Estimating a Model of Decentralized Trade with Asymmetric Information*

Tristan Hennig†

October 2020

Please click here for the latest version

Abstract

This paper demonstrates how to quantify search and information frictions in OTC markets. I use transaction data for the U.S. corporate bond market to estimate a model featuring both a search and information friction via the simulated method of moments. The data support the notion that trades are informative and uncertainty about the fundamental value diminishes as time passes. Trading in bonds with different time to maturity exhibit different characteristics which is reflected by differences in the estimated parameters. As a result, changes in the trading frictions through regulatory interventions or technological progress would have heterogeneous effects on the trading in different bonds. Increasing the probability of finding a trading partner by 20% decreases spreads by 18 percent on average (25 percent for bonds with long time-to-maturity), increases welfare by 21.8 percent on average (25 percent), and decreases price volatility by 40 percent on average (51 percent). However, these improvements come at the cost of a substantial slow-down in price discovery. The speed of convergence of the price to the true value decreases by up to 26 percent. This result serves as a caution regarding the impact of recent regulation (such as MiFID II) that mandates some trading to be on-exchange rather than OTC.

Keywords: Trade frictions, OTC markets, asymmetric information

JEL codes: D53, D82, D83, G12, G14, G18

*I am indebted to Hamid Sabourian and Robert Evans for continuous guidance and advice. I would also like to thank Ben Lester, Julius Vainora, and Venky Venkateswaran as well as a seminar audience at the University of Cambridge for helpful comments. Funding from the Cambridge Trust and Trinity Hall is gratefully acknowledged.

†Faculty of Economics, University of Cambridge, Sidgwick Avenue, Cambridge, CB3 9DD, United Kingdom. Email: th486@cam.ac.uk, Website: www.tristanhennig.com

1 Introduction

Many financial products trade in so-called over-the-counter (OTC) markets. Examples are corporate bonds, municipal bonds, derivatives and mortgage backed securities. These markets are large and important: The gross market value of OTC derivatives exceeded $12 trillion in 2019[1]. Trade on OTC markets occurs directly between two counterparties, without the involvement of a central exchange. As a result, a transaction occurring on an OTC markets may not immediately be visible to all market participants, but only to those who formed part of that transaction. Moreover, those wishing to trade must search for a trading partner who is willing to take the opposite side of the trade.

The markets in the above-mentioned products have undergone substantial changes over recent years. The introduction of mandatory reporting has meant that market participants are able to observe aggregate trading in many over-the-counter markets, albeit sometimes with a small delay. One instance of this is the Financial Industry Regulatory Authority's TRACE system. More recently, regulators have also begun to mandate that at least some trades be conducted on an exchange. Regulations to this effect are, for example, included in the European Union's Markets in Financial Instruments Directive II (MiFID II). At the same time, technological advances such as electronic quoting and trading systems have reduced the cost of finding a trading partner.

Another characteristic of assets traded in OTC markets is that they usually involve non-negligible fundamental risk. For instance, a company may default on the bonds it has issued to investors. This fundamental risk affects trading behaviour. When trading frictions are low, investors might be happy to buy a bond they know will default eventually. Holding such a bond allows them to collect regular coupon payments or satisfy other liquidity needs. When trading frictions are high instead, investors would be reluctant to hold such a bond as they don't know when they can next get rid of it. As a result, a decrease in trading frictions would make investors' trading behaviour less dependent on the fundamental value and more dependent on their liquidity needs. This reduces the informativeness of trades and slows down

[1]Bank for International Settlement Statistical Release: OTC derivatives statistics at end-June 2019, available at https://www.bis.org/publ/otc_hy1911.pdf

learning.

In this paper, I employ the model in Lester, Shourideh, Venkateswaran, and Zetlin-Jones (2018) to quantify the effects of a reduction in search frictions. Contrary to existing work, this model does not only include the search friction characteristic of OTC markets (in the spirit of Duffie, Gârleanu, and Pedersen (2005)), but also asymmetric information in the spirit of Glosten and Milgrom (1985). It thus manages to capture the notion that the way investors' optimal trading strategy depends on the fundamental value of the asset may change with the severity of the search friction.

I use FINRA's TRACE dataset which contains all transactions in US corporate bonds. My sample period is the last 4 years available, October 2015 to September 2019. After cleaning the data, I split the transactions by time to maturity of the bonds. I then estimate the model using the simulated method of moments. As moments, I use the bid-ask spread, the fraction of the asset traded per week, the trade imbalance, the variance of the price, the variance of the spreads and the variance of the price relative to the trading volume.

The results indicate that the drivers of the trading behaviour – aggregate liquidity shocks, individual liquidity shocks, and superior information about the fundamental value all contribute differently depending on the bond grade and time to maturity. Consistent with the notion that learning is a major driver of trading in bonds, prices vary most relative to trading volume in bonds with high time to maturity and least in those closest to maturing. Dealers are least willing to take bonds with long time-to-maturity into their inventory.

The structural approach then allows me to conduct a counterfactual analysis and quantify the effects of a reduction in trading frictions. I find that a decrease in trading frictions of 20% decreases spreads by 18 percent on average. However, it also slows down learning. By one measure, the chance that the price of a bond converges to within 5% of its true value within 1000 periods drops by 26%. The effects are biggest for bonds with long time-to-maturity and smallest for bonds closest to maturing. In the appendix, I also report results for a complete removal of search frictions where the direction of the results is the same, but the magnitude is larger.

The rest of this paper is structured as follows. Section 2 describes the data together with the features motivating the analysis in this paper. Section 3 contains a very brief summary of the Lester, Shourideh, Venkateswaran, and Zetlin-Jones (2018) model. This is followed by a description of the estimation methodology in Section 4. Section 5 presents and discusses the results which allow the counterfactual analysis in Section 6. Section 7 concludes.

2 Related literature

The theoretical literature on decentralized markets using random search starts with Duffie, Gârleanu, and Pedersen (2005) and now consists of a vast body of work. As the contribution of this paper is empirical, the reader is referred to Weill (2020) for a survey. Within this literature, the model presented by Lester, Shourideh, Venkateswaran, and Zetlin-Jones (2018) (henceforth LSVZ) is unique in the sense that it combines a model of trade based on random search with a Glosten and Milgrom (1985)-style pricing mechanism. It therefore offers a unified framework to study the corporate bond market under both a trading friction and asymmetric information. The present paper takes this framework to the data.

This paper relates to the growing empirical literature studying over-the-counter markets. Brancaccio, Li, and Schürhoff (2017) look at how dealers acquire information through experimentation in municipal bond markets. Li and Schürhoff (2019) and Hagströmer and Menkveld (2019) examine the network structure of OTC markets. Edwards, Harris, and Piwowar (2007) and Bessembinder, Maxwell, and Venkataraman (2006) study transaction costs and the effect of mandatory reporting requirements in the corporate bond market. This paper makes use of the data collected through mandatory reporting, but abstracts from experimentation motives and network effects.

This paper aims to quantify the effects of a reduction in search frictions in a decentralized market which may come about through the rise of electronic trading. O'Hara and Zhou (2019) examine the implications of electronic trading on OTC markets whereas Vogel (2019) derives conditions which ensure that the presence of electronic trading is beneficial. Another potential reason for reduced search cost is regulation which mandates to be moved from over-the-counter markets to a centralised venue where the search problem disappears. Papers comparing the two market structures include Biais (1993), Glode and Opp (2018), Dugast, Üslü, and Weill (2019). This paper contributes to the literature by focusing on asymmetric information.

In terms of methodology, I use a structural estimation to uncover the fundamental parameters of the Lester, Shourideh, Venkateswaran, and Zetlin-Jones (2018) model from transaction data. Compared to a reduced form approach, this allows me to estimate parameters that are not readily observable and then use the estimates to perform a counterfactual analysis. Due to these possibilities, structural estimations have been used to study many different markets with search frictions such as the housing market (Carrillo (2012)), labour market (Eckstein and Wolpin (1990)), and commercial aircraft market (Gavazza (2016)). Focusing on financial markets, Feldhütter (2012) demonstrates how to identify periods of high selling pressure in the corporate bond market. Liu (2020) performs a structural estimation of the fully decentralized model in Hugonnier, Lester, and Weill (2020), but with endogenous dealer search intensities. To the best of my knowledge, the present paper is the only paper estimating a model that takes not just the search friction, but also the information friction into account.

3 Data

I combine data from several sources. First, I use the Mergent FISD database to obtain characteristic information (such as issue date, maturity, amount outstanding, etc.) on all corporate bonds issued in the U.S. As is standard in the literature, I focus only on bonds that do not have special characteristics. In particular, for a bond to be included in my sample, I require that the bond pay a fixed coupon, that it be non-convertible, non-exchangeable, non-putable, not private-placed (Rule 144a), not asset-backed and not perpetual. I set the sample period to be the most recent 4 year period available which is October 2015 to September 2019. I also exclude bonds which were traded for less than 3 months inside this window. This returns a total of 46,477 bonds.

My main data source is the enhanced version of the "Trade Reporting And Compliance Engine" (TRACE), maintained by the Financial Industry Regulatory Authority (FINRA). This database contains the universe of all transactions in U.S. corporate bonds. One advantage of the enhanced TRACE dataset is that volume information is not top-coded. In the standard version of the dataset, the volume of transactions is capped at USD 5 million for investment grade and USD 1 million for high yield bonds. I merge the list of bonds obtained from the Mergent FISD database with the TRACE Masterfile which contains information on the "grade" of each bond. This allows me to classify each bond as either "high yield" ($N = 6,755$) or "investement grade" ($N = 39,722$).

Finally, I query the enhanced TRACE database to obtain all transactions in these bonds in the sample timeframe. Broker-dealers have to report every transaction to the TRACE within 15 minutes of the trade being executed. However, entries in the TRACE database cannot be amended or modified once they have been entered into the system. If there is a correction or cancellation of a trade, a separate report has to be filed, followed by another report with the correct transaction. Dick-Nielsen (2014) describes how to clean the data so as to avoid double or triple counting of some transactions. After applying this cleaning procedure, my sample has roughly 46.3 million transactions left. Of those, 33 million are in investment grade and 13.3 million in high yield bonds.

Tables 1 and 2 contain summary statistics split by grade and maturity. It reveals that most bonds are traded quite infrequently with the median number of trades per week between 3 and 10. It also reveals heavily skewed distributions for the number of trades, trade size and to a smaller extent the amount outstanding. That is, most weeks see few trades, but some see many; most trades are "small", but some are very large. These findings are in line with existing studies of the corporate bond market, some of which study different sample periods of the same dataset.

To estimate the model in the next section, I also require an estimate of the bid-ask spreads for each bond. As the TRACE data only lists transactions and their prices, I use the literature standard "Imputed Roundtrip" (IRT) method developed by Feldhütter (2012) to infer the spread. An IRT consists of at least two transactions in the same security for the same par value amount that occur within 15 minutes of each other. These will be a customer selling to a dealer, followed by possibly several inter-dealer trades, and concluded by the last dealer selling to a customer. The inferred spread is the difference between the highest and lowest price in the sequence of transactions.

High yield	Time to maturity < 1 year	1-3 years	3-10 years	> 10 years
Mean amount outstanding	642m	572m	498m	364m
Median amount outstanding	500m	468m	400m	200m
Mean trade size	552k	450k	531k	548k
Median trade size	25k	25k	27k	25k
Mean number of trades per week	17	15	19	13
Median number of trades per week	9	8	10	5

Table 1: Summary statistics on the trading activity in high yield bonds

Investment grade	Time to maturity < 1 year	1-3 years	3-10 years	> 10 years
Mean amount outstanding	660m	651m	589m	514m
Median amount outstanding	500m	500m	400m	350m
Mean trade size	615k	421k	428k	789k
Median trade size	30k	25k	25k	50k
Mean no. of trades per week	16	16	16	7
Median no. of trades per week	8	8	7	3

Table 2: Summary statistics on the trading activity in investment grade bonds

By trade size	$< 25k$	$25k - 100k$	$100k - 500k$	$> 500k$
HY	47	51	28	13
IG	53	53	26	11

By Maturity	$<$ 1 year	1-3 years	3-10 years	$>$ 10 years
HY	15	27	48	73
IG	13	23	46	77

Over time	15Q4 - 16Q3	16Q4 - 17Q3	17Q4 - 18Q3	18Q4 - 19Q3
HY	52	44	39	36
IG	57	49	44	38

Table 3: "Imputed Roundtrip" spreads for the corporate bonds in my sample. Values are in USD cents.

Table 3 contains summary statistics on the IRT-based spreads in my sample. Again, we can observe that some of the well established stylised facts for the corporate bond market also hold in this sample. Spreads for investment grade bonds are overall slightly lower than spreads for high yield bonds. Spreads are largest for small trade sizes and fall in the trade size. In contrast, they are monotonically increasing in the time to maturity. The fact that spreads in bonds with different maturities differ substantially indicates that the effect of moving trading from OTC to an exchange will likely also have differential effects on these bonds.

So far, the literature on trade in corporate bond markets has paid little attention to the existence of asymmetric information and the effects of its presence. Existing models have instead focused on dealers' market power arising due to the search friction. An investor who wishes to trade a large sum is likely a well-connected and sophisticated investor such as a hedge fund rather than a retail investor. Sophisticated investors have better outside options. If the dealer they have matched with doesn't offer competitive prices, the investor will find them easy to refuse due to their high ability to find another trading partner. Put another way, dealers have market power over investors due to the search friction in the market. This power (and therefore the ability to set spreads) is highest vis-a-vis unsophisticated investors who trade smaller sums and have higher costs to finding an alternative trading partner. Overall, this gives rise to the pattern present in Table 3 – transaction spreads decrease in trade size. This stands in stark contrast to centralized limit order book markets where large trades would incur higher spreads than small trades.

In this paper, I abstract from the trade size pattern induced by market power and focus instead on the pattern with respect to time to maturity. Bonds with low maturities contain a relatively smaller amount of uncertainty regarding their fundamental value than bonds with more time to maturity. If dealers are concerned about being adversely selected, they should therefore set larger spreads on bonds with long times to maturity which is exactly what we observe in the data. The model reflects this by making investors choose only between trading or not trading and by employing a price mechanism as in Glosten and Milgrom (1985).

4 The corporate bond market

This section (i) presents evidence of search and information frictions in the corporate bond market, (ii) outlines how the corporate bond market has changed in recent years, and (iii) argues that the analysis in this paper is relevant for future policy.

Not just one, but several recent developments in corporate bond markets have meant that search costs have been reduced. The first development is the rise of electronic trading. The corporate bond market is traditionally voice-operated and has severely lagged behind other markets (most notably the equity market) in terms of the percentage of trade conducted electronically. Electronic trading systems make it easier for investors to survey the market despite the absence of a central exchange. For instance, through so-called "Request-for-Quote" (RFQ) systems investors can send a trade request to multiple dealers at the same time. Dealers can respond with a quote and the investor can then pick the most attractive option.

O'Hara and Zhou (2019) examine the rise of electronic trading in the corporate bond market in detail. Among other things, they find that dealers who conduct more electronic trading than competitors offer lower prices in voice-based trades as well. They argue that, as electronic trading allows dealers to better find customers, dealers need to rely less on the inter-dealer market to offload their positions. This reduces cost, which the dealer can then use to make their quotes more competitive.

Given its apparent benefits, and the much higher market share of electronic trading in other markets, O'Hara and Zhou (2019) also investigate why electronic trading is not even more prominent. One reason appears to be the market structure. Many bonds are illiquid and trade infrequently. Contrary to equities, large trades are not typically split up into many small trades. Dealers want to retain control of these trades, particularly for high-yield bonds, where trades are more likely to be informationally sensitive. Nonetheless, electronic trading now takes up a substantial share of the market and continues to grow. According to a recent report by SIFMA[2], electronic trading made up 30% of all trading volume in investment grade bonds in 2019, up from 19% the year before. This has brought down search cost in the corporate bond market, intensified dealer competition, and improved access to relevant information for all market participants.

The second development is the introduction of regulation, partly in response to the global financial crisis of 2008/2009. In Europe, the European Union's Market in Financial Instruments Directive II (MiFID II) has expanded pre-trade and post-trade transparency requirements to non-equity financial instruments such as corporate bonds. Pre-trade transparency refers to the availability of information (most notably, quotes) that market participants have access to before they engage in the search for a trading partner. Post-trade transparency requires the reporting of trades (i.e. the transaction price, volume, and time) that have taken place over the counter to the regulator or a private company approved for collection and dissemination of trade data. The regulation now in force in both Europe and the U.S. permits the model assumption that market participants can observe anonymised, aggregate trading activity despite the decentralized nature of the market. Lastly, regulation like MiFID II and its global counterparts contain rules that require some trades to be conducted on-exchange which further reduces search cost.

[2] SIFMA Insights Electronic Trading Market Structure Primer, October 2019, available at https://www.sifma.org/resources/research/electronic-trading-market-structure-primer/

5 Model

This section contains a very condensed version of the model in Lester, Shourideh, Venkateswaran, and Zetlin-Jones (2018). The only modification I make is that I treat all meetings between dealers and investors as "competitive" (to use the terminology in their paper) and do away with "monopoly" meetings[3]. Readers familiar with the model can skip this section.

The environment. There are two states of the world, $j \in \{h, l\}$. There is a single risky asset with fundamental value v_j, $v_h > v_l$. Time t is discrete and infinite. However, at the start of every period the game terminates with chance δ.

The agents. There are two types of agents in the model, a mass N of traders (investors), and a mass 1 of dealers. All agents are risk-neutral, live forever, and do not discount the future. Dealers can take unrestricted positions in the asset, but traders are only allowed to hold at most one unit of it at any time (that is, they are either "owners" or "non-owners").

Payoffs. When the game terminates, an agent holding the asset receives payoff v_j. Additionally, for every period she holds the asset, trader i receives a flow payoff of $\omega_t + \epsilon_{it}$ where $\omega_t \overset{iid}{\sim} \mathcal{N}(0, \sigma_\omega)$ is an aggregate liquidity shock and $\epsilon_{it} \overset{iid}{\sim} \mathcal{N}(0, \sigma_\epsilon)$ is an idiosyncratic liquidity shock. Denote the cdfs of these distributions by F and G respectively. Dealers do not receive flow payoffs.

Trading. Every period, a trader meets a dealer with probability π. The dealer then quotes the trader two prices. The "Ask" price A_t at which the dealer is willing to sell to the trader and the "Bid" price B_t at which the dealers is willing to buy from the trader. After observing the prices, the trader decides whether to buy a unit of the asset, sell a unit, or walk away. Due to the restriction on traders' holdings, a trader can only sell if he currently owns the asset and buy if he does not currently own the asset.

Information. Traders perfectly know the state of the world whereas dealers do not. Dealers also do not observe traders' flow payoffs. However, dealers can observe traders' aggregate behaviour and use this information to learn about the state. They have a common prior $Pr(j = h) = \mu_0$ and their beliefs do not disperse over time. That

[3]The model presented here therefore obtains by setting $\alpha_c = 1$ in the model presented in Lester, Shourideh, Venkateswaran, and Zetlin-Jones (2018)

is, because observing aggregate behaviour is a superset of observing an individual traders' behaviour, all dealers have the same belief at the start of every period.

Traders' optimal behaviour. If a trader does not meet a dealer, she has no choice but to hold her position until the next period. If she does meet a trader, she observes the bid and ask prices and optimally decides whether to trade or not. Formally, let $W^q_{j,t}$ be the maximal value of holding $q \in \{0,1\}$ at the beginning of period t when the true state is j. Then,

$$W^0_{j,t} = (1-\delta)\cdot 0 + \delta \mathbb{E}_{\omega,\epsilon}\left[\pi \max(\omega_t + \epsilon_{i,t} + W^1_{j,t+1} - A_t, W^0_{j,t+1}) + (1-\pi) W^0_{j,t+1}\right]$$

and

$$W^1_{j,t} = (1-\delta)\cdot v_j + \delta \mathbb{E}_{\omega,\epsilon}\left[\pi \max(-\omega_t - \epsilon_{i,t} + B_t + W^0_{j,t+1}, W^1_{j,t+1}) + (1-\pi) W^1_{j,t+1}\right].$$

Subject to meeting a dealer, a non-owner will decide to buy iff $\epsilon_{i,t} > \bar{\epsilon}_{j,t}$ and an owner will decide to sell iff $\epsilon_{i,t} < \bar{\epsilon}_{j,t}$ where the thresholds are found by solving

$$\omega_t + \bar{\epsilon}_{i,t} + W^1_{j,t+1} - A_t = W^0_{j,t+1}$$

$$-\omega_t - \underline{\epsilon}_{i,t} + B_t + W^0_{j,t+1} = W^1_{j,t+1}.$$

This is equivalent to

$$\bar{\epsilon}_{i,t} = A_t - \omega_t - W^1_{j,t+1} + W^0_{j,t+1} = A_t - \omega_t - R_{j,t+1}$$

$$\underline{\epsilon}_{i,t} = B_t - \omega_t - W^1_{j,t+1} + W^0_{j,t+1} = B_t - \omega_t - R_{j,t+1}$$

where we have defined the *reservation value* $R_{j,t} = W^1_{j,t} - W^0_{j,t}$ so that

$$R_{j,t} = (1-\delta)v_j + \delta \mathbb{E}_{\omega,\epsilon}(R_{j,t+1}) + \delta\pi\Omega_{j,t} \tag{1}$$

with

$$\Omega_{j,t} = \mathbb{E}_{\omega,\epsilon}\left(\max(-\omega_t - \epsilon_{i,t} - R_{j,t+1} + B_t, 0) - \max(\omega_t + \epsilon_{i,t} + R_{j,t+1} - A_t, 0)\right).$$

Dealers' pricing. As in Glosten and Milgrom (1985), dealers are assumed to face unmodelled price competition so that all their profits are competed away. This drives the price to equal the expectation of the value of the asset conditional on all available information at that point. Formally, the ask price A_t and bid price B_t must satisfy

$$A_t = \frac{\mathbb{E}_{j,\omega}\left(v_j(1 - G(\bar{\epsilon}_{j,t}(\mu,\omega))\right)}{\mathbb{E}_{j,\omega}\left((1 - G(\bar{\epsilon}_{j,t}(\mu,\omega))\right)} \tag{2}$$

$$B_t = \frac{\mathbb{E}_{j,\omega}\left(v_j G(\underline{\epsilon}_{j,t}(\mu,\omega)\right)}{\mathbb{E}_{j,\omega}\left(G(\underline{\epsilon}_{j,t}(\mu,\omega)\right)}. \tag{3}$$

Note that the thresholds $\bar{\epsilon}_{j,t}, \underline{\epsilon}_{j,t}$ depend on the price A_t and B_t respectively. A spread between bid and ask arises due to the adverse selection problem facing dealers. A trader wishing to buy from a dealer may do so due to having received a large aggregate or idiosyncratic shock or because he knows that the true state is $j = h$. The dealer loses out to information-based trades, but recoups those losses from liquidity-induced trades. The spread is such that these gains and losses cancel out in expectation.

Demographics. Tomorrow's owners are today's owners who did not meet a dealer or met a dealer, but decided not to sell upon observing the prices, plus today's non-owners who met a dealer and decided to buy. Let $N^q_{j,t}$ be the mass of traders holding $q \in \{0, 1\}$ units at time t when the true state is j.

$$N^1_{j,t+1}(\mu,\omega) = N^1_t\left(1 - \pi + \pi(1 - G(\underline{\epsilon}_{j,t}(\mu,\omega))\right) + N^0_t \pi\left(1 - G(\bar{\epsilon}_{j,t}(\mu,\omega))\right) \tag{4}$$

Similarly, tomorrow's non-owners are today's non-owners who did not meet a dealer or met a dealer, but decided not to buy upon observing the prices, plus today's owners who met a dealer and decided to sell.

$$N^0_{j,t+1}(\mu,\omega) = N^1_t \pi G(\underline{\epsilon}_{j,t}(\mu,\omega)) + N^0_t\left(1 - \pi + \pi G(\bar{\epsilon}_{j,t}(\mu,\omega))\right) \tag{5}$$

I assume that the initial population is split equally between owners and non-owners and that this is common knowledge. Note that the evolution depends on the true state of the world (when $j = h$ more traders will buy, all else equal). Dealers know the exact distribution of owners and non-owners in all periods. However, as there are multiple values of the aggregate shock consistent with the observed evolution, dealers cannot perfectly infer the true state. Instead, they use it to learn as described above.

Learning. Dealers learn by observing aggregate trading activity, which is equivalent to observing the thresholds. Furthermore, all thresholds contain the same information as dealers know which prices are offered in equilibrium. To be precise, observing the thresholds is equivalent to observing

$$S_t = R_{t+1} + \omega_t$$

where $R_{t+1} = R_{j,t+1}$ iff the true state is j. At time t, dealers receive a noisy observation of traders' reservation values at $t+1$. However, traders' reservation values at $t+1$ naturally depend on dealers' beliefs at $t+1$, μ_{t+1}. To disentangle this, consider first how dealers update upon observing the signal S_t. As there are only two states of the world, there are only two values for the aggregate shock that are consistent with observing S_t. Denote these by $\omega_h = S_t - R_{h,t+1}(\mu_{t+1})$ and $\omega_l = S_t - R_{l,t+1}(\mu_{t+1})$. Dealers now update using Bayes' rule

$$\mu_{t+1} = \frac{\mu_t f(\omega_h)}{\mu_t f(\omega_h) + (1-\mu_t) f(\omega_l)} \tag{6}$$

$$= \frac{\mu_t f(S_t - R_{h,t+1}(\mu_{t+1}))}{\mu_t f(S_t - R_{h,t+1}(\mu_{t+1})) + (1-\mu_t) f(S_t - R_{l,t+1}(\mu_{t+1}))} \tag{7}$$

which is clearly a fixed-point problem in μ_{t+1}. Let the solution be denoted by $\mu_{t+1}(\mu_t, S_t)$.

Lastly, the traders take into account the way dealers will update. As they know the true state and the aggregate shock, they can therefore perfectly forecast dealers' beliefs at $t+1$, μ_{t+1}. This forecast $\tilde{\mu}_{j,t+1}(\mu_t, \omega_t)$ is the solution to

$$\mu_{t+1} = \frac{\mu_t}{\mu_t + (1-\mu_t)\frac{f(\omega_t + R_{j,t+1}(\mu_{t+1}) - R_{l,t+1}(\mu_{t+1}))}{f(\omega_t + R_{j,t+1}(\mu_{t+1}) - R_{h,t+1}(\mu_{t+1}))}} \tag{8}$$

As individual traders' behaviour depends on both the aggregate and idiosyncratic shock, dealers can always at most learn $R_{t+1} + \omega_t + \epsilon_{i,t}$ from the interaction with a single trader. However, by observing aggregate behaviour they learn $R_{t+1} + \omega_t$. As a result, dealers will always have the same beliefs at the beginning of each period.

Equilibrium. A Markov equilibrium is described by

1. Traders' optimal decisions are completely described by the thresholds

$$\underline{\epsilon}_j(\mu, \omega) = B(\mu) - \omega - R_j(\tilde{\mu}(\mu, \omega))$$

$$\bar{\epsilon}_j(\mu, \omega) = A(\mu) - \omega - R_j(\tilde{\mu}(\mu, \omega))$$

where $R_j(\mu)$ is given by equation (1).

2. Dealers post zero-profit prices. That is, bid and ask prices satisfy equations (2) and (3).

3. Given j, μ, ω, the population of owners and non-owners evolves according to equations (4) and (5).

4. For all $S \in \mathbb{R}$, dealers' beliefs evolve according to a function $\mu^+(\mu, S)$ that solves equation (6)

5. Traders forecast dealers' beliefs using a function $\tilde{\mu}(\mu, \omega)$ that solves equation (8). Furthermore, their forecast is correct. That is,

$$\tilde{\mu}(\mu, \omega) = \mu^+ (\mu, R_j(\tilde{\mu}(\mu, \omega) + \omega))$$

Welfare. Taking welfare to be the total payoff of dealers and investors, we can simplify by noting that dealers' and investors' payments between each other cancel out. Investors' payoffs depend then only depend on whether potential gains from trade are realised or not. Let $Q_{j,t}$ be investors' period t payoff net of payments to dealers when the true state is j. This payoff can be split up into owners' payoff $Q^1_{j,t}$ (of which there are $N^1_{j,t}$) and non-owners' payoff $Q^0_{j,t}$ (of which there are $N^0_{j,t}$). In total, we have

$$Q_{jt} = N^1_{j,t} Q^1_j(\mu, \omega) + N^t_0 Q^0_j(\mu, \omega),$$

and

$$
\begin{aligned}
Q^1_{j,t}(\mu,\omega) &= (1-\pi)\left(\int_{-\infty}^{+\infty}(\omega+\epsilon)dG(\epsilon)\right) + \pi\left(\int_{\underline{\epsilon}_{jt}}^{+\infty}(\omega+\epsilon)dG(\epsilon)\right) \\
&= (1-\pi)\omega + \pi\left(\omega(1-G(\underline{\epsilon}_j)) + \int_{\underline{\epsilon}_{jt}}^{+\infty}\epsilon dG(\epsilon)\right), \\
Q^0_{j,t}(\mu,\omega) &= \pi\int_{\bar{\epsilon}_{jt}}^{+\infty}(\omega+\epsilon)dG(\epsilon) \\
&= \pi\left(\omega(1-G(\bar{\epsilon}_j)) + \int_{\bar{\epsilon}_{jt}}^{+\infty}\epsilon dG(\epsilon)\right).
\end{aligned}
$$

Owners of the security may not find a dealer in period t in which case their payoff is simply the sum of aggregate and idiosyncratic shock. If they do find a dealer, they sell if their realisation of ϵ is low enough, and hold otherwise. Similar reasoning applies to the payoff of non-owners. As with spreads, there are two effects of a change in π. The direct effect is that the number of meetings per period increases as more investors find a dealer, increasing the possibility to realise potential gains from trade. However, the dynamic effect is that learning slows down which increases spreads in the long run. Higher spreads mean fewer gains from trade are realised. Which one of these two effects dominates is unclear a priori.

6 Methodology

I use the Simulated Method of Moments (SMM), developed by McFadden (1989) and Pakes and Pollard (1989), to estimate the parameter vector $\beta = \{\pi, \sigma_\omega, \sigma_\epsilon\}$. The principle underlying SMM is the same as with GMM; to match model moments with data moments. The crucial difference is that the moments implied by the model are computed by simulation as closed-form expressions for the moments are not available. The model used in this paper does not admit a closed-form solution, making SMM the natural choice.

To be precise, let $m(\beta)$ be a vector of model-implied moments obtained by simulating the model using parameter values β. to create a simulated dataset of the same size as the actual data. To ensure precision, this is carried out S times with the s-th sample moments denoted $m^s(\beta)$. Further, let m_D denote the corresponding data moments. The SMM estimator is

$$\hat{\beta} = \arg\min_{\beta} \left(\frac{1}{S} \sum_{s=1}^{S} m^s(\beta) - m_D \right)' W \left(\frac{1}{S} \sum_{s=1}^{S} m^s(\beta) - m_D \right).$$

where W is a weighting matrix. I set W to be the inverse of the variance-covariance matrix of the data moments which is the efficient weighting matrix. I obtain this matrix by resampling with replacement ("bootstrapping") from the data, computing the moments for each such sample, and then computing the covariance of those moments. The results in Michaelides and Ng (2000) indicate that a simulated dataset that is ten times as big as the actual data ensures a good performance of the SMM estimator. I therefore set $S = 10$. Lastly, to avoid the scale of moments having an effect on the estimation, I use the percentage deviation, $\frac{\frac{1}{S}\sum_{s=1}^{S} m^s(\beta) - m_D}{m_D}$, instead.

For each parameter vector β, I solve the model, use the solution to simulate S datasets, and compute the moments of the simulated data. I then use the Nelder-Mead algorithm to search the parameter space for the minimum of the objective function. In general, the objective function is well behaved. The Nelder-Mead algorithm converges to the optimum quickly, and does so from a wide range of starting points. I also check that the result is indeed the global optimum using a genetic algorithm.

The standard errors at $\hat{\beta}$ are the square roots of the diagonal elements of the covariance matrix for the parameter estimates Q where

$$Q = \left(1 + \frac{1}{S}\right) \left[\frac{\partial m(\hat{\beta})}{\partial \beta}' W \frac{\partial m(\hat{\beta})}{\partial \beta} \right]^{-1}.$$

The procedure also delivers a way to test the over-identifying moment conditions (J-test):

$$J = \frac{NS}{S+1} \min_{\beta} \left(\frac{1}{S} \sum_{s=1}^{S} m^s(\beta) - m_D \right)' W \left(\frac{1}{S} \sum_{s=1}^{S} m^s(\beta) - m_D \right)$$

$$\sim \chi^2(\text{\#moments - \#parameters}).$$

I use 6 moments to identify the three parameters.

1. The "competitive" spread, that is, the average imputed round-trip spread for transactions over USD 100,000 or more,

$$\mathbb{E}\left[A_t - B_t\right] = \mathbb{E}\left[\frac{\mathbb{E}_{j,\omega}\left(v_j(1-G(\bar{\epsilon}_{j,t}(\mu,\omega))\right)}{\mathbb{E}_{j,\omega}\left((1-G(\bar{\epsilon}_{j,t}(\mu,\omega))\right)} - \frac{\mathbb{E}_{j,\omega}\left(v_j G(\underline{\epsilon}_{j,t}(\mu,\omega)\right)}{\mathbb{E}_{j,\omega}\left(G(\underline{\epsilon}_{j,t}(\mu,\omega)\right)}\right].$$

2. The fraction traded. The average weekly trading volume in all transactions divided by N,

$$\mathbb{E}\left[\frac{N_t^0\pi\left(1-G(\bar{\epsilon}_{j,t}(\mu,\omega))\right) + N_t^1\pi G(\underline{\epsilon}_{j,t}(\mu,\omega))}{N}\right].$$

3. The imbalance. The average absolute difference between the fraction of volume in buy transactions and the fraction of volume in sell transactions,

$$\mathbb{E}\left[\frac{N_t^0\pi\left(1-G(\bar{\epsilon}_{j,t}(\mu,\omega))\right) - N_t^1\pi G(\underline{\epsilon}_{j,t}(\mu,\omega))}{N}\right].$$

4. The variance of the prices. The variance of the average weekly transaction price,

$$Var\left(\frac{\mathbb{E}_{j,\omega}\left(v_j(1-G(\bar{\epsilon}_{j,t}(\mu,\omega))\right)}{\mathbb{E}_{j,\omega}\left((1-G(\bar{\epsilon}_{j,t}(\mu,\omega))\right)}\right).$$

5. The change in prices relative to volume. The ratio between the average absolute between-week change in prices and the fraction traded,

$$\mathbb{E}\left[\left(\frac{\mathbb{E}_{j,\omega}\left(v_j(1-G(\bar{\epsilon}_{j,t}(\mu,\omega))\right)}{\mathbb{E}_{j,\omega}\left((1-G(\bar{\epsilon}_{j,t}(\mu,\omega))\right)}\right) - \left(\frac{\mathbb{E}_{j,\omega}\left(v_j(1-G(\bar{\epsilon}_{j,t-1}(\mu,\omega))\right)}{\mathbb{E}_{j,\omega}\left((1-G(\bar{\epsilon}_{j,t-1}(\mu,\omega))\right)}\right)\right] \Big/ \left(\frac{N_t^0\pi\left(1-G(\bar{\epsilon}_{j,t}(\mu,\omega))\right) + N_t^1\pi G(\underline{\epsilon}_{j,t}(\mu,\omega))}{N}\right).$$

6. The variance of the spreads,

$$Var\left(\frac{\mathbb{E}_{j,\omega}\left(v_j(1-G(\bar{\epsilon}_{j,t}(\mu,\omega))\right)}{\mathbb{E}_{j,\omega}\left((1-G(\bar{\epsilon}_{j,t}(\mu,\omega))\right)} - \frac{\mathbb{E}_{j,\omega}\left(v_j G(\bar{\epsilon}_{j,t}(\mu,\omega))\right)}{\mathbb{E}_{j,\omega}\left(G(\bar{\epsilon}_{j,t}(\mu,\omega))\right)}\right).$$

The IRT spreads and, more specifically, the difference in spreads between bonds with different maturities indicates that the trading behaviour of investors differs depending on which bond they are trading. Moreover, it may be that the fundamental quantities I am attempting to recover using the structural estimation (that is, the ease of finding a trading partner, the variance of the aggregate shock, and the variance of the idiosyncratic shock) are different for bonds with different characteristics. I therefore estimate the model in subsamples. In particular, I split the sample by time to maturity ($<$ 1 year, 1-3 years, 3-10 years, $>$ 10 years) and by grade (high yield or investment grade) as these appear to be the main characteristics affecting the trading patterns of any given bond. For completeness, I also report results for the entire sample.

Before running the estimation, I need to calibrate a few parameters that are not identified by the data. First, the initial belief μ_0. As the price in our model reflects the market default probability and almost all bonds are issued at or close to par value, I set $\mu_0 = 0.9$. Setting $\mu_0 = 1$ would obviously make the model uninteresting. I also run the estimation with the more standard $\mu_0 = 0.5$ and the results are qualitatively the same. Second, I set $\delta = 0.99$. Third, I set $v_h = 1$ (the bond does not default) and $v_l = 0$ (the bond defaults). Bond prices are reported as a percentage of face value, so I divide all prices in the data by 100 to ensure that the prices in the data and the model-implied prices are on the same scale.

I also have to calibrate N, the number of traders in the economy. I simulate the model for $N = 1$ and normalise all data moments to reflect this. Recall that πN is the number of traders who meet with a dealer in any given period (they may decide not to trade). I follow Hugonnier, Lester, and Weill (2020) and set $N = 55$ million. This is based on the number of households in the US who are potential bond holders. I then divide the weekly amount traded by N to get the fraction of households trading. For instance, if a week sees a trading volume of USD 5 million, I determine the fraction of investors who traded to be just under 10%.

A few words on the identification to conclude this section. While in general all parameters affect all moments in this model (making the structural approach necessary), the following key channels should help getting some intuition for the identification:

- π is the fraction of investors meeting a dealer in a given period. Higher values of π will therefore increase the fraction traded. π also affects the spreads. If π rises, investors' trading decisions are now more dependent on the realisations of the idiosyncratic and aggregate shocks and less dependent on the fundamental value as it now becomes easier to offload or acquire the asset in future periods if needed. This slows down learning as trades are now less informative on average.

- σ_ω is the observation noise around the reservation value. The reservation value is state dependent and noisily observing it therefore allows the dealers to learn about the state. Higher values of σ_ω decrease the speed of learning.

- Less informative trades have two effects: Dealers are less likely to be adversely selected. This reduces spreads. However, they also learn at a slower rate. This means that spreads take longer to converge to zero as information is released through trade. The average spread over a period of time may be higher or lower than before.

- σ_ϵ is the standard deviation of the idiosyncratic shocks. For any given set of buy and sell thresholds an increase in σ_ϵ reduces the difference between the proportion of investors buying and the proportion selling.

7 Results

Table 4 presents the estimated parameters for our 8 subsamples. The standard errors are good which supports the view that the variation in our choice of moments identify the model parameters reasonably well. The estimated value for π fluctuates around the 0.2 mark. This value implies that investors meet a dealer on average every 5 weeks. In terms of the search friction in our model, the absence of substantial variation in π implies that the chance of finding a trading partner does not seem to depend too much on the rating or maturity of the bond being traded. A small exception is the value for bonds with maturity 1-3 years. Looking at the data, these bonds seem to trade slightly less frequently than the other maturities and this is likely the driver of this feature of the results.

High yield	< 1 year	1-3 years	3-10 years	> 10 years
π	0.23 (0.06)	0.15 (0.01)	0.23 (0.02)	0.22 (0.01)
σ_ω	1.28 (0.36)	1.05 (0.08)	0.50 (0.03)	0.43 (0.02)
σ_ϵ	10.43 (2.56)	8.55 (2.46)	2.91 (0.35)	1.65 (0.15)

Investment grade	< 1 year	1-3 years	3-10 years	> 10 years
π	0.26 (0.13)	0.17 (0.03)	0.18 (0.02)	0.20 (0.01)
σ_ω	2.14 (1.41)	2.10 (0.31)	0.99 (0.09)	0.57 (0.04)
σ_ϵ	15.13 (4.64)	10.99 (2.22)	4.42 (0.51)	1.83 (0.14)

Table 4: Results for the non-stationary version of the model. Standard errors in parentheses.

In terms of the variability of the liquidity shocks, the estimated standard deviation of the aggregate shock is monotonically falling in the maturity for both high yield and investment grade bonds. The same holds true for the standard deviation of the idiosyncratic shock, but at a higher level.

Table 5 shows how the model fits the data. Several points are worth pointing out:

1. The model fits both the spread and the fraction traded reasonably well. While it slightly overestimates the level of the spreads, it manages to replicate the increase in spreads with maturity as well as the lower spread level for investment-grade bonds. The fraction traded does not exhibit a clear pattern.

2. The monotonically falling estimates for σ_ω reflect the pattern in the data regarding the variance of prices. In terms of the model, bonds with low maturities have already been trading for a longer time than bonds with high maturities. As information about the fundamental state is revealed through trading, there is less uncertainty about the value of low-maturity bonds. This implies lower price variability.

3. For both high-yield and investment-grade bonds, dealers are most willing to take on inventory of low maturity bonds. Again, this is consistent with the learning story (least uncertainty about the value). Due to the comparative statics explained in the methodology section, we would therefore expect to see low levels for the idiosyncratic shock standard deviation for bonds with high maturities. However, the opposite is the case. The model doesn't manage to match the empirical pattern for the trade imbalance.

4. In terms of change of prices relative to volume, the model manages to correctly match the pattern of the data. The values for high maturity bonds are higher than for low maturity bonds. However, it does not quite manage to match the magnitude, especially for high maturity bonds.

5. The fit reflects the structure of the covariance matrix of the data moments. The spread, fraction traded and Var(Prices) moments receive relatively higher weight than the Var(Spread), imbalance, and Δ(Prices)/Vol.

I use more moments than there are parameters to estimate the model. I can therefore evaluate the fit by testing the validity of the over-identifying restrictions using a J-test. The null hypothesis, $\left(\frac{1}{S}\sum_{s=1}^{S} m^s(\beta) - m_D\right) = 0$, is that the restrictions are valid. The test statistic is asymptotically distributed χ^2 with degrees of freedom equal to the number of moment conditions exceeding those needed for exact identification. In my case, this is three. As can be seen from the table, the model is rejected by the data at varying levels of confidence. The literature consensus is not to place much emphasis on such a result as any model would be rejected given enough data (see e.g. Taylor (2010) for a structural estimation with a similar fit). Nonetheless, the fit is certainly not perfect.

8 Counterfactual analysis

Table 6 reports the results of the following exercise. For each rating-maturity group (e.g. high yield bonds with 1-3 years time-to-maturity), I take the estimated parameters from table 4, and increase π by 20%. I then simulate the model at this new set of parameters in the same way as I did for the SMM estimation and report the

High yield	< 1 year		1-3 years		3-10 years		> 10 years	
	Data	Fitted	Data	Fitted	Data	Fitted	Data	Fitted
Spread	5.31	5.93	10.26	11.38	18.35	23.77	39.26	49.63
Fraction traded	14.72	11.50	10.70	7.59	13.35	11.70	10.65	10.88
Imbalance	4.03	0.95	3.32	0.61	3.51	1.36	2.57	1.91
Var(Prices)	1.96	1.79	9.59	9.23	28.60	27.05	64.78	58.97
Var(Spread)	0.01	0.00	0.03	0.00	0.05	0.01	0.13	0.04
Δ(Prices)/Vol	9.73	2.36	13.10	7.03	22.92	7.21	65.45	11.87
Test of overident. restr. (J-test)	$\chi^2 = 13.13$ ($p = 0.004$)		$\chi^2 = 17.98$ ($p = 0.000$)		$\chi^2 = 18.92$ ($p = 0.000$)		$\chi^2 = 10.46$ ($p = 0.015$)	

Investment grade	< 1 year		1-3 years		3-10 years		> 10 years	
	Data	Fitted	Data	Fitted	Data	Fitted	Data	Fitted
Spread	3.22	3.54	6.04	7.47	13.37	18.24	28.64	42.82
Fraction traded	16.79	13.05	10.92	8.35	11.25	8.94	10.38	10.02
Imbalance	4.99	1.24	3.88	1.04	4.13	1.31	3.48	2.05
Var(Prices)	0.51	0.46	1.24	1.14	7.49	7.04	31.36	28.41
Var(Spread)	0.00	0.00	0.01	0.00	0.04	0.00	0.06	0.02
Δ(Prices)/Vol	3.49	1.08	7.89	2.66	19.44	5.27	60.33	8.6
Test of overident. restr. (J-test)	$\chi^2 = 12.10$ ($p = 0.007$)		$\chi^2 = 14.68$ ($p = 0.002$)		$\chi^2 = 22.25$ ($p = 0.000$)		$\chi^2 = 14.78$ ($p = 0.002$)	

Table 5: Empirical vs fitted moments

resulting model-implied moments. This allows me to study the effect of a reduction in trading frictions.

The table shows that for all rating-maturity groups, a decrease in trading frictions prompts a decrease in the spreads. To be precise, spreads decline by 18.89 percent on average with the range being 15.93 to 24.88 percent. They decline most for the high maturity bonds and least for the bonds closest to maturity. From the discussion of the model, we know that a decrease in search frictions may increase or decrease the spread. Observing that the spread decreases for all rating-maturity pairs is therefore a result in itself. It shows that at the estimated parameters, the effect of the decrease in the adverse selection issue faced by dealers outweighs the increase due to slower learning.

To check how robust this relationship is to changes in the estimated parameters, I perform the following sensitivity analysis: Figure 1 shows this plot for the high yield bonds with over 10 years time to maturity and for investment grade bonds with less than one years time to maturity. I also perform this analysis for all other rating-maturity pairs and find that the only pair where spreads increase for at least some region of π are the high yield bonds with more than ten years' time to maturity. In general, an increase in π seems to increase spreads whenever both σ_ω and σ_ϵ are low.

Table 6 also allows us to quantify some more changes. The fraction traded trivially increases by 20 percent. Welfare increases by 21.8 percent on average with the biggest gains seen in the high maturity bonds. Δ(Prices)/Vol declines by 34.03 percent on average with the biggest changes again in the high maturity bonds. Var(Prices) declines by 40 percent on average again with the biggest change in the high maturity bonds. The imbalance increases by an average of 21.35 percent with no substantial differences between maturities.

For all of spreads, welfare, variance of prices, and Δ(Prices)/Vol the effects of the increase in π are biggest for high-maturity bonds and smallest for bonds close to maturity, both in absolute and in relative measure. One possible explanation consistent with the learning arguments brought forward in this paper would be that bonds with long time-to-maturity have the highest amount of uncertainty. Relative to bonds closer to maturity, less time has passed in which trading could have revealed information about the fundamental value and thereby reduced some of that

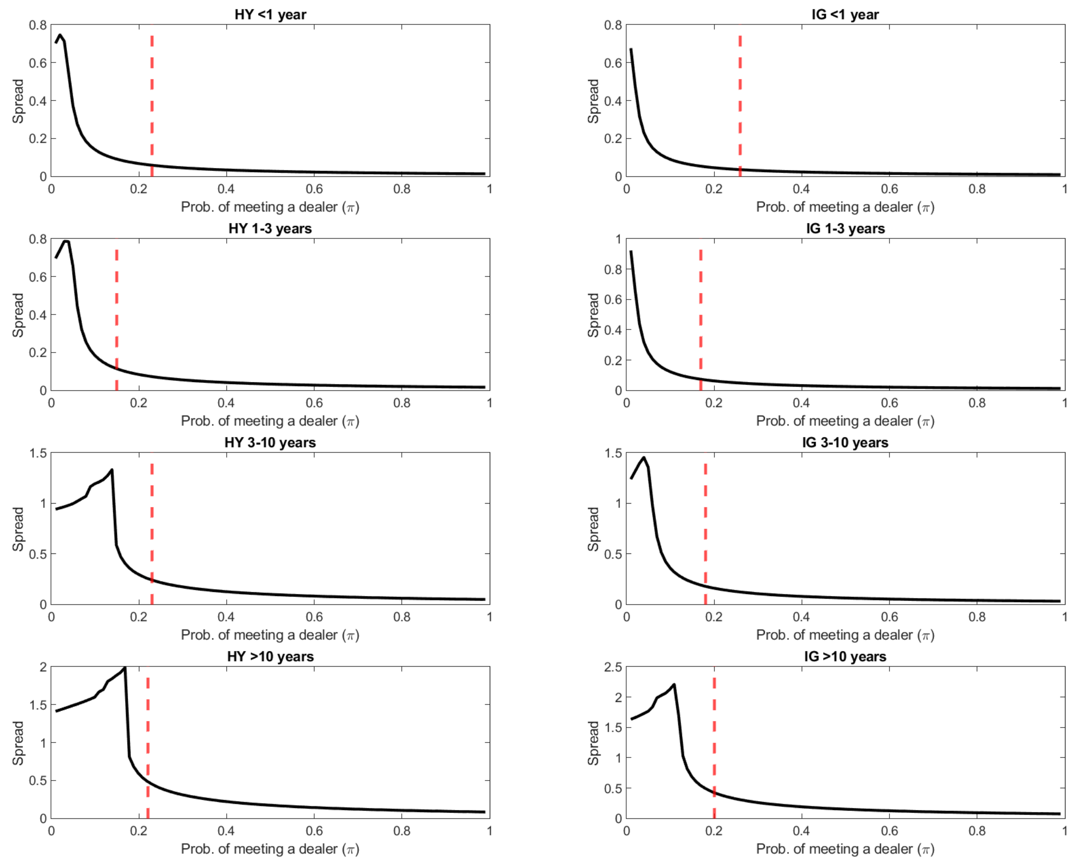

Figure 1: Sensitivity analysis: Model-implied spreads for different values of π. All other parameters are fixed at their estimated value. The top figure is for high yield bonds with over 10 years time-to-maturity and the bottom figure is for investment grade bonds with less than one years time-to-maturity. The vertical line is drawn at the SMM estimate for π (see table 4).

uncertainty.

From the discussion of the model we know that an increase in π makes traders' behaviour less dependent on the fundamental value. This reduces the information content of trades and slows down learning over time. One way to quantify the slow-down in learning is to look at how many price paths converge to their true value within a certain amount of time. I simulate 10,000 paths and look at the proportion of paths which are within 5% of the true value after 1000 periods. For instance, for the investment-grade bonds with more than 10 years time-to-maturity, the proportion of paths that reach this threshold is 66.3%, this drops to 49.1% when π increases by 20%. It drops to 0 when $\pi = 1$. Figure 2 illustrates this point by showing three such simulated price paths, all of which are for the same random shocks. However, the reduced informativeness of prices means that as π increases, beliefs take longer to converge to the true value.

9 Conclusion

In this paper, I have demonstrated how to use transaction data for a decentralised market to structurally estimate a model that features not just a trading friction, but also an information friction. While the trading friction is well established in the literature, the information friction as well as the interaction of the two remain understudied. The TRACE data used for this paper support the notion that trading behaviour reveals information as time passes. This allows market participants to learn about the fundamental value of the asset over time. The results reflect this and indicate that trading differs substantially depending on the time to maturity of any given bond. When time to maturity is long, prices react more strongly to trades. Trade imbalances are lower as dealers want to take on less inventory of assets with uncertain value. The structural estimation also allows me to perform a counterfactual analysis. I show that a reduction in trading frictions would improve liquidity, but slow down learning.

There are several avenues for future research. To better capture the nature of trading in OTC markets one would need a fully decentralized model that does not assume a frictionless inter-dealer market as the model used in this paper does. Fur-

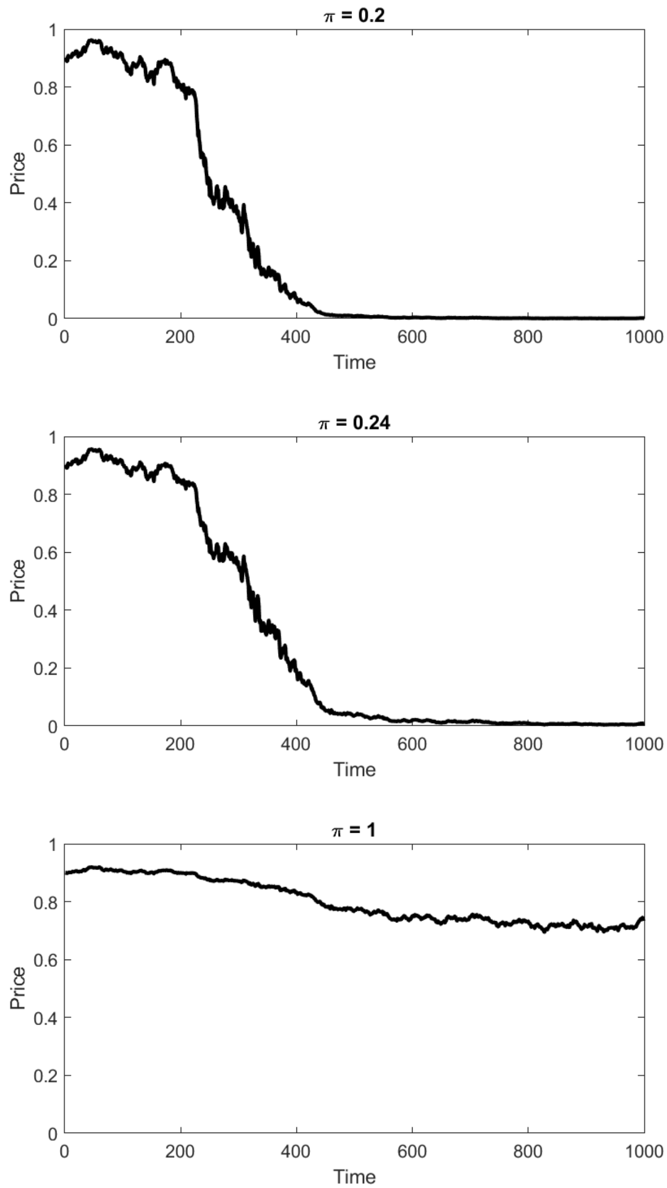

Figure 2: Simulated price paths for $\pi = 0.2$ (top), $\pi = 0.24$ (middle), and $\pi = 1$ (bottom). Remaining parameters fixed at estimated values for investment-grade bonds with > 10 years time-to-maturity.

High yield	< 1 year		1-3 years		3-10 years		> 10 years	
	Fitted	$\pi = 0.28$	Fitted	$\pi = 0.18$	Fitted	$\pi = 0.28$	Fitted	$\pi = 0.26$
Spread	5.93	4.93	11.38	9.34	23.77	18.81	49.63	37.28
Fraction traded	11.50	13.79	7.59	9.11	11.70	14.04	10.88	13.07
Imbalance	0.95	1.15	0.61	0.74	1.36	1.65	1.91	2.31
Var(Prices)	1.79	1.18	9.23	5.52	27.05	14.56	58.97	29.26
Var(spread)	0.00	0.00	0.00	0.00	0.01	0.00	0.04	0.01
Δ(Prices)/Vol	2.36	1.63	7.03	4.75	7.21	4.56	11.87	6.78
Welfare	1.00	1.20	0.54	0.65	0.29	0.36	0.16	0.20

Investment grade	< 1 year		1-3 years		3-10 years		> 10 years	
	Fitted	$\pi = 0.31$	Fitted	$\pi = 0.20$	Fitted	$\pi = 0.21$	Fitted	$\pi = 0.24$
Spread	3.54	2.97	7.47	6.28	18.24	15.03	42.82	33.86
Fraction traded	13.05	15.66	8.35	10.02	8.94	10.72	10.02	12.02
Imbalance	1.24	1.51	1.04	1.27	1.31	1.59	2.05	2.48
Var(Price)	0.46	0.31	1.14	0.78	7.04	4.25	28.41	15.33
Var(Spread)	0.00	0.00	0.00	0.00	0.00	0.00	0.02	0.01
Δ(Prices)/Vol	1.08	0.75	2.66	1.86	5.27	3.59	8.6	5.44
Welfare	1.65	1.98	0.77	0.93	0.34	0.41	0.17	0.21

Table 6: Counterfactual analysis, increase probability of finding a trading partner by 20%.

thermore, the trading between investors and dealers could also be generalised. For instance, removing the restriction of holdings to be 1 or 0 would make the model more realistic and allow the dynamics of trading volume to better match those observed in the data. The learning process currently assumes that investors know the state perfectly. In reality, this is likely too strong and the less strict assumption that some investors obtain an imperfect signal about the true state is more appropriate. However, this would require modelling not just the evolution of beliefs of dealers, but also those of investors and therefore probably presents a significant modelling challenge.

References

Bessembinder, Hendrik, William Maxwell, and Kumar Venkataraman (2006). "Market transparency, liquidity externalities, and institutional trading costs in corporate bonds". In: *Journal of Financial Economics* 82.2, pp. 251–288.

Biais, Bruno (1993). "Price formation and equilibrium liquidity in fragmented and centralized markets". In: *The Journal of Finance* 48.1, pp. 157–185.

Brancaccio, Giulia, Dan Li, and Norman Schürhoff (2017). "Learning by trading: The case of the us market for municipal bonds". In: *Unpublished paper. Princeton University.*

Carrillo, Paul E (2012). "An empirical stationary equilibrium search model of the housing market". In: *International Economic Review* 53.1, pp. 203–234.

Duffie, Darrell, Nicolae Gârleanu, and Lasse Heje Pedersen (2005). "Over-the-counter markets". In: *Econometrica* 73.6, pp. 1815–1847.

Dugast, Jérôme, Semih Üslü, and Pierre-Olivier Weill (2019). "A theory of participation in OTC and centralized markets". In: *National Bureau of Economic Research Working Paper 25887.*

Eckstein, Zvi and Kenneth I Wolpin (1990). "Estimating a market equilibrium search model from panel data on individuals". In: *Econometrica: Journal of the Econometric Society*, pp. 783–808.

Edwards, Amy K, Lawrence E Harris, and Michael S Piwowar (2007). "Corporate bond market transaction costs and transparency". In: *The Journal of Finance* 62.3, pp. 1421–1451.

Feldhütter, Peter (2012). "The same bond at different prices: identifying search frictions and selling pressures". In: *The Review of Financial Studies* 25.4, pp. 1155–1206.

Gavazza, Alessandro (2016). "An empirical equilibrium model of a decentralized asset market". In: *Econometrica* 84.5, pp. 1755–1798.

Glode, Vincent and Christian C Opp (2018). "Over-the-Counter vs. Limit-Order Markets: The Role of Traders' Expertise". In: *Limit-Order Markets: The Role of Traders' Expertise (November 8, 2018).*

Glosten, Lawrence R and Paul R Milgrom (1985). "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders". In: *Journal of financial economics* 14.1, pp. 71–100.

Hagströmer, Björn and Albert J Menkveld (2019). "Information revelation in decentralized markets". In: *The Journal of Finance* 74.6, pp. 2751–2787.

Hugonnier, Julien, Benjamin Lester, and Pierre-Olivier Weill (2020). "Frictional intermediation in over-the-counter markets". In: *The Review of Economic Studies* 87.3, pp. 1432–1469.

Lester, Benjamin et al. (2018). "Market-Making with Search and Information Frictions". In: *NBER Working Paper* w24648.

Li, Dan and Norman Schürhoff (2019). "Dealer networks". In: *The Journal of Finance* 74.1, pp. 91–144.

Liu, Shuo (2020). "Dealer's search intensity in U.S. corporate bond markets". In: *Working Paper, UCLA.*

McFadden, Daniel (1989). "A method of simulated moments for estimation of discrete response models without numerical integration". In: *Econometrica: Journal of the Econometric Society*, pp. 995–1026.

Michaelides, Alexander and Serena Ng (2000). "Estimating the rational expectations model of speculative storage: A Monte Carlo comparison of three simulation estimators". In: *Journal of econometrics* 96.2, pp. 231–266.

O'Hara, Maureen and Xing Alex Zhou (2019). "The Electronic Evolution of Corporate Bond Dealers". In: *Journal of Financial Economics (JFE), Forthcoming*.

Pakes, Ariel and David Pollard (1989). "Simulation and the asymptotics of optimization estimators". In: *Econometrica: Journal of the Econometric Society*, pp. 1027–1057.

Taylor, Lucian A (2010). "Why are CEOs rarely fired? Evidence from structural estimation". In: *The Journal of Finance* 65.6, pp. 2051–2087.

Vogel, Sebastian (2019). "When to Introduce Electronic Trading Platforms in Over-the-Counter Markets?" In: *Available at SSRN 2895222*.

Weill, Pierre-Olivier (2020). "The Search Theory of OTC Markets". In: *Annual Review of Economics* 12.

Appendix

A Counterfactual no search friction

Table 7 contrasts the model-implied moments when search frictions are completely removed ($\pi = 1$) with the data moments. The direction of change is the same as in table 6, where search frictions are reduced by 20%, but the magnitude of effects is bigger.

B Results for the full sample

The main text groups bonds by maturity and rating. This section contains the results for the full data sample. Table 8 presents the parameter estimates. Again, we obtain a value for π around 0.2 which implies that investors find a trading partner on average once every 5 weeks. The estimate for σ_ϵ is larger than the estimate for σ_ω. The data reject the model at the 1% confidence level, but as argued before this is not a disastrous result. Table 9 shows how the model fits the data and also performs

High yield	< 1 year		1-3 years		3-10 years		> 10 years	
	Fitted	$\pi = 1$	Fitted	$\pi = 1$	Fitted	$\pi = 1$	Fitted	$\pi = 1$
Spread	5.93	1.35	11.38	1.65	23.77	4.82	49.63	8.36
Fraction traded	11.50	50.00	7.59	50.00	11.70	50.00	10.88	49.98
Imbalance	0.95	5.37	0.61	5.41	1.36	7.54	1.91	11.28
Var(Prices)	1.79	0.08	9.23	0.12	27.05	0.56	58.97	0.76
Var(spread)	0.00	0.00	0.00	0.00	0.01	0.00	0.04	0.00
Δ(Prices)/Vol	2.36	0.12	7.03	0.15	7.21	0.32	11.87	0.37
Welfare	1.00	4.36	0.54	3.58	0.29	1.27	0.16	0.74

Investment grade	< 1 year		1-3 years		3-10 years		> 10 years	
	Fitted	$\pi = 1$	Fitted	$\pi = 1$	Fitted	$\pi = 1$	Fitted	$\pi = 1$
Spread	3.54	0.93	7.47	1.26	18.24	3.13	42.82	7.42
Fraction traded	13.05	50.00	8.35	50.00	8.94	50.00	10.02	49.99
Imbalance	1.24	6.21	1.04	8.38	1.31	9.76	2.05	13.37
Var(Price)	0.46	0.03	1.14	0.03	7.04	0.14	28.41	0.42
Var(Spread)	0.00	0.00	0.00	0.00	0.00	0.00	0.02	0.00
Δ(Prices)/Vol	1.08	0.07	2.66	0.08	5.27	0.16	8.6	0.28
Welfare	1.65	6.31	0.77	4.62	0.34	1.91	0.17	0.85

Table 7: Counterfactual analysis, complete removal of search friction

counterfactual analyses for $\pi = 1$ and for a 20% decrease in trading frictions ($\pi = 0.23$).

	π	σ_ω	σ_ϵ
Estimate	0.19	0.78	4.61
Standard Error	0.0141	0.0631	0.7595
Test of overident. restr.	$\chi^2 = 19.44$ ($p = 0.000$)		

Table 8: Results for the aggregate data

Moment	Data	Fitted	$\pi = 0.23$	$\pi = 1$
Spread	14.14	17.58	13.65	3.03
Fraction traded	11.97	9.26	11.57	50.00
Imbalance	3.89	1.03	1.30	7.36
Var(Prices)	13.25	12.47	6.34	0.22
Var(Spread)	0.03	0.00	0.00	0.00
Δ(Prices)/Vol	23.39	6.59	4.01	0.20

Table 9: Empirical and fitted moments for the entire data set