

Should Corporate Bond Trading Be Centralized?

Sébastien Plante

University of Pennsylvania*

January 16, 2018

Abstract

This paper shows that centralizing the US corporate bond market would yield large gains in efficiency. By studying two markets where corporate bonds are successfully traded on central limit order books, I estimate that the transaction costs of US corporate bonds would decrease by 70% on average if trading migrated from over-the-counter markets to limit order markets. To study the social value of reforming the corporate bond market, I build a parsimonious model of centralized and decentralized trading. The model implies that the optimal market structure can be determined by appropriately scaling the transaction costs associated with each market structure. Estimating the scaling factors reveals that moving to limit order markets would generate a social surplus equal to 21 bps of total par value. Large bond issues with low credit ratings and long time to maturity would benefit the most.

*Plante: splante@wharton.upenn.edu

1 Introduction

In the US, corporate bond transactions are intermediated by dealers in decentralized over-the-counter (OTC) markets. Unlike US equities, which are largely traded in transparent central limit order books (CLOB), corporate bonds are traded in fragmented markets with little pre-trade transparency since dealer quotations are not consolidated and publicly available. It is often suggested that the opaqueness associated with this market structure is responsible for the lack of liquidity and high transaction costs of corporate bonds (see for example Biais and Green (2007), Harris (2015), and Harris et al. (2015)).

How efficient is the current market structure? Should corporate bonds instead be traded in centralized limit order markets, just like US equities are? Would mandated changes regarding the consolidation and disclosure of quotes improve liquidity? Given the importance of the \$9 trillion US corporate bond market as a source of capital formation, these questions have captured the interest of policymakers, practitioners, and academics alike. For example, when speaking about the role of the SEC in the fixed-income market, former SEC Commissioner Gallagher (2015) said, “At the Commission level, there has been a lot of discussion about issues such as enhanced pre and post trade transparency . . . For example, the Commission should be looking at all options for facilitating electronic and on exchange transactions.” At this point, however, there is no consensus on whether a centralized market structure would actually improve the liquidity of corporate bonds.

Several market participants and observers are skeptical about the viability of a centralized market structure, for at least three reasons. A first argument is that corporate bonds are primarily illiquid because investors don’t need to trade them often. As a result, monitoring and updating firm quotes would be too costly and centralized venues would fail to attract competing limit orders. A second argument is that corporate bonds are mostly traded by large institutions that would not find the liquidity they need on a CLOB. A third argument is that the sheer number of bond issues, and their bespoke contractual features, make them unsuitable for centralized trading.¹ In other

¹These arguments are frequently mentioned in the financial press, see for example Childs (2014), Ng and Grind (2012), and BIS (2014)

words, the concern is that bonds are intrinsically illiquid such that changing the market structure would not improve market quality.

Other observers have a more positive outlook. For example, Harris et al. (2015) make the case for a centralized corporate bond market and suggest, “The US Securities and Exchange commission (SEC) could rapidly and substantially improve bond market efficiency by simply requiring brokers to post their customers’ limit orders to an electronically accessible broker platform or alternative trading system (ATS), where one customer’s limit order could trade against another customer’s order without dealer intermediation.” Their argument is supported by the successful implementation of comparable market reforms on the Nasdaq in 1997, a market that operated in a similar way as the modern US corporate bond market. In summary, the merits of a centralized market structure for corporate bonds are still being debated.

Surprisingly, given the importance of the issue for policy, I am not aware of any study that quantifies whether a centralized market structure would be more efficient than the current OTC structure. This paper fills that gap and makes two main contributions. First, by studying two markets where corporate bonds trade successfully on central limit order books, I estimate that transaction costs would decrease by 70% on average if US bonds traded on similar venues. This implies that, over the life of a bond, investors would on average save the equivalent of 0.66% of a bond’s total par value in transaction costs. Second, I evaluate the welfare implications of centralizing the market structure through the lens of a structural model and obtain large and positive estimates. For the average bond, a benevolent social planner would spend up to 21 bps of total par value to migrate trading to centralized venues. In the cross-section of bonds, I document that larger issues with lower credit ratings and longer time to maturity would benefit the most.

To construct counterfactual transaction cost estimates, I rely on evidence from the Israeli market and the early 20th century US market where corporate bonds are/were successfully traded on exchange. In Israel, corporate bonds have been actively traded on the CLOB of the Tel Aviv Stock Exchange for several decades. Abudy and Wohl (2017) and Protnick and Gur-Gershgoren (2011) discuss the institutional setting of this market. In order to measure the liquidity of Israeli

bonds and stocks, I obtain trade and quotation data from the Tel Aviv Stock Exchange (TASE).² In the first half of the century, US corporate bonds were also actively traded on exchange, most prominently on the NYSE. Biais and Green (2007) provide a detailed account of the microstructure of the US bond market in the 20th century and document the demise of exchange trading in the mid to late 1940s. In order to measure the liquidity of US bonds and stocks in the early 20th century, I hand collect historical data on quotation, volume, and security characteristics for a large number of securities from 1917 to 1921.

The evidence from these two markets reveals that exchange traded corporate bonds have similar, if not lower, transaction costs than equities. This observation holds both unconditionally and after controlling for known determinants of the bid-ask spread, such as trading volume and return volatility. In contrast, corporate bond transaction costs in the modern US market are several times larger than equity transaction costs.

In order to estimate the transaction costs that modern US corporate bonds would have on a centralized venue, I propose an estimator based on equity transaction costs and observable bond characteristics. I test the accuracy of the equity-implied estimator in the Israeli and historical US markets, where estimated and measured transaction costs can be compared. In both samples, equity-implied and measured transaction costs are highly correlated, and the equity-implied estimates are on average larger than measured transaction costs. As a result, applying the estimator to modern US bonds provides a conservative estimate of how much transaction costs would decrease if trades were centralized. The estimates suggests that transaction costs would decrease by 70% relative to their current OTC levels.

While obtaining these counterfactual transaction cost estimates is informative, the exercise does not allow us to conclude that a centralized market structure would be more efficient. Indeed, transaction costs are wealth transfers between liquidity providers and liquidity seekers, not deadweight losses. To estimate the social value of centralizing the market structure, I build a parsimonious model of centralized and decentralized trading that delivers, in closed-form, a for-

²I am indebted to Avi Wohl for helping me gain access to this data.

mula for the relative efficiency of the two market structures. This is done by adapting the seminal framework of decentralized trading introduced by Duffie et al. (2005) (DGP) in such a way that both market structures can be represented in the same economic environment.

As in DGP, the decentralized market structure is modeled as a request for quote trading procedure, where finding counterparties is time consuming and subject to search frictions. Upon contact, bargaining between customers and dealers determines transaction prices. These modeling assumptions are in line with the institutional setting of the modern US corporate bond market; in practice, customers must often sequentially contact dealers and negotiate over the terms of trade. The main departure from the original DGP framework has to do with the market participation of investors and market makers.³ Instead of assuming exogenous measures of investors and market makers, I model their endogenous and dynamic participation.⁴ Remarkably, enriching the DGP framework with endogenous participation delivers a tractable model that can be solved in closed form solution, both in and out of steady state. More importantly, the modified framework can seamlessly be adapted to model centralized trading.

The centralized market structure is modeled as a limit order book, where liquidity providers post executable quotes that are publicly observable. Despite the transparency of the centralized venue, trades are not executed instantly. In practice, concerns related to price impact lead investors to split and execute their orders over time. For example, Admati and Pfleiderer (1988) show that strategic liquidity traders may delay the execution of their orders to minimize transaction costs. The model accommodates these frictions through a trade processing function that induces delays in order execution. Formally, I model the decentralized market structure within a random search framework, as in DGP, while the centralized market structure is modeled within a directed search framework, in the spirit of Moen (1997) and Lester et al. (2015). Beside the difference in trading procedure, the economic environment is identical in the two versions of the model.

³I also bring a few technical modifications to the original DGP setup: I replace the increasing return to scale matching function with a constant return to scale matching function, and abstract from the interdealer market.

⁴In Section 7, DGP study an economy where dealers make a static decision on the intensity of their market making activity and abide by it. They write, "A full dynamic analysis of the optimal control of market making intensities with small switching costs would be interesting, but seems difficult." The modified framework presented here provides a tractable way to model the dynamic liquidity provision of buy-side investors and sell-side dealers.

From a quantitative perspective, the interest of the model lies in the simple and intuitive formulas it delivers to evaluate the relative efficiency of the two market structures. Specifically, the model implies that the difference between welfare on the centralized market structure, W^c , and welfare on the decentralized market structure, W^d , is proportional to

$$W^c - W^d \propto \frac{S^d}{\theta_m} - \frac{S^c}{\beta^c}, \quad (1)$$

where S^d and S^c are the respective round-trip transaction costs on the decentralized and centralized venues, and θ_m and β^c are the shares of trade surplus that accrue to market makers on the decentralized and centralized venues, respectively.

The welfare formula highlights that comparing transaction costs is not sufficient to determine which market structure is optimal. To compare efficiency, transaction costs must be scaled by the fractions of trade surplus that accrue to market makers, θ_m and β^c . Intuitively, transaction costs are the product of two components. The first one is related to the degree of trading frictions, and is relevant for efficiency: trading costs are larger when the asset is difficult to trade, and more frictions in the trading process generates misallocations which are costly from a social perspective. The second component is related to the fraction of trade surplus that market makers extract, which corresponds to a wealth transfer, and isn't relevant for efficiency. In order to compare efficiency, transaction costs must be adjusted to reflect the effect of the later component. Section 5 formalizes this intuition.

Since the current transaction costs of US corporate bonds, S^d , can directly be measured in the data, and since the equity-implied estimator provides counterfactual estimates for S^c , we are left to estimate the two remaining parameters, θ_m and β^c . The estimation indicates that market makers extract a larger share of surplus on over-the-counter markets, $\hat{\theta}_m > \hat{\beta}^c$, such that transaction costs must decrease by at least 39% of their OTC level for a centralized market structure to be optimal. Estimating the proportionality constant in equation (1), I find that a social planner would pay up to 21 bps of total par value on average to migrate trading to centralized venues. In the cross-section of bonds, larger issues with lower credit ratings and longer time to maturity would benefit

the most.

The rest of the paper is organized as follow. Section 2 reviews the related literature. Section 3 describes the data. Section 4 constructs the equity-implied estimates. Section 5 introduces the model and derives the welfare formula. Section 6 estimates the welfare implications. Section 7 concludes.

2 Related Literature

This paper contributes to the empirical literature on corporate bond liquidity by providing counterfactual estimates of the transaction costs that corporate bonds would have on limit order markets. This paper also contributes to market structure theory by building an empirically tractable framework that delivers a simple formula to evaluate the relative efficiency of centralized and decentralized market structures.

2.1 Empirical Literature

Several papers have studied the transaction costs of modern US corporate bonds. The salient findings are that transaction costs are high, particularly for small orders; and that the implementation of TRACE, which introduced some level of post-trade transparency, has reduced transaction costs for all trade sizes (Schultz (2001), Bessembinder et al. (2006), Edwards et al. (2007), Goldstein et al. (2007)). The observation that small orders receive worst prices is usually attributed to the ability of OTC dealers to extract rents from less sophisticated investors. The impact of TRACE on transaction costs suggests that post-trade transparency has been beneficial to investors.

Using a proprietary dataset of intraday quotations, Harris (2015) finds many instances of trade-through – i.e., trades occurring outside the quoted spread – suggesting that the market would likely benefit from additional pre-trade transparency and price protection rules. O’Hara et al. (2017) and Hendershott et al. (2017) find that smaller and less active insurance companies systematically receive worst execution prices relative to more active insurance companies, suggesting that dealers

can extract rents from less active investors. Hendershott and Madhavan (2015) study the transaction costs of corporate bond orders executed on MarketAxess's electronic platform. Although this platform also operates on a request for quote basis, the venue allows customers to contact multiple dealers at once which they show reduce transaction costs. They argue that such venues might pave the way toward centralized and continuous trading.

Biais and Declerck (2007) study the OTC European corporate bond market and similarly observe that large orders obtain better prices. They also report that, for all trade sizes, transaction costs are lower for Euro bonds than for US bonds, a finding they attribute to fiercer competition among the dealers from different Eurozone countries. Harris and Piwowar (2006) and Green et al. (2007) document similar patterns in the US municipal bond market, another decentralized OTC market.

While these papers suggest that opaque OTC markets generate large transaction costs, they do not address how liquid bonds would be on limit order books. One of the main contribution of my empirical analysis is to provide such counterfactual estimates.

A few papers have investigated the trading of corporate bonds on limit order markets. Abudy and Wohl (2017) and Protnick and Gur-Gershgoren (2011) study the case of the Israeli market, where corporate bonds have been traded on the CLOB of the Tel Aviv Stock Exchange for several decades. They report that, despite a relatively small market capitalization, the Israeli corporate bond market is quite liquid. Abudy and Wohl (2017) also show that Israeli corporate bonds are on average less expensive to trade than stocks. Odegaard (2017) studies the liquidity of corporate bonds on the CLOB of the Oslo Stock Exchange and also finds that corporate bonds are less expensive to trade than stocks.

Biais and Green (2007) study the historical US corporate bond market and report that, until the late 1940s, bonds were actively traded on exchange. They document that in the early 1940s, exchange traded bonds were fairly liquid, with transaction costs similar to their modern OTC counterparts. They document that trading left the exchange in the late 1940s, when insurance

companies owned close to 80% of the aggregate bond market.⁵ Biais and Green (2007) also document that corporate bond exchange trading never gained traction again, despite a growing share of retail ownership in the 1970s, likely because of order flow externalities. Once a market has captured most of the order flow, investors individually have no incentive to deviate and submit their order to a potentially more efficient, but currently illiquid, market.

While these papers show that, corporate bonds have been successfully traded on limit order markets in other environments, they do not speak to how liquid would modern US bonds be if trading were centralized today. The estimates I provide in this paper suggest that modern US bond transaction costs would decrease significantly. Moreover, my paper contributes to the literature by hand collecting and analyzing a novel historical dataset of US securities. I study a sample of 76 stocks and 188 bonds from the beginning of 1917 to the end of 1921, while Biais and Green (2007) study a sample of 6 bonds from the beginning of 1943 to the end of 1947.⁶

2.2 Theoretical Literature

From a theoretical perspective, my paper is related and contributes to the literature on intermediated OTC markets that evolved from Duffie et al. (2005) (DGP).⁷ In the original DGP framework, an exogenously specified measure of investors participate in an OTC market where they must search for counterparties and bargain over the terms of trade upon contact. I modify the original random search framework of DGP by modeling the endogenous and dynamic participation of investors and market makers, while keeping the analysis highly tractable. Importantly, the modified framework can easily be adapted to accommodate a centralized market structure.

I model a central limit order book by substituting the random search framework of DGP for

⁵Evidence from the *Kimber's Record of Insurance Company Security Purchases*, dating back to 1915, show that insurance companies traded bonds and stocks almost exclusively OTC, even when these securities were actively traded on exchange. A possible explanation for the reluctance of insurance companies to trade on exchange is that the NYSE had a mandated floor on broker commissions, while these commissions were fully negotiable OTC.

⁶The sample of Biais and Green contains intraday transaction data; my sample consists of monthly quotation data.

⁷Examples of paper who extended the original DGP framework include Lagos and Rocheteau (2009) and Lester et al. (2015) who introduce heterogeneity in the framework and Lagos et al. (2011) who study liquidity crises in a related setup.

a directed search framework. Directed search was first introduced by Moen (1997) in the context of the labour market. Lester et al. (2015) apply the concepts of directed search to model an OTC market where dealers compete for order flow by publicly committing to firm prices.⁸ In this paper, directed search is used to model how orders are handled in a central limit order book.

My paper is also related to the literature on market structure optimality. Pagano (1989) and Rust and Hall (2003) develop models where investors choose between trading on an exchange and searching for counterparties. Rust and Hall (2003) find that the existence of a centralized trading venue improves welfare; Pagano (1989) obtains ambiguous welfare implications when some of the order flow is executed off-exchange. Babus and Parlatore (2017) develop a model where a decentralized market structure may prevail in equilibrium, even though a centralized market structure is optimal. Glode and Opp (2017) show that a decentralized market structure may improve allocative efficiency by limiting harmful screening.

My paper contributes to this literature by developing a model that delivers sufficient statistics for the relative efficiency of centralized and decentralized trading venues. For a survey of the literature on sufficient statistics, see Chetty (2009).

3 Data

This paper studies the markets for equities and corporate bonds in three different environments: the modern US market (2013–2015), the early 20th century US market (1917–1921), and the modern Israeli market (2013–2015). The main objective of the empirical analysis is to obtain counterfactual estimates for the transaction costs that modern US corporate bonds would have on limit order markets. Since these estimates are inferred from equity transaction costs, the samples are composed of firms that have both listed equities and publicly traded corporate bonds.

Data availability dictates the time coverage of my samples. The Israeli data I have access to covers the period 2013–2015, thus for comparability I select the same time period for the modern

⁸In the US corporate bond market trades are typically executed well within dealer quoted spreads (Table 1 shows that effective spreads are about 40% the size of quoted spreads in my sample). Thus a price setting mechanism that involves bargaining seems to provide a more accurate representation of the current trading procedure.

US sample. The early 20th century data must be collected manually from sources publicly available up to the end of 1921. I choose to collect data for the last five years available since this is when the cross sectional coverage is the most comprehensive.

Table 1 contains relevant descriptive statistics for the transaction costs and general characteristics of the securities included in my three samples. The attributes of the samples are compared and contrasted below.

3.1 Modern US Market

Despite a large number of active market centers, trading in listed US equities is quite centralized. The bids and offers from all trading venues are consolidated in real time to form the National Best Bid and Offer (NBBO), which contains the highest bid and the lowest offer across all market centers. In addition, the price priority of the NBBO is enforced market wide: the order protection rule of Regulation NMS ensures that investors receive an execution price at least as good as the standing NBBO. For a detailed account of US equities trading environment, see Angel et al. (2011) and Angel et al. (2015).

To measure the liquidity of US equities, I obtain intraday trade and quotation data from the daily TAQ database. This database contains a complete history of all trades and quotations in listed equities time-stamped to the millisecond. The transaction data include execution prices and quantities; the quotation data include quoted bid and ask, as well as the quantities offered. This high frequency data is complemented with security level characteristics from CRSP.

Unlike equities, nearly all US corporate bonds transactions are executed on decentralized over-the-counter markets offering little pre-trade transparency. While some dealers advertise their quotes on electronic platforms such as Bloomberg and Tradeweb, these quotes are often simply indicative and used as the starting point of negotiations. Even when executable quotes are advertised, there is no systematic effort to consolidate and disseminate the information. Although pre-trade transparency is quite limited, the market offers some degree of post-trade transparency since transactions are disseminated through the TRACE database within 15 minutes of execution.

To measure the liquidity of US corporate bonds, I obtain intraday transaction data from the enhanced TRACE database. As opposed to the standard version of the database, the enhanced version reports uncensored trade sizes, which allows for accurate computation of trading volume. End of day corporate bond quotations are from Bloomberg. The trade and quotation data is complemented with bond characteristics data from Mergent FISD.

To be admissible in the sample in a given year, a security must: be covered by CRSP or Mergent FISD, have an amount outstanding of at least 50 million dollars, have an annual trading volume of at least 2 million dollars, and have enough quotation midpoints to compute at least 40 daily returns. Convertible bonds are excluded from the sample. The final sample is composed of all the admissible securities issued by firms that have an admissible common stock listed and at least one admissible corporate bond issue. The final sample is composed of 457 firms that collectively have 1,801 bonds outstanding. The equity sample contains 5.5 billion transactions and 110 billion intraday quotes, while the corporate bond sample contains 2.5 million transactions and 0.4 million end of day quotes.

Using the trade and quotation data, I compute two annual measures of liquidity: the quoted spread and the effective spread. For stocks, the quoted spread at time t is defined as

$$\text{quoted spread}_t = \frac{A_t - B_t}{M_t}, \quad (2)$$

where A_t and B_t are the respective best ask and bid standing at time t , and M_t is the average of B_t and A_t or midpoint. For a given stock in a given year, the quoted spread is defined as the time-weighted average of quoted spread $_t$ over the year. The effective spread for a given transaction k is defined as

$$\text{effective spread}_k = 2 \left| \frac{P_k - M_k}{M_k} \right|, \quad (3)$$

where P_k is the price of the k^{th} transaction, M_k is the standing midpoint at the time of the transaction. For a given stock in a given year, the effective spread is defined as the dollar-volume-weighted average of effective spread $_k$ over the year.

These two liquidity measures must be defined slightly differently for corporate bonds since quotations are only observed at a daily frequency. For a given bond in a given year, the annual measure of quoted spread is defined as the simple average of end-of-day quoted spread over the year. The annual effective spread measure is computed as the dollar-volume-weighted average of the transaction level effective spreads calculated as in (3), where the closing mid point of the previous day is used for M_k .

The first two columns of Table 1 present descriptive statistics for the modern US sample. The first two rows show that the transaction costs of US corporate bonds are several times larger than those of equities: the average quoted and effective spreads are respectively 0.56% and 0.23% for bonds, compared to 0.09% and 0.07% for equities. This is striking since the standard deviation of equity returns (32%) is on average several times larger than the standard deviation of bond returns (5%), and volatility is known to be positively associated with transaction costs. On the other hand, the corporate bond market is thinner than the equity market, which could arguably explain their high transaction costs: the average equity (bond) issue has a market capitalization of 22.3 (0.7) billion dollars, a daily trading volume of 130.38 (2.11) million dollars and an annual turnover of 268% (86%). The lower trading activity of the bond market is often argued to be the reason why corporate bonds are not suitable for centralized trading. I provide evidence against this view in the next two subsections.

The corporate bonds in my sample have standard characteristics. The average bond was issued with 12 years to maturity, and has 8 years left until maturity. Callable debt is the norm with 96% of the bond having the feature. About one-fifth of the bonds are secured. None of the bonds have sinking funds or CPI indexation clauses. In terms of credit ratings, 71% of the bonds are investment grade (IG), 20% are high yield (HY), and the remaining 9% are not rated (NR).

3.2 Israeli Market

In Israel, corporate bonds are actively traded on exchange along with equities on the Tel-Aviv Stock Exchange (TASE), the only exchange in the country. Abudy and Wohl (2017) show that

only 6% of the aggregate trading volume in corporate bonds occurs off exchange, meaning that investor liquidity needs are mostly satisfied on the exchange.

To measure the liquidity of Israeli equities and corporate bonds, I obtain intraday trade and quotation data from the TASE. Similar to the US TAQ data, the Israeli trade and quotation data contains a complete intraday history of all trades and quotations on the exchange, including execution prices, transaction volume, quoted bid and ask, and the quantities available at those prices. The high frequency data is complemented with security level characteristics from Bloomberg, and with Moody's credit ratings obtained from the website www.valuation.co.il.⁹

To be included in my samples, Israeli securities must go through the same filters that were applied to US securities.¹⁰ The final Israeli sample is composed of 51 firms that collectively have 139 bonds outstanding. The equity sample contains 7 million transactions and 140 million intraday quotes. The bond sample contains 4 million transactions and 321 million intraday quotes. This high frequency data is used to compute yearly aggregates of quoted and effective spreads, following the procedure outlined for US equities.

The third and fourth columns of Table 1 present descriptive statistics for the Israeli bond and equity samples. When applicable, amounts expressed in New Israeli Shekel (NIS) have been converted to US dollars assuming that 1 USD is worth 3.7 NIS, which corresponds to the average exchange rate over 2013-2015. The first two rows of the table show that, in sharp contrast with the modern US sample, Israeli equities are about three times more expensive to trade than Israeli bonds: the average quoted and effective spreads are respectively 0.22% and 0.20% for bonds, and 0.73% and 0.65% for equities. Moreover, the transaction costs of exchange traded Israeli bonds are about one-third those of OTC traded US bonds. This is remarkable since the Israeli corporate bond market is much thinner than its US counterpart: the average Israeli (US) bond has a market capitalization of 228 (690) million dollars, a daily trading volume 0.55 (2.11) million dollars, and an

⁹I thank Eran Ben-Horin for providing me with the data.

¹⁰I require securities to be covered by Bloomberg (instead of CRSP or Mergent FISD). To compute the market capitalization and volume filters, I convert the New Israeli Shekels (NIS) amounts to US dollars using an exchange rate of 3.7, the average exchange rate over 2013-2015. Otherwise, the sample selection procedure is the same as for modern US securities.

annual turnover of 61% (86%). The Israeli example shows that corporate bonds can be successfully traded on a centralized venue in market that is thinner than the current US bond market.

In terms of their characteristics, Israeli bonds are similar to their US counterparts in several respects: they have similar time to maturity, similar volatility, and a similar fraction of them are secured. In both samples, over 90% of the bonds have early redemption features, with the difference that US bonds are mostly callable while Israeli bonds mostly have sinking funds. Israeli and US bonds are different in that CPI indexation is quite frequent in Israel, while nonexistent in the US sample. In addition, a larger fraction of Israeli bonds are unrated.

3.3 Early 20th Century US Market

During the first half of the century, US bonds and stocks were actively traded on organized exchanges and on over-the-counter markets. Hickman (1960) reports that, of all the active exchanges, the New York Stock Exchange (NYSE) had the largest listed capitalization and total trading volume.

To measure the liquidity of corporate bonds and equities in the historical setting, I hand collect monthly quotations and volume data for NYSE listed securities from the Bank and Quotation section of the *Commercial & Financial Chronicle*. The coverage of this publication is comprehensive; it includes end of month quotations and transaction volume for all NYSE listed equities and corporate bonds. This data is used to compute yearly quoted spread measures by averaging the monthly quotations. I was unable to obtain high frequency transaction level data for this time period, which prevents me from computing effective spread measures.¹¹ At the time I collected the data, digitized copies of the publication were publicly available from 1895 until the end of 1921 at the HathiTrust Digital Library.¹² Given the costs involved, I decided to limit data collection to five years. I chose the most recent years available since this is where the cross-sectional coverage is the most comprehensive.

¹¹Interestingly, a record of all transactions in US equities and corporate bonds as well as daily bid and ask quotations were published by Francis Emory Fitch. These records are available at the archive of the NYSE, but I was unable to access and digitize a comprehensive sample.

¹²<https://www.hathitrust.org/>

To obtain data on securities characteristics, I hand collected data from the *Moody's manuals*. The *Moody's manuals* contain useful information on bond characteristics, such as issuance and maturity date, credit ratings, embedded options, as well as information on the outstanding amount of stocks and bonds. Moreover, the *Moody's manuals* contain a list of all the securities issued by a given firm, which streamlines the identification of firms that have both an issue of common stock and at least one issue of corporate bond outstanding.

To be included in my samples in a given year, securities must go through to the same filters that were applied to Israeli and modern US securities.¹³ The final sample is composed of 76 firms that collectively have 188 corporate bonds outstanding. The equity sample contains 3,728 monthly quotations, while the final bond sample has 7,522.

The last two columns of Table 1 present descriptive statistics for my historical US sample. The quantities expressed in dollars have been CPI adjusted by a factor of 16 to convert them in 2015 USD. Consistent with the Israeli findings, the first row of the table shows that exchange traded US corporate bonds used to have smaller transaction costs than equities: the average quoted spread was 2.09% for corporate bonds compared to 2.72% for equities. This is despite the fact that the US corporate bond market was thin compared to the equity market: the average stock (bond) had a market capitalization of 982 (368) million dollars, a daily trading volume of 3.10 (0.07) million dollars, and an annual turnover of 112% (6%).

While US corporate bonds in the early 20th century had characteristics similar to their modern counterparts, a few differences are noteworthy. First, bonds typically had longer time to maturity: the average time to maturity at issuance used to be 51 years compared to 12 years for the modern US bonds. In addition, a large fraction of the bonds were secured, which probably explains the large fraction of investment grade bonds. In the era of the gold standard, 84% of the bonds were carrying gold clauses. This contractual feature protected investors against the inflation that would follow a depreciation of the dollar against gold. In this sense, these clauses were similar in spirit

¹³I require securities to be covered by the *Moody's manuals* (instead of CRSP, Mergent FISD, or Bloomberg). To compute the market capitalization and volume filters, the dollar amounts have been CPI adjusted by a factor of 16 to convert them in 2015 USD. Instead of requiring 40 daily midpoint returns, I require 7 monthly midpoint returns. Otherwise, the sample selection procedure is the same as for the modern US securities and the Israeli securities.

to modern CPI indexation clauses. See Edwards et al. (2015) for a detailed account of the history of gold clauses.

4 Counterfactual Transaction Cost Estimation

The observations reported in Table 1 suggest that corporate bonds trading in a centralized venue have lower transaction costs than equities on average, while the opposite holds when bonds are trading OTC. In this section, I show that this observation still holds after controlling for known determinants of transaction costs and bond characteristics. The regression results suggest that, conditional on volume and volatility, corporate bonds have similar or lower transaction costs than equities when they both trade on limit order books. I use this result to infer counterfactual transaction cost estimates for US corporate bond based on the transaction costs of US equities.

4.1 Transaction Costs Analysis

Table 2 reports the results of quoted spread regressions on trading volume, return volatility, and corporate bond indicator variables. In each specification, the continuous variables are log-transformed. The logarithm of volume and volatility are demeaned by their average across corporate bonds. Standard errors are clustered at the firm level.

Specification (1) includes trading volume, return volatility, and a corporate bond indicator as independent variables. The results for the modern US sample suggest that, even after controlling for differences in trading activity and volatility, OTC traded corporate bonds are more expensive to trade than equities. The coefficient on the bond indicator is 1.88 and statistically significant at the 1% level. This implies that, conditional on having the same trading volume and volatility, the transaction costs of OTC traded US corporate bonds are 555% larger (6.55 times larger) than the transaction costs of exchange traded US equities.¹⁴

On the contrary, the results for the Israeli and the historical US sample suggest that, conditional

¹⁴The average effect is given by $(e^{1.88} - 1) \cdot 100\% = 555\%$.

on volume and volatility, exchange traded corporate bonds have similar or lower transaction costs than equities. In the Israeli sample, the coefficient on the bond indicator is 0.03, implying that bond transaction costs are on average 3% larger than equity transaction costs once we control for volume and volatility. That coefficient, is not statistically different from zero. In the historical sample, this coefficient is -0.89 and significant at the 1% level, implying that bond transaction costs are on average 59% lower for bonds.

To investigate how the previous results relate to bond characteristics, specification (2) replaces the corporate bond indicator variable with a set of bond characteristic indicators. The estimates show that the previous results are consistent across all credit ratings, maturities, and contractual features. In the modern US sample, the coefficients on bond characteristics are all positive and statistically significant, with the exception of the coefficients on the callable and secured indicators which are not significant. In the Israeli and historical US sample, the coefficients are all negative or statistically insignificant.

In order to test if the elasticities of transaction costs to volume and volatility are different between bonds and stocks, specification (3) adds bond-volume and bond-volatility interaction terms to the regression model. The results show that, when bonds trade on exchange, these elasticities are similar between the two asset classes. For both the Israeli and historical US sample, the coefficients on bond-volume (0.05 for Israel and 0.03 for the US) and bond-volatility (0.04 for Israel and -0.07 for the US) are small relative to the coefficient on volume (-0.55 for Israel and -0.43 for the US) and volatility (0.60 for Israel and 0.46 for the US). None of the interaction terms are statistically different from zero. In terms of magnitudes, the coefficients on volume and bond-volume for the historical US sample imply that a 10% increase in trading volume is associated with a decrease of 3.49% in equity transaction costs and a decrease of 3.30% in bond transaction costs.

In contrast, the results from the modern US sample suggest that the elasticity of transaction costs to trading volume is lower for OTC corporate bonds than for centralized equities. The coefficient on volume is -0.39 and the coefficient on bond-volume is 0.17. These coefficients imply

that a 10% increase in trading volume is associated with a 3.2% decline in transaction costs for equities and a 1.9% decrease in transaction costs for bonds. This evidence suggests that volume has a larger impact on transaction costs when trades are centralized. This supports the view that, when investors need to trade often, consolidating the market eliminates more frictions. The coefficient on the bond-volatility interaction is -0.09 and marginally statistically significant. Unlike the previous results, this estimate is not robust to the choice of the transaction cost metrics, as shown in Table 3.

In Table 3, I repeat the previous analysis for the Israeli and the modern US sample using effective spreads as transaction cost measure. With the exception of the coefficient on bond-volatility in specification 3 of the modern US sample which is now positive and statistically insignificant, all the other results are qualitatively similar.

To summarize, the evidence suggests that, after controlling for differences in volume and volatility, corporate bonds trading in a centralized venue have similar or even lower transaction costs than equities. The results are consistent across bonds with different characteristics. Moreover, the elasticities of transaction costs to volume and volatility are similar between bonds and equities when they trade on a centralized venue. The next subsection makes use of these observations to build counterfactual estimates of what would be the transaction costs of US corporate bonds if trading was centralized.

4.2 Equity-implied transaction costs

Based on the evidence presented so far, I propose to infer the transaction costs that corporate bonds would have on a centralized venue from the transaction costs of equities. The estimation of equity-implied bond transaction costs is done in two steps. The first step consists of regressing equity transaction costs (either quoted spread or effective spread) on trading volume and volatility

$$\log(S_{eit}) = \beta_0 + \beta_1 \log(\text{volume}_{eit}) + \beta_2 \log(\text{volatility}_{eit}) + \epsilon_{eit}, \quad (4)$$

where S_{eit} is the equity transaction cost of firm i in year t . The fitted values of this regression can be used to infer what would be the transaction cost of a stock having the same volume and volatility as a given bond. Under the assumption that bonds and stocks have the same expected transaction costs conditional on volume and volatility, the equity-implied estimates are given by

$$\log(\hat{S}_{bjt}) = \hat{\beta}_0 + \hat{\beta}_1 \log(\text{volume}_{bjt}) + \hat{\beta}_2 \log(\text{volatility}_{bjt}), \quad (5)$$

where \hat{S}_{bjt} denotes the counterfactual transaction cost estimates for bond j in year t , and volume_{bjt} and volatility_{bjt} are the volume and volatility that bond j would have on the centralized venue in year t .

Two remarks are in order. First, we should suspect the estimator to be upward biased. The results of the previous section suggest that, after controlling for differences in volume and volatility, corporate bonds in the centralized Israeli and historical US markets have similar or lower transaction costs than equities. Second, the estimator requires knowledge of the volume and volatility that a given bond would have on a centralized venue. This is challenging to obtain since all we currently observe in the US are the corresponding decentralized quantities. I address this issue indirectly. I argue, based on the experience of the 1997 market reforms of the Nasdaq, that applying the estimator using the volume and volatility measured OTC provides an upper bound \bar{S}_{bjt} on the proper equity-implied estimate \hat{S}_{bjt} .

In early 1997, the SEC implemented several new regulations that significantly changed how the Nasdaq processed orders. Before the introduction of the new rules, the Nasdaq had operated as a fragmented dealer market, akin to the current US corporate bond market in many respects. For example, investors were unable to compete directly with market makers by posting limit orders, and the best available prices were not always publicly accessible. Under the new rules, dealers had to display customer limit orders, and the best available prices were made available to all market participants (see Barclay et al. (1999) for a detailed discussion of the reform).

Weston (2000) shows that following these policy changes, transaction costs decreased by 30% and trading volume increased by 30%. Moreover, Sapp and Yan (2003) show that the volatility

of Nasdaq stocks decreased following the reforms. Given the similarities between the reforms implemented on the Nasdaq and the ones proposed for the bond market, moving bond trading to centralized venues would likely increase trading activity and decrease volatility as well. Since the coefficient on volume, $\hat{\beta}_1$, is negative and the coefficient on volatility, $\hat{\beta}_2$, is positive, applying the estimator using decentralized measures of volume and volatility provides an upper bound \bar{S}_{bjt} on \hat{S}_{bjt} .

Before applying the estimator to the modern US bond market, I test its performance in the Israeli and historical US samples. In both markets, volume and volatility are measured on centralized venues, and the proper equity implied estimates can be compared against measured transaction costs to test the accuracy of the estimator. As expected, the coefficients obtained in the first step of the estimation procedure are nearly identical to those reported in Tables 2 and 3, thus I do not report them separately.

Figures 1 and 2 illustrate the performance of the estimator in the cross-section of bonds. Each figure plots measured transaction costs on the y-axis against equity implied estimates on the x-axis; each point represents a bond-year. For bonds lying under the 45-degree line, equity-implied estimates overshoot measured transaction costs, and vice versa. Each figure reports the cross-sectional quartiles of the relative error (RE) defined as

$$RE = \frac{\hat{S}_{bjt} - S_{bjt}}{S_{bjt}}, \quad (6)$$

where \hat{S}_{bjt} is the equity-implied estimate of bond j in year t , and S_{bjt} is the measured transaction cost of the same bond.

Figure 1 displays the Israeli results for quoted and effective spreads. Both graphs indicate that the estimator performs quite well in the Israeli sample. Measured and predicted spreads are highly correlated, with a correlation coefficient of 83% for quoted spreads and 72% for effective spreads. The median relative error is 12% for quoted spreads, and 22% for effective spreads. Hence the equity implied estimator tends to overshoot measured transaction costs in the Israeli sample.

Figure 2 displays the estimator's performance for quoted bond spreads in the historical US

sample. While measured and predicted transaction costs are still highly correlated, with a correlation coefficient of 71%, the figure shows that the equity-implied estimates overshoot measurements substantially. The median relative error is 149%, meaning that equity implied estimates for the median bond is more than twice as large as the measured transaction cost. Together, these results confirm that the equity-implied estimator tends to overestimate the actual transaction costs of corporate bonds on centralized venues.

I now apply the equity-implied estimator to the modern US bond market. Again, the coefficient obtained in the first step of the estimation procedure are virtually identical to those reported in Tables 2 and 3, thus I do not report them separately. Figures 3 illustrates the results of the second step of the estimation. In this figure, the x-axis represent the equity-implied estimate obtained using OTC measures of volume and volatility, and the y-axis displays the transaction costs measured OTC. The figure reports the cross-sectional quartiles of the predicted percent change (PC) defined as

$$PC = \frac{\bar{S}_{bjt} - S_{bjt}}{S_{bjt}}, \quad (7)$$

where \bar{S}_{bjt} is the equity-implied estimate of bond j in year t , and S_{bjt} is the OTC transaction cost for the same bond.

According to the estimates, moving to a centralized market structure would reduce the median quoted spread by 74%, and the median effective spread by 63%. The quoted spread of all bonds is predicted to decrease while the effective spread of 94% of the bonds is expected to decrease. Recall that those are conservative estimates of the transaction decrease since the equity-implied estimator tends to overestimate actual transaction costs.

Table 4 reports descriptive statistics for the projected transaction costs savings. In this table, the bond-year estimates and measurements of each bond have been averaged over the sample period. We see that quoted spreads would on average decrease by 43bp, from 56bp to 13bp, and effective spreads would decrease by 16bp, from 23bp to 7bp.

To get a sense of the dollar value of the transaction cost savings implied by these estimates,

Table 4 also reports statistics on dollar annual savings (DAS) defined as

$$\text{ADS}_{bj} = \text{annual volume}_{bj} \left(\frac{\text{measured ES}_{bj} - \text{predicted ES}_{bj}}{2} \right), \quad (8)$$

and on relative annual savings (RAS) defined as

$$\text{ARS}_{bj} = \frac{\text{ADS}_{bj}}{\text{total par value}_{bj}}, \quad (9)$$

where $\text{annual volume}_{bj}$ is a bond j 's average annual OTC trading volume over the sample period, measured ES_{bj} is the bond's average OTC effective spread, predicted ES_{bj} is the bond's average equity-implied effective spread, and $\text{total par value}_{bj}$ is the bond's average outstanding total par value. The table also reports the present value of the transaction cost savings over a bond's life, $\text{PV}(\text{DAS})$ and $\text{PV}(\text{RAS})$, defined as

$$\text{PV}(\text{ADS})_{bj} = \left(\frac{1 - e^{-rT}}{r} \right) \text{DAS}_{bj}, \quad (10)$$

and

$$\text{PV}(\text{ARS})_{bj} = \left(\frac{1 - e^{-rT}}{r} \right) \text{RAS}_{bj}, \quad (11)$$

where T is the time to maturity at issuance of the bond, and r is the risk-free rate, which I assume equal to 2%.

The table shows that, for the average bond, annual transaction cost savings amount to \$0.42 million, or 6 basis point of the total par value. Over a bond's life, the present value of the projected transaction cost savings is \$4.74 million, or 66 basis points of total par value. Aggregating the savings over the sample, I obtain that a centralized market structure would generate transaction cost savings of \$681 million per year. This amount represents 6 bp of the 1.2 trillion in bond par value. The present value of the aggregate savings over the life of those 1,642 bonds is \$6.7 billion or 66 bp of total par value. Table 4 also reveals that there is substantial cross-sectional variation in the predicted reduction in transaction costs. For example, the effective spread of the bond at the 5th percentile would decrease by only 1 bp, while the effective spread of the bond at the 95th

percentile would decrease by 54 bp.

Table 5 reports the regression results of transaction cost savings on various bond characteristics. In each regression the logarithm of total par value and year to maturity have been demeaned. Standard errors are clustered at the issuer level.

The dependent variables in the first two columns of Table 5 are respectively the percentage change in quoted and effective spreads, as defined in equation 7. The intercept -0.72 in the first column means that, if trading were to migrate to centralized venues, the quoted spread of an unsecured and non-callable investment grade bond of average market capitalization and average time to maturity would decrease by 72% on average. The intercept -0.47 in the second column has a similar interpretation for effective spreads. Both intercepts are statistically significant.

High yields bond would experience even larger reduction in transaction costs. In the first column, the coefficient -0.11 on the high yield dummy implies that the expected decrease in quoted spread is 11% larger for high yield bonds than for investment grade bonds. For effective spreads, the corresponding coefficient is -0.21. Both coefficients are statistically significant. Similarly, issues with longer time to maturity would experience larger decrease in transaction costs. In the first column, the coefficient on the logarithm of year to maturity is -0.02, which implies that a 1% increase in year to maturity implies yields an additional 0.02% decrease in quoted spread. The corresponding coefficient for effective spread is -0.01. Only the coefficient on quoted spread is statistically significant. Larger issues would also experience larger decrease in their effective spread; the coefficient -0.18 implies that a 1% increase on total par value yields an additional 0.18% decrease in effective spread. The coefficient is statistically significant. The corresponding coefficient for quoted spread is close to zero and not statistically significant. Lastly, the coefficients on the callable and secured dummies are insignificant.

The dependent variable in the third column is the logarithm of annual dollar savings (ADS). The results show that the projected annual transaction cost savings are greater for large high yield issues with longer time to maturity. The coefficient 1.26 on total par value implies that a 1% increase in total par value is associated with an increase in annual transaction cost savings of

2.52%.¹⁵ Similarly, the coefficient 0.49 on year to maturity implies that a 1% increase in year to maturity yields an increase of 0.63% in annual transaction cost savings. The coefficient 1.36 on the high yield dummy implies that high yield bonds would experience annual transaction savings 290% greater than investment grade bonds, keeping everything else constant. The coefficients on the callable and secured dummies are here again insignificant. Unreported regressions of the other transaction cost saving metrics (ARS, PV(ADS), and PV(ARS)) yield qualitatively similar results.

In this section, we have seen that the transaction costs of corporate bonds would decline substantially if corporate bonds trading migrated from the current decentralized OTC market structure to a centralized market structure. The change would generate significant transaction cost savings for investors, particularly for large speculative issues with long time to maturity. While interesting in themselves, these results do not directly address the central question of this paper: would a centralized market structure be more efficient? Indeed, transaction costs are not deadweight losses, but rather a transfer of wealth between liquidity providers and liquidity seekers.

In the next section, I build a structural model to evaluate the welfare implications of a change in market structure. Interestingly, the model delivers sufficient statistics for the relative efficiency of the two market structures that connect welfare to transaction costs. The empirical work done in this section turns out to be useful for welfare calculations.

5 Model

Time is continuous and runs forever, $t \in [0, \infty)$. The economy features two types of agent, investors and market-makers. All agents are risk-neutral and infinitely lived, with time preference determined by a constant discount rate $r > 0$. Agents have access to a bond, available in fixed supply x . The bond pays a coupon at rate c until it matures, at which point the bond holder receives principal p . Maturity occurs randomly at Poisson rate η . The bond either trades on a centralized market structure, modeled as a central limit order book, or a decentralized market structure, modeled as a dealer market operating on a request for quote basis. The two market

¹⁵The marginal effect of one percent increase in total par value is given by $(e^{1.26} - 1) \% = 2.52\%$.

structures are subject to imperfect trading processes that impose costs and delays on market participants. In both market-structures, market-makers are necessary to intermediate trades; they act as passthrough intermediaries and never hold the bond in inventory.

5.1 Investors and Market Makers

Investors are characterized by their holding of the bond, restricted to either 0 or 1, and their intrinsic liquidity type that is either "h" or "l". Type l owners bear a holding cost δ per unit of time, whereas type h owners do not. Type h agents switch to type l at Poisson rate ξ . The transition process are independent across investors. Once an investor is hit by a "liquidity shock", he remains of type l forever. While I do not explicitly model the nature of these liquidity shocks in this paper, there are many possible ways to interpret them. For instance, they might represent changes in investor's subjective valuation of the asset, a sudden need for cash, changes in hedging need, or any conceivable event that lead investors to reallocate their portfolio.

The full set of investor types is then $\mathcal{T} = \{h0, h1, l0, l1\}$, where "h" and "l" designate the investor's subjective valuation of the asset, and "0" and "1" represent the quantity of asset owned. There is a continuum of non-atomistic agents; the measure of investor of type $\kappa \in \mathcal{T}$ at time t is denoted by $\mu_{\kappa,t}$. I assume throughout the paper that there is, at all times, an infinite measure of type h0, $\mu_{h0,t} = \infty$. There are gain from trade between investors of type h0 and type l1, however trade execution is not instantaneous and participating in the market involves flow costs.

The flow participation cost of buyers (sellers) in market structure $k \in \{c, d\}$ is denoted by χ_b^k (χ_s^k), where the letter "c" designates a centralized market structure, and the letter "d" designates a decentralized market structure. These parameters stand for the opportunity cost of investor's time, and for the inconvenience of executing a trade: in a decentralized venue, these costs account for the effort exerted to contact and negotiate with dealers, while in a centralized venue, they represent the cost of monitoring the market and splitting orders to ensure fair execution. Participation in the market is voluntary. Investors dynamically choose, at each instant, whether to participate or stay out of the market. They may suspend and resume their market activity at no cost.

Market makers face a similar participation decision. Active market makers have the opportunity to intermediate trades and earn the bid-ask spread, however, market-making involves a flow cost χ_m^k . This parameter stands for the labor and capital costs associated with market-making (hiring traders, financing the balance sheet, IT investment, the cost of settling and clearing trades.) In a decentralized venue, this parameter also embeds the cost of maintaining relationships with customers. In a centralized venue, it accounts for the cost of continuously monitoring the market and updating quotations. There is an infinite measure of potential entrant in the market-making industry. The market structure specifies the trading procedure by which market makers intermediate transactions.

5.2 Market Structures

In the decentralized market structure, the asset trades on a request for quote basis. Investors privately contact market-makers and negotiate over the price – i.e., market-makers do not post firm quotes publicly. In a centralized market structure, however, the quotations of market-makers are publicly disseminated in a central limit order book. I assume that, under both market structures, technological constraints prevent trades from being executed instantly. In practice, there is a variety of reasons that can justify this assumption.

In decentralized market, delays can occur because investors need time to contact and negotiate with dealers, and dealers need time to find suitable counterparties. In a centralized market, delays may happen when market depth is insufficient such that orders must be split to prevent front running, or when quotes are stale at the time liquidity is needed. In both market structures, market participants might need to gather and process information about the asset before committing to a transaction, (Abel et al., 2013). There may also be technological and institutional constraints that limit the rate at which market-makers can clear and settle trades.

In this paper, I abstract from modeling the exact source of trading delays. Instead, these frictions are subsumed under trade processing functions. Formally, suppose that at time t there is a measure b_t of buyers, a measure m_t of market-makers, and a measure s_t of sellers. I assume that

the rate at which trades are executed on the venue $k \in \{c, d\}$ is equal to

$$g_t^k = g^k(b_t, m_t, s_t) = \lambda^k b_t^{\alpha^k} m_t^{\beta^k} s_t^{1-\alpha^k-\beta^k},$$

a constant return to scale Cobb-Douglas function.¹⁶ For convenience, denote the the execution rate of a market participant of type $\phi \in \{b, m, s\}$ on market structure k by $\rho_{\phi_t}^k = \frac{g_t^k}{\phi_t}$.

The law of motion of the measure of low valuation owners in market structure k follows

$$\dot{\mu}_{l1,t}^k = -g_t^k + \xi \mu_{h1,t}^k, \quad (12)$$

where the first term on the right-hand side accounts for the flow g_t^k of type-l1 investors who sell their asset, and the second term reflects that type-h1 are hit by liquidity shocks at rate ξ . Combining the last equation with the market clearing condition

$$\mu_{h1,t}^k + \mu_{l1,t}^k = x, \quad (13)$$

delivers the law of motion of the measure of high valuation owners $\mu_{h1,t}^k$.

5.3 Decentralized Equilibrium

In the decentralized market structure, investors and market makers first make their entry decision, and later negotiate on the price. Denote by $V_{\kappa,t}^d$ the value function of an investor of type $\kappa \in \mathcal{T}$ when the bond trades OTC. The value function of a potential buyer, $V_{h0,t}^d$, satisfies the HJB equation

$$rV_{h0,t}^d = \dot{V}_{h0,t}^d + \max \left\{ 0, \rho_{b_t}^d \left(V_{h1,t}^d - V_{h0,t}^d - A_t^d \right) - \chi_b^d \right\} + \xi \left(V_{l0,t}^d - V_{h0,t}^d \right), \quad (14)$$

where A_t^d is the ask price a buyer can negotiate at time t . This equation implies that a buyer benefits from participating whenever the rate of execution of his order, $\rho_{b_t}^d$, times the private surplus he obtains when he purchases the asset, $\left(V_{h1,t}^d - V_{h0,t}^d - A_t^d \right)$, is greater than the flow participation

¹⁶As Lester et al. (2015) write, "Whether or not the order execution technology has constant return to scale remains an open question in the context of asset markets. Unfortunately, in contrast to the labor market—where reliable data exist for the number of unemployed workers, the number of vacancies, and the number of matches that form in a particular labor market—the analogous data for financial markets is elusive."

cost, χ_b^d . Free entry, combined with the assumption that there is an infinite measure of potential buyers, implies that in equilibrium $\rho_{b_t}^d (V_{h1,t}^d - V_{h0,t}^d - A_t^d) \leq \chi_b^d$, otherwise additional buyers would enter the market. This last expression holds with equality whenever $b_t > 0$, otherwise some buyers would strictly prefer to leave the market.

The value functions of the other types of investor satisfy the following HJB equations

$$rV_{h1,t}^d = \dot{V}_{h1,t}^d + \xi (V_{l1,t}^d - V_{h1,t}^d) + \eta (p + V_{h0,t}^d - V_{h1,t}^d) + c, \quad (15)$$

$$rV_{l1,t}^d = \dot{V}_{l1,t}^d + \max \left\{ 0, \rho_{s_t}^d (V_{l0,t}^d - V_{l1,t}^d + B_t^d) - \chi_s^d \right\} + \eta (p + V_{l0,t}^d - V_{l1,t}^d) + c - \delta, \quad (16)$$

$$V_{l0,t}^d = 0. \quad (17)$$

Intuitively, equation (15) reflects that high valuation owners receive the coupon c until they are hit by a liquidity shock, which occurs at rate ξ , or until the asset matures, which occurs at rate η . Equation (16) shows that, once hit by a liquidity shock, investors receive a reduced utility flow $c - \delta$ from owning the asset and choose between becoming active sellers or staying out of the market. Note that the measure s_t of active seller is bounded above by the measure of low-valuation owners, $s_t \in [0, \mu_{l1,t}^d]$. Equation (17) reflects that low-valuation investors leave the market forever once they liquidate their holdings. For future reference, let $\Delta V_{i,t}^k \equiv V_{i1,t}^k - V_{i0,t}^k$ denote the reservation value of type $i \in \{h, l\}$ in market structure $k \in \{c, d\}$, and let $\Delta_{h,l}^k \equiv \Delta V_{h,t}^k - \Delta V_{l,t}^k$ denote the total surplus a trade generates.

The value function of market makers in a decentralized market must satisfy

$$rV_{m,t}^d = \dot{V}_{m,t}^d + \max \left\{ 0, \rho_{m_t}^d (A_t^d - B_t^d) - \chi_m^d \right\}, \quad (18)$$

which implies that market makers choose to intermediate trades whenever the rate at which they complete transactions $\rho_{m_t}^d$ times the spread they earn in the process $(A_t^d - B_t^d)$ is large enough to compensate them for the cost χ_m^d they incur. The zero-profit condition of market makers is $\rho_{m_t}^d (A_t^d - B_t^d) \leq \chi_m^d$, with equality whenever $m_t > 0$.

Active market makers do not publicly commit to firm quotes, they instead negotiate the price with investors when they are contacted. At the outcome of the bargaining process, buyers obtain

a share θ_b of the surplus, market makers a share θ_m , and sellers a share θ_s , with the restriction that $\theta_b + \theta_m + \theta_s = 1$. The resulting ask and bid prices are

$$A_t^d = (1 - \theta_b) \Delta V_{h,t}^d + \theta_b \Delta V_{l,t}^d, \quad (19)$$

$$B_t^d = (1 - \theta_s) \Delta V_{l,t}^d + \theta_s \Delta V_{h,t}^d, \quad (20)$$

which implies that the market maker commission is indeed a share θ_m of the surplus, $A^d - B^d = \theta_m \Delta_{hl}^d$. We now define and characterize a competitive equilibrium in a decentralized market structure.

Definition 1. *Given an initial asset allocation, $\{\mu_{l1,0}^d, \mu_{h1,0}^d\}$, a decentralized equilibrium is a list composed of asset allocation processes, $\{\mu_{l1,t}^d, \mu_{h1,t}^d\}_{t \geq 0}$; processes for the measures of buyers, market makers, and sellers $\{b_t, m_t, s_t\}_{t \geq 0}$; a collection of value functions, $\{V_{h0,t}^d, V_{h1,t}^d, V_{l1,t}^d, V_{l0,t}^d, V_{m,t}^d\}_{t \geq 0}$; and processes for the bid and ask prices, $\{B_t^d, A_t^d\}_{t \geq 0}$, satisfying (12)–(20) for all $t \geq 0$.*

Let's first note that an inactive market, where $b_t = m_t = s_t = 0$ for all $t \geq 0$, can always be supported in equilibrium. It is not optimal for buyers and market makers to participate if sellers stay out of the market; it is not optimal for sellers to participate if buyers or market makers stay out of the market. Such a self fulling equilibrium equilibrium is sustainable for all parameter values, and its characterization is trivial. For an equilibrium with active market to exist, the following condition must be satisfied

$$\frac{\delta}{r + \xi} \geq \frac{1}{g^d \left(\frac{\theta_b}{\chi_b^d}, \frac{\theta_m}{\chi_m^d}, \frac{\theta_s}{\chi_s^d} \right)}. \quad (21)$$

I show in the Appendix that this condition ensures that low valuation owners have sufficient incentive to participate in the market given the costs involved. The following proposition characterizes the equilibrium. The proof is relegated to the Appendix.

Proposition 1. *Provided that condition (21) is satisfied, there exists an equilibrium with the following properties. All type l1 investors participate in the market, $s_t = \mu_{l1,t}^d$ for all $t \geq 0$. The*

surplus is constant, $\Delta_{hl,t}^d = \Delta_{hl}^d$, and given by the unique positive solution to

$$(r + \eta + \xi) \Delta_{hl}^d + \chi_s^d \left[g^d \left(\frac{\theta_b}{\chi_b^d}, \frac{\theta_m}{\chi_m^d}, \frac{\theta_s}{\chi_s^d} \right) \Delta_{hl}^d \right]^{\frac{1}{1-\alpha^d-\beta^d}} = \delta + \chi_s^d. \quad (22)$$

The buyer-to-seller ratio and market maker-to-seller ratio are constant over time, which in turn implies that execution rates are constant as well, $\rho_{\phi_t}^d = \rho_{\phi}^d$ for all $\phi \in \{b, m, s\}$. The bid and ask prices are also constant, $A_t^d = A^d$ and $B_t^d = B^d$, and satisfy

$$\begin{aligned} A^d &= \frac{c}{r + \eta} + \frac{\eta p}{r + \eta} - \left(\frac{\xi}{r + \eta} + \theta_b \right) \Delta_{hl}^d, \\ B^d &= \frac{c}{r + \eta} + \frac{\eta p}{r + \eta} - \left(\frac{\xi}{r + \eta} + \theta_b + \theta_d \right) \Delta_{hl}^d, \end{aligned}$$

and the bid-ask spread is $S^d = A^d - B^d = \theta_m \Delta_{hl}^d$. The equilibrium measure of type-1 investor as a function of time is

$$\mu_{l1,t}^d = \left(\mu_{l1,0}^d - \frac{\xi x}{\xi + \rho_s^d} \right) e^{-(\xi + \rho_s^d)t} + \frac{\xi x}{\xi + \rho_s^d}, \quad (23)$$

thus the steady state level of misallocated asset is $\mu_{l1,\infty}^d = \frac{\xi x}{\xi + \rho_s^d}$.

The fact that execution rates are constant over time is a direct consequence of free entry and a constant return to scale trading technology. The prices are equal to the present value of the coupon and principal, $\frac{c}{r+\eta} + \frac{\eta p}{r+\eta}$, minus a liquidity discount. The discount is function of the frequency at which investor require liquidity, ξ ; the bargaining power of the participants, θ_b and θ_d ; and the difference between buyer and seller reservation values, Δ_{hl}^d , which summarizes the severity of the trading frictions. Intuitively, following a liquidity shock, an owner's subjective valuation decrease by more when selling the asset is difficult.

The comparative statics of the model are intuitive. The surplus Δ_{hl}^d is increasing in the holding cost δ and in the participation costs χ_b^d , χ_m^d , and χ_s^d ; but decreasing in the efficiency of the trade execution technology λ : following a liquidity shock, an owner's subjective valuation of the asset decrease by more when the holding cost is high, when participation in the market is costly, and when the trade execution technology is less efficient. The surplus is decreasing in ξ and η since

the gain from trade is lower when buyers are hit by liquidity shocks at a higher rate, and when the asset is expected to mature in the near future. The comparative statics on bargaining powers and the shares of the order processing function are ambiguous. These results were obtained by implicitly differentiating equation (22). Equation (23) implies that, in steady state, the measure of misallocated asset increases with the frequency at which liquidity shocks occur, ξ , and decreases with the speed at which seller trade, ρ_s^d .

I conclude this section by defining and characterizing the social welfare of this economy. In this risk-neutral framework, social welfare is the present value of the bond payoff net of holding costs and participation costs,

$$W^d = E_\tau \left[\int_0^\tau e^{-rt} (xc - \delta\mu_{l1,t}^d - \chi_b^d b_t - \chi_m^d m_t - \chi_s^d s_t) dt + e^{-r\tau} xp \right], \quad (24)$$

where $\tau \sim \exp(\eta)$ corresponds the time at which the bond matures. Using the results in Proposition 1, this integral can be calculated explicitly

$$W^d = W^* - \left(\mu_{l1,0}^d + \frac{\xi x}{r + \eta} \right) \Delta_{hl}^d, \quad (25)$$

where $W^* = \frac{xc}{r+\eta} + \frac{\eta xp}{r+\eta}$ corresponds to the welfare achieved in a frictionless economy.

Thus the welfare loss associated with an OTC market structure, $W^* - W^d$, is composed of two terms: the first one, $\mu_{l1,0}^d \Delta_{hl}^d$, accounts for the initial misallocation of the asset; the second one, $\frac{\xi x}{r+\eta} \Delta_{hl}^d$, is the discounted value of all future utility lost to trading frictions. For both terms, the magnitude of the welfare loss is proportional to the surplus Δ_{hl}^d . Interestingly, most of the model's parameters do not explicitly enter the welfare expression. To quantify the welfare loss caused by trading frictions, it is sufficient to measure the difference in subjective valuation between buyers and sellers; individually estimating the deep structural parameters that generates these frictions is not necessary. In other words, the total surplus is a sufficient statistic for the impact of trading frictions on welfare.

5.4 Centralized Equilibrium

In the centralized market structure, market makers compete for order flow by posting firm bid and ask quotations on central limit order books (CLOB). There are potentially several CLOB, or market centers, in operation. Each market center publicly states a contract $\sigma_t = (A_t, B_t)$ that specifies the ask price A_t and the bid price B_t at which it will accept limit orders from market makers. Denote by $\Sigma_t = \mathbb{R}^+ \times \mathbb{R}^+$ the set of feasible contracts at time t and by Σ_t^* the set of contracts active in equilibrium at time t – i.e., the set of contracts where at least some market makers send limit orders. Investors observe the set of active contracts and decide at which venue to send their market orders.

Denote the measures of active buyers, market makers, and sellers on a contract $\sigma_t \in \Sigma_t$ by $b_t(\sigma_t)$, $m_t(\sigma_t)$, and $s_t(\sigma_t)$, respectively. Correspondingly, denote the aggregate trade execution rate for this contract by $g_t^c(\sigma_t) \equiv g_t^c(b_t(\sigma_t), m_t(\sigma_t), s_t(\sigma_t))$, such that the execution rates experienced by market participants are $\rho_{\phi_t}^c(\sigma_t) = \frac{g_t^c(\sigma_t)}{\phi_t(\sigma_t)}$, where $\phi \in \{b, m, s\}$. The value function of a potential buyer must then satisfy

$$rV_{h0,t}^c = \dot{V}_{h0,t}^c + \max_{\sigma_t \in \Sigma_t^*} \left\{ 0, \rho_{b_t}^c(\sigma_t) (V_{h1,t}^c - V_{h0,t}^c - A_t) - \chi_b^c \right\} + \xi (V_{l0,t}^c - V_{h0,t}^c). \quad (26)$$

This equation implies that investors send their orders to the contract that delivers the highest utility flow, provided that this flow is weakly positive. The zero utility condition of buyers implies that $\rho_{b_t}^c(\sigma_t) (V_{h1,t}^c - V_{h0,t}^c - A_t) = \chi_b^c$ for all $\sigma_t \in \Sigma_t^*$. Similarly, the value functions of the other types of investor satisfy

$$rV_{h1,t}^c = \dot{V}_{h1,t}^c + \xi (V_{l1,t}^c - V_{h1,t}^c) + \eta (p + V_{h0,t}^c - V_{h1,t}^c) + c, \quad (27)$$

$$rV_{l1,t}^c = \dot{V}_{l1,t}^c + \max_{\sigma_t \in \Sigma_t^*} \left\{ 0, \rho_{s_t}^c(\sigma_t) (V_{l0,t}^c - V_{l1,t}^c + B_t) - \chi_s^c \right\} + \eta (p + V_{l0,t}^c - V_{l1,t}^c) + c - \delta, \quad (28)$$

$$V_{l0,t}^c = 0, \quad (29)$$

and the value function of market makers satisfies

$$rV_{m,t}^c = \dot{V}_{m,t}^c + \max_{\sigma_t \in \Sigma_t^*} \left\{ 0, \rho_{m_t}^c(\sigma_t)(A_t - B_t) - \chi_m^c \right\}. \quad (30)$$

Market maker's zero-profit condition implies that $\rho_{m_t}^c(\sigma_t)(A_t - B_t) = \chi_m^c$ for all $\sigma_t \in \Sigma_t^*$.

Following Moen (1997), the equilibrium allocation is required to be a no-surplus allocation. Assuming there is active trading, the set of contracts offered in equilibrium must maximize low valuation owners utility subject to the participation constraint of buyers and market makers

$$\begin{aligned} \Sigma_t^* \supseteq \arg \max_{\sigma_t \in \Sigma_t} \quad & \rho_{s_t}^c(\sigma_t) (V_{l0,t}^c - V_{l1,t}^c + B_t) \\ \text{subject to} \quad & \rho_{b_t}^c(\sigma_t) (V_{h1,t}^c - V_{h0,t}^c - A_t) = \chi_b^c, \\ & \rho_{m_t}^c(\sigma_t) (A_t - B_t) = \chi_m^c. \end{aligned} \quad (31)$$

If such a contract was not offered in equilibrium, a market center could propose the value maximizing contract and require a participation fee from sellers. For a small enough participation fee, the market center would manage to attract sellers, and generate profit. Competition among market centers bid this fee down to zero. We now define and characterize a competitive equilibrium on a centralized market structure.

Definition 2. *Given an initial asset allocation, $\{\mu_{l1,0}^c, \mu_{h1,0}^c\}$, a centralized equilibrium is a list composed of asset allocation processes, $\{\mu_{l1,t}^c, \mu_{h1,t}^c\}_{t \geq 0}$; a collection of contracts, $\{\Sigma_t^*\}_{t \geq 0}$; processes for the measures of buyers, market makers, and sellers in each active contract, $\{b_t(\sigma_t), m_t(\sigma_t), s_t(\sigma_t)\}_{t \geq 0}$ for each $\sigma_t \in \Sigma_t^*$; and a collection of value functions, $\{V_{h0,t}^c, V_{h1,t}^c, V_{l1,t}^c, V_{l0,t}^c, V_{m,t}^c\}_{t \geq 0}$ satisfying (12)–(13), and (26)–(31) for all $t \geq 0$.*

As in the decentralized framework, an equilibrium where the centralized market is inactive exists for all parameter values. For an active market equilibrium to exist, the following condition must be satisfied

$$\frac{\delta}{r + \xi} \geq \frac{1}{g^c \left(\frac{\alpha^c}{\chi_b^c}, \frac{\beta^c}{\chi_m^c}, \frac{1 - \alpha^c - \beta^c}{\chi_s^c} \right)}. \quad (32)$$

This condition is analogous to the condition we obtained in the decentralized market, with the exception that the bargaining power parameters in equation (21) have been substituted for the shares of the matching function. I expand on the reason for this similarity below. The following proposition characterizes the active market equilibrium.

Proposition 2. *Provided that condition (32) is satisfied, there exists an equilibrium with the following properties. A single time invariant contract is offered, $\Sigma_t^* = \{\sigma^*\} = \{(A^c, B^c)\}$. All type l1 investors direct their order flow toward that contract, $s_t(\sigma^*) = \mu_{l1,t}^c$ for all $t \geq 0$. The total surplus is constant, $\Delta_{hl,t}^c = \Delta_{hl}^c$, and given by the unique positive solution to*

$$(r + \eta + \xi) \Delta_{hl,t}^c + \chi_s^c \left[g^c \left(\frac{\alpha^c}{\chi_b^c}, \frac{\beta^c}{\chi_m^c}, \frac{1 - \alpha^c - \beta^c}{\chi_s^c} \right) \Delta_{hl,t}^c \right]^{\frac{1}{1 - \alpha^c - \beta^c}} = \delta + \chi_s^c. \quad (33)$$

The buyer-to-seller ratio and the market maker-to-seller ratio are constant, which implies that the trade execution rates are constant as well, $\rho_{\phi_t}^c(\sigma^*) = \rho_{\phi}^c$ for each $\phi \in \{b, m, s\}$. The bid and ask prices satisfy

$$\begin{aligned} A^c &= \frac{c}{r} - \left(\frac{\xi}{r} + \alpha^c \right) \Delta_{hl}^c, \\ B^c &= \frac{c}{r} - \left(\frac{\xi}{r} + \alpha^c + \beta^c \right) \Delta_{hl}^c, \end{aligned}$$

which implies that the bid-ask spread is $S^c = A^c - B^c = \beta^c \Delta_{hl}^c$. The equilibrium measure of type l1 investor as a function of time is

$$\mu_{l1,t}^c = \left(\mu_{l1,0}^c - \frac{\xi x}{\xi + \rho_s^c} \right) e^{-(\xi + \rho_s^c)t} + \frac{\xi x}{\xi + \rho_s^c}, \quad (34)$$

thus the steady state level of misallocated asset is $\mu_{l1,\infty}^c = \frac{\xi x}{\xi + \rho_s^c}$.

I show in the Appendix that there is a unique contract that maximizes low valuation owners utility. This contract is the only one offered in equilibrium since any other contract would fail to attract seller's order flow. The remainder of Proposition 2 mirrors Proposition 1 with the difference that the bargaining power parameters have been substituted for the corresponding shares of the matching function. This is because, in the centralized venue, the shares of the matching function

endogenously determine surplus allocation among market participants. For example, market maker commission in the decentralized market structure is $S^d = \theta_d \Delta_{hl}^d$ while their commission is $S^c = \beta^c \Delta_{hl}^c$ in the centralized market structure. The fact that the share of the matching function determines the surplus allocation among market participants is a standard result of the directed search literature, see Moen (1997) and Lester et al. (2015) for examples.

Welfare in the centralized market structure is defined in a similar way as in the decentralized market structure. A direct calculation shows that

$$W^c = W^* - \left(\mu_{l1,0}^c + \frac{\xi x}{r + \eta} \right) \Delta_{hl}^c \quad (35)$$

where here again $W^* = \frac{xc}{r+\eta} + \frac{\eta xp}{r+\eta}$ corresponds to the welfare achieved in a frictionless economy. This expression is analogous to the one we obtained for the decentralized venue. With these results on hand, we can now discuss the relative efficiency of the two market structures.

5.5 Relative Efficiency

To simplify the exposition, I assume throughout the rest of the paper that $\mu_{l1,0}^d = \mu_{l1,0}^c = 0$. A natural interpretation of this assumption is that when the asset is initially offered to the public at $t = 0$, only high valuation investors purchase it. Given this assumption, the difference in efficiency between the centralized and decentralized market structure is

$$W^c - W^d = \frac{\xi x}{r + \eta} \left(\Delta_{hl}^d - \Delta_{hl}^c \right). \quad (36)$$

This equation implies that the optimal market structure is the one associated with the lowest total surplus Δ_{hl} . This is intuitive since we saw earlier that the difference in valuation between buyers and sellers summarizes the severity of the trading frictions on either market structure. The multiplicative factor $\frac{\xi x}{r+\eta}$ magnifies the social benefit of trading on the optimal venue. While the welfare expression is simple and intuitive, the surpluses are not directly measurable and must be inferred from observable quantities.

In Proposition 1 and Proposition 2, we established that transaction costs on the decentral-

ized and centralized venues are $S^d = \theta_m \Delta_{hl}$ and $S^c = \beta^c \Delta_{hl}$, respectively. Substituting these expressions back in 36 yields

$$W^c - W^d = \frac{\xi x}{r + \eta} \left(\frac{S^d}{\theta_m} - \frac{S^c}{\beta^c} \right). \quad (37)$$

The formula shows that, while transaction costs are informative about the relative efficiency of the two market structures, they must be appropriately scaled by the share of surplus that market makers earn on the venues in order to study welfare. The next section maps the welfare expression to data.

6 Welfare Analysis

Equation (37) highlights that the sign of the term in parenthesis, $\frac{S^d}{\theta_m} - \frac{S^c}{\beta^c}$, determines the optimal market structure. The transaction costs of OTC traded corporate bonds, S^d , can directly be measured in the data. In Section 4, we constructed counterfactual estimates of the transaction costs that bonds would have on limit order markets, S^c . Hence we are left to estimate the parameters θ_m and β^c . To evaluate the magnitude of the welfare gain or loss, we must also estimate the frequency of investor liquidity needs ξ . The other parameters are directly measurable: x is the total supply of asset, r is the risk-free rate, and $\frac{1}{\eta}$ is the time to maturity.

6.1 Estimation of θ_m

I estimate the bargaining power parameter by matching the liquidity premium of corporate bonds reported in Cui et al. (2018). Assuming that buyers and sellers have a similar amount of bargaining power when negotiating with dealers, we have from Proposition 1 that the liquidity premium ϕ^d is given by

$$\phi^d = P^* - M^d = \left(\frac{1}{2} + \frac{\xi}{r + \eta} \right) \frac{A^d - B^d}{\theta_m},$$

Table 6 reports the fit of the model. The bargaining power estimate is $\hat{\theta}_m = 0.73$.

6.2 Estimation of β^c

I estimate the coefficient β^c that applies to US stocks using data on trading volume and book depth. I then show that, in Israel, the estimate for bonds is of the same magnitude as the one for stocks. As a result, the estimate for US stocks is used in the welfare calculations.

Recall that the coefficient β^c is the market maker's share in the matching function,

$$g_t^c = g^c(b_t, m_t, s_t) = \lambda^c b_t^{\alpha^c} m_t^{\beta^c} s_t^{1-\alpha^c-\beta^c},$$

where m_t is the measure of active market makers. Since each market maker posts limit orders to buy and sell exactly one unit of the asset, m_t also corresponds to the quantity of asset available at the bid and ask, also known as the depth of the book. The quantity g_t^c corresponds to trading volume at time t . Thus we have that

$$\frac{\partial \text{Volume}_t}{\partial \text{Depth}_t} = \frac{\partial \log g_t^c}{\partial \log m_t} = \beta^c.$$

Hence β^c is equal to the elasticity of trading volume with respect to the the depth of the book. For a given security i , I estimate β_i^c by regressing daily trading volume on the average depth of the book during the day. Since the time series of volume and depth are non stationary (both increase over time over the sample period), I first difference the data

$$\Delta \log \text{Volume}_{i,t} = \beta_i^c \Delta \log \text{Depth}_{i,t} + \epsilon_{i,t}$$

From this exercise, I obtain individual estimates $\hat{\beta}_i^c$ for all modern US stocks, and for all Israeli bonds and stocks. I estimate β^c as the cross-sectional average of the security level estimates

$$\hat{\beta}^c = \frac{1}{N} \sum_{i=1}^N \hat{\beta}_i^c. \quad (38)$$

Table ?? reports the estimates for US stocks, and Israeli bonds and stocks, separately. In Israel, the estimates are 0.65 for stocks, and 0.60 for bonds. The difference between the two estimates, 0.05, is significant at the 10% level. This difference is potentially caused by the minimum order amount rule of the Tel Aviv Stock Exchange. The minimum order size is 2,000NIS for stocks

and 10,000 for bonds.¹⁷ Such constraints make the order flow more lumpy, which might mute the response of volume to change in book depth.

For US stocks, the estimate is 0.45. The smaller US estimate is probably related to the fact that a larger fraction of trading in US equities is done off exchange (9% in Israel versus 33% in the US). Hence we should expect volume to be less sensitive to book depth in the US. Since the estimates for bonds and stocks are close to each other in the Israeli sample, I assume that the estimate β^c obtained for US equities would also hold for bonds when I estimate the welfare implications of centralizing corporate bond trading.

6.3 Estimation of ξ

The last parameter to estimate is the frequency of liquidity shocks, ξ . Intuitively, the value of this parameter should be tied to the frequency at which the asset is traded. In steady state, the annual turnover of the asset, $T = \frac{\int_t^{t+1} g_s^k ds}{x}$, is given by

$$T = \frac{\rho_s^k \xi}{\rho_s^k + \xi} = \frac{\xi}{1 + \frac{\xi}{\rho_s^k}} \leq \xi. \quad (39)$$

Hence the asset's turnover, measured on either market structures, provides a lower bound on the liquidity needs of investors. Provided that $\frac{S^d}{\theta_m} > \frac{S^c}{\beta^c}$, this inequality implies that

$$W^c - W^d \geq \frac{Tx}{r + \eta} \left(\frac{S^d}{\theta_m} - \frac{S^c}{\beta^c} \right). \quad (40)$$

When instead $\frac{S^d}{\theta_m} < \frac{S^c}{\beta^c}$, the opposite inequality holds. In the welfare calculations below, I substitute the parameter ξ for the bond's turnover.

¹⁷See Abudy and Wohl (2017) for an in depth discussion of the institutional details of the Tel Aviv Stock Exchange.

6.4 Welfare Calculations

In this section, we analyze the welfare implications of centralizing the corporate bond market structure in the cross section of bonds. For each bond i , I calculate the quantity

$$\Delta W_i = \frac{T_i x_i}{r + \eta_i} \left(\frac{S_i^d}{\hat{\theta}_m} - \frac{S_i^c}{\hat{\beta}^c} \right),$$

where T_i is the bond's turnover, x_i is number of bond outstanding, $r = 0.02$ is the risk-free rate, $\frac{1}{\eta_i}$ is the time to maturity at issuance, S_i^d is the OTC transaction cost, S_i^c is the equity-implied transaction cost estimate, $\hat{\theta}_m = 0.73$ is the estimate of dealer bargaining power in OTC markets, and $\hat{\beta}^c = 0.45$ is the estimate of the market maker share in the matching function on the centralized venue. Table 8 reports the estimates across credit ratings.

The cross sectional differences in the benefit of centralizing the market structure echoes the findings of Section 4 where we showed that lower credit rated bonds would experience larger decline in transaction costs. In terms of magnitude, the welfare gains correspond to 14 bps of the par value of Aaa/Aa rated bonds and 29 bps of the total par value of Ba rated bonds.

7 Conclusion

This paper showed that moving corporate bond trading to a centralized market structure would generate large benefits. On average, effective spreads would decrease by 70% while quoted spread would decrease by 77%. Over the life of the average bond, this implies transaction cost savings equal to 0.66% of the bond total par value. Based on the insights of a new model of centralized and decentralized trading, I find that moving to a limit order market would generate welfare gains equal to 21 bps of total par value for the average bond. Larger bond issues, with lower credit ratings and longer time to maturity would benefit the most from a centralized market structure. The results have obvious policy implications. Future research could use the methodological framework develop in this paper to evaluate if other OTC markets would benefit from moving to centralized venues.

A Appendix A

A.1 Proof of proposition 1

Investor's and market maker's Bellman equations make clear that at each point in time, two outcomes are possible: either the market is active and the measures b_t , d_t , and s_t are all strictly positive, or the market is inactive $b_t = d_t = s_t = 0$. Let's conjecture that the former outcome holds for all $t \geq 0$. This implies that the zero-profit condition of market makers and zero-utility condition of buyers hold with equality at all times. Substituting in these two conditions the expressions for the bid and ask prices from (19) and (20), we obtain

$$\rho_{b_t}^d \theta_b \Delta_{hl,t}^d = \chi_b^d, \quad (\text{A.1})$$

$$\rho_{m_t}^d \theta_m \Delta_{hl,t}^d = \chi_m^d, \quad (\text{A.2})$$

for all $t \in [0, \infty)$. Manipulating (A.1) and (A.2), we solve for the buyer-to-seller ratio and the market maker-to-seller ratio

$$\frac{b_t}{s_t} = \left[\lambda^d \left(\frac{\theta_b}{\chi_b^d} \right)^{1-\beta^d} \left(\frac{\theta_m}{\chi_m^d} \right)^{\beta^d} \Delta_{hl,t}^d \right]^{\frac{1}{1-\alpha^d-\beta^d}}, \quad (\text{A.3})$$

$$\frac{d_t}{s_t} = \left[\lambda^d \left(\frac{\theta_b}{\chi_b^d} \right)^{\alpha^d} \left(\frac{\theta_m}{\chi_m^d} \right)^{1-\alpha^d} \Delta_{hl,t}^d \right]^{\frac{1}{1-\alpha^d-\beta^d}}. \quad (\text{A.4})$$

Substituting the prices (19) and (20), as well as the ratios (A.3) and (A.4) back into the investor's HJB equations (14)–17 shows after manipulation that the surplus must satisfy $\dot{\Delta}_{hl,t}^d = f^d(\Delta_{hl,t}^d)$, where

$$f^d(\Delta_{hl,t}^d) = (r + \eta + \xi) \Delta_{hl,t}^d + \theta_s \left[\lambda^d \left(\frac{\theta_b}{\chi_b^d} \right)^{\alpha^d} \left(\frac{\theta_m}{\chi_m^d} \right)^{\beta^d} \Delta_{hl,t}^d \right]^{\frac{1}{1-\alpha^d-\beta^d}} - (\delta + \chi_s^d). \quad (\text{A.5})$$

Since this equation doesn't explicitly depends on time, we conjecture that $\Delta_{hl,t}^d = \Delta_{hl}^d$. Under this conjecture, the gain from trade must solve $f^d(\Delta_{hl}^d) = 0$, which clearly has a unique positive solution.

We now characterize the condition under which our initial conjecture that an active market

exists holds. Substituting the expression of the bid price back into the HJB equation of type 11 reveals that low valuation owners become active sellers if and only if $\rho_s^d \theta_s \Delta_{hl}^d \geq \chi_s$ (assuming that seller participate when they are indifferent between participating and staying out of the market). This condition is verified if and only if $f\left(\frac{\chi_s^d}{\rho_s^d \theta_s}\right) \leq 0$. A direct calculation shows that this condition is equivalent to

$$\frac{\delta}{r + \xi} \geq \frac{1}{g^d\left(\frac{\theta_b}{\chi_b^d}, \frac{\theta_m}{\chi_m^d}, \frac{\theta_s}{\chi_m^d}\right)}.$$

When this condition is satisfied, all low-valuation investors become active in the market, such that $s_t = \mu_{l1,t}^d$. Solving the values function in terms of the gain from trade yields

$$\begin{aligned} V_{h0,t}^d &= V_{l0,t}^d = V_{d,t}^d = 0, \\ V_{h1,t}^d &= \frac{c}{r + \eta} + \frac{\eta p}{r + \eta} - \left(\frac{\xi}{r + \eta}\right) \Delta_{hd}^d, \\ V_{l1,t}^d &= \frac{c}{r + \eta} + \frac{\eta p}{r + \eta} - \left(1 + \frac{\xi}{r + \eta}\right) \Delta_{hd}^d. \end{aligned}$$

The bid and ask prices are then

$$\begin{aligned} A_t^d &= \frac{c}{r + \eta} + \frac{\eta p}{r + \eta} - \left(\frac{\xi}{r + \eta} + \theta_b\right) \Delta_{hl}^d, \\ B_t^d &= \frac{c}{r + \eta} + \frac{\eta p}{r + \eta} - \left(\frac{\xi}{r + \eta} + \theta_b + \theta_d\right) \Delta_{hl}^d, \end{aligned}$$

The trade execution function can be written as $g_t^d = \rho_s^d \mu_{l1,t}^d$, and the law of motion for the measure of type-11 investor written as

$$\dot{\mu}_{l1,t}^d = -\rho_s^d \mu_{l1,t}^d + \xi (x - \mu_{l1,t}^d).$$

Solving the ODE yields

$$\mu_{l1,t}^d = \left(\mu_{l1,0}^d - \frac{\xi \rho_s^d}{\xi + \rho_s^d}\right) e^{-(\xi + \rho_s^d)t} + \frac{\xi \rho_s^d}{\xi + \rho_s^d}.$$

□

A.2 Proof of proposition 2

Assuming that a contract σ_t is active for all $t \geq 0$, the zero-profit condition of market makers and the zero-utility condition of buyers for that contract must hold with equality

$$\rho_{b_t}^c(\sigma_t) (V_{h1,t}^c - V_{h0,t}^c - A_t) = \chi_b^c, \quad (\text{A.6})$$

$$\rho_{m_t}^c(\sigma_t) (A_t - B_t) = \chi_m^c. \quad (\text{A.7})$$

Substituting these two conditions back into the no-surplus condition (31), taking the FOC and solving for the buyer-to-seller and market maker-to-seller ratios yields

$$\frac{b_t}{s_t} = \left[\lambda^c \left(\frac{\alpha^c}{\chi_b^c} \right)^{1-\beta^c} \left(\frac{\beta^c}{\chi_m^c} \right)^{\beta^c} \Delta_{hl,t}^c \right]^{\frac{1}{1-\alpha^c-\beta^c}}, \quad (\text{A.8})$$

$$\frac{d_t}{s_t} = \left[\lambda^c \left(\frac{\alpha^c}{\chi_b^c} \right)^{\alpha^c} \left(\frac{\beta^c}{\chi_m^c} \right)^{1-\alpha^c} \Delta_{hl,t}^c \right]^{\frac{1}{1-\alpha^c-\beta^c}}. \quad (\text{A.9})$$

Substituting these ratios back into the free entry conditions reveals that

$$V_{h1,t}^c - V_{h0,t}^c - A_t = \alpha^c \Delta_{hl,t}^c, \quad (\text{A.10})$$

$$A_t - B_t = \beta^c \Delta_{hl,t}^c, \quad (\text{A.11})$$

$$V_{l0,t}^c - V_{l1,t}^c + B_t = (1 - \alpha^c - \beta^c) \Delta_{hl,t}^c. \quad (\text{A.12})$$

Substituting the buyer-to-seller and market maker-to-seller ratios back into the investor HJB shows after manipulation that the gain from trade must satisfy $\dot{\Delta}_{hl,t}^c = f^c(\Delta_{hl,t}^c)$ where

$$f^c(\Delta_{hl,t}^c) = (r + \eta + \xi) \Delta_{hl,t}^c + (1 - \alpha^c - \beta^c) \left[\lambda^c \left(\frac{\alpha^c}{\chi_b^c} \right)^{\alpha^c} \left(\frac{\beta^c}{\chi_m^c} \right)^{\beta^c} \Delta_{hl,t}^c \right]^{\frac{1}{1-\alpha^c-\beta^c}} - (\delta + \chi_s^c). \quad (\text{A.13})$$

Following from there the steps we followed in the proof of Proposition 1, we obtain the remaining results of Proposition 2.

References

- Abel, A., Eberly, J., and Panageas, S. (2013). Optimal inattention to the stock market with information costs and transaction costs. *Econometrica*, 81(4):1455–1481.
- Abudy, M. and Wohl, A. (2017). Corporate bond trading on a limit order book exchange. Forthcoming in *Review of Finance*.
- Admati, A. and Pfleiderer, P. (1988). A theory of intraday patterns: Volume and price variability. *The Review of Financial Studies*, 1(1):3–40.
- Angel, J. J., Harris, L. E., and Spatt, C. S. (2011). Equity trading in the 21st century. *Quarterly Journal of Finance*, 1(1):1–53.
- Angel, J. J., Harris, L. E., and Spatt, C. S. (2015). Equity trading in the 21st century: An update. *Quarterly Journal of Finance*, 5(1):1–39.
- Babus, A. and Parlatore, C. (2017). Strategic fragmented markets. Working Paper.
- Barclay, M., Christie, W., Harris, J., Kandel, E., and Schultz, P. (1999). Effects of market reform on the trading costs and depths of nasdaq stocks. *The Journal of Finance*, 54(1):1–33.
- Bessembinder, H., Maxwell, W., and Ventakaraman, K. (2006). Optimal market transparency: Evidence from the initiation of trade reporting in corporate bonds. *Journal of Financial Economics*, 82(2):251–288.
- Biais, B. and Declerck, F. (2007). Liquidity, competition and price discovery in the european corporate bond market. Working Paper.
- Biais, B. and Green, R. (2007). The microstructure of the bond market in the 20th century. Working Paper.
- BIS (2014). Market-making and proprietary trading: Industry trends, drivers and policy implications. CGFS Papers No 52.

- Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, 1:451–487.
- Childs, M. (2014). Blackrock says 'broken' company bond market needs change now. Retrieved from <https://www.bloomberg.com>.
- Duffie, D., Garleanu, N., and Pedersen, L. H. (2005). Over-the-counter markets. *Econometrica*, 73(6):1815–1847.
- Edwards, A. K., Harris, L. E., and Piwowar, M. S. (2007). Corporate bond market transaction costs and transparency. *The Journal of Finance*, 62(3):1421–1451.
- Edwards, S., Longstaff, F. A., and Garcia Marin, A. (2015). The u.s. debt restructuring of 1933: Consequences and lessons. Working Paper.
- Glode, V. and Opp, C. (2017). Over-the-counter vs. limit-order markets: The role of traders' expertise. Working Paper.
- Goldstein, M. A., S., H. E., and R., S. E. (2007). Transparency and liquidity: A controlled experiment on corporate bonds. *The Review of Financial Studies*, 20(2):235–273.
- Green, R., Hollifield, B., and Schurhoff, N. (2007). Financial intermediation and the costs of trading in an opaque market. *Review of Financial Studies*, 20(2):275–314.
- Harris, L. (2015). Transaction costs, trade throughs, and riskless principal trading in corporate bond markets. Working Paper.
- Harris, L., Kyle, A. S., and Sirri, E. R. (2015). Statement of the financial economists roundtable, april 2015: The structure of trading in bond markets. *Financial Analysts Journal*, 71(6):5–8.
- Harris, L. and Piwowar, M. (2006). Secondary trading costs in the municipal bond market. *The Journal of Finance*, 61(3):1361–1397.

- Hendershott, T., Li, D., Livdan, D., and Schurhoff, N. (2017). Relationship trading in otc markets. Working Paper.
- Hendershott, T. and Madhavan, A. (2015). Click or call? auction versus search in the over-the-counter market. *The Journal of Finance*, 70(1):419–447.
- Hickman, B. (1960). Statistical measures of corporate bond financing since 1900. NBER.
- Lagos, R. and Rocheteau, G. (2009). Liquidity in asset markets with search frictions. *Econometrica*, 77(2):403–426.
- Lagos, R., Rocheteau, G., and Weill, P.-O. (2011). Crises and liquidity in over-the-counter markets. *Journal of Economic Theory*, 146(6):2169–2205.
- Lester, B., Rocheteau, G., and Weill, P.-O. (2015). Competing for order flow in otc markets. *Journal of Money, Credit and Banking*, 47(2):77–126.
- Moen, E. (1997). Competitive search equilibrium. *The Journal of Political Economy*, 105(2):385–411.
- Ng, S. and Grind, K. (2012). Large institutions discuss new marketplace for bonds. Retrieved from <http://www.wsj.com>.
- Odegaard, B. A. (2017). Bond liquidity at the oslo stock exchange. *Norges Bank Staff Memo*, (5):1–49.
- O’Hara, M., Wang, Y., and Zhou, X. (2017). The execution quality of corporate bonds. Forthcoming in *Journal of Financial Economics*.
- Pagano, M. (1989). Trading volume and asset liquidity. *The Quarterly Journal of Economics*, 104(2):255–274.
- Protnick, D. and Gur-Gershgoren, G. (2011). Corporate bond market liquidity in israel at heterogeneous trade sizes. *Koret-Milken Institute Fellows Program*, (37):1–49.

- Rust, J. and Hall, G. (2003). Middlemen versus market makers: A theory of competitive exchange. *Journal of Political Economy*, 111(2):353–403.
- Sapp, T. R. A. and Yan, X. (2003). The nasdaq-amex merger, nasdaq reforms, and the liquidity of small firms. *The Journal of Financial Research*, 26(2):225–242.
- Schultz, P. (2001). Corporate bond trading costs: A peek behind the curtain. *The Journal of Finance*, 56(2):677–698.
- Weston, J. P. (2000). Competition on the nasdaq and the impact of recent market reforms. *The Journal of Finance*, 55(6):2565–2598.

Table 1: SAMPLE DESCRIPTIVE STATISTICS

Variable	US (2013-2015)		Israel (2013-2015)		US (1917-1921)	
	Bonds (N=1,841)	Stocks (N=462)	Bonds (N=139)	Stocks (N=51)	Bonds (N=188)	Stocks (N=76)
Common variables						
Quoted spread (%)	0.56	0.09	0.22	0.73	2.09	2.72
Effective spread (%)	0.23	0.07	0.20	0.65		
Daily volume (2015 \$mm)	2.11	130.38	0.55	1.04	0.07	3.10
Market capitalization (2015 \$bn)	0.69	22.32	0.23	1.03	0.37	0.98
Annual Turnover (%)	86	268	61	20	6	112
Volatility (%)	5	32	6	26	10	29
Inst. ownership (%)	89	61	70	35		
Off exchange volume (%)	100	33	6	10		
Bond specific characteristics						
Time to maturity (years)	8		8		34	
Time to maturity at issuance (years)	12		11		51	
Callable (%)	96		5		38	
Sinkable (%)	0		90		21	
Secured (%)	20		12		91	
CPI-linked/gold clause (%)	0		69		84	
IG (%)	71		47		91	
HY (%)	20		2		9	
Unrated (%)	9		51		0	

This table reports liquidity and security characteristic averages for the equities and bonds that compose the three samples. For a given security in a given year, the daily value of the variables are first averaged over the year. The table reports cross-sectional averages across all security-year annual aggregates. Quoted spread is defined in equation (2) and effective spread in equation (3); details regarding time aggregation are provided in section 3. Daily volume is the daily trading volume expressed in 2015 million USD. Market capitalization is the total market value of the outstanding issue expressed in 2015 million USD. Amounts originally expressed in Israeli new shekel were converted to USD using an exchange rate of 3.7, the average over the sample period. The dollar amounts of the historical sample have been CPI adjusted by a factor of 16. Annual turnover is defined as the total annual dollar volume divided by total market capitalization. Volatility is the annualized standard deviation of daily midpoint returns. Time to maturity corresponds to the number of years left before the principal of the bond is due for repayment. Time to maturity at issuance represent the number of years a bond had until maturity at the time of issuance. A bond is callable if it can be redeemed by the issuer prior to its maturity. A bond is sinkable if the issuer has committed to periodically repurchase a fraction of the issue. A bond is secured if it is collateralized or guaranteed by a third party. Gold clauses refer to a contractual feature of historical US bonds that pegged the value of a bond's coupons and principal to the price of gold at the time of issuance (see subsection 3.3 for additional details). A bond is classified as investment grade (IG) if it has a Moody's credit rating of Baa or above, lower credit rated bonds are classified as high yield (HY), and bonds without credit ratings are classified as unrated.

Table 2: QUOTED SPREAD REGRESSIONS

Dependent Variable: Log(QS)	US (2013-2015)			Israel (2013-2015)			US (1917-1921)		
	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
Bond	1.88*** (0.05)			0.03 (0.07)			-0.89*** (0.05)		
Log(volume)	-0.29*** (0.01)	-0.33*** (0.01)	-0.39*** (0.01)	-0.52*** (0.02)	-0.53*** (0.02)	-0.55*** (0.02)	-0.41*** (0.01)	-0.41*** (0.01)	-0.43*** (0.01)
Bond \times log(volume)			0.17*** (0.02)			0.05 (0.03)			0.03 (0.02)
Log(Std returns)	0.46*** (0.01)	0.37*** (0.02)	0.41*** (0.03)	0.71*** (0.04)	0.62*** (0.05)	0.60*** (0.07)	0.46*** (0.03)	0.42*** (0.03)	0.46*** (0.05)
Bond \times log(std returns)			-0.09** (0.04)			0.04 (0.08)			-0.07 (0.06)
IG \times ST		1.28*** (0.09)	1.11*** (0.11)		-0.13 (0.16)	-0.13 (0.18)		-1.08*** (0.11)	-1.11*** (0.13)
IG \times LT		1.47*** (0.07)	1.33*** (0.11)		0.04 (0.19)	0.02 (0.19)		-0.96*** (0.11)	-0.97*** (0.13)
HY \times ST		2.01*** (0.09)	1.86*** (0.12)		0.00 (0.12)	-0.04 (0.15)			
HY \times LT		1.84*** (0.09)	1.67*** (0.12)		0.23 (0.15)	0.18 (0.20)		-0.76*** (0.13)	-0.75*** (0.15)
NR \times ST		1.21*** (0.09)	1.05*** (0.12)		-0.15 (0.16)	-0.16 (0.18)			
NR \times LT		1.44*** (0.08)	1.29*** (0.12)		0.11 (0.14)	0.09 (0.16)			
CPI linked/Gold clause					0.12* (0.06)	0.12* (0.06)		0.07 (0.05)	0.06 (0.05)
Callable		0.02 (0.05)	0.06 (0.04)		-0.32** (0.13)	-0.30** (0.13)		-0.02 (0.04)	-0.02 (0.04)
Sinkable					-0.21* (0.12)	-0.21 (0.13)		0.01 (0.04)	0.00 (0.04)
Secured		0.20*** (0.04)	0.17*** (0.04)		0.05 (0.06)	0.07 (0.07)		-0.01 (0.08)	-0.00 (0.09)
Observations	3547	3378	3378	414	414	414	989	989	989
R-squared	0.85	0.88	0.89	0.90	0.91	0.91	0.81	0.81	0.81

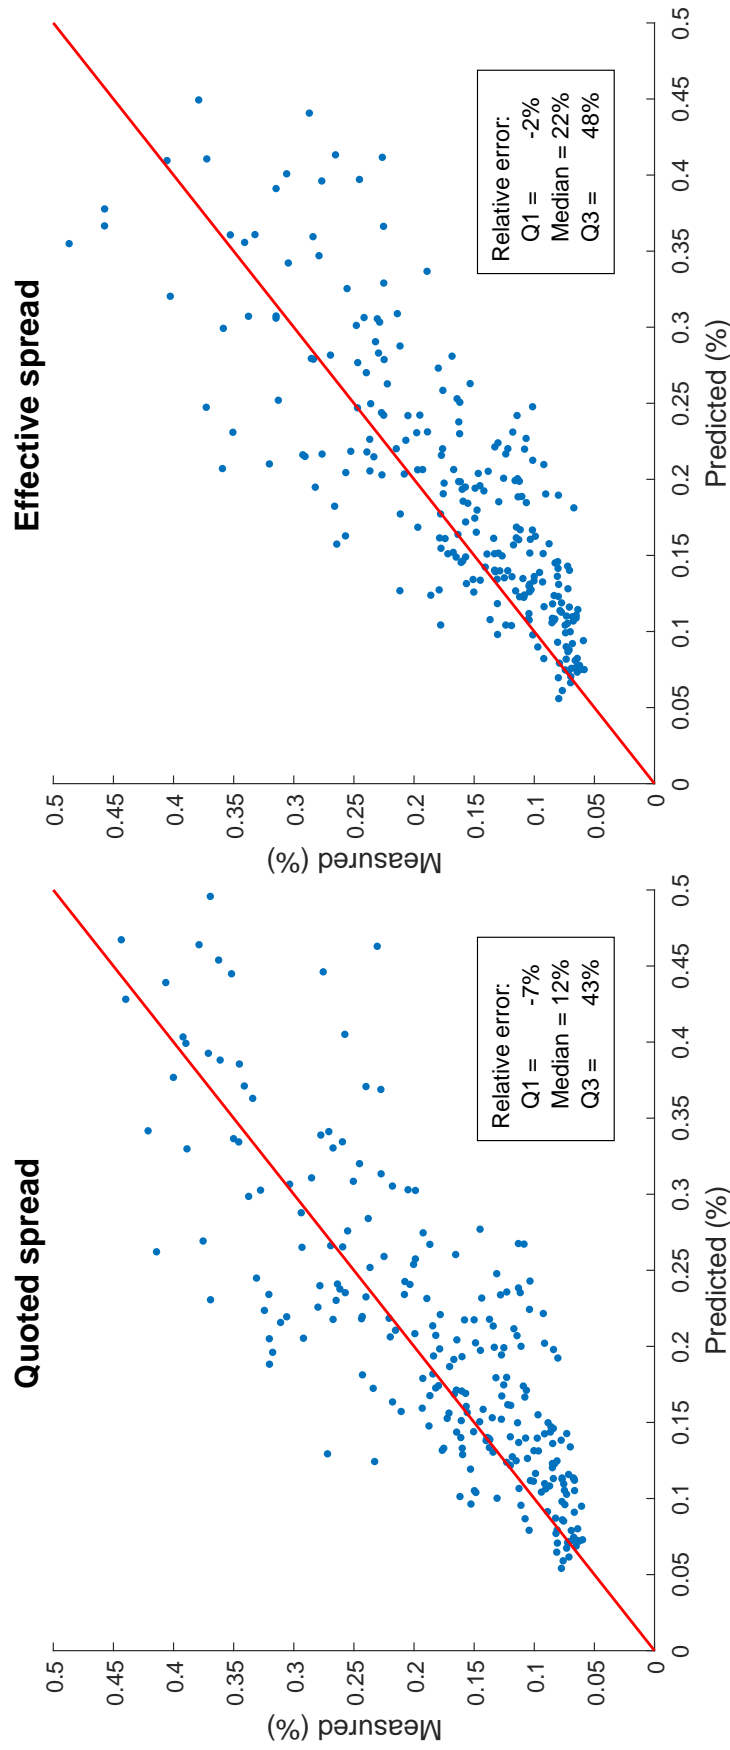
The dependent variable is the annual average quoted spread (QS), calculated according to equation (2); details regarding time aggregation are provided in Section 3. A bond is classified as short term (ST) if it has five years or less until maturity, and long term (LT) otherwise. The remaining variables are defined in Table 1. In all specifications, standard errors are clustered at the firm (or issuer) level. Standard errors are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 3: EFFECTIVE SPREAD REGRESSIONS

Dependent Variable: Log(ES)	US (2013-2015)			Israel (2013-2015)		
	(1)	(2)	(3)	(1)	(2)	(3)
Bond	1.60*** (0.07)			-0.00 (0.08)		
Log(volume)	-0.20*** (0.01)	-0.22*** (0.01)	-0.31*** (0.01)	-0.48*** (0.01)	-0.48*** (0.01)	-0.49*** (0.02)
Bond \times log(volume)			0.23*** (0.03)			0.03 (0.03)
Log(Std returns)	0.54*** (0.02)	0.49*** (0.03)	0.45*** (0.05)	0.72*** (0.04)	0.66*** (0.06)	0.60*** (0.06)
Bond \times log(std returns)			0.03 (0.05)			0.10 (0.08)
IG \times ST		1.22*** (0.14)	0.75*** (0.15)		-0.01 (0.18)	-0.07 (0.17)
IG \times LT		1.25*** (0.11)	0.76*** (0.15)		0.17 (0.20)	0.09 (0.19)
HY \times ST		1.73*** (0.14)	1.24*** (0.17)		0.10 (0.13)	0.00 (0.14)
HY \times LT		1.80*** (0.13)	1.28*** (0.17)		0.29** (0.13)	0.17 (0.15)
NR \times ST		1.17*** (0.14)	0.69*** (0.16)		0.02 (0.17)	-0.04 (0.17)
NR \times LT		1.28*** (0.13)	0.77 (0.17)		0.21 (0.15)	0.13 (0.14)
CPI linked/Gold clause					0.05 (0.06)	0.04 (0.06)
Callable		0.04 (0.09)	0.10 (0.09)		-0.32** (0.14)	-0.31** (0.14)
Sinkable					-0.25** (0.12)	-0.26 (0.13)
Secured		-0.01 (0.06)	-0.05 (0.06)		0.00 (0.07)	0.00 (0.07)
Observations	3539	3371	3371	414	414	414
R-squared	0.56	0.60	0.62	0.89	0.90	0.90

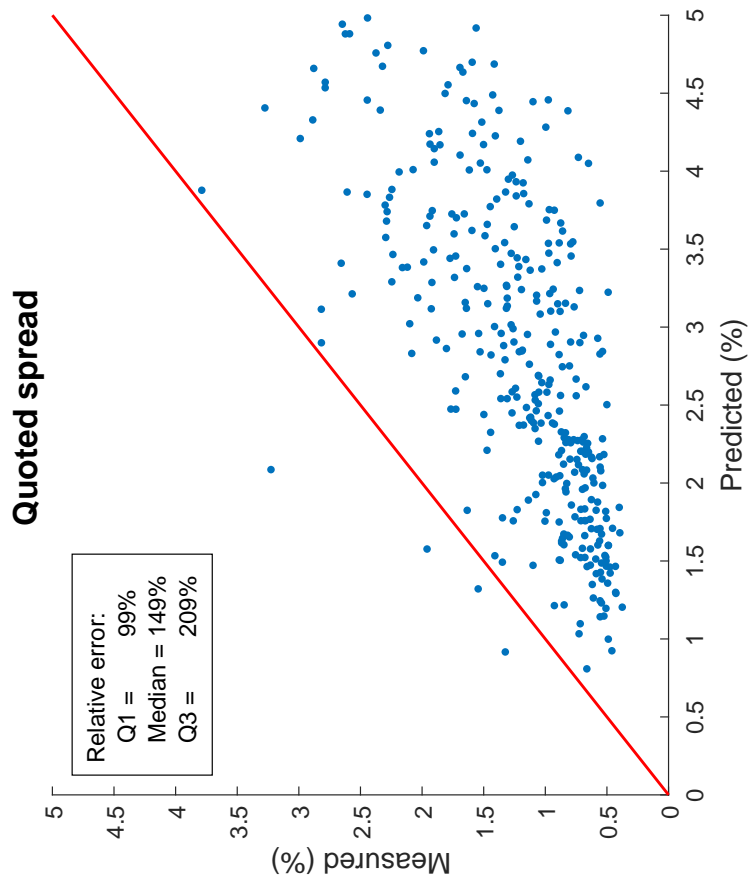
The dependent variable is the annual average effective spread (ES), calculated according to equation (3); details regarding time aggregation are provided in Section 3. A bond is classified as short term (ST) if it has five years or less until maturity, and long term (LT) otherwise. The remaining variables are defined in Table 1. In all specifications, standard errors are clustered at the firm (or issuer) level. Standard errors are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Figure 1:
Measured vs. predicted transaction costs
Israel (2013 – 2015)



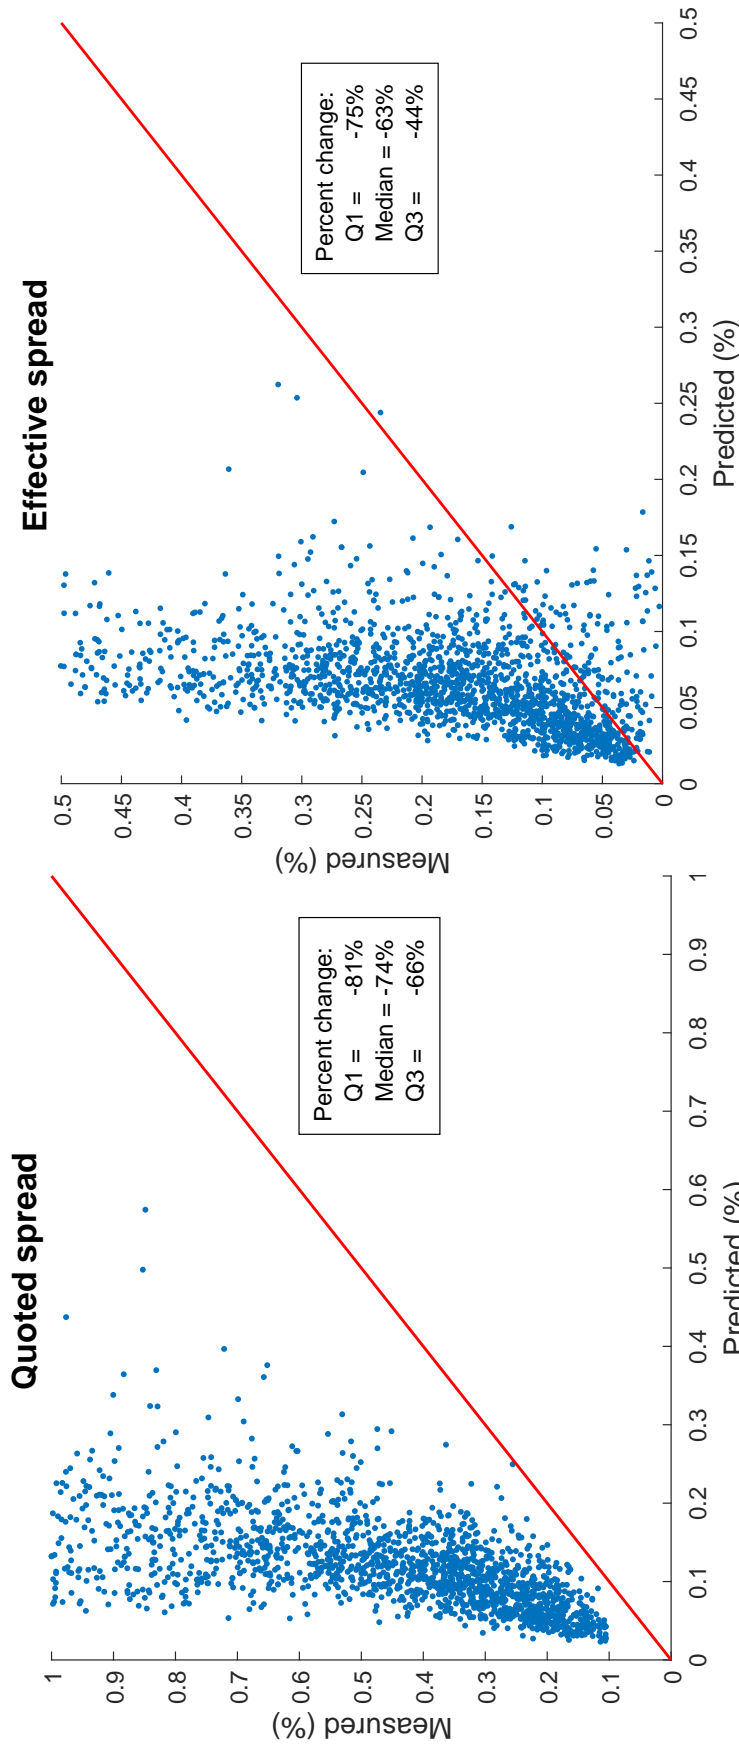
These two figures plot measured transaction costs, on the y-axis, against the corresponding equity-implied estimates, on the x-axis, for bonds in the Israeli sample. Each point represents a bond-year. The equity-implied estimates are calculated following the two-step procedure characterized by equations (4) and (5). The left figure reports the results for quoted spreads, while the right figure reports the results for effective spreads. For bonds lying under the red 45 degree line, equity-implied estimates overshoot measured transaction costs, and vice-versa. The figure reports the three quartile of the relative mean error, as defined in equation (6). A positive relative error means that the equity-implied estimate is larger than the measured transaction cost.

Figure 2:
 Measured vs. predicted transaction costs
 US (1917 – 1921)



This figure plots measured quoted spread, on the y-axis, against the corresponding equity-implied estimates, on the x-axis, for bonds in the historical US sample. Each point represents a bond-year. The equity-implied estimates are calculated following the two-step procedure characterized by equations (4) and (5). For bonds lying under the red 45 degree line, equity-implied estimates overshoot measured transaction costs, and vice-versa. The figure reports the three quartile of the relative mean error, as defined in equation (6). A positive relative error means that the equity-implied estimate is larger than the measured transaction cost.

Figure 3:
 Measured vs. predicted transaction costs
 US (2013 – 2015)



These two figures plot OTC measured transaction costs, on the y-axis, against the corresponding equity-implied estimates, on the x-axis, for bonds in the modern US sample. Each point represents a bond-year. The equity-implied estimates are calculated following the two-step procedure characterized by equations (4) and (5). The left figure reports the results for quoted spreads, while the right figure reports the results for effective spreads. For bonds lying under the red 45 degree line, equity-implied estimates overshoot measured transaction costs, and vice-versa. The figure reports the three quartile of the relative percent change, as defined in equation (7). A negative percent change means that the equity-implied estimate predicts that the transaction cost of the bond would decrease on a centralized venue.

Table 4: Measured vs. Predicted Transaction Costs
US(2013 – 2015)

	Mean	Std	P5	Q1	Median	Q3	P95
Quoted spread							
predicted (bp)	13	6	5	9	12	15	24
measured (bp)	56	34	17	29	46	77	121
difference (bp)	-43	-28	-12	-20	-34	-62	-97
percent change (%)	-77	-82	-71	-69	-74	-81	-80
Effective spread							
predicted (bp)	7	3	3	5	6	9	13
measured (bp)	23	23	4	10	17	28	67
difference (bp)	-16	-20	-1	-5	-11	-17	-54
percent change (%)	-70	-87	-25	-50	-65	-61	-81
Transaction cost savings							
Annual dollar savings (\$mn)	0.42	0.53	0	0.07	0.21	0.57	1.51
Annual relative savings (bp of par)	6	7	0	1	3	8	20
PV dollar savings (\$mn)	4.74	7.63	0	0.45	1.66	5.65	21.81
PV relative savings (bp of par)	66	95	0	8	29	81	256

This table reports cross-sectional descriptive statistics of the predicted change in transaction costs from moving to a centralized market structure. Quoted spread is defined in equation (3), and effective spread in equation (2). Predicted spread measures are obtained following the two-step procedure characterized by equations (4) and (5). Annual transaction cost savings are defined in equations (8) and (9), and the present values of the transaction cost savings over the lifetime of a bond are defined in equations (10) and (11).

Table 5: TRANSACTION COST SAVINGS IN THE CROSS SECTION OF BONDS

Dependent Variable	Δ Quoted spread	Δ Effective spread	Log(ADS)
Intercept	-0.72*** (0.02)	-0.47*** (0.02)	-2.11*** (0.14)
Log(total par value)	0.00 (0.01)	-0.18*** (0.02)	1.25*** (0.07)
Log(time to maturity)	-0.02*** (0.00)	-0.01 (0.01)	0.49*** (0.05)
HY	-0.11*** (0.01)	-0.21*** (0.02)	1.36*** (0.10)
NR	0.01 (0.01)	-0.02 (0.02)	0.36** (0.14)
Callable	0.02 (0.02)	-0.04 (0.02)	0.22 (0.14)
Secured	-0.02 (0.02)	-0.01 (0.02)	0.09 (0.09)
Observations	1,642	1,642	1,642
R-squared	0.24	0.20	0.41

The dependent variable in the first column is the ratio of predicted quoted spread to measured quoted spread. The quoted spread is calculated according to equation (2). The dependent variable in the second column is the ratio of predicted effective spread to measured effective spread. The effective spread is calculated according to equation (3). The dependent variable in the third column is the annual dollar savings defined in equation (8). The independent variables are all defined in Table 1. In all specifications, standard errors are clustered at the firm (issuer) level. Standard errors are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 6: FIT OF THE LIQUIDITY PREMIUM WHEN $\theta_m = 0.73$

	Aaa/Aa	A	Baa	Ba
$A - B$ (bps)	14	15	17	19
ξ (%)	31	32	37	41
η (year)	7.4	7.3	7.8	7.3
Liquidity premium (bps)				
target	50	57	67	83
model	51	53	72	80
r=2%				

Table 7: ESTIMATES OF β^c

	$\hat{\beta}^c$	SE
Israel		
stocks	0.65	0.02
bonds	0.60	0.02
diff	0.05*	0.03
US		
stocks	0.45	0.01

This table reports the estimate of the structural parameter β^c following the procedure outlined in Section 6. The columns labeled "Estimate" contains the point estimate of the parameters. Standard errors are reported in the column SE. The row diff reports the results of a difference in means test between Israeli bonds and stocks. The standard error of the test is calculated using Satterthwaite's approximation. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 8: Welfare analysis

	Aaa/Aa	A	Baa	Ba
<i>A</i> – <i>B</i> (bps)				
decentralized	14	15	17	20
centralized	5	6	6.5	7
difference	9	9	10.5	13
liquidity premium (bps)				
decentralized	51	53	72	80
centralized	34	36	44	45
difference	17	17	28	35
Welfare change (bps total par)	14	13	23	29