



Large Order Size Liquidity in Treasury Markets

by

Eleni Gousgounis, Esen Onur and Bruce Tuckman

This version: June, 2020

OCE Staff Papers and Reports, Number 2020-001

**Office of the Chief Economist
Commodity Futures Trading Commission**

Large Order Size Liquidity in Treasury Markets

Eleni Gousgounis, Esen Onur and Bruce Tuckman*

June 29, 2020

ABSTRACT

We measure market liquidity for large-sized orders in the ten-year treasury futures market estimating mean-variance frontiers for their execution cost during the period of 2012 to 2017. We identify large orders from regulatory transaction data and introduce a methodological innovation to infer the urgency of a large order from the pattern of its execution. We find that the mean-variance frontier becomes significantly worse as order size increases, but that the frontier has improved over the time period studied. We also find that the costs of executing large orders on behalf of customers are significantly greater than the costs of executing orders for house accounts.

JEL classification: G10, G13

Keywords: Large Orders; Implementation Shortfall; Treasury Futures

*U.S. Commodity Futures Trading Commission, Washington, D.C. 20581. Tel: (+1) 202-418-5000. Email: egousgounis@cftc.gov. Onur: U.S. Commodity Futures Trading Commission, Washington, D.C. 20581. Tel: (+1) 202-418-5000. Email: eonur@cftc.gov. Tel: (+1) 202-418-5000. Tuckman: U.S. Commodity Futures Trading Commission, Washington, D.C. 20581. Tel: (+1) 202-418-5000. Email: btuckman@cftc.gov. Tel: (+1) 202-418-5000. The analyses and conclusions expressed in this paper are those of the author(s) and do not reflect the views of other Commission staff, the Office of the Chief Economist, or the Commission. All errors and omissions, if any, are the authors' own responsibility. Corresponding author: Eleni Gousgounis.

I. Introduction

It is widely recognized that liquidity provision across many financial markets has been shifting from dealers to principal trading firms (PTFs) in both the futures and inter-dealer cash market¹. In the generic, bygone business model, dealers traded large positions with customers and then worked out of those positions over time. Customers paid for the service, but dealers bore the execution risk. It was almost exclusively the dealers who had to optimize the execution of large trades, breaking them up into smaller orders and buying or selling strategically over time. More recently, PTFs have become the largest providers of liquidity, passively or actively buying and selling in relatively small quantities and usually transacting through automated systems. They adjust their prices and quantities, dynamically so, as to earn small profits per trade on a large number of trades. At the same time, traditional customers, like asset managers, trading large positions on their own or through an intermediary, have to bear the execution risk themselves. As a result of this relatively new paradigm, many more market participants now have to focus on the execution of large trades.

A vigorous and wide-ranging debate has accompanied the changes just described, both in the popular press and the academic literature. How has market liquidity changed for large orders? Most of the debate has been framed in terms of bid-ask spreads, market depth, the Amihud measure, and the price impact of individual transactions. Relying on these measures, the literature does not find that liquidity deteriorated over time in the 10-Year U.S. Treasury Note futures market². Furthermore, according to the

¹Clark and Mann (2016), “A deeper look at Liquidity Conditions in the Treasury Market”, Treasury Notes Blog, U.S. Department of Treasury. Retrieved on January 8th, 2020 from <https://www.treasury.gov/connect/blog/Pages/A-Deeper-Look-at-Liquidity-Conditions-in-the-Treasury-Market.aspx>

²Adrian et al. (2017), Adrian et al. (2015), Adrian et al. (2016), Bessembinder et al. (2016), DTCC (2015), Joint Staff Report (2015), Trebbi and Xiao (2019). TBAC (2013) notes that bid-ask spreads are “spiky,” so that liquidity as measured by bid-ask spreads may not be consistently available.

Chicago Mercantile Exchange (CME), there is even evidence of liquidity improvement³. However, such liquidity metrics are best suited to address the costs of executing small orders, while as pointed out in the 2019 Financial Stability Report, markets are liquid when market participants can trade large quantities without triggering outsized price changes⁴. Consequently, one might be able to gain a better view of market liquidity by focusing specifically on the execution of large orders. This approach, which is typically limited by data availability, would also address anecdotal concerns⁵, raised by market participants, over an alleged increase in difficulty to conduct large trades in treasury markets. Rick Rieder, global fixed income chief investment officer at Blackrock, provided support for this claim, while speaking at the Fed’s Third Annual Conference on the Evolving Structure of the U.S. treasury markets; he argued that “the markets look like they’re on a certain price level, but when you go to transact they are not at that level and then people get hit or lifted and then try and unwind their risk; you can create a staircase dynamic as people are trying to get out of their risk in smaller size.”⁶ In the same conference, David Tepper, co-founder of Appalosa Management, countered that “there is an extra cost to it, but there is always a way to fool the machines.” He also noted that “Markets are liquid; the futures markets are fantastically liquid.”⁷

Our paper uses a unique dataset and aims to shed some light to this debate, focusing

³CME Group (2016), “The New Treasury Market Paradigm. Treasury Futures.” Retrieved on January 8th, 2020 from <https://www.cmegroup.com/education/files/new-treasury-market-paradigm.pdf>

⁴Board of Governors of the Federal Reserve System (2018a), Financial Stability Report, Retrieved on January 8th 2020 from <https://www.federalreserve.gov/publications/files/financial-stability-report-20191115.pdf>

⁵Bao and Zhou (2015), on the Global Financial System (2016), Blackrock (2015), Blackrock (2016), Committee on Capital Markets Regulation (2015), Dick-Nielsen and Rossi (2018), Papanyan (2015), Board of Governors of the Federal Reserve System (2018b), and Wood (2015).

⁶Hall (2018), “Buy side reports price movement risk in Treasury trading,” Fi-Desk. Retrieved on January 8th, 2020 from <https://www.fi-desk.com/buy-side-reports-price-movement-risk-in-treasury-trading>

⁷Hall (2018), “Buy side reports price movement risk in Treasury trading,” Fi-Desk. Retrieved on January 8th, 2020 from <https://www.fi-desk.com/buy-side-reports-price-movement-risk-in-treasury-trading>

on the study of the execution cost of large orders in the 10-year U.S. Treasury futures market across time. We specifically distinguish between large orders traded for house accounts or on behalf of customers, while the literature focuses mostly on orders traded on behalf of customer accounts.⁸ To assess the execution costs of large orders, we follow [Engle et al. \(2012\)](#). In particular, we study the mean-variance frontier of market impact, measured by implementation shortfall (IS), in the 10-year U.S. Treasury futures market. Traders may choose to execute large orders relatively rapidly, by, for example, hitting existing bids or lifting existing offers. This choice, will, on average, result in relatively high average IS, but relatively low variance around that average. By contrast, traders might choose to execute large orders relatively slowly, by, for example, strategically using limit orders to avoid paying bid-ask spreads. This choice will result in a relatively low average IS, but a relatively high IS variance. The mean-variance combinations for various levels of execution urgency are depicted on the mean-variance frontier, which is a downward sloping convex curve.

While we are not the first to consider the piecemeal execution of large orders or the resulting mean-variance frontier of IS, this is the first paper to study large orders in the Treasury futures market. We use unique regulatory transaction data on the 10-year Treasury Note futures contract, which allow us to examine the complete universe of large orders in this market during the observation period. Similar to [Korajczyk and Murphy \(2018\)](#), we construct the unobservable large, “parent orders,” by consolidating “child orders,” which are labeled in the data by account. Details will be provided later in the paper, but, roughly speaking, sequential child orders on behalf of the same participant, which are executed within a relatively short time frame, and which are on the same side of the market, are taken to be part of a single parent order. We find that there is a

⁸For a comprehensive description of papers referencing implementation shortfall of customer orders, see [Hu et al. \(2018\)](#).

great variation across parent orders in size, number of child orders and trades, and time to execution.

Averaging and computing the variance of IS across all orders is not sufficient, of course, to generate a mean-variance frontier. [Engle et al. \(2012\)](#), had a special data set from a particular bank, which included the relative urgency of customer orders, i.e., how important it is to the customer to transact quickly at high but certain cost as opposed to transacting more strategically at a lower expected cost but with higher variance. In general, however, such data does not exist. We, therefore, introduce the methodological innovation of inferring a parent order's urgency based on its observed execution strategy. With this measure of urgency, we are essentially able to estimate the expected value and variance of IS for each urgency, which is exactly what constitutes the mean-variance frontier of execution costs.

Our contribution to the literature is threefold. First, our analysis reinforces the findings of [Engle et al. \(2012\)](#). The estimated mean variance frontiers in the highly liquid 10-year Treasury futures market behave similarly to the mean-variance frontiers in [Engle et al. \(2012\)](#), who study equity markets. More specifically, we find a significant trade-off between the mean and the variance of IS, and this trade-off is much more pronounced for larger orders. While [Engle et al. \(2012\)](#), have a more accurate exogenous urgency measure, we introduce a methodology to extract urgency from the market participants' trading behavior. However, our study is more inclusive, as it utilizes all large orders in the market, while the data set used in [Engle et al. \(2012\)](#) contains just orders, initiated by Morgan Stanley traders on behalf of their clients or by buy side traders on behalf of a portfolio manager. Therefore, our results reinforce the mean-variance frontiers presented in [Engle et al. \(2012\)](#)⁹.

⁹We would like to thank Joel Hasbrouck for pointing out that the efficient trading frontier was originally suggested by [Kissel and Glantz \(2003\)](#)

Second, in an effort to address the debate over the liquidity of the treasury futures markets, we track how the mean-variance frontiers of large orders evolve over time. We find that the mean-variance frontier for large orders in this market has improved over the sample period. This finding, which indicates that liquidity has improved over the period of 2012-2017, is in agreement with the literature, which uses more conventional liquidity measures. Moreover, our results provide evidence against the concerns, often raised by market participants, over the relatively increased execution difficulty of large orders.

Third, we compare mean-variance frontiers separately for house and customer accounts. In our sample over half of large orders are initiated done on behalf of house accounts, and it is therefore worthwhile comparing them to customer orders. This is the first paper, to our knowledge to conduct such a comparison. We find that customer orders appear to face higher execution costs, compared to orders executed for house accounts. This finding is robust to restricting the set of traders to those who routinely trade both customer and house accounts. Moreover, orders executed by traders who exclusively trade for house accounts enjoy the lowest mean-variance frontier. This is consistent with a market in which certain traders have greater skills or access to better execution technology.

The plan of the paper is as follows. Section II discusses the literature review related to transaction costs of large orders and liquidity. Section III outlines the conceptual approach that underlies our methods and offers a brief description of the data set used. Section IV provides summary statistics for our sample, develops the empirical models used and explains our results. Section V offers our concluding remarks.

II. Literature Review

Our paper fits in the literature exploring the cost of trading large orders in various markets. Most of the papers in the literature, similar to our approach, use the idea of implementation shortfall (IS) to measure this cost, which was initially suggested by [Perold \(1988\)](#). A challenging part of working on IS of large orders is finding the right data that identifies executions belonging to that order. Some researchers use proprietary data that identifies transactions belonging to large, institutional orders ([Van Kervel and Menkveld \(2019\)](#), [Sağlam \(2018\)](#), [Engle et al. \(2012\)](#)), while others use the Aber Noser data to gain access to this information . Finally, similar to our methodology, a few of the papers in the literature back out large orders from transaction data under certain assumptions and then measure IS.¹⁰

In terms of research utilizing proprietary institutional investor data, [Van Kervel and Menkveld \(2019\)](#) make use of order executions by four institutional investors and combine that data with public HFT transaction data in Swedish index stocks. Their research question is designed to understand the actions of HFTs around these large institutional orders and they find that HFTs need considerable time to detect these large, informed institutional orders. While HFTs supply liquidity to institutional orders during this process, they start trading in the same direction of institutional orders once they catch on, and this makes it more costly for institutional investors to execute their large orders.

Similarly, [Sağlam \(2018\)](#) analyzes predictable patterns in large order execution strategies and how those relate to the cost of execution. Their empirical analysis uses large orders submitted by 146 investors, mostly institutional portfolio managers, and reveal that the cost of trading is positively correlated with predictable trading patterns. They conclude that their findings are consistent with predatory trading (back-running)

¹⁰[Korajczyk and Murphy \(2018\)](#) and [Putnins and Barbara \(2017\)](#) are a few examples of such papers.

idea rather than sunshine trading.

Abel Noser data allow researchers answer various research topics. While [Barbon et al. \(2019\)](#) focus on how brokers diffuse (leak) information to their best clients, [Ben-Rephael and Israelsen \(2017\)](#) document the trading desk skills of management companies can vary depending on which institutional client they are serving. In terms of the analysis of execution cost of institutional orders using Abel Noser data, [Chakrabarty et al. \(2017\)](#) find that majority of short-term institutional trades lose money, especially in more volatile markets and in small stocks. [Brogaard et al. \(2014\)](#) analyze whether HFTs increase the execution costs of institutional investors and they find no evidence suggesting that HFTs causally increase or decrease these costs.

The approach in our paper is more in line with other studies backing out large orders from proprietary transaction level data. Similar to ours, [Korajczyk and Murphy \(2018\)](#) make use of a transaction and message level data provided by a Canadian markets regulator to analyze changes around a regulatory change in the marketplace. They find that this exogenous change reduced HFT order activity, which also coincided with increased bid-ask spreads and decreased price impact for institutional orders. Another study making use of regulatory data is [Putnins and Barbara \(2017\)](#), where authors make use of a transaction level data from the Australian market regulator, ASX. They identify a rich heterogeneity among algorithmic and high-frequency traders in their data and while some of these traders increase institutional trading costs, others decrease them. Overall, they find that the negative effect of toxic traders is offset by the positive effect of beneficial traders in the markets they analyze. Similar to those two papers, [Chen and Garriott \(2020\)](#) also use transaction and message data from the Montreal Exchange with trader identifiers, which they use to identify large institutional trades as well as high frequency traders. They find that existence of HFTs in the Government of Canada Bond futures market actually improves transaction costs for institutional traders.

Our findings also complement existing literature on variation on execution costs. [Anand et al. \(2011\)](#) find significant variation in execution costs across trading desks of various management companies, showing that even when the average execution cost is low, variation can be high. [Ben-Rephael and Israelsen \(2017\)](#) study the differences in execution costs across clients within management companies and find there to be systemic differences across clients for a subset of management firms. We measure the execution cost of large house orders and compare them to the execution cost of similarly sized customer orders and show that execution costs for customer orders are significantly higher than those of house orders.

In terms of methodology, our paper follows [Engle et al. \(2012\)](#), [Almgren and Chriss \(2000\)](#) and [Almgren \(2003\)](#) which suggest that a reduction in execution costs by taking longer to execute the order corresponds to added risk or trading. [Engle et al. \(2012\)](#) estimates a risk-return frontier, which allows them to calculate a risk-adjusted cost of execution. We utilize this approach in our paper as well.

III. Data and Methodology

A. Data

Our data set comprises transactions of 10-year Treasury Note futures contracts from February 15th, 2012, to November 30th, 2017. These contracts are traded electronically on Globex, the electronic platform of the Chicago Mercantile Exchange (CME). We limit our attention to outright trades (e.g., excluding calendar trades) in the front month futures contracts and to transactions that originate from market or limit orders.

Our data set is constructed from the Transaction Capture Report database of the U.S. Commodity Futures Trading Commission (CFTC), which contains detailed transaction

information, including the transaction time, price and quantity of every futures trade. The database also includes information on the order from which each trade originated, namely the order entry time, the order type (e.g. market, limit, stop-order), a flag for automated trades, and an order identifier, which allows us to identify which transactions are part of which order. Finally, the database has fields pertaining to the counterparties to each transaction and the traders executing that transaction. In particular, the database identifies the counterparties, indicates the buyer and the seller, indicates which side initiated the trade, provides an identification number for each trader, and contains a customer type indicator (CTI), which allows us to distinguish customer from proprietary trades¹¹.

B. Methodology

B.1. Parent Orders

The first step in our analysis is to identify large orders placed by market participants. Our data set contains an order identifier, which along with market participant information, can accurately aggregate transactions into “child orders.” These trades represent successive partial fills of the same child order. Unfortunately, our data does not provide any information on cancellations and modifications. However, these child orders are most likely to have come from market participants who are slicing larger, “parent orders” into these smaller, child orders. While there is no way to know with certainty which child orders belong to which parent orders, we can try to infer parent orders from the trading pattern of a market participants’ child orders.

¹¹This allows us to identify customer and proprietary accounts. We can also identify which traders execute orders just for their proprietary accounts, just for customer accounts or for both their proprietary accounts and some customer accounts. However, trader to customer association is not one-to-one. A trader generally has more than one customer whose orders she executes and technically a customer can have their orders handled by more than one trader

To identify parent orders, we first aggregate all transactions on the same date, originating from the same account, and with the same order id into child orders. Next, we aggregate these child orders into parent orders based on the following rules: sequential child orders belong to the same parent order if they were placed on the same day by the same account and trader, and were all on the same side of the market; exclude parent orders where the average time between the entry of child order exceeds one hour; exclude orders taking longer than one day to execute, as well as orders entered or executed on weekends and holidays. Furthermore, since our intention is to study large orders, we include in our sample only those parent orders for at least one thousand contracts, which correspond to less than the largest 1% of identified parent orders.

B.2. Urgency

As pointed out by [Engle et al. \(2012\)](#), the cost of executing an order depends on the level of risk that the agent is willing to assume. Risk averse agents trade rapidly, at relatively high average costs with little variance, while those willing to tolerate a higher level of risk execute slower, more opportunistic trading strategies, which are characterized by relatively low average costs but relatively higher variance of costs. The data set used by [Engle et al. \(2012\)](#) comes from a particular broker and contains information on the instructions provided by the owner of an order to the executing broker. Our data set has the advantage of including all large orders placed in the market, but does not have any information as to the intended trading strategies.

To proxy for the urgency of each order, we contrast the distribution of its transactions during the day with an estimate of the distribution of market volume. More specifically, we define an urgency measure as the difference between the volume-weighted execution time of a given parent order and the volume-weighted execution time of a hypothetical VWAP order of the same size. A VWAP – volume-weighted average price – order is one

in which a customer wants to realize a price equal to the volume-weighted average price in the market over the time frame of the order¹². Furthermore, the most typical way to execute a VWAP order is to trade over time in the same proportion as market volumes. Say, for example, that over the three-period time frame of an order, market volume is distributed 20%, 30%, and 50%. In that case, 20% of a VWAP order would be executed in the first period, 30% in the second, and 50% in the third. Mathematically, urgency, u , is given by:

$$u = VWAT_{market} - VWAT_{order}, \quad (1)$$

$$VWAT_{market} = \sum_{t=1}^n \frac{v_{market,t}t}{V_{market}}$$

$$VWAT_{order} = \sum_{t=1}^n \frac{v_{order,t}t}{V_{order}}$$

where n is the number of minutes from the entry of the first child order of a given parent order to the time of the market close, $v_{market,t}$ is the aggregate quantity transacted in the market during minute t , as estimated from the time-pattern of daily market volumes over the last 30 days, V_{market} is the aggregate quantity transacted in the market from minutes 1 to n , as estimated from the time-pattern of daily market volumes over the last 30 days, $VWAT_{market}$ is the volume-weighted execution time of a hypothetical VWAP order over n minutes, $v_{order,t}$ is the aggregate quantity of the order transacted during minute t , V_{order} is the aggregate quantity of the order transacted during minutes 1 to n , and $VWAT_{order}$ is the volume-weighted execution time of parent order.

To repeat, urgency is how much faster the order is executed, in minutes, than a hypothetical VWAP trade of the same size from the time the first child order is entered in the book until the market close. The choice of market close as the end of the time

¹²Effectively, lacking any detailed information on the actual urgency of customers, we are assuming that the execution time frame ends at the market close.

frame implies that trades placed early in the day that are executed relatively quickly, will be measured as having more urgency than trades placed later in the day that are executed just as quickly. To normalize for the time of the execution of the first child order, we normalize urgency by scaling it by its upper and lower boundaries.

More specifically, since $VWAT_{order}$ ranges between 1 and n minutes, the boundaries for u are given by: $VWAT_{market} - n \leq u \leq VWAT_{market} - 1$.

The lower boundary ($VWAT_{market} - n$) is never positive and is the minimum order urgency in minutes, corresponding to the case of a trader executing the entire order at the close. Similarly, the upper boundary ($VWAT_{market} - 1$) is never negative and is maximum possible urgency in minutes, corresponding to the case of a trader executing the entire order in the first minute.

Using these boundaries, we define a normalized urgency measure, U , as follows:

$$\begin{aligned} \text{when } u < 0 \text{ then } U &= \frac{u}{n - VWAT_{market}}, \\ \text{when } u = 0 \text{ then } U &= 0, \\ \text{when } u > 0 \text{ then } U &= \frac{u}{VWAT_{market} - 1}. \end{aligned}$$

By construction, therefore, normalized urgency measure ranges from -1 to 1 , where negative (positive) values correspond to orders executed at a slower (faster) pace than a VWAP order. For our empirical analysis, we multiply normalized urgency by 100, which puts it in the range of -100 to 100 .

B.3. Execution Costs

Our data allows us to estimate execution costs for each parent order. Consistent with the literature, we estimate the implementation shortfall (IS) for each parent order as:

$$IS = 10,000 * D * \frac{P_{order} - P_b}{P_b}, \quad (2)$$

where $P_{t,order}$ is the volume weighted average price of the parent order, P_b is the benchmark price, measured by the average price of trades occurring in the one minute interval preceding the entry of the first child order of a given parent order, and D is the trade direction indicator, which is equal to 1(-1) for a buy(sell) order.

In words, IS compares the realized execution price of the order to the price just before the arrival of the order in the market, and is expressed in this paper in basis points. An order would be considered buy-initiated (sell-initiated) if it is a buy (sell) order and the IS is positive (negative).

B.4. The Model

We model execution costs, following [Engle et al. \(2012\)](#), accounting for both the expected execution cost (IS) and the risk of trading the order. More specifically, the execution cost of each parent order i is given by:

$$IS_i = \exp(X_i' \beta) + \exp\left(\frac{1}{2} X_i' \gamma\right) \epsilon_i, \quad \epsilon_i \sim N(0, 1) \quad (3)$$

where X_i is a vector of conditioning information, which in our base model consists of order and market characteristics. Order characteristics include the logarithm of the order size, the urgency of the order (U_i), and a dummy indicating whether the order belongs to a customer ($cust_dummy_i$). Market characteristics include the logarithm of

the volume of the specific contract on the day of a given parent order execution, and a measure of volatility defined as the logarithmic difference of the maximum and minimum prices of the contract on that day.

The model restricts expected IS and the variance to be positive. This is obviously a reasonable assumption for the variance, and, because our analysis focuses on large orders, is also a reasonable assumption for the mean. This does not imply, of course, that all realizations of IS are positive. In our data set, implementation shortfall is positive for over 60% of the parent orders, and this is roughly the proportion for both customer and proprietary orders.

Under this model, the expected execution cost and variance of each order are estimated jointly using a maximum likelihood estimator under the normality assumption for ϵ_i , where:

$$E(IS_i) = \exp(\beta_1 volatility_i + \beta_2 logvolume_i + \beta_3 logsize_i + \beta_4 U_i + \beta_5 cust_dummy_i), \quad (4)$$

$$E(IS_i) = \exp(\gamma_1 volatility_i + \gamma_2 logvolume_i + \gamma_3 logsize_i + \gamma_4 U_i + \gamma_5 cust_dummy_i)$$

The coefficients of the variables appearing linearly in the exponent can be interpreted as the percentage change in the expected IS for one unit change in X , while the coefficients of the variables appearing as logarithms in the exponent can be interpreted as the elasticity of the corresponding non logged variable.

We also extend the model below to include in the explanatory variables a time trend and its interaction with the logarithm of order size in order to evaluate whether execution costs have changed over time.

IV. Empirical Analysis

A. Descriptive Statistics

Table I presents descriptive statistics for the large parent orders constructed for this analysis. On average, these large parent orders are for about 2,000 contracts, are divided into about 40 child orders, are executed through 200 transactions, and take a total of about 1 hour and 20 minutes to execute. There is significant variation, however, around these averages. For example, at the 95th percentile, parent orders are for about 4,500 contracts, are divided into nearly 200 child orders and over 500 transactions, and take 6.25 hours to execute. Manually executed orders comprise about 40% of the sample.

The urgency of the parent orders, calculated as described in the previous section, average 155 minutes faster than a hypothetical VWAP execution, which seems proportionately quite urgent given the average parent order execution time of 78 minutes. In normalized terms, average urgency is 81% of the way from VWAP execution to immediate execution of the entire order.

Finally, the implementation shortfall of our large parent orders is, on average, 0.27 basis points. The variation around this average, however, is quite large, with a standard deviation of 7 basis points and a 5th to 95th percentile of observations ranging from between -9 and 9 basis points. This large variation relative to the mean supports this paper's focus on both the mean and variance of implementation shortfall to measure liquidity. Our dataset allows us to differentiate the type of accounts from which parent orders originate. Market participants have to declare this type for every order through a CTI code. Table II explains the CTI codes in detail and gives the distribution of parent orders across codes. Briefly, CTI 1 applies to orders initiated and executed for the account of an individual member, but not for proprietary trading. CTI 2 applies to orders initiated and executed for the proprietary accounts of a member firm. CTI 3

applies to orders that a member executes on behalf of another member. CTI 4 applies to orders entered by or on behalf of nonmember entities. According to table II, 99% of the parent orders, which comprise 98% of trading volume, are CTI 2 or CTI 4. For the rest of this paper, therefore, we focus exclusively on these two order types, and refer to them as proprietary and customer orders, respectively. Note that about 40% of the orders in the sample are proprietary and about 60% are for customers.

Table III presents the same summary statistics as in table I, but separated by proprietary and customer parent orders. Generally, customer orders are larger, broken up into a greater number of child orders and executions, take longer to execute, and are more often executed manually. Another key difference, which will be explored further later in the paper, is that customer orders tend to have larger and more variable costs, as measured by *IS*.

B. Implementation Shortfall - Base Model

Our base model, described in the previous section, aims to identify the drivers of transaction costs for large orders in the 10-year treasury futures market. It jointly estimates the expected value and variance of *IS*. We use explanatory variables similar to those in Engle et al. (2012). More specifically, to capture order characteristics, we use the logarithm of the order size, the urgency of the order (U_i), and a dummy indicating whether the order belongs to a customer. To control for market conditions, we use the logarithm of the market volume of the contract on the day of the parent order and price volatility, as measured by the logarithmic difference of the maximum and minimum contract prices on that day. Table IV gives the estimated coefficients of the equations for the mean and variance of the *IS* of large orders in the sample.

The results from the baseline model are for the most part in line with those from the

literature (e.g., [Engle et al. \(2012\)](#)). We find that expected IS increases with market volatility, order size, and urgency, but decreases with market volume. The variance of IS increases with market volatility and order size, and decreases with urgency, as anticipated, but—contrary to priors—increases with market volume.¹³

The effect of order size on the expected value and variance of the costs of trading is economically as well as statistically significant. Relative to an order for 1,000 contracts, an order for 1,650 contracts incurs a 46% increase in expected cost and a 21% increase in the variance of the cost.

The estimated coefficients on the customer dummy variable indicate that, on average, customers face both a higher expected cost and a higher variance of transacting. This finding will be discussed in greater detail below, but is also economically significant. Relative to a proprietary order, a customer order incurs nearly double the expected cost and a one-third increase in variance.

The economic significance of urgency is best illustrated by the mean-variance frontiers of the costs of trading. Using the estimated coefficients of the base model, along with an intraday volatility of 45% and log market volume of 14 (which are approximately sample means), [figure 1](#) shows the frontier facing a customer with an order of 1,000 contracts. Each point on the frontier corresponds to a different level of urgency, and [figure 1](#) highlights two such points, namely, urgency values of 80 and 95. The graph as a whole ranges from an urgency of 100 at the leftmost point to an urgency of 76 at the rightmost point, which range covers the upper 75% of parent orders in our sample. The resulting frontier is downward sloping and convex, confirming that relatively urgent orders execute at higher mean cost but lower variance, while relatively non-urgent orders execute at lower mean cost but higher variance. Furthermore, the mean-variance

¹³This negative estimate is different than what is observed in the literature. We believe this might be because of how we define market volume. While we add the trading volume of the contemporaneous day in our regressions, [Engle et al. \(2012\)](#) adds the lagged 21 day median daily volume to their regression.

trade-off is economically significant: by choosing across the range of urgencies shown, a customer with a 1,000 contract order can choose over a wide range of mean-variance IS combinations, from 0.54 mean and 6.28 variance to 0.16 mean and 15.7 variance.

Figure 2 shows the IS mean-variance frontiers for each of the following four order profiles: customer orders of 1,000 and 5,000 contracts, and proprietary orders of 1,000 and 5,000 contracts. As in figure 1, each graph covers urgency levels from 100 to the far left and 76 to the far right. The economic significance of the differences are striking, both of customer vs. proprietary orders and of 1,000 vs. 5,000 contracts. Order sizes of 5,000 contracts face far worse frontiers than orders of 1,000 contracts, and customer orders face significantly worse frontiers than proprietary orders.

C. Implementation Shortfall - Time Trend

We now explore how the transaction costs of large orders have changed over time, that is, over our sample period from early 2012 until late 2017. For this purpose, we add to the base model a time trend and an interaction term of time trend with parent order size. This exercise can be viewed as testing the proposition that, while execution in the treasury futures market might have become less costly from the growing presence of algorithmic traders, those savings might not have accrued to large orders.

The estimated coefficients in table V reveal a complex relationship between IS , time, and order size. The time trend indicates that mean IS has increased, but its variance has decreased. Relative to smaller orders, however, larger orders have experienced lower IS mean but higher variance. To sort out the relative importance of these effects, figure 3 graphs the mean-variance frontier for customer orders of 1,000 contracts over time. More specifically we present the corresponding mean-variance frontiers for the first quarter of 2012 (Q1 2012), the first quarter of 2013 (Q1 2013), the first quarter of 2014 (Q1 2014)

and the first quarter of 2015 (Q1 2015).

While the interpretation of the regression coefficients, presented in table V, are challenging to interpret, Figure 3 makes it clear that the mean-variance frontier has shifted down and to the left over time. In other words, transaction costs have fallen or liquidity has improved, as measured by the mean-variance frontier available to market participants. Furthermore, the shapes of the frontiers are such that, for trades executed relatively slowly, those in the lower right, there has been a sizeable decrease in variance and a marginal increase in mean. By contrast, for trades executed relatively quickly, those in the upper left, there has been a more modest decrease in variance in addition to a somewhat more pronounced increase in mean. Finally, while not illustrated here, results are similar for proprietary trades. For very large orders (i.e. 5000 contracts), the mean also decreases, as evidenced by the negative interaction of time and parent order size in table V.

D. Implementation Shortfall - House or Customer Accounts

As revealed by our base model, both with and without a time trend, the cost difference between house and customer accounts is sizeable. On average, the cost experienced by customers is almost double of that experienced by house accounts, and the variance of cost for customers is about a third larger. This section explores these sizeable differences in greater detail.

One possible explanation of cost differences across orders is that some traders in these markets are more sophisticated along certain dimensions than others and might even possess better trading technology and be able to trade faster.¹⁴ Furthermore, these highly skillful traders might choose to keep these superior abilities to themselves and

¹⁴Haynes and Roberts (2019) describe the differences in speed across various contracts traded at the CME. Their results indicate that a few of the traders enjoy the highest trading speeds.

not trade at all on behalf of customers.

To pursue this possible story, we define *P-traders* as those that trade large orders on behalf of their own accounts, that is, trades exclusively coded CTI 2, at least 95% of the days in the sample. Analogously, we define *C-traders* as those who trade large orders on behalf of customer accounts, that is, trades exclusively coded CTI 4, at least 95% of the days in our sample. We define all other traders as *P&C-traders*, to be thought of as trading for both house and customer accounts¹⁵. Finally, we exclude traders who trade coded CTI 1 and CTI 3, when those CTI 1 and CTI 3 large orders represent more than (or equal to) 25% of the number of large orders or volume they trade. This is to avoid having traders with considerable trading activity in CTI 1 and CTI 3 orders be classified as *P-traders*, *C-traders* or *P&C-traders*. Such traders are relatively small and rare and when we exclude them we are able to keep about 99% the original sample of large orders.

Table VI presents summary statistics for these three different groups of traders. The top panel, which includes all traders, shows that the number of *C-traders* and *P-traders* are similar, while *P&C-traders* number less than half than those in the other two groups. In terms of average daily activity and total trading activity, however, *P&C-traders* are the most active on average, trading more than four million contracts per day. Finally, the last two columns in the panel show that a third of the trading done by *P&C-traders* is proprietary and two-thirds is on behalf of customers.

Since some traders are active for only a few days in our sample, the bottom panel presents analogous statistics only for traders who are active, i.e., traders who have transacted at least 20 out of the 1460 days in our sample. The count of active traders across these three trader groups is relatively balanced. And, for these active traders,

¹⁵Note that we apply no exclusivity restriction in our definition. Our *P&C-traders* traders can handle customer and house accounts within the same day.

P&C-traders are still the group with highest volume of trades, twice that of *P-traders* and about ten times more than *C-traders*, indicating, again, that most customer-trading is done by *P&C-traders*.

Returning to the hypothesis that the higher cost of customer orders is due to their being executed by less skilled traders, table VII shows the results of an implementation shortfall regression, like that of our earlier base model, but with dummy variables for customer orders executed by *P&C-traders*, customer orders by *C-traders*, and proprietary orders by *P&C-traders*. Hence, all of the estimated coefficients on these dummies are relative to a benchmark corresponding to the omitted dummy, namely, proprietary orders by *P-traders*.

The estimated coefficients in table VII show that the lowest cost of trading is experienced by proprietary trades executed by *P-traders*, followed, in ascending order, by proprietary trades by *P&C-traders*, customer trades by *C-traders*, and customer trades by *P&C-traders*. The estimated variance of implementation shortfall is the smallest for proprietary trades by *P-traders*, followed by the three other categories in the same order as the mean, but are much more similar in value. The greater mean and variance of *IS* for the three groups of trades relative to those of proprietary trades by *P-traders* is also economically significant. Customer trades executed by *P&C-traders* cost 109 percent more on average and have a 62 percent higher variance. For customer trades by *C-traders*, the average and variance are 56 percent and 60 percent higher, respectively, while for proprietary trades by *P&C-traders*, they are 38 percent and 58 percent higher.

Figure 4 shows the mean-variance frontier for the cost of proprietary and customer trades of 1,000 contracts through *P-traders*, *C-traders*, and *P&C-traders*. Consistent with the regression results of table VII, proprietary trades by *P-traders* has the lowest frontier. The frontiers for the three other cases are higher by economically significant magnitudes. Together, then, table VII and Figure 4 are consistent with the hypothesis

that the level of transaction costs depends on trader type, and that, in our sample, *P-traders* - who, by definition, do little to no customer trades - are able to realize the lowest costs.

Next, we ask two detailed questions in order to understand the comparison among the more costly groups in our analysis. First, how do costs among customer order types vary across trader types? Second, we explore how costs for between customer and proprietary orders compare across the same group of traders, namely the *P&C-traders*.

Table VIII shows the estimates from a regression run solely on customer orders. The estimates of the *C-trader* dummy variables in the regression are both positive and significant, reinforcing the previous findings that customer trades done by *C-traders* are more costly, on average, and have higher variance. The average execution cost estimate is also economically significant; namely the implementation shortfall of customer orders traded by *C-traders* is more than 40 percent higher than that of customer orders traded by *P&C-traders*. The difference for the variance of implementation shortfall is much smaller, only higher by 3 percent.

Table IX turns the focus on *P&C-traders* only and evaluates the costs of customer orders separately from proprietary orders using a *Customer* dummy in the regression. Within the same group of traders, those who trade both customer and proprietary orders, we show that customer orders are statistically and economically more costly, on average, and also have a higher estimated variance. Specifically, customer orders experience on average 43 percent higher implementation shortfall and their variance is higher by about 5 percent. We can only speculate that the cost difference must be due the divergence between the nature of these two types of orders, such as their information content. Understanding the exact causes of these differences is an exercise we plan to undertake in the future.

V. Concluding Remarks

This paper develops a mean-variance framework, similar to [Engle et al. \(2012\)](#), to study the execution costs of large orders in the 10-year Treasury Note futures market, through time, across proprietary and customer orders, and across traders that typically trade for their account, for customer accounts, or for both. Our results indicate that from 2012 to 2017 the mean-variance frontier for large orders has shifted downwards, meaning that orders of similar size could be executed over time with lower mean IS at a fixed variance or with lower variance of IS at a fixed mean at lower execution costs towards the end of our sample. We find that orders executed by traders, who trade exclusively for their own accounts, face the lowest frontier, followed, in ascending order by proprietary orders executed by traders that trade for both house and customer accounts, customer orders executed by traders who trade almost exclusively for customers, and customer orders executed by traders who execute for both house and customer accounts. Determining the causes underlying this hierarchy of costs is left for future research.

References

- Adrian, T., M. Fleming, M. Shachar, and E. Vogt.** 2015. “Has U.S. Corporate Bond Market Liquidity Deteriorated?” Federal Reserve Bank of New York Liberty Street Economics.
- Adrian, T., M. Fleming, M. Shachar, and Z. Wojtowicz.** 2016. “Corporate Bond Market Liquidity Redux: More Price-Based Evidence.” Federal Reserve Bank of New York Liberty Street Economics.
- Adrian, Tobias, Michael Fleming, Or Shachar, and Erik Vogt.** 2017. “Market Liquidity After the Financial Crisis.” *Annual Review of Financial Economics*, 9(1): 43–83.
- Almgren, Robert F.** 2003. “Optimal execution with nonlinear impact functions and trading-enhanced risk.” *Applied Mathematical Finance*, 10(1): 1–18.
- Almgren, Robert F., and Neil Chriss.** 2000. “Optimal execution of portfolio transactions.” *Journal of Risk*, 3(2): 5–39.
- Anand, Amber, Paul Irvine, Andy Puckett, and Kumar Venkataraman.** 2011. “Performance of Institutional Trading Desks: An Analysis of Persistence in Trading Costs.” *The Review of Financial Studies*, 25(2): 557–598.
- Bao, Jack, Maureen O’Hara, and Alex Zhou.** 2015. “The Volcker Rule and Market-Making in Times of Stress.” , (2016-102).
- Barbon, Andrea, Marco Di Maggio, Francesco Franzoni, and Augustin Landier.** 2019. “Brokers and Order Flow Leakage: Evidence from Fire Sales.” *The Journal of Finance*, 74(6): 2707–2749.
- Ben-Rephael, Azi, and Ryan D Israelsen.** 2017. “Are Some Clients More Equal Than Others? An Analysis of Asset Management Companies’ Execution Costs.” *Review of Finance*, 22(5): 1705–1736.
- Bessembinder, Hendrik, Stacey E Jacobsen, William F Maxwell, and Kumar Venkataraman.** 2016. “Capital commitment and illiquidity in corporate bonds.”
- Blackrock.** 2015. “Addressing Market Liquidity.” Blackrock Viewpoint.
- Blackrock.** 2016. “Addressing Market Liquidity: A Broader Perspective on Today’s Bond Markets.” Blackrock Viewpoint.
- Board of Governors of the Federal Reserve System.** 2018a. “Financial Stability Report.” Board of Governors of the Federal Reserve System Report.

- Board of Governors of the Federal Reserve System.** 2018*b*. “Financial Stability Report.” Board of Governors of the Federal Reserve System Report.
- Brogaard, Jonathan, Terrence Hendershott, Stefan Hunt, and Carla Ysusi.** 2014. “High-Frequency Trading and the Execution Costs of Institutional Investors.” *Financial Review*, 49(2): 345–369.
- Chakrabarty, Bidisha, Pamela C. Moulton, and Charles Trzcinka.** 2017. “The Performance of Short-Term Institutional Trades.” *Journal of Financial and Quantitative Analysis*, 52(4): 1403–1428.
- Chen, Marie, and Corey Garriott.** 2020. “High-frequency trading and institutional trading costs.” *Journal of Empirical Finance*, 56: 74 – 93.
- Clark, J., and G. Mann.** 2016. “A deeper look at Liquidity Conditions in the Treasury Market.” U.S. Department of the Treasury Treasury Notes Blog.
- CME Group.** 2016. “The New Treasury Market Paradigm.” CME Group Treasury Futures.
- Committee on Capital Markets Regulation.** 2015. “Nothing but the Facts: U.S. Bond Market Liquidity.” Committee on Capital Markets Regulation.
- Dick-Nielsen, Jens, and Marco Rossi.** 2018. “The Cost of Immediacy for Corporate Bonds.” *The Review of Financial Studies*, 32(1): 1–41.
- DTCC.** 2015. “Trends and Risks in Bond Market Liquidity.” DTCC Discussion Paper.
- Engle, Robert, Robert Ferstenberg, and Jeffrey Russell.** 2012. “Measuring and Modeling Execution Cost and Risk.” *The Journal of Portfolio Management*, 38(2): 14–28.
- Hall, C.** 2018. “Buy side reports price movement risk in Treasury trading.” Fi-Desk Editorial.
- Haynes, R., and J. S. Roberts.** 2019. “Automated Trading in Futures Markets - Update 2.” U.S. Commodity Futures Trading Commission Research Paper.
- Hu, Gang, Koren M. Jo, Yi Alex Wang, and Jing Xie.** 2018. “Institutional trading and Abel Noser data.” *Journal of Corporate Finance*, 52: 143 – 167.
- Joint Staff Report.** 2015. “The U.S. Treasury Market on October 15, 2014.” Department of the Treasury, Board of Governors of the Federal Reserve System, Federal Reserve Bank of New York, Securities and Exchange Commission, Commodity Futures Exchange Commission Report.

- Kissel, Robert, and Morton Glantz.** 2003. *Optimal Trading Strategies: Quantitative Approaches for Managing Market Impact and Trading Risk*. AMACOM Publishers.
- Korajczyk, Robert A, and Dermot Murphy.** 2018. “High-Frequency Market Making to Large Institutional Trades.” *The Review of Financial Studies*, 32(3): 1034–1067.
- on the Global Financial System, Committee.** 2016. “Fixed Income Market Liquidity.” BIS CGFS Papers 55.
- Papanyan, Shushanik.** 2015. “Heightened Bond Liquidity Risk is the New Normal.” BBVA Research U.S. Economic Watch).
- Perold, André F.** 1988. “The implementation shortfall.” *The Journal of Portfolio Management*, 14(3): 4–9.
- Putnins, Talis J., and Joseph Barbara.** 2017. “Heterogeneity in How Algorithmic Traders Impact Institutional Trading Costs.” Mimeo, Available at SSRN: <http://ssrn.com/abstract=2813870>.
- Sağlam, Mehmet.** 2018. “Order anticipation around predictable trades.” *Financial Management, forthcoming*, n/a(n/a).
- TBAC.** 2013. “Assessing Fixed Income Market Liquidity.” US Department of the Treasury Presentation to Treasury Borrowing Advisory Committee (TBAC).
- Trebbi, Francesco, and Kairong Xiao.** 2019. “Regulation and Market Liquidity.” *Management Science*, 65(5): 1949–1968.
- Van Kervel, Vincent, and Albert J. Menkveld.** 2019. “High-Frequency Trading around Large Institutional Orders.” *The Journal of Finance*, 74(3): 1091–1137.
- Wood, Duncan.** 2015. “GFMA, IIF, ISDA Plan Liquidity Lobbying Push.” Risk.net.

Figure 1: Mean-variance frontier for a customer parent orders for 1,000 contracts

Figure 1 shows the mean-variance frontier for the *IS* of a customer order of 1,000 contracts, using estimated coefficients from the base model, an intraday volatility of 40%, and a log market volume of 14. The graph shows urgency levels between 76 and 100.

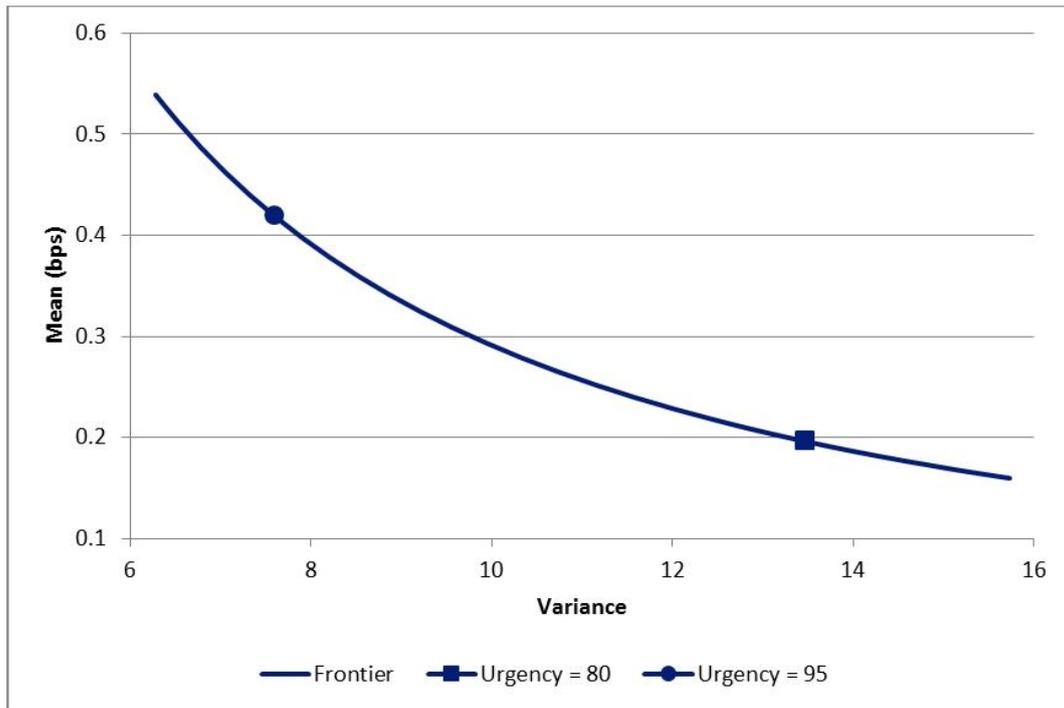


Figure 2: Mean-variance frontier: customer and proprietary orders for 1,000 and 5,000 contracts

Figure 2 shows the mean-variance frontier for the *IS* of a customer order of 1,000 contracts, a customer order of 5,000 contract, a proprietary order of 1,000 contracts, and a proprietary order of 5,000 contracts. As in figure 1, the frontiers are constructed using estimated coefficients from the base model, an intraday volatility of 40%, and a log market volume of 14, and show urgency levels between 76 and 100.

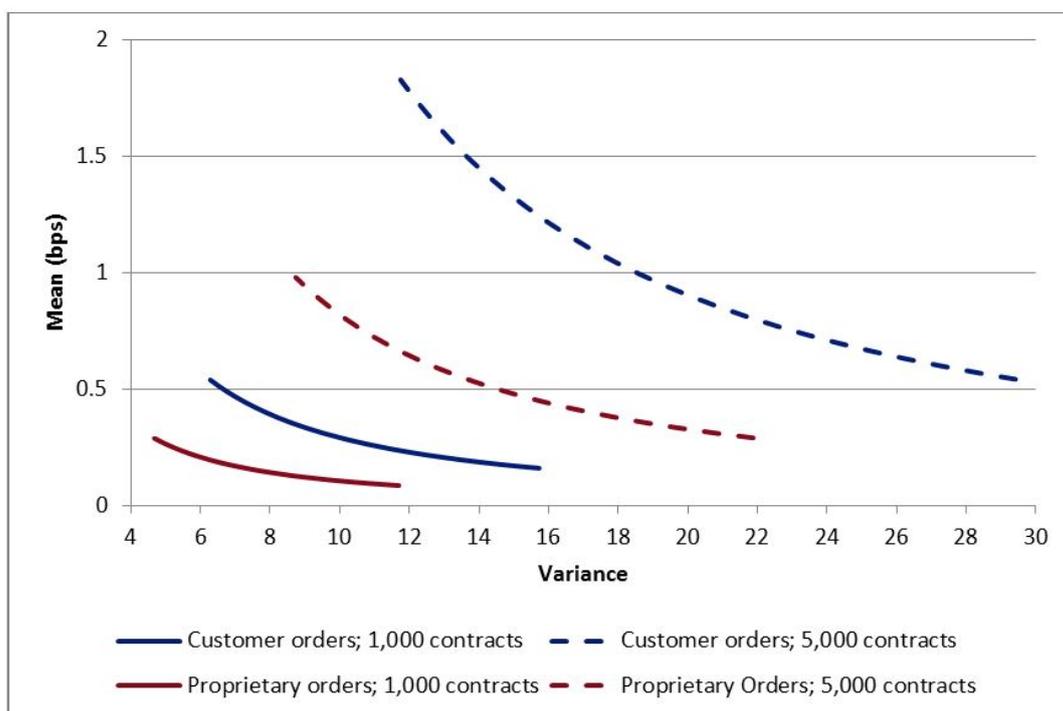


Figure 3: Mean-variance frontiers for customer orders of 1,000 contracts, across time

Figure 3 shows the mean-variance frontier for the *IS* of a customer order of 1,000 contracts over time. The four frontiers are constructed using estimated coefficients from the base model with the time trend and time-order size interaction term, an intraday volatility of 40%, and a log market volume of 14. We estimate the frontier for the first quarter of 2012 (2012 Q1), the first quarter of 2013 (2013 Q1), the first quarter of 2014 (2014 Q1) and the first quarter of 2015 (2015 Q1). As in the previous figures, the graphs show urgency levels between 76 and 100.

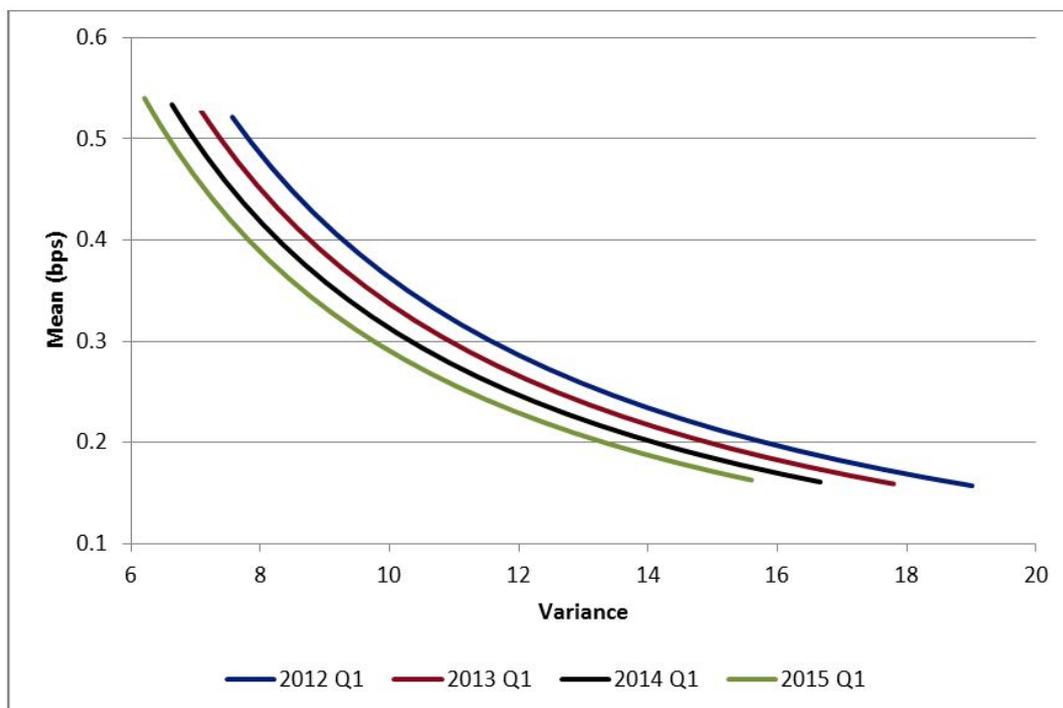


Figure 4: Mean-variance frontier for customer and proprietary orders by trader type

Figure 4 shows the mean-variance frontier for customer and proprietary trades executed by *C-traders*, *P-traders* and *P&C-traders* type. The mean implementation shortfall is shown on the Y-axis and variance of implementation shortfall is on the X-axis. Blue lines in the graph represent the mean-variance frontier for customer orders, while the red lines represent the mean-variance frontier for proprietary orders. Orders executed by a *C-traders* are on the red solid line, while orders executed by a *P-traders* are on the blue solid line. Dashed lines represent order executed by *P&C-traders*.

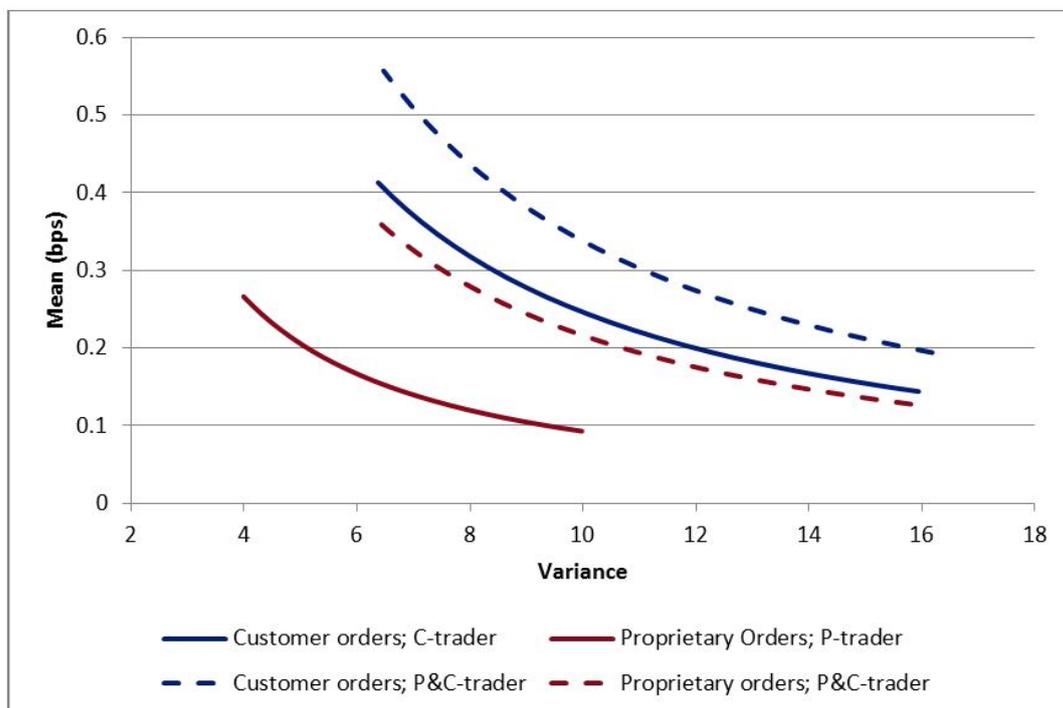


Table I: Summary Statistics for Parent Orders

Table I gives summary statistics for parent orders in the data set. Parent order size is the total size of the identified parent order, measured in number of contracts. Number of child orders is the number of child orders associated with a given parent order. Number of trades is the total number of transactions associated with a given parent order. Total time to execution is the number of minutes it takes to execute a given parent order, measured from the time the first child order is entered in the order book to the time of the last execution of that parent order. Time between entry of child orders for a given parent order is the average time between the entry of subsequent child orders, across all child orders in that parent order, measured in minutes. Manual trades is the proportion of transactions associated with a given parent order that are traded using manual access to the market. Initiated trades shows is the volume-weighted proportion of transactions associated with a given parent order in which the trader is the aggressor. Urgency is the difference between the time to executing a hypothetical VWAP order and the time to executing a given parent order, measured in minutes. Higher values indicate greater urgency in the execution of that parent order. Normalized urgency normalizes Urgency to a variable between -100% and 100%, where -100% indicates full execution at the time of the last transaction of a given parent order and 100% indicates full execution at the time of the first transaction. Finally, implementation shortfall is measures the difference between the contract price just before the start of the execution of a given parent order and the volume-weighted transaction prices of that parent order, measured in basis points.. The sample contains 292,436 parent orders, although there are only 251,431 observations for the time between entry of child orders, due to a number of parent orders that contain a single child order.

Variable	Mean	Median	5 th Percentile	95 th Percentile	Std Dev
Parent order size	1,957	1,447	1,000	4,584	1,782
Number of child orders	41	6	1	189	149
Number of trades	200	139	39	558	240
Total execution time (minutes)	78	13	0	375	143
Time between entry of child orders (minutes)	7	1	0	36	12
Manual trades (%)	40%	0%	0%	100%	0
Initiated trades (volume weighted %)	64%	72%	0%	100%	0
Urgency (minutes)	155	144	10	411	131
Urgency (normalized)	81%	97%	9%	100%	0
IS (bps)	0.27	0.47	-9.04	8.92	7

Table II: Distribution of parent orders by CTI code

Table II shows the distribution of parent orders by CTI code: CTI 1: Electronic Trading, Open Outcry and Privately Negotiated - Applies to transactions initiated and executed by an individual member for his own account, for an account he controls, or for an account in which he has an ownership or financial interest. However, transactions initiated and executed by a member for the proprietary account of a member firm must be designated as CTI 2 transactions. CTI 2: Electronic Trading, Open Outcry and Privately Negotiated – Applies to orders entered or trades executed for the proprietary accounts of a member firm. CTI 3: Electronic Trading – Applies to orders entered by a member or a nonmember terminal operator for the account of another individual member or an account controlled by such other individual member. CTI 3: Open Outcry and Privately Negotiated - Applies to orders that a member executes on behalf of another individual member, or for an account such other member controls or in which such other member has an ownership or financial interest. CTI 4: Electronic Trading Open Outcry and Privately Negotiated - Applies to all orders and transactions not included in CTI categories 1, 2 or 3. These typically are orders entered by or on behalf of nonmember entities.

Parent orders by CTI code				
CTI code	No. of orders	Total volume (contracts)	No. of orders (%)	Total volume (%)
1	3,469	13,317,492	1%	2%
2	173,692	304,929,539	59%	53%
3	243	661,921	<0.1%	<0.1%
4	115,032	253,477,370	39%	44%
All	292,436	572,386,322	100%	100%

Table III: Summary statistics for proprietary and customer parent orders

Table III presents descriptive statistics separately for proprietary and customer trading accounts. The variables are as described in table I. The data set has 174,692 proprietary and 115,032 customer parent orders, 156,114 and 92,593 of which, respectively, have strictly more than one child order.

Variable	Proprietary					Customer				
	Mean	Median	5 th Per- centile	95 th Per- centile	Std Dev	Mean	Median	5 th Per- centile	95 th Per- centile	Std Dev
Parent order size	1,756	1,380	1,000	3,766	1,211	2,204	1,507	1,000	5,482	2,216
Number of child orders	25	6	1	112	71	65	5	1	325	218
Number of trades	174	131	40	446	179	238	157	38	701	305
Total execution time (minutes)	59	7	0	304	120	107	32	0	439	168
Time between entry of child orders (minutes)	5	1	0	30	10	10	3	0	43	14
Manual trades (%)	28%	0%	0%	100%	0	57%	100%	0%	100%	0
Initiated trades (volume weighted %)	66%	74%	0%	100%	0	62%	69%	0%	100%	0
Urgency (minutes)	149	144	18	374	116	165	145	-1	449	152
Urgency (normalized)	85%	98%	18%	100%	0	76%	92%	0%	100%	0
IS (bps)	0.22	0.41	-7.24	7.13	6	0.37	0.62	-11.53	11.30	8

Table IV: Execution costs of parent orders (base model)

Table IV gives coefficients from the jointly estimated equations for the expected value and variance of IS. The independent variables are intraday price volatility, as measured by the logarithmic difference of the maximum and minimum prices of the contract on the day of the order; the logarithm of the market volume of the contract on that day; the logarithm of the parent order size; the normalized urgency of the order, and a dummy variable indicating whether the order belongs to a customer. There are 288,724 parent orders in the data sample.

Variable	Mean			Variance		
	Estimate	Std. error	t-value	Estimate	Std. error	t-value
Constant	-8.072	0.624	-12.94	-4.058	0.136	-29.76
Intraday volatility	0.728	0.082	8.90	3.311	0.015	213.89
Log volume	-0.272	0.037	-7.43	0.375	0.010	38.73
Log parent order size	0.759	0.019	39.57	0.389	0.005	70.99
Urgency (normalized)	0.051	0.004	14.37	-0.038	0.000	-305.86
Customer dummy	0.624	0.024	26.09	0.296	0.005	54.54

Table V: Execution costs of parent orders across time

Table V gives coefficients from the jointly estimated equations for the expected value and variance of IS . The independent variables are intraday price volatility, as measured by the logarithmic difference of the maximum and minimum prices of the contract on the day of the order; the logarithm of the market volume of the contract on that day; the logarithm of the parent order size; the normalized urgency of the order; a dummy variable indicating whether the order belongs to a customer; a time trend, expressed in quarter years; and an interaction of that time trend with the logarithm of parent order size. There are 288,724 parent orders in the data sample.

Variable	Mean			Variance		
	Estimate	Std. error	t-value	Estimate	Std. error	t-value
Constant	-8.560	0.673	-12.72	-4.633	0.161	-28.86
Intraday volatility	0.716	0.083	8.65	3.189	0.016	201.61
Log volume	-0.266	0.038	-7.07	0.468	0.010	46.69
Log parent order size	0.823	0.037	21.95	0.325	0.012	26.89
Urgency (normalized)	0.050	0.003	14.40	-0.038	0.000	-306.75
Customer dummy	0.625	0.024	26.19	0.287	0.005	52.88
Time (quarter years)	0.040	0.020	1.96	-0.052	0.006	-8.42
Time x log parent order size	-0.005	0.003	-2.01	0.005	0.001	6.20

Table VI: Summary Statistics by Trader Type

Table VI gives descriptive statistics on three types of traders. *P-traders* execute proprietary orders at least 95 percent of the days they appear in our sample of large orders. *C-traders* execute customer orders 95 percent of the days they appear in the sample. The rest of the traders are *P&C-traders*. The first panel of the table provides descriptive statistics for all traders and the second panel for traders who are active (traders who have transacted at least 20 out of the 1460 days in our sample) on at least 20 days in our sample.

Trader type	No. of traders	Avg. No. active days	Avg. Daily volume	Total volume	Customer volume (%)
All traders					
P-trader	178	129	1,130,077	201,153,696	0%
C-trader	164	72	241,403	39,590,024	99%
P&C-trader	69	455	4,551,477	314,051,946	67%
Active traders (>20 active days)					
P-trader	79	284	2,529,240	199,809,974	0%
C-trader	68	166	568,350	38,647,786	99%
P&C-trader	58	540	5,411,889	313,889,565	67%

Table VII: Execution costs of parent orders by order and trader type

Table VII gives coefficients from the jointly estimated equations for the expected value and variance of IS. The independent variables are intraday price volatility, as measured by the logarithmic difference of the maximum and minimum prices of the contract on the day of the order; the logarithm of the market volume of the contract on that day; the logarithm of the parent order size; the normalized urgency of the order; and three dummy variables indicating order and trader type, namely: (a) customer parent orders executed by *P&C-traders*; (b) customer orders executed by *C-traders*; and (c) proprietary orders executed by *P&C-traders*. There are 286,842 parent orders in our sample.

Variable	Mean			Variance		
	Estimate	Std. Error	t-value	Estimate	Std. Error	t-value
Constant	-7.792	0.642	-12.13	-4.162	0.137	-30.41
Intraday volatility	0.702	0.080	8.75	3.223	0.016	207.72
Log volume	-0.248	0.036	-6.83	0.390	0.010	40.12
Log parent order size	0.758	0.019	40.19	0.356	0.006	63.83
Urgency (normalized)	0.044	0.004	11.96	-0.038	0.000	-304.31
Indicator: Customer order by P&C-trader	0.738	0.027	27.12	0.483	0.006	77.52
Indicator: Customer order by C-trader	0.439	0.055	8.02	0.468	0.011	42.49
Indicator: Proprietary order by P&C-trader	0.319	0.045	7.16	0.460	0.007	61.61

Table VIII: Execution costs of customer parent orders by trader type

Table VIII gives coefficients from the jointly estimated equations for the expected value and variance of IS for customer parent orders only. The independent variables are intraday price volatility, as measured by the logarithmic difference of the maximum and minimum prices of the contract on the day of the order; the logarithm of the market volume of the contract on that day; the logarithm of the parent order size; the normalized urgency of the order; and a dummy variable equal to 1 when the parent order is executed by a *C-traders* and equal to 0 otherwise. The sample contains 113,881 parent orders.

Variable	Mean			Variance		
	Estimate	Std. error	t-value	Estimate	Std. error	t-value
Constant	-6.539	0.685	-9.55	-4.921	0.207	-23.77
Intraday volatility	1.047	0.089	11.80	2.998	0.023	128.34
Log volume	-0.305	0.043	-7.06	0.523	0.015	35.46
Log parent order size	0.793	0.020	39.36	0.243	0.008	31.40
Urgency (normalized)	0.039	0.003	11.41	-0.034	0.000	-200.69
C-trader dummy	0.345	0.054	6.44	0.030	0.011	2.70

Table IX: Execution costs for *P&C-traders*

Table IX gives coefficients from the jointly estimated equations for the expected value and variance of IS for parent orders executed by P&C-traders. The independent variables are intraday price volatility, as measured by the logarithmic difference of the maximum and minimum prices of the contract on the day of the order; the logarithm of the market volume of the contract on that day; the logarithm of the parent order size; the normalized urgency of the order; and a dummy variable equal to 1 when the parent order is on behalf of a customer and equal to 0 otherwise. The sample contains 146,909 parent orders.

Variable	Mean			Variance		
	Estimate	Std. error	t-value	Estimate	Std. error	t-value
Constant	-8.722	0.785	-11.10	-5.353	0.182	-29.40
Intraday volatility	1.114	0.086	12.96	3.036	0.021	146.72
Log volume	-0.316	0.043	-7.29	0.527	0.013	40.65
Log parent order size	0.834	0.021	39.82	0.272	0.007	38.76
Urgency (normalized)	0.059	0.005	12.29	-0.033	0.000	-217.66
Customer dummy	0.359	0.040	8.94	0.052	0.008	6.74