

Sequential Search for Corporate Bonds*

Mahyar Kargar[†] Benjamin Lester[‡] Sébastien Plante[§]
Pierre-Olivier Weill[¶]

January 30, 2023

Abstract

In over-the-counter (OTC) financial markets, customers search for trades by making repeated inquiries to dealers. Yet, there is little direct empirical evidence of this sequential search process, since existing transaction data only provides information about the times customers complete their trade, but no information about the times they search for a trade. In this paper, we shed new light on customers' sequential search process by leveraging a complete record of inquiries—successful and not—made on the leading electronic trading platform for corporate bonds. We obtain estimates of time to trade and trading costs, conditional on observable trade characteristics and the number of previously unsuccessful inquiries. We find that after the first failed inquiry, it takes two to three days for a customer to purchase an investment-grade bond. When interpreted through the lens of a sequential search model, our estimates highlight the importance of both observed *and unobserved* heterogeneity across customers. Overall, these estimates can serve as useful inputs into quantitative applications of search models and guide future theoretical explorations of sources of search frictions in OTC markets.

Keywords: Over-the-counter markets, corporate bonds

JEL Classification: G11, G12, G21.

*We thank Briana Chang, Tatyana Deryugina, Julia Fonseca, Mark Garmaise, Tim Johnson, Kumar Venkataraman, and conference and seminar participants at AEA/CEANA 2023, NBER Big Data and Securities Markets Spring 2023, INSEAD, UC Irvine, UCLA Macro-Finance Lunch, UC Riverside, and UIUC for comments and suggestions. We thank MarketAxess for providing the main data set used in this paper. The views expressed here are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. All errors are our own responsibility.

[†]University of Illinois at Urbana-Champaign. Email: kargar@illinois.edu

[‡]Federal Reserve Bank of Philadelphia. Email: benjamin.lester@phil.frb.org

[§]University of Wisconsin-Madison. Email: splante@wisc.edu

[¶]UCLA, NBER and CEPR. Email: poweill@econ.ucla.edu

1 Introduction

Over-the-counter (OTC) markets play a key role in the U.S. financial system: they include most fixed income securities, asset-backed securities, repurchase agreements, and various types of derivatives, along with a significant fraction of equity trading volume. OTC markets are different from exchange-based markets because they are decentralized: participants must first find a willing counterparty and then agree on the terms of trade. Search models have emerged as one of the workhorse theoretical and quantitative frameworks to analyze these markets, following the seminal work of [Duffie, Gârleanu, and Pedersen \(2005\)](#) (see [Weill, 2020](#) for a survey). However, the main assumption that investors search for counterparties is not grounded in direct empirical evidence about the sequential trade process. The reason is simple: existing data from OTC markets is comprised of transaction records, which contain information about the time and price at which a trade occurs, but not about the time investors spend searching for a counterparty.¹

In this paper, we leverage a proprietary data set to offer a unique window into the sequential search process of investors in one of the most studied OTC financial markets—the market for U.S. corporate bonds. The data provides a complete record of all inquiries made by customers, and the corresponding replies from dealers, on the leading electronic trading platform for corporate bonds, MarketAxess (MKTX). Crucially, by observing both successful *and unsuccessful* inquiries, the data allows us to estimate how long it takes a customer to execute a trade and how this length of time depends on the properties of the order and the characteristics of the customer. Moreover, by studying the behavior of both the customer and dealers over the course of the sequential search process, our analysis also offers new insights into the sources of delays in the trading process.

We start by documenting that inquiries fail to result in trade quite often—about a third of the time—which is consistent with the findings of [Hendershott and Madhavan \(2015\)](#) from an earlier time period. We go beyond this earlier work by analyzing the behavior of customers who, shortly after a failed inquiry, return to the market to make new inquiries for the same quantity of the same

¹This limitation stands in stark contrast with other applications of search models. For example, data on unemployment spells is informative about workers’ sequential job-finding process, while observations of time-on-the market for houses are informative about the sequential process for selling a home.

bond. In fact, by combining the data from MKTX with additional data from the Trade Reporting Compliance Engine (TRACE), we can observe when customers make additional electronic inquiries on MKTX for the same trade, when they complete the trade on MKTX, when they complete the trade *outside* of the electronic platform (via the traditional voice channels), and when they abandon the trade altogether.

Studying the details of the search process at such a granular level leads to novel estimates of the trading frictions that exist in the corporate bond market. For example, we find that it takes two to three days for a customer to complete the purchase of an investment-grade bond after an initial inquiry fails. Hence, given that approximately 70% of requests are filled at the first inquiry, a lower bound for the unconditional time to trade is about one day.² In addition, our analysis reveals that time to trade varies systematically across different characteristics of the order: time to trade is shorter for sells (relative to buys), small trades (relative to larger trades), and investment-grade bonds (relative to high-yield bonds). We also find that time to trade differs significantly across customers, as more “connected” investors get a larger number of responses to their inquiries and trade more quickly.

We also document the characteristics of contact rates and dealers’ replies over the course of the sequential search process. We find that customers appear to make inquiries on MKTX more frequently as the number of failed inquiries increases, but these inquiries get fewer replies, the best offer gets worse, and the probability of trading falls. These dynamics could be an indication that the terms of trade worsen over the course of the search process, or they could reflect selection based on unobservables. We find evidence of the latter. When interpreted through the lens of a sequential search model, our estimates suggest that customers are heterogeneous in the intensity with which they make inquiries and in the number of responses they are able to elicit from dealers.

We believe that our analysis generates three main contributions to the existing literature. Our first contribution is to organize the data in a way that reveals customers’ sequential search process. We do so by observing that when a customer becomes active on MKTX, she often submits a cluster

²This estimate is a lower bound because we do not have information about how long the customer was searching prior to submitting their first inquiry on MKTX.

of inquiries for a particular bond within a short period of time, rather than submitting one large inquiry. Following the practice in the equity market, we call this cluster a “parent order.” Some of the inquiries in a given parent order are for different quantities of the same bond, a form of order splitting.³ However, other inquiries can be identified as repeated attempts to trade a specific quantity; we call these “child orders.” Because order splitting may be viewed as evidence of asymmetric information rather than search frictions, we focus on the sequential trade process for child rather than parent orders.

Summary statistics about child orders immediately reveal several interesting insights. First, sequential search is, indeed, a prominent feature of the trading process: as noted above, the probability that an inquiry is unsuccessful is about 30 percent, and customers routinely submit repeat inquiries. Conditional on the first inquiry failing, the median number of inquiries in a child order is 2, the 75th percentile is 3, and the 99th percentile is 9. Second, trade is non-exclusive: a customer may eventually trade on the voice market instead of MKTX, which we can observe using the enhanced version of the TRACE data set. The third takeaway is that the probability a child order ends without trade is significant, either because the investor abandons the trade altogether or creates a new child order by submitting an inquiry for a different amount. Lastly, organizing trades into parent and child orders has a significant effect on estimates of trading probabilities and time to trade. For example, we estimate that approximately 80% of child orders are eventually filled, whereas the existing literature (e.g., [Hendershott and Madhavan, 2015](#)) arrives at an estimate of approximately 67% when individual inquiries for trade are treated independently.

After organizing the data into parent and child orders, our second contribution is properly measuring the time it takes to successfully trade child orders. Even with our granular observations of inquiries and trade, measuring the time to trade remains a nontrivial exercise because of two potential sources of bias. The first is survivor bias created by “competing risks.” For example, the measured average time to trade on MKTX is biased downwards because it is based on trades that have occurred relatively quickly, before the arrival of other events such as trading with a dealer outside

³Czech and Pintér (2020) provide evidence of informed investors splitting their orders across multiple dealers in the UK corporate bond market.

the platform or deciding to abandon this particular child order. The second concern is selection bias: since different types of inquiries (and customers) trade with different probabilities, we must account for potential changes in the composition of inquiries—both observable and unobservable—over the course of the sequential search process. We attempt to correct for these biases via Maximum Likelihood estimation: we assume that successful inquiries on MKTX, unsuccessful inquiries on MKTX, voice trades, and exit (i.e., abandoning the child order and/or beginning a new one) occur at independent exponential times with intensities that depend on the characteristics of the child order.

Equipped with these estimated intensities, we calculate the expected time it takes for a child order to trade, either on MKTX or via voice trade, in the event that the customer does not exit the child order. We find that it takes between two to three days to complete a trade *after a first failed inquiry*. Trade is much faster, by about a day, for sales than purchases. Block trades (with size above \$5 million) take about one day longer to trade than micro-size ones (with size below \$100,000). Bonds with amounts outstanding below the median take half a day more to trade. Bond age, turnover, and credit rating all have statistically significant impacts on time to trade, but with economically small effects. Customer connectedness, measured by the average number of responses that a particular customer elicits on MKTX, has an economically significant impact on time to trade. Moreover, we observe a large increase in time to trade, by more than a day, during the pandemic-induced crisis of March 2020, when the corporate bond market suffered severe liquidity disruptions (Kargar, Lester, Lindsay, Liu, Weill, and Zúñiga, 2021). Finally, comparing time to trade on MKTX and voice, we observe that child orders trade much faster on MKTX for all size categories except block trades, suggesting customers use the platform for execution quality rather than price discovery.

These estimates are helpful for at least two reasons. First, they can be directly applied to quantitative analyses based on search-theoretic models, since the arrival rates we estimate are crucial, yet controversial inputs that are typically identified via indirect inference. Second, the correlations we find between our estimates of time to trade and other observable outcomes provide

a natural starting point for additional empirical and theoretical work. For example, the fact that it takes longer for a customer to buy a bond than it does to sell it suggests that it might take time for dealers to locate (or “source”) a bond. Alternatively, it could indicate that sellers are more distressed, on average, than buyers.

Our third contribution derives from studying the dependence of various trading outcomes on the number of previously unsuccessful inquiries, which helps us unpack the evolution of customers’ and dealers’ behavior over the course of the sequential search process. After controlling for various observed characteristics of child orders, we find that, with each additional failed inquiry, the number of dealer responses declines, the terms of trade deteriorate, and the time to trade increases. This dependence on the number of failed inquiries could, in principle, occur for two different reasons. First, dealers could adjust their behavior when they recognize a customer has attempted multiple, unsuccessful inquiries, perhaps because of information leakages or what [Zhu \(2011\)](#) calls the “ringing-phone curse.” Alternatively, this dependence could derive from the fact that child orders differ in characteristics unobserved by the econometrician. We find evidence in support of the latter, rather than the former. In particular, when we repeat our estimation with child-order fixed effects, which control for all characteristics, observed and unobserved, the dependence of outcome variables on the number of failed inquiries largely disappears.

But what accounts for these child order fixed effects? In principle, it could be an unobserved characteristic of the customer or of the market at the time of the child order. We propose a simple theory that incorporates unobserved heterogeneity into the standard sequential search model in the tradition of [McCall \(1970\)](#), and argue that the dependence of outcome variables on the number of failed inquiries is consistent with child orders differing in their inquiry intensity and in the average number of responses they can elicit from dealers. More specifically, some child orders have low inquiry intensity and many responses, while others have high intensity and low responses.

Related literature

Our work is most closely related to the few other papers that have used the proprietary data from MKTX to analyze the impact of electronic trading on corporate bond market conditions (e.g., [Hendershott and Madhavan, 2015](#); [O’Hara and Zhou, 2021](#); [Hendershott, Livdan, and Schürhoff, 2021](#)). Our analysis differs from these papers in both our focus and our approach. More specifically, we are the first to organize the MKTX data into parent and child orders in order to offer new evidence about the sequential search process, including novel estimates of the time to trade conditional on the characteristics of the trade (size, direction, bond rating, and customer connectedness).

Our work also contributes to the vast empirical literature that studies corporate bond market liquidity based on transaction data. Some prominent examples include [Schultz \(2001\)](#), [Bessembinder, Maxwell, and Venkataraman \(2006\)](#), [Edwards, Harris, and Piwowar \(2007\)](#), [Goldstein, Hotchkiss, and Sirri \(2007\)](#), [Bao, Pan, and Wang \(2011\)](#), [Bessembinder, Jacobsen, Maxwell, and Venkataraman \(2018\)](#), and many others.⁴ Our contribution relative to this literature is our attempt to measure the time to trade and to empirically investigate the sequential search process of customers in the corporate bond market. Our estimates of time to trade provide new empirical evidence on a dimension of liquidity at the center of search-theoretic models. [Hendershott, Li, Livdan, and Schürhoff \(2020\)](#) pursue similar goals but for a different dimension of liquidity (the cost of trade failures) in a different market (the market for collateralized loan obligations).

Our attempt to measure time to trade is related to earlier works in the OTC search literature which have proposed strategies to identify investors’ search intensities. For example, according to the model of [Afonso and Lagos \(2015\)](#), [Üslü \(2019\)](#), and [Brancaccio and Kang \(2021\)](#), when search is random and the distribution over agents’ state is continuous, every meeting results in a trade. This allows one to identify the search intensity from the trading intensity. While this identification strategy is reasonable for dealers, it is problematic for customers who presumably spend long periods of time out of the market: clearly, observing that a customer trades once a year does not imply that it takes a year to find a counterparty. [Gavazza \(2016\)](#) addresses this issue using

⁴See [Bessembinder, Spatt, and Venkataraman \(2020\)](#) for a survey

a structural model, taking advantage of aggregate information about the total number of real assets (in his case, aircraft) for sale at a time. [Pintér and Üslü \(2021\)](#), alternatively, use joint observation of trade size and frequency to indirectly identify search intensities. We propose a more direct approach, based on granular observations, which does not rely on the restrictions imposed by a specific structural model.

Finally, our approach is related to the large literature that attempts to estimate the key objects of interest in the standard sequential search model of [McCall \(1970\)](#), which was first used in financial economics by [Garbade and Silber \(1976\)](#). Early attempts to do so in a labor market context include [Kiefer and Neumann \(1979\)](#) and [Flinn and Heckman \(1982\)](#), among others. As in labor economics, this simple partial equilibrium model is a natural starting point for interpreting micro data, as it helps rationalize failed inquiries, repeated attempts to trade, and price dispersion.⁵ However, while we find it useful to formulate a search-theoretic model to motivate our empirical exercise and interpret its findings, it's important to note that our measurement does not impose theoretical restrictions from the model.

2 Data

Our main source of data is MarketAxess (MKTX), the leading electronic trading platform in the corporate bond market. Prior to the introduction of MKTX, in 2000, the corporate bond market operated almost exclusively under a “voice-based” trading system, whereby customers would sequentially contact dealers (via telephone or chat) one at a time to solicit a quote. Stepping into this market, MKTX offered a trading platform allowing buy-side traders (henceforth customers) to query multiple dealers at once via an electronic request for quote (RFQ), thus reducing the time-consuming process of gathering quotes and potentially increasing competition across dealers.

⁵Naturally, understanding the process that generates dealers' offers requires expanding the model to include an explicit analysis of the market structure of the dealer sector, along with the optimal strategies of (potentially heterogeneous) dealers. Given the scope of the current paper, we leave this extended analysis for future work.

As of the third quarter of 2022, MKTX accounts for approximately 21% of total trading volume in the corporate bond market.⁶

When requesting a quote on the MKTX platform, customers specify the bond they wish to trade, the desired quantity, the trade direction or “side” (buy or sell), and the duration of the auction (usually between 5 and 20 minutes). Once submitted, a customer inquiry is sent to a list of pre-authorized dealers.⁷ On the receiving end, dealers observe the details of the inquiry, including the customer’s identity. The receiving dealers may respond to the inquiry with a quote, but are not obligated to do so. At the end of the auction, customers observe the terms of the replies (if any), and can choose to either accept one of the offers or pass.⁸

Our sample from MKTX covers all trading activity from January 3, 2017 to March 31, 2021. The data contain detailed information on customer inquiries, dealer responses, and customer trading decisions. More specifically, for each inquiry, we observe the submission time (stamped at the second), an anonymized customer identifier, the CUSIP (Committee on Uniform Securities Identification Procedures) number of the requested bond, the requested quantity, the trade side (buy or sell), the number of dealers who received the request, and several other attributes. For every response to an inquiry, we observe the anonymized identifier of the responding dealer together with his quote. For inquiries that result in a transaction, we observe the time at which trade occurs and the terms of trade. Note that we observe all inquiries, including those that do *not* result in a trade, either because the inquiry receives no responses or because the customer chooses to reject all responses.

Importantly, when an inquiry fails to trade on MKTX, a customer may trade outside the platform via voice. In the next section, we describe how we attempt to identify these trades using the enhanced version of the Trade Reporting Compliance Engine (TRACE) data set provided by

⁶Source: MarketAxess quarterly report for 2022Q3, available from: <https://investor.marketaxess.com>.

⁷Starting in 2012, MKTX initiated Open Trading, a trading protocol that enables all-to-all trading in the corporate bond market. This protocol allows other investors as well as non-pre-authorized dealers to respond to requests for quotes. Approximately 15% of MKTX auctions are won by responses submitted through Open Trading. For a comprehensive analysis of the Open Trading protocol, see [Hendershott, Livdan, and Schürhoff \(2021\)](#).

⁸The main variation in dealers’ offers is price. In principle, dealers can respond to an offer with a different quantity, but in practice more than 97% of dealer responses are at the quantity level requested by the customer.

FINRA. The TRACE database contains detailed reports of every successful trade, whether it has an electronic or voice origin. When working with TRACE, we filter the data following the standard procedure laid out in [Dick-Nielsen \(2014\)](#). We merge the cleaned data set with the Mergent Fixed Income Securities Database (FISD) to obtain bond fundamental characteristics (e.g., credit ratings, amount outstanding, coupon rates, etc.) Following the bulk of the academic literature, we exclude variable-coupon, convertible, exchangeable, and puttable bonds, as well as asset-backed securities, privately placed instruments, and foreign securities. We also exclude primary market transactions.

Finally, we measure transaction costs as a markdown or markup relative to the benchmark provided by MKTX, called Composite+ (CP+).⁹ CP+ is the proprietary algorithmic pricing engine for corporate bonds from MKTX. It is designed to provide an unbiased two-sided market forecast for institutional-size trades. The engine outputs reference bid and ask prices at a high frequency (every 15 to 60 seconds). These forecasts can be used to benchmark a significant fraction of TRACE records: 90% of high-yield TRACE records can be matched to a standing CP+ forecast; that figure goes up to 95% for investment-grade bonds.

The construction of the forecasts follows two steps. First, MKTX trains a machine learning (ML) algorithm using three distinct sources of bond trading data: (1) historical TRACE prints; (2) indicative bond price data streamed by dealers; and (3) request for quote responses sent by liquidity providers on the MKTX trading platform. Beyond trading data, MKTX uses bond level information and other broad market data, such as CDX levels, to train the prediction engine. The engine is recalibrated overnight at a daily frequency. Second, the calibrated engine is used over the next trading day to generate real-time reference bid and ask prices of individual bonds using all available intraday information.

2.1 The query process: parent and child orders

To give the reader a sense of how the query process works and to motivate the way we organize and analyze the data, we believe it is helpful to present some representative examples of inquiries.

⁹For more details about Composite+, see <https://www.marketaxess.com/price/composite-plus>.

First, panel (a) of Table 1 provides an example of a successful inquiry. In this example, a customer submitted an inquiry to buy \$300,000 in par value of an investment-grade bond issued by Bank of America. The customer received six replies from dealers, whose anonymized identifiers are provided in column (6). Note that, because the bond in question is investment-grade, dealer responses in column (7) are in terms of yield spread relative to a benchmark Treasury bond (a higher yield spread implies a lower purchasing price). As we can see from this column, dealers' quoted yield spreads vary between 126.37 and 129.70 basis points. In the second row of column (9), the entry "Done" shows that the customer accepted the best (highest) offer.¹⁰ In our sample, 67% of all inquiries result in a successful trade.¹¹

[Table 1 about here.]

Panel (b) of Table 1 provides an example of unsuccessful inquiry. This inquiry was submitted by the same customer and for the same bond as the inquiry reported in panel (a), but this time, the customer requested to purchase an amount of \$490,000 in par value instead of \$300,000. A total of nine dealers responded to the customer's new request. By comparing the identifiers of responding dealers for both inquiries, we see that five of the six dealers who responded to the first inquiry also responded to the second untraded inquiry. Four additional dealers, who had not replied to the first inquiry, replied to the second inquiry. However, the customer decided to pass on the best offer (a yield spread of 127.01), as indicated by the "did not trade" (DNT) flag in the last column. In our sample, 16% of inquiries that receive at least one response do not trade. An additional 18% of inquiries do not receive any response.

While customer inquiries are informative about the trading process in and of themselves, a careful examination of the data reveals that individual inquiries are often parts of larger trading orders. As a result, individual inquiries should not always be treated as independent observations. To help the reader see why, Table 2 reports all the inquiries that the customer in our previous

¹⁰In the last row of column (9) in Table 1, the entry "Cover" identifies the second best offer. MKTX informs dealers who submit the second best offer of the rank of their quote. Dealers who submit lower-ranked offers do not learn their relative position in the auction.

¹¹While we examine a different time period, we find a fraction of successful trade consistent with the findings of Hendershott and Madhavan (2015).

examples (Table 1) submitted to purchase this particular Bank of America bond over a six month period. To save space, we do not report the responses that each inquiry received, and report only whether or not a given inquiry resulted in a trade (see column 7). Note that the first and second inquiries reported in Table 2 correspond to the inquiries reported in panel (a) and panel (b) of Table 1.

Notice immediately that the customer made repeated *successful* purchase inquiries for the same bond over an eight day period. Of the six inquiries, four were successful and led to the purchase of 300, 490, 290, and 680 bonds (with \$1,000 par value) for a total of 1760 bonds. This anecdotal evidence suggests that customers sometimes execute large orders by submitting a sequence of smaller inquiries. To give further credence to this interpretation, Figure 1 plots the daily number of purchase inquiries submitted by this customer over a six month horizon. The figure makes clear that the customer's inquiries over that horizon are concentrated in mid-August 2017, which supports the view that the individual inquiries are part of a larger order and not independent events.

The second noteworthy feature of Table 2 is that the customer twice followed an *unsuccessful* inquiry by resubmitting an identical inquiry (same bond, quantity, and trade side) soon afterward. This phenomenon is first observed after the second inquiry and again after the fourth. While both of these unsuccessful inquiries received multiple dealer responses, the customer chose to pass.¹² Hence, the example in Table 2 suggests that even when customers are able to simultaneously contact a large number of dealers, sequential search remains a feature of the trading process in the U.S. corporate bond market.

These patterns of trade are widespread. For example, about a third of trading volume can be attributed to a parent order with two or more child orders, and approximately a quarter of child orders have at least two inquiries. Hence, we argue that a natural first step is to organize RFQs into clusters, representing the total quantity of a particular bond that a customer is attempting to trade, which we refer to as the “parent” order. Within each parent order, we further partition the set of inquiries into sets of “child” orders in which the customer requests a specific quantity of the bond.

¹²For the second inquiry, this can be seen in panel (b) of Table 1. To save space, the responses associated with the fourth inquiry are not reported.

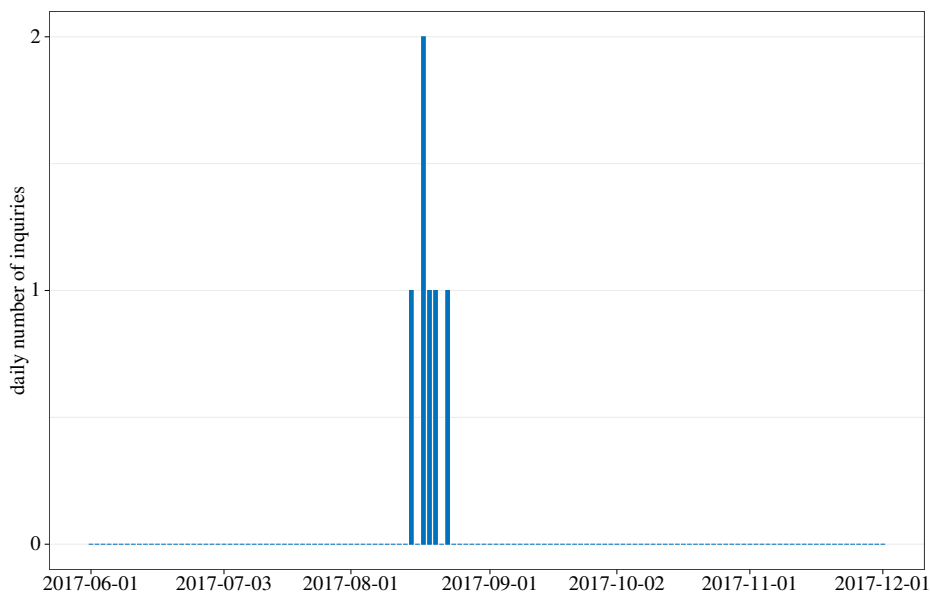


Figure 1. Inquiry cluster

This figure shows an inquiry cluster for a customer purchase for an 11-year, 3.824% investment-grade bond issued on January 17, 2017 by Bank of America over a six-month period in 2017.

We borrow the parent and child order terminology from the equity market literature on institutional trading where large (parent) orders are often split into smaller (child) orders for execution. In the example above, as one can see in columns (8) and (9) of Table 2, all six inquiries make up a single parent order—where the customer attempts to trade 1760 units of this particular bond over an eight day period—and this parent order is split into four smaller child orders.

[Table 2 about here.]

More precisely, since the data itself does not explicitly identify parent and child orders, we employ the following procedure to construct the appropriate identifiers. First, to construct parent orders, we group all inquiries made by a specific customer for a given bond and trade side until we do not observe a new inquiry with the same characteristics (customer, bond, trade side) for N_p days since the last inquiry. The time cutoff delimiting parent orders is admittedly arbitrary. In our main specification, we use a cutoff of five days. However, our main results are not sensitive to this choice; we obtain qualitatively similar results with a cutoff of ten days.

Second, we construct child orders by looking at repeated inquiries from a given customer for the same bond, the same trade side, *and the same requested quantity*. We consider all inquiries

with these characteristics as part of the same child order until either (i) the most recent inquiry of the child order led to an electronic trade on MKTX; (ii) the customer submitted a new inquiry requesting a different quantity, in which case we initiate a new child order with the updated quantity; or (iii) there is no new inquiry with the same characteristics (same customer, bond, trade side, and trade size) for more than N_c days, where $N_c \leq N_p$. When no new inquiry has been submitted for more than N_c days, we consider the execution of the child order unsuccessful on MKTX. Here again, the threshold N_c is arbitrary. While we use a cutoff of five days in our main specification, our main results are not sensitive to this choice.

There are two reasons why a child order may be unsuccessful on MKTX. First, the customer might have given up trading the bond. Second, the customer might have traded the bond via voice. These two outcomes have different economic implications and should be distinguished. Ideally, we would match customer inquiries on MKTX that result in a voice trade using the corresponding TRACE record. However, since TRACE does not report customer identities, it is impossible to match a child order that is traded via voice to its corresponding TRACE record with certainty. Fortunately, this issue can partially be overcome since most corporate bonds trade only a few times a day or less. As a result, the likelihood that two different customers would trade the same quantity of the same bond within a few days is arguably low. We thus infer the occurrence of a voice trade by verifying if there exists a record in TRACE with the same characteristics as the unsuccessful child order (same bond, traded quantity, trade side) within five days of that child order's last on MKTX. In the rare cases where there are multiple matches, we select the closest one in time.

2.2 Summary statistics

We could, in principle, conduct our analysis in two ways: at the level of parent orders, or at the level of child orders. However, the splitting of a parent order into child orders may be driven by considerations other than search, such as asymmetric information (as in, e.g., [Kyle, 1985](#)). For this reason, we find it more natural to study the sequential search process using child orders as our main unit of observation. We begin this section by presenting some summary statistics, explaining

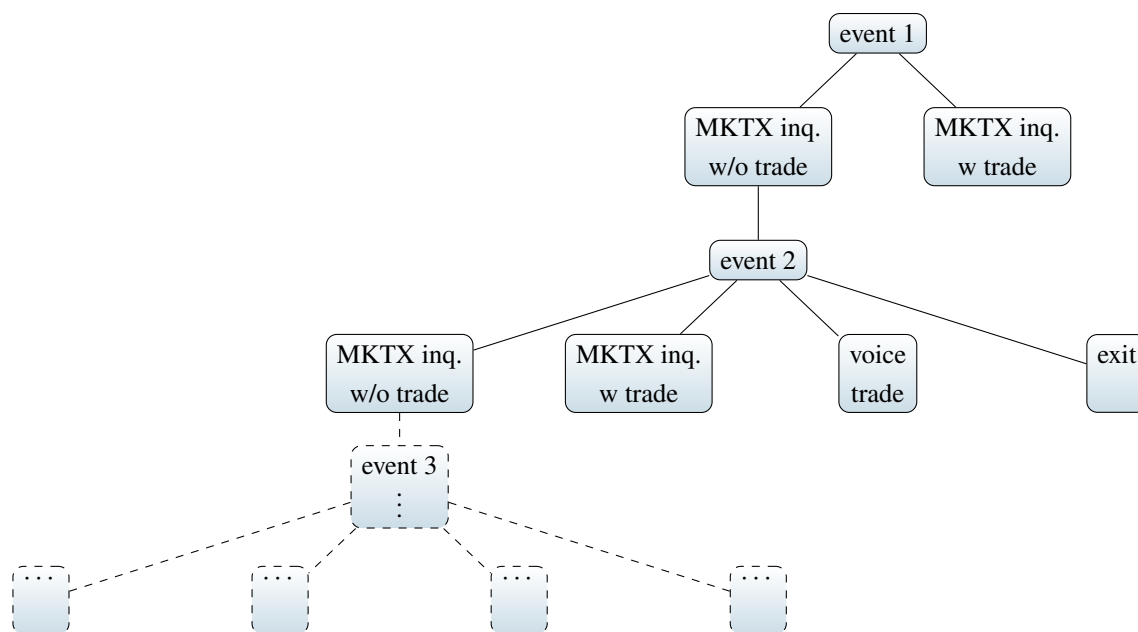


Figure 2. A child order event tree

A child order can be viewed as a sequence of events. Each element of the sequence is one of four possible events: a MKTX inquiry without trade, a MKTX inquiry with trade, a voice trade and, if the child order ends without a trade, an exit. By construction, the first event is always an inquiry on MKTX, either without or with trade.

how our sample differs from previous studies on the corporate bond market, and providing some preliminary evidence about the sequential trade process.

A child order is a sequence of events. Each element of the sequence is one of four possible events. First, the customer may make an inquiry on MKTX that fails to produce a trade. Second, the customer may make an inquiry on MKTX that results in a trade with one of the dealers that responded. Third, we may find that the customer traded the desired bond-quantity pair outside of the MKTX platform, via voice trade, within a short period of time. Fourth, the customer may give up on the trade and exit — either sending an inquiry for a different amount or abandoning the trade altogether. By construction, the first event in any child order that we observe is always an inquiry on MKTX, without or with trade. We can measure the time elapsed to the next event, unless it is an exit. Figure 2 illustrates a child order event tree.

Our focus on child order sets us apart from previous studies, such as [Hendershott and Madhavan \(2015\)](#) or [O’Hara and Zhou \(2021\)](#), who consider the universe of all inquiries and/or of all trades on MKTX. A simple way to illustrate the conceptual difference between child orders and inquiries

is to calculate trade probabilities. Since child orders include repeated inquiries on MKTX, they are naturally associated with a larger trading probability than inquiries alone. The difference is economically significant: in our sample, approximately 75% of child orders are traded on MKTX, while the trade probability at the inquiry level is 67%.

In Table 3 and Figure 3, we present the results of a logit regression. The dependent variable is whether trade occurs on MKTX and the independent variables are indicator functions for customer and trade characteristics. The “baseline category” is a round lot, investment-grade, buy request, for an above-median turnover and amount outstanding, and below-median time to maturity bond, from a well connected customer. Column (1) and (2) of Table 3 present estimates at the child order and inquiry levels, respectively. The unit of the coefficients is log odds ratio of trade. For example, the intercept in the column (1) shows that the odds ratio of trade for the baseline category is $\exp(2.462)$, leading to the probability of trade of $\exp(2.462)/(1 + \exp(2.462)) = 92\%$. In other words, at the child order level, the probability of trade for the baseline category is 92%. At the inquiry level in column (2), the odds ratio is smaller by about 18 percentage points, corresponding to a trade probability of 88%. So, one can see that child orders are executed with higher probability than inquiries. In Figure 3, the blue bars represent inquiry-level trade probabilities, and the combined blue and grey bars represent child-order trade probabilities.

[Table 3 about here.]

The estimates for covariates in Table 3 are interesting as well. For example, sequential trade matters a great deal for the “least connected” customers, defined as those who elicit a relatively low number of responses from dealers. The trade probability at the inquiry level is about 45%, but it is about 55% at the child order level. One can also see that the probability of trade fell at both inquiry and child order levels during the COVID-19 crisis in March 2020, but that the fall was much less dramatic at the child order level: at the inquiry level, the trade probability falls to about 75%, but at the child order level it falls much less, to 83%. This suggests another way sequential trade matters: during stressful events, it is harder for customers to obtain good quotes on MKTX, but investors could compensate for it by waiting.




Figure 3. Estimated trade probability on MKTX at inquiry and child order levels

This figure compares the estimated trade probability on MKTX using logit regression estimates from Table 3. The blue bars present trade probabilities at the inquiry level. The gray bars shows the extra trade probability for a child order, taking into account the option to make repeat inquiries. The top panel shows trade probabilities for different size categories. The bottom panel presents trade probabilities for non-size categories. Indicators for size and non-size categories are defined in Table 8. The baseline category is an odd-lot purchase of an investment-grade bond, with high turnover, low time-to-maturity, and high amount outstanding, during normal times, for a connected investor.

One may wonder whether our sample significantly differs from its inquiry- or trade-level counterparts in other dimensions as well.¹³ Table 4 shows that this is not the case: child-order and inquiry-level summary statistics are broadly the same for trade direction and size and bond characteristics. As in previous studies, we find that trade sizes on MKTX are smaller and bond credit risk is lower than in the market at large. To measure the inter-arrival time of trading opportunities,

¹³For example, suppose that high-yield bonds trade after twice as many inquiries as investment-grade bonds. Then we would find that the number of high-yield inquiries is twice that of high-yield child orders.

we will need to restrict the sample further to child orders with at least one failed inquiry. Column (3) of Table 4 shows that the summary statistics remain similar, though the sample is now more selected towards high-yield bonds and inquiries of larger size since both are less likely to trade at the first MKTX inquiry.

[Table 4 about here.]

Table 5 offers some descriptive statistics about child orders. The first row shows that the probability that a child order does not trade at the first inquiry is nontrivial, about 0.28. The following rows provide the frequency distribution over the next event in the child order, conditional on the number of failed inquiries to date. For example, the second row shows that if the first inquiry fails, the probability that the following event is a failed inquiry on MKTX is 0.17, the probability that there is a successful inquiry on MKTX is 0.09, the probability that there is a voice trade is 0.24, and the probability of an exit is 0.49. We will argue later that these probabilities are not obvious to interpret because of competing risk and selection biases. Notwithstanding these issues, there are a few takeaways. First, trade is sequential: the probability of failing an inquiry is nontrivial, and customers often submit repeat inquiries. Second, trade is non-exclusive: if the first inquiry fails, the child order may eventually trade on MKTX or voice.¹⁴ The third takeaway is that exit is nontrivial: the probability that a child order ends without trade is large. Fourth, the summary statistics in Table 5 show that the frequency distribution over the four events depends on the number of failed inquiries – a form of duration dependence.

[Table 5 about here.]

Table 6 presents inter-arrival times between events in child orders. For example, after one failed inquiry, the average time to the next traded inquiry on MKTX is 0.65 business days. However, as we argue below, this estimate is clearly biased downwards, since observing this event requires that none of the other events occur first.

¹⁴While the probability of a voice trade is larger than that of an MKTX trade, the ratio is not as large as the relative volume of voice to MKTX volume. This suggests that, although trade is nonexclusive, the customers in our sample are using MKTX more intensely than the general population.

[Table 6 about here.]

3 The sequential trade process: theory and evidence

3.1 A McCall (1970) model of a child order

In this section, we formulate and solve a sequential trade model of a child order in the style of [McCall \(1970\)](#), which was first applied to financial markets by [Garbade and Silber \(1976\)](#). This theoretical detour serves two purposes. First, it is a simple and natural theoretical framework to interpret our child order data; in particular, it helps clarify competing risk and selection biases in child order statistics and motivates the statistical models we estimate later. Second, since the [McCall \(1970\)](#) model is the workhorse, partial equilibrium model of sequential search, it constitutes a key building block in virtually all search-based models of OTC markets. Interpreting our empirical evidence through the lens of this sequential search model offers guidance for the quantitative values of key parameters—allowing us to infer, e.g., the time it takes to trade—and highlights which dimensions of the model fit the data well, and which dimensions must be enriched in order to match certain features of the data.

The model. Time is indexed by $t \in [0, \infty)$. We consider a child order to sell a perpetual par bond, that is, a perpetuity with a coupon rate that is equal to the interest rate, r . We assume that the seller is risk-neutral with discount rate r and is distressed, in that she values the bond below its par value of 1. Specifically, when she holds the bond, she derives a flow utility $r - c$, for some distress cost $c > 0$. The seller recovers from distress with intensity γ . Upon recovering, we assume that the seller's continuation value is equal to the par value of the bond, she stops searching, and exits the market.¹⁵ We focus here on a customer looking to sell, for simplicity, but the analysis of a purchase is symmetric.

Consistent with the child order tree of [Figure 2](#), we take $t = 0$ to represent the time at which

¹⁵We discuss alternative assumptions after [Proposition 1](#).

the seller makes her first inquiry on the electronic market.¹⁶ If the first inquiry is unsuccessful, the seller makes inquiries on the electronic or the voice market with Poisson intensities λ_e and λ_v , respectively. At this level of abstraction, the arrival rate of trading opportunities could represent frictions that derive from either side of the market. For example, the source of the friction could be that the customer simply can't find a dealer willing to buy the asset at an acceptable price, as in the literature following [Duffie, Gârleanu, and Pedersen \(2005\)](#). Alternatively, it could be that the customer is busy with other tasks and not actively trading in the market at all times, as in [Biais and Weill \(2009\)](#) and [Biais, Hombert, and Weill \(2014\)](#).

After an inquiry in the electronic market, the seller receives $j \in \{0, 1, 2, \dots\}$ offers with probability q_j . We represent an offer as a bid $1 - m$, where m is the markdown over the bond par value of 1. We assume further that each offered markdown is drawn independently according to the cumulative distribution function (CDF) $G_e(m)$. Correspondingly, when she makes an inquiry in the voice market, the seller receives just one offer, drawn according to the CDF $G_v(m)$. For simplicity we assume that, for both distributions, the lower bound of the support is 0. As will be clear below, the optimal trading strategy of the seller depends on two “sufficient statistics.” First the *total* Poisson intensity of inquiries,

$$\lambda = \lambda_e + \lambda_v,$$

and, second, the CDF over the *lowest* markdown, conditional on an inquiry,

$$F(m) = \frac{\lambda_e}{\lambda_e + \lambda_v} \sum_{j=0}^{\infty} q_j [1 - (1 - G_e(m))^j] + \frac{\lambda_v}{\lambda_e + \lambda_v} G_v(m).$$

The first term in this equation is the probability of making an inquiry on the electronic market, multiplied by the probability that the smallest markdown among j offers is less than m . The second term has the same interpretation, but for the voice market.

Given this notation, the Hamilton Jacobi Bellman (HJB) equation for the seller's value at any

¹⁶Alternatively, one may interpret $t = 0$ as the time at which the seller becomes distressed.

time $t > 0$ is

$$rV = r - c + \lambda \int \max\{1 - m - V, 0\} dF(m) + \gamma(1 - V). \quad (1)$$

The first term on the right-hand side is the flow value of holding the asset, i.e., the coupon r net of the distress cost c . The second term is the option value of search: the seller makes an inquiry with intensity λ , her best offer is distributed according to $F(m)$, and she accepts if the price $1 - m$ is larger than the value of continuing search, V . The third and last term is the expected flow utility if the seller recovers/exits. As is standard, the HJB shows that the optimal trading strategy of the seller is entirely characterized by the reservation markdown

$$m^* \equiv 1 - V. \quad (2)$$

That is, when she makes an inquiry, the seller trades if and only if the lowest markdown she receives is less than m^* . To obtain an equation for m^* , we substitute (2) in the HJB and obtain, after integration by parts, our version of [McCall's](#) celebrated equation, summarized in the following proposition.

Proposition 1 *The reservation markdown of a seller is the unique solution to*

$$m^* = \frac{c}{r + \gamma} - \frac{\lambda}{r + \gamma} \int_0^{m^*} F(m) dm. \quad (3)$$

The reservation markdown m^ increases with the distress cost c , decreases with the interest rate, r , decreases with the exit rate, γ , decreases with the inquiry intensity, λ , and increases in response to a first-order stochastic dominance shift in the distribution of the best markdown, $F(m)$.*

The first term in Equation (3), $c/(r + \gamma)$, is the expected present value of the seller's distress cost. It represents the monopsony markdown: the maximum markdown a seller would be willing to accept if she received just one take-it-or-leave-it offer by a dealer, and no offer forever after. The

optimal reservation markdown is less than the monopsony markdown because of the option value of searching for another offer.

The comparative statics for the reservation markdown are similar to those obtained in the classical job-search setting, with the exception of the effect of varying $r + \gamma$. The reason is that, in our setting, increasing $r + \gamma$ impacts the seller's problem in two ways. First, as in job-search models, increasing $r + \gamma$ reduces the option value of search which, all else equal, increases the reservation markdown. Second, and new to this setting, it decreases the present value of the seller's distress costs, which decreases the reservation markdown. The second effect, it turns out, always dominates in our setting.¹⁷

Next, we use this simple model as an aid to interpret our child order data. Recall the child order tree of Figure 2, where a child order is viewed as a sequence of events. According to the model, there is a new event in the child order tree with intensity

$$\lambda_e + \lambda_v G_v(m^*) + \gamma.$$

Conditional on an arrival, the new event is drawn independently from the arrival time according to the following distribution. The new event is an inquiry without trade on the electronic market with probability

$$\pi_1 = \frac{\lambda_e \sum_{j=0}^{\infty} q_j (1 - G_e(m^*))^j}{\lambda_e + \lambda_v G_v(m^*) + \gamma},$$

¹⁷For now, we have assumed that the seller exits the market when she recovers from distress: as shown in the HJB equation (1), her continuation value is set to the par value of the bond (1) upon recovery. But one may consider other plausible assumptions: for example, an exit in our data could occur because the seller goes back to the market with a different inquiry, e.g., for another quantity or a closely substitutable bond. In the analysis of the [McCall \(1970\)](#) model above, this amounts to changing the continuation value. Assume, for example, that when an exit occurs, the seller submits another child order for almost the same quantity or a nearly identical bond. Then, in the HJB equation, the continuation value is V instead of 1, and the reservation markdown equation is almost the same: the appropriate discount rate is now r instead of $r + \gamma$.

it is an inquiry with trade on the electronic market with probability

$$\pi_2 = \frac{\lambda_e \sum_{j=0}^{\infty} q_j [1 - (1 - G_e(m^*))^j]}{\lambda_e + \lambda_v G_v(m^*) + \gamma},$$

it is a trade on the voice market with probability

$$\pi_3 = \frac{\lambda_v G_v(m^*)}{\lambda_e + \lambda_v G_v(m^*) + \gamma},$$

and it is an exit with probability $\pi_4 = 1 - \pi_1 - \pi_2 - \pi_3$.

The formulae above illustrate two sources of bias that make interpreting child order statistics difficult. We discuss these two sources of bias below.

Competing risk bias. First, since the event type is drawn independently from the event arrival time, it follows that the *observed* expected arrival time of any of the four events is given by

$$\bar{\tau} = \frac{1}{\lambda_e + \lambda_v G_v(m^*) + \gamma}.$$

Notice that this observed expected arrival time is *lower* than the actual arrival time of the event. For example, the actual arrival time of a voice trade is $1/(\lambda_v G(m^*))$. This is a classical survivor bias induced by competing risk (e.g., [Flinn and Heckman, 1982](#); [Katz and Meyer, 1990](#); [Honoré and Lleras-Muney, 2006](#)) created by the arrival of other events. Imagine for example, that sellers exit the market very fast. Then the only trades on the voice market we would observe are those that occur sufficiently quickly, before an exit shock.

The formulae above show that there is a simple way to correct for this survivor bias. For example, the true expected time to trade on voice is equal to the ratio $\bar{\tau}/\pi_3$. As we will show below, this correction can be made more generally using a Maximum Likelihood approach, conditional on observable child-order characteristics.

Selection bias. In the data, we can control for several observed characteristics of child orders, such as bond type, trade size, and a natural measure of customer connectedness. But there may be other characteristics that are difficult to control for based on observables, including the distressed cost of a seller, c ; her inquiry intensities, λ_e or λ_v ; her ability to elicit responses from dealers, $\{q_j\}$; or her exit intensity, γ . Such unobserved characteristics create classical selection issues that could explain the apparent dependence of event probabilities on the number of failed inquiries, shown in Table 5.

To fix ideas formally, suppose that heterogeneity in child orders can be summarized by a one-dimensional type variable $x \in [x, \bar{x}]$. This specification allows for heterogeneity in all structural variables ($\lambda_e, \lambda_v, \gamma, c, r, G_e, G_v$, and so on), provided that there is a fixed relationship between the type x and each structural variable in the cross-section of child orders. Then, the measure of type- x child orders with $n \geq 1$ failed inquiries, $\mu(x | n)$, satisfies the following inflow-outflow equation:

$$\begin{aligned} \lambda_e(x) \left[\sum_j q_j(x) [1 - G_e(m^*(x) | x)]^j \right] d\mu(x | n-1) \\ = [\lambda_e(x) + \lambda_v(x)G_v(m^*(x) | x) + \gamma(x)] d\mu(x | n). \end{aligned}$$

The left-hand side is the inflow generated by child orders that make unsuccessful inquiries on the trading platform. Similarly, the right-hand side is the outflow generated by child orders that make inquiries on the trading platform, trade on the voice market, or exit. Taken together, these inflow-outflow equations imply that

$$d\mu(x | n) = \pi_1(x)^n d\mu(x | 0), \quad \text{where} \quad \pi_1(x) \equiv \frac{\lambda_e(x) \left(\sum_j q_j [1 - G_e(m^*(x) | x)]^j \right)}{\lambda_e(x) + \lambda_v(x)G_v(m^*(x) | x) + \gamma(x)}. \quad (4)$$

According to (4), the measure of type- x child orders with n failed inquiries declines geometrically with n . As discussed above, the geometric coefficient, $\pi_1(x)$, is simply the probability that a type- x inquiry on the electronic trading platform fails to trade (the left-most branch of event 2 in the child-order tree of Figure 2).

Next we show that the direction of the selection bias depends on the geometric coefficient, $\pi_1(x)$. Namely, let

$$dH(x | n) = \frac{d\mu(x | n)}{\int_{\underline{x}}^{\bar{x}} d\mu(y | n)},$$

denote probability distribution of x across child orders conditional on n failed inquiries. The following Lemma reports a key property of this distribution.

Lemma 1 *If $\pi_1(x)$ is an increasing (decreasing) function, then $H(x | n)$ first-order stochastically dominates (is first-order stochastically dominated by) $H(x | n - 1)$.*

Lemma 1 shows that as the number of failed inquiries, n , increases, the sample of child order becomes more selected towards those investors who, in their child order tree, fail inquiries on the trading platform with higher probability. As a result, if x is unobservable to the econometrician, any outcome variable which is monotonically related to x will appear to be monotonically related to the number of failed inquiry.

For example, suppose child orders differ in terms of the customer's distress cost, c , but are otherwise identical. Then $\pi_1(c)$ is decreasing in c since more distressed sellers have a higher reservation markdown, m^* . As a result, as the number of failed inquiries increases, the sample gets more and more selected towards less distressed customers. It follows that we should observe two key outcome variables, the trading probability and the transaction markdown, decline with the number of failed inquiries n .

3.2 Evidence about time to trade

We propose below a statistical framework to measure the time it takes customers to trade after their first inquiry on MKTX, correcting for the competing risk bias discussed above, and controlling for observable trade characteristics.

Maximum Likelihood Estimation. Our unit of observation i is an event node in the child order tree of Figure 2: specifically, the type and time of the event that follows an unsuccessful inquiry on MKTX. We index the $K = 4$ possible events by $k \in \{1, \dots, K\}$. Event $k = 1$ is an inquiry on MKTX without trade, $k = 2$ is an inquiry on MKTX with trade, $k = 3$ is a voice trade, and $k = 4$ is an exit. We assume further that these events arrive at independent exponential times with intensity $\lambda(\theta'_k x_i) = \exp(\theta'_k x_i)$, where x_i is a vector of covariates for that child order. These covariates include trade size, bond characteristics, customers' characteristics, *and* the number of failed inquiries on MKTX; the latter is particularly important, in that it allows us to identify potential duration dependence.

Given this statistical framework, conditional on x_i , the event k occurs at time $\tau_i = t$ with probability density

$$\mathbb{P}(\tau_i = t, \omega_i = k \mid x_i) = \lambda(\theta'_k x_i) e^{-\sum_{\ell} \lambda(\theta'_\ell x_i) t}.$$

This formula is the product of the probability that event k occurs at time t , $\lambda(\theta'_k x_i) e^{-\lambda(\theta'_k x_i) t}$, and the probability that all other events, $\ell \neq k$, occur *after* time t , $e^{-\sum_{\ell \neq k} \lambda(\theta'_\ell x_i) t}$. This is the sense in which there are “competing risks”: the probability density accounts for the fact that we observe event k only if the other events $\ell \neq k$ have not occurred before. Aggregating across events and the number of inquiries, the likelihood function is, evidently:

$$\prod_{i=1}^n \left(\sum_k \mathbb{I}_{\{\omega_i=k\}} \lambda(\theta'_k x_i) e^{-\sum_{\ell} \lambda(\theta'_\ell x_i) \tau_i} \right).$$

Recall that we never observe the time of an exit in our data set; rather, we observe only whether or not an exit occurred. Therefore, integrating with respect to τ_i when $\omega_i = K$, we obtain the likelihood for our actual observations:

$$\prod_{i=1}^n \left(\sum_{k \neq K} \mathbb{I}_{\{\omega_i=k\}} \lambda(\theta'_k x_i) e^{-\sum_{\ell} \lambda(\theta'_\ell x_i) \tau_i} + \mathbb{I}_{\{\omega_i=K\}} \frac{\lambda_K(\theta'_K x_i)}{\sum_{\ell} \lambda(\theta'_\ell x_i)} \right).$$

Taking logs, after a few lines of algebra, we obtain that the log-likelihood is $\sum_i L(\omega_i, \tau_i, x_i, \theta)$, where:

$$L_i(\omega_i, \tau_i, x_i, \theta) = \sum_k \mathbb{I}_{\{\omega_i=k\}} \theta'_k x_i - \mathbb{I}_{\{\omega_i \neq K\}} \left(\sum_{\ell} \exp(\theta'_{\ell} x_i) \right) \tau_i - \mathbb{I}_{\{\omega_i=K\}} \log \left(\sum_{\ell} \exp(\theta'_{\ell} x_i) \right).$$

We first gain some qualitative and quantitative intuition by deriving the *unconditional* Maximum Likelihood Estimator (MLE), i.e., the special case in which the only control is a constant.

Lemma 2 *Let $\hat{\pi}_k$ denote the empirical frequency of event k and $\hat{\tau}$ the empirical average inter-arrival time of an event $k \neq K$. Then, the MLE of θ_k is $\hat{\theta}_k = \log(\hat{\pi}_k/\hat{\tau})$.*

This is the same estimate that we intuitively derived in the previous section, when discussing the competing risk bias. Indeed, after a failed inquiry, the expected arrival time of any event is $\bar{\tau} = 1/(\sum_{\ell} \lambda_{\ell})$, and the probability of event k is $\pi_k = \lambda_k/\sum_{\ell} \lambda_{\ell}$. This shows that $\lambda_k = \pi_k/\bar{\tau}$ and $\theta_k = \log(\lambda_k)$, which is the population counterpart of the estimator in Lemma 2. The estimation results are shown in Table 7.

[Table 7 about here.]

The results offer some guidance about the orders of magnitude of arrival times for different events. For example, the unconditional intensity of a voice trade is $e^{-3.40} = 0.0333$ per business hour, corresponding to an average time of $1/0.0333 = 29.96$ business hours, or about 3.3 business days (assuming 9 hours of trading per day). Importantly, the estimates clearly show that competing risk creates a significant bias in calculating time to trade: indeed, 3.3 business days is much larger than the observed average inter-arrival times shown in Table 6 above.

Next, we move to the *conditional* MLE, with controls for trade characteristics (coefficients shown in Table 8) and for the number of failed inquiries in the child order to date (coefficients shown in Table 9). All controls are dummies. The baseline category is an odd-lot purchase of an investment-grade bond, with high turnover, during normal times, for a connected investor, after one failed inquiry. There is no closed form solution for the estimators. However, since the likelihood

function is concave in the vector of coefficients $\theta = (\theta_k)_{1 \leq k \leq K}$, it can be maximized reliably using existing optimization packages.

Table 8 shows the manner in which the intensities of each event, $\lambda(\theta'_k x)$, vary with trade characteristics. The intensities for the baseline category are obtained by taking the exponential of the intercept. The marginal effect of other trade characteristics is given by the exponential of their respective coefficient. In particular, when the coefficient is sufficiently small, it approximates the marginal effect in percentage term: e.g., from the fourth row in column (2) of Table 8, the intensity of trade with MKTX for a bond rated Ca to C is approximately $-(e^{-0.4} - 1) \simeq 33\%$ lower than for an investment-grade bond.

The estimates in Table 8 demonstrate that intensities vary significantly with trade characteristics. Consider, for example, trade size. We observe that the intensity of trade with MKTX for micro size trades (with size $< \$100,000$) is larger than for odd lot trades (our baseline category with size between $\$100,000$ and $\$1$ million). The intensity for odd lots is larger than for round lots (with size between $\$1$ and $\$5$ million), which is larger than for block trades (with size larger than $\$5$ million). Interestingly the intensity of voice trade is not monotonic in trade size: for example, block trades trade faster on voice. Bonds with low turnover, and high-yield bonds, also have lower trading intensity, both on MKTX and the voice market. Interestingly, sales and purchases are asymmetric: customers trade faster when they sell, on average, than when they buy.

The last rows of Table 8 show the impact of customer connectedness on MKTX. To derive a measure of customer connectedness, we first regress the average number of dealer responses elicited by a particular customer on several control variables, including the customer's average inquiry size, the fraction of his requests that were sell vs. buy, and the fraction of requests that were for investment-grade vs high-yield bonds. We then rank customers into deciles based on residuals of this regression. This measure aims to proxy for customers' existing relationships with dealers or other unobserved characteristics of connected clients. We find that this measure of connectedness creates significant differences in trading intensity on MKTX. This finding is intuitive because a more connected customer receives more offers on average and so is more likely to obtain one that

falls below her reservation markdown. Finally, the fifth row of Table 8 reveals that the COVID-19 crisis (identified by inquiries submitted in March 2020) had a significant negative impact on the trading intensity.

[Table 8 about here.]

Table 9 shows that, after controlling for trade characteristics, the number of failed inquiries retains predictive power for the intensity of each event. The intensity of an inquiry on MKTX that doesn't result in trade increases with the number of failed inquiries. In contrast, the intensity of successful inquiries—i.e., inquiries on either MKTX or via voice that result in trade—decreases as the number of failed inquiries increases. Through the lens of the McCall (1970) model outlined in the previous section, this evidence suggests a role for unobserved child order characteristics, such as heterogeneity in distress cost or the arrival rate of trading opportunities.

[Table 9 about here.]

Time to trade. We define time to trade as the expected time a customer takes to trade, either on MKTX or on voice, if she is not subject to exit shocks. That is, we study a hypothetical world in which the investor never exits in the child order tree—say, because she continues to search when she receives an exit shock.

If the intensities did not depend on the number of failed inquiries, calculating time to trade would be simple. For example, from the intercepts in columns (2) and (3) in Table 8 or 9, the time to trade for our baseline category would be $1/(e^{-3.65} + e^{-3.38}) \simeq 16.65$ business hours, or 1.8 business days. However, the dependence of intensities on the number of failed inquiries requires us to modify this simple formula.

Formally, consider a child order after n failed inquiries. With a slight abuse of notation, let x_n denote the corresponding vector of covariates, where n stands for the number of failed inquiries to

date. Then, the expected time to trade satisfies the following recursive formula:

$$\begin{aligned}
T(x_n) &= \mathbb{E} [\tau' \mid x_n] + \mathbb{P} [\omega' = 1 \mid x_n] \times T(x_{n+1}) \\
&\quad + \mathbb{P} [\omega' = 2 \mid x_n] \times 0 \\
&\quad + \mathbb{P} [\omega' = 3 \mid x_n] \times 0 \\
&\quad + \mathbb{P} [\omega' = 4 \mid x_n] \times T(x_n).
\end{aligned}$$

The first term is the expected time to the next event. The other terms add up to the expected continuation time to trade after the next event. Specifically, if the next event is an unsuccessful inquiry on MKTX, $\omega' = 1$, then there is one additional failed inquiry and the continuation time to trade is $T(x_{n+1})$. If the next event is $\omega' = 2$ or $\omega' = 3$, then trade occurs so the continuation time to trade is zero. The last line corrects the bias induced by the competing risk of exit: specifically, if the next event is an exit ($\omega' = 4$), we assume that the investor continues to search for a trade instead of exiting, so the continuation time to trade is $T(x_n)$.

Bringing the last term from the right-hand to the left-hand side, and using the exponential formula for expected inter-arrival time and event probability, we obtain the following recursion:

$$T(x_n) = \frac{1}{\lambda(\theta'_1 x_n) + \lambda(\theta'_2 x_n) + \lambda(\theta'_3 x_n)} + \frac{\lambda(\theta'_1 x_n)}{\lambda(\theta'_1 x_n) + \lambda(\theta'_2 x_n) + \lambda(\theta'_3 x_n)} T(x_{n+1}). \quad (5)$$

We can use this formula to calculate the time to trade. Moreover, differentiating (5) with respect to x , we obtain a corresponding recursive formula for the gradient of time to trade, which allows us to apply the Delta method and obtain standard errors for the time to trade estimates. We illustrate our results in a sequence of figures, where we plot the expected time to trade, conditional on the number of failed inquiries and specific trade characteristics using estimates from the MLE. We represent the 95% confidence intervals by shaded areas surrounding the conditional expectation.

Figure 4 shows that, for our baseline category, the time to trade increases from about two trading days after one failed inquiry to nearly four trading days after ten failed inquiries. High-yield, old,




Figure 4. Estimated conditional time to trade from the MLE: observed trade characteristics

This figure plots the estimated time to trade from Equation (5), conditional on the number of failed inquiries and on observed trade characteristics except trade size and customer connectedness categories. “Sell” takes the value of 1 for a sale request, and zero otherwise; “COVID” takes the value of 1 if the RFQ is submitted in March 2020, and zero otherwise; “old age” takes the value of 1 if the bond’s age is above the 75th percentile of the distribution, and zero otherwise; “low turnover” takes the value of 1 if the bond’s quarterly turnover is below median, and zero otherwise; “high TTM” takes the value of 1 if the bond’s time to maturity is above the sample median, and zero otherwise; “low amt out” takes the value of 1 if the bond’s amount outstanding is below the sample median, and zero otherwise. The baseline category is an odd-lot purchase of an investment-grade bond, with high turnover, during normal times, for a connected investor, after one failed inquiry.

and low turnover bonds have a longer time to trade, though the difference is small relative to other covariates.

In Figure 5, we study the impact of trade size on time to trade. We observe that smaller trades are faster on MKTX. For example, after one failed inquiry, it takes 1.5 days to trade a micro-size bond, while the time it takes to trade a block-size inquiry is almost twice as long. This evidence is consistent with prior studies that show electronic trading is concentrated on smaller trades (e.g., [Hendershott and Madhavan, 2015](#); [O’Hara and Zhou, 2021](#)).

Figure 6 shows that less connected customers, classified as customers that receive fewer offers from dealers, trade much slower on MKTX. For example, in the baseline category, the most connected customers (in the tenth decile of connectedness) trade after approximately 2.5 days

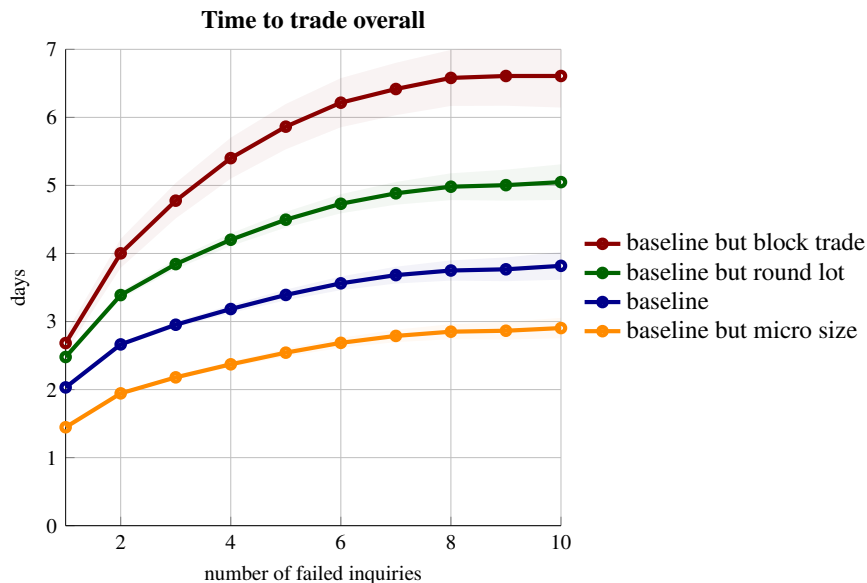


Figure 5. Estimated conditional time to trade from the MLE: impact of size

This figure plots the estimated time to trade from Equation (5), conditional on the number of failed inquiries and on trade size categories, and controlling for other observed trade characteristics. “Micro size” takes the value of 1 if the quantity of dealer response is below \$100,000, and zero otherwise; “odd lot” takes the value of 1 if the quantity of dealer response is between \$100,000 and \$1 million, and zero otherwise; “round lot” takes the value of 1 if the quantity of dealer response is between \$1 million and \$5 million, and zero otherwise; “block trade” takes the value of 1 if the quantity of dealer response exceeds \$5 million, and zero otherwise. The baseline category is an odd-lot purchase of an investment-grade bond, with high turnover, during normal times, for a connected investor, after one failed inquiry.

following two failed inquiries. For the least connected customers, in deciles 1 to 7, it takes almost 3.5 times more to trade on MKTX.

In Figure 7, we compare time to trade on MKTX to the one on voice for different trade size categories. The first takeaway is that, except for block trades, child orders trade much faster on MKTX than voice. This finding may be explained by the fact that customers initiate their first inquiries on MKTX and prefer to trade on the electronic platform, possibly for its execution quality rather than price discovery. Next, micro-size trades are faster than odd and round lots in both MKTX and the voice market, but block trades are much slower on MKTX. Again, this is not surprising, since, as mentioned above, smaller trades are more likely to be traded on electronic platforms.

It is important to keep in mind that our measurements are only descriptive. For example, the time to trade is presumably an endogenous outcome resulting from choices made by both sides of the market. The fact that customer purchases have longer time to trade could either indicate that

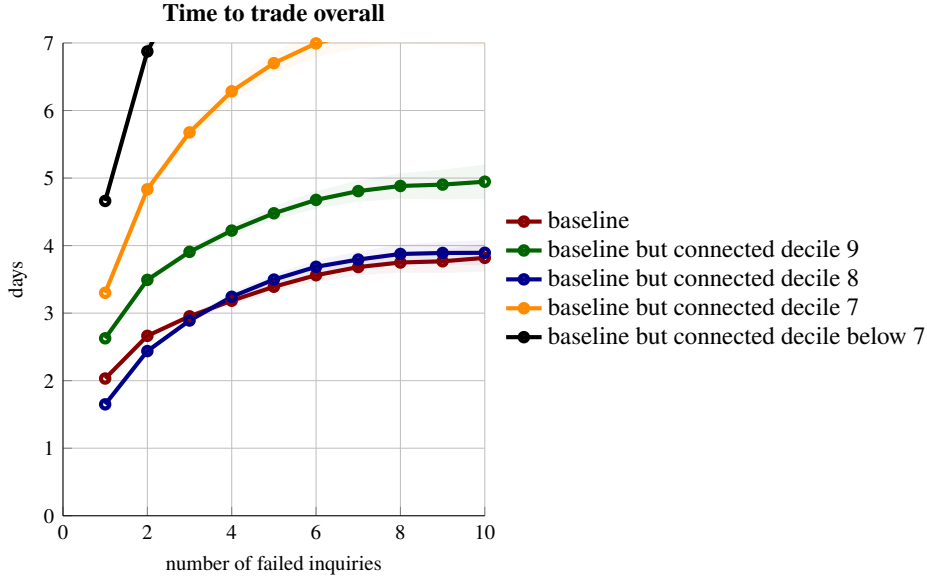


Figure 6. Estimated conditional time to trade from the MLE: impact of customer connectedness

This figure plots the estimated time to trade from Equation (5), conditional on the number of failed inquiries and on customer connectedness categories, and controlling for other observed trade characteristics. We first regress the average number of dealer responses elicited by a particular customer, on that customer’s average inquiry size and fractions of requests for sell trades and high-yield bonds. We then rank customers into deciles based on residuals of this regression. “Connected decile 9” is an indicator for the customer being in decile 9, and similarly for other “Connected” indicators. The baseline category is an odd-lot purchase of an investment-grade bond, with high turnover, during normal times, for a connected investor (in decile 10), after one failed inquiry.

it takes time for dealers to source or locate a bond, or it could indicate that buyers are less eager to trade than sellers, so more willing to continue searching for a better price. The fact that time to trade increased during COVID could indicate that dealers were reluctant to accumulate inventories and changed their bidding behavior as a result.

3.3 The dependence of outcomes on the number of failed inquiries

We have found above that, after controlling for observed trade characteristics, the intensities estimated via MLE continue to depend on the number of failed inquiries. Figure 8 shows that this is true for other outcome variables as well. Panel (a) plots the trading probability on MKTX, calculated based on the MLE:

$$\frac{\lambda(\theta'_2 x)}{\lambda(\theta'_1 x) + \lambda(\theta'_2 x)}.$$

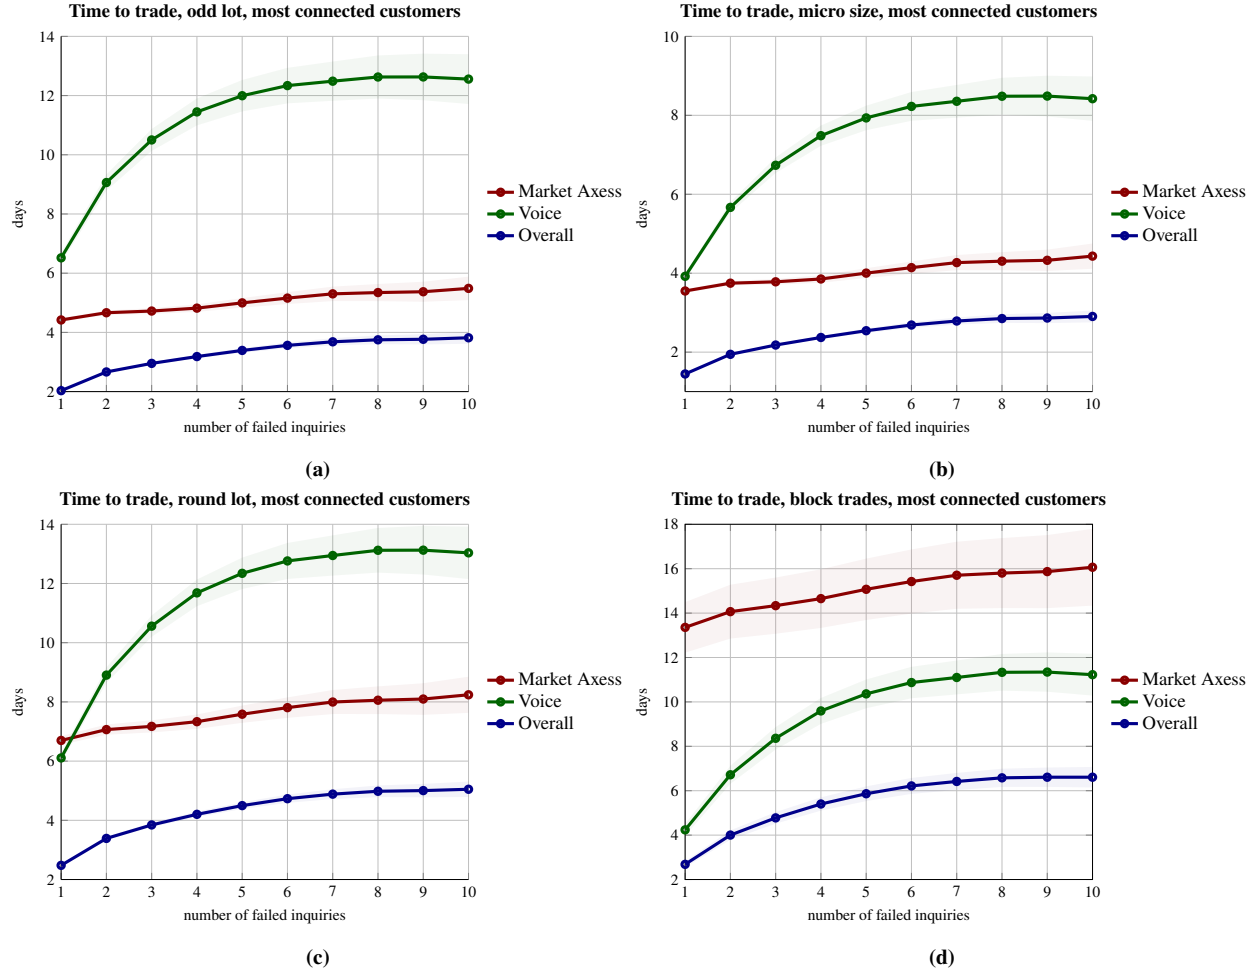


Figure 7. Estimated conditional time to trade from the MLE: MKTX vs. voice

This figure compares the estimated time to trade from Equation (5), conditional on the number of failed inquiries in MKTX vs. voice for the baseline (the top left panel), and different size categories. “Micro size” takes the value of 1 if the quantity of dealer response is below \$100,000, and zero otherwise; “odd lot” takes the value of 1 if the quantity of dealer response is between \$100,000 and \$1 million, and zero otherwise; “round lot” takes the value of 1 if the quantity of dealer response is between \$1 million and \$5 million, and zero otherwise; “block trade” takes the value of 1 if the quantity of dealer response exceeds \$5 million, and zero otherwise.

It shows that the trading probability goes down as the number of failed inquiries increases. Panel (b) shows the inquiry intensity with MKTX, calculated based on the MLE, $\lambda(\theta'_1 x) + \lambda(\theta'_2 x)$. Panel (c) plots the best markdown or spread. It reveals that the best markdown increases as the number of failed inquiries increases. Panel (d) shows the expected number of dealer responses, estimated by Poisson regression.¹⁸ It shows that, as the number of failed inquiries increases, the expected number of dealer responses falls.

¹⁸The associated regression results are presented in column (1) of Table 10.

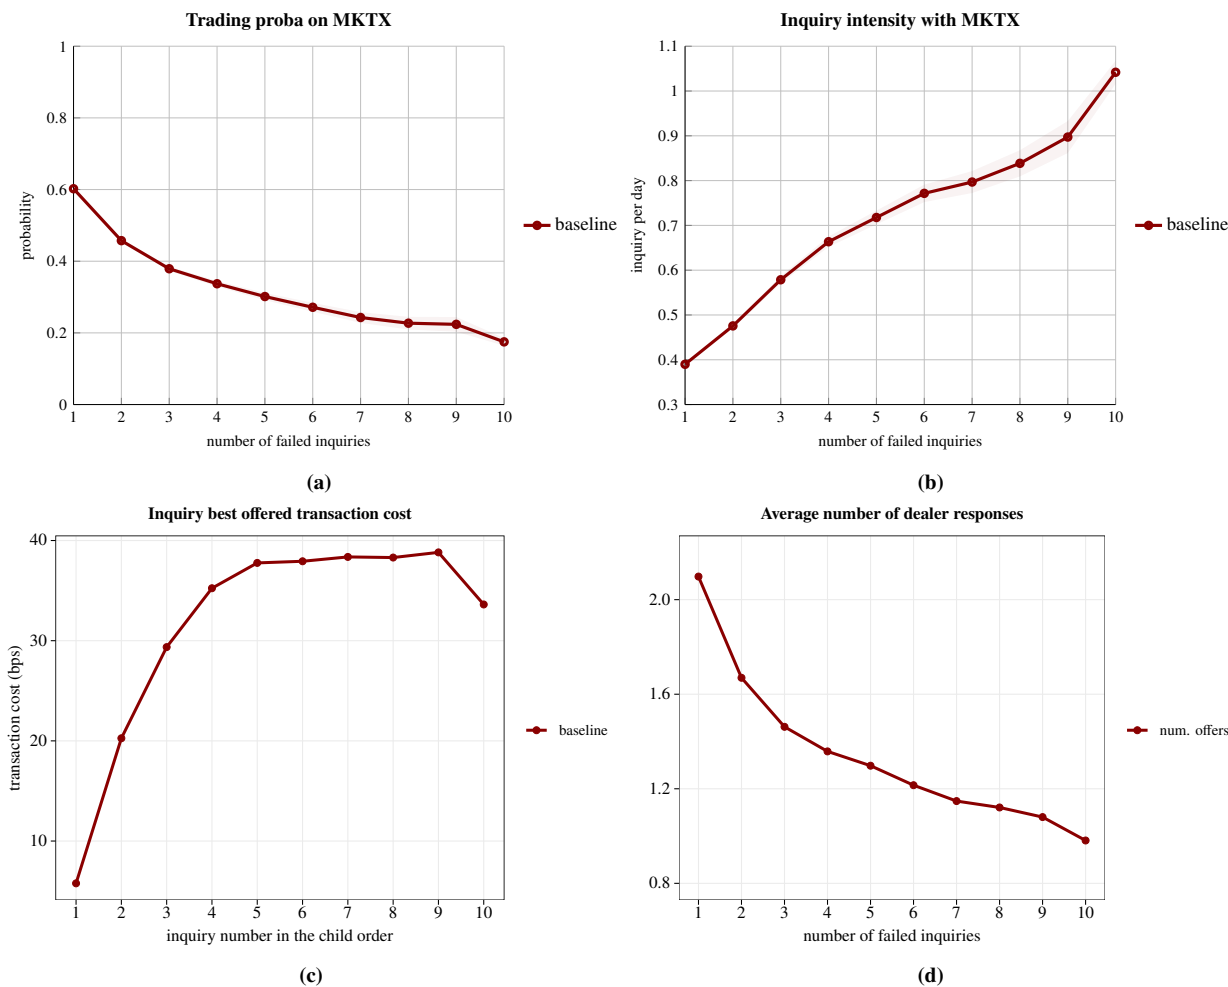


Figure 8. Unobserved characteristics

This figure plots, for the baseline category, the trading probability on MKTX in panel (a), the contact intensity with MKTX in panel (b), the markdown or best offered transaction cost for traded inquiries in panel (c), and the expected number of dealer responses in panel (d) as a function of the number of failed inquiries in child orders. The baseline category is an odd-lot purchase of an investment-grade bond, with high turnover, during normal times, for a connected investor, after one failed inquiry. As discussed in Section 2, we measure transaction cost as a markdown or markup relative to the benchmark provided by MKTX, called Composite+.

The dependence of outcome variables on the number of failed inquiries can be interpreted in two ways. The first explanation is that, as a customer fails more and more inquiries, the trading environment changes. Just to give one example, the increase in the spread for traded inquiries in panel (c) could be consistent with change in dealers’ bidding behavior because of information leakage about this customer, as noted by [Hendershott and Madhavan \(2015\)](#), or because dealers learn that the customer was not able to elicit competitive offers, as in the “ringing-phone curse” described in [Zhu \(2011\)](#).

The second explanation is instead that child orders are heterogeneous in characteristics unobserved by the econometrician. Indeed, according to Lemma 1, the distribution of child orders along such unobserved characteristics changes systematically with the number of failed inquiries. For example, suppose that child orders differ in distress costs, c . Then, since child orders with higher distress costs have larger reservation markdown, m^* , they trade with higher probability. This implies that, conditional on a larger number of failed inquiries, the sample becomes selected towards low distress cost child orders that trade with lower probability. Hence, heterogeneity in distress costs could explain why a child order time to trade appears to increase with the number of failed inquiries.

To tell these two hypotheses apart, we study the dependence of outcome variables on the number of failed inquiries in two ways: controlling for observed trade characteristics and controlling for child-order fixed effects. If unobserved child order characteristics explain the dependence of outcome variables on the number of failed inquiries, then the dependence should disappear after controlling for child order fixed effects. Indeed, when we control for child order fixed effects, we keep *all* child order characteristics fixed, whether they are observed or not.

[Table 10 about here.]

Table 10 shows the Poisson regression results when the outcome variable is the number of dealer responses. In column (1), we control for observed trade characteristics. We find that holding all observed trade characteristics constant, increasing the number of inquiries from 1 to 2, reduces the number of dealer responses by approximately 27% ($= 1 - e^{-0.311}$). Second, in column (2), we use child order fixed effects instead of trade characteristics. Now changing the number of inquiries has much more muted impact on the number of dealer responses: increasing the number of inquiries from 1 to 2, actually *increases* the number of dealer responses by 3.7% ($= 1 - e^{-0.0361}$). The results in Table 10 provide evidence in favor of the hypothesis that child orders differ in unobserved characteristics.

In Table 11, we do the same but for another variable: the spread (transaction cost) of traded inquiries. As discussed in Section 2, we measure transaction cost as a markdown or markup relative

to the CP+ benchmark provided by MKTX.¹⁹ The evidence in column (1) is consistent with panel (b) of Figure 8, showing that spreads rise as the number of inquiries increases. However, when we control for potentially unobserved characteristics with child order fixed effects, in column (2), we obtain a very different picture: the spreads of traded inquiries are much more stable as the number of inquiries within a child order changes and, if anything, go slightly in the opposite direction.

[Table 11 about here.]

These regression results suggest that unobserved characteristics are a likely explanation of the dependence of outcome variables on the number of failed inquiries. But they do not shed light on the nature of these unobserved characteristics. However, our model can help make some indirect inference about these. For example, we can easily reject heterogeneity in distress cost: while this would explain why the trading probability decreases in the number of failed inquiries, as in panel (a) of Figure 8, it would be hard to reconcile with the observation that the inquiry intensity increases.

Which specific unobserved characteristics could be qualitatively consistent with the evidence in the four panels of Figure 8 and with the evidence on time to trade presented above? Panel (b) suggests that, as the number of failed inquiries increase, the sample becomes more and more selected towards child orders with high inquiry intensity on MKTX. According to Lemma 1, this is consistent with a McCall (1970) model in which investors are heterogeneous in their inquiry intensity on the electronic market, λ_e , because, in the child-order tree, the probability of a failed inquiry

$$\pi_1 = \frac{\lambda_e \left(\sum_j q^j [1 - G_e(m^*)]^j \right)}{\lambda_e + \lambda_v G_v(m^*) + \gamma},$$

is increasing in λ_e , for two reasons. First, the probability of an inquiry on the electronic market, $\lambda_e / (\lambda_e + \lambda_v G_v(m^*) + \gamma)$, evidently increases with λ_e . Second, from Proposition 1, the reservation markdown, m^* , decreases with λ_e , since customers who make more frequent inquiries have a larger

¹⁹As an alternative measure, we also compute the trading cost measure in Hendershott and Madhavan (2015), which uses the last inter-dealer trade as the reference price for a given bond instead of CP+. Results remain qualitatively similar using this alternative transaction cost measure.

option value of continuing their search. Panel (a) is consistent with this heterogeneity in λ_e too, since the trading probability on the electronic market declines with the number of failed inquiries.

However, Figure 8 is inconsistent with the hypothesis that child orders are *only* heterogeneous in their inquiry intensity λ_e . Indeed, panel (c) shows that the best offered transaction cost in a given inquiry on MKTX increases with the number of failed inquiries. This suggests that child orders are heterogeneous in another dimension: the distribution of their best offer on the electronic market. Panel (d) suggests a precise reason why this distribution may differ across child orders: the average number of dealer responses declines with the number of failed inquiries. Taken together, panels (a) through (d) of Figure 8 suggest that, after controlling for observable characteristics, child orders differ in two dimensions: first in their inquiry intensity, and second in the number of responses they elicit from dealers.

To formally establish that these two dimensions of heterogeneity are qualitatively consistent with the empirical observations in this paper, consider the [McCall](#) model where, for simplicity, child orders trade only on MKTX. Assume the distribution of dealer responses is Poisson with parameter μ . Suppose there are two types of child orders: type A with high inquiry intensity and low number of responses, and type B with low inquiry intensity and high number of responses. Let the associated parameters for these two types be $\lambda_A > \lambda_B$ and $\mu_A < \mu_B$. Denote the associated distribution of best offers by $F(m | \mu_A)$ and $F(m | \mu_B)$, where

$$F(m | \mu) = \sum_{j \geq 0} e^{-\mu} \frac{\mu^j}{j!} [1 - (1 - G(m))^j] = 1 - e^{-\mu G(m)}.$$

Finally, let the associated reservation markdowns be m_A^* and m_B^* . Then we obtain the following Lemma.

Lemma 3 *Suppose that $\lambda_A > \lambda_B > \lambda_A F(m_A^* | \mu_A)$. Then, as $\mu_B \rightarrow \infty$:*

$$\pi_{1A} > \pi_{1B}$$

$$F(m_A^* | \mu_A) < F(m_B^* | \mu_B)$$

$$\mathbb{E}_{\mu_A} [m] > \mathbb{E}_{\mu_B} [m].$$

Taken together with Lemma 1, the first inequality in Lemma 3 implies that, conditional on a larger number of failed inquiries, the sample becomes more selected towards child orders with high inquiry intensity, but low number of dealer responses. This is consistent with both panel (b) and (d) of Figure 8. The second inequality in Lemma 3 makes it consistent with panel (a) of Figure 8, since child orders with high inquiry intensity but low number of dealer responses have lower trading probability. The third inequality makes it consistent with panel (c), since these child orders also trade at worse spread. Finally, the restriction that $\lambda_B > \lambda_A F(m_A^* | \mu_A)$ ensures that these child orders have longer time to trade.

4 Conclusion

In this paper, we use data from a leading electronic trading platform to provide new and direct empirical evidence about search frictions in the OTC market for corporate bonds. We start from the observation that when a customer's inquiry on the platform fails to trade, the same customer often returns to the market shortly after to make subsequent inquiries for the same quantity of the same bond. We argue that the resulting sequence of repeated inquiries sheds light on the customers' sequential search process. We estimate that, after a failed inquiry, it takes customers between two and three days to trade. We show that this time to trade depends systematically on trade characteristics and trading venue (electronic vs. voice). We provide evidence consistent with unobserved characteristics being a likely reason for the dependence of outcome variables on the number of prior failed attempts to trade. Overall, our estimates can serve as useful inputs

into future quantitative applications of search models while also providing guidance for future theoretical explorations of the micro-foundations of search frictions in OTC markets.

References

- Afonso, Gara, and Ricardo Lagos, 2015, Trade dynamics in the market for federal funds, *Econometrica* 83, 263–313.
- Bao, Jack, Jun Pan, and Jiang Wang, 2011, The illiquidity of corporate bonds, *Journal of Finance* 66, 911–946.
- Bessembinder, Hendrik, Stacey Jacobsen, William Maxwell, and Kumar Venkataraman, 2018, Capital commitment and illiquidity in corporate bonds, *Journal of Finance* 73, 1615–1661.
- Bessembinder, Hendrik, William Maxwell, and Kumar Venkataraman, 2006, Market transparency, liquidity externalities, and institutional trading costs in corporate bonds, *Journal of Financial Economics* 82, 251 – 288.
- Bessembinder, Hendrik, Chester Spatt, and Kumar Venkataraman, 2020, A survey of the microstructure of fixed-income markets, *Journal of Financial and Quantitative Analysis* 55, 1–45.
- Biais, Bruno, Johan Hombert, and Pierre-Olivier Weill, 2014, Equilibrium pricing and trading volume under preference uncertainty, *Review of Economic Studies* 81, 1401–1437.
- Biais, Bruno, and Pierre-Olivier Weill, 2009, Liquidity shocks and order book dynamics, Working paper, TSE and UCLA.
- Brancaccio, Giulia, and Karam Kang, 2021, Search frictions and product design in the municipal bond market, Working paper, NYU and CMU.
- Czech, Robert, and Gábor Pintér, 2020, Informed trading and the dynamics of client-dealer connections in corporate bond markets, Working paper, Bank of England.
- Dick-Nielsen, Jens, 2014, How to clean enhanced TRACE data, Working paper, CBS.

- Duffie, Darrell, Nicolae Gârleanu, and Lasse H. Pedersen, 2005, Over-the-counter markets, *Econometrica* 73, 1815–1847.
- Edwards, Amy K., Lawrence E. Harris, and Michael S. Piowar, 2007, Corporate bond market: Transaction costs and transparency, *Journal of Finance* 62, 1421–1451.
- Flinn, Christopher, and James Heckman, 1982, New methods for analyzing structural models of labor force dynamics, *Journal of Econometrics* 18, 115–168.
- Garbade, Kenneth D., and William L. Silber, 1976, Price dispersion in the government securities market, *Journal of Political Economy* 84, 721–740.
- Gavazza, Alessandro, 2016, An empirical equilibrium model of a decentralized asset market, *Econometrica* 84, 1755–1798.
- Goldstein, Michael A., Edith S. Hotchkiss, and Erik R. Sirri, 2007, Transparency and liquidity: A controlled experiment on corporate bonds, *Review of Financial Studies* 20, 235–273.
- Hendershott, Terrence, Dan Li, Dmitry Livdan, and Norman Schürhoff, 2020, True cost of immediacy, Working paper, UC Berkeley, SFI, and FRB.
- Hendershott, Terrence, Dmitry Livdan, and Norman Schürhoff, 2021, All-to-all liquidity in corporate bonds, Working paper, UC Berkeley and SFI.
- Hendershott, Terrence, and Ananth Madhavan, 2015, Click or call? auction versus search in the over-the-counter market, *Journal of Finance* 70, 419–447.
- Honoré, Bo H., and Adriana Lleras-Muney, 2006, Bounds in competing risks models and the war on cancer, *Econometrica* 74, 1675–1698.
- Kargar, Mahyar, Benjamin Lester, David Lindsay, Shuo Liu, Pierre-Olivier Weill, and Diego Zúñiga, 2021, Corporate bond liquidity during the COVID-19 crisis, *Review of Financial Studies* 34, 5352–5401.

- Katz, Lawrence F., and Bruce D. Meyer, 1990, Unemployment insurance, recall expectations, and unemployment outcomes, *Quarterly Journal of Economics* 105, 973–1002.
- Kiefer, Nicholas M, and George R Neumann, 1979, An empirical job-search model, with a test of the constant reservation-wage hypothesis, *Journal of Political Economy* 87, 89–107.
- Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1336.
- McCall, John J., 1970, Economics of information and job search, *Quarterly Journal of Economics* 84, 113–126.
- O’Hara, Maureen, and Alex Zhou, 2021, The electronic evolution of corporate bond dealers, *Journal of Financial Economics* 140.
- Pintér, Gábor, and Semih Üslü, 2021, Comparing search and intermediation frictions across markets, Working paper, Bank of England and Johns Hopkins University.
- Schultz, Paul, 2001, Corporate bond trading costs: A peek behind the curtain, *Journal of Finance* 56, 677–698.
- Üslü, Semih, 2019, Pricing and liquidity in decentralized asset markets, *Econometrica* 87, 2079–2140.
- Weill, Pierre-Olivier, 2020, The search theory of over-the-counter markets, *Annual Review of Economics* 12, 747–773.
- Wooldridge, Jeffrey M., 2010, *Econometric Analysis of Cross Section and Panel Data* (MIT Press).
- Zhu, Haoxiang, 2011, Finding a good price in opaque over-the-counter markets, *Review of Financial Studies* 25, 1255–1285.

Appendix

A Omitted proofs

A.1 Proof of Lemma 1

Using the definition of H , we obtain that

$$dH(x | n) = \frac{\pi_1(x) d\mu(x | n-1)}{\int_{\underline{x}}^{\bar{x}} \pi_1(y) d\mu(y | n-1)} = \frac{\pi_1(x) dH(x | n-1)}{\int_{\underline{x}}^{\bar{x}} \pi_1(y) dH(y | n-1)}$$

where the first equality follows from the recursion $d\mu(x | n) = \pi_1(x) d\mu(x | n-1)$, and the second equality follows from dividing both the numerator and the denominator by $\int_{\underline{x}}^{\bar{x}} d\mu(x | n-1)$. Therefore:

$$\begin{aligned} \text{sign}(H(x | n) - H(x | n-1)) &= \text{sign}\left(\frac{\int_{\underline{x}}^x \pi_1(y) dH(y | n-1)}{\int_{\underline{x}}^{\bar{x}} \pi_1(y) dH(y | n-1)} - \int_{\underline{x}}^x dH(y | n-1)\right) \\ &= \text{sign}\left(\int_{\underline{x}}^x \left[\pi_1(y) - \int_{\underline{x}}^{\bar{x}} \pi_1(z) dH(z | n-1)\right] dH(y | n-1)\right). \end{aligned}$$

Recall that $\pi_1(y)$ is strictly increasing. This implies that $\pi_1(y) - \int_{\underline{x}}^{\bar{x}} \pi_1(z) dH(z | n)$ is strictly increasing as well, strictly negative when $y = \underline{x}$, and strictly positive when $y = \bar{x}$. It follows that there is an x_0 such that $\pi_1(y) - \int_{\underline{x}}^{\bar{x}} \pi_1(z) dH(z | n-1) < 0$ for all $y < x_0$, and strictly positive for all $y > x_0$. Hence,

$$x \mapsto \int_0^x dH(y | n-1) \left[\pi_1(y) - \int_0^{\bar{x}} \pi_1(z) dH(z | n-1)\right]$$

is first decreasing and then increasing. Since this function is obviously equal to zero at the upper bound of its domain, $x = \bar{x}$, it follows that $H(x | n) \leq H(x | n-1)$, and we have established first-order stochastic dominance.

A.2 Proof of Lemma 3

Clearly, an increase in μ creates a first-order negative shift in $F(m | \mu)$ (i.e., markdown are lower). Hence the reservation markdown is decreasing in μ . Going to the limit in the reservation markdown equation (3), noting that $\lim_{\mu \rightarrow \infty} F(m | \mu) = 1$ for all $m > 0$, we obtain that

$$\lim_{\mu \rightarrow \infty} m^* = \frac{c}{r + \gamma + \lambda} > 0.$$

But $F(m | \mu) \rightarrow 1$ for $m > 0$, i.e., the trading probability becomes arbitrarily close to 1, and the expected markdown becomes arbitrarily close to zero. The result follows.

Tables

Table 1. Responses of a traded and an untraded inquiry

Panel (a) provides dealers' disclosed responses for a traded inquiry submitted on 08/15/2017 to buy \$300,000 of an 11-year, 3.824% investment-grade (USHG) bond issued on 01/17/2017 by Bank of America. The customer received 6 responses, all from dealers, whose anonymized IDs are provided in column (6). Response level (spread over Treasuries for USHG in MKTX) for each dealer response is reported in column (7). In column (10), the response status "Done" flags the response that the submitter accepted, the response status "Cover" flags the second best offer, and the response status "Missed" flags the rest of the responses that the submitter rejected. Panel (b) provides dealer disclosed responses for an untraded inquiry submitted on 08/17/2017 to buy \$490,000 of the same bond in panel (a). The customer received 9 responses, all from dealers, whose anonymized IDs and response levels are reported in columns (6) and (7), respectively. The response status "DNT" for this inquiry in column (9) indicates that the inquiry did not trade.

Panel (a): Responses to a traded inquiry on 08/15/2017								
Cust. ID (1)	Bond CUSIP (2)	Trade Side (3)	Submit Time (4)	Resp. ID (5)	Dealer ID (6)	Resp. Level (7)	Resp. Quant. (8)	Resp. Status (9)
127	06051GGF0	Buy	08:07:06	1	15420	126.37	300	Missed
127	06051GGF0	Buy	08:07:06	2	16323	129.70	300	Done
127	06051GGF0	Buy	08:07:06	3	11595	128.00	300	Missed
127	06051GGF0	Buy	08:07:06	4	16664	128.05	300	Missed
127	06051GGF0	Buy	08:07:06	5	10392	128.32	300	Missed
127	06051GGF0	Buy	08:07:06	6	12867	128.70	300	Cover

Panel (b): Responses to an untraded inquiry on 08/17/2017								
Cust. ID (1)	Bond CUSIP (2)	Trade Side (3)	Submit Time (4)	Resp. ID (5)	Dealer ID (6)	Resp. Level (7)	Resp. Quant. (8)	Resp. Status (9)
127	06051GGF0	Buy	09:56:49	1	15420	125.32	490	DNT
127	06051GGF0	Buy	09:56:49	2	11122	125.70	490	DNT
127	06051GGF0	Buy	09:56:49	3	16377	124.70	490	DNT
127	06051GGF0	Buy	09:56:49	4	12867	125.70	490	DNT
127	06051GGF0	Buy	09:56:49	5	16323	126.20	490	DNT
127	06051GGF0	Buy	09:56:49	6	16664	125.31	490	DNT
127	06051GGF0	Buy	09:56:49	7	10392	125.32	490	DNT
127	06051GGF0	Buy	09:56:49	8	11684	127.01	490	DNT
127	06051GGF0	Buy	09:56:49	9	13910	126.71	490	DNT

Table 2. Cluster of inquiries

This table provides details on the inquiries composing the cluster for an 11-year, 3.824% investment-grade bond issued on 01/17/2017 by Bank of America over a six-month period in 2017, depicted in Figure 1.

Inquiry ID (1)	Cust. ID (2)	Bond CUSIP (3)	Trade Side (4)	Submit Time (5)	Requested Quantity (6)	Inquiry Traded? (7)	Parent Order # (8)	Child Order # (9)
1	127	06051GGF0	Buy	08/15/2017 08:07:06	300	Yes	1	1
2	127	06051GGF0	Buy	08/17/2017 09:56:49	490	No	1	2
3	127	06051GGF0	Buy	08/17/2017 13:57:19	490	Yes	1	2
4	127	06051GGF0	Buy	08/18/2017 08:35:20	290	No	1	3
5	127	06051GGF0	Buy	08/21/2017 08:45:43	290	Yes	1	3
6	127	06051GGF0	Buy	08/23/2017 11:11:38	680	Yes	1	4

Table 3. Trade probability on MKTX: inquiry vs. child order level

This table presents logit regression results of whether trade occurs as the dependent variable and indicators for trade and customer characteristics as independent variables, defined in Table 8. Column (1) presents the regression at the child order level and the corresponding inquiry level estimates are presented in column (2). We rank customers into deciles according to the number of dealer responses they receive, after controlling for inquiry size, fraction of requests for sell trades and HY bonds. “Connected decile 9” is an indicator for the customer being in decile 9, and similarly for other “Connected” indicators. Heteroskedasticity-robust standard-errors are reported in parentheses.

Dependent Variables: Model:	child is traded (1)	inq is traded (2)
<i>Variables</i>		
(Intercept)	2.462*** (0.0029)	2.019*** (0.0025)
Micro size	0.4815*** (0.0021)	0.4187*** (0.0018)
Round lot	-0.3408*** (0.0035)	-0.1795*** (0.0032)
Block trade	-0.6760*** (0.0109)	-0.3941*** (0.0101)
Sell	0.6565*** (0.0020)	0.6514*** (0.0017)
HY	-0.4866*** (0.0024)	-0.4732*** (0.0021)
Covid	-0.8599*** (0.0053)	-0.9448*** (0.0046)
Old age	-0.1863*** (0.0021)	-0.1803*** (0.0018)
High time-to-maturity	-0.0659*** (0.0020)	-0.0227*** (0.0017)
Low turnover	-0.3229*** (0.0027)	-0.2827*** (0.0024)
Low amt outstanding	-0.5826*** (0.0020)	-0.6091*** (0.0017)
Connected decile < 7	-2.296*** (0.0028)	-2.186*** (0.0024)
Connected decile 7	-1.800*** (0.0030)	-1.681*** (0.0027)
Connected decile 8	-1.204*** (0.0030)	-1.159*** (0.0025)
Connected decile 9	-0.6484*** (0.0030)	-0.7128*** (0.0024)
<i>Fit statistics</i>		
Observations	8,680,700	9,441,617
Squared Correlation	0.16503	0.17491
Pseudo R ²	0.16729	0.15811
BIC	6,969,783.3	8,986,339.2

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

Table 4. Summary statistics

This table presents summary statistics for size, bond age and maturity, rating, and trade direction for all child orders (column 1), all inquiries (column 2), and child orders with at least one failed inquiry (column 3). “Sell” takes the value of 1 for a sale request, and zero otherwise; “HY” takes the value of 1 if the bond is high-yield, and zero otherwise; “Dealer-submitted” takes the value of 1 if the inquiry is submitted by a dealer, and zero otherwise.

	Child orders (all) (1)	Inquiries (all) (2)	Child orders (≥ 1 failed inq.) (3)
HY	0.17	0.18	0.26
Sell	0.52	0.51	0.42
Dealer-submitted	0.10	0.11	0.23
<i>Size</i>			
micro size (< \$100k)	0.49	0.48	0.37
odd lot (\$100k–1 million)	0.42	0.43	0.52
round lot (\$1–5 million)	0.09	0.08	0.10
block trade (> \$5 million)	0.01	0.01	0.01
<i>Bond age distribution</i>			
Average bond age	3.85	3.91	4.43
< 2 years	0.35	0.34	0.31
2–5 years	0.39	0.39	0.38
5–20 years	0.26	0.27	0.30
> 20 years	0.01	0.01	0.02
<i>Bond maturity distribution</i>			
Average maturity	12.43	12.53	13.71
< 2 years	0.002	0.002	0.002
2–5 years	0.07	0.07	0.06
5–20 years	0.73	0.73	0.68
> 20 years	0.20	0.20	0.25
Observations	9,861,143	11,020,815	2,774,478

Table 5. Child order event statistics

This table presents summary statistics about child order events. A child order can be viewed as a sequence of events, as depicted in Figure 2. Each element of the sequence is one of four possible events: an untraded inquiry on MKTX, a MKTX inquiry with trade, a voice trade, and, if the child order ends without a trade, an exit. By construction, the first event is always either an inquiry on MKTX, without or with trade. The first row shows the probability of a failed and successful inquiry on MKTX. The following rows provides the frequency distribution over the next event in the child order, conditional on the number of failed inquiries to date.

Event	Prob. MKTX inq. w/o trade (1)	Prob. MKTX inq. w trade (2)	Prob. voice trade (3)	Prob. exit (4)
First inquiry	0.28	0.72	N/A	N/A
After 1 failed inquiry	0.17	0.09	0.24	0.49
After 2 failed inquiries	0.34	0.10	0.17	0.39
After 3 failed inquiries	0.46	0.09	0.13	0.31
After 4 failed inquiries	0.55	0.08	0.10	0.26
After 5 failed inquiries	0.61	0.08	0.08	0.22
After 6 failed inquiries	0.66	0.07	0.07	0.20
After 7 failed inquiries	0.69	0.06	0.07	0.18
After 8 failed inquiries	0.73	0.06	0.06	0.16
After 9 failed inquiries	0.74	0.05	0.05	0.15
After 10 failed inquiries	0.77	0.05	0.05	0.14

Table 6. Child orders statistics: Inter-arrival times.

This table presents summary statistics about time between child order events (in business days). A child order can be viewed as a sequence of events, as depicted in Figure 2. Each element of the sequence is one of four possible events: an untraded inquiry on MKTX, a MKTX inquiry with trade, a voice trade, and, if the child order ends without a trade, an exit. Columns (1)–(3) present time, in business days, to an untraded inquiry on MKTX, a MKTX trade, and a trade on voice across child orders, conditional on the number of failed inquiries to date.

	Time to MKTX inq. w/o trade (1)	Time to MKTX inq. w trade (2)	Time to voice trade (3)
After 1 failed inquiry	0.82	0.65	1.04
After 2 failed inquiries	0.87	0.82	1.34
After 3 failed inquiries	0.85	0.85	1.46
After 4 failed inquiries	0.84	0.88	1.53
After 5 failed inquiries	0.82	0.85	1.56
After 6 failed inquiries	0.80	0.87	1.56
After 7 failed inquiries	0.78	0.89	1.63
After 8 failed inquiries	0.77	0.84	1.59
After 9 failed inquiries	0.75	0.86	1.44
After 10 failed inquiries	0.72	0.88	1.44

Table 7. The unconditional Maximum Likelihood Estimator

This table presents estimation results for the unconditional MLE, where the only control is a constant for event $k \in \{1, \dots, K\}$, with $K = 4$. Event $k = 1$ is an inquiry on MKTX without trade, $k = 2$ is an inquiry on MKTX with trade, $k = 3$ is a voice trade, and $k = 4$ is an exit. Robust standard errors as explained in Chapter 12.5.1 of [Wooldridge \(2010\)](#) are reported in parentheses. Our sample has $N = 2,383,637$ observations.

Event	MKTX inq. w/o trade (1)	MKTX inq. w trade (2)	voice trade (3)	exit (4)
	-3.59*** (2.76×10^{-6})	-4.211*** (4.67×10^{-6})	-3.40*** (2.11×10^{-6})	-2.93*** (1.71×10^{-6})

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

Table 8. The estimated coefficients of the MLE, part 1: trade characteristic dummies

This table presents the first part of our estimation results for the MLE, conditional on trade characteristics (this table) and the number of failed inquiries in the child order to date (in Table 9). “Sell” takes the value of 1 for a sale request, and zero otherwise; “Ba1 to Caa3” takes the value of 1 if the bond’s Moody’s rating is between Ba1 and Caa3; “Ca to C” is similarly defined; “COVID” takes the value of 1 if the RFQ is submitted in March 2020, and zero otherwise; “Old” takes the value of 1 if the bond’s age is above the 75th percentile of the distribution, and zero otherwise; “Turnover below median” takes the value of 1 if the bond’s quarterly turnover is below median, and zero otherwise; “High time-to-maturity” takes the value of 1 if the bond’s time to maturity is above the sample median, and zero otherwise; “Low amt outstanding” takes the value of 1 if the bond’s amount outstanding is below the sample median, and zero otherwise; “Micro size” takes the value of 1 if the quantity of dealer response is below \$100,000, and zero otherwise; “Odd lot” takes the value of 1 if the quantity of dealer response is between \$100,000 and \$1 million, and zero otherwise; “Round lot” takes the value of 1 if the quantity of dealer response is between \$1 million and \$5 million, and zero otherwise; “Block trade” takes the value of 1 if the quantity of dealer response exceeds \$5 million, and zero otherwise. We rank customers into deciles according to the number of dealer responses they receive, after controlling for inquiry size, fraction of requests for sell trades and HY bonds. “Connected decile 9” is an indicator for the customer being in decile 9, and similarly for other “Connected” indicators. Robust standard errors as explained in Chapter 12.5.1 of Wooldridge (2010) are reported in parentheses. Our sample has $N = 2,383,637$ observations.

Event	MKTX inq. w/o trade (1)	MKTX inq. w trade (2)	voice trade (3)	exit (4)
(Intercept)	-4.06*** (0.0056)	-3.65*** (0.0065)	-3.38*** (0.0049)	-2.91*** (0.0044)
Sell	0.121*** (0.0036)	0.632*** (0.0048)	0.354*** (0.0033)	-0.029*** (0.003)
Ba1 to Caa3	-0.00544* (0.0042)	-0.0246*** (0.0056)	0.0769*** (0.0039)	-0.203*** (0.0035)
Ca to C	0.00523 (0.041)	-0.4*** (0.065)	0.259*** (0.036)	-0.255*** (0.036)
COVID	-0.0705*** (0.0074)	-0.483*** (0.0098)	-0.412*** (0.0065)	-0.167*** (0.0059)
Old	0.0157*** (0.0036)	-0.108*** (0.0047)	-0.0549*** (0.0032)	0.045*** (0.0029)
Turnover below median	0.00936** (0.0041)	-0.106*** (0.0056)	-0.0617*** (0.0041)	0.164*** (0.0034)
High time-to-maturity	-0.0111*** (0.0036)	0.0539*** (0.0048)	-0.0852*** (0.0033)	0.0804*** (0.0029)
Low amt outstanding	0.155*** (0.0036)	-0.264*** (0.0047)	-0.294*** (0.0032)	0.179*** (0.0029)
Micro size	0.0277*** (0.0035)	0.213*** (0.0047)	0.399*** (0.0034)	-0.311*** (0.003)
Round lot	-0.186*** (0.0075)	-0.407*** (0.0099)	-0.0374*** (0.0071)	0.366*** (0.0056)
Block trade	-0.436*** (0.03)	-1.07*** (0.04)	0.112*** (0.026)	0.594*** (0.021)
Connected decile < 7	0.0436*** (0.0049)	-1.63*** (0.0072)	-0.219*** (0.0044)	0.0469*** (0.0038)
Connected decile 7	0.0322*** (0.0059)	-1.14*** (0.0078)	-0.0113** (0.0052)	0.0297*** (0.0046)
Connected decile 8	0.286*** (0.006)	-0.503*** (0.0076)	0.611*** (0.0057)	0.0966*** (0.0051)
Connected decile 9	0.113*** (0.0052)	-0.318*** (0.006)	-0.135*** (0.0049)	-0.123*** (0.0043)
Failed inquiry controls	Yes	Yes	Yes	Yes

Robust standard-errors in parentheses

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

Table 9. The estimated coefficient of the MLE, part 2: the failed inquiries dummies

This table presents the second part of our estimation results for the MLE, conditional on trade characteristics (in Table 8) and the number of failed inquiries in the child order to date (this table). Event $k = 1$ is an inquiry on MKTX without trade, $k = 2$ is an inquiry on MKTX with trade, $k = 3$ is a voice trade, and $k = 4$ is an exit. “Failed j ” takes the value of 1 if the number of failed inquiries in the child order to date is equal to j , and zero otherwise. Robust standard errors as explained in Chapter 12.5.1 of Wooldridge (2010) are reported in parentheses. Our sample has $N = 2,383,637$ observations.

Event	MKTX inq. w/o trade (1)	MKTX inq. w trade (2)	voice trade (3)	exit (4)
(Intercept)	-4.06*** (0.0056)	-3.65*** (0.0065)	-3.38*** (0.0049)	-2.91*** (0.0044)
Failed 2	0.509*** (0.0045)	-0.0764*** (0.0063)	-0.497*** (0.0045)	-0.345*** (0.0039)
Failed 3	0.841*** (0.0062)	-0.0685*** (0.01)	-0.729*** (0.0081)	-0.567*** (0.0068)
Failed 4	1.04*** (0.0081)	-0.0485*** (0.015)	-0.943*** (0.013)	-0.738*** (0.011)
Failed 5	1.17*** (0.01)	-0.082*** (0.022)	-1.08*** (0.02)	-0.858*** (0.015)
Failed 6	1.29*** (0.012)	-0.114*** (0.029)	-1.22*** (0.028)	-0.949*** (0.021)
Failed 7	1.36*** (0.015)	-0.193*** (0.038)	-1.23*** (0.035)	-1.1*** (0.028)
Failed 8	1.43*** (0.017)	-0.21*** (0.048)	-1.35*** (0.047)	-1.21*** (0.036)
Failed 9	1.5*** (0.019)	-0.156*** (0.056)	-1.42*** (0.06)	-1.24*** (0.044)
Failed ≥ 10	1.71*** (0.011)	-0.253*** (0.036)	-1.34*** (0.033)	-1.3*** (0.027)
Trade char. controls	Yes	Yes	Yes	Yes

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

Table 10. Poisson model for the number of dealer responses

This table presents Poisson regression estimates for number of dealer responses on indicators for the inquiry number in child orders. “Inquiry j ” takes the value of 1 if it is the j th inquiry in the child order. In column (1) we include trade characteristics described in Table 8. In column (2), we control for the unobserved child order characteristics by adding child order fixed effects to the regression. The sample excludes inquiries submitted by dealers.

Dependent Variable:	number of dealer responses	
Model:	(1)	(2)
<i>Variables</i>		
(Intercept)	1.903*** (0.0003)	
Inquiry 2	-0.3110*** (0.0008)	0.0361*** (0.0006)
Inquiry 3	-0.4241*** (0.0016)	0.0607*** (0.0012)
Inquiry 4	-0.4724*** (0.0027)	0.0670*** (0.0019)
Inquiry 5	-0.4799*** (0.0038)	0.0812*** (0.0026)
Inquiry 6	-0.4990*** (0.0051)	0.0867*** (0.0034)
Inquiry 7	-0.5211*** (0.0066)	0.0837*** (0.0044)
Inquiry 8	-0.5163*** (0.0081)	0.1017*** (0.0053)
Inquiry 9	-0.5156*** (0.0097)	0.1065*** (0.0064)
Inquiry ≥ 10	-0.4973*** (0.0055)	0.1055*** (0.0069)
Trade char. controls	Yes	
<i>Fixed-effects</i>		
child order		Yes
<i>Fit statistics</i>		
Observations	9,455,325	9,108,063
Squared Correlation	0.33738	0.99172
Pseudo R ²	0.14526	0.36693
BIC	45,117,005.9	165,520,283.0
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Table 11. Unobserved heterogeneity: Transaction costs

This table reports the estimates of regressing inquiry spreads on indicators for inquiries in child orders. In column (1), we include trade characteristics described in Table 8 and year-month fixed effects. In column (2), in addition to indicators for inquiry number in child orders, we add year-month and child order fixed effects to control for unobserved heterogeneity. As discussed in Section 2, we measure transaction cost as a markdown or markup relative to the benchmark provided by MKTX, called Composite+.

Dependent Variable:	inq. best offered transaction cost (bps)	
Model:	(1)	(2)
<i>Variables</i>		
Inquiry 2	8.008*** (0.4365)	-2.307*** (0.0600)
Inquiry 3	12.12*** (0.8524)	-2.254*** (0.1282)
Inquiry 4	13.26*** (1.536)	-2.154*** (0.2092)
Inquiry 5	14.18*** (2.415)	-1.311*** (0.2935)
Inquiry 6	12.81*** (3.402)	-2.073*** (0.3876)
Inquiry 7	11.23** (4.410)	-1.924*** (0.5107)
Inquiry 8	11.32** (5.146)	-0.6653 (0.5839)
Inquiry 9	11.47** (5.077)	-0.1492 (0.6896)
Inquiry ≥ 10	7.668 (4.992)	-1.017 (0.8269)
Trade char. controls	Yes	
<i>Fixed-effects</i>		
child order		Yes
year-month	Yes	Yes
<i>Fit statistics</i>		
Observations	8,212,803	8,218,662
R ²	0.20256	0.95339
Within R ²	0.18961	0.04535
<i>Clustered standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		