Sequential Search for Corporate Bonds*

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

September 23, 2022

Preliminary & Incomplete — Please Do Not Circulate

Abstract

A key feature of over-the-counter financial markets is that customers need to search for the best price by making repeated inquiries to dealers. However, since most existing data is gathered exclusively from successful transactions, there are few direct estimates of this search process, such as how frequently a customer makes price inquiries to dealers and the distribution of price offers that these dealers make in response. We shed new light on this sequential search process using data from the leading electronic trading platform for corporate bonds that contains direct observations about every inquiry, offer, and trade. We obtain estimates of time to trade conditional on observable trade characteristics as well as on the number of previously unsuccessful inquiries. For example, we find that it takes between two and three days for a customer to complete, on average, the purchase of an investment grade bond after first unsuccessful inquiry. We show that time to trade is much shorter for small quantities than for large ones, much shorter for sales than purchases, and varies systematically across trading venues (electronic vs. traditional market). When interpreted through the lens of a sequential search model, our estimates of time to trade highlight the importance of both observed *and unobserved* heterogeneity across customers. In particular, our results suggest that customers are very different in the intensity with which they make inquiries and in the number of responses they are able to elicit from dealers. Overall, these 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.

Keywords: Over-the-counter markets, corporate bonds

JEL Classification: G11, G12, G21.

*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.

[†]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, such as Treasuries, corporate and municipal bonds, commercial paper, 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 framework to analyze these markets, following the seminal work of Duffie, Gârleanu, and Pedersen (2005) (see Weill, 2020 for a survey). However, this literature has suffered from a crucial shortcoming: the main assumption that investors search for counterparties is not grounded in direct empirical evidence about the sequential trade process. The reason behind this shortcoming is simple: existing data from OTC markets is made up of transaction records. As a result, existing data only provides information about the time at which a trade occurs, but not about the time investors spend searching for a trade.[1]

In this paper, we leverage a proprietary data set to offer a unique window into the sequential trade process in one of the most studied OTC financial markets—the market for U.S. corporate bonds. The data provides a complete record of the inquiries made by customers on the leading electronic trading platform for corporate bonds, MarketAxess (MKTX). As in the earlier work of Hendershott and Madhavan (2015), we find that inquiries fail to trade about 30 percent of the time. 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 bond. Indeed, our data allow us to observe when customers make additional electronic inquiries on MKTX for the same trade, when they complete the trade on MKTX or on the traditional voice market, and when they abandon the trade altogether and exit the market. Taken together, these observations about repeated inquiries lead to novel estimates of search frictions in the corporate bond market.

[1]This 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.

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. In addition, our analysis reveals that time to trade is much shorter for small trades than large ones, much shorter for sales than purchases, and varies systematically across trading venues (MKTX vs. voice trades). 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.

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 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. 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 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 Trade Reporting Compliance Engine (TRACE) data. 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.

After organizing the data into parent and child orders, our second contribution is properly

measuring the time it takes to successfully trade a child order. Even with our granular observations of inquiries and trade, measuring the time to trade remains a non-trivial exercise because of the bias created by “competing risks". For example, the measured average time to trade on MKTX is biased downward, because it is based on trades that have occurred relatively quickly, before the arrival of other competing events that would have ended the child order, such as trading with a dealer outside of the MKTX platform, or deciding to abandon this particular child order and exit. We correct for this bias via Maximum Likelihood: 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 observed characteristics of the child order.

Equipped with these estimated intensities, we calculate the expected time it takes a child order to trade, either on MKTX or via voice trade, in the event that the customer doesn’t 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). Bond age, turnover, and credit rating all have statistically significant impacts on time to trade, but with economically small effects. Customer sophistication, measured by the average number of responses that a particular customer elicits on MKTX, has an economically significant impact on time to trade. Finally, 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).

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. More specifically, after controlling for various observed characteristics of child orders, we find that the number of dealer responses declines with each additional failed inquiry, the time to trade increases, and the terms of trade deteriorate. This dependence on the number of failed inquiries could in principle occur for a 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.

Finally, we make indirect inferences about these unobserved characteristics through the lens of a standard sequential search model in the tradition of McCall (1970). We 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. Intuitively, the former are consistent with a "desperate" customer, while the latter could represent a more "patient" customer.

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 sophistication).

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.[2] 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 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. Pintér and Üslü (2021) address this issue in a structural model, using joint observations of trade size and trade frequency. We propose a more direct approach, based on granular observations, which does not rely on the restrictions imposed by specific structural model.

While we find it useful to formulate a search-theoretic model to motivate our empirical exercise and interpret its findings, our measurement does not impose theoretical restrictions from the model. We specifically build on the sequential search problem of McCall (1970), which was first applied

[2] See Bessembinder, Spatt, and Venkataraman, 2020 for a survey

to the financial market by Garbade and Silber (1976). As in labor economics, this simpler partial equilibrium model is a more natural theoretical framework to interpret micro data – for example, it helps rationalize failed inquiries, repeat attempts to trade, and price dispersion.