysis we also find the subset of the most actively traded German corporate bonds to be significantly more liquid than comparable U.S. bonds.

Chapters 3 through 7 contain the original contributions of this thesis. Each of them is self-contained and in principle can be read separately.

Contents

List of publications

The following articles are part of this thesis:

- Schneider, M., F. Lillo, and L. Pelizzon (2018). Modelling illiquidity spillovers with Hawkes processes: an application to the sovereign bond market. *Quantitative Finance 18*(2), 283–293[1]
- Schneider, M. and F. Lillo (2018). Cross-impact and no-dynamic-arbitrage. *forthcoming in Quantitative Finance*
- de Roure, C., E. Moench, L. Pelizzon, and M. Schneider (2018). OTC discount. *in preparation*
- Gündüz, Y., G. Ottonello, L. Pelizzon, M. Schneider, and M. G. Subrahmanyam (2018). Lighting up the dark: liquidity in the German corporate bond market. *in preparation*

[1] An earlier version of this article was circulated as Schneider, M., F. Lillo, and L. Pelizzon (2016). How has sovereign bond market liquidity changed?-an illiquidity spillover analysis. *SAFE Working Paper No. 151*.

Acknowledgments

I would like to thank and express my deepest gratitude to the many people who have helped me over the last years.

First and foremost I want to thank my PhD supervisors and coauthors Fabrizio Lillo and Loriana Pelizzon. I am very grateful for the chance to work with you, for your guidance and the opportunities you have been offering to me, as well as the chance to meet you as the kind and splendid persons that you are. From you I have learned to work with precision and dedication to detail, how to narrate the story of a paper, and about the business of research. Thank you for your invaluable help!

Similarly I am deeply indebted to all my coauthors Calebe de Roure, Emanuel Moench, Giorgio Ottonello, Marti G. Subrahmanyam and Yalin Gündüz. Thank you for your commitment, encouragement and support as well as for bearing with me.

SNS Pisa has become a second home for me because of its people. I am very thankful for the guidance of Stefano Marmi, who as a program coordinator has always been also a personal mentor. I value immensely how Giacomo Bormetti has made the QuantFinance group also an inclusive circle of dear friends and colleagues. Thanks to all of you. Most of all, thanks to my colleagues Giulia Livieri and Marcello Rambaldi for helping me from the very beginning and for your deep friendship! A big thank you also goes to the many friendly staff of SNS for all their support.

That also Frankfurt has become a second home for me is thanks to many dear colleagues at SAFE-Goethe University and the Research Centre of Deutsche Bundesbank. I owe special thanks and lots of respect to Mario Bellia for many favors and being available to help at any moment of time.

Again and again I was pleasantly surprised by the feedback, encouragement and warm welcome received at workshops, conferences and seminars as well as emails seemingly out of the blue. A collective thanks to the research community, to all discussants, referees and colleagues for their feedback on the research presented here, and a special thanks to Davide Tomio, Florian Klöck, Jun Uno, Sebastian Vogel and many more. I also acknowledge kind financial support from the Association of Foundations of Banking Origin that benefited the research presented in Chapters 4 and 5.

Going the way from physics to finance would not have been possible without the encouragement, advice and support of Werner Krauth, Jean-Philippe Bouchaud, Marc Potters, Wei-Xing Zhou and Fabrizio Lillo. Thank you for believing in me.

It is impossible to imagine the last years without my family and friends. I especially thank my parents Martina and Wolfgang for their unconditional love and support. Marco, Matte, Ale, and many friends all over the globe, thank you for all the shared moments. Finally, thank you Maria Chiara for being with me.

Chapter 1

Presented research

This thesis presents a selection of studies on the microstructure of European fixed income markets from a number of different viewpoints, focusing on interconnections among securities, trading structures and different markets. While a lot of parallels to e.g. equity markets exist, fixed-income markets are more complex. This is on one hand due to structural reasons, since bonds are traded mostly in opaque over-the-counter markets. On the other hand it is also related to the higher degree of fragmentation in bond markets, where a single issuer might have hundreds of bonds outstanding and the number of trades in a single bond is typically low. The scope of this thesis is to explore the resulting market structure and help the understanding of this under-investigated market. In this chapter we summarize the main contributions presented in Chapters 3 through 7 and relate them to the existing literature.

1.1 Empirical market microstructure

This thesis contributes in describing and exploring for the first time datasets that have not previously been accessed by academic researchers for the scope of market microstructure studies. Careful preparation of such datasets and the understanding of market rules are elementary towards robust results and ultimately serve the validation of datasets and findings. Other studies in this direction are e.g. Fleming (1997) for the U.S. Treasury market while Dick-Nielsen (2009, 2014) have helped making the U.S. corporate bond TRACE data accessible for a wide range of researchers.

Our contribution I. Chapters 3 and 4 are the first microstructure studies on the retail bond platform MOT. Especially Chapter 3 contributes in describing the market rules of MOT and the available data, validating the dataset and exploring stylized facts of trading and quoting activity. We describe intraday patterns of liquidity, the autocorrelation of order signs and the shape of the limit order book, also in the context of tick size changes.

Our contribution II. Along the same line the work presented in Chapters 6 and 7 is the first to employ a regulatory dataset of transactions by German banks, including over-the-counter trades, for market microstructure studies. Chapter 7 describes the regulatory origin and scope of the data and recommends procedures for cleaning the raw dataset. On this base Chapter 6 contributes further in providing a first study of the cash market for German sovereign bonds (*Bunds*), while Chapter 7 does so for German corporate bonds.

The main contribution of Chapter 7 lies in studying the liquidity of the opaque German corporate bond market and comparing it with the benchmark U.S. market where there is post-trade transparency through the TRACE database. Therefore our study also provides a unique cross-sectional view on the impact of transparency on over-the-counter markets. While the market for U.S. corporate bonds is far more active, we find that the most liquid and actively traded German bonds are actually more liquid, in terms of transaction costs, than comparable U.S. bonds. We posit that this is related to a 'crowding' effect where trading activity and liquidity is concentrated in a small set of liquid bonds.

Let us stress that the studies in Chapters 6 and 7 are also particularly relevant in view of the MiFID II regulations in force since January 2018. First, as a guidance for data preparation, as MiFID II post-trade transaction reports will be structured similarly. And second, as a reference point for future studies on the impact of MiFID II regulations, where we provide the 'before' picture.

1.2 Price impact

Price impact, i.e. the impact of order flow on price dynamics, is a central theme of market microstructure and has received increased attention from researchers and the financial industry in recent years (Bouchaud et al., 2008). This is, from a fundamental point of view, since market impact is essential towards understanding the price formation process, while from an applied point of view it is crucial for optimal execution and transaction cost analysis. Despite its importance and the interest, market impact remains not yet fully understood and its modeling continues to be a challenge for academics and practitioners alike.

Market impact refers to several similar yet different notions and a careful distinction is necessary. First, there is the impact of the aggregated order flow in a certain period (e.g. the averaged order sign of trades in a 1 minute interval). A second notion refers to the aggregate impact of a large *meta-order* that is executed incrementally. In this thesis we refer to a third notion of impact, that of a single trade (i.e. not aggregating into time intervals and regardless of whether it is part of a meta-order).

The majority of research on market impact has naturally focused on *self-impact*, i.e. the impact of trading an asset on the price of the same asset. *Cross-impact*, i.e. the impact of trading an asset on the price process of other assets, has only very recently become the subject of empirical studies (Benzaquen et al., 2017; Wang et al., 2016a,b).

Appropriate models of market impact are an indispensable ingredient for transaction cost analysis and optimal execution problems. Crucially such models should not only be realistic representations of observed impact, but also free of opportunities for price manipulation and arbitrage. This insight has lead to a rich literature already for the single-asset case with self-impact (e.g. Huberman and Stanzl (2004); Gatheral (2010); Alfonsi et al. (2012); Gatheral and Schied (2013a); Curato et al. (2016, 2017)). As optimal execution is taking into account also cross-impact (e.g. Schöneborn (2016); Tsoukalas et al. (2017); Mastromatteo et al. (2017)) it is equally important to understand the limitations that arise in a multi-asset framework with cross-impact.

Our contribution I. The work presented in Chapter 4 contributes both to the theoretical and the empirical literature on cross-impact. In the theoretical part of the chapter we derive necessary conditions for the absence of dynamic arbitrage in the sense of Huberman and Stanzl (2004). While Alfonsi et al. (2016) show that absence of arbitrage is equivalent to the decay kernel being a positive definite matrix-valued function, this is difficult to verify in practice when a decay kernel is obtained coordinate-wise from estimations. We contribute by establishing some easily verifiable conditions that are necessary for absence of arbitrage.

In order to do so we expand the framework of Gatheral (2010). Our model for the price S_t^i of asset i at time t reads

$$S_t^i = S_0^i + \sum_j \int_0^t f^{ij}(\dot{x}_s^j) G^{ij}(t-s)\mathrm{d}s + \int_0^t \sigma^i \mathrm{d}Z_s^i \tag{1.1}$$

where f^{ij} and G^{ij} capture the impact of order size and decay with time respectively from trades in other assets j, σ^i is the volatility and $\boldsymbol{Z}_s$ a correlated noise process (e.g. a multivariate Wiener process). For a set of simple round-trip strategies we calculate the cost according to the model and absence of arbitrage requires that such costs are non-negative. Assuming further that the decay kernel $\boldsymbol{G}(\tau)$ is bounded, non-increasing and continuous around $\tau = 0$, we arrive at the following necessary conditions:

1. **Odd market impact:** $\boldsymbol{f}$ must be an odd function of the trading rate, i.e.

$$f^{ij}(v) = -f^{ij}(-v) \qquad \forall\, i, j\,. \tag{1.2}$$

2. **Symmetric cross-impact:** Cross-impact must be symmetric from $i \leftrightarrow j$ in the sense that

$$v_i f^{ij}(v_j) = v_j f^{ji}(v_i) \quad \forall\, i, j\,. \tag{1.3}$$

3. **Linear cross-impact:** Cross-impact must be linearly dependent on order size, i.e. we require that

$$f^{ij}(v) = \eta^{ij} v \quad \forall\, i, j \tag{1.4}$$

   with η^{ij} constant.

Furthermore we estimate self- and cross-impact among sovereign bonds traded on the electronic platform MOT. To the best of our knowledge we are the first to estimate cross-impact at the level of single trades instead of aggregated order flow, where in order to do so we adapt the multi-asset version (Benzaquen et al., 2017) of the transient impact model by Bouchaud et al. (2004). Analyzing the impact of single trades allows us to explore the origins of cross-impact. Panel (a) of Figure 1.1 shows that there are contributions from contemporaneous order flow as well as quote revisions. Panel (b) illustrates that cross-impact decays similarly to self-impact and is weaker by roughly an order of magnitude.[1]

Finally we also examine whether the symmetry condition is empirically verified. While we do find significant violations thereof, we show that these are not arbitrageable due to the presence of the bid-ask spread and an upper limit on the speed of trading.

Our contribution II. Appendix 6.A of Chapter 6 also relates to price impact. There we study cross-venue response based on a dataset that includes trades from both the over-the-counter (OTC) and the exchange segment of the German sovereign bond (Bunds) market. We do so by comparing the response of mid-quotes on the exchange to OTC trades (that are only observed by the involved counterparties) to the response of exchange trades. We find evidence that response to OTC trades is weaker than for exchange trades, even when taking into account potential mis-classifications of the order sign, and that price changes due to OTC trades occur slowly, with a timescale of one hour or longer.

[1] In Chapter 4 we provide a detailed discussion why observed market impact is non-linear (conditioning of order sizes on available volumes), the initial increase of the estimated decay kernel and the estimation procedure for the decay kernel.

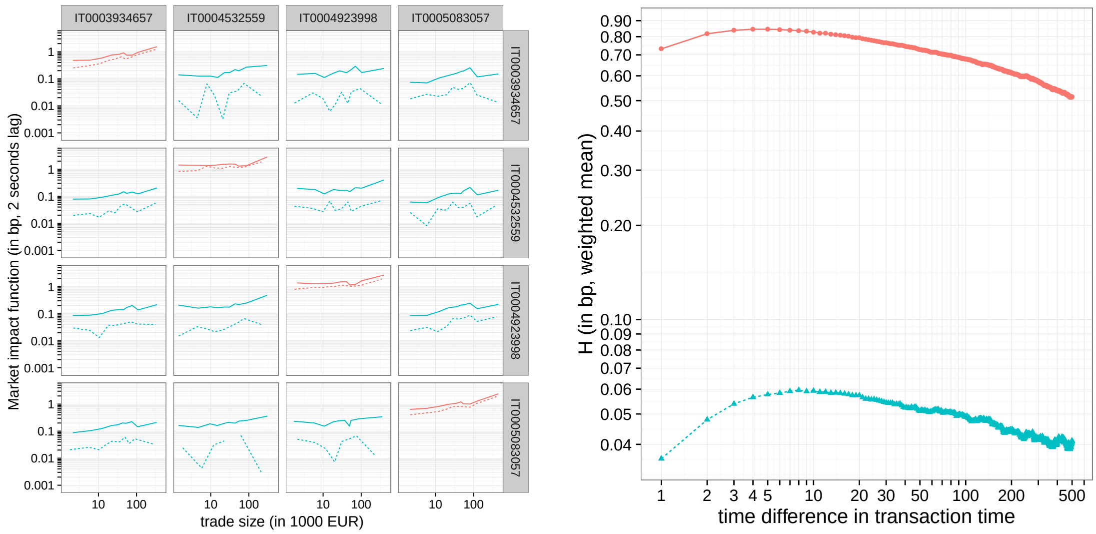

(a) Market impact function **(b)** Decay kernel

Figure 1.1: Estimated shape of self- and cross-impact: Self-impact is depicted in red and cross-impact in blue. Panel (a) shows the average self- and cross-impact function among four bonds where impact is from the column on the row. Solid lines show the market impact function based on all trades, dotted lines show market impact based on isolated trades, i.e. when there was no other transaction from 3 seconds before to 2 seconds after the triggering trade. Panel (b) depicts the average decay kernel for self- and cross-impact among all bonds in our sample. For self- (cross-) impact we show the mean over all bonds (pairings) weighted by the number of transactions in the triggering bond.

1.3 Illiquidity dynamics

Market liquidity is crucial for the functioning of financial markets. At the same time its absence, i.e. illiquidity, has been identified as an important channel for financial contagion (see e.g. Kyle and Xiong (2001); Garleanu and Pedersen (2007); Cespa and Foucault (2014)) and concerns have been raised especially about tail events, i.e. liquidity crashes (Brunnermeier and Pedersen, 2009; Huang and Wang, 2009; Easley et al., 2011). While there is a vast empirical literature on commonality in liquidity pioneered by Chordia et al. (2000), documenting the correlation of liquidity across assets, markets and asset classes (Chordia et al., 2005; Brockman et al., 2009; Karolyi et al., 2012), this literature mostly focuses on the average correlation in liquidity, including therefore normal and tail conditions.

Our contribution. The work presented in Chapter 5 focuses instead on tail events of market liquidity. We construct a liquidity factor from limit order book data and we identify *illiquidity shocks* by adapting to factor increments the peak-over-threshold method which is well-established in the context of price dynamics and exceedances of the Value-at-Risk (VaR) threshold (Chavez-Demoulin et al., 2005; Embrechts et al., 2011; Chavez-Demoulin and McGill, 2012). The innovation of our approach is to abstract from a discrete time grid and use the continuous time-resolution of the underlying limit order book data.

This makes the illiquidity event arrival process particularly suited for modeling with Hawkes processes (Hawkes, 1971a,b). Hawkes processes are a class of self-exciting point processes now widely used in Finance (Bacry et al., 2015) to describe discontinuous processes, such as price jumps or limit order book events.

Here we apply the illiquidity event detection method to the limit order book of MTS, the

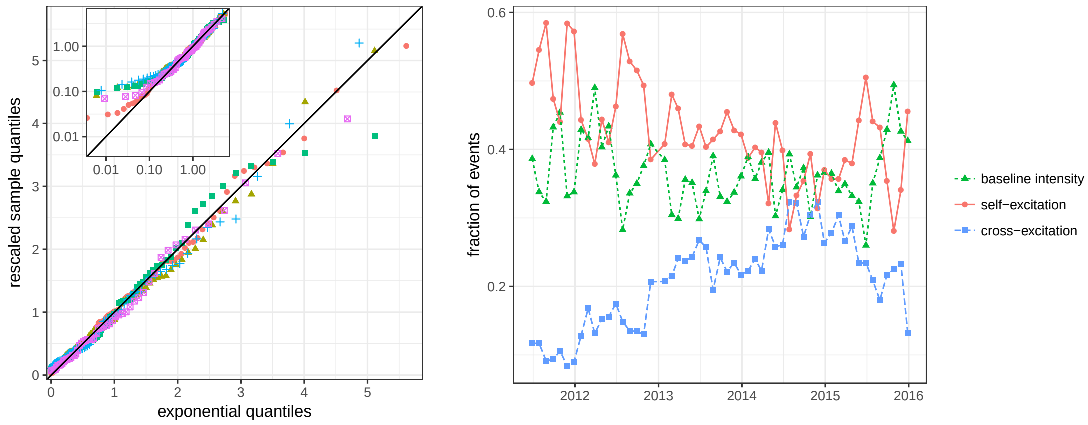

(a) QQ-plot of inter-event durations **(b)** Fraction of events due to illiquidity spirals and spillovers

Figure 1.2: Modeling of illiquidity events with Hawkes processes: Panel (a) shows an example quantile-quantile plot of sample inter-event durations rescaled by the estimated intensity and exponential quantiles. The inset is in logarithmic scale on both axis. Panel (b) shows the fraction of illiquidity events attributed to random arrivals (baseline intensity), illiquidity spirals (self-excitation) and illiquidity spillovers (cross-excitation).

leading interdealer exchange for European sovereign bonds, and we parametrically model their arrival process as a Hawkes process that allows for both self- and cross-excitation. In this framework we can identify the fraction of such illiquidity shocks that is either due to a constant arrival intensity of events, self-excitation, or cross-excitation. We associate self-excitation of illiquidity events with illiquidity spirals and cross-excitation across assets with illiquidity spillovers.

Panel (a) of Figure 1.2 shows the quantile-quantile plot of inter-event durations between illiquidity events rescaled by the estimated arrival intensity and exponential quantiles, illustrating that our approach is indeed well-suited to the modeling of these shocks. In Panel (b) we demonstrate a strong presence of fast illiquidity spirals and illiquidity spillovers throughout the investigated period from 2011 to 2016. We also observe that spillovers occur at the time-scale of a few seconds and significantly faster for bonds with similar maturities.

To our knowledge we are the first to use Hawkes processes to investigate extreme illiquidity events and to document intraday contagion of illiquidity shocks for bonds. As the event detection approach described in this chapter can also be applied to other limit order book markets, potential studies of illiquidity contagion across markets are a natural extension of this work.

1.4 Venue choice

Equities are typically traded on exchanges that are organized as electronic limit order book markets. However many exchanges also have so-called upstairs markets that effectively resemble over-the-counter (OTC) structures. This has lead to a rich theoretical literature on the motivations of traders to be active in one segment or the other (Seppi, 1990; Grossman, 1992) as well as vast empirical literature exploring the *hybrid market* structures and testing the theoretical predictions (cf. e.g. Smith et al. (2001); Bessembinder and Venkataraman (2004); Carollo et al. (2012)).

Regulatory initiatives to shift trading in fixed income instruments and especially bonds from OTC markets towards electronic platforms have created a new interest for hybrid markets (Lee and Wang, 2017; Vogel, 2017). Related studies on bond markets are few, with the notable exception of Barclay et al. (2006), who study the choice between electronic and voice brokerage for U.S. Treasuries that go off-the-run, and Hendershott and Madhavan (2015), who analyze a request-for-quote platform for U.S. corporate bonds.

Our contribution. Chapter 6 contributes by empirically studying a hybrid market with a dominant OTC segment and an exchange that is organized as a fully electronic limit order book. We investigate the German sovereign bond market, using a regulatory dataset of transactions matched with the full limit order book of the leading interdealer exchange MTS. First, we contribute to the literature on venue choice by considering two determinants that, to the best of our knowledge, have not been considered so far. These are immediacy and transparency. The exchange provides an outside option to the OTC search process and, crucially, immediacy. Indeed we find exchange trading more likely in market situations where the demand for immediacy is higher. However we also find evidence that in some cases traders forego profits to benefit from the opacity of OTC trading and avoid the post-trade transparency inherent to the exchange.

Second, we compare transaction costs across the two venues to study cost differences and price discrimination. We do so by calculating the hypothetical price that an OTC trade would have cost on the exchange and consider the difference to the actual trade price. OTC trades provide an average discount of 1.5 basis points to trading on the exchange, which means that OTC round-trip costs are on average 50% lower than on the MTS market. Trades by dealers with access to MTS are on average 0.9 basis points cheaper in the OTC market than those by non-MTS dealers, whereas the impact of immediacy on OTC discounts depends on trade size.

Chapter 2

Introduction

This chapter serves as an overview over the field of market microstructure and aims to be accessible to a broader audience. The seminal textbook on the field O'Hara (1995) defines market microstructure as:

> Market microstructure is the study of the process and outcomes of exchanging assets under explicit trading rules. While much of economics abstracts from the mechanics of trading, the microstructure literature analyzes how specific trading mechanisms affect the price formation process.

Assets in the above definition refers to the goods that are traded, which could be company shares, bitcoin (Donier and Bonart, 2015; Koutmos, 2018) or cattle (Frank and Garcia, 2011). It is customary to distinguish between *asset classes*. The most standard asset classes include *equities* (company shares, also called *stocks*), *fixed income* (debt titles or any contract that obliges the issuer to make scheduled payments, e.g. corporate or sovereign bonds), *foreign exchange* (foreign currencies) and *commodities* (physical goods, e.g. gold, oil, lifestock or electricity). In this thesis we focus on fixed income markets and especially bond markets, however we will often make reference or compare to other asset classes, especially in this chapter.

As market microstructure comprises many diverse subfields, uniting researchers from various disciplines including economics, physics and data science as well as financial practitioners and regulators, a comprehensive overview is beyond the scope of this chapter. Instead we will in the following provide a comprehensible introduction to the topics that are most relevant to this thesis. Section 2.1 illustrates what "explicit trading rules" and "trading mechanisms" in the above definition refer to by presenting examples of financial market structures. Section 2.2 introduces the concept of *liquidity*, which can roughly be transcribed as capturing 'ease of trading.' A particular aspect of this is *price impact*, i.e. how trading activity influences prices, which we highlight in Section 2.3.

2.1 Market structures

Financial markets can be organized in a number of different ways and distinguish themselves in several characteristics. As no two venues or no two asset classes (or even assets) are exactly alike we describe two stylized market structures that are most relevant in the context of this thesis, instead of describing specific markets. That is we consider a market that is organized as an exchange as well as an over-the-counter market:

1. **Exchange:** Strictly speaking an exchange refers to any venue (which could be a physical location or a digital infrastructure) that allows the meeting of traders and

the exchange of assets. Here we refer as *exchange* to such a venue that is organized as an electronic *continuous double-sided auction* market. A double-auction market is built around *orders*, which are binding expressions of an interest to trade, and collected and organized in the *limit order book* (LOB). An order contains information on which asset one wants to trade, whether one wants to buy or sell, at which price, and how much of it. We distinguish between the following *order types*:

- *Market orders* are executed immediately.
- *Limit orders* are collected in the LOB for potential later execution.
- *Cancellations* (or *modifications*) cancel (or modify) existing limit orders. Market orders cannot be cancelled or modified as they execute immediately.

Limit orders are collected in the limit order book until they are executed (i.e. result in a trade) or cancelled. They are organized on two sides according to *price-time priority.* Limit orders for the intention to buy (sell) are grouped on the *bid* (*ask*) side. When a buy (sell) market order arrives it is matched with the limit order(s) with the highest priority on the ask (bid) side. The highest priority is assigned to the limit orders offering the best price (i.e. price priority). In case of a tie (equal prices) the limit order that arrived earlier has priority over later limit orders at the same price (i.e. time priority).

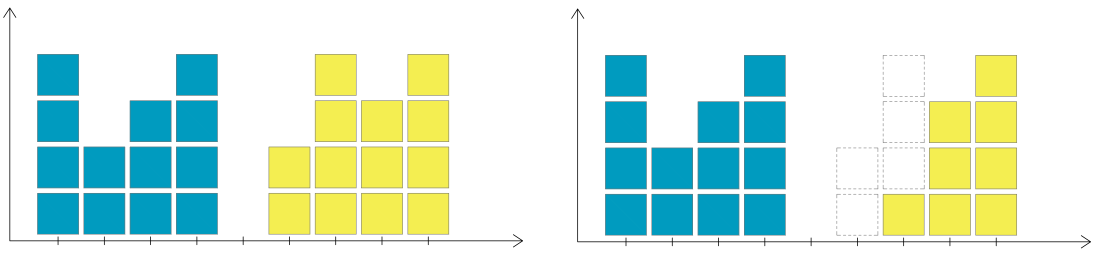

(a) LOB before a trade.   **(b)** LOB after a buy market order.

Figure 2.1: Illustration of a stylized limit order book: Prices are increasing along the x-axis and the y-axis depicts volume, where each block corresponds to one unit of the asset. Limit orders on the bid (ask) side are depicted in blue (yellow). Panel (b) shows the effect of a buy market order on the limit order book from Panel (a): the market order has a size of 5 units and takes out 2 units from the former best ask level and 3 units from the level behind.

Figure 2.1 illustrates these market rules in an example. Panel 2.1a (Panel 2.1b) shows a stylized limit order book right before (after) the arrival of a buy market order for a size of 5 units. The buy market order is executed against the limit orders on the sell side. At the level with price priority there are limit orders for two units of the asset. The remaining 3 units of the market order are executed against limit orders with the time priority at the next level. This implies that the first two units were bought at a cheaper price than the next 3 units, an effect known as *walking up the book* in market jargon.

2. **Over-the-counter:** Trading over-the-counter (OTC) is essentially trading in bilateral negotiations. Someone wanting to e.g. buy an asset starts by contacting potential *counterparties* (i.e. sellers in this case), either by phone or through electronic systems. The trader wanting to buy an asset expresses her interest to do so and for which size (*request-for-quote*, RFQ). The contacted trader in response quotes a price at which she is willing to sell and the potential buyer has the option to accept or not. In

some cases big banks provide *indicative* quotes (i.e. non-binding prices, e.g. through Bloomberg). The buyer might have to contact multiple potential counterparties in order to find a counterparty willing to trade or she could do so trying to get a better price. This is termed the *search process*.

This is only a small subset of possible market structures. Other platforms for example offer to send RFQs to multiple dealers at the same time (a one-sided auction effectively). Some so-called *dark pools* instead take orders and match buy and sell orders only based on volume, while the price is determined from an external reference, e.g. a double auction market. In the following, we compare the exchange and the over-the-counter structures along several dimensions related to market microstructure:

- **Pre-trade transparency:** *Pre-trade transparency* is given when any trader knows *before* initiating a trade the conditions thereof. The exchange offers this, since contracts are standardized and any trader knows the price she would achieve from observing the limit order book. This is not the case in an OTC market. There a trader needs to send a RFQ in order to learn about the conditions she could trade at.

- **Price discrimination:** Whereas all traders receive the same price in the exchange market, the OTC market features *price discrimination*. That is two traders directing the same RFQ to a counterparty might receive different quotes in return.

- **Post-trade transparency:** Trading on an exchange is typically *post-trade transparent*, that is a trade is observed by all market participants after it has happened (but typically not who was involved in the trade). The level of post-trade transparency in OTC markets depends on the rules set by the regulator. A priori only the involved parties know of the trade. However in some asset classes and countries reporting is mandatory. For example data on U.S. corporate bond transactions is collected in the TRACE database and made accessible to researchers.

- **Immediacy:** A market order on the exchange is immediate. The search process in OTC markets can be time-consuming depending on the asset and market conditions.

A basic familiarity with these characteristics is crucial towards understanding the topics presented in the previous section and in the following chapters. However, already this level opens up several exciting research questions:

1. **Market design:** both regulators and exchanges have an interest to set the rules of a market in a way that is 'optimal'. Optimal in the sense of the provider of the exchange might be so that trading activity is maximized, whereas optimal in the sense of the regulator might mean that trading is fair and stable. Parameters of exchange trading that can be altered relatively easy are for example the minimum size of a trade or the *tick size*, i.e. the minimum allowed price change. Several studies examine the effects of tick size changes in the light of various market microstructure aspects and in Section 3.6.2 we investigate the impact of a tick size reduction (i.e. a transition to a finer price grid) on the shape of the limit order book.

2. **Market choice:** So far we have implied that one asset is traded in only one market, e.g. only on the exchange. However for most assets multiple diversely organized platforms and OTC structures compete for order flow and market shares. In Chapter 6 we study the special case of a *hybrid market* where traders have the choice between an exchange venue and an over-the-counter segment to trade the same asset. The

research question is twofold: What determines the choice of trading venue? What determines the cost difference between the two options?

3. **Post-trade transparency:** How does regulator-imposed post-trade transparency impact financial markets? Chapter 7 adds to this open discussion by comparing liquidity in a market with post-trade transparency to a market without this characteristic.

2.2 Market liquidity

Financial market liquidity captures the 'ease' of trading an asset.[1] A security is considered *liquid* if one is able to trade even large quantities of it quickly and at a low cost without affecting prices. On the other hand an asset is *illiquid* if trading is expensive, if one is only able to trade small orders, trading is slow or if trading has a considerable impact on prices. While this is still fairly intuitive, a more precise definition and the measurement of liquidity is far from being straightforward. Already O'Hara (1995) remarks that "liquidity, like pornography, is easily recognized but not so easily defined." This is in part since liquidity has multiple dimensions and no measure is able to capture all aspects of liquidity. In this section we present different approaches to measuring liquidity, acknowledging its multi-dimensionality.[2]

2.2.1 Round-trip cost

A *round-trip trade* is e.g. buying an amount of an asset and subsequently selling the same amount so that the final position in the asset is equal to the initial position of the trader. The cost associated with such a trade is the *round-trip cost* and captures the transaction cost aspect of liquidity.

The round-trip cost for small volumes is measured fairly easily in limit order book markets through the *bid-ask spread*. There the bid-ask spread is defined as the price difference between the best prices on the ask and bid side of the order book :

$$bid\text{-}ask\ spread = price_{best\ ask} - price_{best\ bid}\,. \tag{2.1}$$

A higher bid-ask spread indicates that an asset is less liquid. In the example in Figure 2.1 the bid-ask spread is two ticks in Panel (a) (two price levels from the leftmost ask limit orders in yellow and the rightmost bid limit orders in blue) and three ticks in Panel (b). Note that here we assume a round-trip trade of one unit of the asset, as it is customary to measure bid-ask spread for the smallest possible order size. If we were to calculate bid-ask spread in the same example for e.g. a trade of size five units, the result would be higher.

Measuring round-trip cost in OTC markets is a more intricate matter for two reasons. First, price discrimination implies that round-trip costs differ from agent to agent. Quotes that are available from market services are only indicative. Second, even indicative quotes are often not available for a large number of assets. Therefore it is customary to estimate liquidity based on datasets that record trades.

Effective bid-ask spread, proposed by Hong and Warga (2000), is defined as the difference

[1] Market liquidity is to be distinguished from *funding liquidity*, the availability of credit to finance the purchase of assets. Both are ultimately related, see e.g. Brunnermeier and Pedersen (2009).

[2] For recent overviews on the topic see also Goyenko et al. (2009) and Schestag et al. (2016).

between the average buy and the average sell price[3] (and normalized by their midprice):

$$\textit{effective bid-ask spread} = \frac{\overline{price}_{buy} - \overline{price}_{sell}}{1/2(\overline{price}_{buy} + \overline{price}_{sell})} \,. \tag{2.2}$$

Other measures do not require the knowledge of whether a trade was a buy or a sell. For example *imputed round-trip cost*, developed in Feldhütter (2011) and applied for OTC markets in Dick-Nielsen et al. (2012), proxies the bid-ask spread by comparing the highest to the lowest price of a set of transactions. A simplified definition is

$$\textit{imputed round-trip cost} = 1 - \frac{\min(price)}{\max(price)} \tag{2.3}$$

while the actual calculation averages this measure over sets of trades with identical size.[4]

The *Roll measure*, developed in Roll (1984), relates the autocorrelation of returns to round-trip cost. It is obtained as twice the square root of the negative autocovariance of returns:

$$\textit{Roll measure} = 2\sqrt{-\operatorname{cov}(return_j, return_{j-1})} \tag{2.4}$$

where $return_j$ is the return due to trade j, i.e. $return_j = \frac{price_{after\ trade\ j}}{price_{before\ trade\ j}} - 1$.

2.2.2 Trading activity

Another dimension of liquidity is trading activity which also provides a proxy for how quickly one is able to conclude a large trade.

Turnover is a measure of how much of an asset was traded in a certain period of time. It can be defined quite generally as

$$turnover_{[t_{start}, t_{end}]} = \sum_{t_j \in [t_{start}, t_{end}]} \textit{order size}_j \tag{2.5}$$

where t_{start} and t_{end} are the start and end point of the reference time period respectively and the sum is over all trades j happening in that period. Where there is a well-defined amount outstanding for a security (e.g. for company shares or bonds, but not for e.g. currencies) turnover is often rescaled by the outstanding amount.

A different approach is to capture trading activity is to determine how often a security is traded or how often its price changes. If a security is not traded in a day, then its price (recorded as the price of the last trade) does not change. Such time series of daily prices are often easier to obtain than data on single trades. Lesmond et al. (1999) use the share of zero-return days in a period t:

$$\textit{proportion of zero return days}_t = \frac{\#\textit{zero return days}}{\#\textit{days}} \,. \tag{2.6}$$

Based on a similar idea Bandi et al. (2017) defines *excess idle time* via the occurrence of intraday periods without price changes that last longer than expected.

[3]The average is taken over all sell (buy) transactions in the respective time period, that is for buys $\overline{price}_{buy} = 1/n_{buy} \sum_{j=1}^{n_{buy}} price_{buy,j}$, where n_{buy} is the number of buy trades.

[4]See Chapter 7 and especially Appendix 7.B therein for a precise definition and further liquidity measures.

2.2.3 Price impact as a liquidity dimension

Price impact is at the core of market microstructure and the question how order flow and prices are related to one another. We will touch on this in more detail in the next section and instead focus here on the aspects related to liquidity. If a small trade shifts the price of the traded asset by a large amount, then the asset is considered illiquid.

One metric to measure price impact is the *lambda* proposed in Hasbrouck (2009). It is obtained as the coefficient λ obtained from regressing returns on order volume:

$$return_t = \lambda \epsilon_t \sqrt{order\ volume_t} + \varepsilon_t. \tag{2.7}$$

The return is defined e.g. over a five-minute interval, i.e. $return_t = \frac{price_t}{price_{t-5mins}} - 1$, $order\ volume_t$ is the absolute value of aggregated order volumes of trades during the same interval and ϵ_j is customarily defined as $+1$ (-1) if the majority of order volume stemmed from buy (sell) trades. The dependence on the square root of order volume is motivated by empirical studies that find that the price impact of metaorders is well-described by the so-called square-root law (see e.g. Tóth et al. (2011)).

Another standard measure for lower frequencies is the *Amihud measure*, proposed in Amihud (2002). It is computed as the mean ratio of absolute returns to trade volumes:

$$Amihud_t = \frac{1}{\#trades} \sum_{j \in t} \frac{|return_j|}{order\ size_j} \tag{2.8}$$

where again the sum is over all trades j in period t.

Liquidity is a central theme in market microstructure as well as in this thesis. As this section shows, already the measurement of liquidity is a research topic in its own. For example in Chapter 3 we describe several aspects related to liquidity on the bond exchange MOT. Chapter 7 is primarily concerned with the study of liquidity in the over-the-counter market for German corporate bonds. One key question therein is how far some of the above measures, which have been developed with U.S. bond markets in mind, are also applicable to other global markets.

Also the dynamics of liquidity are of considerable interest (Chordia et al., 2000; Karolyi et al., 2012). Financial crisis are typically accompanied with drops in market liquidity (Kyle and Xiong, 2001; Garleanu and Pedersen, 2007; Brunnermeier and Pedersen, 2009). With this in mind Chapter 5 studies how illiquidity shocks, i.e. sudden drops in market liquidity, propagate over time and across assets.

Finally let us point out that liquidity is a crucial ingredient to many studies in economics and finance, further underlining the importance of its truthful measurement. For example a strand of literature explores the *liquidity premium*, i.e. the extent to which market liquidity is reflected in asset prices and returns (Amihud and Mendelson, 1986). Also the analysis we perform in Chapter 6 relies on liquidity as a determining factor.

2.3 Price impact

Already in the previous section we have touched on the interdependence of order flow and asset prices. Simply put, buying (selling) an asset will on average raise (lower) its price. How exactly this happens is subject of ongoing research. This section will illustrate some key concepts and results of this topic that is at the heart of market microstructure: the price formation process.

2.3.1 Response

The response function R_ℓ measures the average price change conditional on a trade. It is defined as

$$R_\ell = \mathbb{E}\left[\left(price_{t+\ell} - price_t\right)\epsilon_t | trade_t\right] \tag{2.9}$$

where t is the time of a trade, $price_t$ is the price just before the trade and ℓ is a lag that could either be in trade time (i.e. ℓ trades ahead) or in physical time (e.g. 10 minutes ahead). ϵ_t is the *market order sign* or *trade sign* of the trade at time t, defined as $+1$ for a buyer-initiated trade and -1 for a seller-initiated trade.[5] Multiplying by ϵ_t ensures that the response to buys and sells is treated symmetrically.

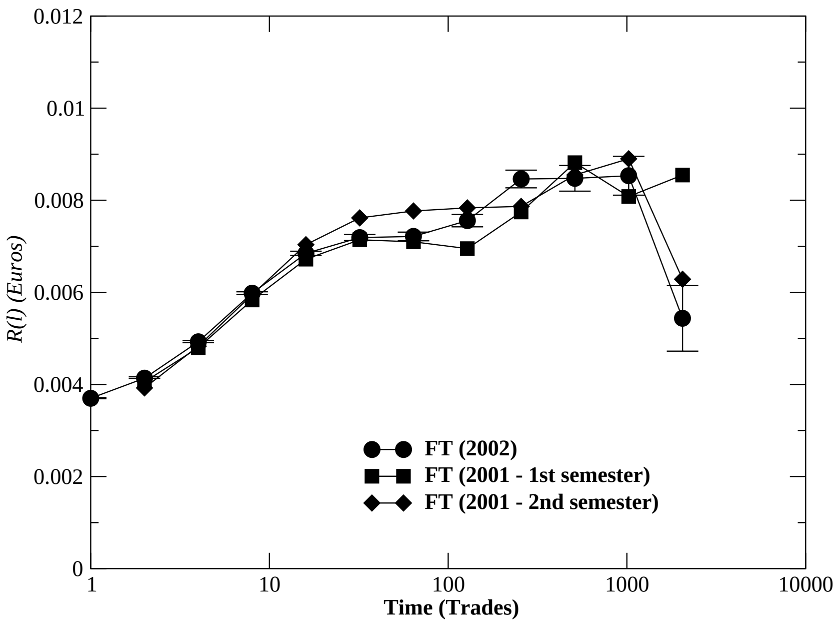

Figure 2.2: Response function: Average price change conditional on a trade. The x-axis shows the lag ℓ since the trade in logarithmic scale (trade time). The asset is the stock 'France Telecom' (FT) and data is from 2001 to 2002. The response R_ℓ on the y-axis has been rescaled for the 2001 data to collapse onto the 2002 data. Source: Bouchaud et al. (2008), adapted therein from Bouchaud et al. (2004).

Figure 2.2 shows an example of a response function for the French stock 'France Telecom'. Response is positive (reflecting the intuition that on average buys move prices upward) and increases by a factor of ~ 2 between $\ell = 1$ and $\ell = 100$. Qualitatively similar response functions have been found for different stocks, periods, markets and asset classes.

Crucially, response is not equal to impact. To see this, consider a simple toy model where buying an asset moves its price up instantly by 1 EUR. Let us assume, for example, that only two buys happen, driving the price up by 1 EUR each. Calculating the response function on this (admittedly small) sample gives a response of 1 EUR for $\ell = 1$ (in line with the model) and a response of 2 EUR for $\ell = 2$, in disagreement with the model. While

[5] While every trade is between a buyer and a seller, in a double-auction market the initiator (or aggressor) of the trade is clearly identified as the party that traded with a market order against an existing limit order.

this example does not carry any statistical significance, it illustrates the limitations of the response function with regard to order flow: if a buy trade is more likely followed by a buy trade, this distorts the observed response from the true impact. In the next section we show that this is indeed the case in practice.

2.3.2 Autocorrelation of order flow

To quantify the notion that a buy (sell) trade is more likely to be followed by another buy (sell) trade than a sell (buy) we use the autocorrelation function (ACF) of the order sign ϵ:

$$ACF_\ell = \frac{\mathbb{E}\left[(\epsilon_t - \mu)(\epsilon_{t+\ell} - \mu)\right]}{\sigma^2} \tag{2.10}$$

where $\mu = \mathbb{E}[\epsilon]$ and $\sigma^2 = \mathbb{E}\left[(\epsilon - \mu)^2\right]$ are the mean and variance of the order sign respectively. ℓ, as above, is the lag, customarily given in number of trades. A positive (negative) autocorrelation function implies that after a buy trade the next trade is more likely a buy (sell).

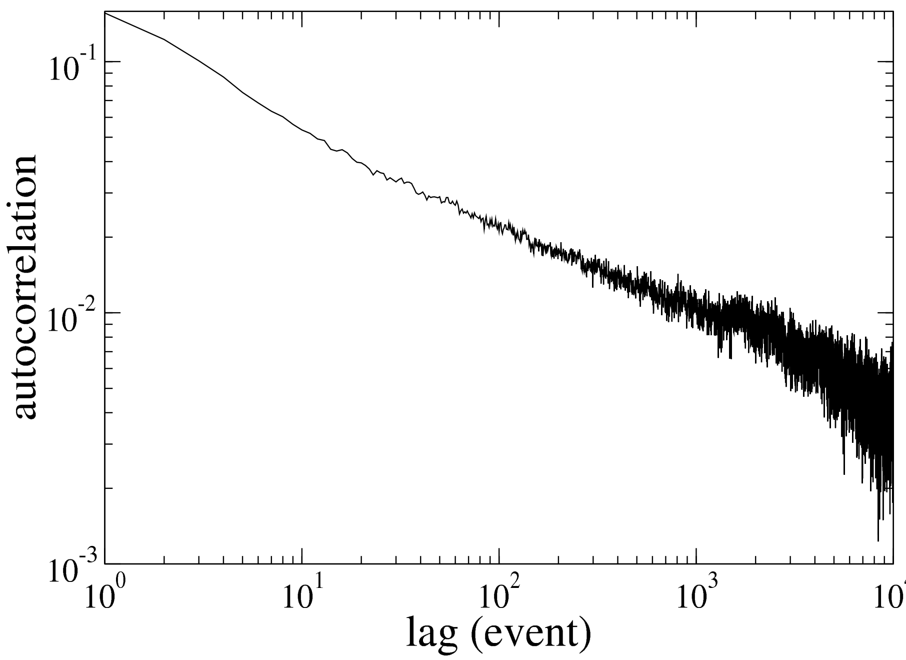

Figure 2.3: Autocorrelation of the market order sign: autocorrelation function of the order sign (± 1 for buy/sell trades) versus lag ℓ in trade events. Data for the stock 'Vodafone' traded on London Stock Exchange in the period 1999-2002. Source: Lillo and Farmer (2004); Bouchaud et al. (2008).

Figure 2.3 shows the autocorrelation function estimated for the stock 'Vodafone' based on trades on London Stock Exchange in 1999-2002. Autocorrelation is clearly positive even for large lags that are longer than the average number of trades in a day and indeed Lillo and Farmer (2004) demonstrate that the order sign is a long memory process. Similar autocorrelation functions have been found also for other order types (limit orders or cancellations), assets or asset classes, markets and time periods.

Two explanations have been proposed for why the order sign is so persistent even over longer periods of time. LeBaron and Yamamoto (2007) suggest *herding*, i.e. different agents trading in a similar fashion. On the other hand Lillo et al. (2005) show that order splitting can explain the long memory of the order sign. *Order splitting* refers to the practice of traders to break up large orders (that would considerably walk up the limit order book if executed in one piece and thus be more expensive) into smaller orders which are then executed incrementally. It has been demonstrated empirically by Toth et al. (2015) that order splitting is the dominating effect causing the persistent autocorrelation of order flow.

2.3.3 Price impact models

We will now seek to arrive at a statistical model for price impact (also termed *market impact*), i.e. a model for the price process as a function of the order flow. We have already presented a simple version of such a model with the toy model in Section 2.3.1. In this simplistic toy model prices are moved up (down) 1 EUR by an arriving buy (sell) market order and otherwise constant. We have also seen in Section 2.3.2 that the order flow shows a long-range autocorrelation and is therefore predictable to some extent. But then in our toy model also prices would be predictable. This is at odds both with observation and the notion that financial markets are statistically efficient, i.e. that market prices are diffusive and thus uncorrelated.

An impact model that is consistent with the long-range memory of the order sign is the propagator model of Bouchaud et al. (2004), a simplified version of which we present here.[6] Crucially impact in the progator model is not permanent (as in our toy model) but decays as a function of time. The price process is defined as

$$price_j = \sum_{k<j} G(j-k)\epsilon_k + \sum_{k<j} \xi_k + price_{-\infty} \tag{2.11}$$

where G is the propagator or decay kernel and ξ is a noise term that captures effects unrelated to the order flow, e.g. news. Price is thus a linear combination of the impacts of previous trades, weighted by the propagator as a function of the time passed since the trade, plus an independent noise term. It is possible to link the response function R_ℓ in equation (2.9) and the autocorrelation function of the order sign ACF_ℓ in equation (2.10) with the price process in equation (2.11). This allows to compute the propagator G from empirically estimated response and autocorrelation functions R_ℓ and ACF_ℓ (Bouchaud et al., 2004; Eisler et al., 2012).

Figure 2.4 shows the propagator $G(\ell)$ estimated for company shares of Apple. The propagator decays with increasing lag ℓ, that is while trades further in the past continue to influence the current price, the strength of this impact reduces the further in the past a trade took place.

The academic interest towards a detailed understanding of price impact is crucial in a multitude of ways. First, a deeper knowledge of the price formation process allows for a better understanding of related issues such as volatility and extreme price movements. Second, estimations of transaction costs in the sense of Section 2.2.3 can benefit from insights about the microstructure of price impact. Finally, understanding price impact is also essential to the field of optimal execution, which is concerned with finding strategies

[6] The simplification is that we neglect the role of order size in the model, as we do throughout this introduction for the sake of accessibility. Note that also other models are consistent with the persistence of the order flow sign. E.g. in the model of Lillo and Farmer (2004) impact is permanent and its strength depends on the expected order sign. Bouchaud et al. (2008) shows that both models are equivalent under additional assumptions.

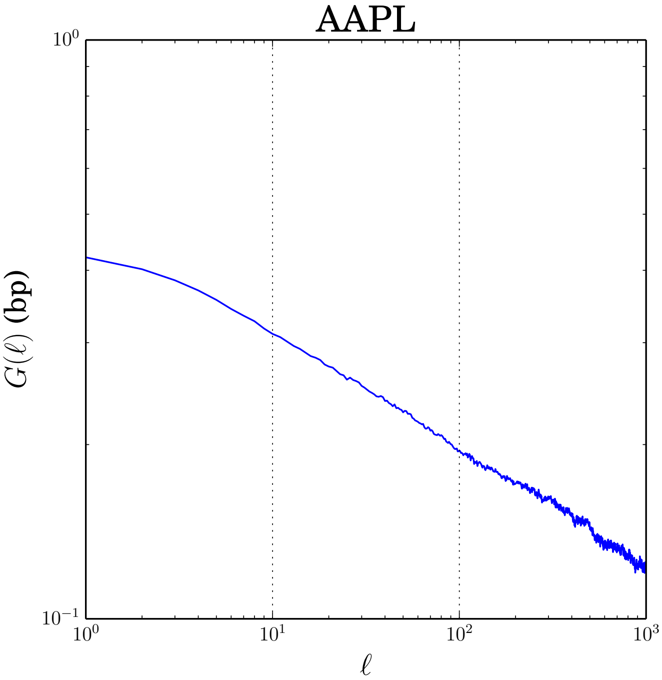

Figure 2.4: Decay kernel of price impact: Propagator of price impact G as defined in equation (2.11) for the stock 'Apple' (AAPL) estimated from NYSE and NASDAQ data for the period February to April 2013. G is given in units of basis points (100 bp = 1%) and the lag ℓ is given in trade time. Both axis are logarithmic. Source: Taranto et al. (2016).

that minimize costs and risks associated with trading (Almgren and Chriss, 2001). Thereby price impact is also of utmost practical importance.

In Chapter 4 of this thesis we consider price impact across multiple assets, i.e. how the order flow of one asset impacts the price process of another asset. This is termed *cross-impact*. We derive conditions for a certain class of models of cross-impact in order to be well-behaved and we estimate cross-impact, comparing the theoretical conditions to our empirical findings.

Chapter 3

Stylized facts of the MOT bond platform

3.1 Introduction

Trading in European fixed income markets takes place, either entirely or predominantly, in over-the-counter markets. Due to the opaque nature of over-the-counter trading in Europe, most studies on e.g. European sovereign bond markets have relied on data from the interdealer exchange MTS. Access to this platform is restricted to dealer banks and minimum order sizes are typically 2 million EUR or higher. In this chapter we focus on a different segment of the market, that has so far received considerably less attention from researchers. That is we study the MOT fixed-income platform,[1] aimed at the retail segment of the market and organized as a fully electronic limit order book market. While also listing a wide range of European fixed income securities, activity on MOT is concentrated on Italian bonds and especially sovereign bonds. According to bi-annual statistics reported by the Italian securities and exchange commission CONSOB MOT accounted for 8.7% (8.8%) of traded value excluding the OTC market in 2014 (2015) as the third-largest platform for Italian sovereign bonds, behind the interdealer exchange MTS and the request-for-quote platform BondVision.[2]

We are aware of only one other study on the MOT market: Linciano et al. (2014) compares the liquidity of dual-listed corporate bonds that are traded not only on MOT but also on the EuroTLX platform. They proxy liquidity through a set of metrics that have mostly been employed to gauge liquidity in over-the-counter markets (turnover ratio, zero trade days, Amihud and Roll measures). Instead here we make use of data that includes information on the limit order book as well as trades. We describe several stylized facts of trading and quoting activity on MOT with the aim of making further analysis accessible. To this end Section 3.2 presents the market structure and rules of MOT and Section 3.3 discusses our data and how we prepare it. In Section 3.4 we discuss intraday patterns in liquidity (proxied by the quoted bid-ask spread and the volume quoted in the limit order book) and trading activity, whereas Sections 3.5 and 3.6 discuss stylized facts in trading and quoting activity respectively. Section 3.7 concludes.

---

[1] `http://www.lseg.com/areas-expertise/our-markets/borsa-italiana/fixed-income-markets/mot`

[2] CONSOB, Bollettino Statistico Nr. 8, March 2016, available at `http://www.consob.it/web/area-pubblica/bollettino-statistico`

3.2 Market rules and mechanisms

In this section we describe the structure and the functioning of the MOT market as described in the official rules and instructions.[3]

3.2.1 Market structure

Traded on MOT are "bonds other than convertible ones, Government securities, euro-bonds, structured bonds, covered bonds, ABS and other debt securities and instruments tradable in the monetary market." The market is divided into two segments distinguished by their settlement mechanisms:

- DomesticMOT for "financial instruments settled via the settlement system managed by Monte Titoli S.p.A."

- EuroMOT for "financial instruments settled via foreign settlement systems managed by Euroclear and Clearstream Banking."

The DomesticMOT segment is further divided into two classes for Italian government securities and other debt securities respectively, while EuroMOT constitutes a single "Eurobond, ABS, securities of foreign issuers and other debt securities class".[4]

3.2.2 Market rules

Trading on MOT starts with an opening auction phase that lasts from 08:00 to 09:00 and then gives way to the continuous trading phase which lasts until 17:30. A closing auction does not take place.

The opening auction is concluded at a random time between 09:00:00-09:00:59 where the random time can differ across bonds. The continuous trading phase starts at the end of the opening auction phase and lasts until 17:30. It is organized as an electronic limit order book with limit and market orders. (Partially) unfilled limit and market to limit orders from the pre-auction phase are automatically transferred to the continuous trading phase as limit orders. The continuous phase also allows for "iceberg orders", limit orders of which only a partial quantity is visible, subject to a minimum displayed size. Once such an iceberg order is fully filled, a new limit order is created for the same partial quantity or the remaining quantity of the order. Execution of all orders follows price-time priority and for newly revealed parts of iceberg orders time priority corresponds to the time of the the new order.

Besides the order types above, the market also permits pre-arranged trades. So called "committed cross" orders can be entered with the aim of concluding contracts with a given counterparty for a price at or between the best bid and ask price. The same applies to "Block Trade Facilities" (BTF) that require a minimum size and can be executed at a wider price range than committed cross orders.[5]

[3] A non-binding English translations of the rules and instructions can be downloaded at `http://www.borsaitaliana.it/borsaitaliana/regolamenti/regolamenti/regolamentoborsa-istruzionialregolamento.en.htm` and `http://www.borsaitaliana.it/borsaitaliana/regolamenti/istruzioni/istruzioni.en.htm` respectively. For consistency with our sample period we refer to the versions effective 16 February 2016. We are not aware of any significant changes to these rules as of January 2018.

[4] Instructions accompanying the Rules of the Markets organised and managed by Borsa Italiana S.p.A., version effective 16 February 2015, Article IA 6.2.1

[5] For government bonds the allowed variation from the best prices is 0.75%.

Bond prices are quoted in percentage points of nominal value and the tick size is determined by the residual life time of a bonds. Bonds with a remaining time to maturity of more than (equal to or less than) two years have a tick size of 1/100 (1/1000) percentage points, i.e. 1 basis point (0.1 basis points) of nominal.

The presence of a specialist or a bid specialist is possible but in practice only the case for a subset of financial sector coporate bonds.

3.2.3 Volatility auctions

If certain price limits are violated, a volatility auction phase is initiated for a duration of 10 minutes plus a random interval shorter than one minute.[6] Price limits are defined in percentage terms of a reference price[7] and there are two reference prices for each bond: the dynamic price and the static price. That is volatility auctions are triggered by either of the following events:

1. If the price of a "contract in conclusion", e.g. a market order entered but not yet executed, would exceed the maximum price variation limit with respect to the static price. The static price is defined as the price of contracts concluded in the last auction phase (or, if there were no trades, the price of the first contract concluded in the continuous trading phase).

2. The price of a contract in conclusion would exceed the maximum price variation limit with respect to the dynamic price. The dynamic price is defined as the price of the last contract concluded in the current session (or the previous days reference price if no contracts have been concluded in the current session).

Volatility auctions may be re-iterated and thus last longer than 10-11 minutes.

Our dataset does not include flags or identifiers for volatility auctions. However we can deduce their presence when there is a crossing between the bid and ask side of the limit order book in the absence of trades. Specifically we detect a volatility auction either when the bid-ask spread is zero or negative or when one or both of best bid and best ask price is recorded as zero, indicating the presence of an unexecuted market order with no limit price.

In Figure 3.1 we illustrate these rules at the example of trading in GR0138014809 on 16 January 2015. The maximum allowed price variation with respect to the static price is 3.5% and 2% with respect to the dynamic price. It is possible that there are further volatility auctions than we can not detect with the above identification criteria, especially during the first half of this day. In total we detect $1,364$ volatility auction periods, $1,058$ of which are for the two Greek bonds in our sample.

Bormetti et al. (2015) identifies volatility auction phases on an equities market with similar rules through the absence of trades in windows of multiples of 10-11 minutes. Since for many of the bonds in our sample trading happens less frequently (and absence of trades for a period of 10 minutes or more would not necessarily imply a high likelihood for the presence of a volatility auction), this method is not suitable for the MOT bond market.

[6]See e.g. Reboredo (2012) for a study on the effect of volatility auctions on trading.

[7]E.g. for a reference price of 100 EUR and a "maximum price variation limit" of 2% the price limits are 98 and 102 EUR.

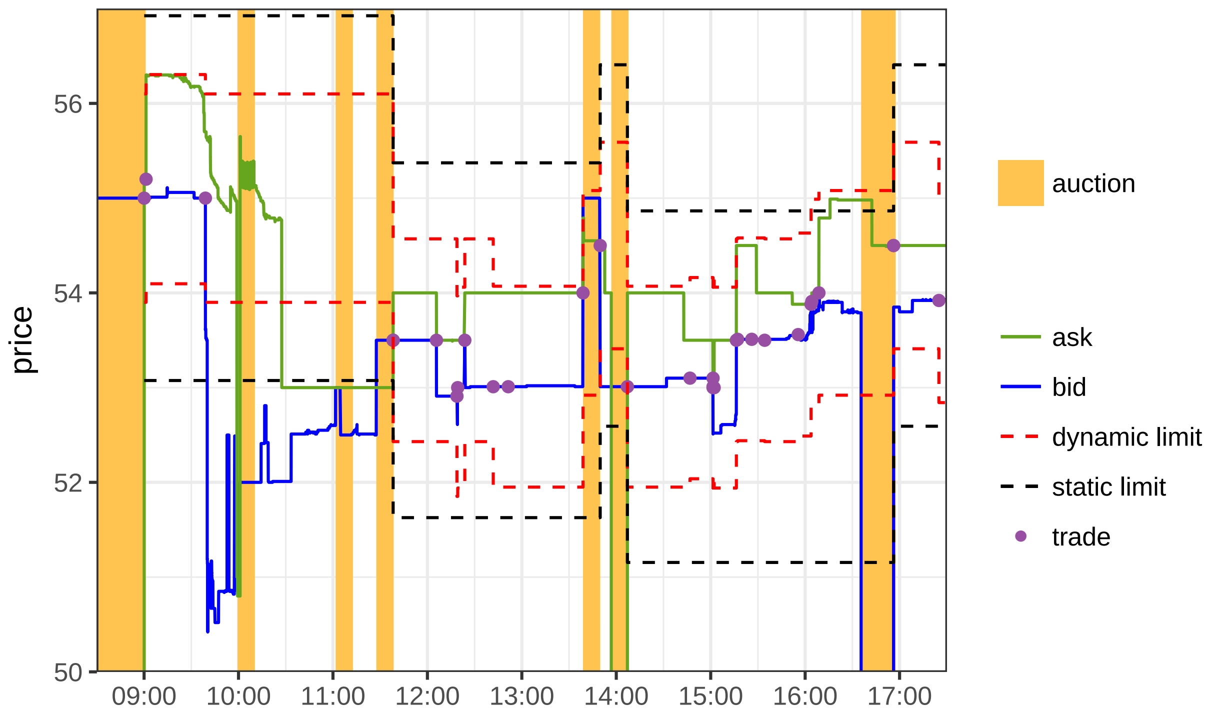

Figure 3.1: Example of volatility auctions in GR0138014809 on 16 January 2015. Best bid and ask price are drawn as solid blue and green lines respectively, and concluded trades as purple dots. The limits corresponding to the dynamic and the static price are shown as dashed red and black lines respectively. Periods with detected volatility auctions are shaded in orange. The first shaded period until 09:01 corresponds to the opening auction. Three volatility auction periods are detected between the trades at 09:38 and 11:38 due to crossed bid and ask quotes, whereas it is likely that more auction phases took place during that time. Volatility auctions are triggered by attempted trades at prices beyond the limits around either the dynamic or the static price. The auction phase at 13:38 is likely due to a violation of the upper dynamic limit, whereas the last auction phase around 17:00 is due to the static limit (an arriving buy market order, indicated by the best bid price dropping to zero, would have been executed against the best ask level which was beyond the limit around the prevailing static price.)

3.3 Data

3.3.1 Data description

Our sample period spans the 194 trading days from 01 Dec 2014 to 27 Feb 2015 and from 13 April 2015 to 16 October 2015, while our bond universe consists of 66 sovereign and corporate bonds. The bond universe consists of mostly Italian sovereign bonds, some European sovereign bonds and Italian corporate bonds, which were hand-picked to represent both a large share of trading activity on the platform as well as a cross-section of the traded asset types. Table 3.2 in the Appendix 3.A lists the ISINs and gives basic descriptives of the coupon rate, maturity at issuance as well as tick size, number of orders in our sample, average order size and average bid-ask spread.

The dataset consists of separate trades and limit-order-book (LOB) datasets, covering all trades and almost all LOB updates.[8] The trades data consists of 2.0 million trades and carries the information of date, ISIN code, intraday time, price and size of a trade. The LOB is recorded as snapshots of the best ten bid and ask side levels. Each entry

[8] In phases of heavy trading multiple updates of the LOB may be summarized into one update of our data. However we are guaranteed at least one update per second whenever there are changes to the LOB and in the vast majority of our sample updates are more frequent.

contains the date, ISIN code and intraday time of the snapshot, as well es the number of occupied levels on the bid- and ask side of the LOB (up to the maximum of 10). For each of these up to 20 levels, the price, available quantity and number of limit orders at the level is recorded. In total our dataset contains 174.8 million snapshots. The intraday time is recorded at millisecond precision as a 9 digit integer.[9]

3.3.2 Data preparation & checks

We prepare our dataset by detecting volatility auction periods as described in Section 3.2.3. All trades and limit orders before 09:01:00 or associated with a volatility auction period are flagged as such and excluded from our further analysis.

When a single market order is executed against more than one limit order, then it is recorded as a number of trades. We aggregate single trades in the same bond to orders when they happen within one millisecond.[10] That is the almost 2 million trades in our sample are remapped to roughly 1.4 million orders. We infer the order sign by comparing to the state of the LOB immediately preceding the trade. An order executed at the best bid (best ask) price is seller-initiated (buyer-initiated) and assigned the order sign -1 (+1).[11] Where an order does not happen at the best price we infer the sign by applying the method of Lee and Ready (1991). By this identification method 93.7% (94.1%) of orders (trades) occur at the best price and of the remaining orders (trades) 9.7%/37.9%/6.8%/35.9%/9.6% (12.5%/35.2%/6.5%/33.5%/12.3%) are below the best bid price/ inside the spread and below the mid price/ at the mid price/ inside the spread above the mid price/ above the best ask price respectively.

3.4 Intraday seasonalities

It is a well-established observation across other markets and venues that trading activity, liquidity and volatility are not uniform across the day, but follow certain intraday patterns, often reflecting e.g. opening and closing effects or lunch hours (see e.g. McInish and Wood (1990, 1992); Bollerslev and Domowitz (1993); Abhyankar et al. (1997); Ito and Hashimoto (2006)). In this section we describe the intraday seasonalities on MOT by looking to the intraday patterns of different measures of liquidity.

Figure 3.2 shows the intraday pattern of trading activity in terms of the share of the total nominal amount traded per day that is traded during bins of 10 minute time. If trading activity were constant throughout the day, each of the 51 bins should account for roughly 2% of daily trading activity. The figure shows instead that some variation is present, with on average roughly $1 - 2.5\%$ of total activity captured in a 10 minute bin. Less active is the period from 13:00 to 14:30 and the last 10 minute bin of the day, whereas the most active time of the day is during the morning hours from after the opening to 12:30. Note that the first 10 minute bin does not include trades concluded during the opening auction. The pattern we observe is present equally for the set of all bonds and for the 10 bonds with the largest trading volume in our sample as well as at the level of

[9]The format is "hhmmssMMM" where h/m/s/M are the hour/minute/second/millisecond digits respectively. E.g. a time of "09:17:25" and 89 milliseconds would be recorded as "091725089".

[10]Most orders that are executed as multiple trades are recorded at precisely the same millisecond timestamp. It can occur in rare cases that the timestamps of trades belonging to the same order span two consecutive milliseconds.

[11]When a market order walks up the LOB, i.e. consumes liquidity from levels deeper in the LOB than the best, the price of the first trade is used to determine the order sign. This order sign is then attributed to the order and all trades therein.

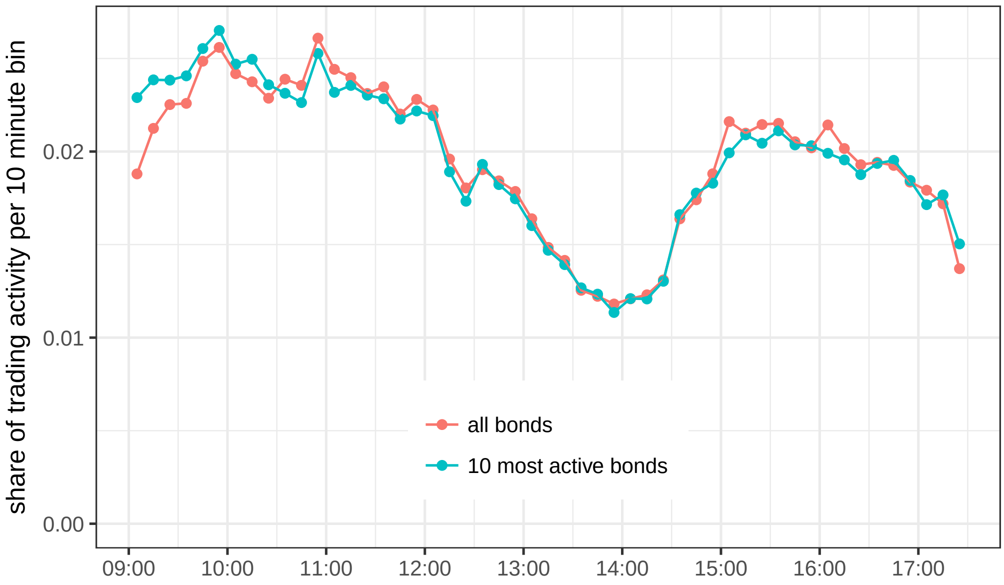

Figure 3.2: Intraday pattern of trading activity: share of daily trading activity in terms of volume traded during 10 minute bins, averaged over days and bonds. *all bonds* refers to all bonds in our sample and *10 most active bonds* are the 10 bonds with the highest trading volume in our sample. Trades occurring during the opening auction or volatility auctions are excluded.

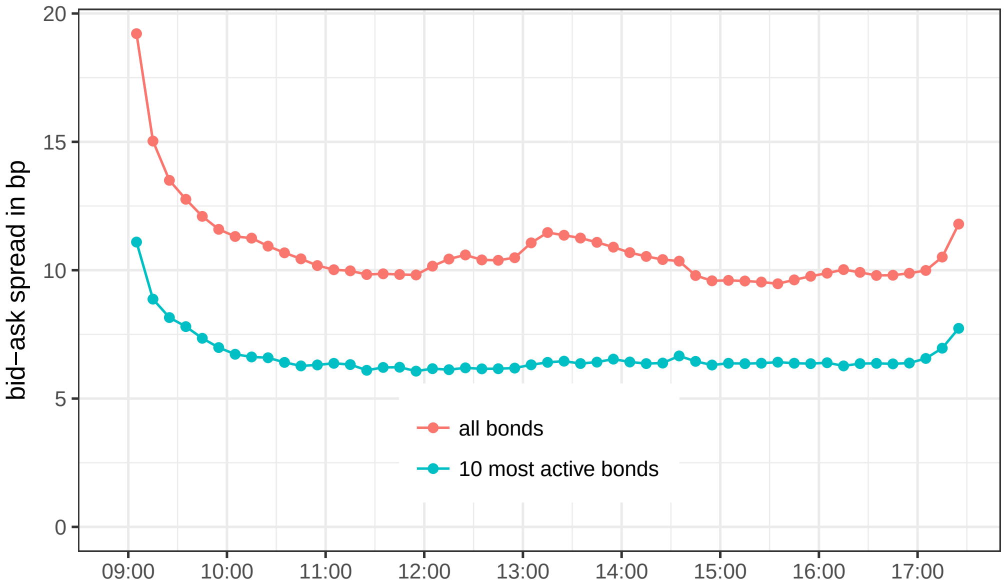

Figure 3.3: Intraday pattern of bid-ask spread: average quoted bid-ask spread during 10 minute bins, given in basis points of nominal value. Averages are computed as time-weighted averages per 10 minute bin, winsorized from above and below at the percentiles corresponding to 1% and 99% and averaged over bonds and days. *all bonds* refers to all bonds in our sample and *10 most active bonds* are the 10 bonds with the highest trading volume in our sample. Auction periods are excluded.

individual bonds.

The intraday pattern of the quoted bid-ask spread in Figure 3.3 follows a typical U-shape, i.e. bid-ask spread is larger in the beginning and in the end of the trading day. During the day bid-ask spread is fairly flat, especially in the sample of more active bonds, whereas from roughly noon to 14:30 a slight increase is visible when considering the full

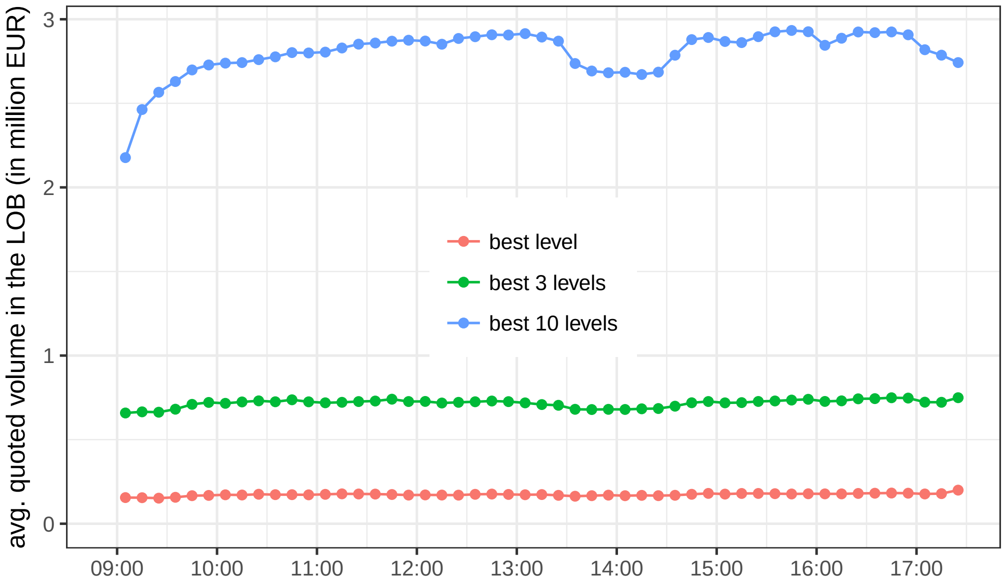

Figure 3.4: Intraday pattern of quoted volume in the limit order book: average quoted volume available at the best price level of the LOB, the three best levels and all observed 10 best levels of the LOB. Quoted volume is given as nominal amount in million EUR and is averaged over bid and ask side. The sample consists of the 10 bonds with the highest trading volume in our sample; results are similar for the sample of all bonds. Averages are computed as time-weighted averages per 10 minute bin, winsorized from above and below at the percentiles corresponding to 1% and 99% and averaged over bonds and days. Auction periods are excluded.

sample of bonds.

Finally Figure 3.4 shows the intraday pattern of the quoted volume at different levels of the limit order book for the sample of the 10 most actively traded bonds. There is no evident pattern when considering the volume quoted at the best or the best three levels on both sides of the book, indicating that the available liquidity for small- and medium-sized trades is roughly constant throughout the day. Only when considering all 10 observable levels of the limit order book there is an inverse U-shape pattern analogous to Figure 3.3, i.e. right after the opening, around 14:00 and towards the closing there is less volume available in the deeper levels of the book.

3.5 Stylized facts of trading activity

This section gives an overview of trading activity by describing the distribution of market order sizes and the sign autocorrelation of trades.

Table 3.1: Distribution of market order size: nominal amount of market orders in EUR (scaled by 1000 EUR). *all bonds* refers to all bonds in our sample and *10 most active bonds* are the 10 bonds with the highest trading volume in our sample. Orders during auction periods are included. The distribution is very similar when considering only orders occuring at the best price.

×1000EUR	mean	std dev	5 Pcl	25 Pcl	median	75 Pcl	95 Pcl	# orders
all bonds	87.4	200.2	2.0	10.0	31.0	80.0	350.0	1,378,730
10 most active bonds	86.8	191.8	2.0	13.0	40.0	81.0	325.0	850,214

Table 3.1 provides summary statistics for the nominal amount of orders. While more than 60% of trading activity in terms of number of trades and traded volume is concentrated

in the most active 10 bonds of our sample, their distribution is very similar, and also when we only consider orders that happen at the best price of the order book. The average order size is only $87,000$ EUR, reflecting that MOT is a retail market, and the 5% and 25% percentiles at $2,000$ and $10,000$ EUR respectively, indicating a large number of small trades. The upper percentiles of trade size reflect the presence of some larger trades, with the 95%, 99% and 99.9% percentiles at $350,000$ EUR, 1 million EUR and 2 million EUR respectively. The latter amount is comparable in size to small trades on the sovereign bond interdealer exchange MTS.

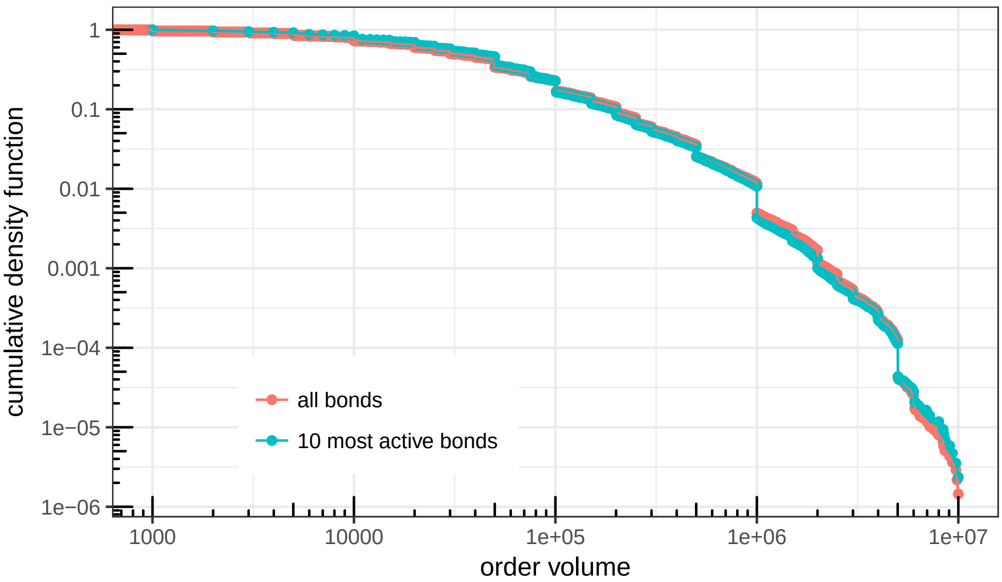

Figure 3.5: CDF of market order size: cumulative density function of market order size, measured as nominal amount of orders in EUR. Log-log scale. *all bonds* refers to all bonds in our sample and *10 most active bonds* are the 10 bonds with the highest trading volume in our sample. Orders during auction periods are included. Vertical "drops", e.g. at 1 million EUR, indicate a preference for certain, typically round, order sizes.

Figure 3.5 shows the cumulative density function of MOT order size. It confirms that there is a large share of small orders and larger orders are increasingly rare, with the largest orders having a size of 10 million EUR. Again this is natural since MOT constitutes a retail market and larger orders mitigate to other trading venues or the over-the-counter market. This makes our setting different from studies on markets that encompass a majority of trading in an asset. E.g. Figure 3 in Lallouache and Abergel (2014) shows a power-law decay of order size over two orders of magnitude, while the decay we observe is much faster. Finally the vertical drops in Figure 3.5 also illustrate a strong preference for "round" order sizes, such as $100,000$ EUR or 1 million EUR.

Next we consider the autocorrelation of the market order sign. A positive autocorrelation at a lag of one trade event means that a buy (sell) order is more likely to be followed by another buy (sell) than one would expect from a random sequence of buy and sell orders. Indeed it has been found in Lillo and Farmer (2004); Bouchaud et al. (2004) that the time series of the order sign is a long memory process and two compatible explanations have been proposed. *Herding* refers to different agents acting in a similar fashion and has been proposed as an explanation in LeBaron and Yamamoto (2007), whereas Lillo et al. (2005) shows that order splitting, i.e. agents breaking large orders into smaller orders that are executed incrementally, can explain the long memory of the order sign. Toth et al. (2015) show empirically that order splitting is the dominating effect.

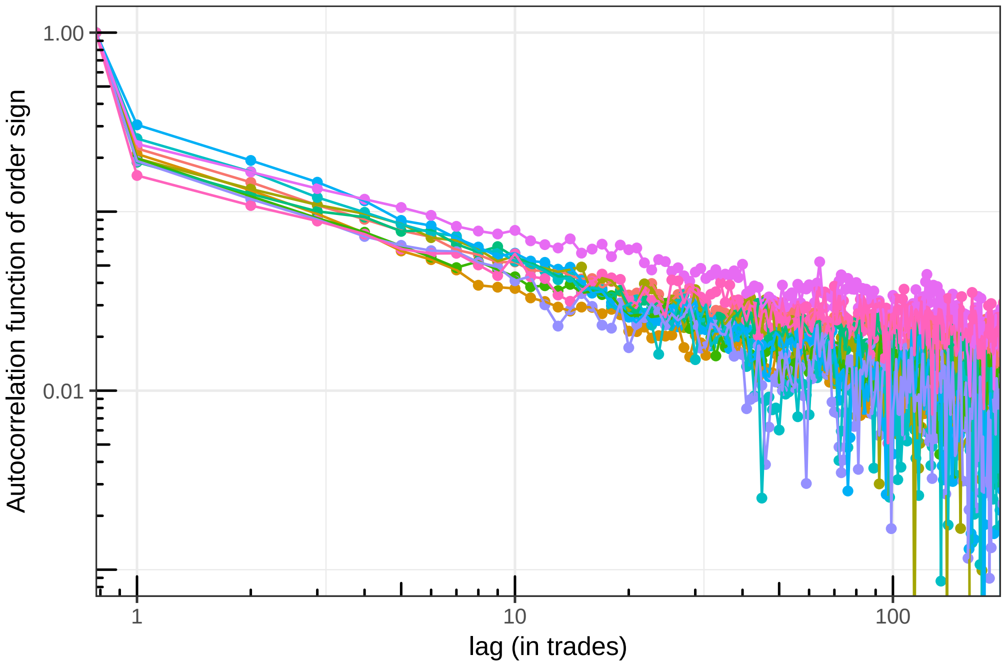

Figure 3.6: ACF of trade sign: autocorrelation function of the market order sign (± 1 for buy/sell orders) for the 10 bonds with the highest trading volume in our sample. Orders during auction periods are excluded. Different colors correspond to autocorrelation functions in different bonds. The result is similar when considering only orders occuring at the best price.

In Figure 3.6 we show the autocorrelation function of the order sign in the 10 most active bonds of our sample. In line with the studies of other markets we also find a positive correlation that is significant for lags of up to hundreds of trades.

3.6 Stylized facts of quoting activity

3.6.1 Shape of the limit oder book

In this section we study the average shape of the limit order book, i.e. how the volume quoted in the limit order book is distributed across price levels. Since tick size plays an important role we now distinguish three sets of bonds: we consider as *large tick* all bonds that have a tick size of 1 basis point or 0.01% of nominal value. The *10 most active bonds* are as before the bonds with the highest trading volume and are a subset of the large tick bonds. By the rules of MOT bonds with a remaining time to maturity of 2 years or less have a tick size of 0.1 basis points or 0.001% of nominal value and we refer to them as *small tick*.

Figure 3.7 shows the average volume quoted in the limit order book as a function of the distance to the best price. Since with our data we observe at most 10 levels of the limit order book, the representation is only exact up to a price difference to the best price of 9 ticks, and increasingly downwards biased for price differences of 10 ticks or more.[12] For

[12] In the case that the price difference between the ten first occupied levels of the LOB is minimal, i.e. one tick each, we have no information for the state of the book at a distance of 10 ticks to the best and beyond and treat this as absence of volume in these levels. That is we take the view of someone who, as us, observes only the best 10 occupied levels of the book. While we could also take into account gaps and consider only states of the book where we know our estimate to be exact, this would induce a bias towards states of the book with large gaps when considering average volumes at higher distances from the best.

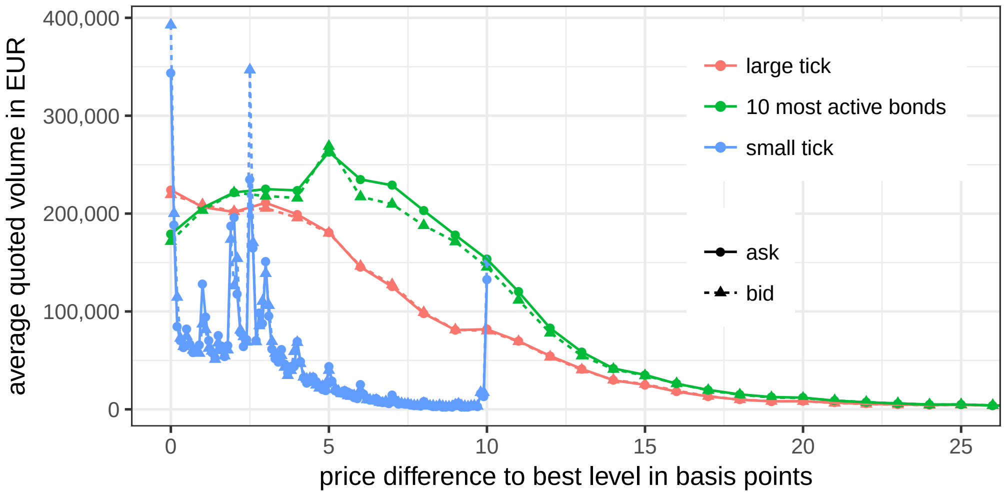

Figure 3.7: Shape of the limit order book: Average quoted volume in the limit order book as a function of the distance to the best quote. The shape is exact only for the first 10 levels. Given as nominal value in EUR. *large tick* refers to all bond-day observations in our sample where the tick size is 1 basis point and *small tick* are observations where the tick size is 0.1 basis points. *10 most active bonds* are the 10 bonds with the highest trading volume in our sample and all large tick. Averages are computed as time-weighted averages per day, winsorized from above and below at the percentiles corresponding to 1% and 99% and averaged over bonds and days. Auction periods are excluded.

computational reasons we calculate the shape of the book up to a distance of 30 ticks (30 basis points) for large tick bonds and up to 100 ticks (10 basis points) for small tick bonds. The graph is symmetrized for the bid and ask side, i.e. a positive price difference to the best corresponds to bid quotes at lower prices than the best bid for the bid side and ask prices above the best ask on the ask side. We do not observe any significant differences between the shape of the book on the bid and ask sides.

Volume in large tick bonds is concentrated in the first few levels. Averaging over all large tick bonds the maximum volume is at the best level and decreasing with distance to the best, whereas for the 10 most active bonds the volume increases initially, peaks at 5 basis points from the best and then decreases. Besides this ‘kink’ the average distribution of volume in large tick bonds fairly smooth. This is different for small tick bonds, where several peaks of volume are visible at ‘round’ distances from the best. The maximum of average volume is at the best level and peaks are clearly discernible at e.g. 1 basis point (10 ticks), 2 basis points (20 ticks), 2.5 basis points (25 ticks) and also 10 basis points (100 ticks) from the best. This implies that some market participants do not make use of the full spectrum of admissible prices and rather prefer price differences that are e.g. multiples of 10 ticks. A similar behavior has been observed and studied on the foreign exchange platform EBS by Lallouache and Abergel (2014) in the context of a tick size change. They find that this is caused by traders that prefer values on the old, coarser, price grid and argue that these are mostly manual traders. Automatic (algorithmic) traders can then take advantage of this behaviour e.g. obtaining price priority by placing a limit order just one tick before the limit orders of manual traders. Tick size changes for three bonds during our sample period and we will revisit this argument in more depth in section 3.6.2 below.

Figure 3.8 gives an alternative representation of the limit order book that takes into account the availability of only 10 levels of the book. Panels (a), (b) and (c) show the average volume quoted at each level, the probability of a level being occupied and the average gap size between two adjacent levels respectively. Unoccupied levels are counted as

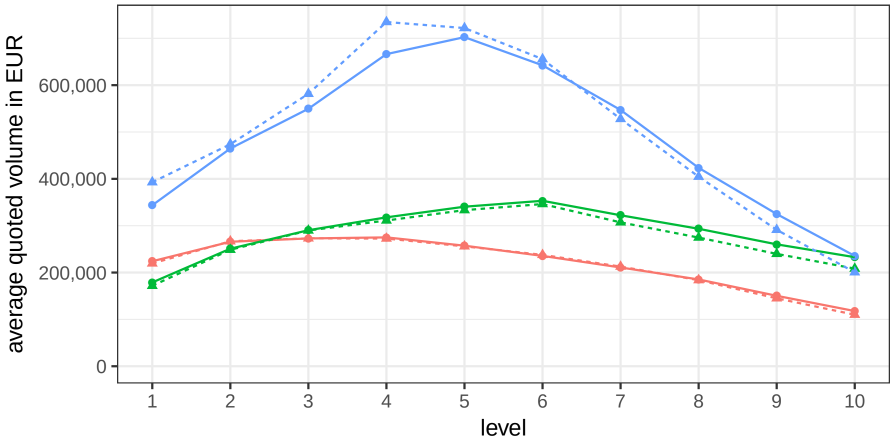

(a) Average quoted volume by level of the LOB in EUR of nominal amount.

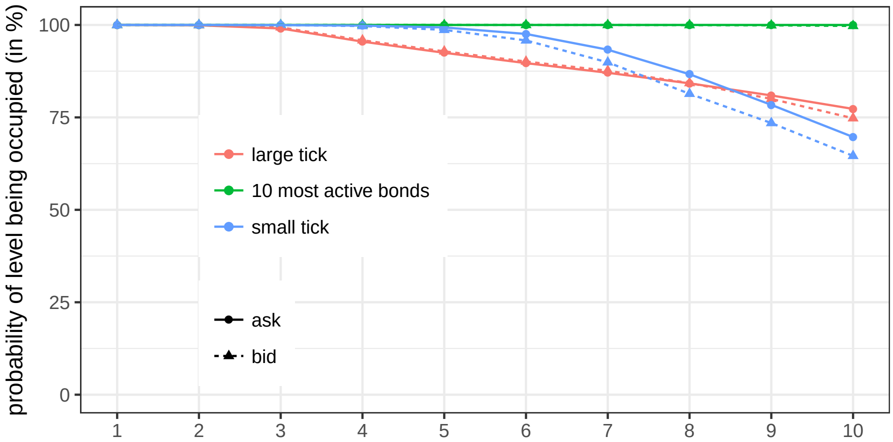

(b) Probability of a level being occupied.

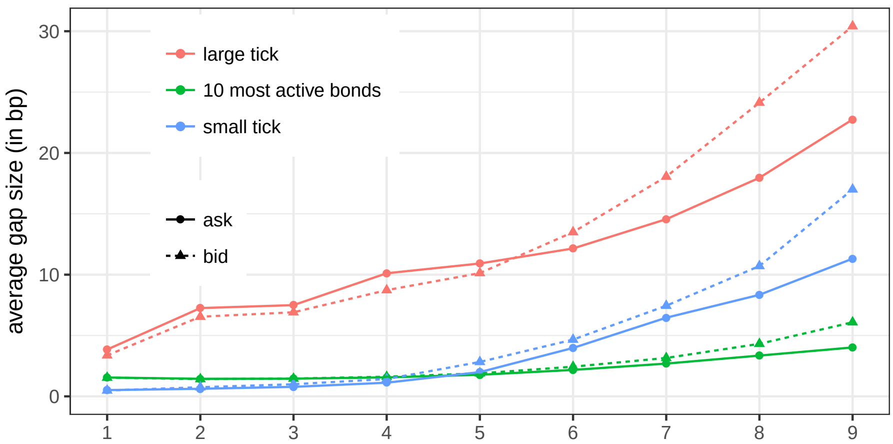

(c) Average gap between adjacent levels of the LOB in basis points of nominal value.

Figure 3.8: Shape of the limit order book: alternative representation that characterizes the shape of the limit order book. *large tick* refers to all bond-day observations in our sample where the tick size is 1 basis point and *small tick* are observations where the tick size is 0.1 basis points. *10 most active bonds* are the 10 bonds with the highest trading volume in our sample and all large tick. Averages are computed as time-weighted averages per day, winsorized from above and below at the percentiles corresponding to 1% and 99% and averaged over bonds and days. Auction periods are excluded.

zero volume in Panel (a). Panel (b) reveals that in the most active bonds all 10 observable levels are occupied virtually all of the time. For the other large tick bonds (for small tick bonds) this holds true only up to the third (fifth) level. Panel (a) of Figure 3.8 overlaps with Figure 3.7 for the first level (0 price difference to the best), whereas for all other levels also the gap size between two occupied levels determines the shape in Figure 3.7. Panel (c) in Figure 3.8 shows that the average gap size is increasing with higher levels. For the first few levels gap sizes are smallest for small tick bonds (measured in basis points, not ticks) and for larger levels the most active bonds have smallest gap sizes. The sample including less active large tick bonds always has the widest gaps between occupied levels. Figure 3.8 is unable however to reproduce the features present in Figure 3.7 especially for small tick bonds. In the following section we will examine how a tick size change affects this shape.

3.6.2 Impact of tick size changes

In Figure 3.7 small tick bonds (where the minimum tick size is 0.1 basis points of nominal value) and large tick bonds (1 basis point minimum tick size) show a strikingly different limit order book shape. This could, in principle, be due to differences in bond characteristics across the two subsamples. To rule out such factors, and for a more detailed analysis we consider in this section three Italian government bonds for which the minimum tick size changed during our sample time as their remaining maturity reached 2 years.

Figure 3.9 shows the effect of the tick size change for these three bonds by comparing the average shape of their limit order book during the two weeks before the tick size change with the average shape over the two weeks after the tick size change. Since we found only negligible differences between the shape on the bid and ask side of the limit order book in Section 3.6.1 we show only the ask side here, the bid side behaving similarly.

Across all bonds it is evident that after the tick size change volume is concentrated at distances from the best corresponding to multiples of 10 ticks, i.e. effectively mimicking the ‘old’ price grid. The shape after the tick size change also resembles the one from before and several features remain unchanged in the new, finer price grid. E.g. IT0005023459 shows a peak at 10 basis points from the best before the tick change which is still present on the finer grid and there even sharper.

The observations from Figure 3.9 suggest that a price grid that is too fine is not always accepted and fully used by market participants. Lallouache and Abergel (2014) find similar externalities to ours, where some traders (termed *manual traders*) stick to the old price grid after a general tick size change. Interestingly in their case the tick size change was later reverted to an interim minimum tick size. Indeed finding an ‘optimal’ tick size is a market design problem that is as complex as it is important to exchanges competing with one another for liquidity. E.g. Goldstein and Kavajecz (2000) reports that a minimum tick size reduction on the New York Stock Exchange led to decreases both in bid-ask spreads and the depth of the limit order book, to the advantage of small traders but at the disadvantage of participants wanting to trade larger order sizes. Huang et al. (2016) applies the approach of Dayri and Rosenbaum (2015) to predict the effects of tick size changes on Tokyo Stock Exchange on an ex-ante basis and define an optimal tick value. Note that the choice of tick size is not only up to the exchange (or regulators) but also companies can influence the effective tick size (defined as the minimum tick value divided by the asset price) of their shares by performing a (reverse) stock-split (Angel, 1997; Schultz, 2000).

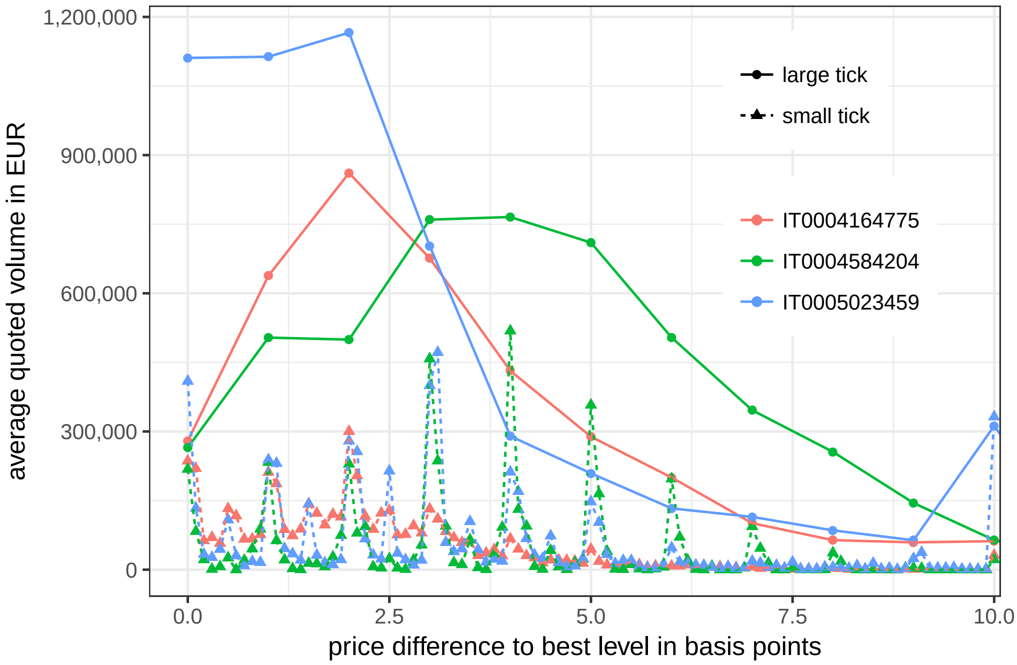

Figure 3.9: Shape of LOB around tick size changes: Average quoted volume in the limit order book as a function of the distance to the best quote. The shape is exact only for the first 10 levels. Given as nominal value in EUR. Shown is the ask side. The sample consists of bonds that underwent a tick size change as their remaining maturity reduced to two years and *large tick* refers to the two weeks before the tick size change when bonds where quoted on a grid of 1 basis points width. *small tick* refers to the two weeks after the tick size change when the minimum price difference between two quotes is 0.1 basis points. Averages are computed as time-weighted averages per day, winsorized from above and below at the percentiles corresponding to 1% and 99% and averaged over days. Auction periods are excluded.

3.7 Conclusion

In this chapter we have explored the bond trading platform MOT, describing its market structure and stylized facts of trading and quoting activity. To the best of our knowledge this is the first such study on a fixed income retail platform. Despite its focus on retail trades and the hybrid market structure of bond markets with a large share of over-the-counter trades we are able to reproduce several stylized facts known from exchange markets in other asset classes, especially equities, on MOT. This suggests it is well suited as a laboratory for the analysis of further topics in market microstructure. Indeed in the next chapter we use the same dataset to empirically test theoretical predictions for price impact across bonds.

Appendix

3.A List of ISINs

Table 3.2 lists the bonds in our sample together with descriptives of bond characteristics, trading parameters and indicators of bond liquidity.

Table 3.2: Descriptives for the set of bonds used in the estimation: *coupon* is the coupon rate in percent. *maturity* is the maturity at issuance of the bond in years. *tick size* is the minimum tick size in percent of nominal value. *changing* indicates that the tick size changes from 1 to 0.1 basis points during our sample. *# orders* is the number of orders for that bond and *order vol.* the average order volume therein in multiples of $1,000$ EUR. *bid-ask* is the average quoted bid-ask spread in basis points of nominal amount, calculated as time-weighted average excluding auction periods.

ISIN	coupon	maturity	tick size	# orders	order vol.	bid-ask
corporate bonds						
IT0004576978	3.5	6.0	0.001	5,075	11.9	6.2
IT0004576994	indexed	6.0	0.001	4,031	12.7	5.6
IT0004645542	5	10.0	0.01	4,577	25.2	16.9
IT0004720436	indexed	10.0	0.01	3,927	27.3	17.5
IT0004966823	5.5	7.0	0.01	6,412	24.4	12.4
IT0005056483	3	5.0	0.01	9,192	32.9	12.4
IT0005075533	6	7.0	0.01	6,894	32.3	22.7
non-Italian sovereign bonds						
DE0001102358	1.5	10.0	0.01	4,739	10.5	16.0
DE0001102366	1	9.9	0.01	555	73.1	12.1
DE0001102374	.5	10.1	0.01	1,416	95.0	12.2
ES00000123X3	4.4	10.5	0.01	3,004	9.4	20.3
ES00000126B2	2.75	10.4	0.01	494	13.5	22.6
ES00000126Z1	1.6	10.3	0.01	931	30.3	19.6
GR0114028534	4.75	5.0	0.01	12,259	18.2	42.8
GR0138014809	2	30.0	0.01	4,498	11.5	89.4
PTOTE5OE0007	4.1	31.1	0.01	23,982	32.2	23.9
Italian sovereign bonds						
IT0001278511	5.25	31.0	0.01	26,372	47.5	15.4
IT0003535157	5	31.0	0.01	24,926	49.3	23.4
IT0003844534	3.75	10.3	0.001	3,862	60.7	2.0
IT0003934657	4	31.5	0.01	223,909	71.9	5.6
IT0004009673	3.75	15.5	0.01	19,909	66.6	5.8
IT0004019581	3.75	10.4	0.001	6,827	59.5	2.4
IT0004164775	4	10.1	changing	5,413	65.0	3.8
IT0004273493	4.5	10.4	0.01	4,321	88.1	4.9
IT0004361041	4.5	10.3	0.01	5,707	110.7	4.9
IT0004423957	4.5	10.5	0.01	5,198	103.6	5.5
IT0004489610	4.25	10.3	0.01	10,132	93.2	5.7

Table continued on next page.

Table 3.2 continued from previous page.

ISIN	coupon	maturity	tick size	# orders	order vol.	bid-ask
Italian sovereign bonds (continued)						
IT0004518715	floating	6.9	0.001	11,725	106.4	1.9
IT0004532559	5	31.0	0.01	43,435	56.2	18.0
IT0004536949	4.25	10.4	0.01	9,510	155.1	5.9
IT0004545890	indexed	31.9	0.01	25,834	56.0	28.7
IT0004584204	floating	7.0	changing	10,996	88.0	2.6
IT0004594930	4	10.4	0.01	13,791	85.2	5.9
IT0004634132	3.75	10.5	0.01	11,710	81.9	6.2
IT0004695075	4.75	10.5	0.01	5,994	86.8	6.9
IT0004759673	5	10.5	0.01	7,343	93.6	7.7
IT0004794142	4.875	6.0	0.01	3,649	24.3	12.0
IT0004801541	5.5	10.5	0.01	5,985	87.9	8.3
IT0004848831	5.5	10.2	0.01	5,128	101.6	8.9
IT0004898034	4.5	10.2	0.01	8,638	96.2	8.4
IT0004917842	5.75	10.0	0.01	6,729	24.5	15.0
IT0004923998	4.75	31.3	0.01	62,991	69.7	15.7
IT0004953417	4.5	10.6	0.01	12,278	169.3	5.5
IT0004969207	indexed	4.0	0.01	39,101	109.3	3.3
IT0005001547	3.75	10.5	0.01	11,496	147.4	6.6
IT0005004426	indexed	10.5	0.01	10,570	95.3	25.4
IT0005022204	zero-cpn	1.0	0.001	807	69.6	0.9
IT0005023459	1.15	3.0	changing	3,016	108.5	5.1
IT0005024234	3.5	15.8	0.01	62,475	84.3	9.8
IT0005028003	2.15	7.5	0.01	19,295	170.6	5.0
IT0005030504	1.5	5.1	0.01	5,976	246.1	5.4
IT0005044976	zero-cpn	2.0	0.001	7,012	135.9	1.9
IT0005045270	2.5	10.3	0.01	49,800	136.9	4.5
IT0005056541	floating	6.2	0.01	17,654	192.9	5.6
IT0005058463	.75	3.3	0.01	4,174	134.9	5.4
IT0005058919	indexed	6.0	0.01	31,676	131.0	5.2
IT0005069395	1.05	5.0	0.01	17,646	184.1	4.3
IT0005083057	3.25	31.6	0.01	191,730	84.0	6.2
IT0005086886	1.35	7.2	0.01	22,041	128.9	4.9
IT0005090318	1.5	10.3	0.01	55,936	108.1	4.4
IT0005094088	1.65	17.0	0.01	102,275	65.0	6.2
IT0005105843	indexed	8.0	0.01	30,321	132.9	3.6
IT0005106049	.25	3.1	0.01	2,618	199.3	5.1
IT0005107708	.7	5.0	0.01	9,544	174.1	5.7
IT0005127086	2	10.3	0.01	7,420	81.6	4.0
IT0005135840	1.45	7.0	0.01	1,849	113.4	4.4

Chapter 4

Cross-impact and no-dynamic-arbitrage

4.1 Introduction

Market impact, i.e. the interplay between order flow and price dynamics, has increasingly attracted the attention of researchers and of the industry in the last years (Bouchaud et al. (2008)). Despite its importance, both from a fundamental point of view (due to its relation with supply-demand) and from an applied point of view (due to its relation with transaction cost analysis and optimal execution), market impact is not yet fully understood and different models and approaches have been proposed and empirically tested.

It is important to note that market impact refers to different aspects of this interplay and that they should be carefully distinguished (see Bouchaud et al. (2008) for a discussion). First, there is the impact of an individual trade or of the aggregated signed[1] order flow in a fixed time period. Second, especially for transaction cost analysis and optimal execution, it is more interesting to consider the impact of a large trade (sometimes termed as *meta-order*) executed incrementally by the same investor with many transactions and orders over a given interval of time. Both these definitions of market impact are typically investigated by considering one asset at a time, i.e. without considering the effect of a trade (or of an order) in one asset on the price dynamics of another asset.

This is the third type of impact, that we study in this chapter, and that is termed *cross-impact*. Understanding and modeling cross-impact is important for many reasons, since it enters naturally in problems like optimal execution of portfolios, statistical arbitrage of a set of assets, and to study the relation between correlation in prices and correlation in order flows. Conceptually, while self-impact, the impact of a trade on the price of the same asset, can qualitatively be understood as the result of a mechanical component (e.g. a market order with volume larger than the volume at the opposite best) and an induced component (resilience of the order book due to liquidity replenishment), the source of cross-impact is less clear. On one side if a trader is liquidating simultaneously two assets one can obviously expect a non-vanishing cross-impact. Since impact measures are typically averages across many measurements, this mechanism produces cross-impact if simultaneous trades and positively correlated order flow are frequently observed. On the other side, liquidity providers and arbitrageurs detect local mispricing between correlated assets and bet on a reversion to normality by placing orders. In other words this induced cross-impact relates to the possibility of identifying price changes due to local imbalances

[1] Conventionally buyer (seller) initiated trades have positive (negative) volume and order sign.

of supply-demand in one asset (rather than to fundamental information) and of exploiting the possibly short-lived mispricing between correlated assets.

Even though cross-impact has already been discussed e.g. in Almgren and Chriss (2001) as an extension of their optimal execution model and in Hasbrouck and Seppi (2001) in a principal component approach, it has only recently been the subject of extensive empirical studies. Pasquariello and Vega (2013) empirically show that order imbalance has a significant impact on returns across stocks and sectors at the daily scale. Wang et al. (2016a,b) present evidence for a structured price cross-response and correlated order flow at the intraday time-scale across stock pairs. Benzaquen et al. (2017) link cross-response and order flow in a multivariate extension of the Transient Impact Model (TIM) of Bouchaud et al. (2004) and show that their model can reproduce a significant part of the well-known correlation structure of asset returns. Mastromatteo et al. (2017) exploits this link between correlation and cross-impact, showing that cross-impact is crucial for a correct estimation of liquidity when trading portfolios. Wang and Guhr (2016) perform a scenario analysis in a model similar to Benzaquen et al. (2017), finding that cross-response is related both to cross-impact and correlated order flow across assets.

It is clear that the cross-impact problem talks naturally to dynamic arbitrage and to the possibility of price manipulation, as already discussed in Huberman and Stanzl (2004). It is therefore natural to ask which constraints the no-price-manipulation assumption imposes on market impact models. There is a large literature on this problem, often focused on the single asset case (Huberman and Stanzl (2004), Gatheral (2010), Alfonsi et al. (2012), Gatheral and Schied (2013a), Curato et al. (2016, 2017)).

In the multi-asset case many articles are concerned with strategies for optimal portfolio liquidation in the presence of volatility risk by expanding the model of Almgren (2003). Schied et al. (2010) show that optimal execution strategies for investors with constant absolute risk aversion are deterministic and for a more general absolute risk aversion setting Schöneborn (2016) finds that the optimal strategies for investors with different risk preferences vary only in the speed of their execution. The case of cross-impact in a lit market when there is also a dark pool is discussed in Kratz and Schöneborn (2015). Tsoukalas et al. (2017) instead develop a limit order book model with cross-impact and find that it can be optimal to temporarily take up positions contrary to the direction of one's trading intent. The paper most related to ours from a theoretical point of view is Alfonsi et al. (2016). They model multi-asset price impact by considering a linear version of the model of Gatheral (2010), extending thus the model already considered in Alfonsi et al. (2012). They show that the absence of no-dynamic-arbitrage on a discrete-time grid corresponds to the decay kernel being described by a positive definite matrix-valued function. Furthermore they formulate further conditions to ensure that resulting optimal strategies are well-behaved, both in discrete and continuous time, and show how such kernels can be constructed. However it is not generally straightforward to establish positive definiteness when a decay kernel is obtained coordinate-wise from estimations and therefore necessary conditions for the absence of dynamic arbitrage that can be verified on estimated decay kernels prove useful.

In this chapter, focusing on the TIM framework in continuous time, we establish some easily verifiable necessary conditions that must be satisfied by self- and cross-impact, in order to avoid the presence of price manipulation. We do this in the same spirit of Gatheral (2010) by explicitly constructing trading strategies that lead to price manipulation and negative expected cost. Some of these relations are simple generalizations to the multi-asset case of the corresponding relations for the single asset case derived in Gatheral (2010). Other relations that we derive here are instead genuinely relative to the multi-asset case.

In particular we formalize in Lemma 4.3.6 that cross-impact must be symmetric, i.e. the return induced in asset i by a trade of volume v in asset j must be equal to the impact of a trade of the same volume v in asset i on the price of asset j.

It is natural to ask whether this symmetry condition is empirically verified. In this chapter we study a market whose microstructure, to the best of our knowledge, has not been explored so far. This is the MOT market for sovereign bonds,[2] a fully electronic limit order book market for fixed-income assets. One of the reasons for our choice is that, due to the nature of the traded assets, we expect cross-impact, especially due to quote revisions, to be very high. In fact, two Italian fixed-rate BTPs differ mostly through the coupon rate and the time-to-maturity - factors which are accounted for in the price, which moves in a very synchronised way since for most purposes both titles are perfectly interchangeable.

Calibrating a multivariate TIM in trade time we find that there exist pairings of bonds where the symmetry condition of cross-impact is violated in a statistically significant way. By comparing the potential profit from a simple arbitrage strategy to transaction costs such as the bid-ask spread, which are neglected in the model, we conclude that arbitraging is not profitable. It is also crucial to point out that the empirical part of this work is important because it is the first application of a TIM model to fixed income markets and to the best of our knowledge it is the first work to consider cross-impact of single market orders and not the order sign imbalance aggregated over fixed time intervals (as done in Busseti and Lillo (2012), Benzaquen et al. (2017) and Wang et al. (2016a)).

The rest of the chapter is structured as follows. Section 4.2 introduces our model and the links to the no-dynamic-arbitrage principle. Section 4.3 discusses some general constraints on cross-impact that arise in our framework for bounded decay kernels and the corresponding proofs are given in Appendix 4.A. In Section 4.4 we study cross-impact empirically and compare to the theoretical results in Section 4.3. Finally Section 4.5 concludes.

4.2 Model setup

The presence of dynamic arbitrage depends on the market impact model. In this chapter we consider the Transient Impact Model (TIM) introduced in Bouchaud et al. (2004) (see Bouchaud et al. (2008) for a discussion). The model has been originally formulated in discrete time, and its continuous time version, that we present in the next section, has been proposed in Gatheral (2010).

4.2.1 Price process and cost of trading

Gatheral (2010) assumes that the asset price S_t at time t follows a random walk with a drift determined by the cumulative effect of previous trades

$$S_t = S_0 + \int_0^t f(\dot{x}_s)G(t-s)\mathrm{d}s + \int_0^t \sigma \mathrm{d}Z_s \tag{4.1}$$

where $f(\dot{x}_s)$ represents the (instantaneous) impact of trading at a rate $\dot{x}_s$ at time $s < t$ weighted by a decay kernel $G(\tau)$ with $\tau = t-s$. Z_s is a noise process, for example a Wiener process, and σ is the volatility. For consistency with equation (4.3) below, the trading rate $\dot{x}$ is given in units of number of shares per unit of time. In our multivariate extension we consider the prices of a set of assets where the drift in asset i not only depends on the

[2] `http://www.lseg.com/areas-expertise/our-markets/borsa-italiana/fixed-income-markets/mot`

trading history of asset i but also on past trades in assets $j \neq i$. Thus the price process of asset i is given by

$$S_t^i = S_0^i + \sum_j \int_0^t f^{ij}(\dot{x}_s^j) G^{ij}(t-s) \mathrm{d}s + \int_0^t \sigma^i \mathrm{d}Z_s^i \tag{4.2}$$

with a correlated noise process $\boldsymbol{Z}_s$ (e.g. a multivariate Wiener process) and where in addition to the *self-impact* terms f^{ii} and G^{ii} we have introduced additive *cross-impact* terms f^{ij} and G^{ij}, $i \neq j$, that represent the impact of trading in asset j on the price of asset i.

For a trading strategy $\Pi = \{\boldsymbol{x}_t\}$, $t \in [0,T]$, where $\boldsymbol{x}_t$ is the vector of asset positions x_t^i in asset i at time t, the expected cost (or implementation shortfall) is

$$\begin{aligned} C(\Pi) &= \mathbb{E}\left[\sum_i \int_0^T \dot{x}_t^i (S_t^i - S_0^i) \mathrm{d}t\right] \\ &= \sum_{i,j} \int_0^T \dot{x}_t^i \mathrm{d}t \int_0^t f^{ij}(\dot{x}_s^j) G^{ij}(t-s) \mathrm{d}s \end{aligned} \tag{4.3}$$

where as in Gatheral (2010) we consider only costs due to price impact, i.e. the price shift induced by our own trading, and neglect *slippage costs*, i.e. costs due to all other market frictions such as bid-ask spreads and trading fees. Also Alfonsi et al. (2016) ignore these costs by arguing that a sophisticated trading strategy consists not only of market orders but also of limit orders, thus on average both paying (earning) a half spread from market (limit) orders. In practice the constraints we derive may well be weakened when slippage costs are unavoidable, e.g. when immediacy requires to execute a trading strategy with market orders. In Section 4.3.7 we introduce a bid-ask spread and discuss its implications on our results.

4.2.2 Principle of no-dynamic-arbitrage

Huberman and Stanzl (2004) define a *round-trip trade* as a sequence of trades whose sum is zero, i.e. a trading strategy $\Pi = \{\boldsymbol{x}_t\}$ with

$$\int_0^T \dot{\boldsymbol{x}}_t \mathrm{d}t = \boldsymbol{0}\,. \tag{4.4}$$

This implies a round-trip in all assets traded in the strategy, i.e. $\int_0^T \dot{x}_t^i \mathrm{d}t = 0\ \forall\, i$. A *price manipulation* is a round-trip trade Π whose expected cost $C(\Pi)$ is negative and the *principle of no-dynamic-arbitrage* states that such a price manipulation is impossible. Formally, the principle requires that for any round trip trade Π it is

$$C(\Pi) \geq 0\,. \tag{4.5}$$

In the one-dimensional case this translates to

$$C(\Pi) = \int_0^T \dot{x}_t \mathrm{d}t \int_0^t f(\dot{x}_s) G(t-s) \mathrm{d}s \geq 0 \tag{4.6}$$

and imposes a relationship on the market impact function $f(\cdot)$ and the decay kernel $G(\cdot)$. The functions $f(\cdot)$ and $G(\cdot)$ are said to be *consistent* if they exclude the possibility of price manipulation. Several papers have studied the consistency of this kind of market impact models (for a review see Gatheral and Schied (2013a)). Gatheral et al. (2011) show that any decay kernel that is non-singular at time zero is inconsistent with non-linear $f(\cdot)$. Moreover Gatheral (2010) sets some necessary constraints for no arbitrage for power law dependence of f and G and Curato et al. (2017) show that inconsistencies can also arise for power-law $f(\cdot)$ and $G(\cdot)$ even when necessary conditions derived in Gatheral (2010) are not violated.

In the multidimensional case the cost requirement is

$$C(\Pi) = \sum_{i,j} \int_0^T \dot{x}_t^i \mathrm{d}t \int_0^t f^{ij}(\dot{x}_s^j) G^{ij}(t-s) \mathrm{d}s \geq 0 \tag{4.7}$$

and we are similarly looking at what forms of $\boldsymbol{f}(\cdot)$ and $\boldsymbol{G}(\cdot)$ are consistent when there is cross-impact.[3] Specifically we are asking what limits there are to cross-impact, i.e. the form of $f^{ij}(\cdot), i \neq j$, and whether the presence of cross-impact leads to possible arbitrages in pairs of $\boldsymbol{f}(\cdot)$ and $\boldsymbol{G}(\cdot)$ that are consistent in the one-dimensional case.

4.3 General constraints on cross-impact for bounded decay kernels

Let us for the remainder of this chapter assume, without loss of generality, that $\boldsymbol{G}(\tau)$ is a dimensionless quantity, i.e. all dimensionality of cross-impact, including the sign, is captured by the instantaneous market impact function $\boldsymbol{f}$.

In this section we also assume that the decay kernel $\boldsymbol{G}(\tau)$ is non-increasing, right-continuous at $\tau = 0$ and bounded in all components, i.e. that there exists an upper bound $U > 0$ so that $|G^{ij}(\tau)| < U$ for all $\tau \in [0, \infty)$ and all i, j. While we do not consider unbounded kernels in this section, we discuss in Appendix 4.B some constraints that arise for the popular class of pure power law kernels for cross-impact. The non-increasing assumption rules out non-zero decay kernels with $G^{ij}(0) = 0$ so that we are able to take $\boldsymbol{G}$ as normalized to 1 for its smallest lag $\tau = t - s$, i.e. $G^{ij}(0) = 1$ for all pairs ij.[4]A special case of such a kernel is exponential decay $G^{ij}(t-s) = e^{-\rho^{ij}(t-s)}$ as in Obizhaeva and Wang (2013).

In the following we consider the two-dimensional case, $i \in \{a, b\}$, i.e. the number of assets $N = 2$. Note that all results from the one-dimensional case still hold since we are free to choose a trading strategy that is active only in one asset, e.g. $\Pi = \left(\Pi^a, \Pi^b\right)^\intercal = (\Pi^a, 0)^\intercal$, $\forall\, t \in [0, T]$ where Π^a is a round-trip trading strategy in asset a.

In this chapter we assume some properties of the decay kernel G and then deduce from the absence of dynamic arbitrage properties of f. An interesting alternative is to fix f and deduce the properties of G. Some results exist when f is linear, both in the single asset case (see Proposition 5 in Gatheral and Schied (2013b)) and in the multi-asset case in a discrete-time framework (see Proposition 1 of Alfonsi et al. (2016)). When f is non-linear the problem of existence of no-arbitrage conditions is still open (see Curato et al. (2017)).

The proofs of the results in this section are given in Appendix 4.A and are obtained following the approach of Gatheral et al. (2011). All results can also be obtained following

[3]Alfonsi et al. (2016) consider a slightly different case where instead of a round-trip strategy, they consider the liquidation of an existing portfolio. This is equivalent in the limit of building up the portfolio infinitely slowly and when impact is purely transient.

[4]The non-increasing assumption is necessary already in the single-asset case to avoid arbitrage opportunities from simple buy-hold-sell strategies. In the multi-asset case we require it e.g. for our symmetry result in Lemma 4.3.6.

Gatheral (2010) under the slightly more restrictive assumptions of the decay kernel $\boldsymbol{G}$ being representable as a suitable series expansion and considering only the first non-zero orders in the limit of $\tau \to 0^+$.

4.3.1 A simple strategy with two assets

In the following we will often make use of a simple strategy in two assets which is split into two phases of trading at constant rates.

Example 4.3.1. A simple in-out strategy.
At first we build up a position at a constant trading rate from time 0 until time Θ, with $0 < \Theta < T$, and then liquidate the position in a second phase from Θ until T.

$$\Pi = \{\boldsymbol{x}_t\} \quad , \quad \dot{\boldsymbol{x}}_t = \begin{cases} (v_{a,\mathrm{I}}, v_{b,\mathrm{I}})^\intercal & \text{for} \quad 0 \leq t \leq \Theta \\ (v_{a,\mathrm{II}}, v_{b,\mathrm{II}})^\intercal & \text{for} \quad \Theta < t \leq T \end{cases} . \tag{4.8}$$

The velocities $v_{i,\mathrm{I}}$, $v_{i,\mathrm{II}}$ are constrained by our choice of the strategy. Since Π is a round-trip strategy, the trading rates $v_{i,\mathrm{I}}$ and $v_{i,\mathrm{II}}$ have opposite signs, i.e. $\kappa = v_{i,\mathrm{I}}/v_{i,\mathrm{II}} < 0$, and the time Θ when the trading direction changes is given as $\Theta = \frac{-v_{i,\mathrm{II}}}{v_{i,\mathrm{I}}-v_{i,\mathrm{II}}}T = \frac{1}{1-\kappa}T$. Let us further fix notation with $\lambda = v_{a,\mathrm{I}}/v_{b,\mathrm{I}} = v_{a,\mathrm{II}}/v_{b,\mathrm{II}}$. Figure 4.1a illustrates a possible realization of this strategy with $\lambda < 0$.

The cost of this strategy can be decomposed as $C(\Pi) = \sum_{i,j=a,b} C_A^{ij} + C_B^{ij} + C_C^{ij}$ where

$$\begin{aligned} C_A^{ij} &= v_{i,\mathrm{I}} f^{ij}(v_{j,\mathrm{I}}) \int_0^\Theta \mathrm{d}t \int_0^t G^{ij}(t-s)\mathrm{d}s \\ C_B^{ij} &= v_{i,\mathrm{II}} f^{ij}(v_{j,\mathrm{I}}) \int_\Theta^T \mathrm{d}t \int_0^\Theta G^{ij}(t-s)\mathrm{d}s \\ C_C^{ij} &= v_{i,\mathrm{II}} f^{ij}(v_{j,\mathrm{II}}) \int_\Theta^T \mathrm{d}t \int_\Theta^t G^{ij}(t-s)\mathrm{d}s \end{aligned} \tag{4.9}$$

In the one-dimensional case the principle of no-dynamic-arbitrage imposes a constraint on the term C_B^{ii} and from equation (4.7) it follows that $-C_B^{ii} \leq C_A^{ii} + C_C^{ii}$ as in Gatheral (2010). For the multi-dimensional case $C(\Pi) \geq 0$ further implies a relationship between the strength of cross-impact and self-impact.

In the following we will try to exploit cross-impact in order to push down the cost of strategies. Supposing that cross-impact is positive for positive trading rates, i.e. $f^{ij}(v) > 0$ for $v > 0$, we can choose $\lambda < 0$, e.g. trading into asset a while contemporaneously trading out of asset b, in order to get a negative contribution from cross-impact.

4.3.2 Cross-impact as odd function of the trading rate

In the one-dimensional case Gatheral (2010) shows that permanent market impact needs to be an odd function in the rate of trading v, i.e. $f(v) = -f(-v)$. We show here that the same holds for cross-impact for decay kernels that are non-singular around $\tau \to 0^+$.

Lemma 4.3.2. *Assume a price process as in (4.2) with a bounded, non-increasing decay kernel $\boldsymbol{G}$ that is continuous around $\tau = 0$. Then such a model admits price manipulation if $\boldsymbol{f}$ is not an odd function of the trading rate, i.e. unless*

$$f^{ij}(v) = -f^{ij}(-v) \qquad \forall\, i, j\,. \tag{4.10}$$

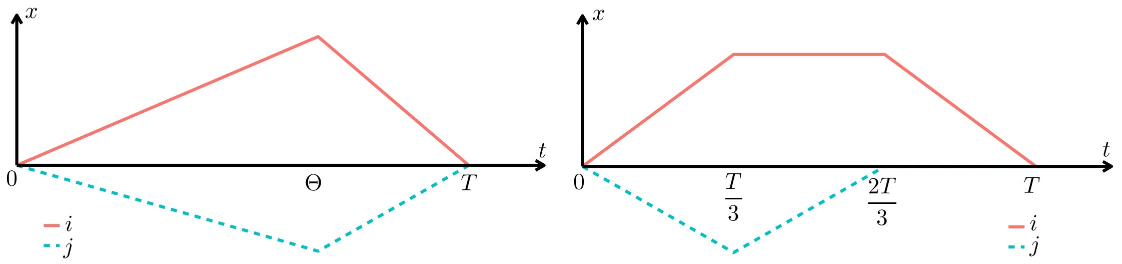

(a) Strategy as in Example 4.3.1 **(b)** Strategy as in Example 4.3.5

Figure 4.1: Schematic of the trading strategies in Example 4.3.1 (left panel) and Example 4.3.5 (right panel).

We will use equation (4.10) for the remainder of this chapter. As a corollary it follows that

Corollary 4.3.3. *Absence of dynamic-arbitrage for a price process as in (4.2) with a decay kernel that is bounded, non-increasing and continuous around* $\tau = 0$ *requires that*

$$f^{ij}(0) = 0 \qquad \forall\, i, j\,. \tag{4.11}$$

4.3.3 Constraints on the strength of cross-impact

The cost constraint in equation (4.7) also imposes a constraint on the relative strength of f^{ij}. Let us consider a simple example at first.

Example 4.3.4. Trading in and out at the same rate

We consider a strategy as above in Example (4.3.1) where we are trading in and out of positions at the same rate, i.e. $v_{i,\mathrm{I}} = -v_{i,\mathrm{II}}$ and therefore $\Theta = T/2$, but in different directions in the two assets, choosing e.g. $v_{a,\mathrm{I}} = v_a > 0, v_{b,\mathrm{I}} = -v_b < 0$ and thus $\lambda < 0$. For simplicity let us assume a uniform decay of market impact, i.e. $G^{ij}(t) = G(t)$ for all pairings ij. The cost is then

$$\begin{aligned} C(\Pi) = &\Big[v_a f^{aa}(v_a) + v_b f^{bb}(v_b) - v_a f^{ab}(v_b) - v_b f^{ba}(v_a)\Big] \\ &\left\{\int_0^{T/2} \mathrm{d}t \int_0^t \left[G(t-s) - G(t+T/2-s)\right] \mathrm{d}s + \int_{T/2}^{T} \mathrm{d}t \int_{T/2}^{t} \left[G(t-s) - G(T-s)\right] \mathrm{d}s\right\} \end{aligned} \tag{4.12}$$

and Gatheral (2010) shows that the term in curly brackets in equation (4.12) is greater than zero when further requiring that $G(\cdot)$ is strictly decreasing. Thus the no-dynamic-arbitrage constraint (4.5) requires that

$$v_a f^{aa}(v_a) + v_b f^{bb}(v_b) - v_a f^{ab}(v_b) - v_b f^{ba}(v_a) \geq 0 \tag{4.13}$$

for any $v_a, v_b \geq 0$, thus constraining the relative size of the cross-impact terms f^{ab} and f^{ba} with respect to self-impact. Note that by setting $v_b = 0$ we recover the one-dimensional case and it follows that $v_a f^{aa}(v_a) \geq 0$.

In the general case, the decay $G^{ij}(\tau)$ is not uniform and we can not factor out the term in curly brackets in equation (4.12), instead we have to weight each of the terms of equation (4.13) with a factor that depends on the decay $G^{ij}(\tau)$. Furthermore we are free to choose a strategy with different trading rates as in Example 4.3.1 or a more sophisticated strategy. Alfonsi et al. (2016) consider this problem in discrete time with linear instantaneous price impact. Their Proposition 2.6 states that absence of arbitrage in the sense of equation (4.7) is equivalent to the condition that the elementwise product of strength of impact and the decay kernel corresponds to a positive definite matrix-valued function.

4.3.4 Symmetry of cross-impact

Let us assume that $\boldsymbol{G}(\tau)$ is bounded, non-increasing and continuous around $\tau = 0$. We will show that in this case impact needs to be symmetric, i.e. $v_i f^{ij}(v_j) = v_j f^{ji}(v_i)$, in order to avoid price manipulations. For illustration we first consider an example where impact is linear and permanent.

Example 4.3.5. An asymmetric strategy with purely permanent and linear impact. Suppose market impact is linear and permanent, i.e. $f^{ij}(v) = \eta^{ij} v$ and $G^{ij}(\tau) = 1 \,\forall\, i, j$. Then the cost of trading in the single-asset case only depends on the initial and final positions x_0 and x_T. If there is cross-impact between two or more assets, there is also an interaction term between the trading rates in different assets, i.e.

$$\begin{aligned} C(\Pi) &= \sum_{i,j} \int_0^T v_{i,t} \mathrm{d}t \int_0^t \eta^{ij} v_{j,s} \mathrm{d}s \\ &= \sum_i \frac{\eta^{ii}}{2}(x_T^i - x_0^i)^2 + \sum_{j \neq i} \eta^{ij} \int_0^T v_{i,t} \mathrm{d}t \int_0^t v_{j,s} \mathrm{d}s \end{aligned} \tag{4.14}$$

and while the first sum with the self-impact terms disappears for a round-trip strategy since $\boldsymbol{x}_0 = \boldsymbol{x}_T$, this is not generally the case for the second sum with the terms due to cross-impact. To see this, let us consider a different round-trip strategy Π in two assets, which is now asymmetric and lasts over three phases:

$$v_{a,t} = \begin{cases} v_a & \text{for} \quad 0 \le t \le T/3 \\ 0 & \text{for} \quad T/3 < t \le 2T/3 \\ -v_a & \text{for} \quad 2T/3 < t \le T \end{cases} , \quad v_{b,t} = \begin{cases} -v_b & \text{for} \quad 0 \le t \le T/3 \\ v_b & \text{for} \quad T/3 < t \le 2T/3 \\ 0 & \text{for} \quad 2T/3 < t \le T \end{cases} , \tag{4.15}$$

with $v_a, v_b > 0$ and as illustrated in Figure 1(b). While self-impact cancels out when we calculate $C(\Pi)$, the asymmetry in our strategy makes for a non-trivial total cost that stems from cross-impact:

$$C(\Pi) = v_a v_b \frac{T^2}{18} (\eta^{ba} - \eta^{ab}). \tag{4.16}$$

If $\eta^{ba} < \eta^{ab}$ this gives a negative cost and likewise when $\eta^{ba} > \eta^{ab}$ by interchanging assets $a \leftrightarrow b$ in the strategy (4.15). Therefore it follows that cross-impact needs to be symmetric with respect to asset pairs in order to exclude arbitrage opportunities, as observed in Huberman and Stanzl (2004).

In fact we can expand this result to the transient impact case for general cross-impact functions f^{ij}:

Lemma 4.3.6. *If decay of market impact $\boldsymbol{G}(\tau)$ is bounded, non-increasing and continuous around $\tau = 0$, absence of dynamic arbitrage requires that*

$$v_i f^{ij}(v_j) = v_j f^{ji}(v_i) \quad \forall\, i, j \,. \tag{4.17}$$

4.3.5 Linearity of market impact

Gatheral et al. (2011) finds that any single-asset market impact model as in equation (4.1) is inconsistent when f is non-linear and G is bounded and non-increasing. We expand this proposition to the multi-asset case with cross-impact, i.e.

Lemma 4.3.7. *Assuming a price process as in (4.2) with a bounded, non-increasing decay kernel* $\boldsymbol{G}$ *that is continuous around* $\tau = 0$ *and a non-linear market impact function* $\boldsymbol{f}$*. Then such a model admits price manipulation.*

As a corollary of Lemma 4.3.7 we can extend Lemma 4.1 in Gatheral (2010) for self-impact to the case with cross-impact.

Corollary 4.3.8. *A price process as in (4.2) with self- and cross-impact that decays exponentially at different rates and instantaneous price impact that is non-linear, admits price manipulations.*

We obtain the same corollary for purely permanent impact by taking the limit $\rho^{ij} \to 0^+$ in the case of exponential decay $G^{ij}(t-s) = e^{-\rho^{ij}(t-s)}$, as already observed in Huberman and Stanzl (2004):

Corollary 4.3.9. *Nonlinear permanent self- and cross-asset market impact is inconsistent with the principle of no-dynamic-arbitrage.*

Let us reconsider Example 4.3.4 taking into account linearity and symmetry of cross-impact as shown above. In this case Equation (4.13) simplifies to

$$v_a^2\eta^{aa} + v_b^2\eta^{bb} - 2v_a v_b \eta^{\text{cross}} \geq 0 \tag{4.18}$$

and minimizing the cost constrains the strength of cross-impact $\eta^{\text{cross}} = \eta^{ab} = \eta^{ba}$ as

$$\eta^{\text{cross}} \leq \sqrt{\eta^{aa}\eta^{bb}}\,, \tag{4.19}$$

in agreement with Proposition 3.7.(b) of Alfonsi et al. (2016) and equivalent to the condition for a symmetric 2×2 matrix to be positive-semidefinite.

4.3.6 Exponential decay

The conditions of linearity and symmetry of cross-impact are necessary for absence of arbitrage, but are they also sufficient?

Example 4.3.10. An asymmetric strategy with symmetric, exponentially decaying linear impact.

Let us re-consider the strategy in Eq. (4.15) with exponentially decaying impact $G^{ij}(t-s) = e^{-\rho^{ij}(t-s)}$ and a linear instantaneous impact function $f^{ij}(v) = \eta^{ij}v$ that is now symmetric with $\eta^{ab} = \eta^{ba} = \eta^{\text{cross}}$. The cost terms for self-impact are now

$$\begin{aligned} C^{aa} &= \frac{\eta^{aa}v_a^2}{(\rho^{aa})^2}\left[-e^{-\rho^{aa}T} + 2e^{-2\rho^{aa}T/3} + e^{-\rho^{aa}T/3} - 2 + \frac{2\rho^{aa}T}{3}\right] \\ C^{bb} &= \frac{\eta^{bb}v_b^2}{(\rho^{bb})^2}\left[-e^{-2\rho^{bb}T/3} + 4e^{-\rho^{bb}T/3} - 3 + \frac{2\rho^{bb}T}{3}\right]\,. \end{aligned} \tag{4.20}$$

and likewise for cross-impact

$$\begin{aligned} C^{ab} &= \frac{\eta^{\text{cross}}v_a v_b}{(\rho^{ab})^2}\left[-\frac{\rho^{ab}T}{3} + 2e^{-\rho^{ab}T/3} - 3e^{-2\rho^{ab}T/3} + e^{-\rho^{ab}T}\right] \\ C^{ba} &= \frac{\eta^{\text{cross}}v_a v_b}{(\rho^{ba})^2}\left[2 - \frac{\rho^{ba}T}{3} - 3e^{-\rho^{ba}T/3} + e^{-2\rho^{ba}T/3}\right]\,. \end{aligned} \tag{4.21}$$

When we develop the terms in squared brackets in (4.20) and (4.21) in powers of $\rho^{ij}T$ all terms of order $(\rho^{ij}T)^0$ and $(\rho^{ij}T)^1$ cancel out, while terms proportional to $(\rho^{ij}T)^2$ sum to 0 thanks to the symmetry of instantaneous cross-impact. The cost to the first non-zero order of $\rho^{ij}T$ is then

$$C(\Pi) = \frac{T^3}{6}\left[\frac{2}{3}\eta^{aa}v_a^2\rho^{aa} + \frac{4}{27}\eta^{bb}v_b^2\rho^{bb} - \frac{5}{27}\eta^{\text{cross}}v_a v_b\left(\rho^{ab}+\rho^{ba}\right)\right] + \sum_{i,j}\mathcal{O}\left((\rho^{ij})^2T^4\right) \tag{4.22}$$

and when ρ^{ab} or ρ^{ba} is large enough[5] compared to the other terms the cost can still be negative. For absence of price manipulations we therefore also require further constraints on the speed of decay described by $\boldsymbol{G}$.

This is in agreement with the results of Alfonsi et al. (2016) in discrete time. Their Proposition 3.7 proves that the conditions of symmetry $\eta^{ab} = \eta^{ba}$, and a non-increasing decay kernel, i.e. $\min(\rho^{ab}, \rho^{ba}) \geq \frac{1}{2}(\rho^{aa}+\rho^{bb})$ and $\frac{1}{4}(\eta^{ab}\rho^{ab}+\eta^{ba}\rho^{ba})^2 \leq \eta^{aa}\rho^{aa}\eta^{bb}\rho^{bb}$, are sufficient for the absence of arbitrage. We complement this result in Lemma 4.3.6 by showing that symmetry $\eta^{ij} = \eta^{ji}$ is indeed necessary for any decay kernel $\boldsymbol{G}(\tau)$ that fulfills our conditions of being bounded, non-increasing and continuous around $\tau = 0$. Note that this excludes kernels where for some ij $G^{ij}(0) = 0$ but $G^{ij}(\tau) > 0$ for some $\tau > 0$. Indeed Example 3 in Alfonsi et al. (2016) has a kernel that is asymmetric for $\tau > 0$ and which does not allow price manipulation.

4.3.7 Symmetry and bid-ask spread

As pointed out in section 4.2.1 we have so far neglected slippage costs. Here we relax this assumption for the case of a trader that can not ignore the bid-ask spread. Let us re-consider example 4.3.5, now with a constant bid-ask spread B_i.

Example 4.3.11. An asymmetric strategy with purely permanent and linear impact and constant bid-ask spread.
Suppose market impact is linear and permanent, i.e. $f^{ij}(v) = \eta^{ij}v$ and $G^{ij}(\tau) = 1\,\forall\, i,j$. A trader pays half the bid-ask spread both when buying and selling an asset, and therefore the cost of a strategy Π reads

$$\begin{aligned} C(\Pi) &= \sum_{i,j}\int_0^T v_{i,t}\mathrm{d}t\int_0^t \eta^{ij}v_{j,s}\mathrm{d}s + \sum_i\int_0^T \frac{1}{2}\left|v_{i,t}\right|B_i\mathrm{d}t \\ &= \sum_i \frac{\eta^{ii}}{2}(x_T^i - x_0^i)^2 + \sum_{j\neq i}\eta^{ij}\int_0^T v_{i,t}\mathrm{d}t\int_0^t v_{j,s}\mathrm{d}s + \sum_i\int_0^T\frac{1}{2}\left|v_{i,t}\right|B_i\mathrm{d}t. \end{aligned} \tag{4.23}$$

We use the same asymmetric round-trip strategy of equation (4.15). Only the self-impact term cancels out, with the remaining terms being due to cross-impact and the bid-ask spread respectively:

$$C(\Pi) = v_a v_b\frac{T^2}{18}(\eta^{ba}-\eta^{ab}) + (v_aB_a + v_bB_b)\frac{T}{3}. \tag{4.24}$$

If $\eta^{ba} < \eta^{ab}$ the first term is negative and likewise if $\eta^{ba} > \eta^{ab}$ by interchanging assets $a \leftrightarrow b$ in the strategy (4.15). Since this negative cost due to cross-impact scales $\sim v^2$ and the cost due bid-ask spread only $\sim v$, we can always choose v_a, v_b (or T) large enough so that $C(\Pi) < 0$. It follows that cross-impact needs to be symmetric also when there is a constant bid-ask spread.

[5] While keeping the product of $\rho^{ij}T$ small for all ij.

We can then generalize Lemma 4.3.6 for the presence of a constant bid-ask spread.

Corollary 4.3.12. *Assume a price process as in (4.2) with decay of market impact $\boldsymbol{G}(\tau)$ that is bounded, non-increasing and continuous around $\tau = 0$, $f^{ij}(v) = \eta^{ij} v$ that is linear and has the expected cost*

$$C(\Pi) = \sum_{i,j} \int_0^T \dot{x}_t^i \mathrm{d}t \int_0^t f^{ij}(\dot{x}_s^j) G^{ij}(t-s) \mathrm{d}s + \sum_i \int_0^T \frac{1}{2} \left| \dot{x}_t^i \right| B_i \mathrm{d}t \tag{4.25}$$

with B_i the constant bid-ask spread of asset i. Then such a model admits price manipulation unless cross-impact is symmetric, i.e. unless $\eta^{ij} = \eta^{ji}$ for all i, j.

Here we have assumed a constant bid-ask spread, which is a good approximation for large-tick assets, where the bid-ask spread is equal to the minimum tick size most of the time. But even where that is not the case, it suffices to assume that there is an upper bound on the bid-ask spread and re-interpret B_i as this upper bound. Our argument then still holds, as we are always able to choose v_a, v_b large enough so that spread costs are outweighed by gains from cross-impact.

Being able to choose the speed of trading v large enough is also the more critical assumption from an applied point of view. In practice, e.g. for too high a value for the trading rate v, linear impact may no longer be realistic, as liquidity in the limit order book would be consumed faster than being replenished and additional price impact costs would arise. Let us consider a simple case of example 4.3.11 where we set $T = 1$, $B_a = B_b = B$ and $v_a = v_b = v$ and we denote $\Delta\eta = \eta^{ab} - \eta^{ba}$ the cross-impact asymmetry, which we assume positive for fixing ideas. The no-arbitrage condition can be rewritten as

$$-\frac{v\Delta\eta}{12} + B \geq 0\,. \tag{4.26}$$

This inequality sets the maximal trading speed for the absence of arbitrage, when the spread and the asymmetry of impact are given. In fact it can be rewritten as $v \leq 12B/\Delta\eta$. Therefore, when we discuss empirical asymmetries of cross-impact in section 4.4.5, we also consider upper bounds for v_i. Similarly, when the spread and the maximal trading velocity $v_{\max}$ are given, the maximal asymmetry of cross-impact is $\Delta\eta \leq 12B/v_{\max}$. Finally, given the maximal trading speed and the cross-impact asymmetry, the spread must be larger than $B \geq \Delta\eta v_{\max}/12$.

4.4 Empirical evidence of cross-impact

4.4.1 Market structure of MOT

For the empirical analysis we consider Italian sovereign bonds traded on the retail platform 'Mercato telematico delle obbligazioni e dei titoli di Stato' (MOT). We choose to estimate cross-impact between bonds instead of equities since we expect the strength of cross-impact among sovereign bonds of the same issuing country, especially of similar maturity, to be bigger than the one between e.g. stocks or indices. Sovereign bonds of one country typically have a very similar underlying risk and their prices are implicitly connected via the yield curve, a link that we deem stronger than e.g. a common factor between stocks of the same sector.

The secondary market for European sovereign bonds is divided into an opaque over-the-counter market (OTC) and an observable exchange-traded market. The Italian securities and exchange Commission CONSOB publishes a bi-annual report listing the share in trading of Italian government bonds separated per trading venue.[6] For the year 2014 (2015) the share of OTC trading has been 58.8% (59.1%), while 45.6% (44.8%) of trading on platforms took place on the inter-dealer platform MTS. MOT is the third-largest platform by traded value with 8.7% (8.8%) of traded value excluding the OTC market in 2014 (2015). Most of the literature for the Italian and European government bonds market focuses on MTS, with the exception of Linciano et al. (2014) who compare the liquidity of dual-listed corporate bonds across MOT and the EuroTLX platform. Darbha and Dufour (2013) review the market microstructure of MTS in the context of the market for European sovereign bonds and discuss several liquidity measures based on the limit order book, trades or bond characteristics. They note that MTS 'normally has a few trades per bond per day, even for the most liquid government bonds'. Indeed, due to large minimum sizes, for most titles there is on average less than one transaction per day on MTS, making studies of market impact difficult. Dufour and Nguyen (2012) overcome this issue by building impulse response functions from regressions of returns on order flow at 10 second intervals to study permanent market impact. In a different approach Schneider et al. (2018) use a measure of (virtual) mechanical price impact along with other liquidity measures calculated from the limit order book to detect illiquidity shocks that can be modeled as a self- and cross-exciting Hawkes process in and across Italian sovereign bonds.

Instead in this chapter we focus on MOT where we observe a sufficient number of (smaller) trades as well as an active limit order book. Italian government bonds are traded on the DomesticMOT segment of MOT where the trading day is divided into an opening auction from 8:00 to 9:00 followed[7] by a phase of continuous trading until 17:30. If certain price limits are violated during the continuous trading, a volatility auction phase is initiated for a duration of 10-11 minutes. MOT is organized as a continuous double auction where besides market and limit orders also partially hidden 'iceberg orders', 'committed cross' orders and 'block trade facilities' are allowed. While the presence of a specialist or a bid specialist is possible, in practice this is only the case for a subset of financial sector corporate bonds not in our sample. The tick size depends on the residual lifetime and is 1 basis point of nominal size or 0.1 basis points if the residual lifetime is less or equal than two years, corresponding to 1 or 0.1 euro cents respectively.

Our dataset contains all trades and limit order book (LOB) snapshots[8] for a selection of 60 ISINs from December 1, 2014 to February 27, 2015 and April 13, 2015 to October 16, 2015 for a total of 194 trading days. For the remainder of this chapter we will focus on a set of $N = 33$ fixed rate or zero-coupon Italian sovereign bonds listed in Appendix 4.C with at least 5,000 trades throughout our sample to ensure sufficient liquidity and statistical significance of our results. To avoid intraday seasonalities we further restrict our data to 10:00 - 17:00 and discard observations when we detect a volatility auction. The average spread is smaller than 10 ticks for most of the bonds with the exception of some very long-term bonds and bonds where the tick size is 0.1 basis points. More than 92% of the orders in our sample are executed at the corresponding best bid or ask quote[9] and thus identified as sell or buy orders respectively, while all other orders are classified according

[6]CONSOB, Bollettino Statistico Nr. 8, March 2016, available at `http://www.consob.it/web/area-pubblica/bollettino-statistico`

[7]The conclusion of contracts from the opening auction happens at a random time between 09:00:00-09:00:59.

[8]In phases of heavy trading multiple updates of the LOB may be recorded as one update in our data. However there is at least one update per second whenever there are changes to the LOB and in the vast majority of our sample updates are more frequent.

[9]The remaining $\sim 8\%$ can either be due to orders that were executed across more than one millisecond (so that they are recorded as two or more orders), missed LOB updates or exotic order types.

to the algorithm of Lee and Ready (1991).

Let us fix notation for the estimations in the following sections. We consider the log-price $X_t^i = \log(S_t^i)$ of the mid-price of the best bid and ask quote for asset i at time t and calculate the return $r^i_{t,t+\Delta t}$ from time t to time $t+\Delta t$ as $r^i_{t,t+\Delta t} = X^i_{t+\Delta t-\varepsilon} - X^i_{t-\varepsilon}$ for $\varepsilon \to 0^+$. ϵ_t^i is the sign of a trade (market order) and $+1$ for a buyer-initiated transaction, -1 for a sell, and undefined when there is no trade in the asset i at time t. I_t^i is an indicator function that is $+1$ when there is a trade in asset i at time t and 0 otherwise and we consider the product of an undefined trade sign with a 0 indicator function to be 0 such that the product $\epsilon_t^i I_t^i$ is always defined and one of $\{-1, 0, +1\}$. The size of a trade V_t^i is given as its nominal value in EUR and the price is reported per one asset (or contract) with a face value of 100 EUR. Unlike e.g. Benzaquen et al. (2017) we do not de-mean the order sign in order to avoid attributing a price impact to the absence of transactions in a bond in the sense of Corollary 4.3.3. However we have verified that our results are qualitatively similar when considering de-meaned order signs ϵ or ϵI and de-meaned returns.

4.4.2 Response function

We define the self- and cross-response function $R^{ij}_{\Delta t}$ as the unconditional Δt-ahead return in asset i controlled for the order sign of asset j, i.e.

$$R^{ij}_{\Delta t} = \mathbb{E}\left[\left(X^i_{t+\Delta t-\varepsilon} - X^i_{t-\varepsilon}\right)\epsilon_t^j I_t^j\right]. \tag{4.27}$$

For $i = j$ we will speak of *self-response* and of *cross-response* for $i \neq j$. Figure 4.2 shows the average self- and cross-response function for all bonds in our sample and their pairings respectively. For positive lags Δt we find that self-response is on average larger than cross-response by a factor of ~ 5, consistent with observations of Benzaquen et al. (2017); Wang et al. (2016a,b). $R^{ij}_{\Delta t=0}$ is zero by definition, whereas for small negative Δt we find that R^{ij} is on average positive, producing a cusp at $\Delta t = 0$. We conjecture that such behavior is not observed in Benzaquen et al. (2017) because of the rather large time lag of 5 minutes, corresponding to ~ 80 units of transaction time in Figure 4.2. In the single asset case this feature is clearly present for the large-tick stock Microsoft in Figure 1 of Taranto et al. (2016). As shown there, the kink could be related to correlations of market order flow with past returns and indicates a forecasting power of current returns on the future order sign imbalance. Interestingly we find that the cross-response measured at negative lags is smaller (i.e. larger in absolute value) than self-response, contrary to the observations in Benzaquen et al. (2017).[10] The figure also shows the prediction from the model of the negative lag impact (see Taranto et al. (2016) for details). We observe a clear difference with the empirical data suggesting also for cross-impact a reaction of order flow to past price dynamics of other bonds.

4.4.3 Instantaneous market impact

We measure the instantaneous market impact function $\boldsymbol{f}(\cdot)$ as

$$f^{ij}(V) = \mathbb{E}\left[r^i_{t-\varepsilon,t+2s}\epsilon_t^j | I^j_{t,V} = 1\right] \tag{4.28}$$

which is the expected return in asset i from just before a trade at time t until 2 seconds after t, multiplied by the trade sign in asset j at time t and conditional on a trade in asset

[10] We suppose that this is related to the fact that many of the bonds considered here are easily substitutable for one another.

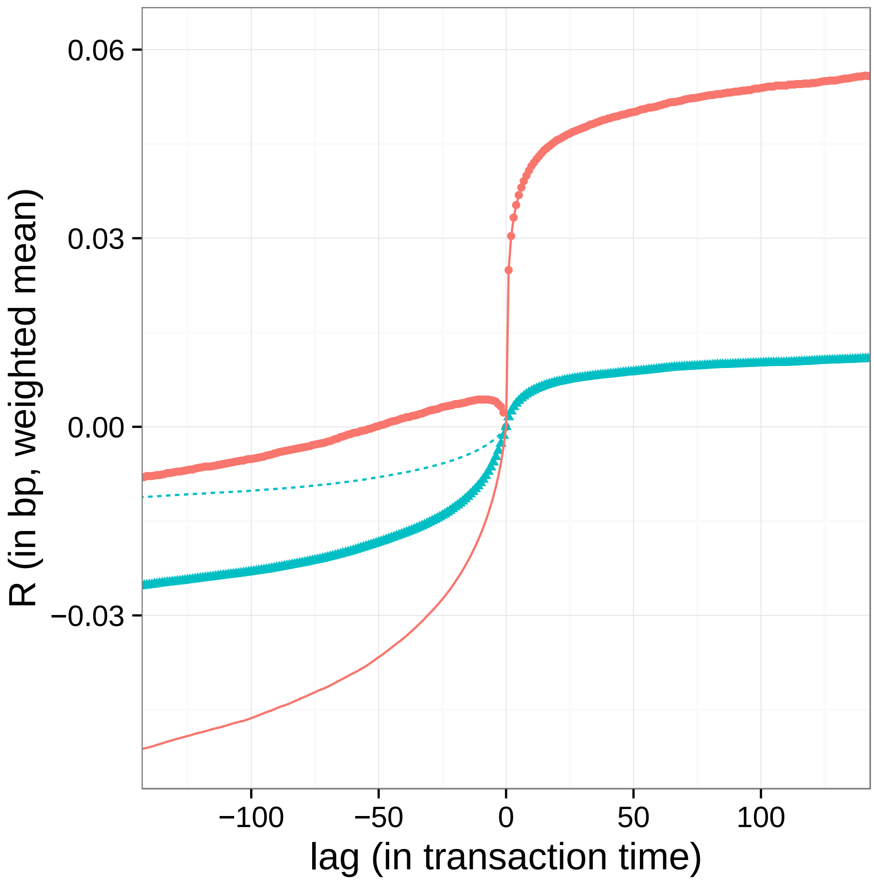

Figure 4.2: Plot of average self- and cross- response function $R^{ij}_{\Delta t}$ in transaction time as defined in Section 4.4.4. Mean over all bonds and pairings in our sample and weighted by the number of trades in the triggering bond j. Self-response is shown as red dots connected by solid lines, cross-response as blue triangles connected by dashed lines. The lines correspond to the prediction from the model in Section 4.4.4.

j at time t of size V. We have chosen the two second interval as twice the maximum time between two updates of the limit order book, i.e. we can rule out that changes in the book were not reported in our data. For measurement purposes we bin similar trade sizes together, with the bin size chosen as a function of the number of trades in the triggering bond j.

Figure 4.3 shows self- and cross-impact between all bonds in our sample as a function of trade size V measured in units of face value. Cross-impact is universally present across our sample and on average smaller than self-impact by roughly one order of magnitude. The cross-impact curves of different pairings ij are very close one to the other when both bonds have a time-to-maturity of at least four years left. For bonds with three or less years left until maturity we do not observe an intense trading activity, thus the curves in the leftmost column in the figure are very noisy. Likely the price-dynamics of these short-term titles are more decoupled from the medium- and long-term bonds with a lifespan of four

[11] In principle this observation suggests the presence of arbitrage opportunities due to the violation of Lemma 4.3.7. However we should remember that what is shown in Figure 4.3 is the observed impact, which might be different from the virtual impact, since the former does not take into account the selection bias due to the fact that traders condition the market order volume to what is present at the opposite best. For a discussion of this point in the self-impact case, see Bouchaud et al. (2008).

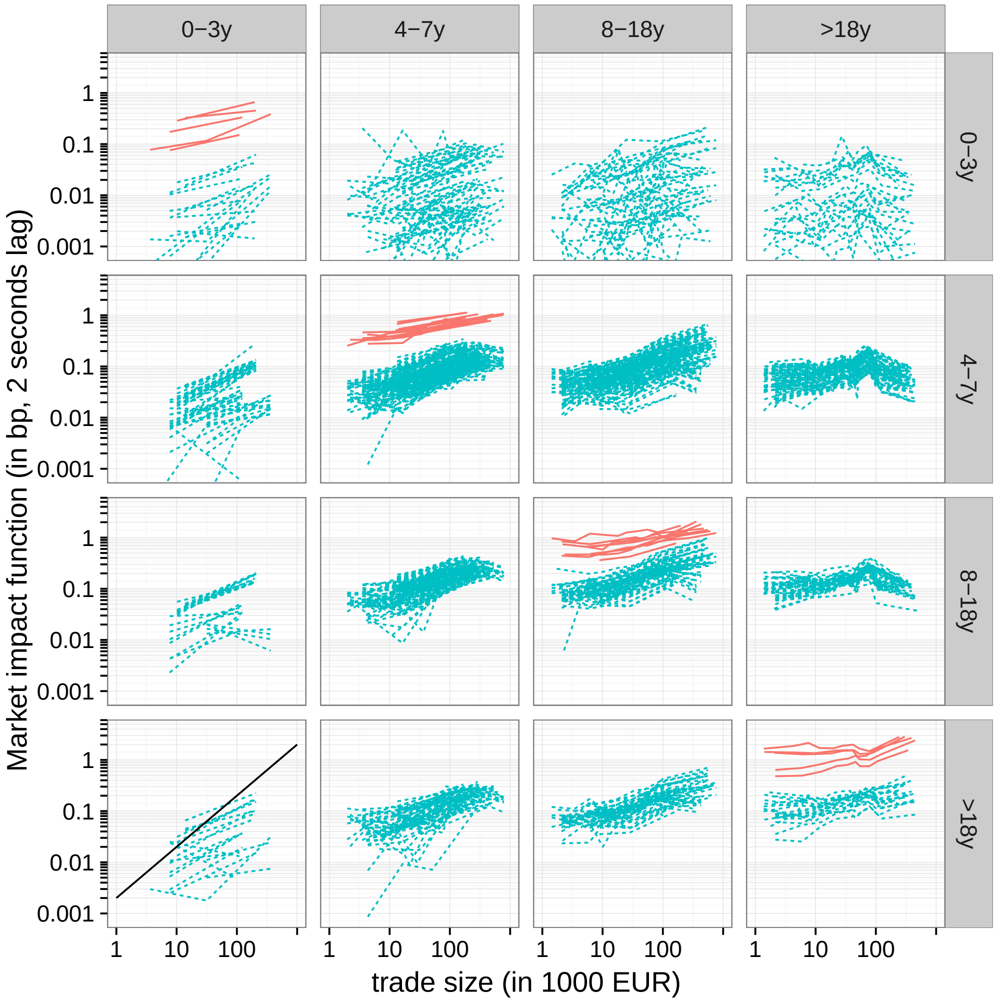

Figure 4.3: Plot of the average self- and cross-impact function among all pairs of bonds in our sample as a function of trade size V measured in units of face value. Each line corresponds to one pairing ij, grouped by time-to-maturity into four categories, where impact is from the column on the row. Self-impact is shown in the diagonal panels as red solid lines, cross-impact is shown as blue dashed lines and present in all panels. Price impact is calculated as average price change (multiplied by the trade sign) after a lag of 2 seconds, the minimum time that ensures we observe an update of the limit order book. Self- and cross-impact is clearly non-linear. For comparison the solid black line in the lower left panel illustrates a linear impact function.

or more years. The figure shows that all the estimated functions $f^{ij}(V)$ are non-linear, being concave and well described by a power law behavior with an exponent smaller than 1. This has been already observed in self-impact (Lillo et al. (2003)) and is extended here to cross-impact.[11]

Having established the evidence for cross-impact, we investigate its possible origin: Is this due to correlated trades across assets (e.g. a strategy trading several bonds simultaneously) or is it mostly due to quote revision following a trade, leading to changes of the mid-price of a bond in the absence of trades? To discriminate between these alternatives, we repeat the analysis in Figure 4.3 and distinguish now whether there were any trades beyond the triggering one in any other bond in our sample during a period from 3 seconds before to 2 seconds after the triggering transaction, which we will call *isolated trades*. For better readability in Figure 4.4 we focus on the four most recently issued 30 year BTPs in

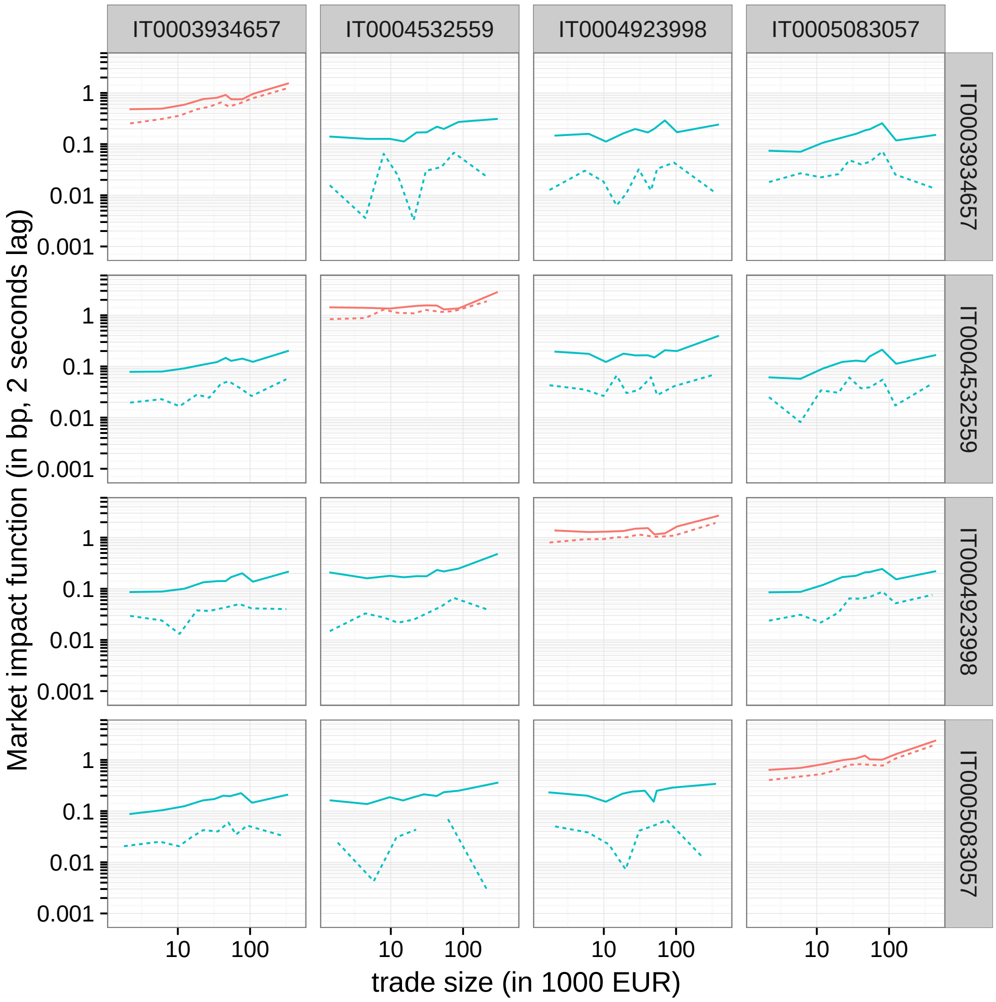

Figure 4.4: Plot of the average self- and cross-impact function among the four most recent 30 year bonds in our sample. Each line corresponds to one pairing ij, where impact is from the column on the row. Self-impact is present on the diagonal panels in red, cross-impact on the off-diagonals in blue. Solid lines show the market impact function based on all trades as in Figure 4.3, dotted lines show market impact based on isolated trades only, i.e. when there was no other transaction from 3 seconds before to 2 seconds after the triggering trade.

our sample, which were shown in the lower right panel of Figure 4.3. Results are similar for all other pairs of bonds. When we consider market impact of isolated trades only, self impact is lower than unconditionally. This is somewhat expected since order signs are positively autocorrelated and we exclude contributions where other trades have on average a positive contribution to impact. However the decrease in market impact is stronger for the cross-impact components, which are smaller by a factor of $\sim 5-10$ on average, whereas self-impact decreases only by a factor ~ 2 on average. We conclude therefore that both an autocorrelation of orders across assets as well as quote revisions play a role in forming cross-impact. In the next section we will take into account the (cross-) autocorrelation of the order sign when we estimate the shape of the decay of market impact.

4.4.4 Decay kernel

To estimate the empirically observed decay function we employ a multivariate version of the transient impact model of Bouchaud et al. (2004) and similarly to Benzaquen et al. (2017); Wang and Guhr (2016). While the advantage of the model lies in the fully non-parametric estimation of the kernel that we obtain, the TIM is typically estimated in event time which is asset-specific. Previous approaches avoid potential pitfalls by estimating the propagator in calendar time and binning trades. The estimation then is sensitive to the bin width. A small bin-width such as 1 second in Wang and Guhr (2016) introduces problems in the treatment of bins without trading activity, while a large bin width such as 5 minutes in Benzaquen et al. (2017) is too coarse to observe effects of single transactions. The main difference of our estimation is that we estimate the propagator in a combined market order time. Specifically our combined trade time is defined to advance by one unit for any unique timestamp at which there is at least one trade recorded, irrespective of the asset(s).[12]

Our model for the (log-) mid-price X_t^i of asset i just before a trade at time t reads

$$X_t^i = \sum_{t'<t} \left\{ \sum_j \left[H^{ij}(t-t')\epsilon_{t'}^j I_{t'}^j \right] + \xi_{t'}^i \right\} + X_{-\infty}^i \tag{4.29}$$

where ϵ_t^i is the order sign and I_t^i an indicator function for a trade in asset i at time t as defined in Section 4.4.1. ξ is a noise term with correlation matrix $\mathbf{\Sigma}^{(\xi)}$ and the empirically observed correlation structure of returns $\boldsymbol{r}$ of $\boldsymbol{X}$ is not $\mathbf{\Sigma}^{(\xi)}$ but the noise component $\mathbf{\Sigma}^{(\xi)}$ plus the component due to the correlated order flow and cross-impact $\mathbf{\Sigma}^{(\boldsymbol{H})}$, as shown in Benzaquen et al. (2017). Finally self- and cross-impact is captured by the propagator matrix $H^{ij}(\delta t)$ which gives the price impact of a trade in asset j on asset i after a positive time lag δt. Note that here we assume that trades of all volumes have the same impact and to avoid confusion with the previous sections we denote the decay kernel $\boldsymbol{H}$.[13] In this model returns r_t^i in asset i from a trade at time t to the next time-step are then defined as

$$\begin{aligned}
r_t^i &= X_{t+1}^i - X_t^i \\
&= \sum_j \underbrace{H^{ij}(1)}_{\mathcal{H}^{ij}(0)} \epsilon_t^j I_t^j + \sum_j \sum_{t'<t} \underbrace{H^{ij}(t+1-t') - H^{ij}(t-t')^{ij}}_{\mathcal{H}^{ij}(t-t')}(t-t')\epsilon_{t'}^j I_{t'}^j + \xi_t^i \\
&= \sum_j \sum_{t'\leq t} \mathcal{H}^{ij}(t-t')\epsilon_{t'}^j I_{t'}^j + \xi_t^i
\end{aligned} \tag{4.30}$$

where $\boldsymbol{\mathcal{H}}(\ell) \equiv \boldsymbol{H}(\ell+1) - \boldsymbol{H}(\ell)$, $\boldsymbol{H}(\ell \leq 0) \equiv 0$ and due to the definition of the price process in equation (4.29) a lag of $\tau = 0$ as the argument of $\boldsymbol{G}$ in equation (4.2) corresponds to $\ell = 1$ for $\boldsymbol{H}$. In practice (both due to computational limitations and to avoid dealing with overnight effects) the sum over t' is performed up to a cutoff lag p. For an estimation of $\boldsymbol{H}$ that is more stable with respect to p (Eisler et al. (2012)) we compute the observable

[12] In other words, each trade advances time by one step, unless when there are two or more trades (in the same or different assets) recorded at exactly the same timestamp (at millisecond resolution). In such a case our combined trade time advances only by 1. In our sample ca. 3% of trades happen at the same time-stamp as another trade in a different bond.

[13] $\boldsymbol{H}$ corresponds to the elementwise product of $\boldsymbol{f}$ and $\boldsymbol{G}$ as defined in equation (4.2), given the assumption of indifference to trade size.

$\tilde{\mathcal{S}}^{ij}(\ell)$

$$\begin{aligned}\tilde{\mathcal{S}}^{ij}(\ell) &= \mathbb{E}[r^i_{t+\ell}\epsilon^j_t I^j_t] && (4.31)\\ &= \sum_k \sum_{n\geq 0} \mathcal{H}^{ik}(n)\mathbb{E}\left[\epsilon^k_{t+\ell-n} I^k_{t+l-n}\epsilon^j_t I^j_t\right] \\ &= \sum_k \sum_{n\geq 0} \mathcal{H}^{ik}(n)\tilde{C}^{kj}(\ell - n) && (4.32)\end{aligned}$$

where $\tilde{\boldsymbol{C}}(\ell-n)$ is the cross-correlation matrix of the modified order sign $\epsilon^i_t I^i_t$ at lag $\ell - n$.[14] To estimate $H^{ij}(n) = \sum_{l=0}^{n-1} \mathcal{H}^{ij}$ we re-write equation (4.32) as a matrix equation

$$\tilde{\boldsymbol{\mathcal{S}}} = \boldsymbol{\mathcal{H}}\tilde{\mathbf{C}} \tag{4.33}$$

where with a slight abuse of notation $\tilde{\boldsymbol{\mathcal{S}}}$ and $\boldsymbol{\mathcal{H}}$ are row vectors of $N \times N$ block matrices, i.e. $\tilde{\boldsymbol{\mathcal{S}}} = (\tilde{\boldsymbol{\mathcal{S}}}(0), \cdots, \tilde{\boldsymbol{\mathcal{S}}}(p-1))$ and $\boldsymbol{\mathcal{H}} = (\boldsymbol{\mathcal{H}}(0), \cdots, \boldsymbol{\mathcal{H}}(p-1))$, and $\tilde{\boldsymbol{C}}$ is a symmetric block-Toeplitz matrix of $p \times p$ blocks of the correlation matrices at different lags, of dimension $Np \times Np$,

$$\tilde{\boldsymbol{C}} = \begin{pmatrix} \tilde{\boldsymbol{C}}(0) & \tilde{\boldsymbol{C}}(1) & \cdots & \tilde{\boldsymbol{C}}(p-1) \\ (\tilde{\boldsymbol{C}}(1))^\intercal & \tilde{\boldsymbol{C}}(0) & \cdots & \tilde{\boldsymbol{C}}(p-2) \\ \vdots & \vdots & \ddots & \vdots \\ (\tilde{\boldsymbol{C}}(p-1))^\intercal & (\tilde{\boldsymbol{C}}(p-2))^\intercal & \cdots & \tilde{\boldsymbol{C}}(0) \end{pmatrix} \tag{4.34}$$

where we use that $\tilde{\boldsymbol{C}}(-m)) = (\tilde{\boldsymbol{C}}(m))^\intercal$. To estimate $\boldsymbol{\mathcal{H}}$ and thus $\boldsymbol{H}$ we invert $\tilde{\boldsymbol{C}}$ and right-multiply equation (4.33) with $\tilde{\boldsymbol{C}}^{-1}$, where both $\tilde{\boldsymbol{S}}$ and $\tilde{\boldsymbol{C}}$ are constructed from (weighted) averages over daily estimations.

Figure 4.5 shows the mean of the decay kernel $H^{ij}(\tau)$ for self- and cross-impact averaged over all the bonds and pairings and weighted by the number of transactions. The mean and median values are not shown here but behave similarly. Both propagators do not decay immediately but reach their peak after ~ 10 transactions. This indicates a market inefficiency which has been observed for self-impact in other markets (see e.g. Figure 1 in Taranto et al. (2016)). In the absence of slippage this inefficiency could be exploited by e.g. a simple buy-hold-sell strategy. However here the expected gain is on the order of ~ 0.1 basis points while spread costs are > 1 basis points so that such a strategy would not be profitable. Further we observe that self- and cross-impact decay rather slowly with average self-impact reaching its initial level after ~ 100 transactions, corresponding to ~ 10 minutes of physical time and cross-impact taking even longer.

4.4.5 Testing for symmetry of cross-impact

We have shown in section 4.3.4 that for a bounded decay kernel the strength of cross-impact must be symmetric across pairs, i.e. $\eta^{ij} = \eta^{ji}$. Here we check whether this is empirically verified. In the estimation of the previous section where we are averaging over the trade volume, effectively regressing returns on trade events, this corresponds to the condition that $\hat{H}^{ij}(1) = \hat{H}^{ji}(1)$, i.e. we are assuming that prices are roughly constant so that absolute returns can be approximated by relative returns and that the average value (trade volume weighted by price) does not differ across bonds. As a robustness check, we repeat the estimation taking into account trading value, i.e. we modify Equation (4.30) to be

$$r^i_t = \sum_k \sum_{t'\leq t} \tilde{\mathcal{K}}^{ik}(t-t')\epsilon^k_{t'} W^k_{t'} I^k_{t'} \tag{4.35}$$

[14] Note that even though we refer to it as such, $\tilde{\boldsymbol{C}}$ is not strictly speaking a correlation matrix, as we do not de-mean nor normalize $\epsilon^i_t I^i_t$.

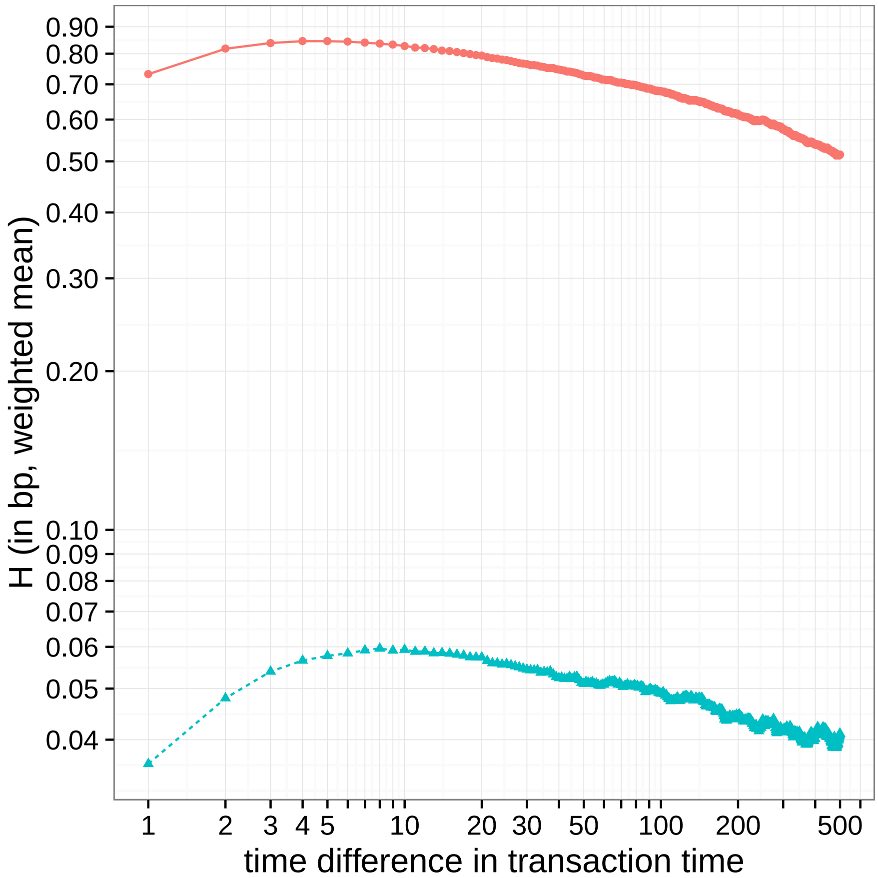

Figure 4.5: Plot of the estimated average decay kernel H^{ij} (in basis points) for self-impact (red dots connected by solid lines) and cross-impact (blue triangles connected by dashed lines) among all bonds in our sample. For self- (cross-) impact we show the mean over all bonds (pairings) weighted by the number of transactions in the triggering bond.

where we are now regressing returns on traded value $W_t^i = S_t^i V_t^i$ and the estimated impact and the decay kernel $\boldsymbol{K}(n) = \sum_{l=0}^{n-1} \tilde{\boldsymbol{\mathcal{K}}}(l)$ is connected to the $\boldsymbol{\eta}$ and $\boldsymbol{G}$ discussed in Sections 4.2 and 4.3 via

$$\tilde{K}^{ij}(t-s) = \tilde{\eta}^{ij} G^{ij}(t-s) = \frac{\eta^{ij}}{S_t^i S_t^j} G^{ij}(t-s) \tag{4.36}$$

when assuming a linear $\boldsymbol{f}$. Clearly the symmetry of η^{ij} of Lemma 4.3.6 must hold also for $\tilde{\eta}^{ij} = \tilde{K}^{ij}(1)$ which is the impact and decay kernel estimated at its smallest lag. Again we assume a roughly constant bond price process $\boldsymbol{S}$. The added accuracy of this estimation due to including the value is countered by the fact that the empirically observed market impact function is non-linear.

While we may easily check for symmetry on the estimated impact matrix $\boldsymbol{\eta}$, this does not allow for any statement on its statistical significance. Therefore we repeat the estimation on a shorter time scale, i.e. we obtain $\tilde{\boldsymbol{\mathcal{S}}}$ and $\tilde{\mathbf{C}}$ by averaging over the days of each calendar week instead of over the whole sample period and estimate the decay kernel $\boldsymbol{H}_w$ (or $\boldsymbol{K}_w$ respectively for the estimation on trade value) for each week w separately. For each of the 41 estimated $\boldsymbol{H}_w$ we compute the asymmetry $\Delta H_w^{ij} = H_w^{ij}(1) - H_w^{ji}(1)$ and for each of the $33 \times 32/2 = 528$ pairs we perform a Student's t-test of the null hypothesis that $\Delta H_w^{ij} = 0$. For robustness we repeat this for three different aggregation periods: weekly as

Table 4.1: Percentage of bond-pairs for which the null of symmetry in cross-impact is rejected according to a t-test on the null $\Delta H_w^{ij} = H_w^{ij}(1) - H_w^{ji}(1) = 0$ ($\Delta K_w^{ij} = K_w^{ij}(1) - K_w^{ji}(1) = 0$). Tests are performed on weekly/bi-weekly/monthly estimations of $\boldsymbol{H}$ ($\boldsymbol{K}$) from regressions of returns on signed trades (value of trades).

Percentage of significantly asymmetric pairs		confidence level		
regression on	aggregation	1%	5%	10%
trade events	weekly	8.0%	16.3%	24.4%
	bi-weekly	6.1%	15.3%	24.6%
	monthly	4.0%	14.0%	24.1%
trade value	weekly	3.0%	12.1%	21.8%
	bi-weekly	3.4%	11.6%	21.0%
	monthly	2.5%	11.0%	21.8%

described above, bi-weekly, and monthly.

In Table 4.1 we report the number of pairs for which the null hypothesis that $\Delta H^{ij} = 0$ ($\Delta K^{ij} = 0$) is rejected. The table reveals that for all scenarios and for all confidence levels the number of bond pairs for which the assumption of symmetric cross-impact is not supported is larger than the number expected under the null hypothesis.[15] This implies that in principle it is possible to exploit this dynamic arbitrage opportunity in at least some pairs, for example by using the strategy presented in Section 4.3.4.

We now check whether such a strategy would also be easily profitable in the bond pairs singled out above when taking into account bid-ask spread costs. While we have shown in section 4.3.7 that cross-impact should be symmetric also when there is slippage, the proof required trading fast and at high trading rates. Here we evaluate the negative (positive) cost term due to asymmetric cross-impact (slippage) for realistic values of the execution duration T and trading rate v:

$$C^{\text{cross}} \simeq v_a v_b \frac{T^2}{18} \Delta\eta \qquad , \qquad C^{\text{slippage}} \simeq (v_a B_a + v_b B_b) \frac{T}{3} \tag{4.37}$$

where $\Delta\eta = |\eta^{ab} - \eta^{ba}|$. In order to make a profit the ratio

$$\frac{C^{\text{cross}}}{C^{\text{slippage}}} \simeq \frac{v_a v_b T \Delta\eta}{6(v_a B_a + v_b B_b)} \tag{4.38}$$

must be larger than one. For the duration T, we need to keep in mind that in the proof of Lemma 4.3.6 we operated in the limit of very fast trading, i.e. under the assumption that the kernel is approximately constant. The most conservative estimate for T then would be 3 units of trade time, as this is the fastest we can execute the three phases of the strategy. On the other extreme, the empirically observed decay of impact in Figure 4.5 suggests that the kernel actually first increases for ~ 10 trade time units and then decays slowly, reaching its initial value only after ~ 100 (~ 500) trade time units for self-impact (cross-impact). Thus the maximal value of T which is consistent with the constant kernel is of the order of ~ 100 trade time units. If we assume too high of a value for the trading

[15] We have been unable to make out any obvious patterns which pairs are significantly asymmetric when ordering by various measures of liquidity and trading activity (time-to-maturity, maturity, bid-ask spread, average number of trades per day, average trade volume, turnover, tick size). This suggests that the asymmetry we observe is not just a mere artifact of any of those measures.

rate v, the assumption of linear impact costs no longer holds, as liquidity in the limit order book would be consumed faster than being replenished and additional price impact costs would arise. We therefore suppose that an arbitrageur would use an average-sized trading rate and assume $v_i T$ as three times the average trade value in asset i as reported in Table 4.2.[16]

In the following we estimate the ratio in equation (4.38) for the set of 6 pairs (1.1%) where symmetry is rejected at the 5% level for at least five of the six aggregation and regression scenarios of Table 4.1.[17] Conservatively assuming T as 3 units of trade time and with the assumptions described above, we get $\frac{C^{\text{cross}}}{C^{\text{slippage}}} \sim 1 \cdot 10^{-4}$. If we are able to maintain this strategy for a longer period T (at the same trading rate), we are getting closer to profitability since gains from cross-impact scale as T^2 and losses due to slippage as T. Assuming $T = 100$, i.e. the timescale when on average impact has decayed beyond its initial timescale, yields a ratio $\frac{C^{\text{cross}}}{C^{\text{slippage}}} \sim 0.005$. Only if we were able to neglect decay and other costs when executing our strategy throughout a whole trading day, it could turn profitable. Given that we do observe a faster decay (but also considering the associated risk and the chance that dominating the trading activity with the strategy might produce a less favorable impact structure) we conclude that dynamic arbitrage from cross-impact is unprofitable at least with our simple trading strategy. However our results also indicate that it is worth taking cross-impact into account when executing other strategies.

4.5 Conclusion

Even though cross-impact has been studied in the theoretical literature on optimal portfolio liquidation, empirical studies have been scarce until very recently. In this chapter we aim to connect the two strands of literature from a no-dynamic-arbitrage perspective.

A desirable market impact model should be free of arbitrage opportunities. In this chapter we focus on the specific class of multi-asset Transient Impact Models (TIMs) of market impact and we derive some necessary conditions for the absence of dynamic arbitrage. In particular, by using specific examples of simple round-trip strategies, we focus our attention on possible constraints on the shape and size of cross-impact.

One such condition is symmetry of cross-impact with respect to its direction between assets and we test it on empirical cross-impact kernels obtained by estimating a TIM in transaction time on Italian sovereign bonds traded in the MOT electronic market. Due to the strongly interrelated nature of these assets, we believe cross-impact plays a much stronger role here than compared to other asset classes such as, for example, equities. We find that while there exist statistically significant violations of the no-arbitrage conditions related to impact symmetry, these are unprofitable because of slippage costs such as the bid-ask spread which are neglected in the theoretical considerations.

In addition to this, we want to stress our contributions in describing the high-frequency market microstructure of the MOT sovereign bond market, applying the TIM to fixed-income markets and presenting evidence for cross-impact at the level of single orders instead of aggregated order flows. While this type of modeling and empirical estimation has been performed on many different types of markets, the application to sovereign bond (electronic) markets is new. This makes our study also relevant from a monetary policy point of view. Recent studies (Schlepper et al. (2017); Arrata and Nguyen (2017); De Santis

[16]To see this, denote the average trade size in asset i as $\bar{x}^i$ shares and take the case of executing the strategy in three trades corresponding to 3 units of trade time. The first phase, i.e. the first trade, lasts $T/3$ and is of size $\bar{x}^i$ shares, therefore $v^i T = 3\bar{x}^i$ shares.

[17]Considering different sets leads to similar results.

and Holm-Hadulla (2017)) that aim to quantify the price impact (measured as decline in the yield to maturity) of Quantitative Easing purchases in the Euro area could benefit from taking into account cross-impact effects.

Appendix

4.A Proofs

Let us recall that for all proofs in this section we assume that the decay kernel $\boldsymbol{G}$ is bounded, i.e. there exists an upper bound $U > 0$ so that $|G^{ij}(\tau)| < U$ for all $\tau \in [0, \infty)$ and all i, j and therefore we take $\boldsymbol{G}$ as normalized to 1 for its smallest lag $\tau = t - s$, i.e. $G^{ij}(0) = 1$ for all pairs ij. Furthermore we assume $\boldsymbol{G}(\tau)$ to be non-increasing and right-continuous at $\tau = 0$, i.e.

$$\forall\, \varepsilon > 0\, \exists\, T_\varepsilon > 0 \text{ such that } \forall\, \tau \text{ with } 0 < \tau < T_\varepsilon \text{ and } \forall\, i, j:\ |G^{ij}(0) - G^{ij}(\tau)| < \varepsilon\,. \tag{4.39}$$

Proof of Lemma 4.3.2. We first show that a non-odd f leads to a price manipulation in the single-asset case. Let us therefore assume an in-out strategy Π as in the first component of Example 4.3.1 with $\kappa = -1$ and therefore $\Theta = T/2$, i.e. both phases of trading last equally long. That is we first accumulate a position at the rate $v > 0$ to then liquidate it at the same negative rate $-v$. Without loss of generality we choose v such that $f(v) > -f(-v) \geq 0$ and that

$$\varepsilon := \frac{1}{4}\frac{f(v) + f(-v)}{f(v)} > 0 \tag{4.40}$$

which is non-zero since f is not an odd function.[18] Then by continuity there exists a $T_\varepsilon > 0$ for which we can bound the cost of our strategy as

$$\begin{aligned}
C(\Pi) &= v f(v) \int_0^{T/2} \mathrm{d}t \int_0^t G(t-s)\mathrm{d}s \\
&\quad - v f(-v) \int_{T/2}^{T} \mathrm{d}t \int_{T/2}^{t} G(t-s)\mathrm{d}s \\
&\quad - v f(v) \int_{T/2}^{T} \mathrm{d}t \int_0^{T/2} G(t-s)\mathrm{d}s \\
&\leq v f(v)\frac{T^2}{8} - v f(-v)\frac{T^2}{8} - v f(v)(1-\varepsilon)\frac{T^2}{4} \\
&= v\frac{T^2}{8}\left[-f(v) - f(-v) + 2\varepsilon f(v)\right] \\
&< 0
\end{aligned} \tag{4.41}$$

when choosing $T = T_\varepsilon$. That is for any trading rate v there is a $T_\varepsilon > 0$ for which there is a price manipulation by our choice of ε. Therefore we conclude that $f(v)$ must be an odd function of v in the single-asset case and the same holds for self-impact f^{ii} in the multi-asset case with cross-impact since we can always execute a strategy in only one asset.

Let us now show that the same holds for cross-impact. We choose $v_a, v_b > 0$ and for

[18] In the case that $f(v) < -f(-v)$ all we need is to change the sign in equation (4.40) to ensure that ε is positive. In the cases that either $f(v) \leq 0 \wedge f(-v) \leq 0$ or $f(v) \geq 0 \wedge f(-v) \geq 0$ the price manipulation arises in a simple in-out or out-in strategy as above respectively. Finally if assuming $vf(v) \leq 0$ the proof is analogous to the one above.

simplicity we assume that $\mathrm{sgn}(\dot{x}_i f^{ij}(\dot{x}_j)) = \mathrm{sgn}(\dot{x}_i \dot{x}_j)$ for all i, j, the proof being analogous in the other cases. We re-define

$$\varepsilon := \frac{v_a \left[f^{ab}(v_b) + f^{ab}(-v_b)\right]}{4v_a f^{aa}(v_a) + 4v_b f^{bb}(v_b) + 3v_a f^{ab}(v_b) - v_a f^{ab}(-v_b) + 3v_b f^{ba}(v_a) - v_b f^{ba}(-v_a)} > 0 \tag{4.42}$$

where again we assume for simplicity that $f^{ab}(v_b) + f^{ab}(-v_b) > 0$, the proof being analogous in the other case with $f^{ab}(v_b)$ and $f^{ab}(-v_b)$ interchanged in (4.42). We assume a first strategy Π_1 as in Example 4.3.1, again with $\kappa = -1$ and therefore $\Theta = T/2$, and $\lambda > 0$, that is we first accumulate both assets at the rates $v_a, v_b > 0$ to then liquidate both positions at the negative rates $-v_a, -v_b < 0$:

$$\Pi_1 = \{\dot{\boldsymbol{x}}_t\} \quad , \quad \dot{\boldsymbol{x}}_t = \begin{cases} (+v_a, +v_b)^\intercal & \text{for} \quad 0 \le t \le T/2 \\ (-v_a, -v_b)^\intercal & \text{for} \quad T/2 < t \le T \end{cases} \tag{4.43}$$

Choosing $T = T_\varepsilon$ we are able to bound the cost of this strategy as above, obtaining

$$\begin{aligned} C(\Pi_1) &\le \frac{T^2}{8} \sum_{i,j=a,b} v_i f^{ij}(v_j) - v_i f^{ij}(-v_j) - 2v_i(1-\varepsilon) f^{ij}(v_j) \\ &= \frac{T^2}{8} \left\{ -v_a \left[f^{ab}(v_b) + f^{ab}(-v_b)\right] - v_b \left[f^{ba}(v_a) + f^{ba}(-v_a)\right] + 2\varepsilon \sum_{i,j=a,b} v_i f^{ij}(v_j) \right\} . \end{aligned} \tag{4.44}$$

Repeating the same estimation for a strategy Π_2 which is anti-symmetric with $\lambda < 0$ but otherwise as above, i.e.

$$\Pi_2 = \{\dot{\boldsymbol{x}}_t\} \quad , \quad \dot{\boldsymbol{x}}_t = \begin{cases} (+v_a, -v_b)^\intercal & \text{for} \quad 0 \le t \le T/2 \\ (-v_a, +v_b)^\intercal & \text{for} \quad T/2 < t \le T \end{cases} . \tag{4.45}$$

The cost is similarly bounded from above:

$$\begin{aligned} C(\Pi_2) \le \frac{T^2}{8} &\left\{ -v_a \left[f^{ab}(v_b) + f^{ab}(-v_b)\right] + v_b \left[f^{ba}(v_a) + f^{ba}(-v_a)\right] \right. \\ &\left. + \varepsilon \left[2v_a f^{aa}(v_a) + 2v_b f^{bb}(v_b) + v_a f^{ab}(v_b) - v_a f^{ab}(-v_b) + v_b f^{ba}(v_a) - v_b f^{ba}(-v_a)\right] \right\} . \end{aligned} \tag{4.46}$$

Combining the cost of the two strategies with identical parameters v_a, v_b we have

$$\begin{aligned} C(\Pi_1) + C(\Pi_2) \le \frac{T^2}{4} \{ & - v_a \left[f^{ab}(v_b) + f^{ab}(-v_b)\right] \\ & + \varepsilon \, \Big[4v_a f^{aa}(v_a) + 4v_b f^{bb}(v_b) \\ & \quad + 3v_a f^{ab}(v_b) - v_a f^{ab}(-v_b) + 3v_b f^{ba}(v_a) - v_b f^{ba}(-v_a)\Big] \Big\} \end{aligned} \tag{4.47}$$

which is negative for our choice of ε and therefore price manipulation is possible.

□

Proof of Lemma 4.3.6. From the single-asset case we use the result that self-impact is linear and denote this as

$$f^{ii}(v) = \eta^{ii} v \quad . \tag{4.48}$$

Given the trading rates $v_a, v_b > 0$ and assuming w.l.o.g. that $v_a f^{ab}(v_b) > v_b f^{ba}(v_a) > 0$, we choose

$$\varepsilon := \frac{1}{2}\frac{v_a f^{ab}(v_b) - v_b f^{ba}(v_a)}{2\eta^{aa}v_a^2 + 2\eta^{bb}v_b^2 + 3v_a f^{ab}(v_b) + v_b f^{ba}(v_a)} \tag{4.49}$$

so that $\varepsilon > 0$.[19] Then by our assumption of continuity there exists a $T_\varepsilon > 0$ for which $|G^{ij}(0) - G^{ij}(\tau)| \le \varepsilon$ for all τ with $0 \le \tau \le T_\varepsilon$ and for all $i, j \in \{a, b\}$. We implement the same asymmetric strategy (4.15) as in Example 4.3.5 with $T = T_\varepsilon$ and calculate the cost of this strategy as $C(\Pi) = C^{aa} + C^{bb} + C^{ab} + C^{ba}$ where C^{aa} and C^{bb} are the self-impact costs of trading in assets a and b respectively and C^{ab} and C^{ba} are the costs due to cross-impact from asset b to a and vice versa. The explicit calculation of the self-impact terms gives

$$\begin{aligned}
C^{aa} &= \eta^{aa}v_a^2 \left\{ \int_0^{T/3} \mathrm{d}t \int_0^t G^{aa}(t-s)\mathrm{d}s + \int_{2T/3}^{T} \mathrm{d}t \int_{2T/3}^{t} G^{aa}(t-s)\mathrm{d}s - \int_{2T/3}^{T} \mathrm{d}t \int_0^{T/3} G^{aa}(t-s)\mathrm{d}s \right\} \\
&\le \eta^{aa}v_a^2 \left\{ \frac{T^2}{18} + \frac{T^2}{18} - (1-\varepsilon)\frac{T^2}{9} \right\} \\
&= \varepsilon\eta^{aa}v_a^2\frac{T^2}{9} \\
C^{bb} &\le \varepsilon\eta^{bb}v_b^2\frac{T^2}{9}
\end{aligned} \tag{4.50}$$

and likewise for cross-impact:

$$\begin{aligned}
C^{ab} &= v_a f^{ab}(v_b) \left\{ -\int_0^{T/3} \mathrm{d}t \int_0^t G^{ab}(t-s)\mathrm{d}s + \int_{2T/3}^{T} \mathrm{d}t \int_0^{T/3} G^{ab}(t-s)\mathrm{d}s - \int_{2T/3}^{T} \mathrm{d}t \int_{T/3}^{2T/3} G^{ab}(t-s)\mathrm{d}s \right\} \\
&\le v_a f^{ab}(v_b) \left\{ -(1-\varepsilon)\frac{T^2}{18} + \varepsilon\frac{T^2}{9} \right\} \\
&= v_a f^{ab}(v_b)\frac{T^2}{18}(-1+3\varepsilon) \\
C^{ba} &= v_b f^{ba}(v_a) \left\{ -\int_0^{T/3} \mathrm{d}t \int_0^t G^{ba}(t-s)\mathrm{d}s + \int_{T/3}^{2T/3} \mathrm{d}t \int_0^{T/3} G^{ba}(t-s)\mathrm{d}s \right\} \\
&\le v_b f^{ba}(v_a) \left\{ -(1-\varepsilon)\frac{T^2}{18} + \frac{T^2}{9} \right\} \\
&= v_b f^{ba}(v_a)\frac{T^2}{18}(1+\varepsilon)\,.
\end{aligned} \tag{4.51}$$

Summing over all terms yields

$$\begin{aligned}
C(\Pi) &\le \frac{T^2}{18}(v_b f^{ba}(v_a) - v_a f^{ab}(v_b)) + \varepsilon\frac{T^2}{18}\left(2\eta^{aa}v_a^2 + 2\eta^{bb}v_b^2 + 3v_a f^{ab}(v_b) + v_b f^{ba}(v_a)\right) \\
&= \frac{T^2}{36}(v_b f^{ba}(v_a) - v_a f^{ab}(v_b)) \\
&< 0
\end{aligned} \tag{4.52}$$

and we conclude that there is a price manipulation unless cross-impact is symmetric, i.e. we require $v_b f^{ba}(v_a) - v_a f^{ab}(v_b)$. □

[19] We can choose an equivalent $\varepsilon > 0$ in all other cases, i.e. for $v_b f^{ba}(v_a) > v_a f^{ab(v_b)} > 0$ by interchanging $a \leftrightarrow b$ in equation (4.49) and below. In the case that the denominator in (4.49) is negative we interchange $a \leftrightarrow b$ in order to ensure $\varepsilon > 0$ while the case of v_a and v_b such that the denominator is exactly zero is resolved by a slight modification of the turnaround points in the strategy.

Proof of Lemma 4.3.7. We prove the claim that that cross-impact needs to be a linear function of volume by contradiction, that is we show that non-linear cross-impact introduces arbitrage opportunities.

Therefore we consider a scenario as in Example 4.3.1 and illustrated in Figure 4.1a, i.e. trading in two assets over two phases, denoted by I and II, at a constant rate $v_{i,\mathrm{I}}$ and $v_{i,\mathrm{II}}$ respectively during each phase. The rates of the first and second phase are related by $\kappa = v_{i,\mathrm{I}}/v_{i,\mathrm{II}} < 0$ and the turn-around point $\Theta = T\frac{-v_{i,\mathrm{II}}}{v_{i,\mathrm{I}}-v_{i,\mathrm{II}}} = T\frac{1}{1-\kappa}$ is common to both assets a, b. From the single-asset case we use the result that self-impact is linear and denote this as $f^{ii}(v) = \eta^{ii}v$. We assume w.l.o.g. that $vf^{ij}(v) \geq 0$ and thus have $\lambda < 0$, i.e. trading in opposite directions.[20]

Let us assume that $f^{ab}(v)$ is a non-linear function. From Corollary 4.3.3 we know that $f^{ij}(0) = 0$ and therefore non-linearity implies that there exist $v_1, v_2 > 0, v_1 \neq v_2$ for which $f^{ab}(v_1)/v_1 \neq f^{ab}(v_2)/v_2$.[21] Therefore we can choose $v_b = v_{b,\mathrm{II}} > 0$ and $\kappa < 0, \kappa \neq -1$ so that

$$-\kappa v_a f^{ab}(v_b) > v_a f^{ab}(-\kappa v_b) > 0 \tag{4.53}$$

where also $v_a = -v_{a,\mathrm{II}} > 0$.[22] Then we can define

$$\varepsilon := \frac{1}{2}\frac{v_a f^{ab}(-\kappa v_b) + \kappa v_a f^{ab}(v_b)}{\kappa\eta^{aa}v_a^2 + \kappa\eta^{bb}v_b^2 - v_a f^{ab}(-\kappa v_b) + \kappa v_a f^{ab}(v_b)} > 0 \quad . \tag{4.54}$$

To see that ε is positive, note that the numerator in equation (4.54) is negative by equation (4.53) and all terms in the denominator in equation (4.54) are negative as well: the first two since $\kappa < 0$ and $\eta^{ii} > 0$ (cf. equation (18) with $v_b = 0$), and the last two terms are negative by equation (4.53) again.

By continuity of $\boldsymbol{G}$ as defined in equation (4.39) we can choose $T = T_\varepsilon$ such that the cost terms for self- and cross-impact can be bounded from above as

$$\begin{aligned}
C^{ii} &= \kappa^2\eta^{ii}v_i^2\int_0^\Theta \mathrm{d}t\int_0^t G(t-s)\mathrm{d}s + \eta^{ii}v_i^2\int_\Theta^T \mathrm{d}t\int_\Theta^t G(t-s)\mathrm{d}s + \kappa\eta^{ii}v_i^2\int_\Theta^T \mathrm{d}t\int_0^\Theta G(t-s)\mathrm{d}s \\
&\leq \eta^{ii}v_i^2\left[\kappa^2\theta^2/2 + (T-\Theta)^2/2 + (1-\varepsilon)\kappa T(T-\Theta)\right] \\
&= \varepsilon\eta^{ii}v_i^2\frac{\kappa^2T^2}{(1-\kappa)^2}
\end{aligned} \tag{4.55}$$

and for $i \neq j$

$$\begin{aligned}
C^{ij} &= \kappa v_i f^{ij}(-\kappa v_j)\int_0^\Theta \mathrm{d}t\int_0^t G(t-s)\mathrm{d}s - v_i f^{ij}(v_j)\int_\Theta^T \mathrm{d}t\int_\Theta^t G(t-s)\mathrm{d}s \\
&\quad + v_i f^{ij}(-\kappa v_j)\int_\Theta^T \mathrm{d}t\int_0^\Theta G(t-s)\mathrm{d}s \\
&\leq v_i f^{ij}(-\kappa v_j)\left[(1-\varepsilon)\kappa\Theta^2/2 + \Theta(T-\Theta)\right] + (1-\varepsilon)v_i f^{ij}(v_j)(T-\Theta)^2/2 \\
&= \frac{T^2}{(1-\kappa)^2}\left\{v_i f^{ij}(-\kappa v_j)\left[-\frac{\kappa}{2} - \varepsilon\frac{\kappa}{2}\right] + v_i f^{ij}(v_j)\left[-\frac{\kappa^2}{2} + \varepsilon\frac{\kappa^2}{2}\right]\right\} \quad .
\end{aligned} \tag{4.56}$$

[20] If $vf^{ij}(v) \leq 0$ then arbitrage arises from a strategy with $\lambda > 0$, i.e. trading in the same direction.

[21] We make the choice $v_1, v_2 > 0$ for simplicity of our arguments, this is without loss of generality since by Lemma 4.3.2 $f^{ij}(v)$ needs to be an odd function of v.

[22] The choice is either $v_b = v_1$ and $-\kappa v_b = v_2 > 0$ or $v_b = v_2$ and $-\kappa v_b = v_1 > 0$ (with $v_1 \leftrightarrow v_2$), chosen such that the first inequality in equation (4.53) is fulfilled.

Summing over all terms for two assets $i, j \in \{a, b\}$ gives

$$
\begin{aligned}
C(\Pi) &= \sum_{ij} C^{ij} \\
&\leq \frac{T^2}{2(1-\kappa)^2} \Big\{ -\kappa \Big[v_a f^{ab}(-\kappa v_b) + v_b f^{ba}(-\kappa v_a) + \kappa v_a f^{ab}(v_b) + \kappa v_b f^{ba}(v_a) \Big] \\
&\qquad + \varepsilon \Big[2\kappa^2 \eta^{aa} v_a^2 + 2\kappa^2 \eta^{bb} v_b^2 - \kappa v_a f^{ab}(-\kappa v_b) \\
&\qquad\qquad - \kappa v_b f^{ba}(-\kappa v_a) + \kappa^2 v_a f^{ab}(v_b) + \kappa^2 v_b f^{ba}(v_a) \Big] \Big\} \\
&= \frac{T^2}{(1-\kappa)^2} \Big\{ \kappa \Big[-v_a f^{ab}(-\kappa v_b) - \kappa v_a f^{ab}(v_b) \Big] \\
&\qquad + \varepsilon\kappa \Big[\kappa \eta^{aa} v_a^2 + \kappa \eta^{bb} v_b^2 - v_a f^{ab}(-\kappa v_b) + \kappa v_a f^{ab}(v_b) \Big] \Big\} \quad , \qquad (4.57)
\end{aligned}
$$

where in the last step we have used symmetry of cross-impact as shown in Lemma 4.3.6. Then by our choice of ε

$$
\begin{aligned}
C(\Pi) &\leq \frac{\kappa T^2}{2(1-\kappa)^2} \Big\{ -v_a f^{ab}(-\kappa v_b) - \kappa v_a f^{ab}(v_b) \Big\} \qquad (4.58) \\
&\leq 0
\end{aligned}
$$

and it follows that non-linear cross-impact f^{ab} admits arbitrage. Therefore, by symmetry, linearity of all f^{ij} is a necessary condition for absence of arbitrage.

□

Proof of Corollary 4.3.12. We follow the same proof as for Lemma 4.3.6 with linear self- and cross-impact, i.e. $f^{ij}(v) = \eta^{ij} v$. The bid-ask spread enters as an additional cost term in equation (4.52) so that

$$
C(\Pi) \leq v_a v_b \frac{T^2}{36} (\eta^{ba} - \eta^{ab}) + (v_a B_a + v_b B_b) \frac{T}{3} \quad . \qquad (4.59)
$$

Since $\eta^{ba} - \eta^{ab} < 0$ we can choose v_a, v_b large enough so that $C(\Pi) < 0$ and therefore price manipulation is still possible unless cross-impact is symmetric. □

4.B Power-law decay and impact

A popular class of unbounded decay kernels are power law kernels, e.g. $\boldsymbol{G}(\tau) \sim \tau^{-\gamma}$, $0 < \gamma < 1$. Their advantage lies in allowing for more realistic parametrizations of the market impact function, such as concave power-law impact $\boldsymbol{f}(v) \sim \mathrm{sgn}(v)|v|^{\delta}$ for $0 < \delta < 1$. While Gatheral (2010) establishes necessary conditions for such a model to be consistent in the one-dimensional case, numerical optimizations reported in Curato et al. (2016) find that violations of the principle of no-dynamic-arbitrage occur even when these conditions are verified, proofing them to be necessary but not sufficient. It remains an open problem whether and under what conditions power law decay kernels and market impact functions are consistent. In this section we do not address this question but consider necessary constraints that arise from the presence of cross-impact. Specifically we show that in this case the shape parameter of the market impact function $\boldsymbol{f}$ needs to be unique for all self- and cross-impact terms.

Lemma 4.B.1. *Assume a price process as in (4.2) where decay of market impact $\boldsymbol{G}(\tau)$ is a power law function, i.e. $G^{ij}(\tau) = \tau^{-\gamma^{ij}}$ with $0 < \gamma^{ij} < 1$ and $f^{ij}(v) = \eta^{ij}\,\mathrm{sgn}(v)|v|^{\delta^{ij}}$ is also power-law with $\eta^{ij} \geq 0$ for all i, j.*[23] *Then absence of dynamic arbitrage requires that*

$$\delta^{ij} = \delta \quad \forall\, i, j\,. \tag{4.60}$$

Proof of Lemma 4.B.1. Consider a strategy of two phases lasting equally long where at first we build up a position at a constant trading rate from time 0 until time $\Theta = T/2$ and then liquidate the position in a second phase from $T/2$ until T, i.e.

$$\Pi = \{\dot{\boldsymbol{x}}_t\} \quad , \quad \dot{\boldsymbol{x}}_t = \begin{cases} (v_a, v_b)^\intercal & \text{for} \quad 0 \leq t \leq T/2 \\ (-v_a, -v_b)^\intercal & \text{for} \quad T/2 < t \leq T \end{cases} . \tag{4.61}$$

which is a special case of Example 4.3.1. Further we use the notation $\lambda = v_a/v_b$. Explicit calculation of the cost terms due to self-impact yields

$$\begin{aligned} C^{ii} &= \eta^{ii} v_i^{1+\delta^{ii}} \left[\int_0^{T/2} \mathrm{d}t \int_0^t \mathrm{d}s \tau^{-\gamma^{ii}} + \int_{T/2}^T \mathrm{d}t \int_{T/2}^t \mathrm{d}s \tau^{-\gamma^{ii}} - \int_{T/2}^T \mathrm{d}t \int_0^{T/2} \mathrm{d}s \tau^{-\gamma^{ii}} \right] \\ &= \frac{\eta^{ii} v_i^{1+\delta^{ii}}}{1-\gamma^{ii}} \left[2 \int_0^{T/2} t^{1-\gamma^{ii}} \mathrm{d}t - \int_{T/2}^T \left\{ (t - \frac{T}{2})^{1-\gamma^{ii}} - t^{1-\gamma^{ii}} \right\} \mathrm{d}t \right] \\ &= \frac{\eta^{ii} v_i^{1+\delta^{ii}} T^{2-\gamma^{ii}}}{(1-\gamma^{ii})(2-\gamma^{ii})} \left[2^{\gamma^{ii}} - 1 \right] \\ &=: \Lambda^{ii}(\eta^{ij}, \gamma^{ij}, T) v_i^{1+\delta^{ii}} \end{aligned} \tag{4.62}$$

and similarly we obtain for cross-impact

$$\begin{aligned} C^{ij} &= -\eta^{ij} v_i v_j^{\delta^{ij}} \left[\int_0^{T/2} \mathrm{d}t \int_0^t \mathrm{d}s \tau^{-\gamma^{ij}} + \int_{T/2}^T \mathrm{d}t \int_{T/2}^t \mathrm{d}s \tau^{-\gamma^{ij}} - \int_{T/2}^T \mathrm{d}t \int_0^{T/2} \mathrm{d}s \tau^{-\gamma^{ij}} \right] \\ &= -\frac{\eta^{ij} v_i v_j^{\delta^{ij}} T^{2-\gamma^{ij}}}{(1-\gamma^{ij})(2-\gamma^{ii})} \left[2^{\gamma^{ij}} - 1 \right] \\ &=: -\Lambda^{ij}(\eta^{ij}, \gamma^{ij}, T) v_i v_j^{\delta^{ij}} \end{aligned} \tag{4.63}$$

with $\Lambda^{ij} > 0 \,\forall\, ij$. Since $\lambda = v_a/v_b$ we substitute w.l.o.g. $v_a = \lambda v$, $v_b = v$. The total cost of the strategy is thus

$$\begin{aligned} C &= \sum_{ij} C^{ij} \\ &= \Lambda^{bb} v^{1+\delta^{bb}} - \lambda^{\delta^{ba}} \Lambda^{ba} v^{1+\delta ba} - \lambda \Lambda^{ab} v^{1+\delta ab} + \lambda^{1+\delta^{aa}} \Lambda^{aa} v^{1+\delta^{aa}} \\ &= \Lambda^{bb} v^{1+\delta^{bb}} - \lambda^{\delta^{ba}} \Lambda^{ba} v^{1+\delta ba} + \mathcal{O}(\lambda) \end{aligned} \tag{4.64}$$

where in the last step we are choosing λ small enough so that linear terms in λ can be neglected. Then if $\delta^{bb} > \delta^{ba}$ we can choose $v > 0$ small enough so that $C < 0$ and likewise if $\delta^{bb} < \delta^{ba}$ we can choose v large enough so that there is a price-manipulation. Therefore we require $\delta^{bb} = \delta^{ba} =: \delta^b$ and likewise $\delta^{aa} = \delta^{ab} =: \delta^a$. Therefore we can re-express the

[23] In the one-dimensional case Gatheral (2010) shows that for absence of dynamic arbitrage it is also necessary that $\gamma \geq \gamma^* = 2 - \frac{\log 3}{\log 2} \approx 0.415$ and $\gamma + \delta \geq 1$.

cost as

$$
\begin{aligned}
C &= \Lambda^{bb} v^{1+\delta^b} - \lambda^{\delta^b} \Lambda^{ba} v^{1+\delta a} - \lambda \Lambda^{ab} v^{1+\delta a} + \mathcal{O}(\lambda^{1+\delta^a}) \\
&= \underbrace{\left(\Lambda^{bb} - \lambda^{\delta^b} \Lambda^{ba}\right)}_{:=\Lambda^b} v^{1+\delta^b} - \lambda^{\delta^a} \Lambda^{ba} v^{1+\delta a} + \mathcal{O}(\lambda^{1+\delta^a})
\end{aligned} \tag{4.65}
$$

now also considering linear terms in λ. Since λ is small we can use that $\Lambda^b > 0$ and by the same arguments as above we conclude that absence of arbitrage requires $\delta^a = \delta^b = \delta$.

$\square$

4.C List of ISINs

Table 4.2 reports basic descriptive statistics and liquidity measures for all the $N = 33$ bonds that were used for estimations. The set of bonds was selected as all fixed rate or zero-coupon Italian sovereign bonds with at least 5,000 trades throughout our sample. Note that some bonds were issued during our sample period and therefore less than 194 trading days were observed.

The description for each bond lists the bond type (BTPs refers to fixed-income treasury bonds, CTZ to zero coupon bonds), the fixed interest rate (where applicable) and the maturity date. Maturity is the time from issuance of a bond to the maturity date in years and time-to-maturity is calculated as the remaining time from the end of our sample (October 16, 2015) to the maturity date. The trade volume measures in Table 4.2 (mean traded volume per day and mean volume per trade) are reported as face volume traded and to arrive at the value one has to multiply by the price. The liquidity measures of the limit order book (mean number of limit order book updates per day, mean spread and ratio of tick size over mean spread) are computed from 10:00 to 17:00 of each day as we restrict our analysis to this period to avoid intraday seasonalities. Bid-ask spread is given in units of basis points of the face value, e.g. an average spread 14.9bp corresponds to a contract with a nominal value of EUR 100 offered at a mean spread of 14.9 euro-cents.

Table 4.2: Descriptives and liquidity measures for the set of bonds used in estimation.

ISIN	description	maturity (in years)	time-to -maturity (in years)	# days with observa- tions	avg. # trades per day	avg. traded volume per day (in million EUR)	avg. volume per trade (in 1,000 EUR)	avg. # LOB updates per day (in 1,000)	avg. spread (in basis points of par)	tick size / avg. spread
IT0001278511	BTPS 5.250 01/11/29	31.0	14.1	194	108.8	5.5	50.3	31.4	14.9	0.07
IT0003535157	BTPS 5.000 01/08/34	31.0	18.8	194	103.3	5.5	52.9	44.3	22.7	0.04
IT0003934657	BTPS 4.000 01/02/37	31.5	21.3	194	882.1	66.5	75.4	31.0	5.5	0.18
IT0004009673	BTPS 3.750 01/08/21	15.5	5.8	194	84.4	5.8	68.5	18.5	5.6	0.18
IT0004019581	BTPS 3.750 01/08/16	10.4	0.8	194	29.2	1.8	60.5	4.9	2.3	0.04
IT0004164775	BTPS 4.000 01/02/17	10.1	1.3	194	22.8	1.5	67.8	3.2	3.7	0.03
IT0004361041	BTPS 4.500 01/08/18	10.3	2.8	194	23.3	2.6	111.9	6.8	5.1	0.20
IT0004423957	BTPS 4.500 01/03/19	10.5	3.4	194	22.4	2.4	108.1	10.8	5.6	0.18
IT0004489610	BTPS 4.250 01/09/19	10.3	3.9	194	42.7	4.1	95.6	12.2	5.7	0.18
IT0004532559	BTPS 5.000 01/09/40	31.0	24.9	194	179.3	10.6	59.1	35.0	17.6	0.06
IT0004536949	BTPS 4.250 01/03/20	10.4	4.4	194	41.2	6.7	162.8	14.3	5.9	0.17
IT0004594930	BTPS 4.000 01/09/20	10.4	4.9	194	57.6	5.1	89.3	15.9	5.8	0.17
IT0004634132	BTPS 3.750 01/03/21	10.5	5.4	194	49.3	4.2	85.3	18.0	6.1	0.16
IT0004695075	BTPS 4.750 01/09/21	10.5	5.9	194	25.3	2.3	89.7	17.8	6.7	0.15
IT0004759673	BTPS 5.000 01/03/22	10.5	6.4	194	30.7	3.1	99.9	18.1	7.5	0.13
IT0004801541	BTPS 5.500 01/09/22	10.5	6.9	194	25.2	2.3	91.4	19.0	8.3	0.12
IT0004848831	BTPS 5.500 01/11/22	10.2	7.0	194	20.9	2.2	107.2	16.8	8.7	0.11
IT0004898034	BTPS 4.500 01/05/23	10.2	7.5	194	35.1	3.6	101.4	18.8	8.3	0.12
IT0004923998	BTPS 4.750 01/09/44	31.3	28.9	194	252.9	18.7	73.9	33.1	15.4	0.07
IT0004953417	BTPS 4.500 01/03/24	10.6	8.4	194	50.9	9.2	180.0	23.4	5.5	0.18
IT0005001547	BTPS 3.750 01/09/24	10.5	8.9	194	47.1	7.3	154.7	22.5	6.5	0.15
IT0005024234	BTPS 3.500 01/03/30	15.8	14.4	194	257.3	22.6	87.8	27.2	9.3	0.11
IT0005028003	BTPS 2.150 15/12/21	7.5	6.2	194	83.1	14.4	173.0	16.7	4.7	0.21
IT0005030504	BTPS 1.500 01/08/19	5.1	3.8	194	25.8	6.7	258.5	10.5	5.4	0.18
IT0005044976	CTZ 14- 30/08/16 24M	2.0	0.9	194	31.8	4.3	135.3	2.5	1.7	0.06
IT0005045270	BTPS 2.500 01/12/24	10.3	9.1	194	202.1	29.0	143.7	25.7	4.4	0.23
IT0005069395	BTPS 1.050 01/12/19	5.0	4.1	194	74.6	14.2	190.3	11.2	4.2	0.24
IT0005083057	BTPS 3.250 01/09/46	31.6	30.9	163	892.4	78.3	87.7	31.4	6.1	0.16
IT0005086886	BTPS 1.350 15/04/22	7.2	6.5	145	124.1	16.5	133.1	18.0	4.7	0.21
IT0005090318	BTPS 1.500 01/06/25	10.3	9.6	135	332.7	36.8	110.7	28.3	4.3	0.23
IT0005094088	BTPS 1.650 01/03/32	17.0	16.4	134	603.3	41.4	68.7	22.9	6.1	0.16
IT0005107708	BTPS 0.700 01/05/20	5.0	4.5	121	65.5	11.7	178.6	12.3	5.6	0.18
IT0005127086	BTPS 2.000 01/12/25	10.3	10.1	35	174.8	14.6	83.3	18.6	3.8	0.26

Chapter 5

Illiquidity spillovers

5.1 Introduction

Market liquidity - the ability to quickly trade large quantities of an asset at a low cost - is crucial to the functioning of financial markets and therefore of great interest to both market participants and policymakers. Concerns about illiquidity - the absence of market liquidity - have recently been raised not only about the *average* level of (il-)liquidity, but especially about *extreme* illiquidity events and the associated risk. In a study on the liquidity of the U.S. treasury market Adrian et al. (2015) conclude: *"perhaps the concerns are not so much about average liquidity levels, as we examined, but about liquidity risk. [...] episodes of sharp, seemingly unexplained price changes in the dollar-euro and German Bund markets have heightened worry about tail events in which liquidity suddenly evaporates."* During such tail events, also referred to as "liquidity crashes" or "liquidity dry-ups" (Brunnermeier and Pedersen (2009); Huang and Wang (2009); Easley et al. (2011)), market participants are unwilling to provide liquidity and as a consequence trading becomes prohibitively expensive.

Studying the dynamics of abrupt illiquidity events is important also because the clustering of several such shocks in the same asset can lead to a significant liquidity crash. In the following we will refer to such a cluster as an *illiquidity spiral.*[1] Moreover dry-ups of market liquidity might not be constrained to one asset but can propagate to other assets. We refer to the contagion of liquidity crashes in one asset to another as *illiquidity spillover.* Indeed market liquidity has become an important channel in the theoretical literature on financial contagion (see e.g. Kyle and Xiong (2001); Garleanu and Pedersen (2007); Brunnermeier and Pedersen (2009); Cespa and Foucault (2014) and the discussion in Bongaerts et al. (2015)). They relate to a large literature on commonality in liquidity pioneered by Chordia et al. (2000) and documenting the correlation between the liquidity of different assets, markets (Brockman et al. (2009)) and asset classes (e.g. Chordia et al. (2005), see also Karolyi et al. (2012) for a survey). The relation with price is studied by Hameed et al. (2010) who show that stock market declines increase the commonality in liquidity. However this literature mostly focuses on the average correlation in liquidity, including therefore normal and tail conditions.

In this chapter we focus instead on time lagged correlations of the occurrence of extreme illiquidity events across different assets. By using limit order book information and by constructing a suitable liquidity factor, we identify these events as an adaptation of the

[1] I.e. illiquidity spirals are downward spirals in market liquidity and not to be confused with liquidity spirals in the sense of Brunnermeier and Pedersen (2009) where downward moves in market and funding liquidity reinforce themselves.

peaks-over-threshold method to the liquidity factor increments. To model the dynamics of illiquidity spillovers and spirals we will use the class of self-exciting point processes named Hawkes processes (Hawkes (1971a,b)). They were initially applied to model earthquake data (Vere-Jones (1970, 1995); Ogata (1988)) and are now widely used in Finance (see Bacry et al. (2015) for a recent review) to describe discontinuous processes, such as price jumps or limit order book events. In fact spirals and spillovers can be naturally related to self- and cross-excitation events, respectively, of a Hawkes process.

We apply this approach to identify bond illiquidity spirals and spillovers. To our knowledge we are the first to use Hawkes processes to investigate extreme illiquidity events and to investigate intraday commonality and contagion of illiquidity shocks for bonds. In fact most of the research concentrates either on the equity market or on bond liquidity at daily or weekly frequency. A recent study by Bongaerts et al. (2015) directly identifies liquidity shocks in different equity markets at 5-minute resolution and does not find any evidence for spillovers in time nor across markets, challenging the concepts of liquidity dry-ups and contagion of liquidity shocks.

Here we follow a different approach, both methodologically and for the choice of a much shorter time scale. We follow the well-established method to model the occurrences of exceedances of the Value-at-Risk (VaR) threshold (Chavez-Demoulin et al. (2005); Embrechts et al. (2011); Chavez-Demoulin and McGill (2012)) by parametrically modeling the arrival process of illiquidity events as a Hawkes process that allows for both self- and cross-excitation. In this framework we can identify the fraction of such illiquidity shocks that is either due to a constant arrival intensity of events, self-excitation, or cross-excitation. We apply our methods to a dataset of the limit order book of the Mercato dei Titoli di Stato (MTS), the leading interdealer platform for European sovereign bonds, selecting a representative sample of Italian sovereign bonds of different maturities in the period from June 2011 through December 2015. Using a condensed measure of liquidity that is based on spread, price-impact and depth, we find significant evidence for the presence of illiquidity spirals and illiquidity contagion. We observe that illiquidity becomes more interconnected throughout our sample and the proportion of illiquidity contagion has roughly doubled from 2011 to 2015. Illiquidity spillover typically occurs within a few seconds (much faster than the sampling interval of 5 minutes in Bongaerts et al. (2015)) and is faster in bonds that are related by a similar maturity. It is notable that trades are very sparse on MTS and our results can be seen as directly due to quoting activity, supporting a channel of "informativeness" as in Cespa and Foucault (2014). Therefore the main contributions of this chapter lie in (i) proposing an event detection method for limit order book data that does not use a discrete time grid and therefore especially suited for modeling with Hawkes processes, and (ii) documenting extreme illiquidity contagion, i.e. cross-excitation of illiquid tail events, as well as self-exciting illiquidity spirals by applying our modeling approach to fixed-income order book data.

The rest of the chapter proceeds as follows. Section 5.2 describes our event detection algorithm and the Hawkes kernel used for modeling. In Section 5.3 we apply our method to the MTS dataset. Section 5.3.1 describes our dataset together with the market structure and the historic context. Our liquidity measure is introduced in 5.3.2 and sections 5.3.3 and 5.3.4 describe the results of our event detection and Hawkes modeling exercise respectively, with some robustness checks in Section 5.3.5. Section 5.4 mentions further possible applications of our method and concludes.

5.2 Method

We assume to observe a process $x^i(t)$ in continuous time that measures the illiquidity of an asset i at time t. This can be a direct liquidity metric (e.g. spread or depth) or a suitable combination of such metrics. For example in the empirical investigation below we will use a liquidity factor derived from a Principal Component Analysis of three liquidity metrics extracted from limit order book data. The important point is that we are able to measure the process, which is typically piecewise constant and changing in correspondence of limit order book events, at any time. In the next section we propose an identification method that yields a counting process $N^i(t)$ of extreme illiquidity events in the asset i up to time t. Then in Section 5.2.2 we introduce Hawkes processes and our modeling and estimation approach.

5.2.1 Event detection

Detection of extreme events is common in empirical analysis of prices because of their importance for risk assessment and the possible connection with jumps. Due to the unobservability of the efficient price and the presence of microstructure noise, the typical approach is to identify returns (rescaled by a suitable volatility estimator) beyond a certain high threshold, set e.g. as a high percentile of the return distribution (Chavez-Demoulin et al. (2005)) or as a multiple of the volatility estimator (Lee and Mykland (2008)). This method is often referred to as peaks-over-threshold (PoT), and can also capture the magnitude of the exceedances over the threshold.

Common to these detection methods is that they are effectively carried out on discrete time intervals, even when high-frequency data are available. For example Chavez-Demoulin and McGill (2012) samples stock-returns over 15-minute periods, Bongaerts et al. (2015) over five minutes or Bormetti et al. (2015) at one minute frequency. However the discrete time grid might miss events that occur at shorter time-scales and, when modeled as Hawkes processes, forces discrete time on the originally continuous Hawkes process. Furthermore discrete grids introduce the problem of contemporaneity of jumps. Indeed Bollerslev et al. (2008); Gilder et al. (2014); Bormetti et al. (2015) observe a significant number of co-jumps in price across multiple stocks.

When considering extreme illiquidity events, we can instead observe a proxy of illiquidity, i.e. $x^i(t)$ in our notation, which is updated at irregular but frequent intervals. The event time $M^i(t)$ advances by one unit with every update of x^i. We detect an illiquidity event, i.e. a shock to liquidity, when there is a large and abrupt increase in illiquidity, that is when the speed of increase in illiquidity is over a threshold θ^i. Specifically the counting process of liquidity shocks is defined as

$$N^i(t) = \sum_{m=0}^{M^i(t)} \mathbb{1}_{\left\{\frac{x^i(t_m)-x^i(t_{m-\ell^i})}{t_m-t_{m-\ell^i}} > \theta^i \wedge N^i(t_m)=N^i\left(t_{m-\ell^i}\right)\right\}} \tag{5.1}$$

where $M^i(t)$ is the number of updates of x^i up to time t, t_m the physical time corresponding to update m, θ^i the threshold and ℓ^i a lag that determines the level of coarse-graining in event time. The speed of increase in illiquidity is measured in physical time over a window length determined as ℓ^i updates in event time. In order to avoid modeling market activity as opposed to liquidity we have added a second condition, keeping an illiquidity shock only if there are at least ℓ^i updates since the previous detection. Thereby we ensure a minimum lag (in event time) between two successive detected illiquidity events. Note that the threshold θ^i and lag ℓ^i depend on the asset i to take into account a potentially different nature of liquidity processes (e.g. of different asset classes) and different update rates. For example, taking ℓ^i as the average number of updates over a certain range of physical time ensures that in the subsequent estimation step there is no bias introduced from different update rates in the liquidity processes of the assets estimated together.

5.2.2 Hawkes processes

Hawkes processes (Hawkes (1971a,b)) are point processes with counting process $N^i(t)$ where the intensity $\lambda^i(t)$ depends on the past history of the process. That is their intensity can be written as (following the notation of Bacry et al. (2015))

$$\lambda^i(t) = \mu^i + \int_{-\infty}^{t} \Phi^{ii}(t-s)\mathrm{d}N^i(s) + \sum_{j\neq i} \int_{-\infty}^{t} \Phi^{ij}(t-s)\mathrm{d}N^j(s) \tag{5.2}$$

where μ^i accounts for an exogenous baseline intensity and $\boldsymbol{\Phi}(\tau)$ is a kernel matrix which for all elements ij is non-negative, i.e. $\Phi^{ij}(\tau) \geq 0\ \forall\ \tau$, causal, i.e. $\Phi^{ij}(\tau) = 0\ \forall\ \tau < 0$, and L^1-integrable. For an asymptotically stationary process we further require that the spectral radius of the L^1-norm of $\boldsymbol{\Phi}$ is smaller than 1. Given stationarity of the kernel we can take the expectation value on both sides of the process in (5.2), i.e.

$$\bar{\boldsymbol{\lambda}} = \boldsymbol{\mu} + \boldsymbol{\gamma}\bar{\boldsymbol{\lambda}} \tag{5.3}$$

where $[\boldsymbol{\gamma}]^{ij} = \int_0^\infty \Phi^{ij}(\tau)\mathrm{d}\tau$ is the kernel norm and $\bar{\boldsymbol{\lambda}}$ the unconditional expectation of the event arrival intensity, which is well-defined under the stationarity condition above:

$$\bar{\boldsymbol{\lambda}} = (\mathbb{1} - \boldsymbol{\gamma})^{-1}\boldsymbol{\mu}\,. \tag{5.4}$$

Re-writing equation (5.3) for a single component i and dividing by $\bar{\lambda}^i$ we get

$$1 = \frac{\mu^i}{\bar{\lambda}^i} + \gamma^{ii} + \sum_{j\neq i} \gamma^{ij}\frac{\bar{\lambda}^j}{\bar{\lambda}^i}, \tag{5.5}$$

which gives an interpretation to each term on the rhs as the fraction of events due to the baseline intensity μ^i, self-excitation, and cross-excitation respectively. In the following we will take $N^i(t)$ to be the counting process of illiquidity events and therefore interpret $\sum_{j\neq i}\gamma_{ij}\bar{\lambda}^j/\bar{\lambda}^i$ as a measure of the frequency of illiquidity spillovers from other assets j to asset i. γ^{ii} is the fraction of illiquidity events explained by self-excitation. A higher γ^{ii} means liquidity shocks in asset i are more likely to amplify and propagate in time and market liquidity therefore is more fragile or less resilient along the time dimension.

5.3 Application

5.3.1 MTS market

The sovereign bond market of the Eurozone is one of the largest in the world with €6.8 trillion outstanding nominal value at the end of 2015.[2] The largest debtor country by the same measure is Italy with €1.8 trillion outstanding. Beyond its size the Italian market is

[2]Nominal Value of outstanding amounts issued by central governments according to ECB: https://www.ecb.europa.eu/stats/ecb_statistics/escb/html/table.en.html?id=JDF_SEC_OAT_DEBT_SECURITIES&period=2015-12

of even bigger importance as a pillar of the Euro, as during and after the sovereign bond crisis, Italy was repeatedly seen as crucial to the survival of the Eurozone.[3]

The secondary market for Italian sovereign bonds is divided in an opaque over-the-counter (OTC) market and an observable exchange-traded market that is organized in different platforms. The market share of the OTC market was 59.1% in 2015, while 44.8% of trading in organized platforms reported to the Italian securites and exchange commission, CONSOB, took place on the interdealer platform Mercato dei Titoli di Stato (MTS).[4] Besides accounting for almost half of the exchange trade activity, MTS is also important because it is used by the Italian treasury to evaluate the performance of primary market participants in terms of liquidity provision.[5] MTS therefore is the leading interdealer trading platform for European and especially Italian sovereign bonds (Dufour et al. (2004); Pelizzon et al. (2016)). It is organized as an electronic limit order book split into a "EuroMTS" market for benchmark bonds and a domestic market and minimum quote and transaction sizes are typically €1 million or larger.

Our dataset contains all trades and limit orders on the MTS platform from June 2011 to December 2015 at millisecond and from the beginning of 2013 at microsecond resolution. From the dataset that includes all bonds we pick for each date the three most recently issued 5, 10 and 30 year fixed-rate Italian government bonds (Buoni del Tesoro Poliennale, BTP). In the U.S. treasury market the on-the-run, i.e. the most recently issued bond of a maturity, enjoys a special status and better liquidity than older off-the-run bonds. This is not the case in Italy (Coluzzi et al. (2008)) where the issued volume of a bond is typically augmented several times in reopenings. Performing our analysis on three bonds per maturity improves our statistics without dispersing too much over different times to maturity. We have verified that there is no significant difference between on- and off-the-run bonds in the analysis below.

We remove duplicated and inactive orders as well as filtering for obvious outliers and reconstruct the full limit order book. Regular trading hours are from 9:00 to 17:30 and we discard the first 30 and last 15 minutes of each day. We further discard all observations from 26 November 2012 through 28 December 2012, as on those days the data have already been recorded with microsecond precision while still being reported at millisecond precision in our dataset. This renders it impossible for us to reliably reconstruct the limit order book at the full precision without introducing artefacts.

Our sample period is also interesting from an economic point of view. It starts amid the European sovereign debt crisis where in August 2011 Spanish and Italian 10 year bond yields breached 6% and led the ECB to purchase mostly Spanish and Italian sovereign debt through the Securities Market Program (SMP) during the same month.[6] Subsequent rating downgrades and international concerns mounted pressure on the Italian government under Berlusconi that was replaced by a technocratic cabinet in November 2011. In December 2011 and February 2012 the ECB used their Long Term Refinancing Operation (LTRO) to infuse credit into the banking system at exceptionally good conditions, which indirectly lowered yields (Crosignani et al. (2015)). The situation subsequently relaxed by mid 2012 to early 2013, in the light of significantly lowered interest rates and the "whatever it takes" speech of ECB president Draghi in July 2012. After cutting the deposit interest rate to

[3]See e.g. "Monti in fight for survival - of Italy and euro", *Financial Times*, June 20, 2012, available at `http://www.ft.com/cms/s/0/1ef8289c-baef-11e1-b445-00144feabdc0.html`

[4]CONSOB, Bollettino Statistico Nr. 8, March 2016, available at `http://www.consob.it/web/area-pubblica/bollettino-statistico`

[5]See `http://www.dt.tesoro.it/en/debito_pubblico/specialisti_titoli_stato/`

[6]See Ghysels et al. (2014) and press release `http://www.ecb.europa.eu/press/pr/date/2011/html/pr110807.en.html`.

negative terrain in June 2014, on 22 January 2015 the ECB announced their Quantitative Easing (QE) program under the name of Public Sector Purchase Program (PSPP). The PSPP was started in March 2015 with a monthly volume of €50 billion of purchases of euro area sovereign debt with a remaining maturity of 2-30 years. For Italy this translates to monthly purchases of €7.7 billion, with so far unknown effects for market liquidity.

5.3.2 Liquidity metrics

Liquidity is a multi-dimensional latent process and for exchange-traded securities typically observed through asset characteristics, trades (Goyenko et al. (2009)) and/or the limit order book. Since trades are sparse on MTS (Darbha and Dufour (2013)), trade-based measures of liquidity could only be sensibly constructed at daily or lower frequencies, incompatible with our high-frequency analysis. We therefore focus on liquidity measures based on the limit order book.

Following Pelizzon et al. (2014) we use three "raw" metrics of liquidity: bid-ask spread, total quoted volume, and inverse depth. *(Bid-Ask) Spread* is defined as the best ask minus the best bid price and captures the round-trip cost of small trades. *Total (Quoted) Volume* is the sum of all the volume quoted in the limit order book, irrespective of price, and taken as the mean of the bid and ask sides. It is highly correlated ($> 95\%$) with the number of active proposals and indicates general willingness to participate in the market and depth of the whole book. *Inverse Depth* on the ask (bid) side is defined as how much a buy (sell) trade of size €15 million would shift the best ask (bid) price at any instant of time.[7] To obtain inverse depth we take the mean of the measures on the bid and ask side. Inverse depth, in the absence of frequent trades, is a measure of (the virtual or mechanical) price impact.[8]

To obtain a condensed measure of illiquidity we perform a principal component analysis (PCA) for each bond on the "raw" liquidity metrics bid-ask spread, total quoted volume and inverse depth as in Fleming (2003) and Mancini et al. (2013). Let $\boldsymbol{L}$ be the $T \times 3$ matrix of the demeaned and standardized time-series (T observations) of the three liquidity metrics above, sampled at 1-minute intervals. Then the empirical covariance matrix is proportional to $\boldsymbol{L}'\boldsymbol{L} = \boldsymbol{V}\boldsymbol{\Lambda}\boldsymbol{V}'$ where $\boldsymbol{\Lambda}$ is the 3×3 diagonal matrix of eigenvalues and $\boldsymbol{V}$ the 3×3 matrix of the eigenvectors of $\boldsymbol{L}'\boldsymbol{L}$. The eigenvector $\boldsymbol{v}_1$ corresponding to the first principal component, capturing the majority of variance in liquidity, has positive eigenvector loadings in the Spread and Inverse Depth component and negative loading in the Total Volume component, thus measuring illiquidity.[9] The time-series of our condensed illiquidity measure *PCA1* is thus constructed as $PCA1 = \boldsymbol{L}\boldsymbol{v}_1$ where now both *PCA1* and $\boldsymbol{L}$ are at the same millisecond or higher time-resolution as our data.

Since our sample period spans from the sovereign bond crisis up to the implementation of Quantitative Easing in the Eurozone, we do not assume that the (co-)variance structure of liquidity remained constant and instead estimate the illiquidity measure for each month on a rolling-window basis. That is for each month we estimate the demeaning and standardization coefficients of the raw liquidity metrics and the PCA loadings based on $\boldsymbol{L}$ sampled in the period from six months before to six months after the current month,

[7] The amount of €15 million was chosen as the 90% percentile of trade sizes. Inverse depth thus reflects the cost of a large trade requiring immediacy (Pelizzon et al. (2014)).

[8] While price impact is typically computed as a regression of price changes on order flow, having at our disposition the complete limit order book, we can also compute the mechanical price response that would arise to a given trade. This virtual or mechanical price impact is however only one component of price impact, as the reaction of market participants to trades also plays a role, which cannot be captured in this way.

[9] The eigenvector is uniquely defined up to a sign. We normalize the Spread component to always be positive, i.e. *PCA1* measures illiquidity as opposed to liquidity.

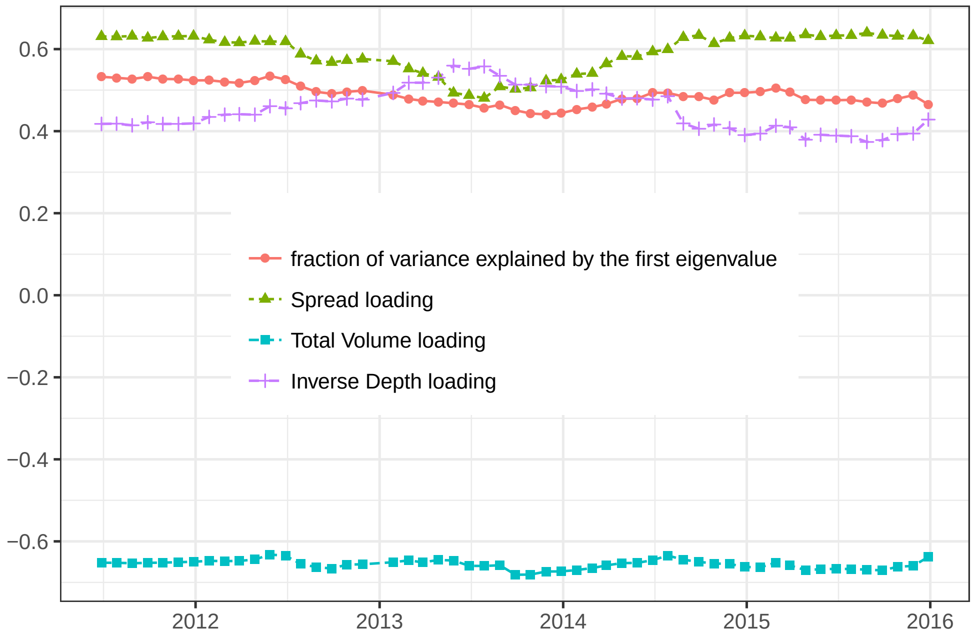

Figure 5.1: Eigenvector loadings and fraction of variance explained by the first eigenvector. Green triangles, blue rectangles and purple crosses show the eigenvector loadings of the first eigenvector in each month for Spread, Total Volume and Inverse Depth respectively, averaged over all bonds in the sample during the respective month. Red circles show the fraction of variance explained by the first eigenvector.

or shorter when new bonds are issued. Figure 5.1 shows the eigenvector loadings in the raw liquidity measures as average over all bonds in the active subsample, as well as the fraction of variance explained by the first eigenvector. Eigenvector loadings are similar over time and we have also verified that they are similar across different bonds, with small deviations occurring mostly in the relatively calm period around 2013, and robust also to different sampling frequencies.

In Figure 5.2 we report the distribution of *PCA1* sampled at one minute frequency for all the bonds in our sample. The distribution is very skewed and clearly shows a heavy right tail, indicating the presence of market liquidity crises. This distribution is similar for all bonds in our sample.

5.3.3 Illiquidity event detection

In this section we apply the event detection developed in Section 5.2.1 to the condensed liquidity measure *PCA1* described in Section 5.3.2 above. The counting process $N^i(t)$ of illiquidity events of bond i at time t is then defined as

$$N^i(t) = \sum_{m=0}^{M^i(t)} \mathbb{1}_{\left\{ \frac{PCA1^i(t_m) - PCA1^i(t_{m-\ell^i})}{t_m - t_{m-\ell^i}} > \theta^i \wedge N^i(t_m) = N^i(t_{m-\ell^i}) \right\}} \tag{5.6}$$

where $M^i(t)$ is the number of updates of $PCA1^i$ up to time t and t_m the physical time corresponding to update m. As for the PCA we evaluate the parameters ℓ^i and θ^i for each month over a running window lasting from six months before to six months after the current month. For each of these windows and for each bond i we determine ℓ^i as the

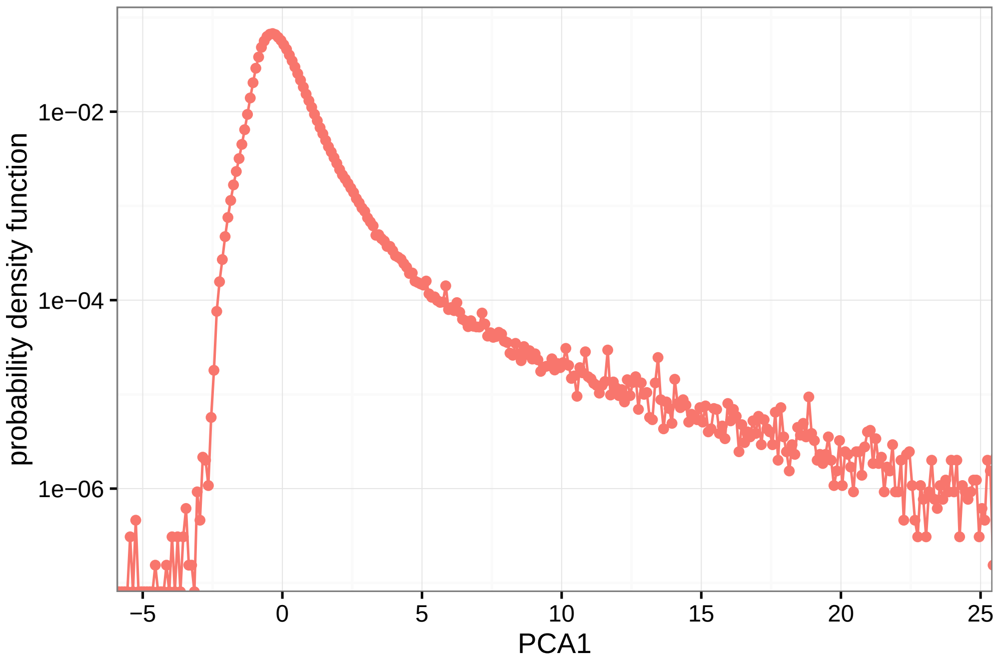

Figure 5.2: Probability density function of illiquidity *PCA1* sampled at 1 minute intervals, averaged over all days and bonds in the sample. y-axis is logarithmic.

average number of $PCA1^i$ updates per minute and θ^i as the 95% percentile of the liquidity returns $\frac{PCA1^i(t_m)-PCA1^i(t_{m-\ell^i})}{t_m-t_{m-\ell^i}}$.

Figure 5.3 illustrates our detection method by showing a snapshot of the full limit order book of a 5 year BTP together with eight illiquidity events detected in quick succession around 15:27 on 7 June 2012. The first event follows a rapid decline in the total quoted volume while the subsequent events correspond to a step-wise widening of the bid-ask spread.[10]

Figure 5.4 describes the distribution of the number of detected illiquidity events per day. While the peak of the distribution is around ~ 20 events per bond per day, there is also a heavy tail of days with many more illiquidity events. Given that we are considering extreme events, the high number of detected observations may seem counter-intuitive. However note that since we are working in event time, a quick succession of liquidity shocks is entirely possible and thus constitutes the actual extreme event. The idea is that one illiquidity event by itself is unlikely to have severe externalities, but the concern is much more about the risk it carries, i.e. whether it gets amplified (resulting in subsequent shocks) or not. Finally note that also in Chavez-Demoulin et al. (2005) the threshold is such that around 10% of the data are exceedances, and the VaR at higher thresholds is modeled based upon these events.

The intraday pattern of illiquidity events in Figure 5.5 can be decomposed into two main components. First there is a background of events throughout the day that is stronger around the opening and weaker during lunch hours. This follows the intraday pattern of

[10]Indeed we often observe such a step-wise widening of the bid-ask spread when the limit order that was previously the best disappears. Our observations suggest that there is a structure of confident participants that are willing to quote at the best price and following participants that make their orders relative to those of other market participants. This has been confirmed in conversations with market participants.

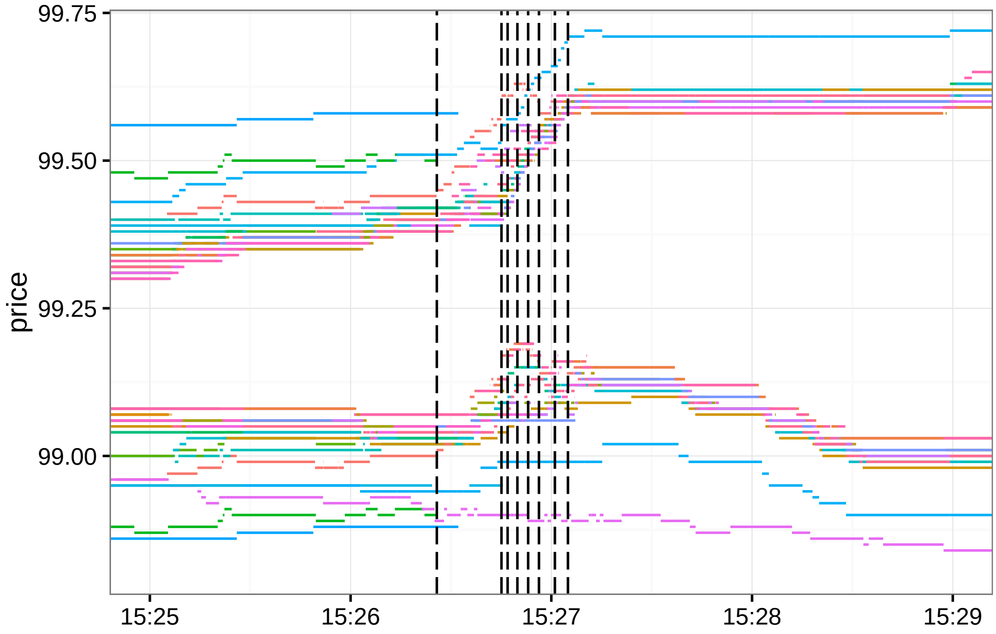

Figure 5.3: Dynamics of the limit order book with the detected illiquidity events of the 5 year BTP IT0004793474 on roughly four minutes of June 7, 2012. Solid horizontal lines represent limit orders, quoted in EUR per 100EUR face value and different colors indicate the limit orders of different market participants. Detected illiquidity events are marked as vertical dashed black lines. No trades took place in the period depicted.

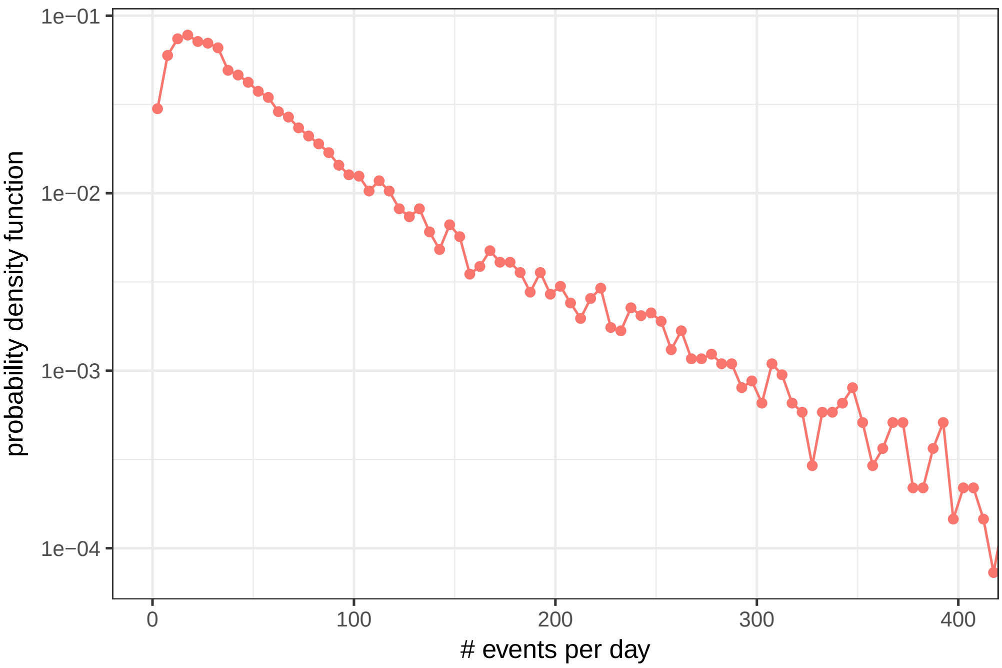

Figure 5.4: Probability density function of the number of illiquidity events per day, averaged over all days and bonds in the sample. y-axis is logarithmic.

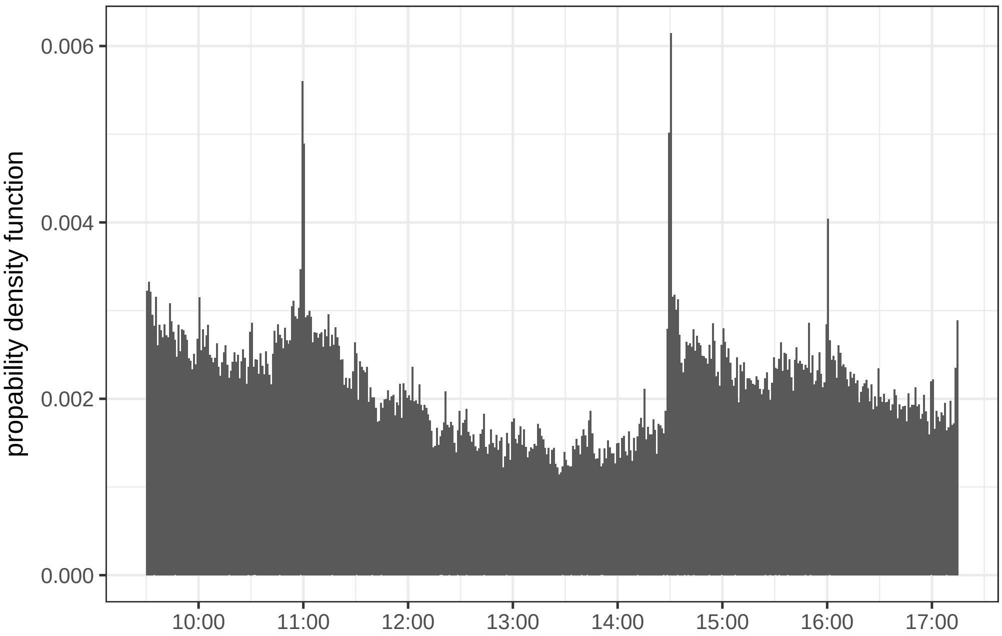

Figure 5.5: Intraday pattern of illiquidity events. Sampled over the whole sample period and all bonds. The binwidth is 1 minute. Results are similar across maturity groups and over time.

market activity since the detection of liquidity shocks is related to market activity. Secondly there are several major spikes at 11:00, 14:30 and 16:00. We attribute these peaks to the arrival of scheduled exogenous news. E.g. the spike at 14:30 corresponds to the release of important macro-economic news in the U.S. leading to a similar spike in illiquidity in the U.S. treasury market on announcement days (Fleming and Remolona (1999)), while Euro-area data is typically released at 11:00. Rambaldi et al. (2015) show that scheduled announcement effects lead to increased trading activity in the foreign exchange market both after and also before the release of macroeconomic news and that activity around the news release can be well described by a Hawkes model. Therefore in Section 5.3.5 we conduct a robustness test showing that our results on self- and cross-excitation are not due to these announcement effects.

5.3.4 Self-excitation and spillover

The first question we ask, is whether the Hawkes modeling of extreme illiquidity events is necessary, or if a Poisson process explains the data well enough. To answer this question, the left panel of Figure 5.6 shows the quantile-quantile plot of sample inter-event durations of illiquidity shocks against exponential quantiles for a subset of days in our sample. If the arrival intensity of illiquidity events were constant (i.e. a Poisson process), then the sample quantiles plotted against exponential quantiles should lie on a line through the origin. This is clearly not the case since the figure shows a significant presence of short inter-event times, indicative of clustering, and justifying the use of Hawkes processes.

Therefore we model the arrival intensity $\lambda^i(t)$ of illiquidity events in bond i as a Hawkes process as described in Section 5.2.2

$$\lambda^i(t) = \mu^i + \sum_j \int_{-\infty}^{t} \Phi^{ij}(t-s)\mathrm{d}N^j(s) \tag{5.7}$$

where $N^i(t)$ is the counting process of illiquidity events as defined in equation (5.6). We choose to parametrize the Hawkes kernel $\mathbf{\Phi}(t)$ in equation (5.7) as an exponential kernel

$$\Phi^{ij}(t-s) = \sum_{k=1}^{P_{ij}} \alpha_k^{ij} \exp(-\beta_k^{ij}(t-s)) \tag{5.8}$$

with a double exponential kernel for the self-exciting component ($P_{ii} = 2$) and a single exponential for the cross-excitation terms ($P_{ij} = 1, i \neq j$). The advantage of the exponential kernel is that the corresponding likelihood function can be computed recursively and thus lends itself to maximum-likelihood estimations (Toke (2011)). We can further interpret the α^{ij} in equation (5.8) as capturing fragility (the amplitude of reaction to an illiquidity shock) whereas the inverse of β^{ij} is the time-scale of the decay of self- and cross-excitation: the larger β the quicker the arrival intensity of induced shocks decays. Both parameters enter in calculating the norm of the kernel

$$\gamma^{ij} = \sum_{k=1}^{P_{ij}} \frac{\alpha_k^{ij}}{\beta_k^{ij}} \, . \tag{5.9}$$

Empirically we proceed by estimating the Hawkes parameters separately for each trading day and each pair of bonds via a maximum-likelihood estimation, leaving us with 14 parameters to estimate. Estimating more bonds together (i.e. summing over more j) would require too many fitting parameters (Bacry et al. (2015)) given our short time series and instead we will average over multiple estimations for the same bond i with different j. The average is either taken over all other bonds $j \neq i$ or over bonds with similar maturity where we are interested in spillover effects between maturities. To avoid over-fitting we discard estimations where there are not sufficient illiquidity events in both bonds[11] (5.1% of estimations) and we discard estimations where the estimated Hawkes process is non-stationary, i.e. has a spectral radius ≥ 1 and thus the stationary state is not well-defined (9.8% of estimations). By choosing trading days as the estimation period we avoid the problem of overnight returns. Alternatively, assuming stationarity of the parameters across days, one could estimate the kernel parameters on a longer time series of concatenated observations (as in e.g. Chavez-Demoulin and McGill (2012)) or by summing the likelihoods of different days.[12]

In the right panel of Figure 5.6 we show the quantile-quantile plot of the inter-event durations rescaled by the estimated intensity for the same days as in the left panel of the same figure. The plot indicates that the Hawkes process describes well the dynamics of liquidity shocks. The inset shows the same plot in a logarithmic scale. For very short durations the data quantiles are larger than the theoretical ones, which we conjecture is due to the minimum duration of ℓ limit order book updates between two events imposed in our event detection scheme. While the finite reaction time of market participants is an alternative explanation, that effect will still be amplified by the minimum lag. In a Ljung-Box test the null of no autocorrelation between rescaled inter-event times is rejected (at a 10% significance level) in 10% of cases for a lag of both 1 and 10 events and similarly for other significance levels, supporting our choice of modelling.

[11] We require at least 20 observations in both bonds, i.e. $\sum_i N_i \geq 20$.

[12] We have verified that our results also hold for estimation on weekly concatenated time-series. The only non-negligible difference is that, due to the longer estimation period, the slower of the two time-scales of self-excitation (i.e. the largest $1/\beta_k^{ii}$) is slower than for daily estimations. Therefore self-excitation captures a larger fraction of events in the weekly estimation whereas less events are attributed to the baseline intensity.

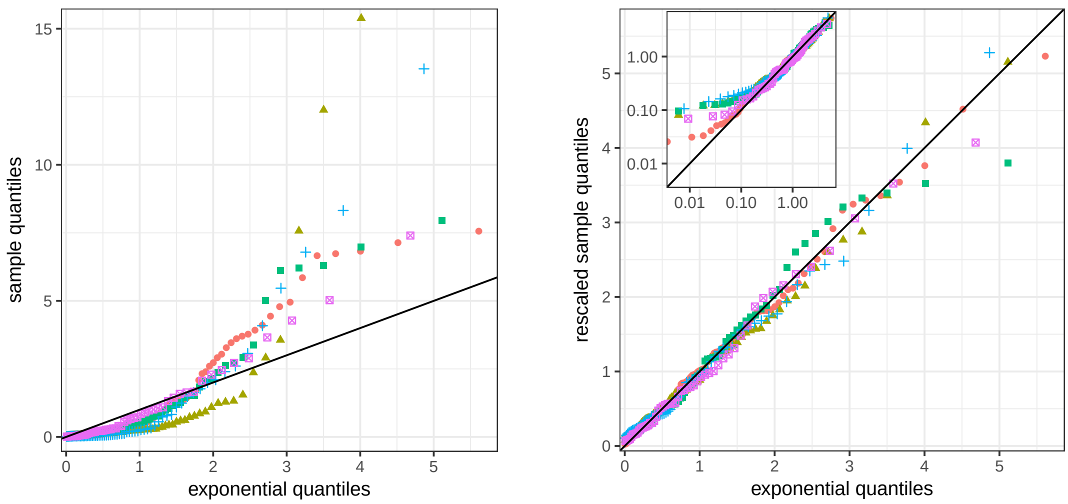

Figure 5.6: Quantile-quantile plot of sample inter-event durations between liquidity shocks and exponential quantiles (left panel) and between sample inter-event durations rescaled by the estimated intensity and exponential quantiles (right panel). The inset is in logarithmic scale on both axis. The inter-event durations are of IT0003934657 (estimated with IT0004286966, both are 30 year BTPs) and different colors and symbols indicate different days from 26 October to 1 Nov 2011.

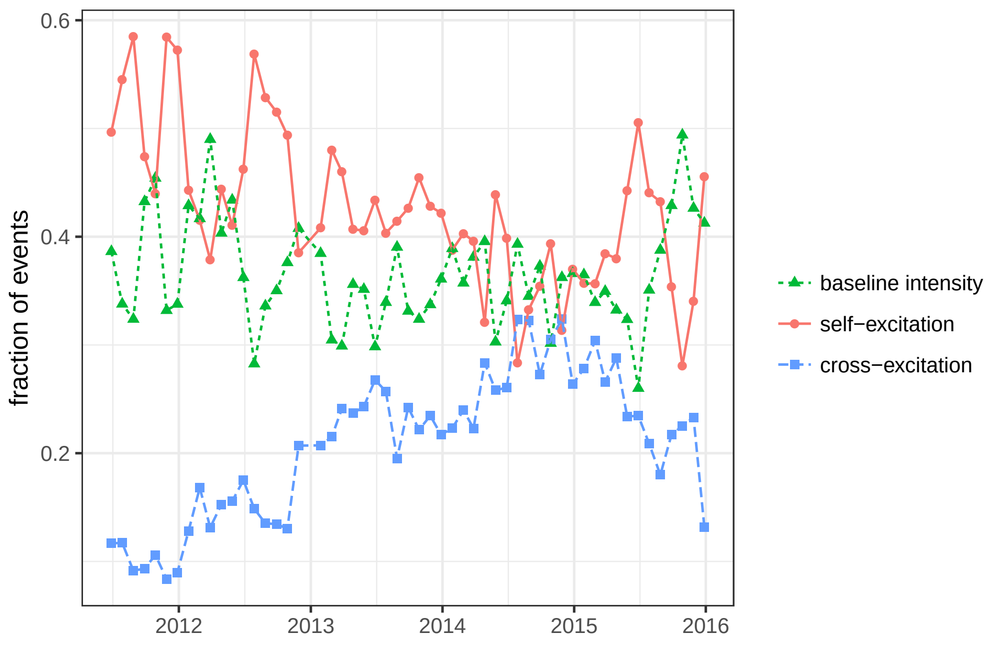

Figure 5.7: Fraction of illiquidity events attributed to random arrivals (baseline intensity), illiquidity spirals (self-excitation) and illiquidity spillovers (cross-excitation).

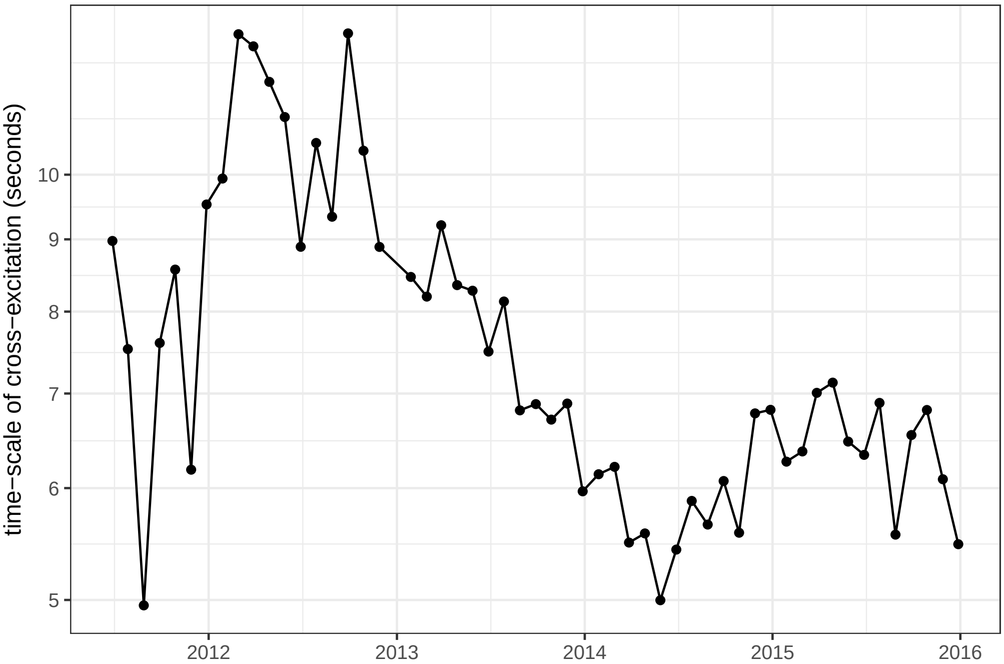

Figure 5.8: Time-scale of illiquidity spillover (cross-excitation). Measured as mean over monthly medians of $1/\beta^{ij}$, $(i \neq j)$.

For each estimation (i.e. day and pair of bonds), we calculate the fraction of illiquidity events of bond i attributed to self-excitation, cross-excitation from the paired bond j, and the baseline intensity. The averaged fractions, displayed in Figure 5.7, indicate that on average approximately one third of the events arrive randomly and are not due to self- nor cross-excitation, irrespective of time or the maturity of bonds. This fraction may be relatively large because of the rather lax threshold we have defined in Section 5.3.3. The fraction of self-excited events is between one half and one third of events, pointing to the presence of illiquidity spirals. While there is a decreasing trend, this is fully compensated by cross-excitation (and thus illiquidity spillovers) since the fractions by the definition in equation (5.5) sum up to 1. It is worth noting that the fraction of events that are due to cross-excitation has roughly doubled during our sample from ~ 0.1 in 2011 to ~ 0.2 in 2015.

Having established the presence of illiquidity spillovers, we study its typical timescale. The parameter $1/\beta^{ij}$ $(i \neq j)$ gives the decay time-scale of the cross-excitation kernel and thus an estimate for the timescale of illiquidity spillovers. Figure 5.8 shows that typically illiquidity spillovers take place on a time-scale of a few seconds. We further observe very fast spillovers during the sovereign bond crisis in 2011 and a general decrease from 2012 to 2015. We conjecture that the first feature is due to a more fragile market during the crisis while the decreasing trend since 2012 is related to advances in technology and therefore faster updating of quotes with respect to information arriving from other bonds. Table 5.1 gives further information on the average time-scale of illiquidity spillovers distinguished by the maturity bin of the originating and the affected bond. Spillover is faster from shorter maturity bonds that are typically more active in terms of number of limit order book events.

Finally we test whether illiquidity spillover is faster for bonds within the same maturity

Table 5.1: Time-scale of illiquidity spillover (cross-excitation) in seconds. Measured as mean over the monthly medians of $1/\beta^{ij}$, $(i \neq j)$ for each maturity pairing.

in seconds		from		
		5y BTP	10y BTP	30y BTP
	5y BTP	4.19	7.36	10.94
to	10y BTP	5.87	6.54	10.16
	30y BTP	7.18	8.46	7.70

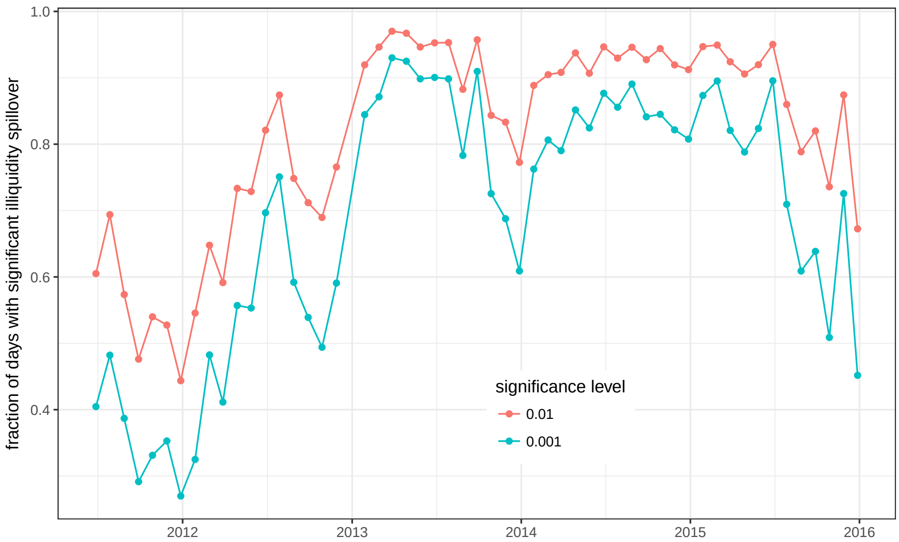

Figure 5.9: Fraction of estimations for which cross-excitation α^{ij} is significantly different from zero at the 1% and the 0.1% confidence level.

bin and we find that spillover across similar maturities is on average approximately two seconds faster than across different maturity bins. The difference in monthly medians is significant at all standard confidence levels.

5.3.5 Robustness checks

In this section we perform two robustness checks to further establish our finding of illiquidity spillover. Specifically, in Section 5.3.5 we show that for a large share of our estimations, illiquidity spillover is statistically significant and in 5.3.5 we show that spillover is still present when controlling for the arrival of exogenous news.

Significance of parameter estimates

Following Bowsher (2007) we estimate confidence intervals of the parameter estimates from the inverse of the negative Hessian of the likelihood function when assuming normality. This assumption is supported by the central limit theorem for maximum likelihood estimators of stationary point processes, including Hawkes processes (see Ogata (1978); Ozaki (1979)). For every estimation in Section 5.3.4 we test whether the cross-excitation parameters α^{ij}, $i \neq j$ are significantly different from zero under the normality assumption. In Figure 5.9 we show the fraction of such significant α^{ij} averaged over all bond pairs and days in a month. Both at the 1% and the 0.1% significance level there is significant illiquidity spillover in at least 20% of observations, and in the period from 2013 to mid 2015 this fraction is on average 80%, strongly supporting our evidence for illiquidity spillover.

Announcement effects

The intraday pattern of liquidity shocks in Figure 5.5 shows that a sizeable fraction of illiquidity events occur close to announcements. To verify that our observations are not a mere artifact of announcement effects, we repeat our analysis for a subsample of 5 and 10 year BTPs in 2012. To avoid the peaks and any other regular effects, we drop from the observations in each day the half hour from 15 minutes before to 15 minutes after 11:00 and 14:30 each and the 10 minutes from 5 minutes before to 5 minutes after each full hour and append the remaining times. The results we obtain are in line with our findings above. There are less days with a sufficient number of events to perform an estimation and for the estimated time-series the quality of the fits is visibly worse due to appending. The fractions reported in Figure 5.7 change only slightly in that there is a shift from self-excited to baseline events by 5-10 percentage points. Also the time-scales of both self- and cross-excitation remain similar.

5.4 Conclusion

Illiquidity risk, especially dry-ups of liquidity and illiquidity spillovers, is a major threat to the functioning of financial markets and thus a concern for investors and regulators alike. We propose a new identification and modeling approach to detect when the reactions of market participants to illiquidity shocks are susceptible to becoming a self-enforcing circle that leads to dry-ups of liquidity and spillovers of illiquidity across different assets. This approach gives a directional measure of illiquidity spillovers and fragility to extreme illiquidity events.

Our empirical analysis focuses on the Italian sovereign bond market from 2011 through 2015, encompassing the core parts of the European sovereign debt crisis, the start of Quantitative Easing by the ECB and most recently the regime of extremely low bond yields. We find a strong presence of fast illiquidity spirals and illiquidity spillovers throughout the investigated period. Moreover the proportion of illiquidity spillovers has roughly doubled since 2011. Spillover happens at the time-scale of a few seconds and is significantly faster for bonds with similar maturities. It is especially fast during the sovereign bond crisis in 2011 but also gradually becoming faster from 2012 to 2015.

Our identification method is not restricted to illiquidity, but can be applied to any process that is based on a limit order book or similarly piecewise constant datasets. Furthermore one could easily adapt our method for forecasting purposes by basing it on estimators that rely only on information up to the current point in time. Variations in the strength of self- and cross-excitation could then serve as a measure for liquidity risk and as a warning indicator.

In this chapter we have considered the case of closely related assets traded on the same market. We leave the extension to spillover across asset classes or markets for future work.

Chapter 6

Venue choice in hybrid markets

6.1 Introduction

German sovereign bonds, generally known as *Bunds*, enjoy benchmark status for Europe as a safe asset and are considered the second most liquid sovereign bond market in the world after the U.S. Treasury bond market. While the vast majority of transactions in the Bund cash market are dealt over-the-counter (OTC), i.e. in bilateral negotiations, it effectively constitutes a *hybrid market*. That is Bund dealers have the choice to execute a trade either in the OTC segment of the market, or send it to an exchange. The scope of our study is twofold. First, we explore the decision process whether to trade in the OTC segment or on the exchange. A major reason why dealers trade OTC is because prices are advantageous. Thus, in a second step, we study differences in transaction costs across segments and their determinants.

The decision to trade in one segment or another is complex since dealers face advantages and disadvantages in executing trades on an exchange compared to OTC. To trade OTC, a dealer needs to enter bilateral negotiations with counterparties, thus incurring search costs and execution lags. On the other hand, OTC trades are typically not observed by other market participants and may thereby satisfy dealers' demand for opacity. Moreover, in a relatively opaque OTC market, different investors may pay quite different prices for the same asset at essentially the same time. The resulting price dispersion may vary in terms of the relative bargaining power of the market participant, their access to alternative trading opportunities and the quality of their information both about the fundamental value of the asset and about recent transactions. Trading on an exchange, in turn, guarantees certain and immediate execution. On the other hand, the trade is also immediately revealed to other market participants. Moreover, transaction costs are increasing with trade size so that large orders need to be split into smaller trades and executed sequentially when trading on an exchange.[1]

We study the relative importance of these driving forces by making use of a unique regulatory dataset which includes *all* transactions on Bunds made by German financial institutions from 2011 to 2016, a dataset very similar to U.S. TRACE for corporate bonds and to the Treasury TRACE data that FINRA is collecting since July 2017. We match these transactions with the full limit order book and trades of the interdealer platform MTS, the largest sovereign bond exchange in Europe. To understand the venue choice between OTC and exchange we look into a group of dealers with access to both MTS and OTC. For a sample of trades where trading in both venues is possible we estimate a probit model

[1] For a detailed overview of the theoretical advantages and disadvantages and complexity of trades in the OTC markets versus exchanges see Duffie (2012a).

for the probability that a dealer will trade on the exchange instead of over-the-counter. The crucial advantage of our dataset and setting over most similar studies is that in doing so we are able to define for any OTC trade the *contemporaneous* conditions of trading on the exchange.

We find that trading on the exchange is less likely (i) when transaction costs are high, i.e. when bid-ask spreads are wide and for large trades, and (ii) for younger bonds, especially the cheapest-to-deliver that is linked closely to the Bund futures contract and carries benchmark status. Additionally, we find that the required immediacy towards a trade, proxied by the occurrence of auctions and intraday volatility, raises the probability of trading on the exchange.

In a second step, we analyse the cross-sectional dispersion of prices negotiated at a particular time between the OTC segment and the exchange. Specifically, we measure the price discount or premium of an OTC trade with respect to potential execution of the same trade on the exchange at the same time. We first calculate the hypothetical trading cost of a given trade on the exchange, taking into account the state of the limit order book. We then define the *OTC discount* as the difference between the hypothetical trade price and the actual observed OTC price.

Our main finding is that the two markets are largely complementary because the MTS limit order book bounds the price discrimination in the OTC market as MTS represents an outside option to OTC trading in the sense of Duffie et al. (2005, 2007). Transaction costs are on average by $1.5 - 2$ basis points lower in the OTC market. This discount is economically significant given that the average half-spread on the exchange is only 3.6 basis points in the 10-year Bund. Since the main difference between both markets is immediacy, we interpret the OTC discount as a proxy for the cost of immediacy. However, we also observe a significant number of trades where dealers leave money on the table by avoiding the exchange. This suggests that transparency or opacity play an important role for venue choice. We assess the drivers of the OTC discount by regressing it on trade and bond-time characteristics. We find that larger trades and those in older bonds enjoy higher discounts with respect to exchange prices. Moreover, only $10 - 20\%$ of increases in the quoted bid-ask spread on the exchange are passed on to OTC trades. We also investigate whether being a dealer in the MTS market provides a competitive advantage. Indeed we find that transaction prices of MTS dealers are on average $0.7 - 0.9$ basis points lower than those of non-MTS dealers.

This chapter makes three main contributions relative to the extant literature. First, our unique data allow us to study the determinants of venue choice in a highly liquid hybrid sovereign bond market. While several previous studies have analyzed related questions for equity markets (Smith et al. (2001); Bessembinder and Venkataraman (2004); Friederich and Payne (2007); Carollo et al. (2012)), there is little to no evidence for bond markets. Two notable exceptions are Barclay et al. (2006) who study the choice between electronic and voice brokerage for U.S. Treasuries that go off-the-run and Hendershott and Madhavan (2015) who analyze U.S. corporate bonds on a request-for-quote platform. In contrast to these studies we can compare between exchange and OTC trading that is in line with the settings described in theoretical studies (Seppi, 1990; Grossman, 1992; Lee and Wang, 2017; Vogel, 2017). Moreover, we can explicitly characterize the importance of immediacy and opacity as determinants of venue choice which have not previously been analyzed for bond markets.

Our second contribution is a detailed analysis of the differences in transaction costs between the OTC and exchange segments. Our dataset allows for a full reconstruction of the limit order book. We can thus disentangle the determinants of OTC transaction costs by relating the OTC discount to trade- and bond characteristics.

Finally, we contribute to the literature by analyzing in detail the microstructure of the Bund market, which so far has been largely unexplored.[2] This is surprising given the benchmark role of the Bund market in the Eurozone and beyond.

Our analysis is relevant for academics, practitioners and regulators alike and especially so in light of MiFID II. The intention of the recently rolled out MiFID II regulation is to improve market conditions in and beyond European fixed-income markets by introducing transparency requirements as well as other measures aimed at shifting trading activity towards organized trading platforms.[3] This study gives insights in the incentives of traders faced with such a venue choice in an affected market and will help understanding the effects of the regulation in the years to come. There is a vast and growing theoretical literature on the design and modeling of financial markets that we relate to, and particularly hybrid markets. While the early literature focused on modeling equity markets with off-book segments (Seppi (1990); Grossman (1992)), there is now a new interest in this direction in the light of hybrid fixed-income markets (Lee and Wang (2017); Glode and Opp (2017)) and the effects of introducing benchmarks (Duffie et al. (2017)) or hybrid markets (Vogel (2017)). Our study provides a natural laboratory and allows to empirically test some of the predictions therein.

The remainder of this chapter proceeds as follows: Section 6.2 provides an overview of related studies and the theoretical literature. In Section 6.3 we introduce and motivate our hypotheses and Section 6.4 presents the market setting and data. Section 6.5 examines the venue choice between trading on-exchange or over-the-counter, whereas in Section 6.6 we study the differences in transaction costs across the two markets as well as their drivers. Section 6.7 concludes. Furthermore in the Appendix 6.A we provide an exploratory analysis of price response to over-the-counter transactions

6.2 Literature review

Most stock exchanges have or used to have so-called *upstairs* or *off-book* segments where (large) trades could be concluded away from the limit order book of the exchange (the *downstairs market* or *on-book*) in bilateral negotiations effectively resembling over-the-counter market structures. This inspired a series of both theoretical and empirical research on hybrid markets that has since been expanded beyond the original equities upstairs market context.

On the theoretical side Grossman (1992) and Seppi (1990) provide concurrently compatible motivations for trading in the upstairs market: in the model of Seppi (1990) an institution has a trading need either because she needs to rebalance her portfolio (termed *liquidity-motivated*) or due to an endowment with private information (*information-motivated*). While a liquidity provider in the downstairs market needs to quote a wider bid-ask spread to insure herself against the risk of being adversely selected by information traders, dealers in the upstairs market can screen their customers and offer better conditions to uninformed clients. That is a liquidity-motivated trader may find it optimal to give up its anonymity by interacting with a dealer in the upstairs market, whereas information-motivated traders resort to the downstairs segment. This implies that dealers are able to infer (at least to some

[2]Upper and Werner (2002) studies the information content of the Bund futures and cash market in 1998.

[3]Since January 3, 2018 directive 2014/65/EU - Markets in Financial Instruments Directive II (MiFID II) is effective for all European markets. It includes provisions for pre- and post-trade transparency, separation of transaction and related services fees among many others. Our sample period ends before the introduction of MiFID II.

degree) the information state of their clients through a relationship build from repeated interactions. Indeed Seppi (1990) describes implied *no-bagging* agreements incentivizing customers not to trade the same asset again too soon after an upstairs trade and therefore to reveal their full liquidity need to the dealer or else face worse conditions on future trades.

Grossman (1992) argues that upstairs trades need not necessarily be uninformed. In his model limit orders in the downstairs market represent *expressed order flow*, whereas some traders, for fear of being speculated against, are unwilling to reveal their full trading needs in the observable limit order book. Dealers in the upstairs market acquire knowledge about this *unexpressed order flow* through interaction with their clients. Thus the upstairs market is able to facilitate trades that would otherwise not happen in the downstairs market. As in the model of Grossman (1992) customers choose whether to be active in the upstairs or downstairs segment, this may lead to externalities where, because of a clustering of traders in the upstairs market, bid-ask spreads in the downstairs market are wide and activity is low.

Similar themes have resurfaced recently in the wake of increased attention for the market structure of fixed-income markets and regulatory initiatives. The model of Lee and Wang (2017) is similar to Seppi (1990) with informed and uninformed investors that can trade on an exchange or over-the-counter. However the model of Lee and Wang (2017) does not focus on differences in trade size and instead features *price discrimination* based purely on the (assumed) reputation of an investor. Therefore dealers *cream-skim* the order flow from uninformed investors on the over-the-counter market by offering a lower bid-ask spread to investors deemed uninformed whereas informed investors are driven to the exchange. Crucially their model explains why smaller orders are traded OTC despite the availability of liquid exchanges and how OTC trading is predominant for standardized and frequently traded assets.

Another set of articles compare OTC and exchange markets but without considering them jointly. Glode and Opp (2017) present a model with a sequential search process for counterparties, arguing that search frictions in OTC markets can promote higher welfare through rents that encourage expertise acquisition. Also Malamud and Rostek (2017) discusses the efficiencies gained from a more general setup of decentralized markets, while Malinova and Park (2013) compare the price impact in dealer and limit order market.

Duffie et al. (2017) study the introduction of benchmarks to OTC markets. While the benchmark lowers dealers' profit margins, it also increases market participation and thus increases welfare. Similarly Vogel (2017) studies the effects of introducing a request-for-quote platform to an OTC market, predicting that it lowers average costs for all traders and providing sufficient conditions for the increased participation due to lower costs to outweigh the gains foregone by the dealers.

This theoretical literature is matched with a vast empirical literature on hybrid equities markets,[4] bond markets (Barclay et al. (2006); Hendershott and Madhavan (2015)) and the index CDS market (Riggs et al. (2017)).

The papers most related to the first part of our analysis on venue choice are Smith et al. (2001); Barclay et al. (2006); Hendershott and Madhavan (2015); Riggs et al. (2017). Smith et al. (2001) finds mechanism choice related to several measures of liquidity as well as trade size. Barclay et al. (2006) focus on the effects of U.S. treasuries going off-the-run and argues that venue choice depends mostly on trading volume on each venue (and thus

[4]E.g. of NYSE (Keim and Madhavan, 1996; Madhavan and Cheng, 1997; Madhavan and Sofianos, 1998), Toronto Stock Exchange (Smith et al., 2001), Helsinki Stock Exchange (Booth et al., 2002), Paris Bourse (Bessembinder and Venkataraman, 2004) and London Stock Exchange (Friederich and Payne, 2007; Carollo et al., 2012)

the ability to match trading interests) and information asymmetry. Since for treasuries going off-the-run the latter can be assumed constant, they identify the matching ability as a main driver. The main difference in our setup with respect to those two papers is that we consider a predominantly over-the-counter market that does not feature on-the-run effects. Instead in Hendershott and Madhavan (2015); Riggs et al. (2017) not only parameters of the market setting are different, but the market structures in themselves. Hendershott and Madhavan (2015) compares transactions in the voice OTC market for U.S. corporate bonds to one-sided electronic auctions on the multi-dealer request-for-quote (RFQ) platform MarketAxess, considering also cross-sectional differences in bond characteristics. While a central limit order book exists in the setting of Riggs et al. (2017) it is so illiquid that their comparison is between RFQ and request-for-streaming (RFS) mechanisms, the latter coming close to an OTC segment. Their ability to identify counterparties involved in a trade allows them to find that e.g. asset managers are more likely to trade via RFQ.

The second part of our study on transaction cost differences compares observed transaction prices to hypothetical exchange prices derived from the knowledge of the full limit order book. We are only aware of a similar analysis in Bessembinder and Venkataraman (2004) who find upstairs trades on Paris Bourse advantageous with respect to exchange prices, thereby supporting the model of Grossman (1992). Our analysis does not suffer from two restrictions present in their setup: first their data does not include information on cancelled limit orders, thus overestimating liquidity present in the book. Second in our case there are no restrictions on the price range or size of OTC trades.

Other papers compare transaction costs without using the full limit order book information. Dunne et al. (2015) compare trades in European sovereign bonds on the retail request-for-quote platform Bondvision to quotes on the interdealer exchange MTS. They find that the retail quotes and trades occur at prices on average 1-2 basis points better than in the interdealer market. Valseth (2015) quantifies the transaction cost differences across mechanisms in Norwegian government bonds by comparing the median bid-ask spread of quotes in the OTC segment to that of the exchange and finds a price improvement in the OTC market of 2-6 basis points.

Specifically with regard to price discrimination, Collin-Dufresne et al. (2017) finds significant differences in transaction costs both across the dealer-to-client and dealer-to-dealer segments and across trading protocols for interdealer trades in the index CDS market.

Let us note that multimarket trading relates also to a huge literature on dark pool trading. For a recent overview we refer to Menkveld et al. (2017). They create a pecking order of equity trading venues by the transparency and immediacy they offer and show how surprises lead to shifts in towards lit markets that offer higher immediacy at a higher cost.

6.3 Hypotheses

In our setup a subset of market participants, referred to as *MTS dealers*, have the choice to trade either over-the-counter (OTC) or on an exchange, the inter-dealer platform MTS. All other participants (*non-MTS dealers*) can only trade on the OTC market. Trading on the exchange is immediate and transparent, i.e. the trade is observed by other market participants. The OTC market instead is opaque, but trading OTC incurs search costs (time delays and resources).

In the first hypothesis we consider the drivers of venue choice, i.e. when is a MTS dealer more likely to trade on the exchange than over-the-counter? Formally we test the following hypothesis:

Hypothesis 1. *MTS dealers are more likely to execute a trade on exchange instead of OTC when a) the bid-ask spread on the exchange is small, b) the size of the trade is small, c) a bond is older and not the cheapest to deliver for the current futures contract, and d) required immediacy is higher.*

The first two drivers relate to the cost of a trade: naturally, the higher the bid-ask spread on the exchange, the more costly is a trade and less the incentive of a trader to transact on the exchange. In addition to the bid-ask spread quoted at the best, trades that are larger than the amount quoted at the best also incur additional costs due to executing at deeper levels of the limit order book, an effect known as 'walking up the book'. By the model of Grossman (1992) OTC dealers are able to provide better prices by tapping into pools of unexpressed liquidity while in Seppi (1990) larger trades are associated with uninformed liquidity traders that receive better quotes from dealers.

The age and cheapest-to-deliver status are related to the price discrimination arguments of Lee and Wang (2017). If a bond has special status by being cheapest-to-deliver for the current Bund future, i.e. linked to a highly liquid instrument, the adverse selection risks associated with it are lower. This allows dealers to offer more competitive spreads in the OTC market and attract more trades.

With regard to immediacy the advantage of the exchange is natural as trading requires no search process at all. We proxy the immediacy in the market through the occurence of auctions and intraday volatility. We posit that when a bond is newly issued or re-opened, this creates need for immediacy through various channels: First only an auction truly reveals the ability of the issuer to place new debt and is thus observed with great attention by traders in related securities, as they might need to adapt their portfolios on the arrival of new information. Second, even in the absence of informational effects, many participants (e.g. insurers) are required to shift or roll over their holdings and thus create additional demand with dealers that might not be fully predictable. Along the same line, also dealers that participate in the auction only learn after the fact about their allotted amounts and might need to correct their holdings on the secondary market within the same day. Higher intraday volatility implies a higher risk of adverse price movements. A MTS dealer might therefore prefer the secure and immediate execution of a trade on the exchange over a potentially longer search across counterparties in the OTC market when intraday volatility is high.

In the next hypotheses we consider only over-the-counter trades. As we have the full information of the order book on the exchange, we are able to compare the price of the OTC trade with the hypothetical cost had the initiator chosen to trade the same amount at the same time on the exchange. We term the difference between both prices (or costs) *OTC discount.* Formally we test:

Hypothesis 2. *Trades that take place OTC on average have a positive OTC discount, i.e. their price is favorable in comparison to the attainable price of executing the same trade on the exchange.*

Hypothesis 2 provides a test for the prediction of Grossman (1992) that the OTC (upstairs) segment is able to make available more liquidity than what is visible on the exchange (downstairs). Furthermore, as observed trades are also a result of the conditions encountered by traders, the hypothesis also talks to the predictions of Seppi (1990) and Lee and Wang (2017) that the OTC segment provides favorable quotes to uninformed investors (a feature themed price discrimination in Lee and Wang (2017)). That is an informed trader will find himself less able to trade OTC (or at worse conditions) and therefore prefer the exchange, while uninformed investors flow to the OTC segment.

We interpret the OTC discount as a proxy for the cost of immediacy, i.e. the price-improvement that dealers forgo when they trade on the exchange.[5] Having established its prevalence, we formally test the drivers of OTC discounts:

Hypothesis 3. *OTC discounts are larger a) when the bid-ask spread is larger, b) when the size of the trade is larger, c) when a bond is older and d) for traders that have access to the exchange market.*

The first two determinants relate to the concept that large transaction costs on the exchange (proxied by wide bid-ask spreads and large orders that incur additional costs from walking up the book) allow for larger discounts in over-the-counter trades. By the arguments of Hypothesis 1 also older bonds tend to be less liquid and thus there is more play for OTC discounts.

Finally MTS dealers always have an outside option to trading OTC (i.e. on the exchange) and thus we also expect them to achieve a higher discount when trading over-the-counter (Duffie et al., 2005, 2007).[6]

6.4 Market setting and data

Our empirical analysis is based on trades in German federal government securities, typically referred to as *Bunds*. In section 6.4.1 we give an overview of the primary and secondary market for these securities and the related Bund futures market. Section 6.4.2 then introduces our dataset and provides some descriptive statistics of quantities of interest in our analysis.

6.4.1 Bund market

German sovereign debt securities enjoy benchmark status in the Eurozone and worldwide as a liquid and safe asset. They exist as 6- or 12-month zero coupon Treasury discount papers ("Unverzinsliche Schatzanweisungen", *Bubills*), 2-year "Bundesschatzanweisungen" (*Schaetze*), 5-year "Bundesobligationen" (*Bobls*) and 10- and 30-year "Bundesanleihen" (*Bunds*).[7] In this study we will focus on 2-year Schaetze, 5-year Bobls and 10- and 30-year Bunds and, where not explicitly mentioning maturities, we will intend all of them when refering to Bunds from here on.

German government securities are issued regularly by the German finance agency ("Deutsche Finanzagentur", *DFA*) either as new issues or as reopenings of already issued bonds. Participants in this primary market are the members of the *Bund issues auction group*, a group of currently 36 international banks that commit to subscribing to a certain minimal amount of the total annual issuance. Auction days are announced well in advance and the tender process runs until 11:30 a.m. on the day of the auction, after which the

[5] Trading at a negative discount is a priori irrational. It can occur nonetheless, e.g. when trading relationships are at play, when one requires opacity and is unwilling to display trading intentions publicly or when a non-MTS dealer does not have access to the exchange and is unable (within her time-constraints) to locate a MTS dealer willing to offer a better price. Additionally trading OTC might give access to added services such as contemporaneous hedging on the futures market.

[6] This could also be related to the characteristics of MTS dealers instead of their access to the exchange. We are unable to disentangle whether such an advantage is due to their MTS status or other factors, such as e.g. the skills of their trading desks.

[7] There are also inflation-linked Bobls and Bunds, which we do not consider in this study. We also do not consider any regional-issued debt, such as *Laender* bonds or debt titles from supranationals with a federal guarantee, e.g. by Kreditanstalt für Wiederaufbau (KfW).

allotment decision is made immediately and the results published.[8] Even though Bunds are actively traded and priced through the secondary market and the Bund futures contracts, information about the quantity and conditions at which the sovereign is able to issue is only revealed during the auction, which is therefore regarded with great interest by traders in all related securities and markets.

The secondary market for German federal debt titles is predominantly an over-the-counter market. A survey by DFA among the members of the Bund issues auction group pegs daily trading volume at more than 17 billion EUR.[9] Besides the over-the-counter market there exist several trading platforms, most of which are aimed at the retail market. Notwithstanding a relatively large number of trades in the retail market, trade sizes are typically too small to carry any economic significance. The most important interdealer platform is MTS, which is operated as a fully electronic limit order book during the hours from 9 a.m. to 5:30 p.m. Trading activity on MTS accounts for only a small market share in terms of traded volume as even for liquid bonds typically only a few trades per bond and day are recorded; however quoting is very active and one usually finds limit orders in excess of 100 million EUR on both the bid and ask side of the book for most securities. This, in conjunction with the availability of MTS data for market participants and researchers, has given MTS a benchmark function for European sovereign bond markets.[10]

A primary contributor to the liquidity of the secondary *cash* market for German Bunds is the even more liquid Eurex futures contract. There exist futures contracts for 2-year Schaetze, 5-year Bobls and 10-year and 30-year Bunds, with most activity in the 10-year Bund futures. Turnover across all futures was almost 32 trillion EUR in 2017, more than seven times the turnover in the cash market, with a minimum size of 100,000 EUR and minimum tick sizes corresponding to $0.5 - 2$ basis points depending on the contract. Trading activity is generally concentrated in the contract with the nearest delivery day, which is around the 10th of each March, June, September and December. By construction ca. 3-5 bonds are *deliverable* for each contract and one bond is the *cheapest-to-deliver* and therefore its price is tied to the one of the futures through a close arbitrage link.[11] It is worth pointing out that *physical delivery* of the futures on the delivery day is rare and most contracts are closed by entering an opposite position. That implies that, notwithstanding the comparatively more active futures market, anyone wanting to own Bunds (e.g. banks or insurances for regulatory reasons) or to enter an arbitrage position still needs to be active on the cash market. The next section describes our data on this market.

6.4.2 Data preparation and descriptives

Our study is based on a regulatory dataset of trades by German financial institutions which we connect to the full limit order book data from the interdealer exchange MTS. In

[8] Auction group members can place competitive and non-competitive bids. The former are alloted in full at the bid price up to the lowest accepted price and the latter at a weighted average price of the accepted competitive bids. For more details regarding the auction process, auction schedule, members of the Bund issues auction group and auction results we refer to the DFA website: `https://www.deutsche-finanzagentur.de/en/institutional-investors/primary-market/`.

[9] See `https://www.deutsche-finanzagentur.de/en/institutional-investors/secondary-market/`. Our sample captures about 15% of this trading activity.

[10] Dufour et al. (2004) provides a detailed description of the MTS dataset and Darbha and Dufour (2013) give an overview over market structure and liquidity. MTS data have been used and validated in numerous studies at the European level, an incomplete list of which includes e.g. Beber et al. (2009); Pelizzon et al. (2016).

[11] Contractual details for the futures can be found at `http://www.eurexchange.com/exchange-en/products/int/fix/government-bonds/Euro-Bund-Futures/14770`. Trading hours last from 8 a.m. to 10 p.m. and thus exceed those of MTS.

order to make full use of both datasets we restrict our analysis to the period from June 2011 through December 2016 and as outlined above our sample consists of 2-year Schaetze, 5-year Bobls and 10- and 30-year Bunds, which we will also collectively refer to as Bunds.

The regulatory transactions data is based on reporting requirements of German financial institutions mandated by the German Securities Trading Act ("Wertpapierhandelsgesetz", *WpHG*) and the respective regulation "Wertpapierhandel-Meldeverordnung" (*WpHMV*). It includes any transaction by the reporting institutions in a wide set of securities, including German government bonds, and contains information on the price, size and time of the trade and a flag that indicates whether a trade was over-the-counter or the platform it was made on. We further have anonymized identifiers for the reporting agent and the counterparty of a trade, where the identifier for the counterparty can be missing when the counterparty is non-German.[12]

Our dataset from the interdealer exchange MTS contains all trades thereon as well as the full limit order book information on all executable quotes. This allows us to combine both datasets and from the MTS dataset we compute for any trade in the WpHMV dataset the quoted bid-ask spread at the same time and the price a trade of the same size would have incurred on MTS had it been transacted there. The WpHMV data do not include information on whether a trade was buyer- or seller-initiated. Hence, we infer the order sign of each trade by comparing its price to the contemporaneous midprice on the exchange, following Bessembinder and Venkataraman (2004); Eisler and Bouchaud (2016).[13] For each bond and for each day, we calculate intraday volatility based on five-minute returns of MTS mid quotes. Where the ID of the initiator of a trade is known, we also append a dummy that indicates whether she has access to the MTS market, which is inferred from her other trades. Finally we also control the exchange flag of the WpHMV data by matching all trades to the full set of trades of MTS. Lastly we add information on auctions, amount outstanding and bond status as eligible or cheapest-to-deliver for the futures contract from DFA, Bloomberg and Thomson Reuters Eikon.

Our full sample contains almost 470,000 trades across 237 German federal bonds. The first row in Table 6.1 reveals that most of these trades are very small, with the 25% percentile at only 100,000 EUR of nominal amount. This is by far too small for trading on MTS, where the required minimum trade size is 2 million EUR. Our intention is to compare trades where there is a choice between trading over-the-counter and on MTS. Therefore we limit our sample to the set of trades where trading in both venues is an economically viable option. We refer to those as trades where *MTS is possible*. More specifically, it requires a minimum trade size of 2 million EUR and the trade must be initiated by a MTS dealer during MTS trading hours.[14] The last two rows of Table 6.1 show the size distribution of trades on MTS, once for the set of trades we observe through the WpHMV dataset and below for the set of all trades on MTS as reported in the MTS dataset. The distribtion is very similar across both sets indicating that, even though we capture only roughly 6% of MTS trades, our dataset is representative of the market as a whole. We also observe that MTS trades tend to be rather small, with the 95% percentile

[12]For a detailed description of the dataset we refer to Chapter 7 and the text of the law and regulation, for which a non-binding English translation is provided at `https://www.bafin.de/SharedDocs/Veroeffentlichungen/EN/Aufsichtsrecht/Gesetz/WpHG_en.html` (Section 9 therein) and `https://www.bafin.de/SharedDocs/Veroeffentlichungen/EN/Aufsichtsrecht/Verordnung/WpHMV_en.html?nn=8379960` respectively.

[13]We are unable to calculate some of these measures e.g. when an OTC trade takes place outside of MTS opening hours or when the size of a trade exceeds the quantity available in the MTS limit order book, either because a trade is too large or because the order book is depleted.

[14]We also exclude trades so large that they would exceed the liquidity available in the limit order book of MTS.

Table 6.1: Statistics of trade size: This table shows the distribution of trade size measured in million EUR. *full sample* refers to all trades in the WpHMV sample. The block *OTC trades where MTS is possible* refers to OTC trades from our WpHMV sample that could have been traded on MTS. The row *all* refers to all such trades and the row $\leq$ *25M EUR* imposes a maximum size of 25 million EUR. The block *MTS trades* refers to trades on the interdealer exchange MTS. *WpHMV data* describes the statistics of MTS trades that we observe in our WpHMV data. *MTS data* refers to all trades on MTS from the MTS dataset.

	Mean	Std Dev	5 Pcl	25 Pcl	50 Pcl	75 Pcl	95 Pcl	# Obs
full sample	7.01	20.56	0.00	0.10	1.00	5.00	32.10	469,689
OTC trades where MTS is possible								
all	16.95	31.64	2.00	4.00	5.59	18.00	58.00	49,519
$\leq$ 25M EUR	7.88	6.41	2.00	3.20	5.00	10.00	25.00	41,888
MTS trades								
WpHMV data	6.92	5.81	2.00	4.00	5.00	10.00	10.00	2,275
MTS data	6.07	4.72	2.00	2.50	5.00	10.00	10.00	37,943

at 10 million EUR compared to 32 million EUR in our full sample. This is expected as larger trades on MTS consume liquidity from higher levels of the limit order book, thus incurring a higher cost and deincentivizing dealers from making very large transactions on the exchange. Therefore we will also consider a subsample of the trades where MTS is possible with an upper limit of 25 million EUR.[15] Indeed we then find that the distribution of trade size for this subset of OTC trades is comparable to the one for trades on MTS.

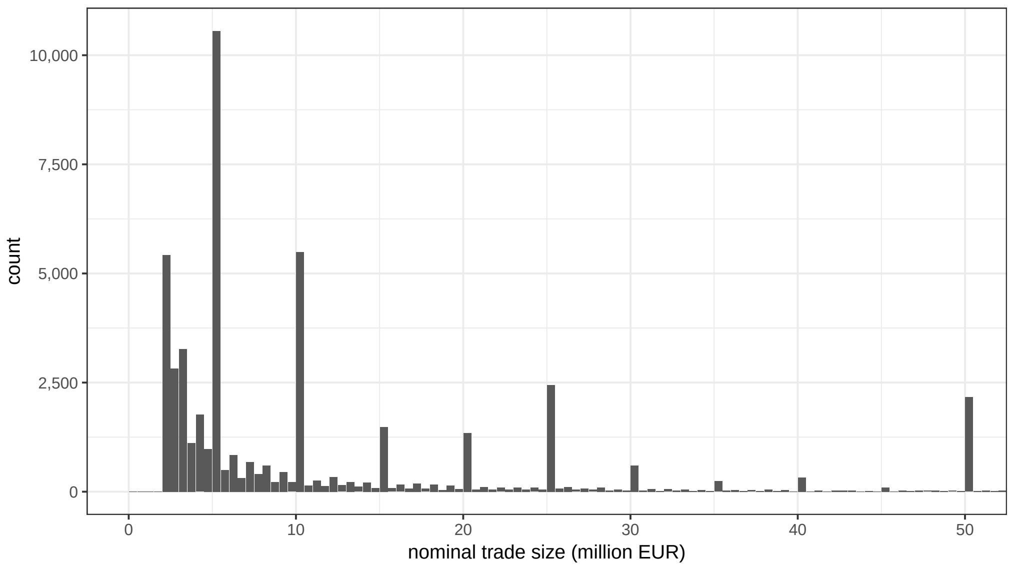

Figure 6.1: Histogram of trade size when MTS is possible: number of trades in bins of trade size measured as the nominal amount of the trade. Binwidth is 0.5 million EUR and the sample consists of OTC trades where MTS trading is possible and MTS trades, therefore imposing a minimum size of 2 million EUR. A preferences for "round" amounts is present for both OTC and MTS trades.

To visualize the distribution Figure 6.1 shows the histogram of trade size for all trades where MTS was possible.[16] Beyond the evidence in Table 6.1, Figure 6.1 clearly reveals

[15] 25 million EUR corresponds to the 99.4% percentile of trade size on MTS.

[16] Both OTC and MTS trades are shown, i.e. Figure 6.1 corresponds to the union of the subsets described in the second and the fourth row of Table 6.1.

a preference for 'round' amounts of trade size, such as e.g. 5, 10 or 25 million EUR. This feature is present for both OTC and MTS trades and is in line with trading in the Bund cash market being still predominantly 'manual', as expected for an OTC-dominated market.[17]

Table 6.2: Statistics of bid-ask spread by maturity: Descriptive statistics of the average daily MTS bid-ask spread of the on-the-run German federal bonds in each maturity bin. Given in basis points and based on average daily values of bid-ask spread and winsorized at the 1% level from above.

	Mean	Std Dev	5 Pcl	25 Pcl	50 Pcl	75 Pcl	95 Pcl
2-year Schatz	3.67	1.86	1.68	2.56	3.14	4.12	7.51
5-year Bobl	6.80	3.75	3.65	4.74	5.87	7.75	12.02
10-year Bund	7.22	3.94	3.91	4.98	6.14	8.15	13.67
30-year Bund	54.90	37.65	23.22	37.55	46.78	60.20	99.30

Finally we also provide summary statistics for the quoted bid-ask spread on the MTS exchange. Table 6.2 and Figure 6.2 give the descriptives and time evolution of the average daily quoted bid-ask spread on MTS for the on-the-run bond of each maturity bin. In line with the previous literature we find that bid-ask spread is increasing with maturity, ranging from on average 3.7 basis points in 2-year Schaetze to 7.2 basis points in the 10-year Bund. 30-year Bunds are less liquid and have a much wider bid-ask spread with a mean of 55 basis points.

The time evolution of bid-ask spread in Figure 6.2 is smoothened to the weekly level and indicates the larger trends during our sample period ranging from June 2011 to December 2016. The European sovereign bond crisis is discernible in 2011-2012 whereas the peak at the end of 2016 might be related to low trading activity and scarcity effects at the end of the regulatory year.

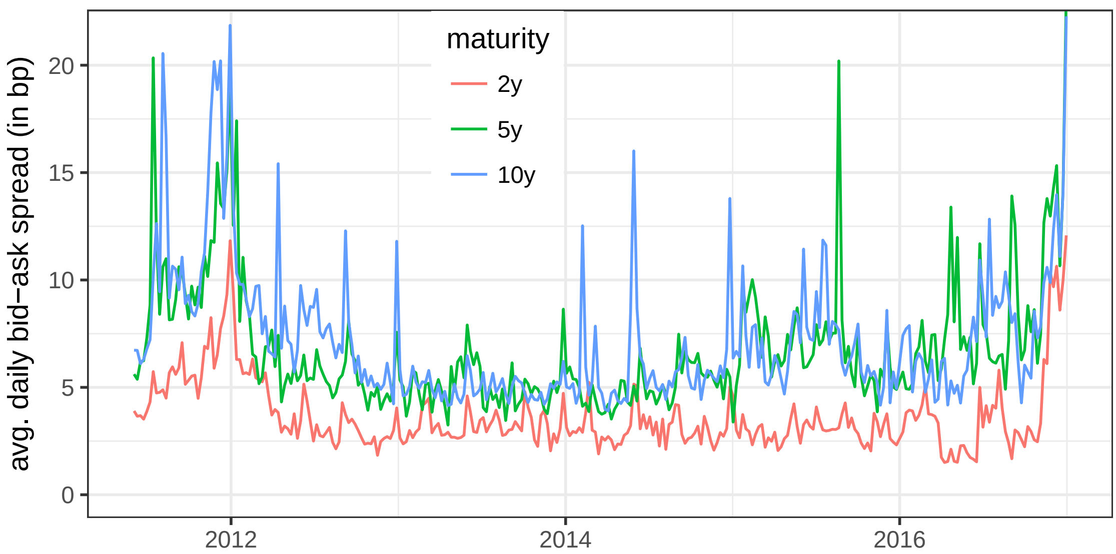

Figure 6.2: Bid-ask spread: Average daily quoted bid-ask spread on MTS of the on-the-run of each maturity bin. Aggregated to weekly level and given in basis points.

[17]Round amounts are associated with manual trading, whereas automated trading algorithms tend to make use of the continuity of possible sizes. See Lallouache and Abergel (2014) for an example in the context of a tick size reduction.

6.5 Venue choice

Hypothesis 1. *MTS dealers are more likely to execute a trade on exchange instead of OTC when a) the bid-ask spread on the exchange is small, b) the size of the trade is small, c) a bond is older and not the cheapest to deliver for the current futures contract, and d) required immediacy is higher.*

In order to test Hypothesis 1, we estimate a probit model

$$MTS_n = f(v_n) \tag{6.1}$$

at the level of single trades (indexed by n), where our dependent variable MTS_n is a dummy that is 0 for a trade in the over-the-counter segment and 1 for trades on the exchange MTS and independent variables include trade, bond and bond-time characteristics summarized in the vector v_n.

Our sample consists only of trades that are able to take place in either section of the market, i.e. we ensure that each trade is qualified for MTS. That is our sample consists of trades with a minimum size of 2 million EUR that took place during MTS trading hours and were initiated by a trader that has access to the MTS market, leaving us with a sample of 51,565 trades of which 2,270 are on MTS (4.4%).[18] While it is in principle possible to trade for the full quantity quoted on the limit order books of MTS, this would be very costly. Therefore, we also repeat our estimations for a subsample of trades below a certain threshold above which trading on MTS is unlikely, chosen as 25 million EUR of nominal value.[19] Finally, we also repeat our estimations for the subsample of trades initiated by MTS-dealers that are obliged to report and where thus we observe the complete set of their trading activity.[20]

Let us stress that we are unable to observe the part of the decision process of whether to trade at all or not. Especially in the OTC market a trader might be unable to find a willing counterparty to a transaction or only at a cost that outweighs the expected benefits of the trade so that the originally intended trade is no longer desirable. E.g. Hendershott and Madhavan (2015) report that on average only $51.0 - 73.4\%$ of electronic auctions (corresponding to multi-dealer requests-for-quote) lead to trades and $8.2 - 14.7\%$ of auctions do not receive a quote from a dealer.

The independent variables in our model include trade characteristics, bond characteristics and bond-time characteristics. Trade characteristics are the size of a trade, given as the natural logarithm of the nominal value of the trade in EUR, and the bid-ask spread prevailing on MTS at the time of the trade, given in basis points. As bond characteristics we include a dummy for the maturity of the bond, i.e. whether it is a 2-, 5-, 10- or 30-year title. Bond-time characteristics are the age of a bond, i.e. the time since its issuance in years, the nominal amount outstanding in billion EUR, a dummy that is 1 when a bond is the cheapest to deliver for the current futures contract on the day of the trade and 0 else, and the intraday volatility in the same bond at the same time and day. We also add a

[18] We also exclude trades when the bid-ask spread on MTS was prohibitively high, setting the threshold as 100 basis points, corresponding to the 99%-percentile across all bonds and the 95%-percentile for 30-year Bunds, the least liquid bonds in our sample.

[19] Our results are robust for the choice of threshold between 15 or 50 million EUR. 25 million EUR corresponds to the 94% threshold of trade sizes in our full sample and the 99.4% threshold of trade sizes on MTS.

[20] That is because our sample also contains trades by institutions that have no reporting obligation but where an identifier is provided. For those instutions we then only observe their trades with German counterparties.

dummy that is 1 when there was an auction affecting the bond on the same day, and 0 for non-auction days.[21]

Table 6.3: Probit for market choice: Marginal effects of a probit model for the choice of trading on exchange or over-the-counter. Sample are trades where MTS is possible, i.e. initiated by MTS dealers during MTS hours with a minimum size of 2 million EUR. Dependent variable is 1 for trades on the exchange MTS and 0 for OTC transactions. Explanatory variables include MTS bid-ask spread at the time of the trade, natural logarithm of nominal trade size in EUR, bond age, amount outstanding of the bond, intraday volatility (on MTS) and dummies for being cheapest-to-deliver and days when the same bond is (re-)issued. The baseline specification corresponds is for 10-year bonds and we including dummies for 2-year, 5-year or 30-year Bunds. Z-scores are given in brackets where standard errors are clustered at the dealer level and *, ** and *** denote significance at the 10%, 5% and 1% level respectively.

			trade size $\leq$ 25M EUR	
				initiator is reporting
bid-ask spread (bp)	-0.0052***	-0.0052***	-0.0059***	-0.0041***
	(-4.7920)	(-4.8270)	(-4.7054)	(-3.0564)
trade size (log)	-0.0199***	-0.0206***	-0.0100	-0.0106
	(-2.8073)	(-2.8378)	(-1.2070)	(-1.1134)
dummy: 2-year Schatz	0.0514***	0.0540***	0.0646***	0.0568***
	(5.5271)	(5.6932)	(5.8192)	(3.7978)
dummy: 5-year Bobl	0.0052	0.0062	0.0072	0.0107*
	(1.2848)	(1.5438)	(1.5176)	(1.7604)
dummy: 30-year Bund	0.4978***	0.4896***	0.5152***	0.4386***
	(4.9039)	(4.7754)	(4.8014)	(2.8907)
bond age (years)	0.0019**	0.0021***	0.0024***	0.0015
	(2.3445)	(2.7468)	(2.8379)	(1.4162)
amount outstanding (billion)	0.0008	0.0011*	0.0011*	0.0009
	(1.5062)	(1.9274)	(1.7066)	(1.1711)
dummy: cheapest-to-deliver	-0.0412***	-0.0424***	-0.0487***	-0.0354***
	(-5.2191)	(-5.1435)	(-4.8376)	(-2.8090)
intraday volatility		1.0316***	1.1130***	0.6573**
		(2.8962)	(2.6742)	(2.4596)
dummy: issuance day		0.0313**	0.0416***	0.0230
		(2.2265)	(2.6590)	(1.2325)
R^2_{Pseudo}	.1166	.1193	.1048	.1051
N	51,565	51,522	43,895	31,213

Table 6.3 shows the marginal effects of the probit estimation, with standard errors clustered at the dealer level.[22] Quoted bid-ask spread is significant in all specifications and on average an increase in bid-ask spread by 4 basis points (roughly one standard deviation of bid-ask spread in 10-year Bunds) makes it 2% less likely that a trade is taking place on MTS. Trade size is also negatively related to trading on MTS when we do not impose an upper cap on trade size: the marginal effect of ca. -0.02 implies that a trade of 10 million EUR is 1.4% less likely on MTS than an otherwise comparable trade half as large, i.e. of size 5 million EUR. When we cap trade size this effect is roughly half

[21] The dummy is set to 1 when a bond was newly auctioned or re-opened (tapped) on the same day. For new issuances we set the dummy to 1 also for the previous on-the-run bond in the same maturity category. That is in the case of a re-opening the dummy is 1 for the bond that is re-opened only and in the case of a new issuance for the newly issued bond and the one that is thus going off-the-run.

[22] Due to the lower number of dealers in the sample with trades only by reporting dealers significance estimates in the rightmost column are based on bootstrapped standard errors.

as strong and no longer significant. The positive coefficients for age suggest that older bonds are more likely to be traded on exchange, whereas the cheapest to deliver bond for the current futures contract is more likely to be transacted over-the-counter.[23] Next we consider immediacy proxies in our specifications. When intraday volatility is high the risk of adverse price movements is larger and a trader will want to fulfill her trading need faster. Beyond volatility, we also consider the required immediacy to be higher on auction days. That is since only on auction days the capability of the issuer to place their debt is revealed, dealers participating in the primary auction learn only during the day of their endowment obtained during the auction and might be faced by added trading needs from clients that wish to invest in the newly issued bond. Indeed we find the coefficient for intraday volatility to be significant across all specifications and consistently pointing towards a higher probability of trading on the exchange. Also on auction days we find that trading on the exchange is $2.3 - 4.2\%$ more likely and significantly so for all but the last specification.

We thus confirm Hypothesis 1, identifying transaction costs, order size, benchmark status and immediacy as determinants of venue choice. It is established that higher transactions costs (Smith et al. (2001); Bessembinder and Venkataraman (2004)) and larger order sizes (Smith et al. (2001); Bessembinder and Venkataraman (2004); Barclay et al. (2006); Hendershott and Madhavan (2015)) are drivers towards bilateral trading arrangements. Our finding that also the cheapest-to-deliver status makes OTC trading more likely confirms such a prediction in Lee and Wang (2017). The effect due to immediacy, to the best of our knowledge, has not been taken into account so far.

6.6 Trading costs

Having established the drivers of market choice in the previous section, we consider in this section the cost differences between trading over-the-counter and on the exchange. To do so we introduce OTC discount in Section 6.6.1 and study its drivers in Section 6.6.2.

6.6.1 OTC discount

Hypothesis 2. *Trades that take place OTC on average have a positive OTC discount, i.e. their price is favorable in comparison to the attainable price of executing the same trade on the exchange.*

To test Hypothesis 2 we examine cost differences between over-the-counter trades and the trading conditions provided through the limit order book. Therefore we compare the actual, observed, price of a trade with the hypothetical price the same trade would have incurred, had it been placed on the exchange. This is feasible because at any point in time we have the knowledge of the full limit order book of MTS. Hence, we are able to determine for any OTC trade how much a virtual trade on MTS would have cost. Formally we define OTC discount for any trade n as

$$OTC\ discount_n = \epsilon_n \left(price_n^{\text{virtual,MTS}} - price_n^{\text{observed, OTC}} \right) , \tag{6.2}$$

that is the price difference between the virtual price a trade would have incurred on MTS and the actually observed price of trade n, symmetrized for buyer- and seller-initiated trades by multiplying with the trade sign ϵ_n ($\epsilon = \pm 1$ for buyer-/seller-initiated trades). By

[23] Due to the construction of the futures contract the cheapest-to-deliver mostly coincides with the on-the-run bond during our sample period as long as the minimal requirement for the amount outstanding is met.

our definition a positive OTC discount implies that executing a trade over-the-counter was cheaper for the initiator than trading on the exchange. Since the discount of MTS trades is by definition equal to zero, we only consider over-the-counter trades in this section. Figure 6.3 shows the histogram of OTC discount for trades in 2-year Schaetze, 5-year Bundesobligationen and 10-year Bunds with a minimum trade size of two million EUR.[24] Our definition of OTC discount already takes into account the effect of walking up the book, i.e. when trades larger than the quantity available at the best execute against limit orders at deeper levels of the book and thus at an additional cost.

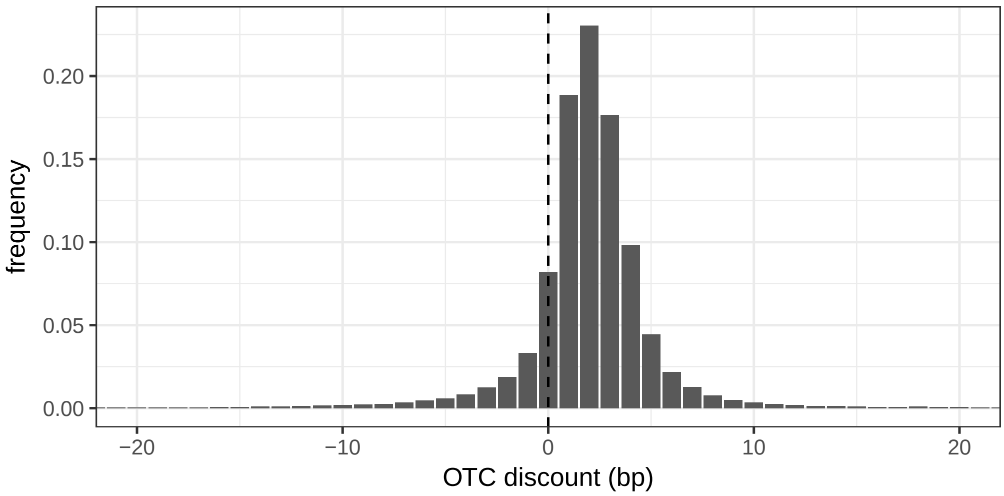

Figure 6.3: Histogram of OTC discount: By how much is trading OTC cheaper than on MTS? Discount received by OTC trades in 2-year Schaetze, 5-year Bundesobligationen and 10-year Bunds with a minimum trade size of two million EUR. OTC discount is the difference of the observed trade price to the hypothetical price of an identical trade on MTS, including the effect of walking up the book, symmetrized for buys and sells. A positive OTC discount implies that trading OTC was cheaper than on the exchange. Given in basis points.

We observe a distribution that is heavy on positive values of OTC discount, i.e. in the majority of cases trading over-the-counter is cheaper than on the exchange. This is natural since trading on the exchange provides the initiator with immediacy, a service that comes at a cost, and we will therefore also refer to the OTC discount as the *cost of immediacy*. Note that the other main difference of trading on the exchange is transparency: trading on the exchange publicly reveals a trading need to all other market participants post-trade, whereas for an over-the-counter trade only the contacted potential counterparties learn of the intention to trade. Since we are unable to disentangle this effect due to transparency or opacity preferences of a trader, we are actually under-estimating the true cost of immediacy. In other words, when one trades on the exchange, one pays not only with higher transaction costs but also with ones private information.

We also observe in Figure 6.3 that some trades have a negative OTC discount. This seems irrational at first, as one could have traded on the exchange at a lower cost. Several factors nevertheless explain such observations. First, the initiator of a trade might wish to protect her valuable private information, accepting a higher cost in the OTC market to avoid post-trade publication of a trade. E.g. Jank et al. (2016) provide evidence of a scenario where investors forego profits to protect their private information in the context of publication thresholds for short-selling. Second, a trader might also accept to trade at a

[24]We are excluding 30-year Bunds from the figure as their much wider bid-ask spread distorts the distribution. Results are nonetheless similar when including also 30-year Bunds.

Table 6.4: Descriptive statistics of OTC discount: OTC discount is the difference of the observed trade price to the hypothetical price of an identical trade on MTS, including the effect of walking up the book, symmetrized for buys and sells and given in units of basis points. A positive OTC discount implies that trading OTC was cheaper than on the exchange. The column p-value gives the p-value for a t-test of the mean being different from zero. Samples are all OTC trades and all OTC trades in 2-year Schaetze, 5-year Bundesobligationen and 10-year Bunds. The latter sample is also subsetted with a minimum size of 2 million EUR of nominal amount and an additional maximum trade size of of nominal 25 million EUR. For the subsample bounded in trade size from above and below we also distinguish between trades initiated by MTS dealers and non-MTS dealers.

OTC trades	mean	stddev	p05	p25	median	p75	p95	num_obs	p-value
all	3.58	8.03	-3.19	0.97	2.10	3.80	21.00	403,415	0.0000
2-year - 10-year bonds									
all	1.53	4.17	-3.00	0.80	1.90	3.00	5.40	338,929	0.0000
min. 2M EUR	2.11	4.68	-3.00	1.00	2.00	3.30	6.58	144,602	0.0000
trade size between 2 - 25M EUR									
all	1.46	3.51	-3.10	0.80	1.90	3.00	5.00	121,798	0.0000
MTS access	1.49	3.40	-3.00	0.90	1.90	3.00	4.90	37,895	0.0000
no MTS access	1.51	3.46	-3.00	0.90	2.00	3.00	5.00	48,822	0.0000

negative OTC discount in order to cultivate a valuable trading relationship. Di Maggio et al. (2017) show that trading relationships affect prices and liquidity and the value of the relationship increases during market turmoil. Lastly a trader might accept a negative discount when trading OTC offers additional services, such as e.g. contemporaneous hedging trades on the futures market.

Hypothesis 2 states that OTC discount is on average positive. In Table 6.4 we give the descriptive statistics of OTC discount. For all of our subsets both the mean and median are positive and in all cases the mean is statistically different from zero at all standard confidence levels, thus confirming the hypothesis. Furthermore, for the subset of trades sized between 2 and 25 million EUR, i.e. of comparable size to MTS trades, the mean (median) OTC discount is 1.5 (1.9) basis points, acting as a lower bound also to the cost of immediacy. We do not find a significant difference in the average OTC discount of trades by MTS dealers or non-MTS dealers, but this is not taking into account potential differences in trade characteristics between these two groups.[25] Our finding supports also the prediction from Grossman (1992) that OTC trading can tap into pools of unexpressed liquidity. This is confirmed even more by the fact that we also observe OTC trades that are too large for the MTS exchange.

6.6.2 Drivers of OTC discount

Hypothesis 3. *OTC discounts are larger a) when the bid-ask spread is larger, b) when the size of the trade is larger, c) when a bond is older and d) for traders that have access to the exchange market.*

In the following we study the drivers behind OTC discount and immediacy costs in order to test Hypothesis 3. Table 6.5 shows the results of estimating the following equation:

$$OTC\ discount_n = \alpha v_n + \Delta_i + \varepsilon_n \tag{6.3}$$

[25] Since for some trades we do not know the identity of the initiator and therefore also ignore her MTS status, these trades do not enter the last two rows of Table 6.4. Therefore the number of observations in the last two rows does not add up to the one in the third row from the bottom.

where the left hand variable is the OTC discount, defined in basis points, v_n is a vector of trade and bond-time characteristics, and Δ_i are bond-fixed effects. We consider several specifications for the estimation: with only a lower bound for trade size of 100,000 EUR, with a lower bound of 2 million EUR, with an additional upper bound of 25 million EUR and considering only trades by MTS dealers. Let us point out that notwithstanding the similarities one should not directly compare the results for venue choice from Section 6.5 to the drivers of OTC discount presented in this section. The estimation in the previous section was for the probability of trading on-exchange or over-the-counter conditional on trading (in either market). Instead the estimation in this section is for the discount of over-the-counter trades, i.e. conditional on a trader being able and willing to trade over-the-counter.

We find that bid-ask spread is a significant driver across all specifications. The magnitude implies that an increase in the quoted bid-ask spread by 1 basis point also leads to an increase in OTC discount by $0.39 - 0.45$ basis points. Given that the bid-ask spread corresponds to the cost of a round-trip, up to 90% of the increased cost of trading on the exchange is rebated in over-the-counter trades. Also trade size significantly influences OTC discount: OTC discount is bigger in larger over-the-counter trades. This is expected since trading costs on the exchange mechanically increase with size whereas liquidity providers in the OTC market can filter based on their perception of a counterparties trading motivation and thus allow for lower transaction costs also in larger trades. The magnitude of size effects depends on the size constraints we impose on our sample: the effect is smaller when imposing an upper limit on trade size, since for these trades the effect of walking up the limit order book is still relatively small. Also when we use a small lower limit for trade size of 100,000 EUR, i.e. including trades that are actually too small for MTS, the size effect is smaller.[26] Also older bonds receive a higher OTC discount, which may reflect increasing bid-ask spreads as bonds are ageing. Being cheapest-to-deliver for the futures is not a significant driver of OTC discount, that is the close pricing linkage to the futures does not confer a lower cost of OTC trading. The effects related to immediacy are more nuanced in this regression. Intraday volatility is only significant at the 10% level in the specifications with a lower limit of 100,000 EUR where it implies that OTC discount is lower on days with higher intraday volatility. The sign of the dummy for issuance days depends on whether we impose an upper bound on trade size or not. With an upper bound the impact of an auction on OTC discount is negative. Considering OTC trades of a size of $2 - 25$ million EUR we find that OTC discount for them is 1.5 basis points lower on days when there is an auction affecting the bond, and even 2.5 basis points lower for OTC trades by MTS dealers. This is in line with our conjecture relating to immediacy and our earlier findings, as in the wake of auctions there is either less OTC discount to be had, or traders are willing to accept a lower OTC discount over the exchange. However for OTC trades typically considered too large for MTS this effect is reversed, implying that there are some very large trades that do not incur such additional costs for immediacy but rather an additional rebate. Finally we also observe that MTS dealers get on average a $0.7 - 0.9$ basis points higher OTC discount, i.e. they are better in their search for counterparties to OTC trades. As we do not control for bank characteristics we are unable to make any claim as to whether this advantage is due to their privilege of having access to the MTS market or rather is a consequence of the skills of their trading desk.

This confirms Hypothesis 3. While drivers relating to transactions costs are naturally

[26]This is mostly a technical effect since small trades do not incur additional costs from walking up the book that are reflected in OTC discount. As almost half of the trades in this specification are smaller than 2 million EUR, this has a substantial influence on the size coefficient.

Table 6.5: Drivers of OTC discount: Ordinary least squares regression of OTC discount (in basis points) on trade- and bond-time characteristics. The sample consists of OTC trades in German sovereign debt titles involving German financial institutions. Regressions include bond-fixed effects and standard errors are clustered at daily time and dealer level. T-values are given in brackets and *, ** and *** denote significance at the 10%, 5% and 1% level respectively.

	trade size $\geq$ 100k EUR		trade size $\geq$ 2M EUR									
						trade size $\leq$ 25M EUR					$\leq$ 25M EUR	
									MTS access			
bid-ask spread (bp)	0.450***	0.450***	0.393***	0.393***	0.393***	0.407***	0.407***	0.407***	0.410***	0.410***	0.395***	0.394***
	(14.995)	(14.995)	(9.074)	(9.074)	(9.075)	(7.662)	(7.661)	(7.663)	(17.308)	(17.230)	(17.015)	(16.905)
trade size (log)	0.581***	0.562***	1.795***	1.756***	1.786***	0.402*	0.423*	0.415*	1.997***	1.946***	0.737***	0.774***
	(4.176)	(4.133)	(8.299)	(8.258)	(8.365)	(1.863)	(1.955)	(1.962)	(9.630)	(9.441)	(3.785)	(4.062)
bond age (years)	0.480***	0.501***	0.351***	0.378***	0.358***	0.319**	0.300*	0.323**	0.324**	0.357**	0.367***	0.336***
	(5.552)	(5.678)	(2.628)	(2.737)	(2.744)	(2.164)	(1.972)	(2.319)	(2.313)	(2.587)	(3.005)	(2.747)
dummy: cheapest-to-deliver	-0.076	-0.075	-0.229	-0.230	-0.190	-0.238	-0.245	-0.204	0.116	0.111	0.642	0.634
	(-0.194)	(-0.192)	(-0.431)	(-0.434)	(-0.351)	(-0.462)	(-0.474)	(-0.397)	(0.162)	(0.155)	(1.186)	(1.171)
intraday volatility	-0.100*	-0.100*	-0.118	-0.118	-0.118	-0.011	-0.011	-0.010	-0.005	-0.004	0.000	0.000
	(-1.758)	(-1.755)	(-1.594)	(-1.591)	(-1.589)	(-0.561)	(-0.567)	(-0.551)	(-0.403)	(-0.392)	(0.024)	(0.017)
dummy: issuance day		1.814**		1.485*			-1.482*			1.555*		-2.491*
		(2.498)		(1.759)			(-1.889)			(1.712)		(-1.896)
dummy: MTS access					0.853**			0.696**				
					(2.361)			(2.170)				
R^2	.7999	.8	.5558	.5558	.5559	.6121	.6122	.6122	.03934	.03943	.04632	.04654
R^2_{adjusted}	.7998	.7999	.5553	.5554	.5554	.6117	.6117	.6118	.03719	.03726	.0438	.044
R^2_{within}	.796	.796	.5456	.5457	.5457	.6024	.6024	.6025	.01044	.01054	.008393	.008621
N	209,540	209,540	112,419	112,419	112,419	96,835	96,835	96,835	49,254	49,254	41,652	41,652

explained by the potential which there is for discounts, other effects are more nuanced. E.g. immediacy seems to play a different role for small, medium-sized and large trades. This is in line with theories as Seppi (1990) that identify large trades as liquidity-motivated, wherefore they receive better quotes, i.e. a higher OTC discount.

6.7 Conclusion

In an environment where academics and regulators increasingly call for a shift from traditional over-the-counter market structures towards electronic platforms and greater transparency, an in-depth understanding of the drivers and motivations behind venue choice is ever more important. The empirical literature has so far focused mostly on either exchange-dominated markets (such as equities) or hybrid markets that involve an over-the-counter market and multi-dealer request-for-quote platforms. A main innovation of this chapter is to study a hybrid bond market with a predominant OTC segment and a liquid limit order book venue. The richness of our data allows us to study at the level of single trades venue choice and transactions costs.

In particular we consider also immediacy and transparency as drivers of venue choice. Our findings suggest that not one single factor is able to capture the dynamics of venue choice on its own. Instead we see significant contributions related to transaction costs, order size, benchmark status of a bond and immediacy as well as evidence that also transparency plays a crucial role.

This study is highly relevant in the context of the recently introduced MiFID II regulation that affects also the Bund market, through e.g. post-trade transparency also for OTC transactions and best-execution rules that encourage trading on regulated platforms. Our results imply that both market structures provide advantages in different circumstances: E.g. the transaction cost analysis reveals that when the exchange is less liquid, the over-the-counter market is still able to provide liquidity at least to some traders. On the other hand the exchange is clearly beneficial in providing immediacy when it is needed. How MiFID II will impact European fixed income markets is an open question to be answered in the years to come, providing numerous opportunities for further research.

Appendix

6.A Cross-venue response

The hybrid setting of the Bund market poses an interesting laboratory also from a price impact point of view. That is, what is the price impact of over-the-counter transactions? Since the OTC segment in our setting is opaque we would expect price impact to be weaker and slower, as less information is revealed and only indirectly observed by most market participants. In this section we explore the potential of our dataset to answer this question. We stop short of applying a transient impact model (TIM) as in Bouchaud et al. (2004) and instead consider price response, i.e. the price movement in a bond conditional on a trade.[27]

Some studies on upstairs markets have also considered price impact of upstairs trades, e.g. Keim and Madhavan (1996); Booth et al. (2002); Carollo et al. (2012). As before we argue that our setting is sufficiently different in having an OTC segment that is independent

[27]See Section 2.3 and also Chapter 4 for a discussion of how price response is related to price impact.

from the exchange. Eisler and Bouchaud (2016) study price impact in the OTC credit index market and show that a TIM is well-suited also for this market structure. Here instead we consider the influence of trades on the quoted mid-price on the exchange MTS. That is we define response as

$$R^i_{\Delta t} = \mathbb{E}\left[\left(X^i_{t+\Delta t-\varepsilon} - X^i_{t-\varepsilon}\right)\epsilon^i_t\right]. \tag{6.4}$$

where X^i_t is the log MTS mid-price of bond i at time t and ϵ^i_t is the order sign of the trade (+1 for buys, −1 for sells) which occured at time t. Thus we term the response to OTC trades *cross-venue* response.[28]

6.A.a Order sign

A crucial input to equation (6.4) is the market order sign ϵ, i.e. whether a trade was a buyer- or seller-initiated. Unfortunately for trades in the *WpHMV* data (transactions data by German banks) the order sign is not given. Therefore we infer the order sign by comparing the price of the trade to the MTS mid price at the beginning of the minute of the trade.[29] A transaction at a price above (below) the mid price is classified as a buy (sell) and trades at the midprice do not enter the calculation of response, following the approach in Eisler and Bouchaud (2016). While trades at the midprice could also be classified following Lee and Ready (1991), we exclude them from this part of our analysis.

To test this classification approach we first apply it to the sample of all MTS trades, where the true order sign is known. We correctly classify 91.5% of trades, better than the 72% reported by Eisler and Bouchaud (2016) for a sample of proprietary trades in credit index instruments and comparing to midquotes from Bloomberg. For robustness we also consider a time lag of plus and minus one minute, i.e. comparing to the mid price from at least one minute before (again at the full minute) and the full minute following the trade respectively. This correctly classifies 85.0% and 69.8% of trades. Classification precision is also roughly consistent across maturities at issuance, with short-term bonds faring slightly better than long-term bonds with typically higher bid-ask spreads. Table 6.6 reports the classification precision for MTS trades in the rightmost column "OS precision".

Table 6.6: Statistics of bid-ask spread by maturity and order sign classification precision: Descriptive statistics of the average daily MTS bid-ask spread of the on-the-run German federal bonds in each maturity bin. Given in basis points and based on values winsorized at the 1% level from above. Order sign classification precision (*OS precision*) is the share of trades where the order sign is correctly inferred from comparison with the MTS mid price in the same minute, given in %.

	Mean	Std Dev	5 Pcl	25 Pcl	50 Pcl	75 Pcl	95 Pcl	OS precision
2-year Schatz	3.67	1.86	1.68	2.56	3.14	4.12	7.51	92.7 %
5-year Bobl	6.80	3.75	3.65	4.74	5.87	7.75	12.02	93.0 %
10-year Bund	7.22	3.94	3.91	4.98	6.14	8.15	13.67	91.9 %
30-year Bund	54.90	37.65	23.22	37.55	46.78	60.20	99.30	88.9 %

For OTC trades from the WpHMV sample we can not compare to the true order sign, as it is not reported and therefore unknown. Instead we check how stable the classification

[28] Cross-venue response is to be distinguished from cross-asset response or impact as discussed in Chapter 4.

[29] The WpHMV data is reported in seconds precision, however the effective resolution is often of lower frequency. This is on the one hand since at some point it becomes difficult to assign a precise timestamp to an OTC trade facilitated via phone and may on the other hand also be due to imprecisions in reporting. By effectively reducing resolution to minutes we take this into account.

of the order sign is with respect to shifts of the reference mid price by one minute back and forth as above. Comparing to the midprice from the previous (following) minute 88.7% (87.8%) of order signs remain unchanged. Finally, Carollo et al. (2012) infer the order sign of trades in the LSE upstairs market as the sign of the client. Our dataset does allow for this analysis when we restrict it to dealer-to-client trades. However we then find only 50.4% of order signs classified in agreement with the mid price comparison above.

6.A.b Response

In Figure 6.4 we show the response function $R^i_{\Delta t}$ averaged over all trades in all bonds for four subsamples: over-the-counter trades from the WpHMV data, MTS trades from the WpHMV data, and MTS trades from the MTS data, where for the latter we once use the true order sign in computing the response and once we use the inferred order sign. We also restrict ourselves to trades of a nominal size between 1 and 50 million EUR as to avoid biases from a large number of small transactions or very large trades.

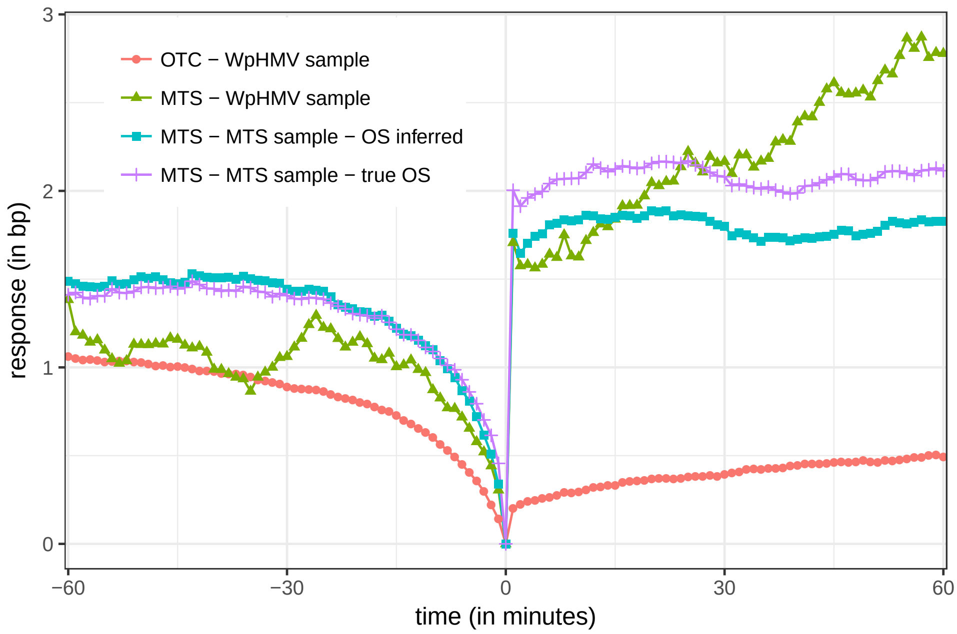

Figure 6.4: Response function: Average log returns of MTS midquotes in response to trades of size 1-50 million EUR nominal value. *OTC- WpHMV sample* indicates response to over-the-counter trades recorded in the regulatory WpHMV dataset and *MTS - WpHMV sample* refers to MTS (exchange) trades included therein. *MTS - MTS sample* refers to all trades on the exchange MTS (of which the MTS trades in the WpHMV sample are a subgroup) and we calculate response once with the inferred order sign (c.f. Section 6.A.a) and once with the true order sign as recorded in the MTS dataset.

Let us first distinguish the response to MTS trades. The differences between the response functions here are due to either sample differences (the WpHMV sample contains only a subset of all MTS trades) or the order sign used for calculating response (for the MTS sample we compute response both using the inferred and the true order sign). Considering only the rather small set of MTS trades recorded in the WpHMV database the response function is rather noisy and we only consider it for illustration purposes here. Looking to the larger MTS sample, naturally the response to MTS trades is higher when inferred with the correct order sign. The ratio corresponds roughly to the classification precision of 91.5%, i.e. immediate response with the true order sign is ca. 2 basis points versus ca. 1.8 basis points with the inferred order sign.

Immediate response is higher for MTS trades by a factor ~ 10, i.e. response to exchange trades is roughly 2 basis points and to OTC trades ca. 0.2 basis points. The response to exchange trades is approximately immediate and appears roughly flat, whereas response to OTC trades keeps increasing after the initial minute for at least one hour, rising from ca. 0.2 basis points to ca. 0.5 basis points on average. Both observations are in line with our initial expectation that information from opaque OTC trades is reflected in quoted prices more slowly and less strongly.

We will in the following distinguish between response in bonds of different maturities in order to expand on this discussion.

Table 6.7: Number of trades by maturity: Number of trades in the over-the-counter (OTC) and exchange (MTS) segment by bond maturity. Includes trades for a nominal amount of 1 - 50 million EUR.

	OTC - WpHMV data	MTS - MTS data	number of bonds
2-year Schatz	15,711	4,782	31
5-year Bobl	36,772	5,403	26
10-year Bund	102,764	16,941	36
30-year Bund	16,697	8,258	14
sum	171,944	35,384	107

Table 6.8: Number of trades per day by maturity: Average number of trades per bond-day if there was at least one trade on said day in said bond. *OTC - WpHMV data* are over-the-counter trades and *MTS - MTS data* are all trades on the interdealer exchange MTS. Includes trades for a nominal amount of 1 - 50 million EUR.

	OTC - WpHMV data	MTS - MTS data	number of bonds
2-year Schatz	2.71	2.87	31
5-year Bobl	3.17	2.27	26
10-year Bund	4.36	2.48	36
30-year Bund	2.20	2.96	14

Tables 6.7 and 6.8 give an overview of the number of OTC and MTS trades per maturity bin and the average number of trades per bond and day (conditional on a trade) respectively. Since we consider the conditional average of number of trades per bond-day in Table 6.8 the average number of OTC and MTS trades is roughly similar. However Table 6.7 reveals that the number of OTC trades is two to six times that of MTS trades, depending on the maturity at issuance of a bond. Most active in the OTC market are 10-year Bunds with a conditional average of 4.4 trades per bond and day and more than $100,000$ trades in total.

Figure 6.5 shows the response function for each maturity bin separately. Note that the strength of response increases with maturity, i.e. response is smallest for short term 2-year Schaetze and largest for very long-term 30-year Bunds. As already noted above, response to MTS trades is mostly immediate and flat, ranging from roughly 0.5 basis points in 2-year bonds to about 5 basis points in 30-year Bunds. Response to OTC trades is in all cases smaller than response to exchange trades by roughly one order of magnitude. Note that this is at least in part due to trades with incorrectly inferred order signs that on average weaken the strength of the estimated response. However we deem our classification to be sufficiently precise to posit that the response to OTC transactions is indeed weaker. Whereas response to exchange trades is immediate we observe that prices continue to increase after OTC trades for at least one hour after the trade to roughly double the initial

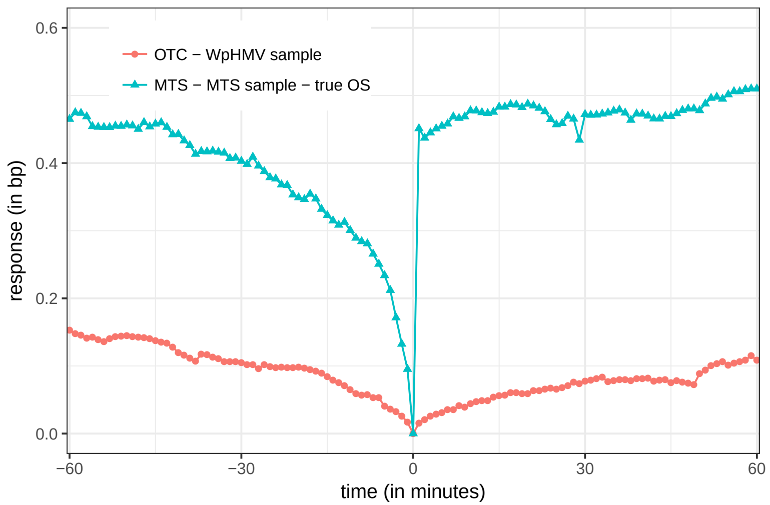

(a) 2-year Schaetze

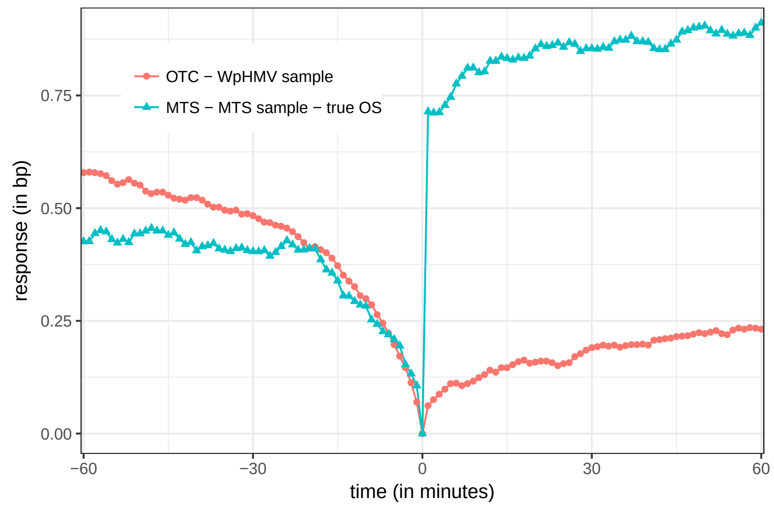

(b) 5-year Bobls

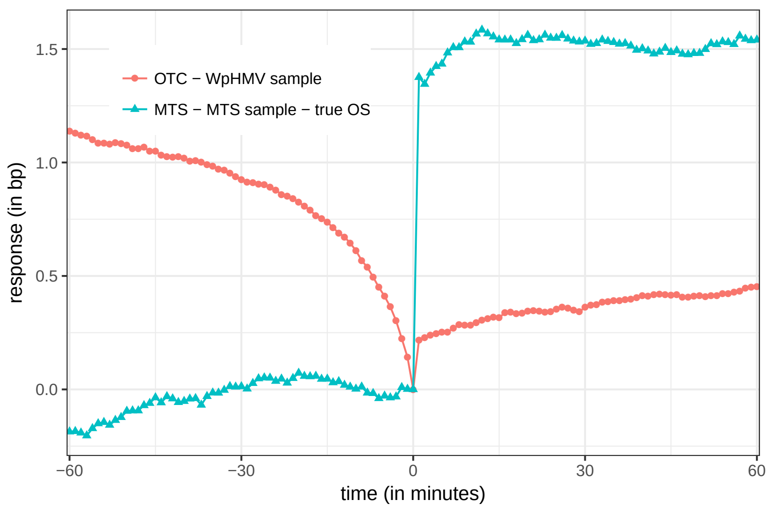

(c) 10-year Bunds

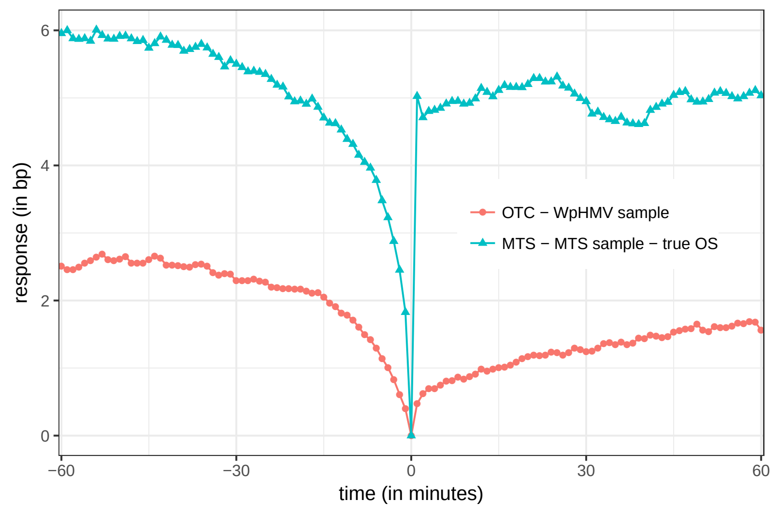

(d) 30-year Bunds

Figure 6.5: Response function: Average log returns of MTS midquotes in response to trades of size 1-50 million EUR nominal value, distinguished by maturity at issuance. *OTC- WpHMV sample* indicates response to over-the-counter trades recorded in the regulatory WpHMV dataset where the ordersign is inferred and *MTS-MTS sample* refers to all trades on the exchange MTS and response is calculated with the true order sign as recorded in the MTS dataset.

response. This could be either due to autocorrelated order signs or, as pointed out above, when information from OTC trades is only slowly incorporated into prices on the exchange, a point that we will discuss in more detail below.

Let us next consider the response at negative lags. In exchange markets the lag at negative responses is typically negative, a feature that is related to the autocorrelation of order flow (cf. Section 4.4.2 and Taranto et al. (2016) for a discussion) and indeed predicted by transient impact models. Here we find that response at negative lags is mostly positive, with the exception of MTS trades in 10-year Bunds. This is evidence that financial agents condition their trading activity on prevailing market conditions, a feature that is not reflected in TIMs. That is a trader is more likely to buy when a bond is relatively cheaper, i.e. on falling prices. This is indeed what we observe in Figure 6.5. Note that this is also observed on exchange markets when observed response at negative lags is larger than what is predicted from the TIM as e.g. in Figure 4.2 in this thesis or Figure 1 in Taranto et al. (2016).

6.A.c Autocorrelation of the order sign

As indicated from Table 6.8 most bonds see only few trades per day and especially so on MTS (cf. also Darbha and Dufour (2013)). This would suggest that herding and order splitting, two main causes of autocorrelation of order signs, are less prevalent and therefore also autocorrelation would be weaker. Furthermore Seppi (1990) describes anecdotal evidence for upstairs markets that traders are discouraged from order splitting and encouraged to reveal their full liquidity needs to the dealer, suggesting that the autocorrelation in the signs of OTC trades should be even lower. At the same time, if there were a significant and strong autocorrelation of OTC order signs, this would provide a competing explanation for the slowly increasing response that we observe in Figure 4.2 for positive lags.

In Figure 6.6 we show the response function for three actively traded 10-year Bunds, examplary for our larger sample. We compute the autocorrelation separately for OTC and MTS trades and therefore the lag on the x-axis corresponds to different physical timescales unless trading activity in both markets were equal.[30] We do find a positive autocorrelation of order signs for the duration of a few trades for both OTC and MTS trades. As expected, this effect is considerably weaker for OTC trades. This supports the explanation that the slow increase of response to OTC trades is due to information from OTC trades being incorporated into public markets only slowly and undermines the alternative explanation that this might be due to order splitting.[31]

We leave the expansion of this study to price-impact for future work. Both applying the TIM of Eisler and Bouchaud (2016) or a modified TIM that identifies each venue are conceivable. The latter approach would then be methodologically similar to the cross-impact TIM used in Chapter 4 of this thesis. The evidence we have collected here suggest that OTC trades indeed have a lower price impact compared to market orders on the exchange.

[30] The OTC market is more active in our case.

[31] Also that OTC response at negative lags is found to be positive is in line with this line of reasoning.

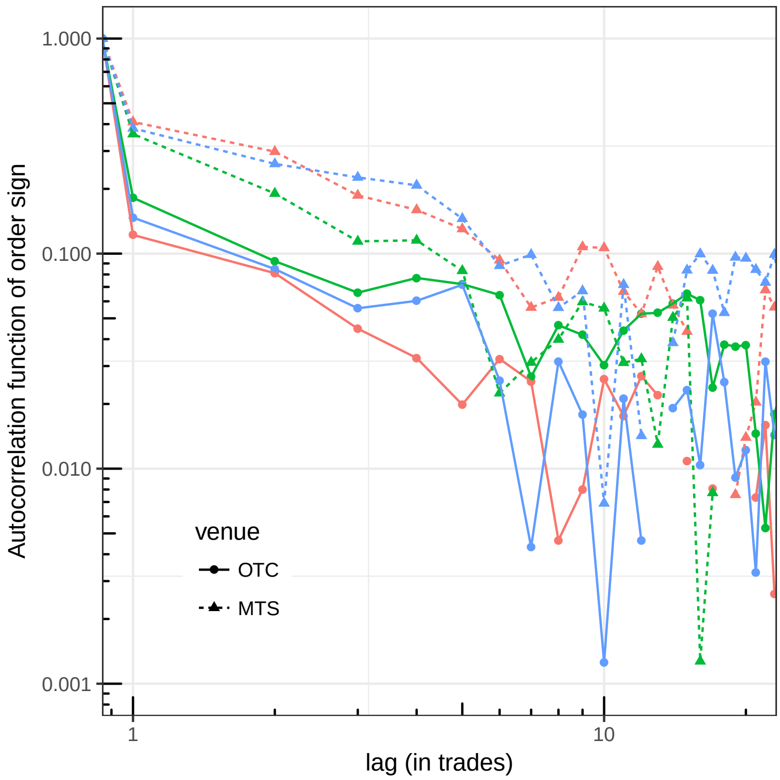

Figure 6.6: ACF of order sign: Autocorrelation function of the order sign (±1 for buy/sell trades) for three actively traded 10-year Bunds. Different bonds are distinguished by different colors and dots connected by solid lines refer to OTC trades while triangles connected with dashed lines refer to MTS trades. Not that the time-scales differ for OTC and MTS trades: lag is given in number of OTC (MTS) trades for the autocorrelation of OTC (MTS) order signs and these correspond to differences in physical time due to trading activity differences. Estimated on the base of trades with a nominal amount of 1-50 million EUR, using the inferred order sign for OTC trades and the true sign for MTS trades.

Chapter 7

Liquidity of German corporate bonds

7.1 Introduction

One of the most interesting developments in the transparency of global financial markets occurred in 2002 with the initial launch of the Trade Reporting and Compliance Engine (TRACE) platform by the Financial Industry Regulatory Authority (FINRA) for the mandatory reporting of transactions in the over-the-counter (OTC) U.S. corporate bond market.[1] However, more than 15 years after the dissemination of TRACE, the effect of transparency on liquidity and investor welfare in OTC markets is still debated. Supporters of OTC market transparency argue that it reduces the asymmetry of information between dealers and investors. Furthermore, transparency encourages the participation of retail/uninformed investors, who can benefit from better price discovery, and obtain a fairer price for their transactions, similar to informed/institutional traders.[2] On the other hand, OTC market transparency could increase transaction costs for some investors, by eliminating dealers' information rents and, thus, incentives to compete or even participate in the market.[3]

In this chapter, we contribute to the debate by providing a unique study of a market without trade transparency, and comparing it to one with full post-trade information dissemination. Specifically, we analyze the liquidity of the German corporate bond market, where there is no mandatory post-trade transparency, and compare it to the U.S. market, where FINRA enforces a strict disclosure protocol. For our analysis, we use a unique regulatory dataset, with a complete set of bond transactions of German financial institutions from 2008 until 2014.[4] To the best of our knowledge, we are the first researchers to use this database to study market liquidity. Therefore, we provide a detailed description of our data cleaning procedures, which determine our sample selection. We focus on straight, unsecured corporate bonds, excluding all bonds with complex optionalities attached. We estimate transaction costs at a weekly frequency by adopting a wide range of liquidity measures.

[1] The TRACE platform was extended to other U.S. fixed income markets, including the structured product market in May 2011 (see Friewald et al. (2017)), and most recently, the Treasury bond market in July 2017.

[2] For theoretical work on OTC market transparency see Pagano and Roell (1996), Duffie et al. (2017), Asriyan et al. (2017). Empirical analysis supporting these arguments can be found in Bessembinder et al. (2006), Edwards et al. (2007).

[3] See Naik et al. (1999), Bloomfield and O'Hara (1999), Holmstrom (2015), Bhattacharya (2016).

[4] Reporting is mandated through the German Securities Trading Act, (*Wertpapierhandelsgesetz*, shortened to "WpHG") and collected by the German federal financial supervisory authority *Bundesanstalt für Finanzdienstleistungsaufsicht*, in short "BaFin".

Starting with a sample of 11,670 corporate bonds, we focus on a relatively liquid sample in which a particular bond trades at least 8 times in a week. Our final sample includes 1,703 corporate bonds for the German market (1,585 issued by financial institutions and 118 issued by non-financial firms). Our study consists of four parts. First, we provide a general description of the German corporate bond market, focusing on the key characteristics of the bonds in our sample and their trading activity. Second, we analyze the time-series evolution of liquidity in the German market. Third, we study the determinants of liquidity in the cross-section of the German market with panel regressions on bond characteristics in the spirit of Edwards et al. (2007). Fourth, we use a matched-sample approach to compare the transaction costs of similar bonds, at the same point in time, in the German and the U.S. market, respectively. The variables that we use for our matching procedure are coupon, rating, time to maturity, size, volume traded and trading frequency.

In our descriptive analysis, we find significant differences between the German and U.S. markets. The former is composed, in great part, of financial bonds (i.e., bonds issued by financial firms), which are ten times as many as non-financial bonds (i.e., bonds issued by non-financial firms). In the U.S. market, financials are also the majority, but just four times as many as non-financials. Bond characteristics present differences, as well, in our sample: German bonds have a higher coupon, a lower time to maturity than their U.S. counterparts, and also, most of the non-financial bonds are unrated. Overall, observed trading activity is much lower in the German market: the bonds that traded at least once 8 times per week are only 17% of the sample, against 74% of the traded sample in the U.S. universe. Looking at the market as a whole, liquidity is clearly much higher in the U.S., with a significantly larger number of securities that trade often, and therefore, likely to provide more informative prices. This result is consistent with various theoretical studies that show that transparency lowers costs for un-sophisticated investors and, therefore, incentivizes participation in the market.[5]

The time-series dynamics of liquidity look similar between the two markets and across different liquidity measures. As expected, we find that transaction costs for German corporate bonds spiked during the 2008-2009 global financial crisis and the sovereign debt crisis in 2011-2012. While the former clearly affected liquidity in the U.S. market as well, as documented by Friewald et al. (2012) and others, the latter did not and, thus, was a shock mostly limited to the Euro-area, and perhaps the rest of Europe. The cross-sectional analysis shows that the relations predicted by search theories of OTC markets are also confirmed in the German market: A bond is more liquid if it has a larger issue size, a better credit rating, a shorter time-to-maturity, a younger age, and a larger volume traded.

The matched sample analysis allows us to compare *frequently* traded German bonds with a group of U.S. bonds that have similar characteristics, at the same point in time. Contrary to our expectations, across all the liquidity measures except imputed round-trip cost, we find that this group of German bonds has significantly lower transaction costs than comparable bonds in the U.S. market. The difference in round-trip costs is within a range of 37-67 basis points, depending on the liquidity measure used. This finding might seem surprising, at first blush. However, it is in line with studies that highlight the potential unintended consequences of an increase in transparency in OTC markets. For example, Naik et al. (1999) show that, in a more transparent market, dealers fail to extract information rents from trading with investors and have less incentive to compete, which could lead to *higher* costs of trading for investors. Bloomfield and O'Hara (1999) provide similar findings in a laboratory experiment. In a recent study, Bhattacharya (2016) shows how post-trade transparency can *increase* transaction costs due to trade delays of investors,

[5] For theoretical studies, see for example Pagano and Roell (1996), Duffie et al. (2017).

who wait longer in order to acquire more information by monitoring disseminated trade prices. In a similar vein, Friewald et al. (2017) document in the U.S. securitized product market that there is an optimal level of detail in disclosure, beyond which there is no improvement in liquidity. A possible explanation for our finding is that, when there is little transparency, investors concentrate their demand into a few well-traded assets, resulting in "crowding." As a consequence, the liquidity of these few bonds is particularly high, while the others are barely traded, resulting in a greater dispersion of liquidity across bonds. On the other hand, when overall transparency increases, investors spread their portfolios across a wider range of assets, given the higher level of information available. While the overall market liquidity improves, there is less relative demand for the previously "well-known assets," and hence their transaction costs, at least in some cases, could be higher.

Overall, our results support the notion that the effects of transparency in OTC markets are multifaceted, and not unambiguously positive. Our analysis shows that transparent markets have greater trading activity and stronger participation, overall. The proportion of securities that is traded frequently is much higher, suggesting better price discovery overall. This indicates that transparent markets are, as a whole, more liquid. However, when restricting the analysis to securities that are most frequently traded in the non-transparent market, the most similar bonds in the transparent markets *could* have higher transaction costs. A possible explanation for this seemingly anomalous result is that, in non-transparent markets, investors concentrate their demand in a few securities, making them more liquid, while barely trading the rest of the assets. An alternative, symmetric explanation could be that market makers find it difficult to provide liquidity beyond a small number of instruments, when transparency is low, overall. In either case, lower transparency leads to "crowding" of demand into a few securities that may be even more liquid, as a consequence.

This chapter makes three main contributions. First, it contributes to the debate on the effect of transparency in OTC markets, by providing a novel analysis of a market without post-trade transparency. Many empirical studies have analyzed the impact of transparency in the U.S. market, using different phases of the TRACE program as identification. However, such analyses cannot overcome the limitation that the assets still belong to the same overall market, allowing for no cross-market comparisons. In contrast, our access to a novel database allows us to analyze, in detail, the German market without transparency, and compare it with the U.S. market with mandatory disclosure. Second, it presents a comprehensive analysis of the liquidity of a rather unexplored market, which is growing and is among the largest European bond markets. Third, we provide a detailed description of a filtering procedure for a new database, which has potential for future research on other illiquid bond markets or for answering policy questions regarding the German bond market.

Our results are of interest for academics and regulators alike. In particular, they speak to the recent debate on the introduction of MiFID II and MiFIR, and provide a perspective on how to critically evaluate the anticipated improvement in transparency in European fixed-income markets.[6]

The chapter is organized as follows. Section 7.2 provides a literature review. Section 7.3 describes the corporate bond market structure in Europe. Section 7.4 introduces our dataset and describes our approach to obtaining samples, for which we provide descriptive statistics. In Section 7.5, we use various liquidity measures from the literature in order to

[6] In an attempt to increase transparency, the European Parliament and the European Council approved in 2014 the Notification 2014/65/EU - Markets in Financial Instruments Directive II (MiFID II) and the Regulation (EU) No 600/2014 – Markets in Financial Instrument Regulation (MiFIR), to enhance pre-and post-trade transparency of both equity and non-equity instruments and derivatives including fixed income bonds, which are applicable to all European markets since January 3, 2018.

examine the time-series evolution of liquidity, study the determinants of liquidity in the cross-section with panel regressions and to compare the transaction costs of similar bonds, at the same point in time, in the German and the U.S. market with a matched-sample approach. Section 7.6 concludes.

7.2 Literature review

7.2.1 European corporate bond market

A number of papers deal with the pricing of European corporate bonds in relation to credit risk and other risk factors and illiquidity risk at an aggregate level. Others provide a description of the vast cross-section of yields. However, none of them provide an analysis of liquidity at an issue level, since they do not employ a dataset nearly as complete and detailed as ours. A few studies use transaction data, e.g., Díaz and Navarro (2002) use data of trades in 1993-1997 on three Spanish bond platforms and Frühwirth et al. (2010) use closing prices from transactions on German exchanges. Other articles (Houweling et al. (2005); Van Landschoot (2008); Castagnetti and Rossi (2013); Klein and Stellner (2014); Utz et al. (2016)) rely on yield quotes, e.g., from Bloomberg, at daily or lower frequencies or consider corporate bond indices (Aussenegg et al. (2015)).

Two papers shed light on the market microstructure of specific trading platforms for corporate bonds: Fermanian et al. (2016) models the request-for-quote (RFQ) process on the multi-dealer-to-client platform Bloomberg FIT, based on a fraction of the RFQs received by BNP Paribas in the years 2014 and 2015. However, their focus is on the behavior of clients and dealers rather than the market as a whole, with the data being used to calibrate their theoretical model. In fact, no statistics on trading volumes or liquidity measures are provided, since that is not their focus. Linciano et al. (2014) study the liquidity of Italian corporate bonds that are listed on two platforms contemporaneously (DomesticMOT or ExtraMOT and EuroTLX) and find a mixed impact of such fragmentation. Also, their analysis is restricted to these platforms and neglects the major market share of OTC trades, which form the majority of trades in corporate bond markets.

A different approach is taken in Bundesbank (2017), which mainly considers the market size of bonds of non-financial corporations in the Eurozone in terms of the total amount outstanding. The report analyzes the bond market in the context of the low-interest-rate environment of the past few years based on supply and demand factors, looking at issued amounts, yields and yield spreads of corporate bond indices.[7] In the context of measuring the impact of ECB bond purchases under the Corporate Securities Purchase Program (CSPP), Grosse-Rueschkamp et al. (2017) show that there was an increase in bond financing in companies whose debt was eligible for purchase by the ECB. Again, there are no statistics provided on trade volumes and liquidity measures.

Biais et al. (2006) investigates liquidity based on a dataset of interdealer trades from 2003-2005 in a set of Euro- and sterling denominated bonds listed in the iBoxx index by looking at quoted and effective bid-ask spreads. Furthermore, the paper considers informational efficiency in the European bond markets and compares it to the early literature on the U.S. TRACE database. However, it refers to the market before the global financial crisis and the European sovereign debt crisis and the consequent changes in market regulation are not reflected therein.

Our study, in contrast, is based on a broader dataset, both in terms of transaction

[7] An earlier study in this direction is Pagano and Von Thadden (2004), in the context of the monetary unification of the Eurozone.

detail and the underlying bond universe, and provides a *direct* comparison with the U.S. TRACE database. Finally, in our evaluation of liquidity, we do not only rely on quoted spreads and prices but employ a range of liquidity measures based on actual transactions that have proven more suitable for the analysis of OTC markets. We are aware of only two other studies that make use of regulatory trade-level data to study liquidity in European bond markets: AMF (2015) studies the French bond market and is aimed primarily at constructing a composite liquidity indicator. Aquilina and Suntheim (2017) provide a similar analysis for the U.K. corporate bond market, quantifying as well the yield spread due to liquidity.[8] It should be emphasized that non of these studies takes into account bond characteristics as drivers of liquidity and their evolution in time, nor is there a comparison against the benchmark U.S. market.

7.2.2 U.S. corporate bond market

Since the inception of TRACE in July 2002, there has been a growing number of empirical studies that analyze the U.S. corporate bond market. Among the first of these after the introduction of post-trade transparency are Bessembinder et al. (2006) and Edwards et al. (2007).[9] Both papers focus on the effect that post-trade transparency has on corporate bond transaction costs, finding that bid-ask spreads significantly reduced after the introduction of TRACE. Edwards et al. (2007) further provide an analysis of trading costs on the cross-section of bonds, showing that those better rated, recently issued and close to maturity are more liquid. In a more recent paper, Bao et al. (2011) focus on the link between illiquidity and pricing in the U.S. corporate bond market, showing that a significant part of the variation of yield spreads can be explained by movements in corporate bond prices. Along the same lines, Lin et al. (2011) find a robust link between illiquidity and corporate bond returns. Liquidity can be an issue especially during periods of financial distress, either at the market or at the security level. The two most prominent papers analyzing U.S. corporate bond market liquidity during the financial crisis are Friewald et al. (2012) and Dick-Nielsen et al. (2012): both show that trading costs spiked during the recent financial crisis, thus having a significant impact on yield spreads, especially those of bonds with high credit risk. Jankowitsch et al. (2014) focus instead on the effects of financial distress on liquidity at the security level, analyzing recovery rates of defaulted bonds. A more recent group of papers focuses on the impact of the Volcker rule on the U.S. corporate bond market liquidity. As pointed out by Duffie (2012b), the restriction imposed by regulators on dealers' trading activity can significantly impact the level of the bid-ask spreads in the market. Bessembinder et al. (2016) study corporate bond liquidity and dealer behavior in the period 2006-2016, finding that trade execution costs have not increased significantly over time, while dealer capital commitment was significantly reduced after the crisis. Bao et al. (2016) analyze the illiquidity of stressed bonds, focusing on rating downgrades as stress events. They find that, after the introduction of the Volcker Rule, stressed bonds are significantly more illiquid, due to Volcker-affected dealers lowering their market liquidity provision. Finally, in a recent working paper, Choi and Huh (2017) demonstrate that customers, such as hedge funds, often provide liquidity in the post-crisis U.S. corporate bond market. Therefore, average bid-ask spreads, which rely on the assumption that dealers provide liquidity, underestimate trading costs that liquidity-demanding investors pay. Finally, Schestag et al. (2016) provide a comprehensive analysis of liquidity measures

[8] Another regulatory report that is concerned with the corporate bond market at the European level is ESMA (2016); however, it relies on data from Markit and Euroclear, both based on market averages, instead of transactions-level data.

[9] Harris and Piwowar (2006) have a similar study for the municipal bond market.

in OTC bond markets using a sample period that covers 2003-2014.

7.3 Market structure and financial market regulation

From the market microstructure point of view, the European corporate bond market is mostly a classic over-the-counter (OTC) market.[10] This OTC market is further differentiated by size and involves an inter-dealer segment (i.e., D2D) and a retail segment (dealer-customer, i.e., D2C), and into voice and electronic markets by trading mechanism.

The voice market is organized around dealers (large banks and securities houses) and their network of clients. Transactions are largely bilateral via the telephone. As stressed in Duffie (2012a), the process of matching buyers and sellers requires a large amount of intermediation in this market, as well as attendant search costs. Since this market is known as a quote-driven market, i.e., executable prices are offered in response to a counterparty's request to trade, prices for the same bond, at the same time, could vary significantly across dealers; hence, traders often contact more than one dealer in search of the best execution price. This fact, and that quotes and transaction prices are usually not publicly known, make bond trading more opaque than many other traded asset classes.

Besides voice, there exist several electronic platforms with different trading protocols at their core. Single-dealer platforms are often a mere electronic version of the voice mechanisms described above, while multi-dealer platforms allow the customer to request quotes to trade from a number of dealers simultaneously and facilitate automated record keeping. Another recent innovation is "all-to-all" platforms that are estimated to account for almost 5% of electronic trading by now. For a survey of the ongoing developments that are affecting the market structure and functioning of the fixed income markets due to electronic trading, see BIS (2016). A study by Greenwich Associates (Greenwich Associates (2014)) indicates that around 50% of trading volume is conducted electronically in the European investment grade corporate bond market, and almost 20% for high-yield bonds.

7.4 Data

7.4.1 Description of the dataset

Our dataset is based on the transaction reporting obligations of German banks mandated by the German Securities Trading Act (*Wertpapierhandelsgesetz*, "WpHG"). Section 9 of the act, further detailed in the respective regulation (*Wertpapierhandel-Meldeverordnung*, "WpHMV"), requires credit or financial services institutions, branches of foreign institutions and central counterparties (only Eurex Clearing AG, in practice) domiciled in Germany to report to the German Federal Financial Supervisory Authority (*Bundesanstalt für Finanzdienstleistungsaufsicht*, popularly known as "BaFin"). The requirement is to report "any transactions in financial instruments which are admitted to trading on an organised market or are included in the regulated market (*regulierter Markt*) or the regulated unofficial market (*Freiverkehr*) of a German stock exchange." The dataset also captures a large set

[10]Even in cases where exchanges are organized around a central limit order book, their market share is minor. For example, in the Italian bond market, where exchange trading is relatively more common, less than 30% of the turnover in Italian non-government bonds takes place on exchanges, according to a report by the Italian securities regulator Commissione Nazionale per le Società e la Borsa (CONSOB) (CONSOB, Bollettino Statistico Nr. 8, March 2016, available at `http://www.consob.it/web/area-pubblica/bollettino-statistico`), with the rest occurring in the OTC market. This number is likely to be much lower for other European countries, including Germany.

of transactions of non-German institutions at German exchanges.[11]

To the best of our knowledge this dataset has only been used in a set of studies in the context of institutional herding in the German equities market (Kremer and Nautz (2013a,b); Boortz et al. (2014)). Since these prior studies offer neither a comprehensive description of the dataset nor a focus on corporate bonds, we initially provide a detailed description of the dataset and the series of filtering steps we apply to the raw data. The transactions dataset contains security information, detailed information on the transaction (for instance, time, price, size, exchange code or indicator for OTC trades) and the parties involved (an identifier for the reporting institution and, where applicable, identifiers of client, counterparty, broker or intermediaries).[12] We augment this information with security characteristics from the Centralized Securities Database (CSDB), which is operated jointly by the members of the European System of Central Banks (ESCB), together with other security information from Thomson-Reuters, Datastream and Bloomberg. For a smaller subset of bonds we also obtained time-series of daily price quotes from Bloomberg.

Our raw dataset contains all reporting in "any interest-bearing or discounted security that normally obliges the issuer to pay the bondholder a contracted sum of money and to repay the principal amount of the debt" as indicated by a CFI-code starting with "DB" (with "D" for debt instruments, and "B" for bonds).[13] Our bond dataset covers the full set of transactions over the period from January 2008 to December 2014; therefore, it initially includes any type of sovereign, guaranteed, secured, unsecured, negative pledge, junior/subordinated and senior bonds reported through WpHG. For this sample selection we adopt a narrower definition of the corporate bond market than Bundesbank (2017) and the capital market statistics of Deutsche Bundesbank, which include other debt-type securities not classified as bonds. Of the total market size of 145 billion EUR amount outstanding of German corporate bonds at the end of 2014 reported in the capital market statistics, we capture about 41 billion EUR, i.e. roughly one-third. In addition, we note that our initial sample includes non-German bonds (traded by German financial institutions) as well.

7.4.2 Data filtering and sample selection

Our dataset is subjected to a careful filtering process, in order to ensure the soundness and reliability of the final sample. We describe below our general procedure, also mentioning considerations for uses of the data other than ours. We then proceed to describe the sample selection filters that are specific to our study.

Panel A of Table 7.1 provides an overview of the observations discarded throughout the cleaning process. In a first cleaning step, we remove entries with invalid ISINs or time-stamps. Moreover, we employ an error code assigned by BaFin to each observation, which takes the integer values from 0 (no errors) to 3 (serious errors - junk), to drop observations with error code 3 in this step. On average, this step filters out only 0.2% of observations and we observe that the data quality improves after 2009. Recall that the initial filtering of our dataset for bond-type securities relied on the CFI-code provided by

[11] Building societies (Bausparkassen) are excluded from the reporting requirement. Moreover, non-German EU banks do not have to report trades in MiFID-securities since they already report these in their home countries. A non-binding English translation of the law is provided at `https://www.bafin.de/SharedDocs/Veroeffentlichungen/EN/Aufsichtsrecht/Gesetz/WpHG_en.html`

[12] For a full list of variables see the Annex to WpHMV. A non-binding English translation is provided at `https://www.bafin.de/SharedDocs/Downloads/EN/Formular/WA/dl_wphmv_anlage_en.html`

[13] This definition excludes any convertible bonds ("DC"), bonds with warrants attached ("DW"), medium-term notes ("DT"), money market instruments ("DY"), asset-backed securities ("DA"), mortgage-backed securities ("DG"), or other miscellaenous debt instruments ("DM").

Table 7.1: Data Cleaning and Sample Selection: Panel A shows the number of observations after each of the cleaning steps described in section 7.4.2. *Before cleaning* is the initial number of observations. We discard observations with *errors* in ISIN or time-stamp, in *non-debt* securities, in minor *currencies*, corresponding to *technical lines* (duplicate lines that are automatically created by some reporting systems when the trade is on hold) or corresponding to *double-reporting* by both parties of a trade. Filtering is applied for *prices* (absolute values and weekly price median filter) and *complete CFI-codes* (for bond classification). Panel B describes number of bonds, observations and traded volume retained in our sample selection process. *Vanilla bonds* are bonds with a complete CFI code with fixed or zero coupon and a fixed redemption date and which are not classified otherwise as hybrid or structured products. Vanilla bonds are distinguished into *secured/guaranteed* bonds (secured through assets or a non-government entity), *treasury-type* bonds (issued or guaranteed by a federal or state government) and *unsecured bonds*. For unsecured bonds we distinguish between *certificates* and *corporate bonds* which are either *financial bonds* or *non-financial bonds* depending on the issuer type.

Panel A: Data Cleaning

year	before cleaning	after removing errors	non-debt	currencies	technical lines	double-reporting	after price-filtering	complete CFI code	vanilla bonds
2008	3,802,701	3,766,085	3,461,772	3,428,490	2,416,096	1,597,236	1,588,700	1,522,319	1,261,461
2009	3,675,045	3,670,634	3,263,093	3,242,910	1,706,814	1,197,506	1,189,079	1,114,034	845,341
2010	4,106,918	4,106,149	3,989,344	3,965,796	2,019,035	1,486,864	1,478,227	1,343,649	935,190
2011	3,770,269	3,769,170	3,678,184	3,658,141	2,072,979	1,429,922	1,421,010	1,297,610	928,948
2012	4,681,385	4,681,196	4,573,035	4,543,243	2,695,704	1,765,453	1,723,423	1,584,035	961,130
2013	4,258,139	4,258,136	4,122,257	4,077,110	2,670,584	1,711,202	1,687,676	1,497,546	965,263
2014	4,077,082	4,077,070	3,818,136	3,747,882	2,426,399	1,574,867	1,554,282	1,351,927	843,248
$\sum$	28,371,539	28,328,440	26,905,821	26,663,572	16,007,611	10,763,050	10,642,397	9,711,120	6,740,581

Panel B: Sample Selection

		vanilla bonds			unsecured bonds		corporate bonds	
	vanilla bonds	secured	treasury-type	unsecured	certificates	corporates	non-financial	financial
WPHG: all trades								
# bonds	81,664	4,873	4,218	72,573	60,817	11,670	817	10,853
# trades	6,740,581	1,188,767	2,906,710	2,645,104	765,781	1,857,777	738,839	1,118,938
traded volume (million EUR)	12,474,668	791,973	11,244,802	437,893	29,901	389,294	57,209	332,085
WPHG: trades in German bonds								
# bonds	65,819	2,803	1,272	61,744	51,938	9,741	178	9,563
# trades	4,216,200	661,076	1,287,824	2,267,300	671,381	1,578,591	581,594	996,997
traded volume (million EUR)	7,572,336	488,579	6,732,304	351,453	20,863	319,454	34,879	284,575
TRACE								
# bonds	9,602	598	274	8,730		8,730	2,314	6,414
# trades	16,810,039	319,396	414,914	16,075,729		16,075,729	3,537,977	12,537,466
traded volume (million USD)	6,865,837	195,225	274,333	6,396,279		6,396,279	1,605,194	4,790,566

the reporting institutions. In the second step, to ensure robustness of our data, we also remove all observations from ISINs where the CFI-code recorded by CSDB does not start with "DB", thus double-checking our sample selection. This removes another 5.0% of our initial observations. Prices are reported in the currency used in the trade and need to be converted to EUR. In the third step we do so, by keeping only trades originally reported in the main currencies: "EUR", "AUD", "CHF", "GBP", "USD", "CAD", "JPY", "DKK", "NOK" and "SEK".

In the fourth step, we remove so-called *technical lines*. These lines are created in some reporting systems e.g., when a trade is on hold while a broker is gathering more of a security she has committed to sell. Technical lines are detected when the reporting entity field is identical to the client field. Discarding them removes 37.6% of the initial number of observations. Even after accounting for technical lines, the same transaction can still be recorded in multiple lines. This happens, for instance, when both counterparties are obliged to report when a central counterparty is involved or when an intermediary is used. Since our focus is on trading activity, in the fifth step, we keep only one observation for each transaction, i.e., we identify duplicates as trades on the same day, in the same security, at the same price, and for the same absolute volume. While the parties involved in the trade are also reported, their reporting style can be inconsistent. We thus ignore the information on the involved parties to avoid false negative duplicate detections, but instead we use it for fine-tuning our filtering parameters. Another crucial variable is the intraday time of the trade. Unfortunately, it is possible that for the same trade, different intraday timestamps are reported, e.g., when one counterparty of an OTC transaction needed additional time to conclude their side of the trade. As a compromise between false positive and false negative duplicate detections, we consider two lines to be duplicates only when their intraday time difference is a maximum of ten minutes. By discarding duplicates, we drop another 18.5% of the initial data-set or 32.8% of the remaining observations.

Finally, we apply price filters. We first remove trades reported at prices of less than 1% or more than 500% of nominal bond value, and then apply a weekly price median filter, filtering out trades that deviate by more than 10% from the weekly median price. For computational reasons, we do not apply a price reversal filter since we find for a smaller subset of actively traded bonds that only a negligibly small number of trades would be flagged. This filtering step drops only 0.4% of observations. Two additional fields in the dataset indicate whether a deal is on behalf of a client or not and whether the deal affects the balance sheet of the reporting institution. These fields are useful when one is interested in the inventory or balance sheet of the reporting institutions.

The cleaning steps described above leave us with about 10.6 million observations, corresponding to single trades, down from an initial 28.4 million observations. Next, we proceed to select our sample of corporate bonds, relying on the CFI code. Thus, in the second column from the right of Panel A in Table 7.1, we consider only bonds where the 1st, 2nd and 3rd attributes (type of interest, guarantee and redemption) in the CFI code (consolidated from WpHG and CSDB data) are well-defined (i.e., non-"X"). This corresponds to dropping another 8.8% of trades. From this sample, we select bonds with either a fixed or zero coupon rate and a fixed redemption date, and that are not classified otherwise as hybrid or structured products; we name these *vanilla bonds*. Our initial sample of vanilla bonds, therefore, consists of 6.7 million trades in $81,664$ bonds for a traded volume of ca. 12.5 trillion EUR.

Panel B of Table 7.1 distinguishes vanilla bonds by bond securization type as inferred from the second attribute of the CFI code. The column *secured/guaranteed* refers to vanilla bonds either secured through assets or guaranteed by a non-government entity (attribute

"S" or "G" respectively), *Treasury-type* bonds are issued or guaranteed by a federal or state government (attribute "T"), e.g., German Bunds and KfW-issued bonds are also part of this category. *Unsecured* bonds do not carry a guarantee or security (attribute "U"). The largest share of trading volume is due to government bonds with 11.2 trillion EUR, while secured or guaranteed bonds make up for 792 billion EUR. Unsecured bonds account for a total trading volume of 438 billion EUR. Based on the CSDB variable "debt_type" we classify unsecured vanilla bonds into *corporate bonds* and *certificates*. While there is a large number of 60,817 certificates, they make up for only 30 billion EUR of traded volume. The set of corporate bonds is further distinguished in the two rightmost columns into *financial bonds* (i.e., bonds issued by financial corporations such as banks, insurance corporations and financial auxiliaries) and *non-financial bonds* (bonds issued by industrial and other non-financial companies). Even though there are only 817 non-financial bonds in our sample, they make up for a traded volume of 57 billion EUR compared to 332 billion EUR traded in 10,853 financial bonds.

We believe that our sample is highly representative of the German (corporate) bond market, but only to a much lesser extent of the whole European market. Therefore, we focus our attention on German corporate bonds and report the statistics of number of bonds, number of trades and traded volume for German-issued bonds in the middle section of Panel B. Throughout this chapter, therefore, we will compare the German corporate bond market, based on the BaFin dataset, to the market for U.S. corporate bonds, based on TRACE data. Therefore, we also provide the corresponding statistics for U.S. bonds at the bottom section of Panel B. It is remarkable that even though the U.S. corporate bond market is much larger in size, we are starting from an even slightly larger number of bond issues in the German market.[14] It is important to stress that our transaction data on German bonds captures only a share of trading activity (that by German financial institutions, essentially), whereas TRACE data can be considered as covering the full U.S. corporate bond market. Any comparison we make should thus be seen with this caveat in mind. We cannot make any final statement on the *absolute levels* of trading volume and trading activity, which are lower than the total market. However, we believe this sample to be representative of the whole market, and hence informative on *relative* levels of liquidity, trading activity and trading volume.[15]

7.5 Liquidity in markets with and without transparency

In the previous section, we have described, in detail, the filtering and selection process to obtain our dataset for analysis. In this section, we present our main analysis, which can be divided into four parts. First, we provide a general description of the German corporate bond market, focusing on the key characteristics of the bonds and their trading activity. Second, we analyze the time-series evolution of liquidity in the German market. Third, we study the determinants of liquidity in the cross-section of the German market with panel regressions on bond characteristics in the spirit of Edwards et al. (2007). Fourth, we use a matched-sample approach to compare the transaction costs of similar bonds, at the same point in time, in the German and the U.S. market. The variables that we use for our

[14] Most of these are financial bonds, a finding that is confirmed by industry reports and data from the Centralised Securities Database (CSDB) of the Eurosystem of Eurpean central banks. Note that our TRACE sample does not include certificates to make the comparison more reasonable.

[15] This type of limited sample has been used even when larger datasets are available as in the US. For example, Di Maggio et al. (2017) study trading relationships in the U.S. corporate bond market by using a random sample of TRACE data which covers approximately 10% of the market.

matching procedure are coupon, rating, time to maturity, size, volume traded and trading frequency.

7.5.1 Descriptive analysis

Full sample

In Table 7.2 we provide summary statistics of bond characteristics and bond-level trading activity, distinguishing both between financial and non-financial corporate bonds, and comparing our German BaFin sample with the U.S. TRACE counterpart.

The coupon rates of German corporate bonds are slightly larger than for U.S. bonds (6.3% compared to 5.9% for non-financials and 5.4% compared to 5.2% for financial bonds), which is somewhat surprising, considering that overall German interest rates were lower than in the U.S. in our sample period. German corporate bonds are typically shorter-lived with an average maturity of 5.9 years for non-financial bonds and 4.0 years for corporate bonds, whereas the corresponding numbers for U.S. bonds are 13.7 and 8.5 years respectively, also featuring greater variation, as indicated by wider quantiles and a higher standard deviation, in relative terms. Most notably U.S. bonds are also, on average, much larger, measured by issue size. The mean amount issued of German non-financial bonds (financial bonds) is 146 million EUR (48 million EUR), compared to 406 million USD (229 million USD) for U.S. bonds. Also, the 95% percentile of the issue amount at 1.25 billion USD is much larger than the corresponding 700 million EUR (139 million EUR) for German bonds. Notice that the EUR/USD exchange rate during the 2008-2014 period fluctuated between 1.19 and 1.60, with a mean of 1.36 USD per EUR.

Turning to indicators of trading activity, there is a clear disparity between the two bond markets. The average number of trades per day in our sample of German non-financials has a mean of 5.22, and ranges from 0.008 to 18.5, from the 5% to the 95% percentile. For U.S. bonds, the mean is 1.23, and the 5% and 95% percentiles range from 0.001 to 2.16. It is important to consider that, in our German sample, the number of non-financial bonds is only 178 versus $2,374$ in the U.S. sample. However, the picture is quite the opposite for financial bonds. In this case, there are $9,563$ German financial bonds, and the mean of the average number of trades per day is 0.18, and the quantiles range from 0.002 to 0.78. Instead for the U.S. sample, where the number of financial bonds is far lower, $6,414$, the mean is 0.70, more than three times larger, and ranges from 0.002 to 3.44, that is the value at the 95th percentile is four times larger. We observe the opposite picture for the average duration between two trades in the same bond, where for German non-financial bonds, the average is 12.75 days, compared to 69.48 days in the U.S., whereas for financial bonds, the average duration between trades is 64.11 days in Germany versus 49.62 days for the U.S., with a large dispersion around these measures. These summary statistics provide evidence of the large heterogeneity between the two markets, and between financial/non-financial issuers. Yet, there is no clear evidence that these two subsamples of German bonds are any less liquid than the corresponding two subsamples of U.S. corporate bonds.

A clearer picture is given by Table 7.3, which reports the distribution of bonds in bins of the number of days per year for which we observe trading activity in the bond. First, the number of bonds in the two markets is rather similar: 9,741 in Germany and 8,728 in the U.S. However, when looking at the distribution of the trading frequency, there is a clear difference between the two markets: 32% of the U.S. sample is traded at least 100 days a year, while only 6% of the German bonds are traded that often. On the other hand, 48% of the U.S. sample is rarely traded (0-50 trading days a year). The group of rarely traded bonds in the German sample amounts to 87%, representing the great

Table 7.2: Descriptive statistics of bond characteristics for German and U.S. corporate bonds, differentiated for financial and non-financial sector bonds. *Coupon rate* is in percent, excluding zero coupon-bonds. *Maturity* is the maturity at issuance, in years. *Time-to-maturity)* is the average remaining maturity at the time of a trade (averaged for each bond), in years. *Amount issued* is the issued amount of bond in million EUR or USD respectively. *Average number of trades per day* is the average number of trades on any given trading day for the lifetime of the security. *Average number of trades per day if trade* is the average number of trades on a trading day with at least one trade. *Average trading interval* is the average amount of calendar days between two consecutive trades. *Average daily volume* is the average volume in million EUR (USD) traded on one day. *Average daily volume if trade* is the average volume in million EUR (USD) traded on a trading day with at least one trade. *Trade volume total* is the summed value of trades recorded in our sample for a given security, in million EUR or USD respectively.

	German bonds, WPHG					U.S. bonds, TRACE				
	mean	stddev	p05	median	p95	mean	stddev	p05	median	p95
non-financial bonds										
Coupon rate (%)	6.32	1.80	3.21	6.61	9.00	5.94	2.48	1.50	6.05	9.88
Maturity (years)	5.89	2.32	3.02	5.00	10.00	13.71	12.44	2.48	10.02	30.08
Time-to-maturity (years)	4.11	2.21	1.14	4.00	7.88	1.53	4.65	0.00	0.02	7.88
Amount issued (millions)	146.63	236.47	2.17	30.00	700.00	406.49	633.17	12.92	250.00	1,250.00
Avg. # trades per day	5.22	7.11	0.01	2.45	18.48	0.42	1.23	0.00	0.04	2.16
Avg. # trades per day | trade	6.55	6.67	1.25	4.20	18.18	11.67	25.75	1.33	3.02	62.00
Avg. trading interval (days)	12.75	34.23	1.44	1.89	56.14	22.06	69.49	1.46	5.00	91.34
Avg. daily volume (millions)	0.33	0.68	0.00	0.09	1.42	0.22	0.63	0.00	0.02	1.06
Avg. daily volume | trade (millions)	0.49	0.82	0.01	0.20	1.71	28.81	399.01	0.06	2.22	122.75
Total trade volume (millions)	195.95	313.37	0.11	54.66	925.15	693.69	2,054.87	0.73	86.78	3,077.17
financial bonds										
Coupon rate (%)	5.36	4.30	1.45	3.87	15.00	5.16	1.70	2.00	5.30	7.55
Maturity (years)	4.02	2.87	1.07	3.25	10.00	8.53	7.99	1.05	5.04	25.06
Time-to-maturity (years)	2.45	2.38	0.22	1.41	7.20	0.98	2.90	0.00	0.01	5.35
Amount issued (millions)	47.89	203.17	0.10	17.01	139.00	229.28	511.28	10.53	24.73	1,250.00
Avg. # trades per day	0.18	0.70	0.00	0.02	0.78	0.70	2.01	0.00	0.14	3.44
Avg. # trades per day | trade	1.60	1.28	1.00	1.20	3.39	4.97	8.37	1.00	2.96	16.15
Avg. trading interval (days)	64.11	125.24	2.82	25.07	245.32	18.75	49.62	1.48	5.41	72.63
Avg. daily volume (millions)	0.06	0.30	0.00	0.00	0.21	0.32	1.41	0.00	0.01	1.72
Avg. daily volume | trade (millions)	3.63	25.48	0.00	0.07	14.60	4.11	25.81	0.04	0.16	15.27
Total trade volume (millions)	29.76	135.91	0.00	1.11	119.58	746.89	2,911.20	0.28	20.63	4,007.97

majority of the market. This major difference is a clear indication of much greater market liquidity in the U.S., with a significantly larger number of securities that trade often, and therefore, likely to provide more informative prices. This result is consistent with various theoretical studies that show that transparency lowers costs for unsophisticated investors and, therefore, incentivizes participation in the market.[16]

Table 7.3: Frequently traded bonds: Number of bonds for which trading activity is reported on a number of days within bins of 50. For each bond we count the most active year.

	German bonds		TRACE	
	non-financial	financial	non-financial	financial
200+	78	102	288	828
151-200	15	147	161	524
101-150	19	248	197	865
51-100	13	579	286	1,356
0-50	53	8,487	1,382	2,841
$\sum$	178	9,563	2,314	6,414

Liquid sample

In an attempt to provide a precise estimate of transaction costs, we concentrate our following analysis on a set of bonds for which liquidity measures can be estimated. In line with the best practices established in the literature, we consider only bond-week observations that have at least 8 transactions.[17] This leaves us with $1,703$ German bonds (118 non-financial and $1,585$ financial bonds) and $6,493$ U.S. bonds ($1,744$ non-financial and $4,749$ financial bonds). Table 7.4 provides the descriptive statistics of bond characteristics and trading activity for this liquid sample.

The sub-sample partially confirms the characteristics that we observe in the larger sample. As for the whole sample, the amount issued is smaller for German corporate bonds. For non-financial bonds, the average number of trades per day is larger for German bonds compared to the U.S., whereas for financial bonds, the distributions are now quite similar. For all these measures, the dispersion is quite large, indicating that both samples feature considerable heterogeneity across bonds. The coupon rates in this case are largely similar and, instead, the maturity is shorter for German bonds (both non-financials and financials).

In Table 7.5, we also report the credit ratings for our sample of liquid bonds. Credit ratings for the German BaFin sample were obtained from Bloomberg, Thomson Reuters and Bundesbank databases, whereas ratings for TRACE are from Mergent FISD. Most bonds are rated investment grade across samples. This is especially true for German financial bonds, due to the relatively better credit rating quality of state banks. The most striking feature is the lack of credit ratings for German non-financial bonds, where we are able to obtain a credit rating only for 28 of the 118 bonds. This low proportion of rated non-financial bonds is explained by the credit market structure in Europe, which not only

[16] For theoretical studies, see for example Pagano and Roell (1996), Duffie et al. (2017).

[17] Estimations based on TRACE data are usually performed at a daily frequency. However given the less complete coverage of trading in German bonds, this would impose too strict a selection criterion on the German sample. However, while we use a weekly frequency for our analysis, robustness checks at a daily frequency show that our results are in line with our findings below.

Table 7.4: Descriptive statistics of bond characteristics for liquid German and U.S. corporate bonds, differentiated for financial and non-financial sector bonds. *Coupon rate* is in percent, excluding zero coupon-bonds. *Maturity* is the maturity at issuance, in years. *Time-to-maturity)* is the average remaining maturity at the time of a trade (averaged for each bond), in years. *Amount issued* is the issued amount of bond in million EUR or USD respectively. *Average number of trades per day* is the average number of trades on any given trading day for the lifetime of the security. *Average number of trades per day if trade* is the average number of trades on a trading day with at least one trade. *Average trading interval* is the average amount of calendar days between two consecutive trades. *Average daily volume* is the average volume in million EUR (USD) traded on one day. *Average daily volume if trade* is the average volume in million EUR (USD) traded on a trading day with at least one trade. *Trade volume total* is the summed value of trades recorded in our sample for a given security, in million EUR or USD respectively.

	German bonds					TRACE				
	mean	stddev	p05	median	p95	mean	stddev	p05	median	p95
non-financial bonds										
Coupon rate (%)	6.00	1.77	2.93	6.50	8.50	6.12	2.27	2.13	6.13	9.88
Maturity (years)	5.81	2.00	4.00	5.00	10.00	14.35	12.61	4.15	10.02	30.10
Time-to-maturity (years)	4.06	1.71	1.31	4.02	6.61	4.76	6.19	0.00	1.85	18.07
Amount issued (millions)	183.42	264.41	5.00	50.00	757.50	477.75	682.65	14.99	300.00	1300.00
Avg. # trades per day	6.92	7.45	0.11	4.34	20.32	0.51	1.42	0.00	0.04	2.94
Avg. # trades per day | trade	7.90	7.04	1.60	5.41	20.28	30.66	38.39	9.00	16.17	99.82
Avg. trading interval (calendar days)	4.52	8.99	1.44	1.73	19.60	76.74	148.82	7.00	27.22	309.63
Avg. daily volume (millions)	0.46	0.80	0.01	0.16	1.66	0.26	0.70	0.00	0.03	1.33
Avg. daily volume | trade (millions)	0.55	0.83	0.02	0.30	1.71	49.57	474.12	0.22	11.31	166.96
Total trade volume (millions)	269.58	353.44	2.91	108.07	948.85	824.26	2365.86	0.68	123.70	3679.08
financial bonds										
Coupon rate (%)	3.75	2.18	1.61	3.38	6.15	5.15	1.62	2.03	5.35	7.50
Maturity (years)	4.57	2.75	1.19	4.09	10.00	11.03	8.13	2.04	10.00	30.02
Time-to-maturity (years)	2.85	2.18	0.41	2.43	6.87	3.85	3.54	0.20	2.71	9.47
Amount issued (millions)	72.58	151.35	2.36	27.00	250.00	311.92	591.74	10.71	31.31	1500.00
Avg. # trades per day	0.73	1.33	0.04	0.31	2.69	0.89	2.40	0.01	0.13	5.14
Avg. # trades per day | trade	2.90	2.17	1.36	2.15	7.15	25.16	32.66	9.25	14.35	82.01
Avg. trading interval (calendar days)	15.36	25.80	2.09	8.14	49.93	44.78	70.41	7.00	25.22	141.08
Avg. daily volume (millions)	0.12	0.36	0.00	0.02	0.46	0.43	1.37	0.00	0.01	2.58
Avg. daily volume | trade (millions)	1.17	4.86	0.03	0.14	5.00	12.45	44.77	0.16	0.58	60.63
Total trade volume (millions)	60.39	158.84	0.79	14.05	267.57	1050.43	3462.93	0.64	22.50	6109.19

heavily relies on bank loan financing, but also involves corporate bonds that are often held until maturity by long-term investors such as insurance companies. As a consequence, many corporate bond issuers do not seek a credit rating.

Table 7.5: Ratings of liquid bonds: Ratings for German corporate bonds are obtained via Bloomberg, Thomson Reuters or Bundesbank databases. Ratings for U.S. bonds are from Mergent ID. The sample is corporate bonds for which we calculate liquidity measures. Bonds with a rating step of 10 and better are investment grade.

	German bonds		U.S. bonds	
rating step	non-financial	financial	non-financial	financial
1 (Aaa/AAA)	0	64	58	368
2 (Aa1/AA+)	0	38	59	348
3 (Aa2/AA)	0	117	17	365
4 (Aa3/AA-)	0	347	82	690
5 (A1/A+)	4	460	110	539
6 (A2/A)	0	134	250	768
7 (A3/A-)	0	260	161	252
8 (Baa1/BBB+)	7	34	145	106
9 (Baa2/BBB)	7	1	164	89
10 (Baa3/BBB-)	7	17	115	99
11 (Ba1/BB+)	1	0	41	262
12 (Ba2/BB)	1	2	59	67
13 (Ba3/BB-)	0	0	70	68
14 (B1/B+)	1	2	55	173
15 (B2/B)	0	0	55	48
16 (B3/B-)	0	0	90	32
17 (Caa1/CCC+)	0	0	46	9
18 (Caa2/CCC)	0	0	16	6
19 (Caa3/CCC-)	0	0	3	6
20 (Ca/CC)	0	0	0	1
21 (C/C)	0	0	2	0
unavailable	90	109	146	453
$\sum$	118	1,585	1,744	4,749

7.5.2 Measuring liquidity in bond markets

Liquidity metrics

To measure liquidity, we employ a range of liquidity metrics that have been tested and verified on U.S. TRACE data. The *Amihud measure*, proposed in Amihud (2002), is a proxy for market price impact, i.e., the average price shift induced by a trade. The other measures we use capture the cost of a round-trip trade: *Price dispersion*, the *Roll measure*, the *imputed round-trip cost* and the *effective bid-ask spread* all estimate the loss associated with buying and immediately selling an asset (which would be the bid-ask spread in the case of an exchange market). All details regarding the calculation of these liquidity measures are provided in Appendix 7.B. It is important to highlight two aspects of our methodology. First, while these measures are typically calculated on a daily basis for U.S. TRACE data, all measures calculated here for both the German and the U.S. samples are based on weekly data. This allows us to include more German bonds in our analysis that are relatively actively traded, but not typically on a daily basis, while maintaining the comparability between the two samples. Second, the calculation of the effective bid-ask spread requires the trade sign, i.e., whether a trade was buyer- or seller-initiated. This information is provided in TRACE but not in our BaFin data on the German market. Instead there we infer the trade sign using the algorithm of Lee and Ready (1991) by comparing our trade prices with quotes from Bloomberg. Since such quotes are not available for all bonds in our sample, this effectively creates a different subsample and we need to be careful when comparing the effective bid-ask spread with other liquidity measures of German bonds. We present visually this in Figure 7.1, where we show the number of bonds for which the liquidity metrics of price dispersion and the effective bid-ask spread could be computed in each week of our sample.[18] The number of bonds for which we compute the effective bid-ask spread is always smaller than that for the effective bid-ask spread due to the reasons mentioned above.

Table 7.6 provides summary statistics of the different liquidity measures. Panel A reports the liquidity measures for both German and U.S. non-financial bonds. The table indicates that the different liquidity measures do have a large cross-sectional variation, both for the German and U.S. bonds. The price dispersion and Roll measures, and the effective bid-ask spread based on our BaFin sample, are, on average, lower for German non-financial bonds than the corresponding bonds in TRACE, whereas the Amihud illiquidity measure and the round trip cost are higher. On the other hand in Panel B, the financial bonds in TRACE are shown to be, on average, always more illiquid than German financial bonds in the BaFin sample. However, we hasten to emphasize that a simple average comparison is misleading. Such a comparison should not induce us to conclude that German corporate bonds are generally more liquid than the U.S. corporate bonds. In fact, financial bonds in TRACE appear to be more illiquid because of the large differences in the number of bonds considered in both samples, i.e., the number of bond-week observations of the BaFin German sample are only as much as one-eighth of the TRACE U.S. sample in Panel A, whereas in Panel B, this ratio is even less than 10 percent. The time series dynamics of these measures are potentially different and will be taken into account in the analysis below. In the following subsections, we perform a finer and more granular analysis of the different liquidity measures from a time series and cross-sectional perspective, and try to investigate potential similarities and differences in the patterns and cross-sectional characteristics of the different liquidity measures between U.S. and German corporate bonds.

7.5.3 Time-series dynamics of liquidity

In describing the evolution of liquidity, we first consider the effective bid-ask spread. This is the liquidity measure that relies on the most complete set of information and, thus, is our benchmark for the other measures. Recall that the only caveat with the effective bid-ask spread is that it requires the information on the initator of the trade, which is only available for a smaller sample of German bonds that have quote information. Figure 7.2 shows the average level of effective bid-ask spread in German (U.S.) bonds in Panel (a) (Panel (b)), separated for financial and non-financial bonds. For both countries, there is a sharp increase in illiquidity, associated with the financial crisis at the end of 2008. Liquidity in U.S. bonds has since improved steadily, and the average level of the effective

[18]The criteria for the calculation of price dispersion are almost identical to those of the Amihud and Roll measures as well as the imputed round-trip cost. In the interest of readability, we only show the line corresponding to price dispersion in the figure, while the other measures behave in a similar manner.

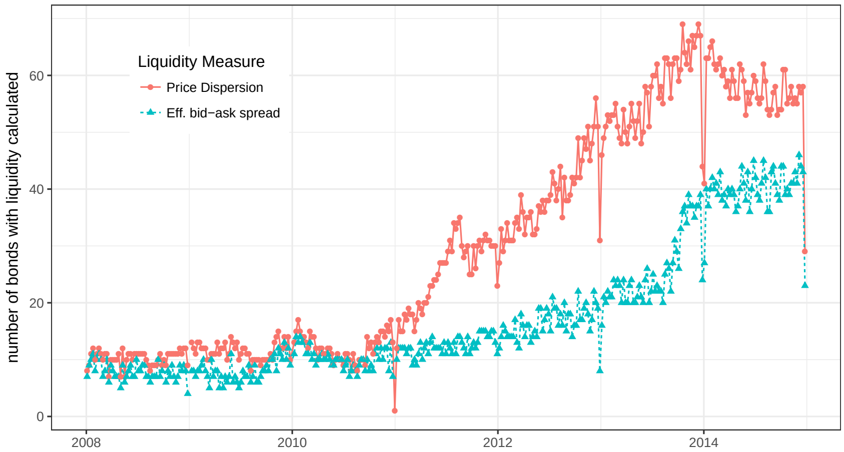

(a) German non-financial bonds

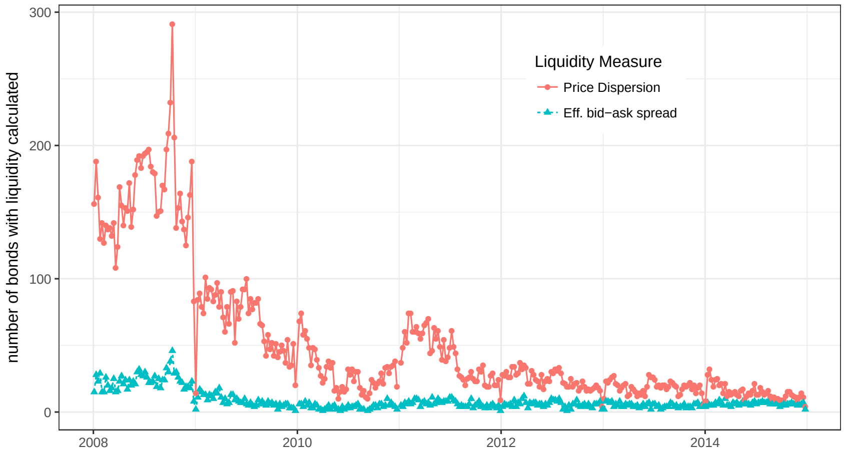

(b) German financial bonds

(c) Number of bonds in TRACE with liquidity calculated at the weekly level.

Figure 7.1: Number of liquid bonds: Number of bonds for which we compute liquidity measures at the weekly level. The sample Panel (a) are German non-financial bonds and German financial bonds for Panel (b). *Price dispersion*, *Roll measure*, *Amihud measure* and *imputed round-trip cost* require 8 observed trades. The calculation of the *effective bid-ask spread* for our German sample further requires the availability of quotes to infer whether a trade was buyer- or seller-initiated, leading to less observations. Year-end effects are clearly discernible in all transaction-based measures of liquidity. For U.S. bonds in Panel (c) the trade sign is included in TRACE and we distinguish only between financial and non-financial bonds.

Table 7.6: Liquidity statistics: Summary statistics of liquidity measures of German and U.S. corporate bonds. Panel A shows statistics for non-financial bonds and Panel B for financial bonds. *Amihud* is the Amihud measure of price impact obtained as the mean ratio of absolute log returns to trade volumes. *Price dispersion* is the root mean squared difference between traded prices and the market valuation proxied by the volume-weighted average trade price. *Roll* is the Roll measure, a proxy for the round trip cost and obtained as twice the square root of the negative auto-covariance of returns. *Effective bid-ask spread* is the difference between the average sell and the average buy price, normalized by their midprice. The trade sign (buy/sell) is inferred by comparing to quotes from Bloomberg. *Imputed round-trip cost* proxies bid-ask spread by comparing the highest to the lowest price of a set of transactions with identical volumes. All measures were computed for every bond and week where there were at least 8 trades with sufficient information available and winsorized at the 0.5% and 99.5% quantile. Units are basis points except for the Amihud measure which is given as in units of basis points per million EUR (USD) for German bonds (U.S. bonds).

Panel A: Non-financial bonds

	mean	stddev	p05	median	p95	# bonds	# bond-weeks
German non-financial bonds							
Amihud (bp per M EUR)	349.47	1447.58	5.17	42.96	1210.15	118	11,057
Price Dispersion (bp)	57.41	59.35	8.61	38.02	177.66	118	11,045
Roll Measure (bp)	67.25	50.61	15.01	54.18	167.02	118	10,674
Effective Bid-Ask (bp)	81.65	109.63	11.00	48.78	282.32	87	6,278
Imputed Roundtrip (bp)	120.18	176.15	9.73	58.38	516.45	118	10,922
TRACE non-financial bonds							
Amihud (bp per M USD)	77.43	115.73	1.41	40.63	273.84	1,658	85,580
Price Dispersion (bp)	90.27	110.88	10.28	57.80	263.07	1,658	85,580
Roll Measure (bp)	128.86	124.76	16.87	94.84	346.51	1,625	79,086
Effective Bid-Ask (bp)	136.04	168.95	5.48	73.91	447.54	1,582	79,119
Imputed Roundtrip (bp)	66.81	69.21	5.32	45.76	199.80	1,602	81,267

Panel B: Financial bonds

	mean	stddev	p05	median	p95	# bonds	# bond-weeks
German financial bonds							
Amihud (bp per M EUR)	29.81	61.17	0.71	11.88	106.85	1,508	18,917
Price Dispersion (bp)	34.69	34.41	3.00	24.85	98.42	1,508	18,744
Roll Measure (bp)	53.84	50.63	3.65	39.24	153.58	1,508	17,945
Effective Bid-Ask (bp)	102.62	92.88	12.33	77.71	270.45	323	3,048
Imputed Roundtrip (bp)	45.83	53.53	2.75	29.26	143.59	1,476	17,253
TRACE financial bonds							
Amihud (bp per M USD)	110.88	190.21	3.45	51.96	401.21	4,318	276,466
Price Dispersion (bp)	123.82	192.00	10.19	73.89	389.02	4,318	276,467
Roll Measure (bp)	164.08	192.58	17.83	110.61	452.80	4,265	256,743
Effective Bid-Ask (bp)	199.76	254.81	6.53	130.67	582.02	4,249	263,809
Imputed Roundtrip (bp)	88.00	100.14	6.63	56.87	263.56	4,218	260,827

bid-ask spread was below 100 basis points. Financial bonds were, on average, less liquid through 2012, a gap that has closed in 2013 and 2014. We observe similar dynamics for German financial bonds, whereas German non-financials seem to become more illiquid from 2013 on.

Again, we cannot distinguish between the two markets by simply looking at the dynamics of the effective bid-ask spread, to ascertain whether this is due to the changing composition of our sample towards more bonds of smaller issue size (which is typically associated with lower liquidity), or it is due to a general deterioration of market liquidity. From Figure 7.1, we know that the number of liquid German bonds has changed significantly in the last part of the period and, instead, the amount outstanding has not increased significantly. From Figure 7.2, it also appears that during the first part of our sample period, German corporate bonds were actually more liquid than their U.S. counterparts. It should be borne in mind, however, that the limited sample for which we compute the effective bid-ask spread is likely to be biased towards more liquid bonds. In this case, the clear differences between the two markets in terms of the distribution of the number of trades per bond in the sample highlighted in Table 7.3 prevent us from comparing the patterns of these figures in Panels (a) and (b).

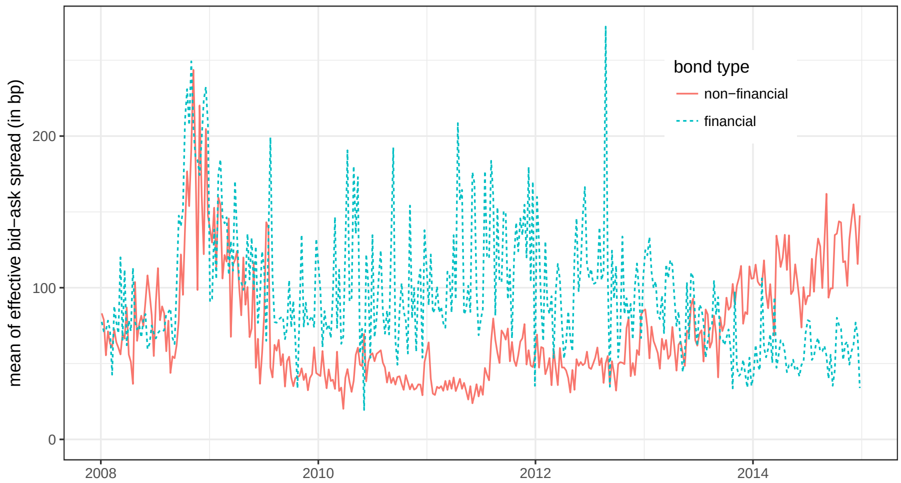

(a) German bonds (WpHG data).

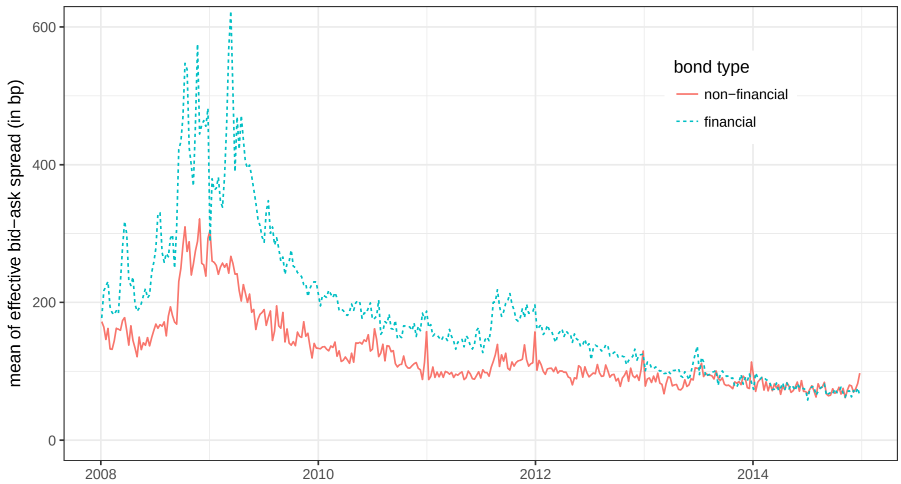

(b) U.S. bonds (TRACE data).

Figure 7.2: Effective bid-ask spread: *Effective bid-ask spread* is the difference between the average sell and the average buy price, normalized by their midprice and given in basis points.

For the same reason that the effective bid-ask spread is based on a smaller sample of bonds, it can result in a rather noisy liquidity measure, as suggested by the spikes (especially for financial bonds) in Panel (a) of Figure 7.2. Importantly, we find these observed trends to be robust in the other liquidity measures, for which we do not have the trade initiation limitation and, thus, can use the full sample of liquid bonds. Panels (a) to (d) of Figure

7.3 show the time series dynamics of the price dispersion, Roll, imputed round-trip cost and Amihud measures respectively for German bonds, and for their counterparts in Figure 7.4 for U.S. bonds. The dynamics of these measures for U.S. bonds coincides with those for the effective bid-ask spread, whereas the dynamics of liquidity in the German market requires more detailed attention.

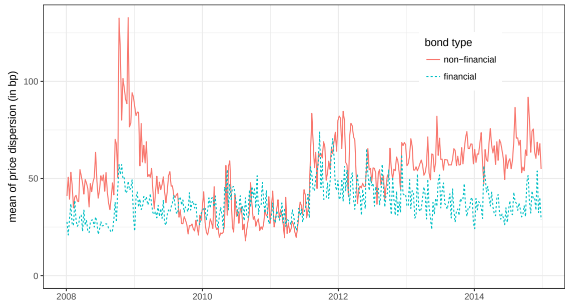

(a) Price dispersion

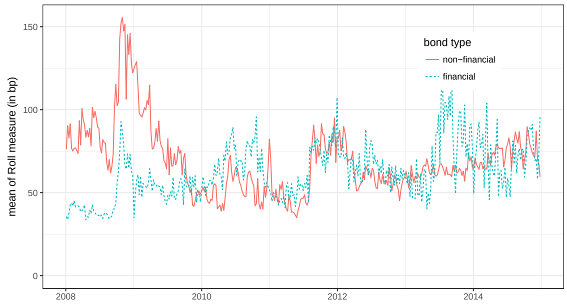

(b) Roll measure

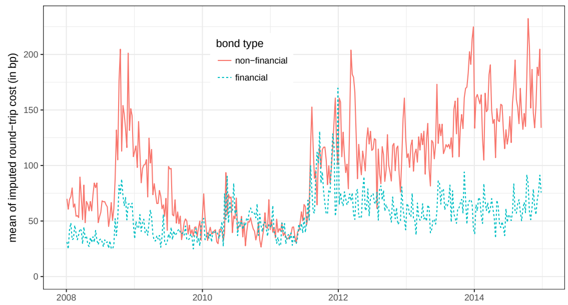

(c) Imputed round-trip cost

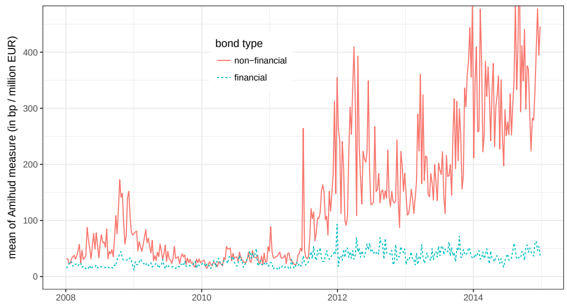

(d) Amihud measure

Figure 7.3: Further liquidity measures - German bonds. *Price dispersion* is the root mean squared difference between traded prices and the market valuation proxied by the volume-weighted average trade price. *Roll* is a proxy for the round trip cost and obtained as twice the square root of the negative auto-covariance of returns. *Imputed round-trip cost* proxies bid-ask spread by comparing the highest to the lowest price of a set of transactions with identical volumes. *Amihud* is a measure of price impact obtained as the mean ratio of absolute log returns to trade volumes. All measures are given in basis points except for the Amihud, which is given in in units of basis points per million EUR of trade volume.

For German corporate bonds, we first note that the qualitative patterns that emerge from the single panels of Figure 7.3 are quite different from another. Price dispersion and imputed round-trip cost behave in a similar manner as the effective bid-ask spread, but several features of the Roll and Amihud measures are not replicated in the other liquidity metrics. Keeping in mind that the Roll measure is based on the auto-covariance of returns, and the Amihud measure on average returns during an observation period, both measures ultimately rely on the pattern of bond returns. The divergence between these two measures may be mostly because not all time stamps and prices are observed in the dataset, and hence returns are not properly defined. We, therefore, focus our attention on the effective bid-ask spread, price dispersion, imputed round-trip cost and, to some extent, the Amihud measure.[19]

We now look at the evolution of liquidity in the German sample and compare it to the U.S. market. A common feature to these time series is the sudden increase in illiquidity at the end of 2008, due to the financial crisis, characterized by a sharp spike at the end of 2008,

[19] We report correlation coefficients for the liquidity measures in Table 7.10 in Appendix 7.B. The results are in line with our more descriptive findings above.

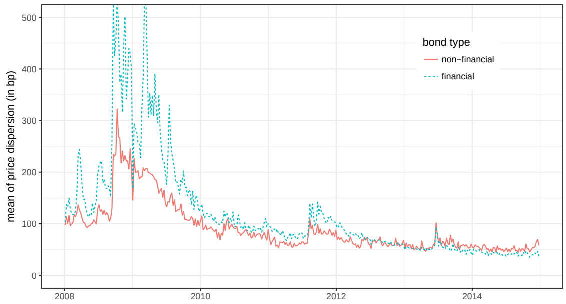

(a) Price dispersion

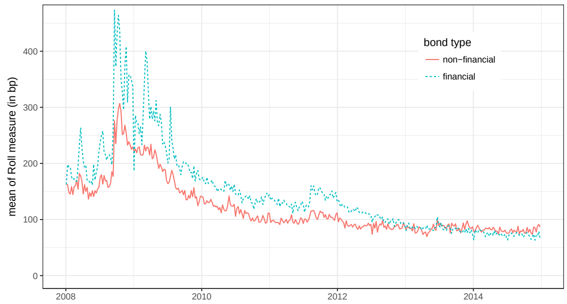

(b) Roll measure

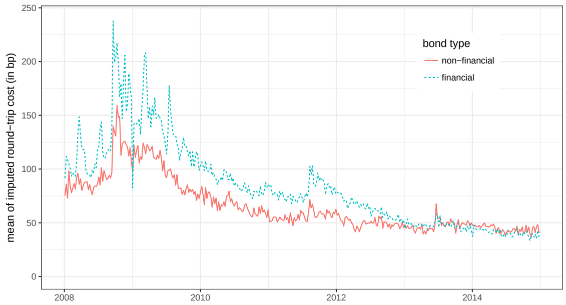

(c) Imputed round-trip cost

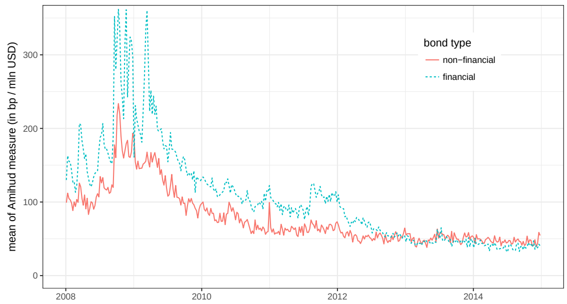

(d) Amihud measure

Figure 7.4: Further liquidity measures - TRACE. Price dispersion is the root mean squared difference between traded prices and the market valuation proxied by the volume-weighted average trade price. Roll is a proxy for the round trip cost and obtained as twice the square root of the negative auto-covariance of returns. Imputed round-trip cost proxies bid-ask spread by comparing the highest to the lowest price of a set of transactions with identical volumes. Amihud is a measure of price impact obtained as the mean ratio of absolute log returns to trade volumes. All measures are given in basis points except for the Amihud, which is given in in units of basis points per million USD of trade volume.

in conjunction with the Lehman Brothers bankruptcy event. The impact of the European sovereign debt crisis of 2012 is also evident in the graphs corresponding to German bonds, whereas this effect is of minor consequence for U.S. bonds. All our measures, with the exception of the Roll measure, show a divergence of German non-financial bonds, which became more illiquid than German financial bonds. As suggested above, the figures we present allow for at least two alternative explanations. First, corporate bonds may have become more illiquid (and also more numerous) in general. Second, trading activity may have expanded to newly issued bonds which are more illiquid and, thus, while the liquidity of individual bonds may not have changed, the overall liquidity level might be lower, since our sample gradually includes more illiquid securities. We aim to distinguish better between these two effects in our panel regressions below, since no such trend is observed for U.S. corporate bonds.

In terms of market liquidity, Table 7.6 provides a simple comparison of the different liquidity measures. In Panel A, the price dispersion, Roll and effective bid-ask are on average lower for German non-financial bonds than those bonds in TRACE, whereas Amihud illiquidity and the round trip cost are higher. On the other hand in Panel B, financial bonds in TRACE are shown to be on average always more illiquid than German financial bonds. However, a simple average comparison is misleading. Notice that the number of bond-week observations of the German sample are only as much as one-eighth of the TRACE sample in Panel A, wheras in Panel B, this ratio is even less than 10 percent.

7.5.4 Determinants of liquidity

In this subsection, we perform a panel analysis that allows us to investigate whether the liquidity measures for German bonds are related to certain bond characteristics, and whether the sensitivity of these measures with respect to these characteristics is of similar magnitude for the U.S. and German samples. The results of the panel analysis are reported in Table 7.7.

Table 7.7 shows that, for German corporate bonds, the amount issued is not an important determinant for the various liquidity measures, the only exception being the effective bid-ask spread, which shows a negative and statistically significant coefficient. In contrast, for the U.S. corporate bond market, this is a highly significant variable for almost

Table 7.7: Determinants of liquidity: Regression of the liquidity measures described in Table 7.6, computed at weekly frequency, on bond and bond-time characteristics. The sample for Panel A are German corporate bonds that carry a rating and U.S. bonds with a rating for panel B. All liquidity measures are calculated for each bond on a weekly basis and winsorized at the 0.5% level from both tails. The explanatory variables are given by bond size (in million EUR/USD), rating step (1 being the best rating on the scale, cf. Table 7.5), time to maturity (in years), age (in years), volume traded in the same bond and week (in million EUR/USD), and a dummy indicating whether the bond is issued by a financial firm. Year-fixed effects are included with the year 2014 as a baseline. Test statistics, derived from standard errors corrected for heteroscedasticity and clustered at the firm level, are given in parenthesis. The significance is indicated as follows: $^{*} < 0.1$, $^{**} < 0.05$, $^{***} < 0.01$.

Panel A: German liquid bonds

	Price Dispersion	Roll	Amihud	Imputed Roundtr.	Eff. Bidask
Bond size (million EUR)	0.005	0.010	−0.008	0.013	−0.032**
	(0.991)	(1.104)	(−1.246)	(1.357)	(−2.294)
Rating step	2.703***	2.754***	2.631	6.267***	3.257
	(3.439)	(2.681)	(1.331)	(3.553)	(1.339)
Time-to-maturity (years)	4.466***	7.853***	1.599	6.307***	4.902**
	(4.661)	(4.631)	(1.275)	(4.469)	(2.324)
Age (years)	1.294**	3.154***	2.664**	2.602**	−1.913
	(2.485)	(3.591)	(2.319)	(2.258)	(−0.886)
Volume (million EUR)	0.029	−0.089*	−0.171*	−0.093**	−0.241*
	(0.989)	(−1.758)	(−1.959)	(−2.538)	(−1.859)
Dummy: financial	17.694***	25.066***	12.430	32.080***	50.080***
	(2.914)	(3.184)	(1.196)	(3.157)	(4.868)
Dummy: year 2008	12.684***	5.508	−0.104	14.300*	66.761***
	(3.179)	(0.566)	(−0.015)	(1.689)	(6.677)
Dummy: year 2009	13.408***	9.120	−3.639	10.278	65.888***
	(3.588)	(0.876)	(−0.572)	(1.121)	(5.379)
Dummy: year 2010	8.307**	9.298	−2.442	5.366	36.318***
	(2.458)	(0.860)	(−0.413)	(0.664)	(4.832)
Dummy: year 2011	8.064*	4.367	0.209	8.857	42.605***
	(1.866)	(0.356)	(0.026)	(0.884)	(4.114)
Dummy: year 2012	10.593**	5.904	9.936	12.692	24.705***
	(2.515)	(0.558)	(1.030)	(1.371)	(3.213)
Dummy: year 2013	3.912	6.279	2.307	6.901	9.088
	(1.506)	(1.060)	(0.568)	(1.562)	(1.470)
Constant	−23.173*	−20.737	−4.062	−53.029**	−2.586
	(−1.856)	(−1.277)	(−0.163)	(−2.271)	(−0.081)
Observations	21,439	20,691	21,599	19,965	5,709
number of bonds	1,391	1,343	1,411	1,341	305
R^2	0.090	0.123	0.029	0.104	0.251
Adjusted R^2	0.089	0.123	0.029	0.103	0.250

Panel B: TRACE liquid bonds

	Price Dispersion	Roll	Amihud	Imputed Roundtr.	Eff. Bidask
Bond size (million USD)	−0.018***	−0.012***	−0.022***	−0.011***	−0.062***
	(−3.631)	(−3.423)	(−5.365)	(−5.478)	(−6.560)
Rating step	12.407***	8.300***	4.849***	3.612***	11.218***
	(8.504)	(7.027)	(4.048)	(5.989)	(3.979)
Time-to-maturity (years)	3.906***	5.946***	4.644***	3.669***	6.181***
	(10.006)	(16.419)	(11.519)	(13.845)	(10.940
Age (years)	0.931	0.845	3.705***	0.301	2.779***
	(0.979)	(1.061)	(3.563)	(1.029)	(2.776
Volume (million USD)	0.032	−0.005	−0.051**	−0.004	−0.153***
	(1.397)	(−0.295)	(−2.316)	(−0.383)	(−2.678
Dummy: financial	44.110***	45.167***	46.378***	27.453***	80.162***
	(7.518)	(9.030)	(7.460)	(9.603)	(8.799
Dummy: year 2008	179.640***	168.019***	142.558***	82.966***	190.142***
	(10.094)	(13.168)	(9.656)	(15.139)	(7.803
Dummy: year 2009	146.160***	125.867***	107.225***	72.495***	165.765***
	(14.781)	(17.150)	(15.604)	(17.675)	(13.161
Dummy: year 2010	25.115***	40.329***	37.112***	25.934***	38.944***
	(3.562)	(6.833)	(6.236)	(8.697)	(4.237
Dummy: year 2011	24.103***	33.459***	30.538***	21.168***	34.972***
	(4.212)	(6.194)	(5.770)	(6.702)	(4.455
Dummy: year 2012	8.979*	13.384***	6.927**	9.131***	21.325***
	(1.856)	(4.925)	(2.535)	(6.217)	(3.976
Dummy: year 2013	1.848	2.119	−1.729	1.491	2.979
	(0.510)	(0.773)	(−0.718)	(0.957)	(0.553
Constant	−90.503***	−52.730***	−56.609***	−21.936***	−57.717**
	(−5.798)	(−3.864)	(−4.055)	(−3.326)	(−2.134
Observations	357,132	357,132	357,132	344,871	357,132
number of bonds	5,660	5,660	5,660	5,617	5,660
R^2	0.322	0.265	0.206	0.281	0.306
Adjusted R^2	0.322	0.265	0.206	0.281	0.306

all the various liquidity measures. A potential explanation for this divergence between the two markets is the relatively lower dispersion in the amounts outstanding in the case of German corporate bonds, in particular due to the absence of bonds with very large amounts outstanding.

The credit rating is also relevant for liquidity, i.e., with more risky bonds being less liquid, since for three out of five liquidity measures, this variable is statistically different from zero at the 1% level. For the U.S., all the liquidity measures are sensitive to the credit rating variable with the respective coefficients being statistically significant. However, if we compare the coefficients between the two panel regressions, we observe that each additional change of one notch in the credit rating is associated with a larger reduction of liquidity, which could be due to a sample selection bias: As we show in Table 7.5, the U.S. corporate bond sample contains a larger fraction of non-investment grade bonds compared to the German counterpart.

The time to maturity variable is highly significant and positive for both German and U.S. corporate bonds, with a similar magnitude, for most of the liquidity measures. The only exception is the Amihud measure for the German sample. The age variable is slightly significant for almost all the liquidity measures for the German corporate bonds and, instead, is significant only for the Amihud and the effective bid-ask spread measures for the U.S. sample. The traded volume of the bond in a given week improves the liquidity of these bonds, but this variable is only marginally significant both for the German and U.S. sample.

The dummy variable for financial bonds is highly significant for both German and U.S. samples, indicating that financial bonds are, on average, more illiquid than non-financial bonds. A comparison of the coefficients for the financial dummy between the two samples indicates that the reduction in liquidity is larger, on average, for U.S. financial bonds than for the German bonds, but this may be again due to the marginally larger proportion of financial bonds considered for the U.S. market.

The panel analysis allows us also to investigate whether liquidity, in general, has improved in recent years. In doing this, we control for bond characteristics that may have changed over time in our sample. The reference year is 2014 and, as the Table shows, dummies from 2008 until 2012 are statistically different from zero and positive at least for three out of five liquidity measures. The dummy variable for 2013 is not statistically significant, indicating that liquidity is quite similar in the years 2013 and 2014. This result indicates that, at least for these three measures, liquidity has improved after controlling for the change in sample composition of bond characteristics. The same applies for the U.S. where, in this case, the variable is positive and significant for all five liquidity measures.

In summary, Table 7.7 provides a panel analysis from which we could observe that the signs of the coefficients of the different determinants are very similar between the two groups: A typical corporate bond is more liquid with a larger bond size, a better rating, a shorter time-to-maturity, a younger age, and a larger volume. Our results are in line with traditional search theories in OTC markets; larger assets are easier to find, and are therefore cheaper to trade. A similar argument applies for the credit rating and the time to maturity: when a bond has a better rating (shorter maturity), dealers will require a lower spread, since it is less risky for them to hold it in their inventory. Trades with larger volumes are most likely being used by larger and more sophisticated investors, who bear lower search costs, have greater bargaining power, and trade with lower spreads, in equilibrium. It is noteworthy to mention that these characteristics less significantly drive the liquidity in the German sample than in the U.S. sample. This is especially true for bond size and volume. Moreover, the yearly dummies indicate that there is an improvement of the liquidity of bonds for both the U.S. and German bonds from 2008 to 2014. Overall, the time dummies confirm the patterns observed in the Figures 7.2 through 7.4.

7.5.5 Matched sample analysis

Our analysis, thus far, has pointed to the conclusion that the U.S. corporate bond market is far more liquid than the German one with a large fraction of bonds that are traded more frequently than in German bonds and a significant larger volume of trade, even adjusting for the differences in the amount outstanding. On the other hand, (i) the simple analysis of the average of the different liquidity measures there is not clear evidence that one market is more liquid than the other for very liquid bonds and (ii) the German and U.S. samples behave similarly: trends in liquidity over time are roughly similar, and so are the magnitudes of liquidity measures and their drivers. However, these inconsistencies could be largely driven by the different compositions of the two samples.

Therefore, in order to provide a more appropriate comparison between the two markets, we perform two additional sets of analysis. The first is a simple horse race where we compare the *most* liquid bonds in the German sample with the corresponding ones we observe in the U.S. sample. This allows us to compare samples with exactly the same number of bonds. This comparison is useful from the point of view of an investor who holds a portfolio of German corporate bonds that are most liquid, without taking into account considerations about maturities, amount outstanding or other bond characteristics. In doing this, we look at the sample of liquid German bonds in each week and take the same number of U.S. bonds that are most liquid in the week, after ranking them from most to least liquid. We then average the liquidity measures across the bond and time dimensions. We perform this analysis for all liquidity measures and find that the most liquid U.S. bonds are far more liquid than the most liquid German bonds. The difference between the two markets is 15.5 basis points for price dispersion and 63.7 basis points for the imputed round-trip cost, both differences being highly significant. Any such comparison based on the most liquid bonds of each country (which could well be the motive of an investor) is clearly in favor of U.S. bonds.

In the second analysis, we investigate whether a German corporate bond is as liquid as a *comparable* U.S. corporate bond. Simply comparing the time-series of liquidity measures shown, for example in Figure 7.2, is not sufficient, since it completely neglects the underlying sample composition, and the changes therein. Instead, in this section, we match German bonds with U.S. bonds of comparable characteristics. This allows us to repeat the panel regression of liquidity from the previous section for the matched sample, and quantify any differences in liquidity through an added dummy variable that controls for whether a bond is part of our German sample or the U.S. one.

Specifically, we make use of our finding from the previous section that the drivers of liquidity are common to both samples and employ a propensity score matching approach to create our matched sample. In order to rule out variations in liquidity over time as a source of differences, our matching of bonds is performed on a weekly basis. In other words, for each week we consider the set of German bonds for which we compute liquidity measures as our treated sample and construct a control sample of U.S. bonds using "nearest neighbor" matching, based on the criteria of amount outstanding (in USD), coupon rate, rating notch, age, maturity at issuance, average trade volume and a rank measure of trading activity. We match only financial bonds to financial bonds and, likewise, for non-financial bonds, and impose minimum closeness criteria for a matched pair to be part of our sample.[20] In order to minimize the impact of any remaining sample differences, we then repeat the analysis of Table 7.7 for the matched sample with two slight modifications: First, we include a dummy variable that is one for German bonds (our treatment group) and zero for U.S. bonds (control group). Second, we drop regression variables that are related to the sample coverage.[21]

Table 7.8 shows the results of the matched panel regression. Drivers of liquidity are consistent with the findings from the previous sections: larger and better-rated bonds are more liquid, and so are non-financial bonds. Illiquidity increases with both age and time-to-maturity, i.e., bonds with shorter maturities and bonds that have been issued recently are more liquid. Most strikingly, the coefficient for German bonds is negative in all liquidity measures, and significantly so for all but the imputed round-trip cost. This implies that liquid German bonds are actually 37 to 67 basis points more liquid than comparable U.S. bonds even after controlling for bond and time effects.

It is worth discussing possible reasons for this surprising finding, which is most likely related to our natural focus on the most liquid German bonds. Potentially the lack of

[20] To be precise, we only keep matches within half a standard deviation of the distance measure from the original observation. Our results are robust to various variations in the matching approach, such as the variable set for matching, matching with or without replacement, matching to a larger control set and other thresholds for closeness.

[21] Otherwise, we would introduce a bias from the more complete coverage of our U.S. sample against the admittedly partial coverage of trading activity in German bonds, in our dataset. In practice this amounts to dropping the variable "volume".

Table 7.8: Propensity Score Matched Regressions: Regression of the liquidity measures described in Table 7.6, computed at weekly frequency, on bond and bond-time characteristics. All liquidity measures are calculated for each bond on a weekly basis and winsorized at the 0.5% level from both tails. The sample is a matched sample of German corporate bonds that carry a rating (treatment group) and a control group of U.S. bonds matched at weekly frequency based on amount outstanding, time-to-maturity, maturity at issuance, rating, coupon rate, average trade volume and classification as financial bond. See section 7.5.5 for details of the matching process. The explanatory variables are given by bond size (in million EUR/USD), rating step (1 being the best rating on the scale, cf. Table 7.5), time to maturity (in years), age (in years), a dummy indicating whether the bond is issued by a financial firm and a dummy indicating whether a bond is part of the German (treatment) sample. Year-fixed effects are included with the year 2014 as a baseline. Test statistics, derived from standard errors corrected for heteroscedasticity and clustered at the firm level, are given in parenthesis. The significance is indicated as follows: $^{*} < 0.1$, $^{**} < 0.05$, $^{***} < 0.01$.

	Price Dispersion	Roll	Amihud	Imputed Roundtr.	Eff. Bidask
Bond size (million USD)	-0.012^{**}	-0.004	-0.018^{***}	-0.005	-0.058^{***}
	(-2.107)	(-0.879)	(-3.414)	(-1.272)	(-6.444)
Rating step	5.143^{***}	4.060^{***}	4.968^{***}	2.474^{***}	5.276^{***}
	(4.344)	(4.448)	(4.658)	(5.117)	(3.181)
Time-to-Maturity (years)	5.052^{***}	7.515^{***}	5.094^{***}	4.647^{***}	7.695^{***}
	(8.015)	(9.389)	(6.628)	(10.995)	(8.282)
Age (years)	3.687^{***}	3.117^{***}	3.857^{***}	1.328^{**}	4.665^{***}
	(3.479)	(3.790)	(3.868)	(2.068)	(3.432)
Dummy: financial	37.398^{***}	30.121^{***}	18.365^{***}	16.940^{***}	60.484^{***}
	(4.291)	(4.134)	(2.829)	(3.452)	(5.294)
Dummy: German sample	-67.428^{***}	-59.602^{***}	-36.856^{***}	-5.631	-41.874^{***}
	(-8.971)	(-9.512)	(-5.174)	(-1.211)	(-3.796)
Dummy: year 2008	81.144^{***}	77.568^{***}	50.724^{***}	33.153^{***}	106.122^{***}
	(6.609)	(6.104)	(4.352)	(5.312)	(6.967)
Dummy: year 2009	68.364^{***}	65.901^{***}	34.600^{***}	28.859^{***}	109.692^{***}
	(7.213)	(6.531)	(3.139)	(4.471)	(7.441)
Dummy: year 2010	22.329^{***}	26.105^{***}	12.850	9.365^{**}	41.348^{***}
	(3.879)	(3.710)	(1.161)	(2.275)	(5.251)
Dummy: year 2011	20.886^{***}	20.763^{***}	8.239	8.746^{*}	40.200^{***}
	(3.675)	(2.744)	(0.716)	(1.733)	(4.403)
Dummy: year 2012	16.042^{***}	13.768^{*}	5.183	6.528	25.173^{***}
	(3.008)	(1.799)	(0.339)	(1.241)	(3.048)
Dummy: year 2013	4.842	5.290	-13.021	-1.153	9.590
	(1.343)	(1.320)	(-1.231)	(-0.357)	(1.553)
Constant	-23.801	-3.454	-15.325	-2.374	-8.688
	(-1.286)	(-0.233)	(-1.039)	(-0.275)	(-0.358)
Observations	22,842	22,038	22,978	21,416	8,648
number of bonds	3,898	3,753	3,942	3,801	1,916
R^2	0.194	0.186	0.064	0.103	0.226
Adjusted R^2	0.193	0.186	0.064	0.102	0.225

transparency in the German market leads to a "crowding"-effect, where market participants concentrate their trading activity in a small set of bonds. These, as a consequence, result more liquid than comparable bonds in a transparent market, where investors are able to diversify their trading across more bonds, which then appear less liquid *at the individual level.*

7.6 Conclusion

In this chapter, we study the impact of transparency on liquidity in OTC markets, the subject of an important debate for academia, industry and regulators. Supporters of OTC market transparency argue that it reduces the asymmetry of information between dealers and investors, and hence encourages the participation of retail/uninformed investors. On the other hand, OTC market transparency could increase transaction costs for investors, by eliminating dealers' information rents and, thus, their incentives to compete, or even participate in the market. We contribute to this debate by providing an analysis of liquidity in a corporate bond market without trade transparency (Germany), and comparing our findings to one with full disclosure (the U.S.). We employ a unique regulatory dataset of transactions of German financial institutions from 2008 until 2014.

Our analysis consists of four parts. First, we provide a general description of the German corporate bond market, focusing on the key characteristics of the bonds and their trading activity. Second, we analyze the time-series evolution of liquidity in the German market. Third, we study the determinants of liquidity in the cross-section of the German market with panel regressions on bond characteristics. Fourth, we use a matched-sample approach to compare the transaction costs of similar bonds, at the same point in time, in the German and the U.S. market.

In our descriptive analysis, we find that overall, observed trading activity is much lower in the German market. The bonds that trade at least once 8 times per week are only 17% of our sample, as against 74% of the traded sample in the U.S. universe. Looking at the market as a whole, overall liquidity is clearly much higher in the U.S., with a significantly larger number of securities that trade often. This result is consistent with various theoretical studies that show that transparency lowers costs for unsophisticated investors and, therefore, incentivizes their participation in the market. Our time-series analysis shows that the average transaction costs for German corporate bonds spiked sharply during the 2008-2009 global financial crisis, but less so during the 2010-2012 sovereign debt crisis. The cross-sectional regressions confirm that, similar to the U.S., the determinants of German corporate bond liquidity are in line with search theories of OTC markets. A bond is more liquid if it has a larger issue size, a better rating, a shorter time-to-maturity, a younger age, and a larger volume traded. Finally, the matched sample analysis reveals that *frequently* traded German bonds have lower transaction costs than comparable bonds in the U.S. market. The difference in round-trip costs is within a range of 37-67 basis points, depending on the liquidity measure. Although surprising at first blush, this finding is in line with studies that highlight the potential unintended consequences of an increase in transparency in OTC markets.[22] A possible explanation for our finding is that, when there is little transparency, investors concentrate their demand into a few well-traded assets. As a result, the liquidity of these few bonds is particularly high, while the others are barely traded. On the other hand, when overall transparency increases, investors spread their portfolios across a wider range of assets, given the higher level of information available. While the overall market liquidity improves, there is less relative demand for the previously "well-traded assets," and hence their transaction costs could become higher.

Our results are of considerable interest to market participants, both on the buy and sell side, as well as to regulators. Furthermore, our findings provide a benchmark for future research on the measurement of liquidity in European fixed-income markets, which could benefit from more detailed data. In fact, from January 2018, MiFID II requires all firms in the European Union to publish details of their OTC transactions in non-equity instruments almost in real time (commonly referred to as OTC post-trade transparency). Hence, the introduction of MiFID II is likely to change the landscape of the European fixed-income

[22]See, for example Naik et al. (1999), Bloomfield and O'Hara (1999), Bhattacharya (2016) and Friewald et al. (2017).

market with new regulations on the provision of trading services and reporting. There are significant concerns in the financial services industry about the cost of fulfilling these requirements, especially since a mandatory transparency could potentially hamper liquidity due to the withdrawal of dealers who would be concerned about "showing their hand." Our results imply that the pre-MiFID opaque structure could be one of the reasons why a selected set of German bonds are revealed to be more liquid than their U.S. counterparts.

Although this study provides some insights into a market that has so far been *dark*, for many of its market participants, from an academic point of view, and even from a regulatory perspective, much still remains to be done to address specific policy questions. This said, our findings lay the foundations for a future academic and regulatory research agenda. Without data regarding the liquidity of the market before the implementation of MiFID II along the lines of our findings, it would be impossible to assess empirically whether pre- and post-trade transparency in the European bond market indeed improved liquidity. While the scope of MiFID II is vast, it remains to be seen how well and how broadly the directive will be implemented. Suffice it to say that even in the U.S., the implementation of TRACE took place over several years, just for the corporate bond market, and its extension to other fixed income markets is still work in progress. It is fair to assume that the dissemination of a reasonable sample of such data will take several years. Our analysis on liquidity has shown, for instance, that given the low number of bonds traded at significant frequency, any increase in the number of these liquid bonds due to MiFID II regulation would be an achievement. Once reliable data from MiFID II become available in the future, our study could serve as a benchmark of the "before-" period. As the construction of reliable transaction databases will be a challenge in the years to come, we expect our findings to remain valid and relevant in the interim.

Another issue that can be addressed based on our evidence is how the Corporate Asset Purchase Program (CAPP), launched by the ECB in June 2016 as part of its Quantitative Easing (QE), has reshaped the liquidity in the European corporate bond markets. This study could be the basis for a proper analysis of the impact of corporate bond purchases in the CAPP. Our study would provide such a benchmark to assess the impact of this specific program by the ECB, which does not have a counterpart in the QE programs of the U.S. Federal Reserve System (FED). Overall, this chapter serves the purpose of shedding light on what was previously unknown and opening up an important topic for further discussion.

Appendix

7.A TRACE data preparation

We use two main sources of data for our analysis on the U.S. corporate bond market. We obtain information on bond characteristics from Mergent FISD, while TRACE enhanced

[23] In addition to the TRACE standard version, TRACE enhanced includes buy-sell indicators for each transaction, and the trading volume is not capped.

[24] We delete duplicates, trade corrections, and trade cancellations on the same day. Moreover, we delete reversals, which are errors detected not on the same day they occurred.

[25] We adopt a median and a reversal filter. The median filter eliminates any transaction where the price deviates by more than 10% from the daily median or from a nine-trading-day median, which is centered at the trading day. The reversal filter eliminates any transaction with an absolute price change that deviates at the same time by at least 10% from the price of the transaction before, the transaction after and the average between the two.

contains bond transactions' prices, which are used for the calculation of the liquidity measures.[23] For comparability with the WpHG data, our sample spans January 2008 to December 2014. For TRACE, we follow standard data cleansing procedures described by Dick-Nielsen (2009).[24] Furthermore, we implement the price filters used in Edwards et al. (2007) and Friewald et al. (2012).[25] We consider only straight (simple callable and puttable) bonds, and exclude any bond with complex structures or optionalities. Details on amount of observations lost in the cleaning process can be found in Table 7.9.

Table 7.9: Data Cleaning Process - TRACE: This table illustrates the data cleaning process and the number and share of observations remaining after each cleaning step for each year of data and the whole dataset. The data sample comes from TRACE enhanced and covers the period January 2008 - December 2014. When necessary, bond characteristics are matched from MERGENT FISD. The initial number of observations is given as *before cleaning*. *TRACE filter* In the first cleaning step we cleanse the transaction data of errors using the algorithm described in Dick Nielsen (2009). In particular, we delete duplicates, trade corrections and trade cancellations on the same day. Moreover, we remove reversals, which are errors detected on a day later than that of the initial trade. *price filter* Additionally, we implement the price filters Friewald et al (2012). Specifically, we adopt a reversal filter, which should eliminate extreme price movements, and a median filter, which identifies outliers in prices reported in TRACE, within a given time period.

year	before cleaning	TRACE filter	% of raw	price filter	% of raw
2008	8,982,733	5,791,024	64%	5,766,619	64%
2009	15,509,609	9,968,885	64%	9,847,259	63%
2010	16,196,597	9,710,084	60%	9,349,861	58%
2011	14,866,634	9,044,960	61%	8,610,110	58%
2012	16,552,442	9,908,571	60%	9,211,178	56%
2013	16,276,111	9,691,278	60%	8,898,621	55%
2014	15,224,322	9,115,658	60%	8,239,993	54%
$\sum$	103,608,448	63,230,460	61%	59,923,641	58%

For subsetting our baseline sample of vanilla bonds, secured bonds are identified as those with SECURITY_TYPE="SS" in Mergent FISD. Treasury-type bonds are those with bond type among "USBD","USBL","USBN","USNT","USSI","USSP","USTC" in Mergent FISD. Unsecured bonds are the remaining ones that do not fall in any of the previous two categories. To divide unsecured bonds into those issued by non-financial firms and those issued by financial corporations we use the industry classification provided by Mergent FISD. Financial bonds are those with INDUSTRY_GROUP=2, while non-financial bonds are those belonging to any other industry group code.

7.B Liquidity measures

We employ a set of liquidity measures that mostly capture the costs associated with price-impact and round-trip trades, following the presentation in Friewald et al. (2017). While these measures are typically calculated on a daily basis for U.S. TRACE data, all measures calculated here are based on weekly data. This allows us to include more bonds in our analysis that are relatively actively traded, but not typically on a daily basis. We define our notation such that Liq_t^i is the liquidity of bond i in week t and N_t^i is the number of trades in bond i in week t.

- The *Amihud measure*, proposed in Amihud (2002), is our proxy of price impact. The more a trade of a given size shifts the observed price, the higher the Amihud measure and the less liquid the bond. The measure is obtained as the mean ratio of absolute log returns to trade volumes:

$$Amihud_t^i = \frac{1}{N_t^i} \sum_{j=1}^{N_t^i} \frac{|r_t^{i,j}|}{V_t^{i,j}} \tag{7.1}$$

where the index j spans all trades in bond i in week t while $r_t^{i,j}$ and $V_t^{i,j}$ are the (log) return and transaction volume associated with the trade j. The measure is given in units of basis points per million EUR (per million USD for our TRACE sample) and we require at least 8 transactions per week in order to calculate it.

All following measures capture the liquidity component that is associated with the cost of a round-trip trade and are given in units of basis points:

- *Price dispersion* was introduced in Jankowitsch et al. (2011). The idea is that the lower the volatility of prices around the consensus price, the more liquid the bond, since agents are more likely to trade the bond at its fair value. It is calculated as the root mean squared (weighted) difference between traded prices $P_t^{i,j}$ and the market valuation P_t^i proxied by the volume-weighted average trade price.

$$PriceDisp_t^i = \sqrt{\frac{1}{\sum_{j=1}^{N_t^i} V_t^{i,j}} \sum_{j=1}^{N_t^i} (P_t^{i,j} - P_t^i)^2 \, V_t^{i,j}} \tag{7.2}$$

  with $P_t^i = \frac{1}{\sum_{j=1}^{N_t^i} V_t^{i,j}} \sum_{j=1}^{N_t^i} P_t^{i,j}$ As for the Amihud measure we require a minimum of 8 transactions per week.

- *Roll* is the Roll measure that relates the autocorrelation of returns to the bid-ask spread, developed in Roll (1984). It is obtained as twice the square root of the negative auto-covariance of returns.

$$Roll_t^i = 2\sqrt{-\mathrm{Cov}(r_t^{i,j}, r_t^{i,j-1})} \tag{7.3}$$

  We require a minimum of 8 transactions per week in order to compute the Roll measure.

- *Imputed round-trip cost*, developed in Feldhütter (2011) and applied for OTC markets in Dick-Nielsen et al. (2012), proxies the bid-ask spread by comparing the highest to the lowest price of a set of transactions with identical volumes. These transactions are assumed to belong to a round-trip trade and the highest (lowest) of their prices thus to correspond to the prevailing ask (bid) price:

$$ImputedRTCost_t^i = \frac{1}{B_t^i} \sum_{b=1}^{B_t^i} 1 - \frac{\min P_t^{i,b}}{\max P_t^{i,b}} \tag{7.4}$$

  where B_t^i is the number of sets with trades of identical size and $P_t^{i,b}$ is the set of prices that belong to the set b. We require a minimum of 8 transactions per week, and at least 2 transactions of the same size, in order to compute the imputed round-trip cost.

- *Effective bid-ask spread*, proposed in Hong and Warga (2000), is the most restrictive of our measures. It is the difference between the average sell and the average buy price, normalized by their midprice:

$$EffSpread_t^i = \frac{2(\bar{P}_t^{i,\text{sell}} - \bar{P}_t^{i,\text{buy}})}{\bar{P}_t^{i,\text{sell}} + \bar{P}_t^{i,\text{buy}}} \tag{7.5}$$

where $\bar{P}_t^{i,\text{sell}} = \frac{1}{N_t^{i,\text{sell}}} \sum_{j=1}^{N_t^{i,\text{sell}}} P_t^{i,j}$ is the average sell price and idem for the average buy price $\bar{P}_t^{i,\text{buy}}$ of bond i in week t. While in TRACE the trade sign (buy/sell) is provided, for our German sample it needs to be inferred using the algorithm of Lee and Ready (1991) by comparing to quotes from Bloomberg. Therefore, for this measure, we not only require Bloomberg quotes, but also 8 trades which must include at least one buy and sell trade each. In our case, it is possible to obtain negative values for the effective bid-ask spread since we infer the trade sign from daily data, but average over one week. We discard all such negative values.[26]

Table 7.10 gives simple descriptives of the correlation of weakly averages of the liquidity measures described above.

Table 7.10: Correlation of liquidity measures: Correlation of weekly means of liquidity measures for German and U.S. corporate bonds. *Amihud* is the Amihud measure of price impact obtained as the mean ratio of absolute log returns to trade volumes. *Price dispersion* is the root mean squared difference between traded prices and the market valuation proxied by the volume-weighted average trade price. *Roll* is the Roll measure, a proxy for the round trip cost and obtained as twice the square root of the negative auto-covariance of returns. *Effective bid-ask spread* is the difference between the average sell and the average buy price, normalized by their midprice. The trade sign (buy/sell) is inferred by comparing to quotes from Bloomberg. *Imputed round-trip cost* proxies bid-ask spread by comparing the highest to the lowest price of a set of transactions with identical volumes. All measures were computed for every bond and week where there were at least 8 trades with sufficient information available and winsorized at the 0.5% and 99.5% quantile.

Panel A: in levels, all German liquid bonds

	Amihud	EffSpread	ImputedRTCost	PriceDisp	Roll
Amihud	1.00	0.25	0.88	0.69	0.51
EffSpread	0.25	1.00	0.26	0.32	0.32
ImputedRTCost	0.88	0.26	1.00	0.91	0.68
PriceDisp	0.69	0.32	0.91	1.00	0.77
Roll	0.51	0.32	0.68	0.77	1.00

Panel B: in differences, all German liquid bonds

	Amihud	EffSpread	ImputedRTCost	PriceDisp	Roll
Amihud	1.00	0.13	0.40	0.14	0.14
EffSpread	0.13	1.00	0.24	0.33	0.15
ImputedRTCost	0.40	0.24	1.00	0.64	0.29
PriceDisp	0.14	0.33	0.64	1.00	0.40
Roll	0.14	0.15	0.29	0.40	1.00

[26]For example, we could observe buys on the first day, following which the bond price falls and the bond is sold again at a lower price later in the same week.

Panel C: in levels, all U.S. liquid bonds

	Amihud	EffSpread	ImputedRTCost	PriceDisp	Roll
Amihud	1.00	0.98	0.98	0.97	0.99
EffSpread	0.98	1.00	0.97	0.98	0.98
ImputedRTCost	0.98	0.97	1.00	0.96	0.98
PriceDisp	0.97	0.98	0.96	1.00	0.98
Roll	0.99	0.98	0.98	0.98	1.00

Panel D: in differences, all U.S. liquid bonds

	Amihud	EffSpread	ImputedRTCost	PriceDisp	Roll
Amihud	1.00	0.81	0.88	0.90	0.90
EffSpread	0.81	1.00	0.76	0.77	0.73
ImputedRTCost	0.88	0.76	1.00	0.93	0.92
PriceDisp	0.90	0.77	0.93	1.00	0.93
Roll	0.90	0.73	0.92	0.93	1.00

7.C Additional figures

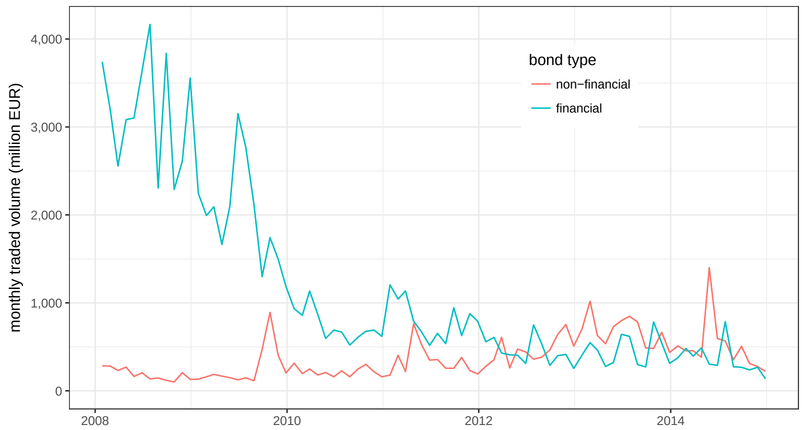

(a) Liquid German bonds.

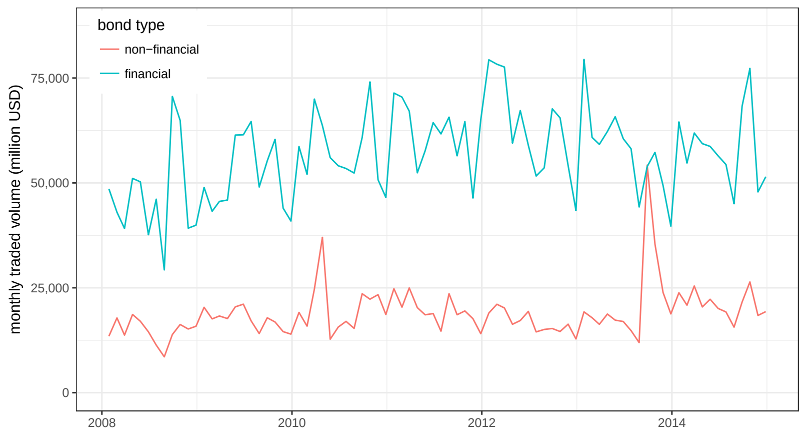

(b) Liquid U.S. bonds.

Figure 7.5: Monthly traded volume: monthly volume of trades in German corporate bonds reported due to WPHG (Panel (a)) and in U.S. corporate bonds (Panel (b)).

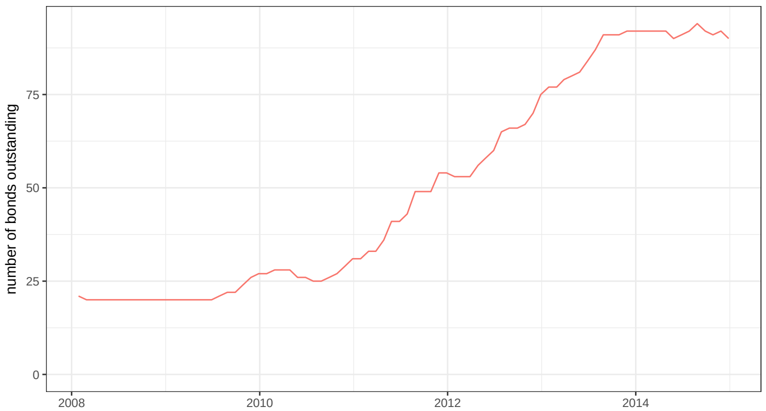

(a) Number of bonds outstanding: German non-financial bonds

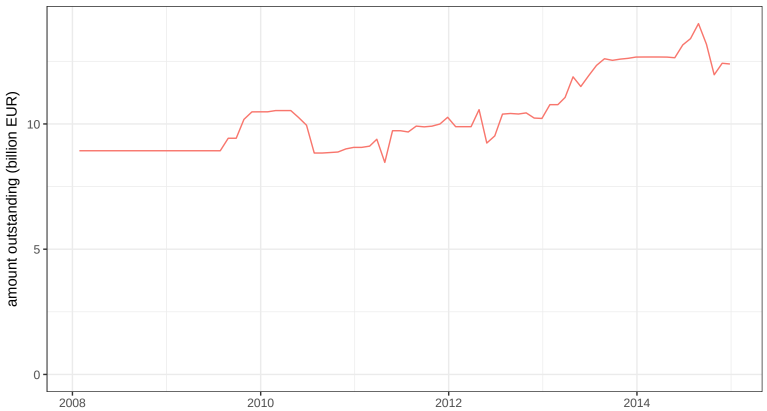

(b) Amount outstanding: German non-financial bonds

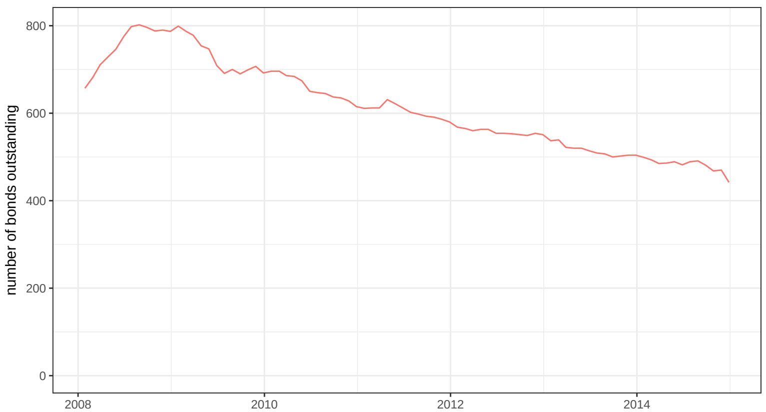

(c) Number of bonds outstanding: German financial bonds

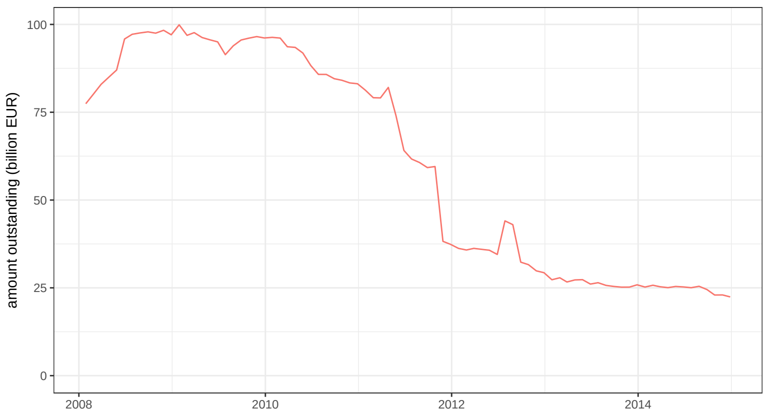

(d) Amount outstanding: German financial bonds

Figure 7.6: Bonds outstanding: Number of bonds outstanding (Panels (a), (c)) and amount outstanding in billion EUR (Panels (b), (d)) across the universe of German corporate bonds for which we compute liquidity measures at least once. Non-financial bonds are in Panels (a) and (b) and financial bonds in Panels (c), (d). Numbers may be inaccurate before April 2009.

Conclusions

In this thesis we have presented a selection of market microstructure studies on fixed-income markets in Europe. Our research spans a wide range of market structures (electronic exchanges, over-the-counter markets, and hybrid markets), market segments (retail as well as interdealer segments), assets (sovereign and corporate bonds) and central topics.

Chapter 4 is concerned with cross-asset price impact. It contributes by extending the work of Gatheral (2010), deriving theoretical limits for the size and form of cross-impact from the condition of absence of dynamical arbitrage. For bounded decay kernels we find that cross-impact must be an odd and linear function of trading intensity and cross-impact from asset i to asset j must be equal to the one from j to i. To test these constraints we use a dataset from the electronic platform MOT, which we have explored in Chapter 3. While we find significant violations of the above symmetry condition in the estimated cross-impact, we show that these are not arbitrageable with simple strategies because of the presence of the bid-ask spread.

In Chapter 5 we study abrupt deteriorations of liquidity conditions. To this end we propose a peak-over-threshold method to identify such illiquidity shocks from limit order book data and we model the time-series of these illiquidity events across multiple assets as a multivariate Hawkes process. This allows us to quantify both the self-excitation of extreme changes of liquidity in the same asset (illiquidity spirals) and the cross-excitation across different assets (illiquidity spillovers). Applying the method to the MTS sovereign bond market, we find significant evidence for both illiquidity spillovers and spirals.

Chapter 6 serves as a bridge between the exchange venues studied in Chapters 3 - 5 and over-the-counter markets (that we study in Chapter 7). It studies the venue choice of dealers in a hybrid bond market with an exchange and a dominant over-the-counter segment and it compares transaction costs across the two venues. Our results indicate that not only transaction costs but also immediacy and post-trade transparency play an important role for this venue choice.

Post-trade transparency is also at the core of our analysis in Chapter 7 where we study liquidity in over-the-counter corporate bond markets. We compare the German market (without transparency) to the U.S. market (with transparency) and find that the latter is much more active. However we also observe that the most liquid and actively traded German corporate bonds are more liquid, in terms of transaction costs, than comparable U.S. bonds. We posit that this is due to transparency, in line with the notion that investors 'crowd' in a small set of the most liquid securities in the absence of transparency.

Let us stress that our methods and findings are far more general than the context of fixed-income markets that we are using here to illustrate them. Especially our results on modeling cross-impact are valid for any two (or more) assets that are related and the methods we use regarding illiquidity spillovers can be applied to any limit order book market.

With respect to fixed-income markets which, particularly in Europe, have received rather little attention until recently, we hope that this thesis serves as a starting point and blueprint for further exploration. The MiFID II regulation, in place since January 2018, is designated also to provide more data on European fixed income markets. Once reliable datasets on this regulatory basis will be available in the years to come, academics will find abundant opportunities for further research.

Bibliography

Abhyankar, A., D. Ghosh, E. Levin, and R. Limmack (1997). Bid-ask spreads, trading volume and volatility: Intra-day evidence from the London Stock Exchange. *Journal of Business Finance & Accounting 24*(3), 343–362.

Adrian, T., M. Fleming, D. Stackman, and E. Vogt (2015). Has U.S. Treasury market liquidity deteriorated? Available online at: `http://libertystreeteconomics.newyorkfed.org/2015/08/has-us-treasury-market-liquidity-deteriorated.html` (accessed 11 June 2016).

Alfonsi, A., F. Klöck, and A. Schied (2016). Multivariate transient price impact and matrix-valued positive definite functions. *Mathematics of Operations Research 41*(3), 914–934.

Alfonsi, A., A. Schied, and A. Slynko (2012). Order book resilience, price manipulation, and the positive portfolio problem. *SIAM Journal on Financial Mathematics 3*(1), 511–533.

Almgren, R. and N. Chriss (2001). Optimal execution of portfolio transactions. *Journal of Risk 3*, 5–40.

Almgren, R. F. (2003). Optimal execution with nonlinear impact functions and trading-enhanced risk. *Applied Mathematical Finance 10*(1), 1–18.

AMF (2015). Study of liquidity in French bond markets. *Report by Autorité des Marchés Financiers*.

Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets 5*(1), 31–56.

Amihud, Y. and H. Mendelson (1986). Asset pricing and the bid-ask spread. *Journal of Financial Economics 17*(2), 223–249.

Angel, J. J. (1997). Tick size, share prices, and stock splits. *The Journal of Finance 52*(2), 655–681.

Aquilina, M. and F. Suntheim (2017). Liquidity in the UK corporate bond market: evidence from trade data. *The Journal of Trading 12*(4), 67–80.

Arrata, W. and B. Nguyen (2017). Price impact of bond supply shocks: Evidence from the Eurosystem's asset purchase program. *Banque de France Working Paper No. 623*.

Asriyan, V., W. Fuchs, and B. Green (2017, July). Information spillovers in asset markets with correlated values. *American Economic Review 107*(7), 579–611.

Aussenegg, W., L. Goetz, and R. Jelic (2015). Common factors in the performance of European corporate bonds–evidence before and after the financial crisis. *European Financial Management 21*(2), 265–308.

Bacry, E., I. Mastromatteo, and J.-F. Muzy (2015). Hawkes processes in finance. *Market Microstructure and Liquidity 1*(01), 1550005.

Bandi, F. M., D. Pirino, and R. Reno (2017). EXcess Idle Time. *Econometrica 85*(6), 1793–1846.

Bao, J., M. O'Hara, and X. A. Zhou (2016). The Volcker Rule and Market-Making in Times of Stress. *Journal of Financial Economics, forthcoming*.

Bao, J., J. Pan, and J. Wang (2011). The Illiquidity of Corporate Bonds. *Journal of Finance 66*(3), 911–946.

Barclay, M. J., T. Hendershott, and K. Kotz (2006). Automation versus intermediation: Evidence from treasuries going off the run. *The Journal of Finance 61*(5), 2395–2414.

Beber, A., M. W. Brandt, and K. A. Kavajecz (2009). Flight-to-quality or flight-to-liquidity? evidence from the euro-area bond market. *Review of Financial Studies 22*(3), 925–957.

Benzaquen, M., I. Mastromatteo, Z. Eisler, and J.-P. Bouchaud (2017). Dissecting cross-impact on stock markets: An empirical analysis. *Journal of Statistical Mechanics: Theory and Experiment 2017*(2), 023406.

Bessembinder, H., S. E. Jacobsen, W. F. Maxwell, and K. Venkataraman (2016). Capital Commitment and Illiquidity in Corporate Bonds. *Journal of Finance, forthcoming*.

Bessembinder, H., W. F. Maxwell, and K. Venkataraman (2006). Market transparency, liquidity externalities, and institutional trading costs in corporate bonds. *Journal of Financial Economics 82*(2), 251–288.

Bessembinder, H. and K. Venkataraman (2004). Does an electronic stock exchange need an upstairs market? *Journal of Financial Economics 73*(1), 3–36.

Bhattacharya, A. (2016). Can transparency hurt investors in over-the-counter markets? *available at SSRN 2746910*.

Biais, B., F. Declerck, J. Dow, R. Portes, and E.-L. von Thadden (2006). European corporate bond markets: transparency, liquidity, efficiency. *CEPR Discussion paper*. available at `http://www.cepr.org/PRESS/TT_CorporateFULL.pdf`.

BIS (2016). Electronic trading in fixed income markets. *Bank for International Settlements Report*.

Bloomfield, R. and M. O'Hara (1999, January). Market transparency: Who wins and who loses? *The Review of Financial Studies 12*(1), 5–35.

Bollerslev, T. and I. Domowitz (1993). Trading patterns and prices in the interbank foreign exchange market. *The Journal of Finance 48*(4), 1421–1443.

Bollerslev, T., T. H. Law, and G. Tauchen (2008). Risk, jumps, and diversification. *Journal of Econometrics 144*(1), 234–256.

Bongaerts, D., R. Roll, D. Rösch, M. A. Van Dijk, and D. Yuferova (2015). The propagation of shocks across international equity markets: A microstructure perspective. *Available at SSRN 2475518*.

Boortz, C., S. Kremer, S. Jurkatis, and D. Nautz (2014). Information risk, market stress and institutional herding in financial markets: New evidence through the lens of a simulated model. *SFB 649 Discussion Paper*.

Booth, G. G., J.-C. Lin, T. Martikainen, and Y. Tse (2002). Trading and pricing in upstairs and downstairs stock markets. *Review of Financial Studies 15*(4), 1111–1135.

Bormetti, G., L. M. Calcagnile, M. Treccani, F. Corsi, S. Marmi, and F. Lillo (2015). Modelling systemic price cojumps with Hawkes factor models. *Quantitative Finance 15*(7), 1137–1156.

Bouchaud, J.-P., J. D. Farmer, and F. Lillo (2008). How markets slowly digest changes in supply and demand. In *Handbook of Financial Markets: Dynamics and Evolution*. Elsevier: Academic Press.

Bouchaud, J.-P., Y. Gefen, M. Potters, and M. Wyart (2004). Fluctuations and response in financial markets: the subtle nature of "random" price changes. *Quantitative Finance 4*(2), 176–190.

Bowsher, C. G. (2007). Modelling security market events in continuous time: Intensity based, multivariate point process models. *Journal of Econometrics 141*(2), 876–912.

Brockman, P., D. Y. Chung, and C. Pérignon (2009). Commonality in liquidity: A global perspective. *Journal of Financial and Quantitative Analysis 44*(04), 851–882.

Brunnermeier, M. K. and L. H. Pedersen (2009). Market liquidity and funding liquidity. *Review of Financial Studies 22*(6), 2201–2238.

Bundesbank (2017). The market for corporate bonds in the low-interest-rate environment. *Deutsche Bundesbank Monthly Report*, 17–32.

Busseti, E. and F. Lillo (2012). Calibration of optimal execution of financial transactions in the presence of transient market impact. *Journal of Statistical Mechanics: Theory and Experiments 2012*(09), P09010.

Carollo, A., G. Vaglica, F. Lillo, and R. N. Mantegna (2012). Trading activity and price impact in parallel markets: SETS vs. off-book market at the London Stock Exchange. *Quantitative Finance 12*(4), 517–530.

Castagnetti, C. and E. Rossi (2013). Euro corporate bond risk factors. *Journal of Applied Econometrics 28*(3), 372–391.

Cespa, G. and T. Foucault (2014). Illiquidity contagion and liquidity crashes. *Review of Financial Studies 27*(6), 1615–1660.

Chavez-Demoulin, V., A. C. Davison, and A. J. McNeil (2005). Estimating value-at-risk: a point process approach. *Quantitative Finance 5*(2), 227–234.

Chavez-Demoulin, V. and J. McGill (2012). High-frequency financial data modeling using Hawkes processes. *Journal of Banking & Finance 36*(12), 3415–3426.

Choi, J. and Y. Huh (2017). Customer liquidity provision: Implications for corporate bond transaction costs. *FEDS Working Paper No. 2017-116*.

Chordia, T., R. Roll, and A. Subrahmanyam (2000). Commonality in liquidity. *Journal of Financial Economics 56*(1), 3–28.

Chordia, T., A. Sarkar, and A. Subrahmanyam (2005). An empirical analysis of stock and bond market liquidity. *Review of Financial Studies 18*(1), 85–129.

Collin-Dufresne, P., B. Junge, and A. B. Trolle (2017). Market structure and transaction costs of index cdss. *available at SSRN 2786907*.

Coluzzi, C., S. Ginebri, and M. Turco (2008). Measuring and analyzing the liquidity of the Italian treasury security wholesale secondary market. *Working Paper, University of Molise*.

Crosignani, M., M. Faria-e Castro, and L. Fonseca (2015). The (unintended?) consequences of the largest liquidity injection ever. *FRB St. Louis Working Paper No. 2017-39*.

Curato, G., J. Gatheral, and F. Lillo (2016). Discrete homotopy analysis for optimal trading execution with nonlinear transient market impact. *Communications in Nonlinear Science and Numerical Simulation 39*, 332–342.

Curato, G., J. Gatheral, and F. Lillo (2017). Optimal execution with non-linear transient market impact. *Quantitative Finance 17*(1), 41–54.

Darbha, M. and A. Dufour (2013). Microstructure of the Euro-area government bond market. In H. K. Baker and H. Kiymaz (Eds.), *Market Microstructure in Emerging and Developed Markets*, Chapter 3, pp. 39–58. Wiley Online Library.

Dayri, K. and M. Rosenbaum (2015). Large tick assets: implicit spread and optimal tick size. *Market Microstructure and Liquidity 1*(01), 1550003.

de Roure, C., E. Moench, L. Pelizzon, and M. Schneider (2018). OTC discount. *in preparation*.

De Santis, R. A. and F. Holm-Hadulla (2017). Flow effects of central bank asset purchases on euro area sovereign bond yields: evidence from a natural experiment. *ECB Working Paper No. 2052*.

Di Maggio, M., A. Kermani, and Z. Song (2017). The value of trading relations in turbulent times. *Journal of Financial Economics 124*(2), 266–284.

Díaz, A. and E. Navarro (2002). Yield spread and term to maturity: default vs. liquidity. *European Financial Management 8*(4), 449–477.

Dick-Nielsen, J. (2009). Liquidity biases in trace. *The Journal of Fixed Income 19*(2), 43–55.

Dick-Nielsen, J. (2014). How to clean enhanced TRACE data. *available at SSRN 2337908*.

Dick-Nielsen, J., P. Feldhütter, and D. Lando (2012). Corporate bond liquidity before and after the onset of the subprime crisis. *Journal of Financial Economics 103*(3), 471–492.

Donier, J. and J. Bonart (2015). A million metaorder analysis of market impact on the bitcoin. *Market Microstructure and Liquidity 1*(02), 1550008.

Duffie, D. (2012a). *Dark markets: Asset pricing and information transmission in over-the-counter markets*. Princeton University Press.

Duffie, D. (2012b). Market Making Under the Proposed Volcker Rule. *Rock Center for Corporate Governance at Stanford University Working Paper No. 106*.

Duffie, D., P. Dworczak, and H. Zhu (2017). Benchmarks in search markets. *The Journal of Finance 72*(5), 1983–2044.

Duffie, D., N. Gârleanu, and L. H. Pedersen (2005). Over-the-counter markets. *Econometrica 73*(6), 1815–1847.

Duffie, D., N. Gârleanu, and L. H. Pedersen (2007). Valuation in over-the-counter markets. *The Review of Financial Studies 20*(6), 1865–1900.

Dufour, A. and M. Nguyen (2012). Permanent trading impacts and bond yields. *The European Journal of Finance 18*(9), 841–864.

Dufour, A., F. Skinner, et al. (2004). MTS time series: Market and data description for the European bond and repo database. Technical report, Henley Business School, Reading University.

Dunne, P. G., H. Hau, and M. J. Moore (2015). Dealer intermediation between markets. *Journal of the European Economic Association 13*(5), 770–804.

Easley, D., M. M. López de Prado, and M. O'Hara (2011). The microstructure of the "flash crash": flow toxicity, liquidity crashes, and the probability of informed trading. *The Journal of Portfolio Management 37*(2), 118–128.

Edwards, A. K., L. E. Harris, and M. S. Piwowar (2007, June). Corporate bond market transaction costs and transparency. *Journal of Finance 62*(3), 1421–1451.

Eisler, Z. and J.-P. Bouchaud (2016). Price impact without order book: A study of the OTC credit index market. *available at SSRN 2840166*.

Eisler, Z., J.-P. Bouchaud, and J. Kockelkoren (2012). The price impact of order book events: market orders, limit orders and cancellations. *Quantitative Finance 12*(9), 1395–1419.

Embrechts, P., T. Liniger, and L. Lin (2011). Multivariate hawkes processes: an application to financial data. *Journal of Applied Probability 48*(A), 367–378.

ESMA (2016). EU corporate bond market liquidity – recent evidence. *ESMA Report on Trends, Risks and Vulnerabilities 2*.

Feldhütter, P. (2011). The same bond at different prices: identifying search frictions and selling pressures. *The Review of Financial Studies 25*(4), 1155–1206.

Fermanian, J.-D., O. Guéant, and J. Pu (2016). The behavior of dealers and clients on the european corporate bond market: the case of multi-dealer-to-client platforms. *Market Microstructure and Liquidity*, 1750004.

Fleming, M. J. (1997). The round-the-clock market for US Treasury securities. *Economic Policy Review 3*(2).

Fleming, M. J. (2003). Measuring Treasury market liquidity. *Economic Policy Review 9*(3).

Fleming, M. J. and E. M. Remolona (1999). Price formation and liquidity in the US Treasury market: The response to public information. *The Journal of Finance 54*(5), 1901–1915.

Frank, J. and P. Garcia (2011). Bid-ask spreads, volume, and volatility: Evidence from livestock markets. *American Journal of Agricultural Economics 93*(1), 209–225.

Friederich, S. and R. Payne (2007). Dealer liquidity in an auction market: Evidence from the London Stock Exchange. *The Economic Journal 117*(522), 1168–1191.

Friewald, N., R. Jankowitsch, and M. G. Subrahmanyam (2012). Illiquidity or credit deterioration: A study of liquidity in the US corporate bond market during financial crises. *Journal of Financial Economics 105*(1), 18–36.

Friewald, N., R. Jankowitsch, and M. G. Subrahmanyam (2017). Transparency and liquidity in the structured product market. *The Review of Asset Pricing Studies 7*(2), 316–348.

Frühwirth, M., P. Schneider, and L. Sögner (2010). The risk microstructure of corporate bonds: a case study from the German corporate bond market. *European Financial Management 16*(4), 658–685.

Garleanu, N. B. and L. H. Pedersen (2007). Liquidity and risk management. *American Economic Review 97*(2), 193–197.

Gatheral, J. (2010). No-dynamic-arbitrage and market impact. *Quantitative Finance 10*(7), 749–759.

Gatheral, J. and A. Schied (2013a). Dynamical models of market impact and algorithms for order execution. In *Handbook of Systemic Risk*, pp. 579–599. Cambridge University Press.

Gatheral, J. and A. Schied (2013b). Dynamical models of market impact and algorithms for order execution. *HANDBOOK ON SYSTEMIC RISK, Jean-Pierre Fouque, Joseph A. Langsam, eds*, 579–599.

Gatheral, J., A. Schied, and A. Slynko (2011). Exponential resilience and decay of market impact. In *Econophysics of Order-driven Markets*, pp. 225–236. Springer.

Ghysels, E., J. Idier, S. Manganelli, and O. Vergote (2014). A high frequency assessment of the ECB Securities Markets Programme. *ECB Working Paper*.

Gilder, D., M. B. Shackleton, and S. J. Taylor (2014). Cojumps in stock prices: Empirical evidence. *Journal of Banking & Finance 40*, 443–459.

Glode, V. and C. Opp (2017). Over-the-counter vs. limit-order markets: The role of traders' expertise. *available at SSRN 2697281*.

Goldstein, M. A. and K. A. Kavajecz (2000). Eighths, sixteenths, and market depth: changes in tick size and liquidity provision on the NYSE. *Journal of Financial Economics 56*(1), 125–149.

Goyenko, R. Y., C. W. Holden, and C. A. Trzcinka (2009). Do liquidity measures measure liquidity? *Journal of Financial Economics 92*(2), 153–181.

Greenwich Associates (2014). Greenwich Associates European Fixed-Income Investors Study.

Grosse-Rueschkamp, B., S. Steffen, and D. Streitz (2017). Cutting out the middleman – the ECB as corporate bond investor. *Available at SSRN 2988158*.

Grossman, S. J. (1992). The informational role of upstairs and downstairs trading. *Journal of Business*, 509–528.

Gündüz, Y., G. Ottonello, L. Pelizzon, M. Schneider, and M. G. Subrahmanyam (2018). Lighting up the dark: liquidity in the German corporate bond market. *in preparation*.

Hameed, A., W. Kang, and S. Viswanathan (2010). Stock market declines and liquidity. *The Journal of Finance 65*(1), 257–293.

Harris, L. E. and M. S. Piwowar (2006). Secondary trading costs in the municipal bond market. *The Journal of Finance 61*(3), 1361–1397.

Hasbrouck, J. (2009). Trading costs and returns for US equities: Estimating effective costs from daily data. *The Journal of Finance 64*(3), 1445–1477.

Hasbrouck, J. and D. J. Seppi (2001). Common factors in prices, order flows, and liquidity. *Journal of Financial Economics 59*(3), 383–411.

Hawkes, A. G. (1971a). Point spectra of some mutually exciting point processes. *Journal of the Royal Statistical Society. Series B (Methodological) 33*(3), 438–443.

Hawkes, A. G. (1971b). Spectra of some self-exciting and mutually exciting point processes. *Biometrika 58*(1), 83–90.

Hendershott, T. and A. Madhavan (2015). Click or call? auction versus search in the over-the-counter market. *The Journal of Finance 70*(1), 419–447.

Holmstrom, B. (2015). Understanding the role of debt in the financial system. *BIS Working Paper No. 479*.

Hong, G. and A. Warga (2000). An empirical study of bond market transactions. *Financial Analysts Journal 56*(2), 32–46.

Houweling, P., A. Mentink, and T. Vorst (2005). Comparing possible proxies of corporate bond liquidity. *Journal of Banking & Finance 29*(6), 1331–1358.

Huang, J. and J. Wang (2009). Liquidity and market crashes. *Review of Financial Studies 22*(7), 2607–2643.

Huang, W., C.-A. Lehalle, and M. Rosenbaum (2016). How to predict the consequences of a tick value change? evidence from the Tokyo Stock Exchange pilot program. *Market Microstructure and Liquidity 2*(03n04), 1750001.

Huberman, G. and W. Stanzl (2004). Price manipulation and quasi-arbitrage. *Econometrica 72*(4), 1247–1275.

Ito, T. and Y. Hashimoto (2006). Intraday seasonality in activities of the foreign exchange markets: Evidence from the electronic broking system. *Journal of the Japanese and International Economies 20*(4), 637–664.

Jank, S., C. Roling, and E. Smajlbegovic (2016). Flying under the radar: The effects of short-sale disclosure rules on investor behavior and stock prices. *Bundesbank Discussion Paper No. 25/2016*.

Jankowitsch, R., F. Nagler, and M. G. Subrahmanyam (2014). The determinants of recovery rates in the US corporate bond market. *Journal of Financial Economics 114*(1), 155–177.

Jankowitsch, R., A. Nashikkar, and M. G. Subrahmanyam (2011). Price dispersion in OTC markets: A new measure of liquidity. *Journal of Banking & Finance 35*(2), 343–357.

Karolyi, G. A., K.-H. Lee, and M. A. Van Dijk (2012). Understanding commonality in liquidity around the world. *Journal of Financial Economics 105*(1), 82–112.

Keim, D. B. and A. Madhavan (1996). The upstairs market for large-block transactions: Analysis and measurement of price effects. *Review of Financial Studies 9*(1), 1–36.

Klein, C. and C. Stellner (2014). The systematic risk of corporate bonds: default risk, term risk, and index choice. *Financial Markets and Portfolio Management 28*(1), 29–61.

Koutmos, D. (2018). Bitcoin returns and transaction activity. *Economics Letters 167*, 81–85.

Kratz, P. and T. Schöneborn (2015). Portfolio liquidation in dark pools in continuous time. *Mathematical Finance 25*(3), 496–544.

Kremer, S. and D. Nautz (2013a). Causes and consequences of short-term institutional herding. *Journal of Banking & Finance 37*(5), 1676–1686.

Kremer, S. and D. Nautz (2013b). Short-term herding of institutional traders: New evidence from the German stock market. *European Financial Management 19*(4), 730–746.

Kyle, A. S. and W. Xiong (2001). Contagion as a wealth effect. *The Journal of Finance 56*(4), 1401–1440.

Lallouache, M. and F. Abergel (2014). Tick size reduction and price clustering in a FX order book. *Physica A: Statistical Mechanics and its Applications 416*, 488–498.

LeBaron, B. and R. Yamamoto (2007). Long-memory in an order-driven market. *Physica A: Statistical Mechanics and its Applications 383*(1), 85–89.

Lee, C. and M. J. Ready (1991). Inferring trade direction from intraday data. *The Journal of Finance 46*(2), 733–746.

Lee, S. S. and P. A. Mykland (2008). Jumps in financial markets: A new nonparametric test and jump dynamics. *Review of Financial Studies 21*(6), 2535–2563.

Lee, T. and C. Wang (2017). Why trade over-the-counter? when investors want price discrimination. *available at SRRN 3087647*.

Lesmond, D. A., J. P. Ogden, and C. A. Trzcinka (1999). A new estimate of transaction costs. *The Review of Financial Studies 12*(5), 1113–1141.

Lillo, F. and J. D. Farmer (2004). The long memory of the efficient market. *Studies in Nonlinear Dynamics & Econometrics 8*(3).

Lillo, F., J. D. Farmer, and R. N. Mantegna (2003). Master curve for price-impact function. *Nature 421*, 129–130.

Lillo, F., S. Mike, and J. D. Farmer (2005). Theory for long memory in supply and demand. *Physical Review E 71*(6), 066122.

Lin, H., J. Wang, and C. Wu (2011, March). Liquidity risk and expected corporate bond returns. *Journal of Financial Economics 99*(3), 628–650.

Linciano, N., F. Fancello, M. Gentile, and M. Modena (2014). The liquidity of dual-listed corporate bonds. Empirical evidence from Italian markets. *CONSOB Working Papers No. 79*.

Madhavan, A. and M. Cheng (1997). In search of liquidity: Block trades in the upstairs and downstairs markets. *Review of Financial Studies 10*(1), 175–203.

Madhavan, A. and G. Sofianos (1998). An empirical analysis of NYSE specialist trading. *Journal of Financial Economics 48*(2), 189–210.

Malamud, S. and M. Rostek (2017). Decentralized exchange. *American Economic Review 107*(11), 3320–62.

Malinova, K. and A. Park (2013). Liquidity, volume and price efficiency: The impact of order vs. quote driven trading. *Journal of Financial Markets 16*(1), 104–126.

Mancini, L., A. Ranaldo, and J. Wrampelmeyer (2013). Liquidity in the foreign exchange market: Measurement, commonality, and risk premiums. *The Journal of Finance 68*(5), 1805–1841.

Mastromatteo, I., M. Benzaquen, Z. Eisler, and J.-P. Bouchaud (2017, 07). Trading lightly: Cross-impact and optimal portfolio execution. *Risk 30*, 82–87.

McInish, T. H. and R. A. Wood (1990). An analysis of transactions data for the Toronto Stock Exchange: Return patterns and end-of-the-day effect. *Journal of Banking & Finance 14*(2-3), 441–458.

McInish, T. H. and R. A. Wood (1992). An analysis of intraday patterns in bid/ask spreads for NYSE stocks. *The Journal of Finance 47*(2), 753–764.

Menkveld, A. J., B. Z. Yueshen, and H. Zhu (2017). Shades of darkness: A pecking order of trading venues. *Journal of Financial Economics 124*(3), 503–534.

Naik, N., A. Neuberger, and S. Viswanathan (1999, July). Trade disclosure regulation in markets with negotiated trades. *The Review of Financial Studies 12*(4), 873–900.

Obizhaeva, A. A. and J. Wang (2013). Optimal trading strategy and supply/demand dynamics. *Journal of Financial Markets 16*(1), 1–32.

Ogata, Y. (1978). The asymptotic behaviour of maximum likelihood estimators for stationary point processes. *Annals of the Institute of Statistical Mathematics 30*(1), 243–261.

Ogata, Y. (1988). Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical Association 83*(401), 9–27.

O'Hara, M. (1995). *Market microstructure theory*, Volume 108. Blackwell Publishers Cambridge, MA.

Ozaki, T. (1979). Maximum likelihood estimation of Hawkes' self-exciting point processes. *Annals of the Institute of Statistical Mathematics 31*(1), 145–155.

Pagano, M. and A. Roell (1996, June). Transparency and liquidity: A comparison of auction and dealer markets with informed trading. *The Journal of Finance 51*(2), 579–611.

Pagano, M. and E.-L. Von Thadden (2004). The European bond markets under EMU. *Oxford Review of Economic Policy 20*(4), 531–554.

Pasquariello, P. and C. Vega (2013). Strategic cross-trading in the U.S. stock market. *Review of Finance 19*, 229–282.

Pelizzon, L., M. Subrahmanyam, D. Tomio, and J. Uno (2014). Limits to arbitrage in sovereign bonds price and liquidity discovery in high frequency quote driven markets. *Mimeo*.

Pelizzon, L., M. G. Subrahmanyam, D. Tomio, and J. Uno (2016). Sovereign credit risk, liquidity, and european central bank intervention: Deus ex machina? *Journal of Financial Economics 122*(1), 86–115.

Rambaldi, M., P. Pennesi, and F. Lillo (2015). Modeling foreign exchange market activity around macroeconomic news: Hawkes-process approach. *Physical Review E 91*(1), 012819.

Reboredo, J. C. (2012). The switch from continuous to call auction trading in response to a large intraday price movement. *Applied Economics 44*(8), 945–967.

Riggs, L., E. Onur, D. Reiffen, and H. Zhu (2017). Swap trading after Dodd-Frank: Evidence from index CDS. *Available at SSRN 3047284*.

Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *The Journal of Finance 39*(4), 1127–1139.

Schestag, R., P. Schuster, and M. Uhrig-Homburg (2016). Measuring liquidity in bond markets. *The Review of Financial Studies 29*(5), 1170–1219.

Schied, A., T. Schöneborn, and M. Tehranchi (2010). Optimal basket liquidation for CARA investors is deterministic. *Applied Mathematical Finance 17*(6), 471–489.

Schlepper, K., H. Hofer, R. Riordan, and A. Schrimpf (2017). Scarcity effects of QE: A transaction-level analysis in the bund market. *Bundesbank Discussion Paper No. 06/2017*.

Schneider, M. and F. Lillo (2018). Cross-impact and no-dynamic-arbitrage. *forthcoming in Quantitative Finance*.

Schneider, M., F. Lillo, and L. Pelizzon (2016). How has sovereign bond market liquidity changed?-an illiquidity spillover analysis. *SAFE Working Paper No. 151*.

Schneider, M., F. Lillo, and L. Pelizzon (2018). Modelling illiquidity spillovers with Hawkes processes: an application to the sovereign bond market. *Quantitative Finance 18*(2), 283–293.

Schöneborn, T. (2016). Adaptive basket liquidation. *Finance and Stochastics 20*(2), 455–493.

Schultz, P. (2000). Stock splits, tick size, and sponsorship. *The Journal of Finance 55*(1), 429–450.

Seppi, D. J. (1990). Equilibrium block trading and asymmetric information. *The Journal of Finance 45*(1), 73–94.

Smith, B. F., D. A. S. Turnbull, and R. W. White (2001). Upstairs market for principal and agency trades: Analysis of adverse information and price effects. *Journal of Finance*, 1723–1746.

Taranto, D. E., G. Bormetti, J.-P. Bouchaud, F. Lillo, and B. Toth (2016). Linear models for the impact of order flow on prices I. Propagators: Transient vs. history dependent impact. *arXiv:1602.02735*.

Toke, M. (2011). An introduction to Hawkes processes with applications to finance. Course slides, available online at `http://lamp.ecp.fr/MAS/fiQuant/ioane_files/HawkesCourseSlides.pdf` (accessed 28 September 2015).

Tóth, B., Y. Lemperiere, C. Deremble, J. De Lataillade, J. Kockelkoren, and J.-P. Bouchaud (2011). Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X 1*(2), 021006.

Toth, B., I. Palit, F. Lillo, and J. D. Farmer (2015). Why is equity order flow so persistent? *Journal of Economic Dynamics and Control 51*, 218–239.

Tsoukalas, G., J. Wang, and K. Giesecke (2017). Dynamic portfolio execution. *Management Science, forthcoming*.

Upper, C. and T. Werner (2002). Tail wags dog? time-varying information shares in the bund market. *Bundesbank Series 1 Discussion Paper No. 2002,24*.

Utz, S., M. Weber, and M. Wimmer (2016). German mittelstand bonds: yield spreads and liquidity. *Journal of Business Economics 86*(1-2), 103–129.

Valseth, S. (2015). Informed interdealer trading in hybrid bond markets.

Van Landschoot, A. (2008). Determinants of yield spread dynamics: Euro versus us dollar corporate bonds. *Journal of Banking & Finance 32*(12), 2597–2605.

Vere-Jones, D. (1970). Stochastic models for earthquake occurrence. *Journal of the Royal Statistical Society. Series B (Methodological)*, 1–62.

Vere-Jones, D. (1995). Forecasting earthquakes and earthquake risk. *International Journal of Forecasting 11*(4), 503–538.

Vogel, S. (2017). When to introduce electronic trading platforms in over-the-counter markets? *available at SSRN 2895222*.

Wang, S. and T. Guhr (2016). Microscopic understanding of cross-responses between stocks: a two-component price impact model. *arXiv:1609.02395*.

Wang, S., R. Schäfer, and T. Guhr (2016a). Average cross-responses in correlated financial markets. *The European Physical Journal B 89*(9), 207.

Wang, S., R. Schäfer, and T. Guhr (2016b). Cross-response in correlated financial markets: individual stocks. *The European Physical Journal B 89*(4), 1–16.