Concentration in Over-the-Counter Markets and Its Impact On Their Performance During Crisis *

Assa Cohen

Sy Syms School of Business, Yeshiva University and The Federal Reserve Bank of Philadelphia †‡

October 5, 2023

Abstract

The paper investigates the connection between concentration and fragility in over-the-counter (OTC) markets. I argue that the increase in spreads in OTC markets during a crisis reflects an increase in dealer *mark-ups* and not merely an increase in dealer *costs* of facilitating trade. Using Regulatory TRACE data on the US Corporate Bonds market, I construct a bond-level HH-index to gauge concentration in the market for each bond. Focusing on the COVID-19 crisis that emerged in the first quarter of 2020, I show that concentrated markets exhibit a greater increase in spreads and a stronger decline in trade volumes during a crisis. I present a model in which trade in a bond is led by dealers who acquire information about it. Systemic distress incentivizes informed dealers to exercise market power more aggressively by submitting low bids that appeal solely to distressed customers. When calibrated to the behavior of spreads across different concentration levels before and during the COVID-19 crisis, the model successfully reproduces the data, including the response of volume to the crisis, which was not targeted. Additionally, the calibration demonstrates that the increased uncertainty during the COVID-19 crisis was key for concentration to exacerbate the severity of the crisis as much as it did.

Keywords: market power, OTC markets, market freeze, information acquisition, corporate bonds.

JEL classification:: E44, G01, G11, G12, G21.

*I am grateful to Guillermo Ordonez, Benjamin Lester, Itay Goldstein, Jesus Fernandez-Villaverde, and Harold Cole for their invaluable advice and continuous support. I am also thankful for the helpful comments and suggestions provided by Enghin Atalay, Luke Taylor, Zvica Eckstein, and the participants of the Penn Macro Seminar. Further, I thank Kerry Rowe from the Federal Reserve Bank of Philadelphia and FINRA for their help in approving the disclosure of results from Regulatory TRACE.

†The views expressed herein are solely those of the author and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors and omissions are my own.

‡Email: assa.cohen@yu.edu

1 Introduction

One of the main functions of financial markets is to provide liquidity. However, often in times of systemic distress, financial markets exhibit a sharp decline in trade volume alongside an increase in the cost of trade (measured by the bid-ask spread). The deterioration in trade conditions limits investors' ability to raise liquidity by selling securities (see Gorton and Metrick (2012)). Furthermore, it often increases the cost of credit, hereby impairing real economic activity (Bao et al. (2011)). Through these channels, it *amplifies* a financial crisis. A prime example of such amplification is the dramatic decline in trade in the market for Mortgage-Backed Securities that was central to the 2008-2009 financial crisis. Other instances of a dramatic decline in liquidity of large and important financial markets during economic downturns include the real-estate bonds and corporate bonds market in the great depression (1931), the junk bonds market in the early 1990s recession (1990), the market for Collateralized Bond Obligations following the burst of the Dot.com bubble (2001), and the markets for Collateralized Debt Obligations (CDO), Asset-Backed Securities (ABS), Commercial Papers and Corporate Bonds in the 2008 crisis, and more (for a full review see Benmelech and Bergman (2018)).

The decline in activity in financial markets is by no means a self-evident implication of systemic distress. In fact, times of distress appear to be times when market participants can gain the most from reallocating liquidity through trade. A key to understanding the dynamics of such a "freeze" in trade is that all markets that were subject to it share a critical trait: they are all . Decentralized markets, also known as *Over-the-Counter (OTC)* financial markets, are markets that lack a central trading platform. Trade in such markets consists mostly of bilateral trades between investors and designated intermediaries, "dealers", who operate as market makers. The claim that the OTC trading protocol itself explains the fragility of financial markets gained much traction after the 2007-2009 crisis, in which, as Chiu and Koeppl (2016) point out: "*there was a stunning difference in how asset markets were affected according to their infrastructure. Markets with centralized trading functioned rather well. To the contrary, in over-the-counter markets... trading came to a halt*".

Some have argued that the fragility of OTC markets originates from the dependence of trade on the performance of the dealer sector (see, for example, Bao et al. (2018) Dick-Nielsen and Rossi

(2019), Duffie (2020), Kargar et al. (2021)). According to such views, heightened demand for liquidity during a crisis exhausted dealers' capacity to absorb securities to their balance sheets as a part of their market-making activities. That, along with greater uncertainty and potential weakening of the balance sheet, leads dealers to require higher compensation for their services, that is, to charge higher spreads.

Such views that emphasize the impact of systemic distress on dealers *holding costs* seem to explain at least some of the dynamics of a "market freeze". However, they cannot account for a salient pattern in the data: during times of distress, a dealer who buys a security from a customer purchases it for a substantial discount compared to what it would have paid if she had bought it from another dealer [1]. Evidently, the holding cost of the security for the dealer who purchased it is the same, regardless of the identity of its seller. Thus, the widening gap seems to reflect a rise in a *mark-up* component. Customers pay higher markups because they are less integrated into the market and have limited access to alternative counter-party for trade.

In this paper, I am guided by the idea that the rise in spreads in OTC markets reflects an increase in the *mark-ups* they charge their customers, rather than merely an increase in the *cost* of facilitating trade. In other words, dealers prey on the dire need for liquidity among their customers to charge higher spreads. As a result, they get to facilitate lower volumes of trade. In this sense, some of the deterioration of trading conditions in OTC markets in times of crisis should be understood in terms of *monopolistic inefficiency*.

Beyond the empirical pattern mentioned above, the claim that market power will create a more substantial disruption in OTC markets during a crisis has a strong intuitive appeal. Consider a dealer who has a monopoly on trading a specific bond. She faces a trade-off between charging higher spreads (bidding low) and increasing the quantity traded. During a crisis, distressed players are eager to sell, even at a significant discount. From the dealer's perspective, this means that an increase in the spreads she charges will have a weaker impact on the volume of trade that she gets to facilitate. This, in turn, makes it optimal to submit lower bids that only appeal to these distressed players. Doing so will result in higher realized spreads and a decline in the

[1]This pattern is implied by papers that used the gap between the inter-dealer and customer-to-dealer prices to gauge spreads. For example, Choi et al. (2021) propose a measure of the gap between inter-dealer pricing and customer-dealer pricing and show that it increased by a factor of five during the 2008 financial crisis. In the following, I document the widening gap between inter-dealer and customer-dealer prices in a crisis more directly, to ensure that it does not emerge from peculiar properties of measures used in past literature.

volume of trades. Note that this is just a manifestation of a more general principle - as demand becomes less elastic, a monopoly (or any player with market power) transitions towards charging higher mark-ups and reducing production.

Naturally, this reasoning raises a second question that will guide this work: *Do dealers have substantial market power, and if so, what is its origin?*. Dealers' market power may seem somewhat surprising, as OTC market making appears, on the face of it, like a competitive field of operation. As I shall discuss below, these markets are typically populated by hundreds of dealers, including many well-known prominent players. Thus, customers will seldom have difficulty finding multiple alternatives for a dealer to trade with.

I explore these questions empirically and theoretically using Regulatory TRACE of the US Corporate Bonds Market, an exclusive version of the TRACE data set that provides dealer identities. Using dealer identities, I uncover a critical feature of the US Corporate Bonds Market - its segmentation. That is, I show that the trade in each specific bond in the market is dominated by a few dealers. For example, for the median bond, the four most prominent dealers trade it account for 60% of transactions and 78% of trade. Hence, although many dealers populate the market, a customer wishing to sell a bond might face very few potential buyers.

The cross-sectional difference in concentration between bonds is used to explore how market power affects the response to systemic distress. Guided by the intuition that dealers exploit customers' urgent need for liquidity to charge higher spreads, I focus on transactions in which customers sell to dealers. Applying regression analysis, I show that given two similar bonds, the one traded in a more concentrated market typically exhibits a greater increase in its spread during a crisis period. This is true even if the same dealer conducts the trade in both bonds. The pattern repeats itself in both the COVID-19 crisis of March-April 2020 and the 2007-2009 crisis. Using regression analysis, I show that it persists when controlling for the most likely confounders, including dealer identity, risk (rating), bond liquidity, and the expected holding time of the bond by the dealer. Even when including these controls, we find sizeable differences in the response of spreads to a crisis between similar bonds traded in markets with varying levels of competition. Specifically, the analysis implies that in a world where all bonds would have been traded in a highly competitive setting, the aggregate increase in spreads during the COVID-19 crisis would have been *at least* 20% higher than it actually was. Although this prediction is not a precious gauge of the impact of competition on trading conditions in times of distress, it indicates that it has a relatively large order of magnitude.

In addition, the data is used to document stylized facts about concentration in OTC markets that may shed light on its origin. I show that most dealers are highly specialized. The lion's share of a dealer's activity consists of trading heavily on a subset of bonds for which the dealer is a dominant player (accounts for much of the trade). At the same time, even large dealers avoid many of the bonds traded on the market. These patterns are reminiscent of an entry cost to the market for each - either a dealer pays it and operates extensively in the market, or he does not and avoids it altogether. At the same time, I find that the only bond attribute with substantial predictive power for the concentration level in its trade is its amount outstanding. Taking the amount outstanding as a proxy for the market size can also be interpreted as the result of an entry cost mechanism, with larger markets allowing more incumbents to cover the cost of entry from their operations.

I suggest interpreting the entry cost as the cost of acquiring information about the traded security. Differently put, dealers that dominate a market for a bond are those that are better informed about the bond's value. This interpretation draws on a large body of literature documenting that OTC markets are pervaded with adverse selection and that often dealers, rather than customers, are the less informed party in the trade (see, for instance, Easley and O'hara (1987) or Chalamandaris and Vlachogiannakis (2020)). Furthermore, this literature demonstrates that dealers protect themselves by trading only for sizeable spreads that cover their expected losses. Thus, informed dealers are expected to bid more aggressively and win a significant market share.

I show that the data is consistent with two implications of this theory. One implication of this claim is that dealers are likely to specialize in bonds that resemble each other, as learning about one of them lowers the cost of becoming informed about the other. I find that this is indeed the case with dealers specializing in bonds issued by the same firm and in firms operating in the same industries. A second implication is that dealers who dominate the market for a specific bond are expected to charge lower spreads. Using regression analysis, I demonstrate that this is indeed the case.

I embed the theory suggested here into a structural model and calibrate it to the behavior of spreads and volume before and during the COVID-19 crisis in the US Corporate Bonds Market in 2019. The calibration has a double purpose. First, it assesses whether the mechanism suggested here can account for the *magnitude* of the differences in the response of spreads and volume to systemic distress across markets for bonds with varying competition levels. Second, it allows the use of counterfactuals to disentangle the role of rising uncertainty, dealers' capacity constraints, and changes in the demand for liquidity in generating the deterioration in trading conditions during the COVID-19 crisis.

The model depicts customers who sell securities to dealers. It assumes that such customers can always find a dealer to trade with. However, most (if not most) dealers are uninformed about the asset's true value. Such dealers will require buying the bond at a discount to compensate for their exposure to adverse selection. As a result, informed players will have an advantage and face competition mostly between themselves. Indeed, the model implies the same pattern in the data, with dominant (informed) dealers charging lower spreads and facilitating most of the trade volume. Following Varian (1980)'s seminal model of sales, I model the competition between informed dealers as an auction with N informed dealers, such that with some probability π each of them may fail to submit a bid. I interpret π as stemming from search frictions and dealers' capacity constraints.

Following Camargo and Lester (2014), I explicitly incorporate asymmetric information into this framework. Thus, there is some probability that a bond is a "lemon". Customers, who hold the security, and informed dealers specializing in trading it, can distinguish a lemon from a "good" asset. Uninformed dealers, on the other hand, cannot. Thus, they are exposed to adverse selection and charge higher spreads to ensure that they are compensated for it. Since customers can always find an uninformed dealer, the value of that spread determines the outside option of a customer who trades with an informed dealer.

The theoretical analysis of the model reveals that the severity of adverse selection can steer the market into one of two distinct states. In the first state, intense adverse selection forces uninformed players to retreat from the market, opting to bid exclusively for less desirable assets, or ''lemons". In a market devoid of uninformed dealers, informed dealers may benefit from submitting low bids that attract only customers desperate for liquidity. This results in a strong transmission of shocks to liquidity demand to spreads and volume. Alongside, we witness a substantial disparity in spreads and volume between markets with varying concentration levels (or, number of informed dealers). In essence, without uninformed dealers, the competition among the informed becomes a more significant determinant of market outcomes. This heightened competition among informed dealers also becomes apparent in the considerable impact of these dealers' capacity constraints on spreads and volume.

In contrast, the second state occurs when adverse selection is less intense, allowing uninformed dealers to compete for orders involving both high-quality assets and 'lemons'. The competition they instigate significantly curtails the capacity of informed players to capitalize on customers' distress by charging higher spreads. Consequently, an uptick in liquidity demand and a tightening of dealers' capacity constraints mildly affect spreads and volume. Furthermore, the differences in spreads between markets with diverse quantities of informed dealers remain minimal.

I calibrate the model to the behavior of spreads and volume of a sample of bonds rated "'BBB-" before and during the COVID-19 crisis. I assume that the crisis results in 3 external changes to the model's parameters: higher demand for liquidity, manifested in lower reservation value of some of the customers, exhaustion of dealers' capacity, reflected in a higher likelihood of failing to submit a bid, and adverse changes in asset composition, embedded in a greater share of "lemons" in the markets and a wider gap between the value between lemons and regular assets.

My focus is on two groups of moments. The first are moments utilized for evaluating changes in asset composition. I estimate the probability of a bond being a 'lemon', meaning less valuable than suggested by its observable traits, using the likelihood of the bond being downgraded contingent on being re-rated. I assess the value of a lemon using the expected inter-dealer price of a bond rated "BBB-" following a downgraded. I find that the COVID-19 crisis period was marked by a dramatic depreciation in the value of a 'lemon' relative to a high-quality bond. Incorporating the exact values into the model reveals a level of adverse selection that would compel uninformed dealers to abstain from trading high-quality assets. This, consequently, enables the shock to customers' liquidity demand to have a sizeable impact on spreads.

The second category of moments serves as targets to calibrate parameters characterizing customers' demand for liquidity and dealers' capacity constraints. The model successfully hits these targets while implying reasonable parameter values. This demonstrates that the mechanism inherent in the model can account for substantial disparities in the response of spreads to systemic distress across markets with varying levels of competition. Moreover, the model replicates the differences in the volume response to the crisis across those markets almost perfectly, although these differences were not used as targets in the calibration. In other words, the concurrent behavior of spreads and volume in the data behave as if they operate under a competitive setting akin to the one represented in the model.

I use the calibrated parameters to study counterfactuals. I show that absent the changes in asset composition that exacerbated the severity of adverse selection in the COVID-19 crisis, we would have witnessed a much milder deterioration of trading conditions, especially in markets with higher concentration levels. In contrast, even if there had been no tightening of dealers' capacity constraints, the rise of adverse selection and the heightened demand for liquidity would result in higher spreads and lower volumes.

The paper's main contribution is in setting forth a novel mechanism underlying the decline in liquidity in OTC markets during times of crisis and providing empirical evidence to support it. In contrast to the literature highlighting that systemic distress hinders trade by imposing higher intermediation costs, the paper emphasizes how a crisis exacerbates the severity of *monopolistic inefficiencies*. It induces dealers to charge higher mark-ups and provide less intermediation services. In that sense, the dealer sector may "clogs" trade even when dealers do not face substantial constraints on their ability to load bonds into their portfolio.

This mechanism implies a novel link between an OTC market structure and financial fragility. I argue that during systemic distress, potential buyers exploit their market power more aggressively to buy securities at substantial discounts from sellers in dire need of liquidity. An OTC trading protocol exacerbates the losses from such a mechanism since it gives buyers more market power. This happens because each dealer serves multiple potential buyers. In other words, trade intermediation *contracts* the size of the buy side of the market. To put this argument into more concrete terms, if there are $n = 20$ potential buyers for a security, and each dealer serves $m = 10$ of them, a customer who wishes to sell it in an OTC setting is facing a market of 2 potential trade counter-parties (the dealers) rather than 20.

Secondly, to the best of my knowledge, this is the first paper to document a correlation between the concentration in the trade of a bond and the rise in spreads in trades of that bond during a crisis. I also show that this correlation holds when we compare bonds with similar attributes and control for potential confounders. Furthermore, the differences in the response of spreads to systemic distress across markets with varying levels are sizeable. If one believes them to imply a causal connection, they suggest that concentration substantially contributes to the overall deterioration of trading conditions in OTC markets in a crisis.

A third contribution of the paper is to document new empirical facts that highlight the segmented nature of OTC markets. It shows that a few dealers typically dominate the market for each bond and that this segmented structure seems to matter for market performance during a crisis. Hence, it indicates that using aggregate measures might often be misleading. Instead, analyzing the market should allow for substantial variation in the trade of different bonds (or bond classes). The paper also characterizes patterns in dealers' behavior that underlie this segmentation. It shows that dealers are highly specialized and that the activity of a typical dealer consists mainly of trading very heavily on a small subset of bonds in which it specializes. This finding sheds new light on what dealers do that may contribute to various agendas in OTC markets research.

The paper also sheds new light on the role of asset composition in shaping OTC market dynamics, specifically in its role in allowing the transmission of shocks to customers' liquidity to market spreads.

Literature Review

The paper is a part of the literature that addresses the question:"Why are OTC markets susceptible to failure during a crisis?". The literature has pointed out two causes that may underlie it. The first is *opacity.* Assets traded in OTC markets are characterized by high levels of heterogeneity and complexity. Also, the decentralized nature of the trade makes it harder to learn from actions taken by others (for instance, quotes are not published). The opacity generates adverse selection. An economic downturn exacerbates it by increasing the likelihood of defaults and imposing a more significant penalty for purchasing a risky asset. Hence, trade is diminished through an Akerlof (1978) lemon market mechanism. Papers that discuss such a mechanism include, for instance, Guerrieri and Shimer (2014), Camargo and Lester (2014), and

Zou (2019)).

The second cause discussed in the literature *limits on dealers' capacity.* Trade in OTC markets is facilitated through dealers. A crisis weakens dealers' balance sheets. As a result, they are less able to bear the risk involved in purchasing securities as a part of their market-making activity. They may also face a shortage of capital. The increase in dealers' costs will prevent them from meeting the heightened liquidity needs that emerge in systemic distress. Much of the literature that discusses limits on dealer capacity evolved in the context of the debate about the post-2008 regulation that placed new restrictions on bank-affiliated dealers. Papers such as Dick-Nielsen and Rossi (2019) or Bao et al. (2018) compare the response of the market to street events (i.e., index exclusion or downgrades) before and after the regulation kicked in. They show that following the regulations spreads increase faster in response to heightened demand for liquidity. In addition, the effect is more pronounced with bank-affiliated dealers, that is, those affected by the regulation.

The paper contributes to this literature by suggesting a third cause: monopolistic inefficiency. This is the first paper to tell that such a mechanism plays a crucial role in generating "market freeze" dynamics. By that, it highlights that competition (or its absence) in OTC markets is critical for understanding their stability. In this context, the paper

The paper is strongly related to the rich and fertile literature that applied search models to the study of OTC markets that originates in the canonical papers of Duffie et al. (2005) and Lagos and Rocheteau (2009). It shares with it the view that spreads in OTC markets embed markups charged by dealers. However, it suggests a different framework of thinking about the origin of the market power that enables markups to emerge. The search literature contends that a dealer's power over customers originates from search frictions that make it hard for the customer to find other dealers to trade with. Note that, at face value, this story does not appear very compelling. In contrast to finding a job (another prime application of search theory) that may require contacting hundreds or even thousands of firms to reach most openings, finding a dealer willing to buy should be more straightforward. More than 90% of the trade in the US Corporate Bonds Market is accounted for by the top 20 dealers. These are large and well-known players, and merely finding them requires no more than a phone call or an email [2]. In contrast, the

[2]Indeed, the Duffie et al. (2005) seminal paper states that the search frictions are an abstraction that is

paper suggests that the obstacle in OTC markets is that few dealers are willing to purchase each security. In other words, each customer faces only very few incumbents. In such circumstances, markups can arise without substantial search frictions (e.g., Cournot competition).

The paper also belongs to the literature that studies the impact of dealer market power on spreads. One segment of this literature consists of empirical papers attempting to disentangle dealer spreads into cost and markup components. For instance, Green et al. (2007) applies a production frontier setting to argue that the markup of a transaction is the difference between the spread taken by the dealer and the smallest spread charged for a similar transaction that occurred at about the same time. The novelty of this paper is inferring market power from *volume*. Since I gauge market power separately from spreads, I can better learn how one affects the other. Specifically, I present evidence of a clear and sizeable correlation between concentration in OTC markets and the response of bids to systemic distress. In this context, the documentation of the concentration in OTC markets has a solid connection to the documentation of a similar pattern among market-makers in stock markets by **?**. Like Schultz, I argue that the concentration results from informed players dominating the market-making activity.

The paper is a part of the literature on information acquisition. It applies Van Nieuwerburgh and Veldkamp (2010) and Veldkamp (2014) theory about under-diversification to explain concentration in OTC. According to the theory, costly information acquisition implies payoffs to specialization. An investor may prefer to hold a narrow portfolio since doing so allows her to become highly informed about her holdings. The seminal paper by Kacperczyk et al. (2005) demonstrates such under-diversification among mutual funds, and documents that funds that specialized in specific industries exhibit better performance. In an OTC setting, a recent paper by Chaderina and Glode (2022) demonstrates how expertise contributes to a dealer's return by increasing its order flow. That, in turn, allows the dealer to extract higher rents from encounters with less sophisticated investors.

The paper is structured as follows: Chapter 2 describes the data. Chapter 3 explores what underlies concentration in OTC markets and suggests that it originates from information acquisition. Then, Chapter 4 presents the model and elicits theoretical results connecting

underpinned by a more complex mechanism, such as a limitation on clearance or time required for dealers to acquire information about the security traded uncertainty with the impact of market power. Chapter 5 calibrates the model. Finally, Chapter 6 concludes and discusses policy implications.

2 Empirical Analysis

The empirical analysis consists of two parts. The first part characterizes concentration in OTC markets. It shows that while the market as a whole is quite competitive, the market for each bond is dominated by a few dealers. It also explores the causes that underlie this concentration. Given the findings, I suggest a theory of market power due to information acquisition. The second part studies the correlation between bond-level concentration and response to systemic distress. Focusing on the Covid-19 crisis, the paper documents that bonds traded by fewer dealers (higher HHI) exhibit a stronger spread increase and a more significant decline in volume during the crisis. Similar results stand for the 2007-2009 crisis.

2.1 Data

I use Regulatory TRACE data on the US Corporate Bonds Market, 2006 - 2020. This data is collected by FINRA, a self-regulatory authority that supervises broker-dealers. It provides detailed information for each secondary market for US corporate bonds trade. The data includes the CUSIP identifier of the bond, the time of the trade, the price, the volume traded, and more. The Regulatory version includes the identities of all dealers participating in the transaction. The identities uncover the share of each dealer in trading each bond.

I apply standard cleaning procedures common in the literature. First, I use **?**'s filter to account for corrections and cancellations of prior reports. Regulatory TRACE includes the report of both parties in a dealer-to-dealer trade. Those trades were removed. The Choi-Hu filter is used to remove transactions with non-member affiliates. The purpose of doing so is to ignore "bookkeeping" transactions a dealer makes with subordinate entities or other institutions owned by the same holding company. While TRACE began in 2001, during 2001 - 2005, many trades were exempt from reporting. Hence, only use data collected after January 2006 is used.

Data on bond attributes attained from Mergent Fixed-Income Securities Database (FISD) is joined with the TRACE data. The FISD data includes a rich specification of the terms of the bond. These include the bond's rating and its amount outstanding.

The sample used starts with all secondary market trades in corporate bonds executed from Jan. 1, 2006, to Dec. 31, 2020. Trades in the primary market and trade in a bond in the first 90 days after its issuance ("on the run") are removed, as well as trades in bonds with less than one year left to maturity. To be included in the sample, it is required that a bond is issued in US dollars by US firms that belong to one of the following three broad FISD industry groups: industrial, financial, and utility. Also, bonds with unique attributes that impact their pricing are filtered out, as is common in the literature. These include perpetual bonds, Yankee, sinking fund, and asset-backed bonds. Last, bonds that cannot be matched to Mergent FISD are removed. After the cleaning procedure, the data sample consists of 121 million observations.

Trades are divided into two types. The first is agency trades. In these trades, a dealer buys security while knowing upfront who he will sell it to later. The second type is principal trades, in which a dealer buys and holds the bond. One can think about the difference between these two types relating to who has the bond until a buyer is found. The bond is left with the customer in agency trades, while the dealer holds it in principal trades. This paper focuses on principal transactions for two reasons. First, these trades are critical for providing customers with an *immediate* injection of liquidity. Second, about three-quarters of the volume in the market is traded in a principal capacity.

2.2 Anatomy of concentration

A bond-level Herfindahl â Hirschman Index (HHI) is constructed for the analysis. The HHI index is one of the most common tools used in economic literature to measure market concentration. It provides a single number that summarizes to which extent is the majority of activity in the market dominated by a few players (in this context, dealers).

Let $V_{b,t}$ be the total volume of bond b traded in year t. Let $v_{d,b,t}$ be the volume of the bonds traded by dealer d in year y. Define the dealer share as

$$s_{d,b,t} = \frac{v_{d,b,t}}{V_{b,t}}.$$

Based, on it, the Herfindahl-Hirschman Index (HHI) at the bond level is defined as

$$HHI_{b,t} = \sum_{d=1}^{n} s_{d,b,t}^2. \tag{1}$$

If trade is conducted by many players, each having a small share, the measure approximates zero. If the trade is dominated by a single player, the HHI is one. If there are n players, each owning a share of $\frac{1}{n}$ of the market, HHI is $\frac{1}{n}$.

An HH-index to segments of the market is also generated. It is computed by applying Equation 1 to a subset of trades. For instance, the bond-level HH-index for principal trades between dealer and customer is computed by applying 1 to a subset of our sample that consists solely of trades between dealers and customers in which the dealer acts in a principal capacity.

Define $CRi_{b,t}, i \in \{1, 2, 3, 4\}$ as the share of trade in a bond b in year t that is facilitated by the i dealers who trade most extensively at the bond. Again, analogous measures for issuers are created and applied to different sub-segments of the market.

The concentration of bonds and the market structure that underlies it are explored in this section. One of the main goals is to understand the reason for the concentration in the data. The analysis in this section focuses on trades in the year 2019 conducted by the largest 50 dealers in the market. That dealers account for constitute more than 99.5% of all trades.

2.2.1 Concentration at the Bond Level

Figure 1 represents the distribution of CR3 - the share of trade in a market for a bond that is accounted for by the three dealers who trade it most intensively (i.e: have the largest market share in terms of volume).

The solid line boxplot represents the bond-level CR3 that is computed based on all trades in the market. For more than 50% of the bond, more than 55% of the volume is traded solely by three

Concentration in the US Corporate Bonds Market

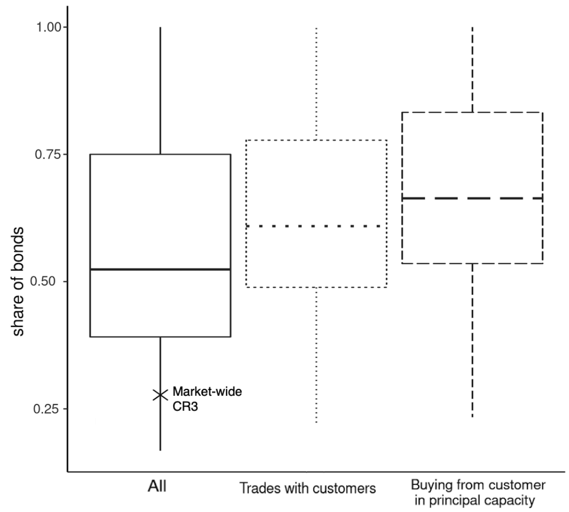

Figure 1: Distribution in the share of trade of the three leading dealers for each bond, 2006-2020.
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

dealers. For the sake of comparison, CR3 of the market as a whole appears on the plot as an X mark. The market-wide HHI is 27% - less than half of the bond-level median CR3. Also, on average, each of the three largest dealers in the market accounts for 9% of the entire volume. That does not seem, at least on the surface, to be a concentrated market.

Looking deeper into the data, it is possible to see that some dealers only operate only in the inter-dealer sector. These are usually alternative trading systems that dealers utilize for small trades. The presence of such dealers does not provide alternatives to a customer who is trying to sell a bond. The dotted plot is obtained by excluding those dealers, indicating a higher level of CR3. Further, some dealers only trade with customers in an agency capacity. They do not provide immediacy, which could be crucial in times of distress. Thus, to get an even better gauge of the actual options that a customer is facing when trying to sell a bond immediately (without waiting for the dealer to find a buyer), the CR3 for trades in which a dealer buys a customer from a bond in a principal capacity is represented by the dashed boxplot. Its implied concentration is even greater, with the median bond having more than 70% of the total trade accounted for by three prominent players.

Figure 2 exhibits an analogous plot for HHI. The median HHI is about 0.38. To get an intuitive understanding of it, imagine a market that only has three traders. The first trader is responsible for 40% of all volume, the second one also for 40%, and the third one for 20%. This highly

Concentration in the US Corporate Bonds Market (measured by HHI)

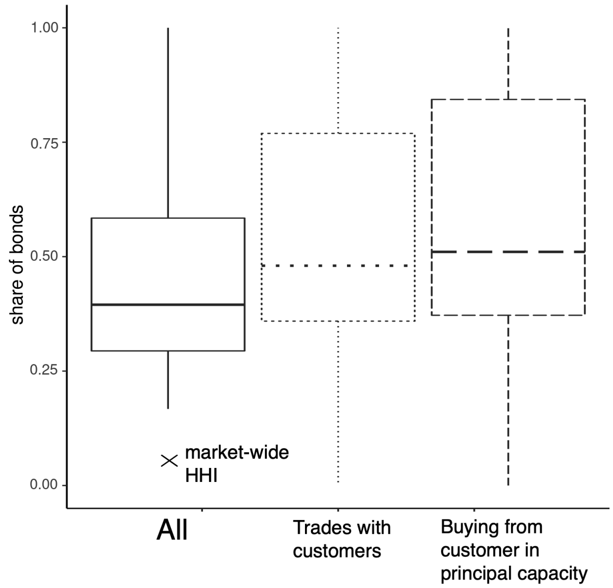

Figure 2: Distribution of the bond-level HHI, 2006-2020.
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

concentrated fictitious market has an HHI of 0.36, which is a lower level than the one found for the median bond market.

Further, it is possible to see that concentrated markets account for a sizeable part of total trade in corporate bonds. Figure 3 plots the cumulative volume (y-axis) that was bought by the dealer in principal capacity when trading bonds with an HH-index of h of below it (x-axis, HHI)). About 50% of the volume sold by customers (bought by dealers) in the year 2019 was generated in trades of bonds with an HH-index of 0.4 or more.

2.2.2 The Causes Underlying Concentration

This section explores patterns in the data that hint at the causes of concentration in the bonds market. It begins with the low HHI of the entire bond market, which implies that the issue at stake is *not* barriers to entry in the dealer sector. Indeed, there are hundreds of dealers, with neither of them holding a share of the market large enough to grant any sizeable market power.

Concentration emerges only when the examination happens at the level of each separate bond. This implies market segmentation. To get a deeper understanding of what underlies it, we take

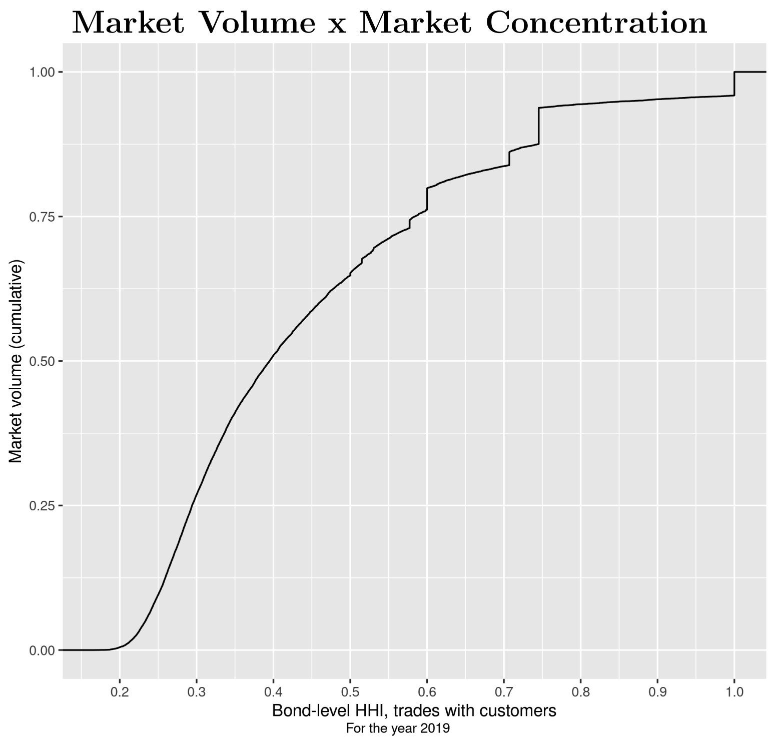

Figure 3: Markets with high concentration account for a substantial share of total market volume.
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

a deeper look into the trade activity of each dealer. Specifically, the goal is to see if dealers are specialized, that is, if they trade heavily on some bonds while being less active in others. For that, one can consider a dealer d as specializing in trading of bond b if its share in volume traded of that bond is at least three times greater than his share in the total volume of trade in the market. Figure 4 plots the cumulative share of dealers (y-axis) for which at x% of their total volume or less was generated in trading bonds they specialize in.

For about 70% of dealers, trades in bonds they specialize in account for at least 50% of the trade. This strongly contrasts with what would be expected based on a standard diversification argument. To mitigate risk, dealers should prefer to hold a broad portfolio that reduces their exposure to the performance of any specific security or issuer. Nonetheless, the exact opposite occurs. The information story suggested below implies this happens for the same reasons explained in Veldkamp (2014). It originates from a trade-off between having a broad and diversified portfolio and being well-informed about the assets that constitute it.

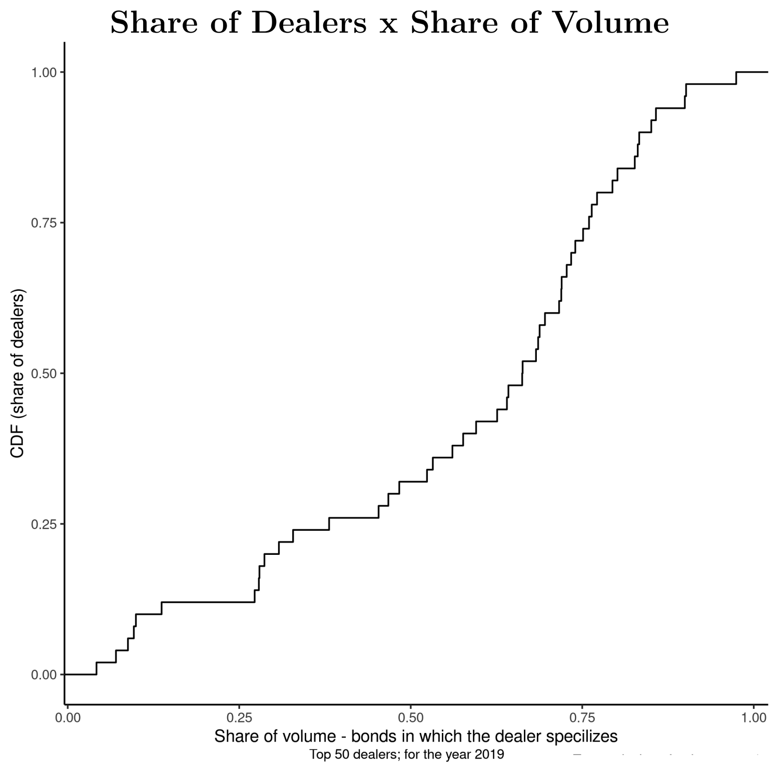

Figure 4: A dealer trades mostly in bonds in which it specializes.
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

To get a closer glance at dealer specialization, define

$$\hat{\delta}_{b,d,t} = \frac{s_{b,d,t} - \bar{S}_{d,t}}{\bar{S}_{d,t}}$$

as the deviation of dealer d share in bond b from its share in the total volume of trade. The latter would have been the expected share if the dealer randomly traded bonds. Its distribution is plotted in Figure 5.

As shown above, the majority of trade of dealers happens in the small subset of bonds (less than 15%) in which they specialize. For other bonds, it is possible to define two groups. One group consists of bonds that the dealer trades sporadically, ($\hat{\delta}_{b,d,t} \in (0.99, 0]$). The other are bonds that the dealer never trades ($\hat{\delta}_{b,d,t} = -1$). From the graph, we see that the latter includes about 45% of all bonds. Indeed, no single dealer ever trades more than 65% of bonds in this market. Further, this result holds if the sample is expanded to start in 2006 and finish in 2020. That is the upper bound on the share of bonds that the dealer trades does not seem to be the result of randomly trading bonds it does not specialize in. Rather, it seems that there are some bonds

Histogram of Dealer Share of Bond as pct. of its Share in Total Volume, 2019

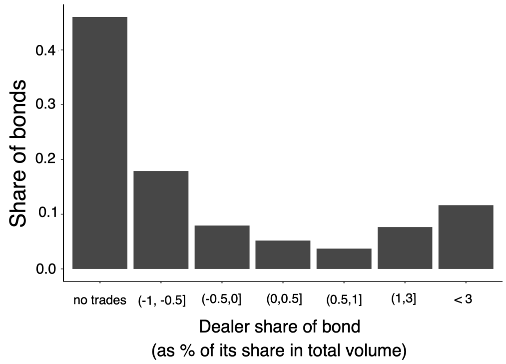

Figure 5: A dealer specializes in a bond, trades it sporadically, or avoids trading it altogether.
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

that dealers intentionally avoid. This pattern has reminiscent of an entry cost, that is, either the dealer pays it and trades in a market, or he does not and does not operate in it at all.

The entry cost flavor of the results gains further support from studying attributes that characterize bonds with high levels of concentration. Table 1 shows the results of the regression model

$$HHI_{b,t} = X_{b,t}\beta + \epsilon$$

, where $X_{b,t}$ is a set of bond attributes in year t. These include the mean rating given to the bond according to Mergent FISD, the coupon rate, year to maturity, the age of the bond (years since it was offered), and a categorical variable for the amount outstanding. Indicator variables for a bond that pays a coupon in a non-standard frequency, a bond that falls within the ambit of rule 144a, and a callable bond are also included. Last, industry group fixed effects and issuer fixed effects in some of the models are considered.

The best predictor of a bond's concentration level is the amount outstanding. Adding it to a regression that does not include controls for the bond's issuer props up the R-squared from 0.08 to 0.21. A low amount outstanding is correlated with a higher level of concentration. The amount outstanding can be interpreted as a proxy for the size of the market. Hence, the regression implies that larger markets have more participants. This is exactly what would be expected in an entry cost setting: the larger the market, the greater the number of players that can enter and share it before doing so no longer justifies paying the entry cost.

Dependent Variable:	bond-level HHI			
Model:	(1)	(2)	(3)	(4)
Variables				
rating	-0.0008	-0.0045***	0.0108**	0.0082**
	(0.0021)	(0.0007)	(0.0026)	(0.0016)
coupon rate	0.0051	0.0080**	0.0077**	0.0079***
	(0.0038)	(0.0018)	(0.0015)	(0.0010)
non-standard coupon frequency	0.0650***	-0.0074	-0.0447*	-0.0423*
	(0.0070)	(0.0257)	(0.0179)	(0.0150)
rule 144a	0.0459***	0.0464***	0.0608***	0.0353***
	(0.0063)	(0.0052)	(0.0056)	(0.0043)
callable	0.0565***	-0.0356	-0.0205	-0.0213
	(0.0050)	(0.0220)	(0.0364)	(0.0333)
time to maturity	0.0007	0.0005*	0.0003	0.0003
	(0.0003)	(0.0002)	(0.0003)	(0.0003)
time since offering	0.0044***	0.0016	0.0045**	0.0029**
	(0.0006)	(0.0010)	(0.0009)	(0.0008)
issue size: 0-100m		0.1501**		0.1095***
		(0.0271)		(0.0143)
issue size: 100m-500m		0.1007***		0.0714***
		(0.0086)		(0.0047)
issue size: 500m-1t		0.0379***		0.0275***
		(0.0025)		(0.0044)
Fixed-effects				
industry	Yes	Yes	Yes	Yes
issuer			Yes	Yes
Fit statistics				
Observations	460,279	460,279	460,279	460,279
R^2	0.08608	0.21349	0.54168	0.56487
Within R^2	0.05716	0.18861	0.04507	0.09338

Clustered (industry) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 1: Regression - attributes of bonds traded in concentrated markets.
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

Another salient thing about the regression above is the impact of controlling for the bond issuer. It dramatically increases the model's predictive power. That suggests that dealers specialize in bonds issued by the same entity. To further explore this possibility, we run the following regression:

$$S_{d,b,t} = \beta_0 + \beta S_{d,i,t,-b} + \epsilon,$$

where $S_{d,b,t}$ is the share of dealer d in the total volume of bond b in year t, and $S_{d,i,t,-b}$ is the share of d in the volume of trade in all securities issued by dealer d besides security b. The results appear in Table 2:

Dependent Variable:	dealer share in trading the bond
Variables	
(Intercept)	0.0123^{***}
	(9.34×10^{-5})
dealer share issuer (but bond)	0.9780^{***}
	(0.0019)
Fit statistics	
Observations	584,798
R^2	0.30679
Adjusted R^2	0.30678

IID standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 2: Trades in the year 2019
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

A 1 bps increase in the share of dealer d in all securities of issuer i besides bond b presages a 0.98 increase in the share of that dealer in trading bond b. That is, if one would like to predict the percentage of trades of a dealer in a bond based on its share in other bonds issued by the same firm, the best guess would be to say that they are about the same. Also, the prediction explains a substantial part of the variance, as implied by the relatively higher R-squared of 0.3.

These results imply that dealers specialize in trading specific issuers rather than bonds. This pattern is consistent with what would be expected if market segmentation originated because of information asymmetries. Knowledge of the issuer allows a dealer to assess its default risk. That, in turn, helps the dealer to price its bonds more appropriately.

If this is true, dealers will specialize in firms that resemble each other. The resemblance means that information that is acquired to assess the resilience of one can help evaluate the resilience of the other. To test this, we run the following regression:

$$S_{d,i,t} = S_{d,j,-i} + \epsilon,$$

Where $S_{d,i,t}$ is the share of dealer d in the volume of all bonds issued b in year t, and $S_{d,j,t,-i}$ is the share of d in the volume of trade in all securities of all issuers from sector j besides i in year t. Issuers are assigned to sectors according to their 5-digits NAICS code from Mergent-FISD. The data includes issuers with 825 different NAICS codes. The regression results appear in Table 3.

Dependent Variable:	dealer share issuer
Variables	
(Intercept)	0.0134^{***}
	(5.4×10^{-5})
dealer share industry (but issuer)	26.27^{***}
	(0.0616)
Fit statistics	
Observations	670,459
R^2	0.21346
Adjusted R^2	0.21346

IID standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 3: Trades in the year 2019.
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

Again, there is a statistically significant positive correlation between the dealer share in trading firms in a particular industry and the likelihood that it trades the bonds of an issuer belonging to this industry. There is also a relatively high R-squared. Knowing only the dealer's share in trading, other issuers in the industry explain much of the variance in the data. Another interesting evidence of specialization comes from comparing each dealer's trade share in the trade in all issuers belonging to a specific industry (by NAICS code), which is $1.79e^{-04}$, to the median of the dealer share in all issuers in general, equaling $1.29234e - 05$.. The former is more than ten times greater than the latter. That is, the distribution of the share of each dealer in each industry has a fat right tail: some dealers trade much more heavily on bonds from a specific sector than all others.

These patterns reverberate a state in which dealers specialize in bonds that they are more familiar with. Such a state need not be surprising. A dealer that overpays for a bond is likely to end up incurring losses when selling it for a lower price at a later date. Considering the small profit margin of dealer activity (most estimates of spreads in regular times are 15 - 30 bps), even a tiny error can result in losses. In this context, papers like Easley and O'hara (1987) and Chalamandaris and Vlachogiannakis (2020) show that OTC markets are pervaded with issues of asymmetric information and that it often happens that the dealer is less informed than the customer. Further, they demonstrate that dealers compensate for their expected losses from adverse selection by charging higher spreads. Thus, if some dealers are better informed about a specific security, we would expect that they would be able to offer better prices. That, in turn, would allow them to win a larger share of the market.

An immediate implication of such logic is that prominent dealers will charge lower spreads. To test it, I run a regression model of spreads on the dealer's prominence level in the market for the bond that was traded:

$$\text{spread} = \beta \text{dealer prominence}_{t-1} + \text{bond} * \text{trade size} * \text{date} + \text{dealer}$$

Where *dealer prominence* is a categorical variable that is assigned with the value "Non-active" if the dealer traded the bond less than five times in principal capacity, "Active", if the dealer traded it more than five times in principal capacity but accounts for less than 10% of total trade volume, and "Prominent" if the dealer traded the bond at least five times and accounts for more than 10% of total trade in the bond. I use the dealer prominence from the previous year to avoid bias due to reverse causality. The *trade size* variable categorizes trades into volume bins of $> \$100,000$, $\$100,000 - \$500,000$, $\$500,000 - \$1,000,000$, $\$1,000,000 - \$5,000,000$, and $\$ > 5,000,000$. The regression is applied to all principal trades occurring in the years 2006-2020.

The regression compares trades of the same bond (identified by cusip), at about the same volume, conducted on the same day, that were made by dealers with varying levels of prominence in the market. It also controls for dealer fixed effects to avoid potential selection bias. Its results appear in table 4

What we can see is that dealer prominence is indeed associated with lower spreads. The differences are sizeable. For instance, we see that when a dealer who is prominent in the market for a specific bond purchases it from a customer, he charges a spread that is 12 bps lower than that charged by a non-active dealer. In other words, it pays for it more. As spreads typically

trade type	C2D		D2C	
Model:	(1)	(2)	(3)	(4)
Active	-0.2227 (0.5272)	-1.178** (0.5568)	2.218*** (0.3878)	-0.4710 (0.3933)
Prominent	-11.58*** (0.5805)	-3.841*** (0.6147)	-8.617*** (0.4351)	-2.611*** (0.4511)
bond - trade size - date	Yes	Yes	Yes	Yes
dealer		Yes		Yes
Observations	10,702,378	10,702,378	12,097,754	12,097,754

Table 4: Dealer prominence and spreads; All principal trades in the years 2006 - 2020

Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

oscillate in a band of 15-30 bps, this is a remarkably substantial difference. When we add controls for dealer fixed effect, the difference is diminished to 3.84 bps. This might indicate a division of roles in the market - some dealers specialize in specific bonds and offer better prices for them, while other dealers are willing to trade anything but are compensated by a higher markup. Yet, also when taking those into account, the differences in pricing are quite large. We see a similar pattern when a dealer sells to a customer and when comparing active and non-active dealers (although there the results are more ambiguous).
The information acquisition logic and the entry-cost patterns seen above imply an interesting depiction of OTC markets. According to it, dealers can pay the cost of acquiring information about the value of the security. The information protects them from adverse selection. Hence, they can compete more aggressively and win a larger market share. This depiction gains further support from a recent paper by Brancaccio et al. (2017). The paper shows that the complexity of a municipal bond presages a more significant market share and higher profits for the bond underwriter. They argue that this is because complexity increases the cost other incumbents need to pay to learn about the value of the security. Further, using close elections as an IV, they show that when underwriters are subject to more lax regulations (republicans win), they tend to issue more complex municipal bonds, allegedly to increase the cost of entering the market for other incumbents.

Lastly, given that dealers specialize by issuer, one might think it is proper to consider the market segmented across issuers rather than bonds. To see if this is the case, Figure 6 focuses on dealers accounting for more than 10% of the trade in an issuer's bonds. It plots the deviation of their share of trade in each bond from their share of trade in all the issuers' securities.

When a dealer trades an issuer's securities, it trades many of them quite heavily. However, it

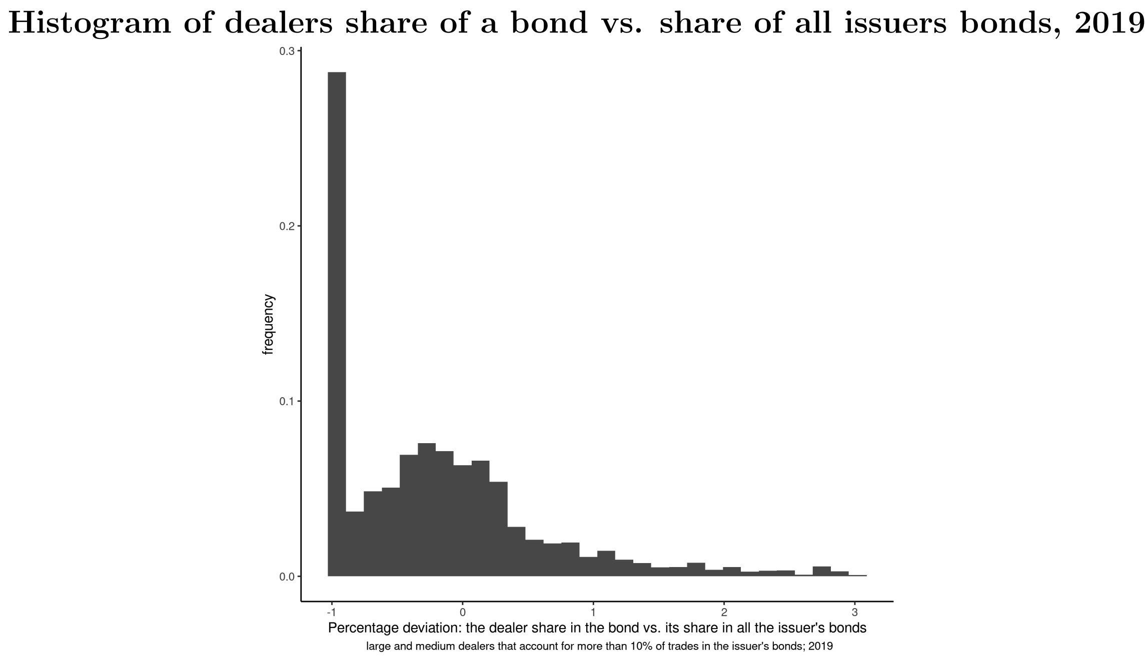

Figure 6: Trading an issuer does not mean trading all its bonds.
Note: The data relied upon to generate this figure was provided by FINRA's TRACE System.

also refrains from trading others altogether. That is, while they focus on securities of a subset of issuers, within each issuer, they focus on a subset of bonds. For instance, that could result from the type of clients the dealer is working with. For example, bonds with one year left to maturity are traded by money-market funds. It makes sense that only dealers connected to such players will intermediate such bonds, even if the dealer is familiar with the issuer and can assess the likelihood of default. Also, bonds can be highly heterogeneous and complex assets. Thus, familiarity with the issuer is insufficient to determine its security value. Given this context, Table 1 shows that bonds with a low amount outstanding will typically be traded by fewer dealers vis-a-vis other bonds issued by the same issuer. Again, the amount outstanding can be considered a proxy for market size, and the larger the market for a bond, the greater the number of dealers that will pay the cost of assessing its actual value. Because of this pattern, markets will be regarded as segmented at the bond level throughout this paper.

2.3 Concentration and Response to Systemic Distress

2.3.1 Markups During Regular and Crisis Times

Before I proceed with the analysis, I would like to provide some additional evidence for a claim that appeared already in the introduction - that is, that in times of crisis, dealers are charging higher markups. I made a claim based on literature that uses measures of spreads based on the difference between the price of a bond in interdealer trades and its price in transactions between customers and dealers. This literature documents a substantial rise in these measures during a crisis (see, for instance, O'Hara and Zhou (2021)). Since, in both cases, the purchasing dealer assigns the same value to the security (considering his holding cost, ability to bear risk, etc.), these measures probably capture a markup component. I will elaborate on the measures from this literature and apply them myself in section 2.3.2. However, before doing so, I would like to present another piece of evidence that demonstrates that dealers buy from (sell to) customers at cheaper (higher) prices compared to those in which they buy from (sell to) other dealers. The purpose of doing so is to ensure that the result does not originate from attributes of the spread measures used in the literature and that they also hold when considering the quantity traded impact on the price.

For that purpose, I run a regression that measures the difference in price between a transaction made by two dealers and a transaction of the same bond traded on the same day at about the same quantity between a dealer and a customer. Note that the simplicity comes with a price - the regression will only include bonds traded by two dealers and a dealer and a customer at about the same volume on a given day. Hence, it selects liquid bonds that typically have lower spreads. Therefore, I advise the reader to focus on the relative sizes of the coefficients appearing in the regression rather than their absolute size.

The regression model is:

$$price = Bond \times Date \times trade\_size + D2D + D2D \times Crisis \tag{2}$$

Where the dependent variable, *price*, is the price at which a transaction occurred. On the right-hand side, I control for an interaction term between the bond, the date of trade, and a trade size categorical variable that divides the trade sizes into the ranges: $[100k, 500k), [500k, 1m), [1m, 2m), [2m, 5m), 5m$ The $D2D$ is a dummy variable that takes the value one if both sides of the trade are dealers (zero otherwise), and the variable $D2D \times Crisis$ takes the value one if both sides of the trade are dealers and the trade occurred during a great recession or the COVID-19 crisis.

I run the regression on a sample of the data that consists of the Great Recession (July 2007 - May 2009) and the period from January 2006 to its beginning, and the COVID-19 crisis (March 5th - April 10th, 2020). and the period between January 2019 and its eruption. Including the periods that proceed the crisis is meant to allow a comparison of crisis periods to near times in which market conditions, absent the distress, are assumed to be similar. I limit my attention to transactions of wholesale trades (volume of more than $100,000) made in a principal capacity. I apply the regression to two subsamples of the data. One includes trades in which customers sell to dealers (C2D) and interdealer trades and is used to compare the price a customer gets paid when selling a bond to the price that a dealer gets for selling it. Similarly, the other subsample consists of trades in which customers buy from dealers (D2C) and interdealer trades and is used to measure the difference between the price a customer pays and the one paid by a dealer.

The regression results appear in 5. It indicates that the difference between the price paid to a customer and the price paid to a dealer in regular times, captured by the $D2D$ dummy variable, is indeed positive when comparing C2D and D2D trades, meaning dealers get paid more for selling the same security on the same day at about the same quantity. The difference of 26 bps is statistically significant. Similarly, when comparing D2C and D2D trades, we find that dealers pay prices 26 bps lower than those paid by customers. More importantly, looking at the discount at which buyers sell a security during a crisis (compared to the prices that dealers get to sell it), captured by the sum $D2D + D2D \times Crisis$, we find that it is 44 bps, or 65% greater than its size in normal times. About the same relative price, change is also found when customers buy from dealers during a crisis, where they pay an additional 43 bps compared to what dealers pay for the same transaction.

In columns (3) and (4) of the table, we see that the result holds also when controlling for a dealer-date fixed effect. That is, it is not driven by the selection of the type of dealers that engage in each type of trade in normal and crisis times.

Dependent Variable:	price			
Trade Type:	C2D	D2C	C2D	D2C
Model:	(1)	(2)	(3)	(4)
Variables				
D2D	0.2603***	-0.2663***	0.2346***	-0.1883***
	(0.0065)	(0.0014)	(0.0031)	(0.0032)
D2D × Crisis	0.1734***	-0.1663***	0.1810***	-0.1506***
	(0.0102)	(0.0123)	(0.0071)	(0.0254)
Fixed-effects				
Bond × Date × trade_size	Yes	Yes	Yes	Yes
Dealer × Date			Yes	Yes
Fit statistics				
Observations	6,981,016	7,966,969	6,981,016	7,966,969
R^2	0.98731	0.98698	0.99154	0.99012
Within R^2	0.00128	0.00270	0.00132	0.00099

Clustered (cusip_id-trd_exctn_dt-t_size) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 5: Regression results, equation 2 applied to the Covid-19 crisis and the period proceeding it.
Note: The data relied upon to generate this table was TRACE Data provided by FINRA's TRACE System.

The takeaway from this simple regression is that at least a part of the increase in spreads that dealers charge customers during times of systemic distress is a markup. It is a markup in the sense that it does not manifest changes in the value that the dealer assigns to holding the security due to capacity constraints, adverse selection, or any other factor. That does not mean that it is not related to such factors. For instance, the capacity constraints of some dealers mean that others are facing weaker competition and can charge higher spreads (and I explore such connections below). However, it does mean that what we are seeing here is not dealers rolling on higher costs of facilitating trade to their prices. Rather, we witness dealers exploiting, if you will, the crisis circumstances to squeeze greater profits from trading with customers. This pattern, in turn, is intimately related to competition in OTC markets.

2.3.2 HHI and Spreads in Regular and Crisis Times

Typically, OTC markets data shows realized prices, rather than asks and bids. Further, many bonds are traded quite infrequently. It may take days, weeks, or months between the time we see the bond sold and the time we see it was bought. These circumstances make it measure bid-ask spreads directly from the data. Thus, it is also a challenge to discern the cost of intermediation and disentangle its changes in the security fundamental value.

Here, a measure of spreads first suggested by O'Hara and Zhou (2021) is used. This measure uses the percentage difference between the price of a bond in a specific transaction and the price of the most recent trade of the same bond in the dealer-to-dealer market.

The underlying assumption made when using this measure as a benchmark is that dealers operate in a more-or-less frictionless market. Hence, the prices of deals between dealers should be close to the fundamental value of the asset. Also, from a more pragmatic point of view, the dealer-to-dealer market is quite active. It is easier to find the last time a bond was traded in that market rather than match a bond buy with the last time it was sold (or vice-versa).

The measure is:

$$spread_j = ln(\frac{\text{trade price}_j}{\text{Benchmark price}_j}) * \text{trade sign}_j \tag{3}$$

The benchmark is the price of the most recent dealer-to-dealer principal trade in the bond. For brevity, I will refer to this trade as the *benchmark trade*. The trade sign is 1 for a customer buying from a dealer and (-1) for a customer selling to a dealer. The spread measures the percentage difference between the price that a customer gets paid (pays) when selling (buying) a bond versus what a dealer gets paid (pays) when selling (buying) the exact same security. Typically, dealers will attain favorable prices, that will be manifested in a positive measure. It is widely held that the differences in pricing originate from the fact that dealers, unlike customers, are well-immersed in the market and are better able to attain multiple offers for their spreads. In other words, the measure captures exactly what the paper focuses on a markup that is imposed on customers by dealers that exercise their monopolistic power. That power, in turn, originates from customers' limited ability to find alternative dealers to trade with. In this context, it is also important to note that the measure is not likely to reflect dealers' holding costs of the security. To see that, focus for the time being on cases in which a customer sells to a dealer. In those cases, the measure is constructed by comparing the price of two trades in which a dealer buys the same bond: one in which he buys it from a customer, and another one, close in time, in which he buys it from another dealer. Trivially, as in both trades, the same security is purchased its holding costs should be the same.

For the coming analysis, I apply a few additional filters to the data. I limit my attention to risky principal trades. I do so for two reasons. First, these trades account for more than 75% of trade volume. Second, as rightly pointed out by Dick-Nielsen and Rossi (2019), the spreads for these trades indicate the cost of immediacy or the price that a customer pays for attaining liquidity in a very short time frame. The rise in the cost of getting liquidity in *timely* manner is a main factor in the dynamics of a crisis in OTC markets. Also, I ignore retail trades and incorporate only transactions with a volume of $100,000 or more. I do so since retail trades constitute the majority of observations but only 10% of total trade volume. Hence, incorporating these observations may generate results that apply to them and actually have very little impact on most of the volume traded. To preserve consistency with previous parts of the work, I focus only on trades conducted by the top 50 dealers, who account for more than 99% of trade volume. I measure concentration using bond-level HHI generated from trades in which dealers buy from customers in a risky principal capacity. To the best of my understanding, this measures the concentration that has an impact according to the theory that underlies this paper, that is, the concentration that dictates the alternatives that a customer who wishes to sell a bond in principal capacity is facing. To avoid bias due to reverse causality, I calculate the HH index for each transaction based on trades in the bond traded in the previous year, $t - 1$.

Alongside this, I apply a few filters to address potential weaknesses of the spread measure. Specifically, I ignore trades in which the time that elapsed between the transaction and the benchmark trade, that is, the last D2D trade at the same bond, is smaller than 15 minutes. Such cases are likely to be agency trades that were mistakenly categorized as principal. Regarding them as such creates a downward bias in our measurements, as agency trades are characterized by lower bid-ask spreads. That is more true for less liquid bonds, as the infrequent trade makes such a rapid sequence of trades in the bonds highly unlikely. At the same time, I ignore trades in which the time that has elapsed since the benchmark trade took place is greater than seven days. I do so to minimize the effect of fundamental changes in the price of the bond over time on the differences that the data exhibits between customer-to-dealer versus dealer-to-dealer bonds. Further, as a robustness test, I run the tests of my hypothesis also on a more limited sample of transactions for which the benchmark trade occurred at the same day as the trade itself. As we shall see, the main results persist when doing so. I winsowrize the spreads to diminish the risk of bias due to outliers. I set the higher bound of the spreads to be the 95% quantile plus 1.5 times the interquartile range, and the lower bound to the 5th quantile minus 1.5 the interquartile range.

Here, I present a set of results attained from regression analysis. The analysis does not meet the gold standard of studying causality in economics. That is, the assignment of the "treatment" (concentration) is not random. Therefore, I cannot negate the possibility that some third element that determines bond-level concentration also underlies its responsiveness to systemic distress. However, the regression analysis allows me to refute well-specified alternative explanations of the correlation by adding controls. Specifically, I show that the strong correlation between concentration and the rise in spreads in a crisis persists when we account for the impact of the bond liquidity, the dealer expected holding period, the dealer identity, the changes in the composition of bonds traded throughout a crisis, and more. The analysis also allows me to demonstrate the overall robustness of the result, that further lowers the likelihood of it originating from a selection bias.

I apply a transaction-level regression model of the form:

$$\begin{aligned} Spread_{i,d,b,t} = \alpha_0 + \beta_1 * HHI_{b,t-1} + \beta_2 * HHI_{b,t-1} * \mathbb{I}\{Crisis\} \\ \gamma * X_{b,t} + \eta_d + \xi_b + \nu_i + \text{trade\_date} + \\ \eta_d * \text{trade\_date} + \xi_b * \text{trade\_date} + \nu_i * \text{trade\_date} + \\ \lambda_b + \lambda_b * \text{trade\_date} \end{aligned} \quad (4)$$

In the LHS of the regression, we have the spread of trade i done by dealer d in bond b at time t. Our main object of interest is β_2 - the coefficient on the interaction term between the concentration level of the bond and the trade being conducted during a crisis. Note that the regression has a dif-in-dif flavor to it, with the bond HHI measuring the size of the treatment and the crisis indicator being the post-treatment dummy. The value of β_2 determines to which extent does the treatment, that is, the HHI, presages a change in the outcome, which is the spread. To diminish bias due to changes in fundamental prices over time, I weigh trade by the inverse of the distance in time to the D2D trade that determines the benchmark price for calculating the spreads.

The regression includes an elaborate set of controls. First, the vector $X_{b,t}$ is a vector of bond attributes that include: age, time to maturity, square root of amount outstanding, the issuer industry, and more (for the full list, see Table 12 in the appendix). Alongside, I include fixed effects for the dealer's identity, ξ_d, the bond rating, ξ_b, the trade size, ν_i, and the bond's liquidity. I transform the trade size to a categorical variable by dividing it into the following bins: [$100-$500), [$500k-$1m), [$1m - $5m), and above $5m. I measure liquidity by the number of days in which the bond was not traded in the previous year - a common measure in the literature. I use only trades in which dealers trade bonds with customers, to avoid a mechanical correlation between this measure and the concentration measure, stemming from the fact that markets with more dealers will have a higher frequency of dealer-to-dealer trades (Below, I show that my results also hold with an alternative liquidity measure that incorporates these dealer-to-dealer trades into the regression). I divide the measure into bins that correspond to five percentile ranges of the liquidity variable.

Further, the regression includes an interaction term of each of the fixed effects with a date categorical variable. If I would have used an interaction term of each fixed effect with the crisis indicator, it would have captured the expected change in the spread of each trade given the attribute represented by the fixed effect. For instance, a dealer-crisis fixed effect would have captured the average addition of each specific dealer to the spreads it charges in crisis times. The interaction of a fixed effect with the date variable does even more. Continuing with the same example, the interaction of a dealer and date fixed effects not only (indirectly) embeds the changes in the spreads dealers charge during a crisis, it also controls for changes in these spreads throughout different phases of the crisis period. To see that, assume that less liquid bonds are traded in more concentrated markets and that they are typically traded only deep into the crisis when spreads are higher. In that case, the interaction term of date and liquidity will capture the expected change in the spread of a bond with a certain liquidity level, given that it was traded at a specific time in the crisis. That, in turn, will prevent these changes from contaminating the estimate of β_2.

In table 6, I report the results of applying the regression model to the COVID-19 crisis and the period proceeding it (columns 1-3) and to the 2007-2009 crisis and the period proceeding it (columns 4-6) [3]. In each, I run the analysis on the full sample, only on trades in which a customer sells to a dealer (C2D) and only on trades in which a dealer sells to a customer (D2C). What we can see is that in both crises period, higher concentration was correlated with a substantial increase in spreads in trades in which a customer sells to a dealer. In the great recession, a 10 bps increase in HHI was correlated with a 1.6 bps higher "jump" in the spreads in customer-to-dealer trades in a crisis. In the COVID-19 crisis it was correlated with a staggering 9 bps rise in the gap between spreads pre-crisis and during the crisis period. In both cases, the result is statistically significant. This pattern is exactly what we predicted if dealers indeed exploit the distress of customers in dire need of liquidity to charge higher markup for inter-mediation. Interestingly, for cases in which a dealer sells to a customer, we see the opposite effect - higher concentration presages a milder increase in the spreads in times of systemic distress. In absolute size, the effect is about half the size of what we see with customer-to-dealer trade. In the next few pages, as we turn to take a closer look at the regression analysis, we see that this correlation, while statistically significant in this specification, is not robust and is likely to be explained away as a spurious correlation originating from limitations of the spread measure.

For that, I begin by running the regression model in equation 4 with one minor change - rather than treating the HHI as a continuous variable, I divide it into bins and regard it as a categorical variable. I choose bins at the length of 10 bps, besides the bins at the edges, 0-0.3 and 0.6-1, that span over a wider range to compensate for these levels being much more sparsely populated. The coefficients and 95% confidence intervals from applying these regressions to the COVID-19 crisis appear in plot 7. Again, they present these coefficients for the entire sample (black), for trades in which a customer sells to a dealer (red), and for trades in which a dealer sells to a customer (blue). Let us begin with trades in which a customer sells to a dealer. We see a pattern that looks almost like a linear upward slope, indicating that higher HHI levels presage a greater increase in spreads. The results are statistically significant, with the 95% confidence interval being far above zero. An important feature to note is the positive and sizeable coefficients at HHI levels of 0.3-0.4 and 0.4-0.5. This indicates that even when comparing a competitive market, that is HHI of 0-0.3, to somewhat less competitive markets, like the one with an HHI of 0.3-0.4, we see a change in the response of spreads to a crisis. This is critical as about 80% of trade

[3] For full results that also report the coefficients on control variables, see table 12 in the appendix

Dependent Variable:	spread					
	COVID-19 Crisis			GRC		
Population:	All	C2D	D2C	All	C2D	D2C
Variables						
HHI	5.7*** (1.4)	13.7*** (1.1)	0.64 (0.92)	-0.01 (5.9)	9.3** (4.1)	-0.21 (3.2)
HHI × crisis indicator	27.6*** (8.4)	92.0*** (9.8)	-43.2*** (9.6)	2.0 (8.4)	16.0*** (5.7)	-8.4* (4.3)
Fixed-effects						
rating	Yes	Yes	Yes	Yes	Yes	Yes
trade size	Yes	Yes	Yes	Yes	Yes	Yes
dealer	Yes	Yes	Yes	Yes	Yes	Yes
dealer-date	Yes	Yes	Yes	Yes	Yes	Yes
date	Yes	Yes	Yes	Yes	Yes	Yes
# days traded (prv. year)	Yes	Yes	Yes	Yes	Yes	Yes
# days traded (prv. year)-date	Yes	Yes	Yes	Yes	Yes	Yes
rating-date	Yes	Yes	Yes	Yes	Yes	Yes
trade size-date	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	1,375,793	656,133	719,660	621,479	271,191	350,288
R^2	0.11938	0.25548	0.21531	0.19310	0.36171	0.36307
Within R^2	0.00693	0.00936	0.01212	0.00485	0.00735	0.00997

Clustered (rating) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 6: Results from applying the standard model (eq 4) to the Covid-19 crisis and the period proceeding it and to the GRC and the period proceeding it; Abbreviated.
Note: The data relied upon to generate this table was TRACE Data provided by FINRA's TRACE System

volume during a crisis is trade in bonds with an HHI of 0.3-0.5 (see 7). Hence, changes in the behavior of these bonds in a crisis can have a sizable impact on the aggregate behavior of the market.

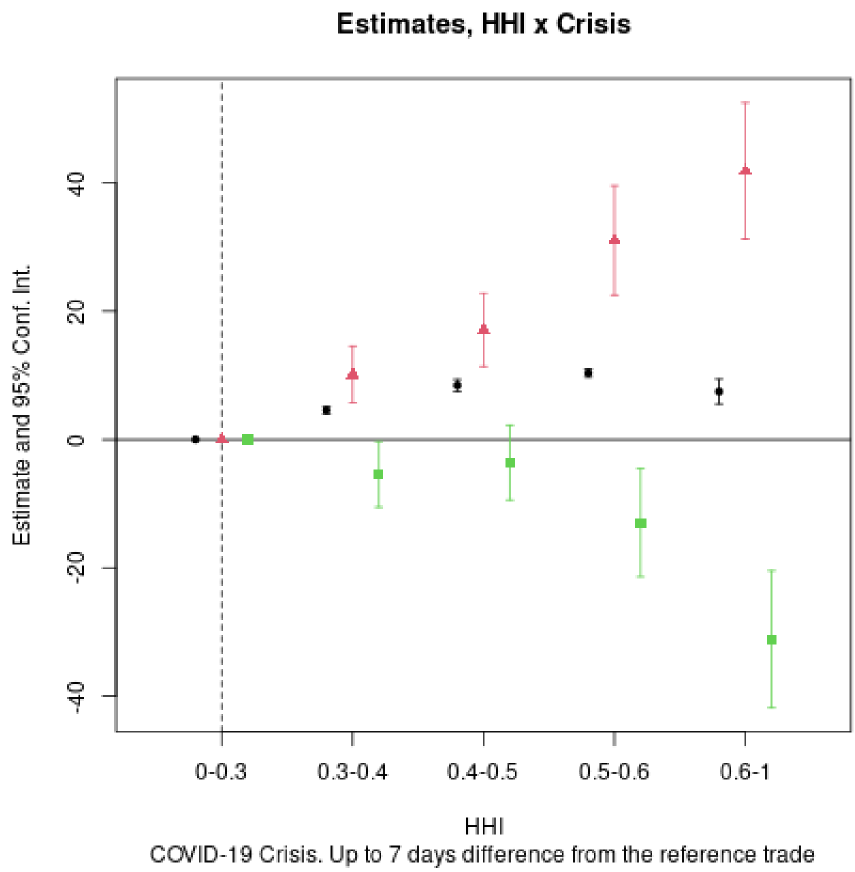

Figure 7: Results from applying the standard model (eq 4) to the Covid-19 crisis and the period proceeding it while treating the bond HHI as a categorical variable.
Note: The data relied upon to generate this table was TRACE Data provided by FINRA's TRACE System

HHI	share of volume
0-0.3	0.11
0.3-0.4	0.59
0.4-0.5	0.19
0.5-0.6	0.06
0.6-0.7	0.02
0.7 - 1	0.02

Table 7: The share of volume bought by dealers in principal capacity by the bond HHI during the Covid-19 crisis.
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

With trades in which dealers sell to customers the pattern is a bit more complex. There, we see a sizable and statistically significant difference only at HHI levels that are greater than 0.5. There are two things to keep in mind about results that persist in this range. First, since there are fewer trades in bonds with this HHI in the sample, the result is more susceptible to error. Second, note that in a market for a bond that with a single market maker dealer may not have a substantial advantage over customers. In such a market, dealers' ability to easily find and trade with any other dealer is less likely to grant them access to multiple potential counter-parties. Thus, the market maker may exploit other dealers' distress to buy cheap from them and sell at a higher price to a customer. That will appear as a negative spread.

A concern that may arise is that the pattern we find in the dealer-to-customer trades reflects changes in the fundamental prices of bonds over time. More specifically, a decline in fundamental prices of bonds will mechanically turn the benchmark price, determined by an earlier trade between two dealers, to a higher price than the current trade price. The change will be more pronounced for less liquid bonds, for which the trades are typically more far apart. As higher concentration is typically correlated with lower liquidity, we will get a bias that operates in the same direction as our coefficients - it will imply that in more concentrated markets, customers sell and buy at a higher discount. That will be manifested in a positive coefficient on the interaction term of HHI and a crisis in customer-to-dealer trades, and a negative coefficient in dealer-to-customer trades.

To address this concern, I run the same regression with HHI bins again, but this time I limit my sample to trades for which the dealer-to-dealer trade that determines the benchmark price to calculate the spreads occured in the same day as the trade itself. The results appear in plot 8. What we can see is that the pattern of higher increase in coefficients for more concentrated markets in trades in which customers sell to dealers persists. The coefficients are about the same sizes as before and are all statistically significant at 1% confidence level. In contrast, the results for trades in which dealers sell to customer are less salient. Besides 0.6-1, neither of the other bins implies a statistically significant decline in spreads across markets with varying levels of competition. In other others, much of the correlation between higher HHI and lower spreads in dealer-to-customer spreads is indeed spurious and driven by changes in fundamental prices. In contrast, such changes do not seem to drive the correlation in trades in which customers sell to dealers.

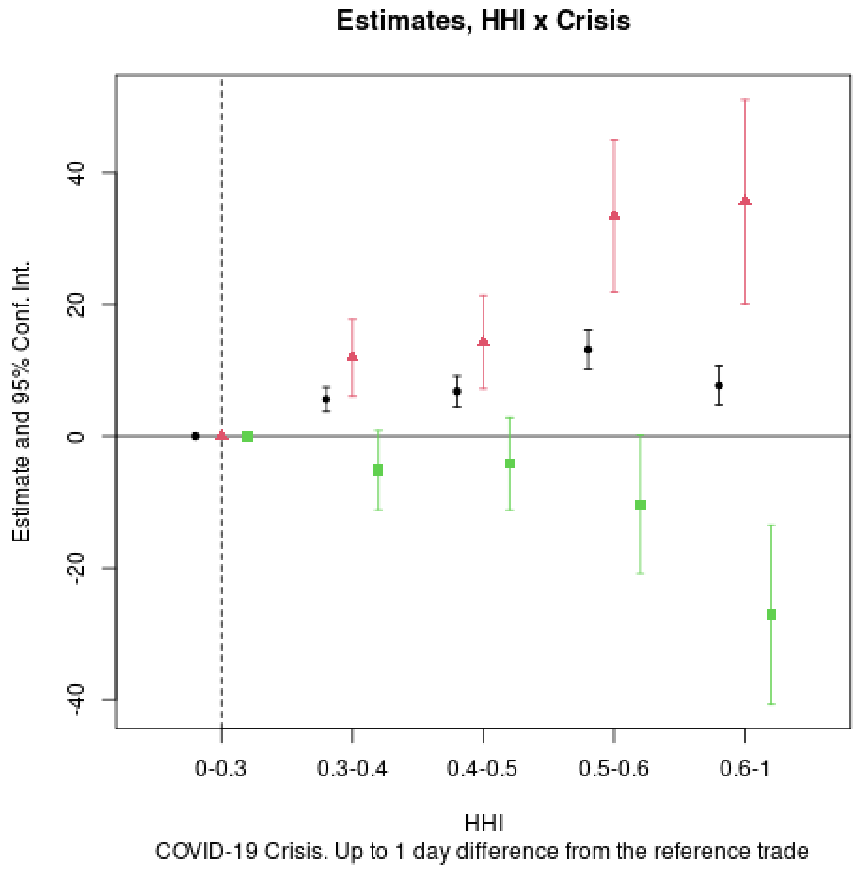

Figure 8: Results from applying the standard model (eq 4) to the Covid-19 crisis and the period proceeding it while treating the bond HHI as a categorical variable. Limiting to cases in which the transaction that generated the reference price occurred at the same day as the trade itself.
Note: The data relied upon to generate this table was TRACE Data provided by FINRA's TRACE System

In plots 9 and 10 we see a similar pattern in the 2007-2009 crisis. Again, we see that the coefficients on the interaction term between HHI and the crisis indicator are all positive, sizeable, and statistically significant. They exhibit an upward trend. There is a break in that pattern as the coefficient on the interaction term on the range between 0.6-1 is smaller than the one in the range of 0.5-0.6. However, even a glance with the naked eye at the confidence intervals of the two we can see that this difference is not statistically significant. These results persist when we limit out attention to cases when the benchmark trade occurred on the same day as the trade itself. In contrast, with trades in which dealers sell to customer the pattern is more complex. The coefficients on the interaction term with the crisis indicator for HHI of 0.3-0.4

and 0.6-1 are close to zero and not statistically significant. Limiting our attention to cases when the benchmark trade occurred on the same day, the interaction term for HHI of 0.5-0.6 losses statistical significant as well and the overall pattern of coefficients does not exhibit a clear monotonic pattern dominating the relationship between HHI and the rise in spreads in a crisis.

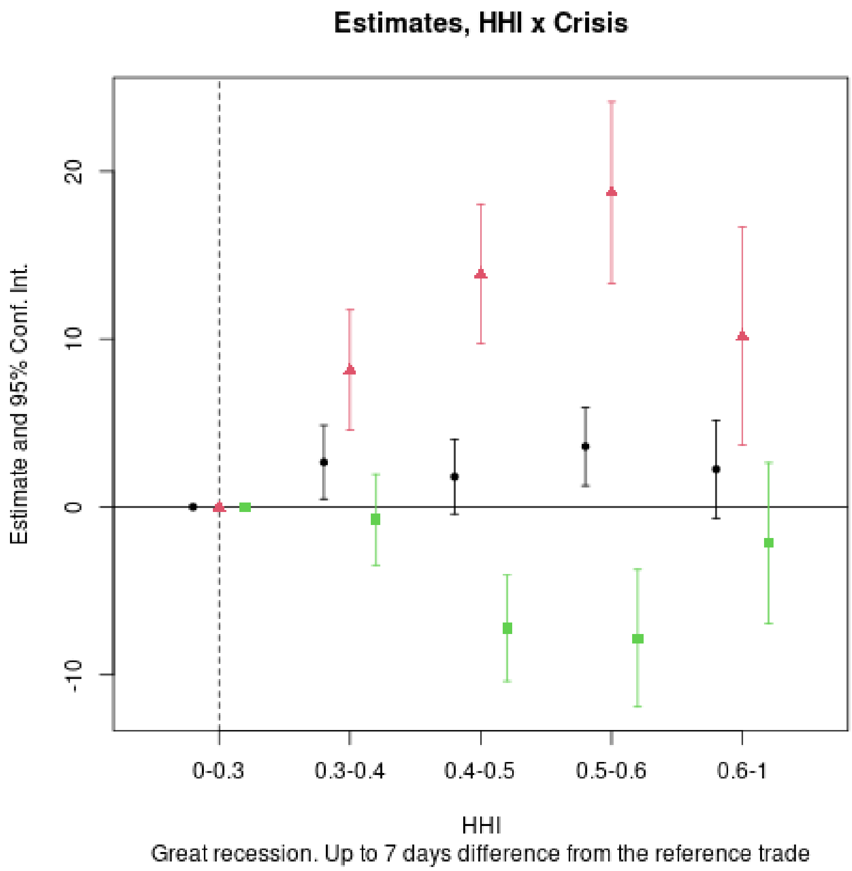

Figure 9: Results from applying the standard model (eq 4) to the 2007-2009 crisis and the period proceeding it while treating the bond HHI as a categorical variable.
Note: The data relied upon to generate this table was TRACE Data provided by FINRA's TRACE System

As a further robustness test, I run an alternative regression model:

$$\begin{aligned} Spread_{i,d,b,t} = \alpha_0 + \rho_b * \eta_d * \nu_i + \beta_2 * HHI_{b,t-1} * \mathbb{I}\{Crisis\} + \\ \eta_d + \text{trade\_date} + \\ \eta_d * \text{trade\_date} + \xi_b * \text{trade\_date} + \nu_i * \text{trade\_date} + \lambda_b * \text{trade\_date} \end{aligned} \quad (5)$$

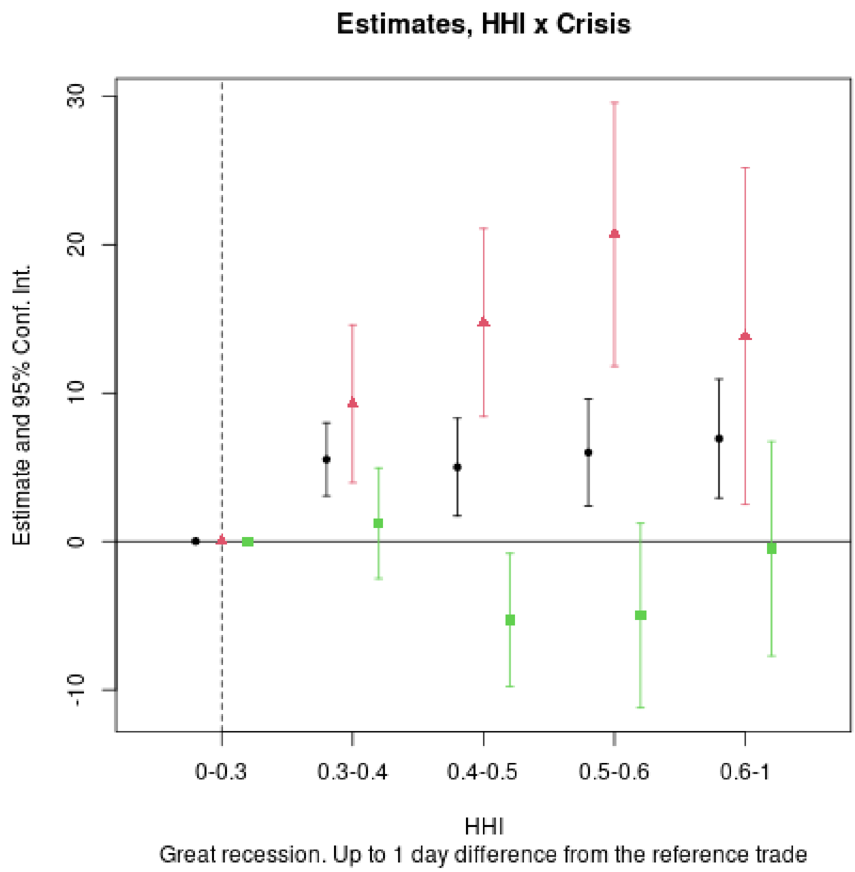

Figure 10: Results from applying the standard model (eq 4) to the 2007-2009 crisis and the period proceeding it while treating the bond HHI as a categorical variable. Limiting to cases in which the transaction that generated the reference price occurred on the same day as the trade itself.
Note: The data relied upon to generate this table was TRACE Data provided by FINRA's TRACE System

In this model, I control for an interaction term of a bond-dealer-volume fixed effect ($\rho_b * \eta_d * \nu_i$). Thus, the model compares trades made by the same dealer, trading the same bond, in the same volume (bin) before and during a crisis. The control for a bond fixed effect makes the use of bond attribute controls redundant (but not the use of interaction terms between those terms and the date). On a deeper level, it prevents potential bias due to the omission of bond attributes. In this regression, the "jump" in the spreads that a dealer will charge consists of an average increase in the spreads due to the crisis time, embedded into the trade date fixed effect, in addition to the term $\beta_2 * HHI_{b,t-1} * \mathbb{I}\{Crisis\}$. The latter captures the same correlation that it embedded in the original model appearing in equation 4: the expected differences in the "jump" for bonds with varying levels of concentration. For the same reasons explained before, I control for the interaction term between dealer, rating, and liquidity fixed effects and the date. The results appear in plot 11.

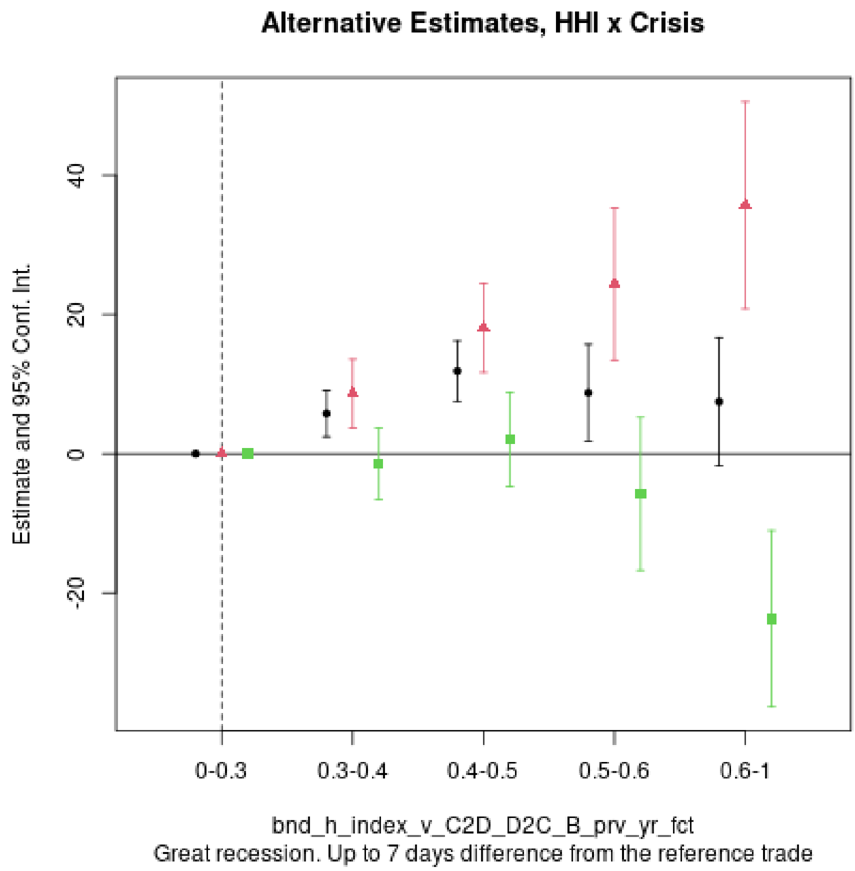

Figure 11: Results from applying the regression model in equation 5 to the COVID-19 crisis and the period proceeding it. HHI is treated as a categorical variable.
Note: The data relied upon to generate this table was TRACE Data provided by FINRA's TRACE System

The result for customer-to-dealer trade remains almost exactly the same as in the original regression model. In contrast, the coefficients on the interaction term between HHI and a crisis for trades in which dealers sell to customers hover around zero for any HHI besides 0.6-1, where they appear negative. That, again, may reflect dealers praying on the distress of other dealers to buy at substantial discounts. Either way, this is evidence of the robustness of the correlation between concentration and the rise in spreads paid by customers that seek to attain liquidity during a crisis. In contrast, the support that the data lends to correlations between the change in spreads in a crisis and the bond HHI for trades in which dealers sell to customers, which are not implied by the logic that guides this paper, seems much more ambiguous. Thus, from this point on, I focus merely on customer-to-dealer trades.

Now, I run a few more robustness tests. First, I run another regression model with an alternative liquidity measure. Rather than using the number of days in which the bond was not traded, I use the median of the difference between the time that a dealer purchases a bond and the next time he gets to sell it. That serves as a proxy of the expected holding period. To avoid bias due to the presence of agency trades, I ignore cases in which the dealer holds the bonds for less than one day. For each trade, I calculate this statistic based on trades of the bond in the previous year to avoid bias due to reverse causality. Further, I change this proxy into a categorical variable by dividing it into bins by percentiles, with each bin covering a range of 5 percentiles.

The main reason to run the regression with this liquidity measure is to address a concern for bias due to the thinness of the inter-dealer market for bonds with high levels of HHI. Concentration implies that a dealer will have greater difficulty selling the bond to another dealer to attain liquidity. That may lead a dealer to require a higher compensation, manifested in greater spread. That is especially true during a crisis when the cost of committing capital to hold a bond for a prolonged period of time in the dealers' portfolio greatly increases. The control for the dealer's expected holding period mitigates this concern. Bonds that do not easily trade in the inter-dealer market will have a longer holding period, as a dealer will be less likely to address transient liquidity needs by trading them. Hence, this liquidity measure will capture costs imposed on the dealer that originate from the breadth of the inter-dealer market.

The results of this regression appear in column (2) of table 8. We see that the inclusion of this liquidity measure diminishes the interaction term of HHI and the crisis indicator from 90 bps to 75 bps. Thus, while this dimension of liquidity may have a role in shaping spreads, it cannot account for the correlation between concentration and the rise in spreads in crisis we find in the data.

Dependent Variable:	spread		
Model:	(1)	(2)	(3)
Variables			
HHI	13.7*** (1.1)	15.3*** (1.5)	
HHI × crisis indicator	92.0*** (9.8)	89.8*** (10.0)	71.0*** (26.0)
rule 144a	-7.1*** (0.60)	-7.7*** (0.90)	-339.5 (1,038.1)
sqrt(age)	-0.11*** (0.01)	-0.11*** (0.02)	0.99** (0.45)
sqrt(time to maturity)	0.20*** (0.01)	0.24*** (0.03)	0.65** (0.26)
sqrt_amtout_issr	$-5.1 \times 10^{-5***}$ (2×10^{-6})	$-4.8 \times 10^{-5***}$ (2.6×10^{-6})	0.001 (0.003)
sqrt(amount outstanding)	-0.0002*** (1×10^{-5})	-0.0002*** (1.5×10^{-5})	-0.0003 (0.01)
coupon rate	2.3*** (0.15)	1.6*** (0.19)	1.1 (0.72)
foreign	-1.3** (0.68)	-2.1** (0.81)	756.7 (2,746.5)
global	-0.52** (0.26)	-0.16 (0.28)	192.7 (752.1)
finance	-1.9*** (0.29)	-1.3*** (0.41)	-1,057.7 (822.0)
utility	-3.0*** (0.57)	-3.4*** (0.64)	-1,266.6 (1,691.3)
Fixed-effects			
rating	Yes	Yes	Yes
trade size	Yes	Yes	Yes
dealer	Yes	Yes	Yes
dealer-date	Yes	Yes	Yes
date	Yes	Yes	Yes
# days traded (prv. year)	Yes		Yes
# days traded (prv. year)-date	Yes		Yes
rating-date	Yes	Yes	Yes
trade size-date	Yes	Yes	Yes
exp. holding time		Yes	
exp. holding time-date		Yes	
HHI			Yes
issuer			Yes
issuer-date			Yes
Fit statistics			
Observations	656,133	874,707	1,270,641
R^2	0.25548	0.23435	0.70151
Within R^2	0.00936	0.00904	0.00024

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 8: Comparing the standard model (eq 4) to two other benchmarks for the COVID-19 crisis: one using expected holding time as a liquidity measure and the other including issuer fixed effects.

Now, we also run our original regression model with an issuer fixed effect. By that, we compare bonds of the same issuer traded in the market with varying concentration levels. That can be caused, for instance, by institutional constraints that prevent certain dealers who are familiar with the issuer from trading some of its bonds. The reason for including the issuer's fixed effects is to control for potential supply shocks. That is, it mitigates the concern that concentration is correlated with features of the bond that make it more appealing to specific types of customers. Thus, the correlation of concentration and spreads may simply represent the fact that those customers were subject to more dire liquidity shocks during the crisis [4].

The results appear in column (3) of Table 8. Again, while a dealer fixed effect diminishes the coefficient on the interaction term from 90 to 71, yet the coefficient remains statistically significant and sizable.

Lastly, in table 6, I gradually add fixed effects to a regression to arrive at the main model appearing in equation 4. What we can see is that the addition (or omission) of fixed effects has very little impact on the regression. Overall, the coefficients remain very stable at a level of about 90 bps. This is further evidence of the robustness of the main result and its independence from the specifics of the regression model that was chosen.

As mentioned, the analysis is not founded on random assignment of treatment, and hence the coefficients cannot be interpreted as proper estimates of the causal impact of concentration on spreads. However, while they cannot provide a precise measure, they may imply an order of magnitude that allows interpreting the relevance of concentration in OTC markets to the rise in the cost of trading (spreads) in a crisis. For that purpose, I calculate the regression prediction for the increase in spreads in times of crisis in a world with an identical composition of bonds but for the fact that all bonds are traded in highly competitive markets with HHI being lower than 0.3. I weigh the importance of the change in spreads for each HHI bin according to the share of volume traded in the crisis accounted for by bonds in that bin. That is, I compute the weighted sum:

[4] I am grateful to Dasol Kim from the Office of Financial Research for raising this concern and offering to add an issuer fixed effect.

$$\sum_{HHI} \beta_{2,HHI} * \text{volume share}_{\text{HHI}} = 0.59 * 10 + 0.19 * 17 + 0.06 * 31 + 0.04 * 42... \approx 12.6$$

As a benchmark, table 9 presents the volume-weighted spread during the COVID-19 crisis and the period preceding it in customer-to-dealers trades. As we can see, the weighted spread increased by about 64 bps. Hence, the regression predicts that in a fully competitive market, the increase in the spread increaseVID-19 crisis would have been 20%

Period	Mean spread
Pre Covid-19 Crisis	18.57
Covid-19 Crisis	82.57

Table 9: Mean spread charged by dealers when buying from customers on principal capacity (trades over $100,000). Weighted by quantity traded.
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

Now, I document the behavior of the volume of trades across different concentration levels. The HHI is divided into categories, and all the volume traded in the bonds included in each bin is summed up. Then, I calculate the percentage change between the total volume of all bonds included in each bin between a time of crisis and a period of similar length that precedes it.

Figure 12 shows the percentage change in volume traded between the Covid-19 crisis and the months that preceded it. One can see that bonds traded by fewer dealers exhibit a more significant relative decline in volume during the Covid-19 crisis. The trend gradually increases with HHI, meaning a more substantial reduction in volume for more concentrated markets. This is important for two reasons. First, this paper's hypothesis implies that dealers will exercise market power more aggressively. They will be more inclined to sacrifice volume to raise the return per transaction. Second, volume is the primary concern of regulators. The problem with dealers' use of market power is not that it generates profits for dealers at the expense of their customers (in fact, some may even see that as a silver lining - improving dealers' access to capital and their ability to keep operating throughout the crisis). Instead, the problem is that when

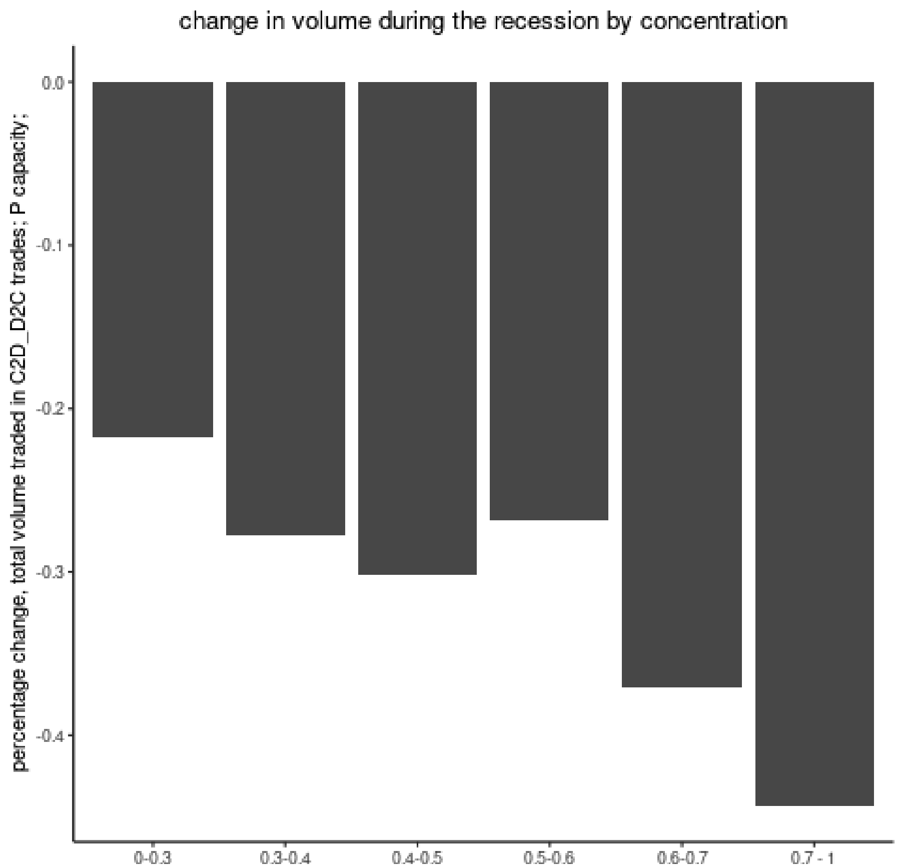

Figure 12: Change in volume bought by dealers in principal capacity by HHI: The Covid-19 crisis and the preceding period.
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

dealers do so, they limit trade and prevent the reallocation of liquidity from customers that are better off to those that are distressed. This can be considered an analogy to the classical monopoly problem, where the concern is the "underproduction" of intermediation services.

The plot implies sizeable differences in the response in volume to distress during a crisis. It is possible to use it, alongside the share of volume traded by HHI level (appearing in Figure 3), to approximate the difference between the decline in volume in a highly competitive market (HHI of 0.3 or less) to what we find in the data. Let Q_{hs} denote the share of trade in bonds with an HHI in the set hs, and let Δ_{hs} denote the decline in volume for bonds in that set. The weighted decline in volume, according to the data, is

$$\sum_{hs} Q_h s{*}\Delta_{hs} = 0.25{*}{-0.22}{+}0.25{*}{-0.28}{+}0.15{*}{-0.32}{+}0.1{*}{-0.28}{+}0.1{*}{-0.37}{+}0.15{*}{-0.47} = -0.3.$$

The change in volume in a highly competitive market is -0.22. That is, the decline in volume

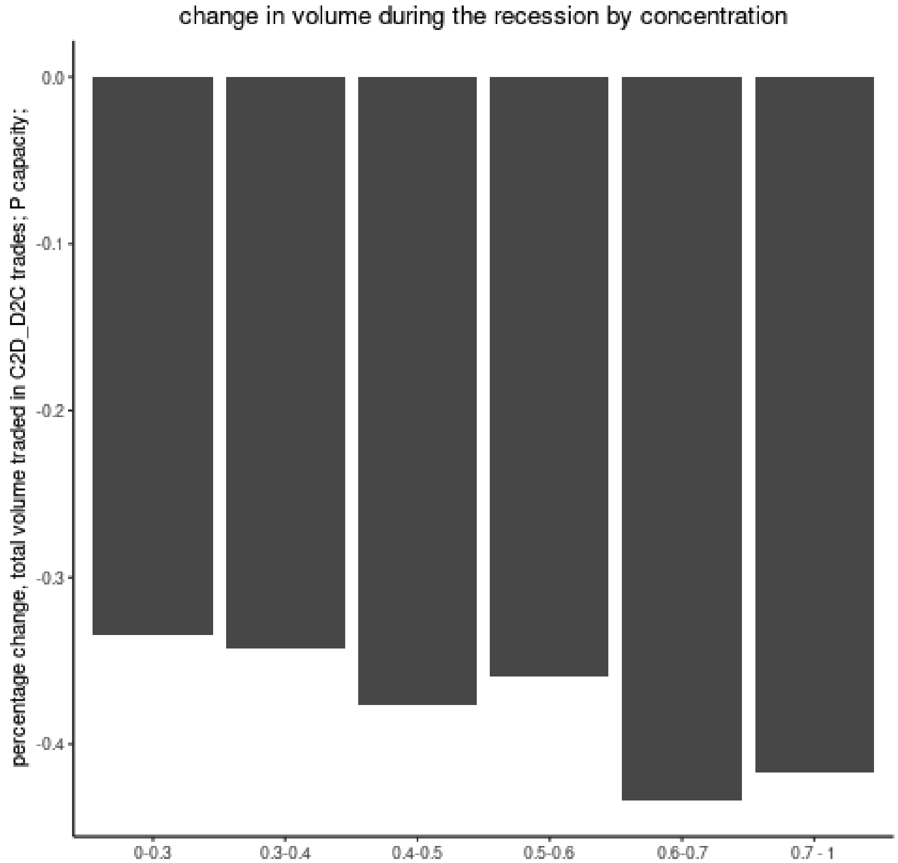

Figure 13: Change in volume bought by dealers in principal capacity by HHI: The 2007-2009 crisis and the preceding period.
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

in the market is 36% greater than what we expect in an economy where all markets have an HHI of 0.3 or less. The calibration exercise assesses which part of this difference originates from changes in the way that dealers exercise their market power.

Figure 13 represent the same phenomenon for the 2008 crisis. Here too, there is a correlation between HHI and the decline in volume. However, it is somewhat more muted compared to the Covid-19 crisis.

3 Model

There is an economy that lasts for two periods, 0 and 1. It is populated by an infinite and discrete number of dealers and an infinite and discrete number of customers. All players are expected utility maximizers and risk neutral. There are infinitely many bonds in the economy.

Each customer holds one bond and may choose to sell it to a dealer. Dealers are assumed to have "deep pockets".

Each bond is of high value v_h with probability q_h, or of low value $v_l < v_h$ with probability $1 - q_h$. A bond value reflects its risk of default. In most cases, the bond's observable attributes, especially its rating, will disclose this risk. In this setting, v_h is the value of a typical bond. In contrast, v_l is the value of a bond that is riskier than what is implied by its observable attributes, and especially rating. Accordingly, q_h is the likelihood that the investor will discover that its rating correctly gauges the risk involved in holding it after acquiring information about the bond.

Customers know the value of the bonds, while only a share of the dealers knows it. This assumption can be thought of as if the market consists of both "insiders", who constantly operate in it, and "outsiders" who can access it but typically do not trade. Customers know the assets that they usually hold and trade. In contrast, they face some dealers who are familiar with that asset and some that are not [5].

Dealers can acquire information. They can perfectly learn the realization of v_j with probability λ by paying a cost of $c(\lambda)$. The cost function $c()$ is increasing, convex, differentiable, and satisfies the Inada condition ($c' > 0; c'' > 0; c'(0) = 0; lim_{\lambda \to 1} c'(\lambda) = \inf$).

In period 0, dealers acquire information about the bond. First, they choose λ. Then, each dealer learns the quality of each bond with probability λ. The probability that n dealers become informed about a quality of a bond will be

$$\lambda \frac{e^{-\lambda}}{n!}$$

.

After the results of the attempts to acquire information are realized, all players in the market

[5] The claim that customers can be equally or more informed than dealers finds support in the literature studying dealers' exposure to adverse selection. Easley and O'hara (1987) is a seminal work that argues that adverse selection manifests in higher volume associated with higher spreads. More recently, Chalamandaris and Vlachogiannakis (2020) documents that institutional investors, accounting for most of trade, typically earn excess returns on their trade in the bonds market. This is especially relevant to this paper, which focuses solely on trades that are larger than \$100,000, and, therefore, more likely to involve institutional investors rather than retail.

learn the number of dealers who became informed about each bond.

At the beginning of period 1, each customer gets hit by a liquidity shock $\delta_i > 0$. Due to the shock, the customer values the asset at $v_j - \delta_i, \quad j \in \{l, h\}$. We assume that:

$$\delta_i = \begin{cases} \delta_r & \text{w/ prob. } (1-\xi) \\ \delta_s & \text{w/ prob. } \xi, \end{cases}$$

where $\delta_s > \delta_r > 0$. The liquidity shock is private information. Dealers do not know it, and customers cannot provide a reliable signal about its value.

Dealers then submit bids to customers for each one of the bonds. All dealers are submitting bids to all customers. Following Lester and Olivier-Weill (forthcoming), it is assumed that, with some probability π, the bid fails to be delivered. π encapsulates the forces that may prevent a customer from receiving from a prominent dealer. A customer may not know of the dealer, or the dealer might not have the liquidity required to serve the customer.

The number of bids a customer gets from informed players has a binomial distribution. It is the number of successes in n trials (with n informed dealers) with a success rate of $(1-\pi)$. At the same time, the customer always gets infinitely many bids from uninformed players. This assumption reflects the notion that one can always sell a security as long as it gives a sufficient discount, that is, a discount that compensates the potential buyer for buying a security he knows nothing about.

After receiving bids, the customer chooses which one to accept, if any. Afterward, payoffs are realized, and the economy comes to an end. The equilibrium notion is:

Definition 3.0.1. *Symmetric equilibrium consists of the following:*

1. *Dealers' choice of the likelihood of becoming informed* λ*;*
2. *The distribution of the number of informed dealers trading each bond, denoted by the* $\Gamma(n)$*;*
3. *The strategy of the uninformed player,* $B^U(n)$ [6]*.*

[6] Note that we assume that the uninformed plays a pure strategy and bids B^U. It is later shown that the assumption is benign.

4. The strategy of the informed, $F_n()$.

That satisfies the conditions:

1. *Given n and $F_n()$, there exists no bid B such that a deviation of a single informed player from $B^U(n)$ to B strictly increases its payoff.*
2. *Given n, $B^U(n)$, and $F_n()$, the bids submitted by the informed maximize its expected payoff.*
3. *The return from buying a bond for $B^U(n)$ is non-negative [7].*
4. *The number of informed dealers has a Poisson distribution with arrival rate λ.*
5. *Dealers' choice of λ maximizes their ex-ante expected payoff.*

The model is solved backward, beginning with finding how trade evolves, and later moving on to studying information acquisition. We assume, WLG, that $v_l < v_h - \delta_s < v_h - \delta_r < v_h$. This section begins by solving for the markets in which $n > 1$. The unique cases in which $n = 0$ and $n = 1$ are studied later. Lastly, this section focuses on symmetric equilibrium, where all informed players choose the same strategy.

3.1 Benchmark model - No Uninformed Dealers

We start by solving for equilibrium in a simplified version of the model. Specifically, we assume that there is no heterogeneity among customers, that is, that $\xi = 0$. Further, we will look at a different environment in which there are no uninformed dealers. Doing so will allow us to focus, for the time being, on the forces that govern the behavior of the informed. Also, it highlights the role of the uninformed dealers by allowing one to compare the outcome of the model with them and without them.

[7] The equilibrium has two optimality conditions for the uninformed: (1) and (3). Adding condition (3) allows us to avoid abnormal behaviors due to our assumption that there is an infinity of uninformed players. Specifically, this assumption implies that each uninformed has a probability of 0 to win and can potentially allow for an equilibrium in which they all make a losing bid. Since the assumption of an infinite number of informed is made simply for mathematical simplicity and reflects the idea that there are many uninformed players, losing bids are unlikely. Adding condition (3) bans them from being a part of an equilibrium.

Let $\bar{R}^i_n$ denote the reservation value of a customer of type $i \in \{r, s\}$ for a high-quality bond. That is, this is the minimum bid a customer who got hit by a liquidity shock of δ_i will accept for it. In the current setting with a single type of customer: $\bar{R}^r = v_h - \delta_r$

Recall that $F_n()$ denotes the CDF of the bids submitted by informed dealers. Following Burdett and Judd (1983), we can show that F is continuous (no atoms), originates at the customer reservation value, $v_h - \delta_r$, and is strictly increasing on the range $[v_h - \delta_r, \bar{B}(n)]$ (connected support) where $v_h > \bar{B}^n > v_h - \delta_r$.

Further, following Weill (2020) and Lester and Olivier-Weill (forthcoming), we can show that the probability that the number of informed dealers that submitted a bid has a binomial distribution with n trials and a success rate of $1 - \pi$. Also, note that the probability that a single dealer bid lower than B is $F(B)$, and hence the probability that all k competitors bid below it is $F(B)^k$.

Using this, we can write the problem of the informed as:

$$\max_{F_n(B)} \int_B (v_h - B) Pr(\text{B is the highest bid}) Pr(B \geq \bar{R}^i) dF(B) =$$

$$\max_{F_n(B)} \int_B (v_h - B) \sum_k^{n-1} \binom{n-1}{k} \pi^{n-1-k} (1-\pi)^k F(B)^k \mathbb{I}\{B \geq v_h - \delta_r\} (v_h - B) \quad (6)$$

By optimality, the informed dealer is indifferent between all bids that are on the support of $F_n()$. We know that the lowest bid on the support is $\bar{R}^i = v_h - \delta_r$, and that this bid wins if and only if all other dealers fail to submit a price. Hence, for any B in the support of $F_n()$:

$$\pi^{n-1} \delta_r = \sum_k^{n-1} \binom{n-1}{k} \pi^{n-1-k} (1-\pi)^k F(B)^k (v_h - B) \quad (7)$$

The last equation, which we will refer to as the informed player optimality condition, implicitly defines the strategy of the informed, $F_n()$. We find that it is fully determined by the following: (i) the value of the security, v_h, (ii) the number of informed players, n, (iii) the likelihood that a dealer fails to submit a bid, π, and (vi) the reservation value of the customer, $\bar{R}^i$.

The bids of the informed are strictly increasing in the number of competitors, n. Or, to be more

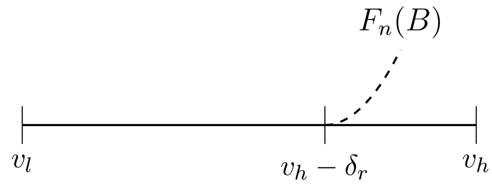

Figure 14: The CDF of bids by informed in a setting with no uninformed and assuming $\xi = 0$.

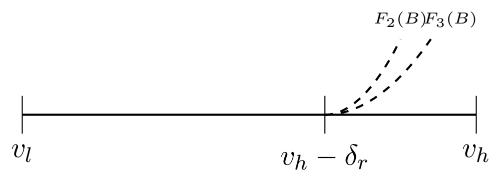

Figure 15: The bids of the informed in a market of two players and three players.

precise, we can state that:

Lemma 3.1. *$F_n(B) < F_{n+1}(B)$ for any B in the support of $F_n()$.*

For the proof, see the appendix. The intuition behind the lemma is simple - the more players are in the market, the greater the decline in the prospects of winning the bond due to submitting a lower bid. Hence, greater competition pushes dealers towards higher bids (or lower spreads).

3.2 Adding Back the Uninformed

Now, we go back to the benchmark model by allowing $\xi > 0$ and letting the uninformed back in. How does that impact the outcome of the model?

Theorem 3.2. *When the security is of high quality, the informed bids are always strictly higher than B^{U^*}.*

The intuition is simple: the uninformed always gets to submit a bid. Hence, a bid below $B^U(n)$ can never win the bond. The informed is better off bidding higher and enjoying some chance of making a return. For the full proof, see the appendix.

That is $B^U(n)$ sets a new lower bound for the bids of the informed that might be accepted by a customer. In other words, now that customers have the option of trading with an uninformed player, their reservation value cannot be below the price that such a dealer will offer.

Further - the lemma already highlights the origin of concentration in our model. Informed dealers will bid more aggressively and win a larger share of the market. In that sense, our model implies that information will be reflected in having a substantial share of the volume of trade.

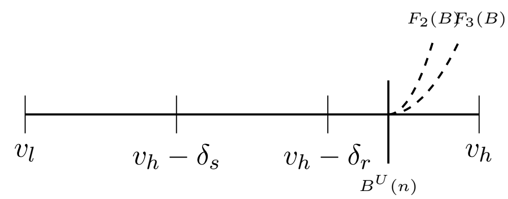

Figure 16: Equilibrium when $B^U(n) > v_h - \delta_r$: two or three informed players.

We proceed to characterize the behavior of the uninformed. By our assumptions, the expected return of the uninformed in case of winning the bond by bidding B^U cannot fall below 0. By optimality, either can it be strictly above zero since in that case, bidding $B^U + \epsilon$ increases the likelihood of bidding the bond from zero to something strictly positive, while still generating a strictly positive return.

Hence, the uninformed always break even, that is:

$$B^U(n) = \hat{q}_h(B,n) * v_h + (1 - \hat{q}_h(B,n))v_l \tag{8}$$

Where $\tilde{B}$ denotes the price at which the security was sold, and $\hat{q}_h(B,n)$ is the likelihood that the bond is of high quality given that is was bought for a bid of B in the presence of n informed players in the market:

$$\hat{q}_h(B,n) = Pr(j = h|\tilde{B} = B, n)$$

We can pin down $\hat{q}_h(B,n)$ using Bayes Rule:

$$\hat{q}_h(B,n) = \frac{Pr(\tilde{B} = B^U(n)|j = h, n)}{Pr(\tilde{B} = B^U(n)|n)} q_h$$

By lemma 3.2, we know that the bid submitted by the uninformed may win the bond when it is of high quality only if all informed players remain out of the market, which happens with probability π^n. In addition, we require that such a bid will be higher than the value that the customer assigns to the bond, $v_h - \delta_i$. Therefore, we can write:

$$Pr(\tilde{B} = B^U(n)|j = h, n) = \pi^n * Pr(B^U(n) \geq v_h - \delta_i)$$

Assume that $B^U(n) \geq v_h - \delta_r$. That is, the uninformed is submitting a bid that appeals to all participants in the market. Thus: $Pr(B^U(n) \geq v_h - \delta_i) = 1$, and: $Pr(\tilde{B} = B^U(n)|j = h, n) = \pi^n$. Plugging these expressions back into our function for $\hat{q}(B, n)$, we find that the probability that a bond bought by the uninformed is of high quality is:

$$\hat{q}_h(B, n) = \frac{\pi^n q_h}{\pi^n q_h + 1 - q_h}$$

Plugging that back into our expression for $B^U(n)$:

$$B^U(n) = \frac{\pi^n q_h}{\pi^n q_h + 1 - q_h} v_h + (1 - \frac{\pi^n q_h}{\pi^n q_h + 1 - q_h}) v_l$$

We check for consistency with our original assumption, $B^U(n) \geq v_h - \delta_r$:

$$B^U(n) = \frac{\pi^n q_h}{\pi^n q_h + 1 - q_h} v_h + (1 - \frac{\pi^n q_h}{\pi^n q_h + 1 - q_h}) v_l \implies \frac{\pi^n q_h}{\pi^n q_h + 1 - q_h} \geq 1 - \frac{\delta_r}{v_h - v_l} \quad (9)$$

That is, the uninformed might submit a bid that appeals to all sellers of the high-quality bond only if:

$$\frac{\pi^n q_h}{\pi^n q_h + 1 - q_h} \geq 1 - \frac{\delta_r}{v_h - v_l} \quad (10)$$

Lemma 3.3. *(bid of the uninformed) The bid submitted by the uninformed is:*

$$B^U(n) = \hat{q}(B^U(n), n) v_h + (1 - \hat{q}(B^U(n), n)) v_l$$

Where if:

$$\frac{\pi^n q_h}{\pi^n q_h + (1-q_h)} > 1 - \frac{\delta}{v_h - v_l} \tag{11}$$

Then:

$$\hat{q}(B^U(n), n) = \frac{q_h \pi^n}{q_h \pi^n + (1-q_h)} \quad \text{, and} \qquad B^U(n) \geq v_h - \delta_r$$

Else, if:

$$\frac{\xi \pi^n q_h}{\xi \pi^n q_h + (1-q_h)} > 1 - \frac{\delta_s}{v_h - v_l} \tag{12}$$

Then:

$$\hat{q}(B^U(n), n) = \frac{\xi q_h \pi^n}{\xi q_h \pi^n + (1-q_h)}, \text{ and} \qquad v_h - \delta_r > B^U(n) \geq v_h - \delta_s$$

If neither of these conditions hold, $\hat{q}(B^U(n), n) = 0$ *and* $B^U(n) = v_l$

For full proof, see the appendix.

Note what the submission of such a high bid by the uninformed depends upon. The likelihood of such a bid is declining in the number of informed dealers, n, and increasing in the likelihood that they cannot bid, π. This is because the presence of informed players makes it less likely that someone selling a high-quality bond will end up trading with the uninformed. In other words, it exacerbates the adverse selection problem it is facing.

Of course, adverse selection is also more significant when the share of high-quality assets, q_h, is small or when the value of a low-quality bond is far smaller than that of a high-quality bond, that is, when the gap $v_h - v_l$ is increasing.

If condition 10 does not hold, one of two things may occur. First, the uninformed may play more prudently by competing only for low-quality bonds. In such a case, $B^U(n) = v_l$. Alternatively, the uninformed might find it optimal to submit a bid in the range $[v_h - \delta, v_h - \delta]$.

The following theorem gives a complete characterization of the behavior of the uninformed:

We noted before that $B^U(n)$ sets a lower bound of the price at which a bond might be sold to an informed dealer. By this, we will enrich our notation slightly by rewriting the customer reservation price (from the informed dealer perspective) as follows:

$$\bar{R}_n^i = \max\{v_h - \delta_i, B^U(n)\}$$

Accordingly, we can restate the problem of the informed dealer by:

$$\max_{F_n(B)} \int_B (v_h - B) \sum_k^{n-1} \binom{n-1}{k} \pi^{n-1-k}(1-\pi)^k F(B)^k \mathbb{I}\{B \geq \bar{R}_n^i\}(v_h - B) \tag{13}$$

Using this new notation, we proceed to state the *shocks' irrelevance theorem*:

Theorem 3.4. *(Shocks irrelevance theorem) If:*

$$\frac{\pi^n q_h}{\pi^n q_h + 1 - q_h} > 1 - \frac{\delta_r}{v_h - v_l} \tag{14}$$

Then:

1. *The outcome of the model is independent of* δ_s, ξ.
2. *All bids are strictly greater than* $v_h - \delta_r$.
3. *All customers always get to trade.*

Proof. When condition 14 holds, the uninformed is submitting a bid greater than $v_h - \delta_r$. Hence, all customers will always prefer to trade with the uninformed over not trading at all. That implies that $\bar{R}^i(n) = B^U(n)$ for both δ_r, δ_s. Using it, we can rewrite the problem of the informed as:

$$\max_{F_n(B)} \int_B (v_h - B) \sum_k^{n-1} \binom{n-1}{k} \pi^{n-1-k}(1-\pi)^k F(B)^k \mathbb{I}\{B \geq B^U(n)\}(v_h - B) \tag{15}$$

The problem is identical regardless of the share of customers hit by a severe liquidity shock, ξ, and the size of that shock, δ_s. Since the informed always bids above the uninformed, its bids will satisfy the following:

$$B \geq B^U(n) > v_h - \delta_r, \quad \forall B \in \text{supp} F_n()$$

Lastly, since the uninformed submit bids that are greater than the customer reservation price regardless of the bond quality and the size of the liquidity shock incurred, all customers get to trade. □

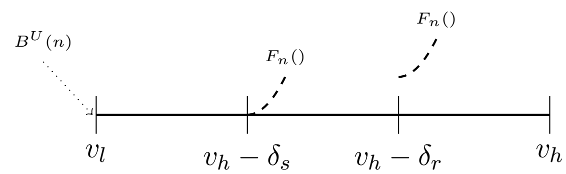

Figure 17: Equilibrium in which the informed submit bids below $v_h - \delta_r$ and the uninformed bids v_l.

In other words, when the adverse selection problem in the market is not very dire, liquidity shocks to the customer sector do not affect it. This is because informed dealers cannot exploit the distress to submit lower bids that solely target customers who got hit by severe shocks. Such bids will be turned down in favor of better offers given by the uninformed. Further, in such circumstances, changes in the number of informed dealers, n, will have a much weaker impact on spreads and volumes. We shall see that in more detail below.

3.3 Response to Liquidity Shocks

From here on, I focus on the case in which condition 14 does not hold unless otherwise stated. In such a case, informed players might choose to submit bids that appeal only to distressed players, that is, bids that are strictly lower than $v_h - \delta_r$. We now turn to characterize the behavior of the informed in such a setting in more detail.

Lemma 3.5. *(The informed bids fall into two separated regions) If* $lim_{\epsilon \to 0} F(v_h - \delta - \epsilon) > 0$*, then, there exists a number,* $v_h - \delta_s < d < v_h - \delta$, $[d, v_h - \delta)$ *such that* $F(x) = F(d) \quad \forall x \in [d, v_h - \delta)$

If the uninformed submits bids below $v_h - \delta_r$, there will be a "hole" in the support of $F()$. The intuition is that bidding $v_h - \delta_r - \epsilon$ (when ϵ is very small) is inferior to simply bidding $v_h - \delta_r$. In the latter case, the returns of buying the security are slightly lower. Still, the likelihood of successfully purchasing it increases substantially due to submitting a bid that appeals to players hit by a liquidity shock of size δ_r. For the full proof, see the appendix.

Next, we find the condition in which the informed submit bids that target only the distressed players:

Lemma 3.6. *(Size of the liquidity shocks impact on the distribution of bids) Assume condition*

14 does not hold. If $\xi > \frac{\delta_r}{(v_h - \bar{R}^s)\pi^{n-1}}$ informed dealers buy only from distressed players in equilibrium. If $\xi < \frac{\delta_r}{(v_h - \bar{R}^s)}$, their bids are always accepted by both distressed and non-distressed. If $\xi \in (\frac{\delta_r}{(v_h - \bar{R}^s)}, \frac{\delta_r}{(v_h - \bar{R}^s)\pi^{n-1}})$ some offers will be accepted by distressed players only while all will accept others.

The theorem states that the more severe the liquidity shock, the more likely the informed are to submit bids that appeal only to distressed players. For full proof, see the appendix.

Now, we compute the mean of the spread charged by the informed. By exploring it, we will get a closer look at the impact of a strategy that targets distressed players on the behavior of the market as a whole.

Theorem 3.7. *([Calculating the average spread) If informed dealers submit bids that are higher than $v_h - \delta_r$ and some that are low, the average spread they charge customers will be:*

$$\hat{S}^i(n) = \frac{n(1-\pi)\pi^{n-1}(v_h - \bar{R}^s)}{V^i(n)}$$

Where $V_i(n)$, the total volume traded by informed players, is given by:

$$V_i(n) = (1-\pi^n)[1 - (1-\xi)\sum_{k=1}^{n} \binom{n}{k} \pi^{n-k}(1-\pi)^k (F(v_h - \delta))^k]$$

If all the bids submitted by the informed exceed $v_h - \delta_r$, then:

$$\hat{S}^i(n) = \frac{n(1-\pi)\pi^{n-1}(v_h - \bar{R}^r)}{1-\pi^n}$$

And if all of these bids fall short of $v_h - \delta_r$, the average spread will be:

$$\hat{S}^i(n) = \frac{n(1-\pi)\pi^{n-1}(v_h - \bar{R}^s)}{1-\pi^n}$$

The full proof appears in the appendix. It uses the equivalence of an informed dealer's expected and realized profit (due to bidding to a continuum of customers).

It is plain to see that the theorem implies that the *share* of the dealer in the surplus generated from trade:

$$\frac{n(1-\pi)\pi^{n-1}}{V_i(n)}$$

It is strictly decreasing in n. Further, the change in the share of the surplus is a major force in shaping spreads. It accounts for the difference in spreads between markets with a different number of informed dealers. As we have shown above, those differences are sizeable. Note that this property of the model separates it from the canonical studies of spreads in OTC using a search-theory framework. In such a framework, the dealer's share is constant (given by the bargaining parameter, typically denoted by θ), and the driving force behind changes in the spreads are changes in the surplus generated from trade.

Using the expressions we got for $\hat{S}^i(n)$, we can explore the predicted impact of concentration on spreads. First, note the case in which some bids by the informed fall below $v_h - \delta_r$. In such cases, competition will likely lead dealers to bid more frequently above $v_h - \delta_r$. Reassigning a greater share of bids from the range below $v_h - \delta_r$ to that above may substantially impact the mean. This is because $F()$ has a "hole" in its support (see lemma 3.5), so there might be a large difference in size between the highest bid that is lower than $v_h - \delta_r$ and $v_h - \delta_r$. This effect manifests in the appearance of $F(v_h - \delta_r)$ in the equation determining the average spread in such cases.

In contrast, when all bids submitted by the informed are above $v_h - \delta_r$, such a channel is not operating. Thus, the impact of increasing the number of competitors is likely to be more muted.

That provides a possible explanation as to why the increase in a bond HHI presages a substantial rise in spreads during distress but not in regular times. Unlike in normal times, during distress, concentration impacts the likelihood that dealers submit bids that appeal only to customers desperate for liquidity. That, in turn, dramatically increases the impact of the concentration average bid.

Further, when we apply similar reasoning to consider theorem 3.4, we can see how *uncertainty* can be a critical factor in augmenting the impact of market power in OTC. Without substantial uncertainty, customers can always trade with uninformed dealers. Informed dealers will shy away from bids that fall below $v_h - \delta_r$, which may significantly attenuate the impact of the number of informed players on the realized bids.

3.4 Endogenous Information Acquisition

Last, we use our model to find dealers' choice of information acquisition. By construction, the uninformed in our model always break even. Hence, the value of acquiring information, that is, the importance of being informed rather than uninformed, is simply the expected profit of the informed player.

In the proof of theorem 3.7, we saw that the expected return of the informed from participating in a market with a total of n informed players is:

$$\mu_s(1-\pi)\pi^{n-1}(v_h - \bar{R}^d)$$

Where $\bar{R}^d$ is the reservation value of the non-distressed player if $B^U(n) \geq v_h - \delta_r$ or if $\xi < \frac{\delta_r}{(v_h - \max\{B^U(n), v_h - \delta_s\})}$. Otherwise, $\bar{R}^d$ is the reservation value of the distressed.

Also, recall that the number of informed dealers has a Poisson distribution with an arrival rate of λ. Thus, the expected value of becoming informed with a probability of λ.

$$\mu_s q_h \sum_{n=1}^{\inf} e^{-\lambda n}(1-\pi)\pi^{n-1}(v_h - \bar{R}^d) - c(\lambda) \tag{16}$$

Taking the FOC, we get the following:

$$\mu_s q_h \sum_{n=1}^{\inf} n e^{-\lambda n}(1-\pi)\pi^{n-1}(v_h - \bar{R}^d) = c'(\lambda) - c(\lambda) \tag{17}$$

The RHS is the marginal cost of increasing the likelihood of successful information acquisition, and the LHS is the marginal return. One immediate implication is that larger markets, that is that markets with a higher μ_s, will see higher λ, and hence greater entry. That is consistent with the finding that bonds with higher amount outstanding are typically traded in markets with lower concentration levels.

4 Calibration

The purpose of the calibration is to examine whether the mechanism described in the paper can account for the magnitude of the differences in spreads and volume change across bonds in markets with varying levels of concentration. In this context, the calibration will also be used to address the question of why concentration presages a dramatic increase in spreads during a crisis but is correlated with a more muted difference in regular times.

I assume that a crisis implies changes in the composition of the assets traded in the market, captured by v_l, q_h, a tightening of dealers' capital constraints, embedded in π, and an increase in the demand for liquidity, manifested in a greater share of distressed customers (higher ξ). I normalize the value of good assets, v_h, to 1 in each period and am left with ten parameters:

$$v_l^g, v_l^b, q_h^g, q_h^b, \pi^b, \pi^g, \xi^g, \xi^b, \delta_r, \delta_s$$

To mitigate selection bias, I calibrate the data to a subsample of transactions in which a dealer buys from a customer bond rated "BBB-" with a par value between $1 M to 5$M. I further limit my attention to trades made by the top 50 dealers and exclude agency trades. Following Kargar et al. (2021), I define the Covid-crisis as the period starting on March 5th, 2020, and ending on April 10th of that year. I regard the "normal" times as the pre-crisis period that starts on Jan. 1st, 2019 and ends with the onset of the Covid-19 crisis.

I begin by calibrating parameters that govern the behavior of the market in normal times, δ_r, ξ^h. For δ_r, I use the measure of Chen et al. (2018), who estimated that sellers' holding costs of bonds rated Ba in normal times are 83 bps, and those rated Baa at 67 bps. Picking the midpoint, where $BBB-$ belongs, I set $\delta_r = 0.75$. Alongside, I assume that $\xi^g = 0$. I will show below that this assumption is benign.

Next, I calibrate parameters that determine asset composition: $q_h^i, v_l^i, i \in g, b$. Recall that we interpret v_h^i as the value of a "typical" bond in a specific market, where a market consists of all bonds that have similar observable attributes. In contrast, v_l^i is the value of bonds that share those observable attributes despite being riskier. I interpret this to imply that if someone had learned its true risk value, they would have assessed it as riskier than it appears.

In other words, such a bond would be downgraded conditional on being audited by a rating agency. I assume that bonds are audited at random, and hence the share of low-quality bonds, $1-q_h$, equals the probability of being downgraded conditional on a re-rating event (downgrade, upgrade, or re-affirmation of its rating). Using the Mergent-FISD rating table, I find that in the pre-crisis period, 13% of re-rating of bonds rated "BBB-" end with a downgrade (that is - $q_h^g = 0.87$). Similarly, during the crisis, the likelihood of being downgraded increases to 18%, so that $q_h^b = 0.82$.

Similarly, I consider v_l as the expected value of a bond conditional on a downgrade. To estimate the value of a bond, I use its price in dealer-to-dealer trades. Given that dealers operate in a relatively frictionless market, the inter-dealer price of a bond should closely reflect its true value. I concentrate on small trades (less than $10,000) to avoid biases stemming from differences in holding costs across trades with varying volumes.

In the pre-crisis period, 54% of cases in which a BBB-" bond is downgraded result in an assigned rating of BB+. In the remaining 46%, the assigned rating is BB. Using these probabilities as weights, the expected value is $\nu^{pre}_{dngrd,BBB-} = 99.8$. Following a similar procedure, I find that in the crisis period $\nu^{crisis}_{dngrd,BBB-} = 0.89$. Remember that I normalized $v_h = 1$. Consequently, v_l is defined based on the relative value of an inferior bond when compared to a standard one. Denote the dealer-to-dealer price for BBB-" bonds in period P as ν^{P}_{BBB-} and note that:

$$\frac{v_l^g}{v_h^g} = v_l^g = \frac{\nu^{pre}_{dngrd,BBB-}}{\nu^{pre}_{BBB-}} = 0.997, \qquad v_l^c = \frac{\nu^{crisis}_{dngrd,BBB-}}{\nu^{crisis}_{BBB-}} = 0.89$$

Note the wide gap separating the expected fall in prices due to a downgrade during the crisis when compared to the pre-crisis period. This is not special to bond-rated "BBB-", but rather reflects a general widening between the price of bonds of different ratings in dealer-to-dealer deals during the crisis. I regard this to indicate an increase *cost of bearing idiosyncratic risk* due to the expected deterioration in market performance. As we shall immediately see, this increase in the cost of risk plays a substantial role in facilitating the impact of concentration on market performance in times of crisis.

Now, I turn to estimate the remaining four parameters: $\delta_s, \xi^c, \pi^c, \pi^g$. I determine these values by targeting the behavior of spreads in markets with varying degrees of competition. Consequently, I classify bonds into three bins: those with an HHI of 0.3-0.4, with an HHI of 0.4-0.7, and with an HHI of 0.7-1. These groups represent markets with three informed dealers, two informed dealers, and a single informed dealer, respectively. For each bin, I calculate the mean spread during the pre-crisis and crisis periods. To minimize measurement error, I use a multi-step process to estimate the mean for each period and concentration bin pair, including calculating the mean for each bond-era pair, winsorizing, and computing the mean for all bonds within each bin [8].

The spreads in the data appear in Table 10. As one can see, there is a clear trend of rising spreads alongside HHI during the COVID-19 crisis. The differences between more and less competitive markets appear substantial, as a transition from a three-informed players market to a one-informed player market increases the average spread by about 70%. In contrast, in regular times, we witness an increase in spreads when transitioning from 3 informed players to 2 informed players market, but then see almost the exact same spreads for $n = 1$ and $n = 2$.

HHI	Spreads (pre-crisis)	Spreads (crisis)	Volume Change
> 0.6	24.41	169.3	-0.71
0.4-0.6	24.25	113.98	-0.33
0.3-0.4	15.37	98.80	-0.15

Table 10: Weighted mean - spreads by HHI, COVID-19 crisis and the period proceeding it

I use the six average spread moments to calibrate the four remaining parameters. Specifically, I choose parameters to minimize the percentage change deviation between the spreads in the data versus those in the model, that is:

[8] I am gauging spreads using the O'Hara-Zhou method used in section 2.3. That is, for each trade, I find the percentage deviation of the price that a dealer paid for it when buying it from a customer vs. the price paid for it in the most recent dealer-to-dealer trade. To avoid bias, I limit my attention only to trades in which the trade used to calculate the reference price occurred more than an hour but less than two weeks than the time when the transaction took place. The purpose of ignoring spreads calculated using a dealer-to-dealer trade that occurred less than an hour before is to exclude agency trades that might have been wrongly categorized as principal. Alongside this, to limit the impact of outliers on my final results, I calculate the mean spread in a few steps. First, I calculate for each bond-era pair a weighted mean of the spreads charged from customers who sold it. As a weight, I used the inverse of the time that elapsed between the dealer-to-dealer trade used to gauge the reference price and the current transaction. Then, I winsowrised the bond-level weighted mean at the 10% and the 90% percentile level, with the purpose of diminishing the impact of outliers on my final result. Lastly, I take a weighted mean over all bonds in each period and concentration category, using the number of trades in the bonds as a weight.

$$\sum_{i\in\{g,c\}}\sum_{n\in\{1,2,3\}}\left(\frac{\hat{\bar{S}}_n^i-\bar{S}_n^i}{\bar{S}_n^i}\right)^2 \tag{18}$$

Where $\bar{S}_n^i$ is the mean spread in state i in a market for a bond with n informed dealers in the data and $\hat{\bar{S}}_n^i$ is its model equivalent. Also, I restrict the parameter space by requiring that it will be harder to get a bid from a dealer in times of distress, that I restrict the solution space to the cases where $\pi^c > \pi^g$.

Note that this is not a convex problem, as multiple factors interact to determine the realized spreads. Their impact is often non-monotonic. For instance, an increase in π pushes spreads up by increasing informed players' monopolistic power, but may also decrease them by emboldening the uninformed to bid more aggressively and pose competition. To address the non-convexity, I search for the missing parameters using a particle swarm algorithm.

In each iteration of the algorithm, I guess the four parameters and add them to the other parameters of the model that were pinned down in earlier stages of the calibration. Then, I plug the parameters into the model to derive the average spread for each state, $i \in \{r, c\}$, and for markets with $n \in \{1, 2, 3\}$ dominant dealers. Specifically, I make use of the informed dealers' optimality condition (eq. 7) to pin down the probability of submitting a bid that appeals only to the distressed players ($F(v_h - \bar{R}^r)$), and plug in the result I attain into the realized spread equation derived in Theorem 3.7. For further details, see the appendix. Lastly, I calculate the 18.

4.1 Calibration Results

The full list of model parameters appear in Table 11. In calibrating the model to replicate the behavior of spreads, I find that $\xi^c = 0.55$, that is, in a crisis, 55% of the sellers are distressed. Further, $\delta_s = 0.00165$, that is, those sellers are willing to sell the asset at a discount of 165 due to pressing liquidity needs. Note that this is about twice the decline in value that prompts regular sellers to sell, of 83 bps ($\delta_r = 0.83$). Alongside, the likelihood that a dealer fails to bid, π, is 0.35 in normal times, and 0.49 in a crisis. Hence, systemic distress raises this likelihood by 42%. Recalling that the likelihood embeds both higher holding costs and search frictions and that the latter are less affected by the crisis, these seem like reasonable estimates.

Variable	**Value**	**Interpretation**	**Corresponding Moment/Source**
q_h^g	0.87	Pr. that a bond is of high quality, normal times	Pr.(downgraded | rating updated) pre-crisis.
q_h^c	0.82	As above, crisis	As above, crisis
v_l^g	0.99	Value of a low quality bond, normal times	Expected change in the D2D price following a downgrade (pre-crisis)
v_l^c	0.89	As above, crisis	As above, (crisis)
δ_r	0.0083	Liquidity shock incurred by non-distressed customers	Chen et al. (2018)
δ_s	0.0165	Liquidity shock incurred by distressed customers	crisis spreads.
ξ^g	0	Share of distressed customers in normal times	(Unidentified).
ξ^c	0.55	As above, crisis	Change in volume
π^g	0.35	Prob. that the dealer does not submit a bid in normal times	Pre-crisis spreads.
π^c	0.49	As above, crisis	crisis spreads.

Table 11: Calibrated model parameters.
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

In figure 18 I plot the spreads in normal times in the data (pre-crisis era) and in the model.

As we can see, the model does very well in replicating the data, both in terms of magnitude and shape. Specifically, both produce a substantial increase in spreads in the transition from a three-informed dealers market to a two-informed dealer market and more or less identical

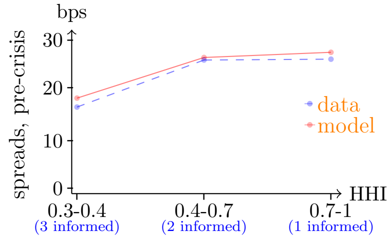

Figure 18: Spreads in the pre-crisis period: model vs. data,
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

spreads in markets with two-informed dealers versus those with a single-informed dealer.

Irrespective of the variables we chose to fit the behavior of spreads, $\delta_s, \xi^c, \pi^g, \pi^c$, we find that v_l is higher than the value assigned to a high-quality bond by a non-distressed seller, $v_h - \delta_r$. Recall that uninformed dealers will always bid weakly above v_l, and hence will submit a bid that is sufficiently high to compensate any customer. That implies that the reservation value of sellers in this setting is not determined by their valuation of the bond, but rather by their outside option of selling it to an uninformed dealer. As a result, the distribution of sellers' liquidity shocks, embedded in δ_s, δ_r, and ξ^c has no impact on spreads in the calibrated model at regular times. In other words, the *shocks irrelevance condition*, presented in Theorem 3.4, holds. This fact also means that the assumption that I made above, that $\xi^c = 0$, is benign - it has no implications for the behavior of the model.

Critically, among the four variables that we calibrated to fit the model to spreads in the data, $\delta_s, \xi^c, \pi^g, \pi^c$, the only one that has any relevance for its behavior in regular times is π^g - the probability that a dealer fails to bid. π^g shapes the spreads through two channels operating in opposite directions. First, alongside the asset composition, embedded in q_h^g, v_l^g, it pins down the bid of uninformed, $B^U(n)$. The higher is π^g, the weaker is the adverse selection, and the higher is $B^U(n)$. That, in turn, improves sellers' outside options and diminishes spreads. At the same time, an increase in π^g weakens competition between informed dealers, and by that, increases spreads.

What happens in the calibrated model is that the transition from two dealers to one lower the exposure of uninformed players to adverse selection and leads them to bid more aggressively. As a result, they pose greater competition to the informed player. That offsets the increase in spreads due to weakened competition among the informed dealers. In the transition from

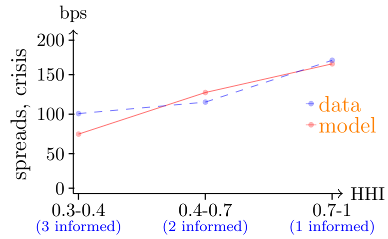

Figure 19: Spreads in a crisis period: data vs. model

Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

two informed dealers to one, the size of each force is about the same, so we witness a negligible change in spreads. In this context, note that the variable π^g is calibrated to fit three different moments - spreads in normal times for markets with one, two, and three informed players. Given that this is the case, the good fit that the model attains cannot be taken for granted.

In figure 19, we see the behavior of spreads in crisis in the model and in the data.

In this case, the shocks irrelevance condition does not hold due to the sizeable differences in values between high and low-quality bonds embedded in $v_l^c = 0.89$. We attain an excellent fit for the data. That should not count for much of an achievement, as we match 3 variables, δ_s, ξ^c, π^c, to three moments characterizing spreads in that state. Specifically, we find that with $n = 1$, the informed dealers use their monopoly power to set the spreads to equal the reservation value of the distressed seller, $v_h - \delta_s$. The other two variables aim to replicate the spreads for markets with 2 or 3 informed players.

The calibration of the model of the behavior of spreads implies an explanation as to why spreads increase so dramatically in concentration in times of distress but change very mildly alongside HHI in the pre-crisis period. The reason lies in the substantial growth in the gap between the value (measured by inter-dealer price) between regular bonds and lemons, embedded in $v_l^c = 0.89$ and $v_l^g = 0.997$. In normal times, absent a large difference between v_h and v_l, the risk in purchasing an unfamiliar bond is not substantial enough to drive out uninformed dealers from submitting *relevant* bids, that is, bids that are high enough to answer the liquidity needs of all customers. Hence, customers always have the outside option of selling their assets at a reasonable price to *some* dealer. This outside option imposes limits on informed dealers' ability to exploit their market power and charge high spreads. Specifically, it prevents them from gaining any business by submitting low bids that only appeal to the distressed.

In other words, in the pre-crisis period, the adverse selection problem of the uninformed is much less dire. Hence, the uninformed can (and does) submit bids that are high enough to appeal to customers who hold a high-quality bond. The customer's outside option of attaining a reasonable deal when trading with the uninformed limits informed dealers' ability to exercise market power and bid low.

Now, I turn to compare the implications of the model for the behavior of volume change in response to the crisis to what we find in the data itself. Note that volume change, in contrast to spreads, was not a target in the calibration.

To allow for this comparison, I first need to create a volume change measure that applies to analogous objects in the model and in the data. The issue here is that while the model assumes that the measure of sellers, μ_s, is the same in the pre-crisis and the crisis period, in the data that need not (and probably is not) the case. Thus, to keep the two comparable, I define the data volume change in trades of bonds in markets with n informed players by:

$$\delta V(n) = V_{r,n} \alpha V_{c,n}$$

Where $V_{r,n}$ denotes the volume traded in such bonds in the pre-crisis period, $V_{c,n}$ the volume traded in those bonds in the crisis period, and $\alpha > 0$ being a multiplier used for adjustment.

One natural candidate to α would be the ratio between the length of the pre-crisis period and the crisis period. While this should improve the measure, However, it ignores the fact that during the crisis, the daily demand for liquidity is probably higher. Another alternative would be to use α so that the aggregate volume in both periods would be about the same. The problem with doing so is that aggregate volume is an endogenous variable and in fact, the variable we want to study. Thus, I set $\alpha > 0$ to be a positive number so that highly competitive markets, that is, those in which HHI < 0.3, experience no volume. This is consistent with the model, in which as n increases, one converges towards competitive pricing in which all demand for liquidity gets answered. I find that $\alpha = 7.5$.

In figure 20, I plot the $\delta V(n)$ for the data and for the calibrated model. As we can see, the model replicates the data with a very high level of precision. That occurs in spite of the fact that volume was not a target for the calibration.

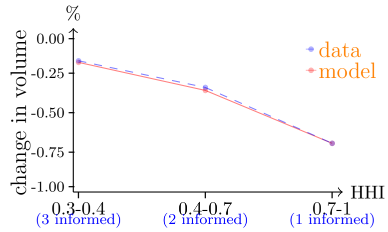

Figure 20: Volume response to the crisis: data vs. model
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

A conservative interpretation of this result is that the differences in the response of volume to systemic distress across markets with varying levels of competition resemble what we would have expected to find if the forces that underlie them are indeed those that govern the model: (i) limited capacity of the informed dealers, and (ii) a strategical choice of dealers to submit bids that target only distressed sellers.

Granted, the co-movement of volume and spreads across markets with varying levels of competition is already baked into the model from the get-go. Yet, while the model structure might determine the shape, it cannot determine exact values. Hence, the success in replicating volume not only in its shape but also in absolute sizes is not, ex-ante guaranteed. Thus, it provides some evidence in favor of the theory that the model embeds.

Furthermore, note the sizeable differences in the response of volume across markets with varying levels of concentration. If the decline in volume indeed originates from concentration, as suggested here, even low concentration levels as HHI $\in [0.3, 0.4]$ lead to a 15% decline in trade volume, while higher levels are associated with an even greater decline. Recalling that 75% of the bonds have an HHI that is greater than 0.3, this can be taken to indicate that concentration has a sizeable contribution to the decline in trade volume ('market freeze') in times of systemic distress.

4.2 Counterfactuals

4.2.1 No Change in Risk, v_l, q_h

To have a better understanding of the critical importance of risk to our mechanism, I run another simulation of the model in which systemic distress is not characterized by a change in the riskiness of assets or in the cost of bearing risk. I do so by setting v_l^c and q_h^c to equal v_l^g, q_h^g correspondingly. Besides that, I am using the calibrated values appearing in Table 11. In this setting, systemic distress means two things: (i) a decline in dealers' capacity to absorb inventories, embedded in the rise in π between the pre-crisis and the crisis period, and (ii) stronger demand for liquidity by customers, reflected in an increase in the share of distressed sellers, ξ. As we shall see, absent a rise in risk these forces do not imply any substantial correlation between concentration and liquidity during distress.

In figure 21 I plot the predicted behavior of spreads during a crisis when assuming no changes in v_l, q_h in response to a crisis. As a benchmark, I plot next to it the prediction from the original model used above, which is the one that incorporates all the calibrated values. In figure 22, I do the same with volume change.

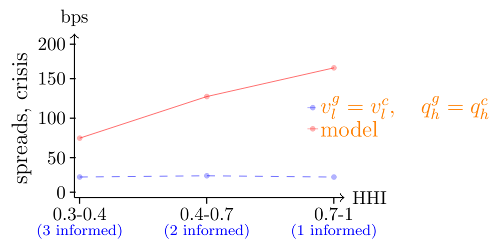

Figure 21: Spreads in the crisis period: model with and without changes in assets composition
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

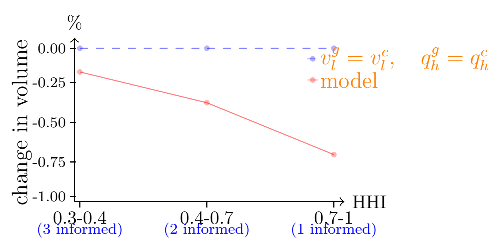

Figure 22: Volume response to the crisis: model with and without changes in asset composition
Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

In the figure, we can see that absent an increase in uncertainty, the simulation implies results that are both quantitatively and qualitatively different. Volume is not affected at all since any customer that is not serviced by the informed can (and does) end up trading with someone that is uninformed. Spreads in crisis times are actually somewhat lower compared to the pre-crisis period. The reason is that the capital shortage among informed dealers increases the likelihood that a customer who trades with the uninformed is holding a high-quality bond. As a result, the uninformed bid more aggressively. That results in shifting the lower bound of the bid of the informed to a higher point. This mechanism can also explain the surprising fact that in this setting, spreads are declining in the transition from 2 informed players to 1 (!).

These results highlight the strong connection between an increase in uncertainty and the impact of market power on the performance of OTC markets. In times of heightened uncertainty, market power has a much more substantial impact. This happens because, in such times, informed dealers are not constrained by the threat of business stealing by the uninformed. We find a very substantial rise in the cost of risk during the Covid-19 crisis. When we examine this rise through the lens of our model, we find that it was an essential condition for the co-movements of market power with spreads and volume that we find in the data.

4.2.2 No Change in Dealers' Capacity, π

Now, let us study what would have happened if there was no change in the likelihood that a dealer successfully submits a bid. That is, consider a case in which $\pi^c = \pi^c$. Hence, a crisis is merely a change in the demand for liquidity, embedded in ξ, and in the level of adverse selection, implied by $v_l^g, v_l^c, q_h^g, q_h^c$.

Note that in our model, π captures the impact of dealers' capital constraints on spreads and volume. As we explained above, dealers' holding costs do not impact the gap between the inter-dealer price and the customer-to-dealer price directly. However, it may affect it indirectly by making it harder for the customer to find a dealer who will have the needed liquidity to buy the security. In our model, this is captured by π. Note that the assumption that capital constraints are a property that our model shares with canonical search-theory models in the spirit of Duffie et al. (2005). In the few models that embed holding cost into a search framework (for instance, Cohen et al. (2022)), an increase in holding costs increases spreads by making it harder to find a dealer that can answer the customer's demand.

In Figures 23 and 24 I plot the behavior of spreads in a crisis and of volume change for a simulation in which I fix π alongside the same variables in the benchmark model. We see that the interaction of capital constraints and market power has a substantial impact on the behavior of spreads during a crisis. Absent a change in dealers' capital constraint, the increase in spreads due to market power is lower but still sizeable at about 50% of what it would have been if the capital constraints were in place. Similarly, the decline in volume across bonds with varying levels of concentration is milder but yet substantial. For bonds with an HHI between 0.3-0.4, the decline without tightening capital constraints is 11%, rather than 15% in their presence, for bonds with an HHI between 0.4-0.6 it is 23% rather than 35%, and for bonds with an HHI greater than 0.6 it is 48% rather than 71%.

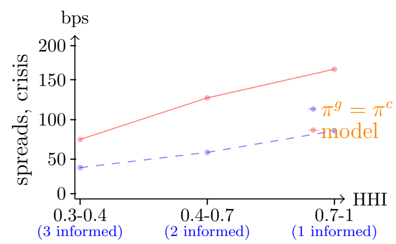

Figure 23: Spreads in the crisis period: model with and without tightening capital constraints

Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

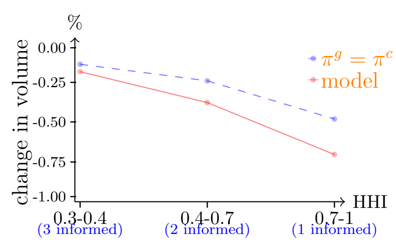

Figure 24: Volume response to the crisis: model with and without tightening capital constraints.

Note: The data relied upon to generate this figure was TRACE Data provided by FINRA's TRACE System.

In other words, without tightening dealers' capital limitations, the impact of concentration on liquidity in dealer markets is lessened. This occurs because such constraints can incapacitate some dealers, reducing competition for the remaining ones. The more moderate shift in volume and spreads during distress in a simulation that sets $\pi^c = \pi^g$ highlights the significance of capital constraints in creating monopolistic inefficiencies in OTC markets during crises. Simultaneously, we see that even without increased capital restrictions, concentration significantly exacerbates the decline in liquidity during systemic distress. This is due to other factors that enhance dealers' competitive edge during a crisis, particularly the heightened uncertainty and the urgent demand for liquidity among their clients. The takeaway is that alleviating dealers' capital constraints is insufficient to restore liquidity.

5 Conclusion

The main takeaway from the paper is that concentration in OTC markets may very well be a concern for financial stability. The dramatic decline in the performance of these markets in a crisis, embedded in higher spreads and lower volumes, seems to be driven, to some extent, by dealers preying on the dire need for liquidity to charge higher markups for their services. In that sense, events of "market freeze" should be understood in terms of *monopolistic inefficiency*

Another significant insight is the highly segmented nature of intermediation in the bond market. It underscores that dealers' capital does not fluidly circulate from one bond to another. The restricted flow of liquid funds among dealers could significantly disrupt the market, while increased trading activities by some dealers might not adequately counterbalance the reduced activities of others. This emphasizes the importance of heterogeneity in dealers' performance, which might be more impactful to the market than previously assumed.

Lastly, this paper brings to light an unanticipated process by which adverse selection amplifies market power and heightens losses from monopolistic inefficiency. The vital role of adverse selection contrasts starkly with the dealers' capacity constraints, which are deemed non-essential for trade to become "clogged" within the dealer sector. Consequently, an unexpected implication arises that merely alleviating the tightening of dealers' capacity constraints might prove insufficient to restore market liquidity.

Dependent Variable:	spread					
		COVID-19 Crisis			GRC	
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
HHI	5.7*** (1.4)	13.7*** (1.1)	0.64 (0.92)	-0.01 (5.9)	9.3** (4.1)	-0.21 (3.2)
HHI × crisis indicator	27.6*** (8.4)	92.0*** (9.8)	-43.2*** (9.6)	2.0 (8.4)	16.0*** (5.7)	-8.4* (4.3)
rule 144a	-2.2*** (0.78)	-7.1*** (0.60)	3.1*** (0.53)	-2.0 (2.1)	-30.0*** (3.1)	18.0*** (2.9)
sqrt(age)	0.08*** (0.02)	-0.11*** (0.01)	0.25*** (0.01)	0.19*** (0.04)	-0.36*** (0.04)	0.55*** (0.03)
sqrt(time to maturity)	0.28*** (0.01)	0.20*** (0.01)	0.35*** (0.01)	0.47*** (0.06)	0.48*** (0.03)	0.49*** (0.02)
sqrt_amtout_issr	$-1.9 \times 10^{-5***}$ (3.3×10^{-6})	$-5.1 \times 10^{-5***}$ (2×10^{-6})	$1 \times 10^{-5***}$ (1.6×10^{-6})	-4.8×10^{-6} (5×10^{-6})	$-1.9 \times 10^{-5***}$ (4.7×10^{-6})	6.1×10^{-7} (3.1×10^{-6})
sqrt(amount outstanding)	-0.0001*** (1.1×10^{-5})	-0.0002*** (1×10^{-5})	-6.8×10^{-7} (8×10^{-6})	-0.0002** (7.3×10^{-5})	-0.0004*** (5.9×10^{-5})	1.6×10^{-5} (4.6×10^{-5})
coupon rate	0.33 (0.26)	2.3*** (0.15)	-1.6*** (0.13)	-0.14* (0.07)	-0.04 (0.06)	-0.23*** (0.05)
foreign	-0.07 (0.64)	-1.3** (0.68)	1.2** (0.58)	-6.3** (2.5)	-7.8*** (2.1)	-1.7 (1.7)
global	0.41 (0.38)	-0.52** (0.26)	1.2*** (0.21)	-2.2*** (0.74)	-10.2*** (0.87)	5.0*** (0.63)
finance	0.28 (0.37)	-1.9*** (0.29)	2.0*** (0.27)	14.1*** (2.2)	9.0*** (1.0)	12.4*** (0.77)
utility	-0.71 (1.2)	-3.0*** (0.57)	1.6*** (0.49)	-5.9*** (2.0)	-13.1*** (1.7)	1.3 (1.5)
Fixed-effects						
rating	Yes	Yes	Yes	Yes	Yes	Yes
trade size	Yes	Yes	Yes	Yes	Yes	Yes
dealer	Yes	Yes	Yes	Yes	Yes	Yes
dealer-date	Yes	Yes	Yes	Yes	Yes	Yes
date	Yes	Yes	Yes	Yes	Yes	Yes
# days traded (prv. year)	Yes	Yes	Yes	Yes	Yes	Yes
# days traded (prv. year)-date	Yes	Yes	Yes	Yes	Yes	Yes
rating-date	Yes	Yes	Yes	Yes	Yes	Yes
trade size-date	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	1,375,793	656,133	719,660	621,479	271,191	350,288
R^2	0.11938	0.25548	0.21531	0.19310	0.36171	0.36307
Within R^2	0.00693	0.00936	0.01212	0.00485	0.00735	0.00997

Clustered (rating) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 12: Results from applying the standard model (eq 4) to the Covid-19 crisis and the period proceeding it and to the GRC and the period proceeding it; Full.
Note: The data relied upon to generate this table was TRACE Data provided by FINRA's TRACE System

6 Empirical Results Appendix:

Dependent Variable:					spread					
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Variables										
HHI	10.1*** (1.1)	16.4*** (1.7)	18.1*** (1.7)	18.4*** (1.5)	17.6*** (1.6)	18.4*** (1.5)	18.5*** (1.5)	18.3*** (1.5)	16.8*** (1.5)	16.3*** (1.5)
HHI × crisis indicator	159.5*** (0.83)	108.6*** (11.6)	108.0*** (11.6)	87.9*** (11.1)	85.9*** (11.5)	89.0*** (10.9)	89.0*** (10.9)	90.9*** (10.6)	91.7*** (10.5)	88.4*** (10.1)
(Intercept)	13.8*** (0.64)									
rule 144a	-9.1*** (0.34)	-9.8*** (1.1)	-9.2*** (1.2)	-9.5*** (1.2)	-7.9*** (1.1)	-7.9*** (1.0)	-7.8*** (1.0)	-7.7*** (1.0)	-7.0*** (0.98)	-7.0*** (0.92)
sqrt(age)	-0.14*** (0.007)	-0.15*** (0.02)	-0.14*** (0.02)	-0.15*** (0.02)	-0.08*** (0.02)	-0.06*** (0.02)	-0.07*** (0.02)	-0.07*** (0.02)	-0.11*** (0.02)	-0.12*** (0.02)
sqrt_amtout_issr	-5.7×10^{-5}*** (1.4×10^{-6})	-5.2×10^{-5}*** (2.8×10^{-6})	-5.1×10^{-5}*** (3×10^{-6})	-5.2×10^{-5}*** (3×10^{-6})	-5.3×10^{-5}*** (3×10^{-6})	-5.3×10^{-5}*** (3.1×10^{-6})	-5.3×10^{-5}*** (3.1×10^{-6})	-5.2×10^{-5}*** (3×10^{-6})	-4.9×10^{-5}*** (2.9×10^{-6})	-4.9×10^{-5}*** (2.9×10^{-6})
sqrt(time to maturity)	0.07*** (0.002)	0.07*** (0.01)	0.07*** (0.010)	0.07*** (0.010)	0.24*** (0.03)	0.24*** (0.03)	0.24*** (0.03)	0.24*** (0.03)	0.24*** (0.03)	0.23*** (0.03)
sqrt(amount outstanding)	-0.0001*** (6.1×10^{-6})	-0.0002*** (1.3×10^{-5})	-0.0002*** (1.9×10^{-5})	-0.0002*** (1.9×10^{-5})	-0.0003*** (2×10^{-5})	-0.0003*** (1.8×10^{-5})	-0.0003*** (1.8×10^{-5})	-0.0003*** (1.8×10^{-5})	-0.0003*** (1.7×10^{-5})	-0.0003*** (1.7×10^{-5})
coupon rate	3.4*** (0.07)	3.3*** (0.19)	3.3*** (0.19)	3.3*** (0.19)	1.4*** (0.22)	1.5*** (0.19)	1.5*** (0.19)	1.5*** (0.19)	1.7*** (0.19)	1.7*** (0.19)
foreign	-5.9*** (0.47)	-5.5*** (0.79)	-5.7*** (0.82)	-5.3*** (0.81)	-1.5* (0.86)	-1.8** (0.81)	-1.8** (0.81)	-1.8** (0.81)	-1.6** (0.81)	-1.8** (0.83)
global	0.18 (0.20)	0.04 (0.40)	-0.03 (0.40)	-0.007 (0.40)	0.04 (0.29)	-0.14 (0.27)	-0.18 (0.27)	-0.21 (0.27)	-0.22 (0.27)	-0.34 (0.28)
finance	-4.8*** (0.22)	-4.5*** (0.43)	-4.5*** (0.42)	-4.5*** (0.42)	-1.2*** (0.42)	-1.1*** (0.42)	-1.1*** (0.42)	-1.2*** (0.41)	-1.2*** (0.41)	-1.3*** (0.40)
utility	-5.7*** (0.39)	-5.6*** (0.93)	-5.3*** (0.89)	-5.4*** (0.83)	-4.6*** (0.73)	-3.8*** (0.63)	-3.7*** (0.63)	-3.7*** (0.63)	-3.3*** (0.62)	-2.9*** (0.62)
Fixed-effects										
date		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# days traded (prv. year)			Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# days traded (prv. year)-date				Yes	Yes	Yes	Yes	Yes	Yes	Yes
rating					Yes	Yes	Yes	Yes	Yes	Yes
rating-date						Yes	Yes	Yes	Yes	Yes
trade size							Yes	Yes	Yes	Yes
trade size-date								Yes	Yes	Yes
dealer									Yes	Yes
dealer-date										Yes
Fit statistics										
Observations	883,600	883,600	883,600	883,600	874,895	874,895	874,895	874,895	874,895	874,895
R^2	0.04993	0.10811	0.10832	0.12570	0.12957	0.17713	0.17758	0.18346	0.19354	0.23565
Within R^2		0.01238	0.01053	0.01030	0.00983	0.01007	0.01008	0.01007	0.00967	0.00941

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

7 Proofs appendix:

7.0.1 Proof of lemma 3.1

Proof. For the proof, we use the informed player optimality condition (eq. 7) to write:

$$\sum_{k}^{n-1}\binom{n-1}{k}\pi^{n-1-k}(1-\pi)^k F_n(B)^k(v_h-B)=\pi^{n-1}\delta_r=\frac{pi^{n-1}\delta_r}{\pi}$$

The denominator of the last term equals $\pi^n\delta_r$, that is, exactly the value of bidding $v_h-\delta_r$ when the market is populated with $n+1$ informed players. Applying eq. 7 once again, we get:

$$\sum_{k}^{n-1}\binom{n-1}{k}\pi^{n-1-k}(1-\pi)^k F_n(B)^k(v_h-B)=\frac{pi^{n-1}\delta_r}{\pi}=\frac{\sum_k^n\binom{n}{k}\pi^{n-k}(1-\pi)^k F_{n+1}(B)^k}{\pi}=$$
$$\sum_{k}^{n}\binom{n}{k}\pi^{n-1-k}(1-\pi)^k F_{n+1}(B)^k$$

Which can be further simplified to:

$$\sum_{k}^{n-1}\binom{n-1}{k}\pi^{n-1-k}(1-\pi)^k F_n(B)^k(v_h-B)=$$
$$\sum_{k}^{n}\frac{n}{n-k}\binom{n-1}{k}\pi^{n-1-k}(1-\pi)^k F_{n+1}(B)^k>\sum_{k}^{n-1}\binom{n-1}{k}\pi^{n-1-k}(1-\pi)^k F_{n+1}(B)^k$$

An inequality that implies: $F_n(B)>F_{n+1}(B)$. □

7.0.2 Proof of lemma 3.2

Proof. Uninformed players outbid any bid below B^U. Hence, such a bid yields zero payoff. Also, no matter how the informed acts, there is always a probability of at least $(1-\pi^n)*(1-q_h)$ that a bond bought by the uninformed is of low quality. Hence $B^U<v_h$. Thus, there exists a $\epsilon>0$ such that bidding $v_h-\epsilon>B^U$ when the quality is good yields a strictly positive return for the informed. Given this option, bidding below B^U and getting zero violates optimality. □

7.0.3 Proof of lemma 3.3

Proof. If $\frac{\pi^n q_h}{\pi^n q_h+(1-q_h)} > 1 - \frac{\delta}{v_h - v_l}$ we find that:

$$\nu(n^*) = \frac{\pi^{n^*} q_h}{\pi^{n^*} q_h + 1 - q_h} v_h + (1 - \frac{\pi^{n^*} q_h}{\pi^{n^*} q_h + 1 - q_h}) v_l > v_h - \delta_r$$

Where $\nu(n^*)$ is the expected value of the bond given that the informed always bids above it. If $B^U(n^*) < \nu$ there exists a bid, $\nu(n^*) - \epsilon > v_h$ such that bidding generates strictly positive expected payoff to the uninformed - contradicting optimality. Of course, if $B^U(n^*)$ is above ν we violate the demand that the uninformed bid is not a losing bet. $B^U(n^*)$ itself satisfies optimality - deviating to bidding something lower yields zero while bidding anything higher means getting less than zero in expectation.

Assume that condition 10 does not hold. It is plain to see that, in this case, the informed cannot break even by giving any bid that may appeal to non-distressed players. If 12 holds, we can show that the uninformed will submit a bid that breaks even that appeals only to distressed customers based on the same logic explained above.

If neither condition holds, $B^U(n^*) < v_h - \delta_s$ and the uninformed bid never wins the high-quality bond. Any bid above v_l will mean paying more than the expected value, while any bid below it will generate a strictly positive payoff, implying the existence of a profitable deviation. □

7.0.4 Proof for lemma 3.5

Proof. The claim states that if informed dealers target only the distressed players with some positive probability, they must entirely refrain from submitting bids on the higher interval range $[v_h - \delta_s, v_h - \delta]$. The intuition is simple: bidding even a bit higher and winning access to another market segment, that is, the non-distressed players, is preferable.

Note that it can never be the case that $F()$ is strictly increasing on an interval $(v_h - \delta - \epsilon, v_h - \delta + \epsilon)$. If there was a range ϵ such that the profits from buying at a bid of $v_h - \delta + \epsilon$ from all customers would strictly dominate buying only from distressed players (ξ of the population) for $v_h - \delta - \epsilon$. □

7.0.5 Proof of lemma 3.6

Proof. Let $\xi > \frac{\delta}{(v_h - \bar{R}^d)\pi^{n-1}}$. Set $\underline{B}^I = \bar{R}^d$. Note that our assumption implies:

$$\xi(v_h - \bar{R}^d)\pi^{n-1} > v_h - (v_h - \delta)$$

The LHS is the expected return from bidding $\bar{R}^d$, while the RHS is the return from submitting the bid $v_h - \delta$ when all other informed players refrain from targeting this market segment. The inequality suggests two things: (i) there is no profitable deviation to soliciting the non-distressed players, as even the best one available is not sufficient, and (ii) by continuity, we can find: $\bar{B}^I$ such that:

$$\xi(v_h - \bar{R}^d)\pi^{n-1} = v_h - (v_h - \bar{B}^I); \quad \text{and} \quad \bar{B}^I < v_h - \delta$$

Next, assume that: $\xi < \frac{\delta}{(v_h - \bar{R}^d)}$. Thus:

$$\xi(v_h - \bar{R}^d)\pi^{n-1} < v_h - (v_h - \delta)\pi^{n-1}$$

Where the RHS is the return of bidding $v_h - \delta$ (recall that informed dealers only submit bids that are higher than $v_h - \delta$). At the same time, the LHS is the return of the most profitable deviation from this equilibrium into the range in which only the distressed buy. We can see that there is no profitable deviation. From here, we can construct the complete equilibrium as usual. We can negate the existence of other equilibria in the following way. If there was an equilibrium in which some dealers were bidding below $v_h - \delta$, these bids would have won with a lower likelihood than $v_h - \delta$ and would not be accepted by non-distressed. Also, someone would always bid $\bar{R}^d$. Using the equation above, we can show the rest.

Last, let $\xi \in (\frac{\delta}{(v_h - \bar{R}^d)}, \frac{\delta}{(v_h - \bar{R}^d)\pi^{n-1}})$. Let $\bar{B}^I_d$ denote the highest bid that an informed player gives that is accepted only by distressed customers. Note that it is the lowest bid only when $v_h - \delta$ is the lowest bid. Thus:

$$(v_h - (v_h - \delta)) = \xi(v_h - \bar{B}^I_s)$$

Or:

$$\bar{B}^I_s = v_h - \frac{\delta}{\xi}$$

Now, we need to ensure that the interval $[\underline{B}^i_s, \bar{B}^I_s]$ is not empty. That is:

$$\bar{B}^I_s = v_h - \frac{\delta}{\xi} \geq \bar{R}^d$$

Which is satisfied by:

$$\xi \geq \frac{\delta}{v_h bybar R^d}$$

Also, note that trivially $\bar{B}^I_s < v_h - \delta$ as required. Thus, we can build an equilibrium in which some bids are accepted by all sellers and some only by the distressed ones.

It is plain to see that we cannot construct an equilibrium with trades only with the distressed guys since that requires that $\xi \geq \frac{\delta}{(v_h - \bar{R}^d)\pi^{n-1}}$ (see above). Similarly, we cannot construct an equilibrium in which all bids are accepted by the non-distressed as this will require $\xi < \frac{\delta}{(v_h - \bar{R}^d)}$ □

7.0.6 Proof of theorem 3.7

Proof. We know from the indifference condition that whenever an informed player bids on a high-quality asset, its expected return is:

$$\pi^{n-1}(v_h - \bar{R}^D)$$

The total expected return is the number of bids that the dealer gets to submit multiplied by the expected return from each bid, or:

$$\text{Expected profit} = \text{nm. of bids} \times \text{expected profit from bidding} = \mu_s(1-\pi)\pi^{n-1}(v_h - \bar{R}^D)$$

The total realized return is the number of trades the dealer conducted multiplied by the average spread per trade. The average spread will be denoted by: $\bar{S}^I(n)$. Note that the average spread measures the realized prices at which bonds were bought rather than the submitted bids.

From symmetry between the informed dealers, each executes $1/n$ of the trades they do as a group. To compute the total volume the group trades, note that there are two scenarios in which a customer is not trading with an informed dealer: (i) None of them submits a bid (prob. of

π^n). or, (ii) the bids they submitted are below the customer reservation value. Specifically, the customer got hit by a mild liquidity shock (δ_r) while the bids are such that they appeal only to players that are in distress (shock of δ_s). The likelihood that any of these scenarios take place is:

$$\pi^n + (1-\pi^n) * (1-\xi) \sum_{k=1}^{n} \binom{n}{k} \pi^{n-k}(1-\pi)^k (F(v_h - \delta_r))^k$$

Recalling that the total measure of customers in the market is μ_s, we find that the total volume of trade facilitated by the informed is:

$$V^i(n) = \mu_s(\pi^n + (1-\pi^n) * (1-\xi) \sum_{k=1}^{n} \binom{n}{k} \pi^{n-k}(1-\pi)^k (F(v_h - \delta_r))^k)$$

Now, we can pin down the realized profit of the informed dealer:

Realized profit = measure of buys by the dealer $\times$ avg. spread earned when selling =

$$\frac{V^i(n)}{n} * \bar{S}^i(n) = (1-\pi^n - (1-\pi^n) * (1-\xi) \sum_{k=1}^{n} \binom{n}{k} \pi^{n-k}(1-\pi)^k (F(v_h - \delta_r))^k * \frac{\mu_s}{n} * \bar{S}^i(n)$$

Since there is a continuum of customers, the dealer's total expected profit from trade equals its total realized profit. Equating the two terms, we get:

$$(1-\pi^n-(1-\pi^n)*(1-\xi) \sum_{k=1}^{n} \binom{n}{k} \pi^{n-k}(1-\pi)^k (F(v_h-\delta_r))^k * \frac{\mu_s}{n} * \bar{S}^i(n) = \mu_s(1-\pi)\pi^{n-1}(v_h-\bar{R}^D)$$

Rearranging:

$$\hat{S}^i(n) = \frac{n(1-\pi)\pi^{n-1}(v_h - \bar{R}^s)}{V^i(n)}$$

Where:

$$V^i(n) = (1-\pi^n)[1-(1-\xi) \sum_{k=1}^{n} \binom{n}{k} \pi^{n-k}(1-\pi)^k (F(v_h - \delta))^k]$$

Using the same logic, we find that when all bids submitted by the informed are more significant than $v_h - \delta_r$:

$$S^i(n) = \frac{n(1-\pi)\pi^{n-1}(v_h - \bar{R}^r_n)}{1-\pi^n}$$

When all the bids it submits are below $v_h - \delta_r$:

$$S^i(n) = \frac{n(1-\pi)\pi^{n-1}(v_h - \bar{R}^s_n)}{1-\pi^n}$$

$\square$

References

Akerlof, G. A. (1978): "The market for âlemonsâ: Quality uncertainty and the market mechanism," in *Uncertainty in economics*, Elsevier, 235–251.

Bao, J., M. O'Hara, and X. (Alex) Zhou (2018): "The Volcker Rule and Corporate Bond Market Making in Times of Stress," *Journal of Financial Economics*, 130, 95–113.

Bao, J., J. Pan, and J. Wang (2011): "The illiquidity of corporate bonds," *The Journal of Finance*, 66, 911–946.

Benmelech, E. and N. K. Bergman (2018): "Credit market freezes," *NBER Macroeconomics Annual*, 32, 493–526.

Brancaccio, G., D. Li, and N. Schürhoff (2017): "Learning by trading: The case of the US market for municipal bonds," *Unpublished paper. Princeton University.*

Burdett, K. and K. L. Judd (1983): "Equilibrium Price Dispersion," *Econometrica*, 51, 955–969.

Camargo, B. and B. Lester (2014): "Trading dynamics in decentralized markets with adverse selection," *Journal of Economic Theory*, 153, 534–568.

Chaderina, M. and V. Glode (2022): "Trading with Expert Dealers," *Available at SSRN.*

Chalamandaris, G. and N. E. Vlachogiannakis (2020): "Adverse-selection considerations in the market-making of corporate bonds," *The European Journal of Finance*, 26, 1673–1702.

Chen, H., R. Cui, Z. He, and K. Milbradt (2018): "Quantifying liquidity and default risks of corporate bonds over the business cycle," *The Review of Financial Studies*, 31, 852–897.

Chiu, J. and T. V. Koeppl (2016): "Trading dynamics with adverse selection and search: Market freeze, intervention and recovery," *The Review of Economic Studies*, 83, 969–1000.

Choi, J., Y. Huh, and S. S. Shin (2021): "Customer liquidity provision: Implications for corporate bond transaction costs," *Available at SSRN 2848344*.

Cohen, A., M. Kargar, B. Lester, and P.-O. Weill (2022): "Inventory, Market Making, and Liquidity: Theory and Application to the Corporate Bond Market," .

Dick-Nielsen, J. and M. Rossi (2019): "The Cost of Immediacy for Corporate Bonds," *The Review of Financial Studies*, 32, 1–41.

Duffie, D. (2020): "Still the World's Safe Haven? Redesigning the U.S. Treasury Market after the COVID-19 Crisis," .

Duffie, D., N. Gârleanu, and L. H. Pedersen (2005): "Over-the-counter markets," *Econometrica*, 73, 1815–1847.

Easley, D. and M. O'hara (1987): "Price, trade size, and information in securities markets," *Journal of Financial economics*, 19, 69–90.

Gorton, G. and A. Metrick (2012): "Securitized banking and the run on repo," *Journal of Financial economics*, 104, 425–451.

Green, R. C., B. Hollifield, and N. Schürhoff (2007): "Financial Intermediation and the Costs of Trading in an Opaque Market," *The Review of Financial Studies*, 20, 275–314.

Guerrieri, V. and R. Shimer (2014): "Dynamic adverse selection: A theory of illiquidity, fire sales, and flight to quality," *American Economic Review*, 104, 1875–1908.

Kacperczyk, M., C. Sialm, and L. Zheng (2005): "On the industry concentration of actively managed equity mutual funds," *The Journal of Finance*, 60, 1983–2011.

Kargar, M., B. Lester, D. Lindsay, S. Liu, P.-O. Weill, and D. Zúñiga (2021): "Corporate Bond Liquidity during the COVID-19 Crisis," *The Review of Financial Studies*, 34, 5352–5401.

Lagos, R. and G. Rocheteau (2009): "Liquidity in Asset Markets With Search Frictions," *Econometrica*, 77, 403–426.

O'Hara, M. and X. A. Zhou (2021): "Anatomy of a liquidity crisis: Corporate bonds in the COVID-19 crisis," *Journal of Financial Economics*, 142, 46–68.

Van Nieuwerburgh, S. and L. Veldkamp (2010): "Information acquisition and under-diversification," *The Review of Economic Studies*, 77, 779–805.

Varian, H. R. (1980): "A Model of Sales," *The American Economic Review*, 70, 651–659.

Veldkamp, L. (2014): *Information Choice in Macroeconomics and Finance.*

Weill, P.-O. (2020): "The Search Theory of Over-the-Counter Markets," *Annual Review of Economics*, 12, 747–773.

Zou, J. (2019): "Information Traps in OTC Markets," .