Institutional Trading Frictions

Chiraphol N. Chiyachantana and Pankaj K. Jain

This Version: March 2009

JEL classifications: G14, G15, G20, G24

Keywords: Institutional trading, implementation shortfall, adverse selection, clean-up costs, price impact, transaction costs.

We are grateful to Abel Noser Corporation for providing us with institutional trading data and to Judy Maiorca for related discussions. Chiyachantana is at Singapore Management University, Lee Kong Chian School of Business, and can be reached at *chiraphol@smu.edu.sg*. Jain is at University of Memphis, Fogelman School of Business and can be reached at *pankaj.jain@memphis.edu*. Suggestions participants at 2008 meeting of National Bureau of Economic Research (NBER) market microstructure group, and finance seminars at University of Memphis are highly appreciated. All remaining errors belong to the authors.

Institutional Trading Frictions

Abstract

We propose and empirically examine a comprehensive measure of institutional trading frictions to include the dimensions of price impact, quantity of execution, return dynamics, speed of execution, and trading commissions. Our empirical analysis reveals that various hidden components of institutional trading frictions such as adverse selection and clean-up costs could add significantly to previously measured directly observable components of transaction costs. Market conditions, order direction, firm size, stock volatility, liquidity demanding or supplying profile, speed of transaction, number of brokers, exchange listing, and analyst coverage explain the variations in trading frictions for different orders. The hidden clean-up cost of non-execution persists in 60 alternative measurement periods. Transaction cost risks are higher for small stocks and volatile stocks.

Introduction

In his AFA presidential address, Stoll (2000) highlights the importance of trading frictions and presents various measures of observed trading costs. He distinguishes between real and information components of friction and also between their static and dynamic measures. He concludes that work remains both in deepening our understanding of friction and broadening the scope of research. This is the task that we undertake in our paper by proposing and empirically examining a comprehensive measure of institutional trading frictions to include the dimensions of price impact, quantity of execution, return dynamics, speed of execution, and trading commissions (acronym PQRST).

We analyze the determinants of PQRST, its persistence, and its randomness. Whereas the vast majority of academic microstructure studies focus on observed transaction cost magnitude, Almgren and Chriss (2000) and Engle and Ferstenberg (2007) suggest in their theoretical models that it is equally important to understand the risk arising from randomness of transaction cost and frequent non-execution of orders. Only with all these tools in hand, institutions optimize their trading strategies in practice to improve alpha capture, reduce price slippage, and prevent orders from being gamed. Carrie (2008) attributes the rapid proliferation of dark pools in alternative trading systems to the increasing importance of implicit trading cost and timing risk, the higher costs of traditional block desk worked orders, potential information leaks, and the backlog effect from slowdown in execution because of transaction costs. We provide a comprehensive empirical analysis of several of these issues in our paper for institutional traders.

Recognizing the concept that reality in financial markets involves the cost of trading and the cost of not trading, Perold (1988) launched the pioneer work on institutional trading costs, defining implementation shortfall as the difference between paper performance and actual performance of a portfolio manager or an investment strategy. The implementation shortfall of institutional trading has two basic components. The first, execution cost, relates to the transactions institutions actually execute and arises from price impact, commission, and other transaction fees and taxes. The second, hidden opportunity cost, relates to transactions that institutions fail to execute and the adverse selection risks of both executed and unexecuted transactions. An order that a fund manager does not fulfill to his satisfaction contains the gain that the investor had to forgo by failing to achieve an ideal investment. Likewise, an increased probability of filling orders with adverse price movements or not filling orders with favorable price movements adds to the institutional trading frictions. Thus, execution quantity and return dynamics are important dimensions of institutional trading frictions. Finally, the speed of execution is especially important in the context of institutional trading because of the role that it plays in balancing the observed price impact and the hidden opportunity costs.

The first component of institutional trading cost is widely studied. For instance, Berkowitz, Logue and Noser (1988), Keim and Madhavan (1995, 1996, 1997), Chan and Lakonishok (1993, 1997), Jones and Lipson (2001), Conrad, Johnson, and Wahal (2001, 2003), Chakravarty, Panchapagesan and Wood (2004) analyze the relationship between investment styles, trade motivations, exchange listing, trading systems and price impact of trades within the U.S. Chiyachantana, Jain, Jiang and Wood (2004) provide new evidence on determinants of price impact and changes in asymmetry of price impact of buys and sells in bull and bear markets in 37 countries. Goldstein, Irvine, Kandel, and Weiner (2009) examine the impact of commission costs on institutional trading patterns.

All these studies focus on observable costs of completely filled orders. However, a portfolio manager would be interested in both the observed and hidden components of transaction costs. Similarly academic studies need to incorporate the deviations of real portfolio performance from the ideal paper portfolio performance. Both observed and hidden components cause frictions. To get a complete picture of implementation shortfall, it is necessary to analyze the interaction between the PQRST dimensions. Wagner and Edwards (1993) implore the inclusion of these costs in best execution measures and suggest that directly observable costs are only the tip of the iceberg. Since their study, institutional trading environment has changed dramatically. A comprehensive and continuous literature on these additional dimensions of institutional trading frictions did not develop largely due to data limitations. Order execution strategy evolves dynamically. Institutions do not always trade their planned order volume. When institutions trade, are they the beneficiaries or victims of the adverse selection in asymmetric information environment? What happens on the road that was planned but not taken by the institutional trader or broker? Does inaction in the financial markets lead to serious opportunity costs?

Our paper fills this gap in the literature by presenting an analysis of such frictions in the context of institutional trading. The dataset spans 776 orders of institutional clients over period of 1999-2005. A sub-set of Abel/Noser's institutional clients provide actual order volume and transaction data to Abel/Noser. We conduct our analysis with both the full sample of orders worth $20 trillion and this reduced subset of orders worth $4.59 trillion and get qualitatively similar results. Whereas the full sample is more representative of overall institutional trading, the sub-set relates a bit more closely to the research questions that we ask. This is because in the latter specification, we eliminate the observations where order level information is merely aggregated or algorithmically generated by Abel/Noser and not directly provided by the client. This enables us to work with actual fill rates instead of estimated fill rates. Although no academic researcher can read the minds of all institutional managers to know their dynamically evolving portfolio and trading decisions, this dataset gets us as close to large measurement of hidden institutional trading frictions as is practically possible.

Our research design and the unique features of the dataset offers several advantages by tracking each institutional buy or sell order down to order placement and trade implementation. Characterizing the distribution of institutional orders that result in trades versus those that do not result in actual transactions, helps us understand the quantity dimension of execution quality. The fact that we don't need to infer trade direction makes our results more accurate compared to research designs that rely on trade only datasets. Quantity dimension interacts with the return dynamics of stocks resulting in a mismatch between actual portfolio performance and ideal performance that academic papers usually measure using price datasets such as CRSP. Order level data also lets us study the time dimension of execution quality. For portfolio managers lost time is lost performance in terms of their assets under management and cash drag. Such hidden aspects of institutional trading frictions can only be studied with detailed information about the direction, timing, and quantity of the institutional investment order, which are missing from most standard microstructure datasets[1].

For a subset of all institutional managers, our dataset enables us to define unfilled shares volume for each institution order as shares in an order minus the actual transaction volume for that order. Filled and unfilled volumes represent the quantity dimension of execution quality. We analyze adverse selection cost or cost recovery from filled shares. Then we analyze clean up costs or benefit for unfilled shares. An adverse difference between a future price and the per

[1] A large body of market microstructure research examining the impact of transaction costs exclusively on observed costs (e.g. spread). The widely used database, such as Trade and Quote (TAQ), only provides information on actual execution of trades quotations for small quantities. We have access to both trading orders and actual transactions for large quantities.

share price on order date represents a cost whereas a favorable price movement represents a benefit. The interaction between quantity and return dimensions captures the hidden opportunity cost, which can add significantly to the directly observable costs. We define total frictions as the sum of proportionate price impact, commissions, adverse selection cost or cost recovery, and clean up costs, all in basis points. Formal definitions and equations for each component are described in more detail later in Section IV of the paper.

The time dimension of execution is measured by looking at the number of days it takes to fill an order with a given set of characteristic. Given the endogenous nature of trade timing, we control for several order characteristics when measuring this dimension of institutional trading friction. Finally, we estimate a two stage simultaneous system of equations for the institutional trading frictions, *PQRST,* capturing various dimensions of execution quality.

Our paper contributes to literature in several ways. First, we show that hidden clean-up costs are comparable in magnitude to the directly observed price impact or commissions, which doubles the estimate of total costs in relation to previous studies. An additional hidden component of institutional friction is the adverse selection cost. Our analysis throws light on the informativeness of institutional trades by establishing that adverse selection cost component is negative. Thus, institutions in our sample are actually more informed than their counterparties and are able to recover part of their cost by using this informational advantage. The dollar magnitude of total transaction costs with its explicit and hidden components in our sample period is nearly $30 billion, nearly a third of which is incurred in form of hidden clean-up costs net of cost recovery. Thus, institutional portfolio managers and academic readers of our paper can have a more accurate idea of the total transaction costs shaved from the paper portfolio return of their investment strategies. For instance, one-way total execution costs average at 67 basis points. It follows from 106% annual turnover computed from CRSP mutual fund dataset from 1999-2005 that one-way transaction costs ate into nearly 15% of the value weighted CRSP return of 4.60% in that period.

Second, we characterize the determinants of both the order fill rate and the dollar clean-up cost. We conjecture that the former are governed by trading strategies where as the latter includes the return loss component which may demonstrate the information superiority of the institution which trades versus that of their counterparties. Third, our control variables let us analyze how bull or bear market condition, listing exchange, and analyst coverage affect both fill rates and total friction. Fourth, we analyze transaction cost risk and opportunity cost persistence.

Firm-specific and order-specific variables that stand out as the main determinants of overall execution quality are firm size, relative order volume, stock volatility, liquidity demand, order duration, multiple brokers, and market conditions, listing exchange, and analyst coverage. Friction is higher for small stocks, large orders, liquidity demanding orders, and orders executed over multiple days or using multiple brokers. Lower fill rates and higher magnitude of adverse returns jointly lead to higher transaction costs. The results have several practical and academic implications. Institutional investors can use them as benchmarks to analyze their own implementation shortfall. The numbers can also provide guidance of whether or not it pays to be aggressive in completely filling the institutional orders. More importantly, the implementation policy can be customized to address the effect of market conditions, firm-specific characteristics, and order-dynamics. From the academic perspective, these measures of transaction costs provide the evidence of the limit to arbitrage and also implore asset pricing models to include the transaction costs because they can lead to significant deviations from the ideal performance of a paper portfolio that we so often see in theoretical and empirical papers.

The remainder of this study is organized as follows. In Section I, we provide the background information about hidden components of institutional trading frictions and discuss the potential variables that can affect order fill rates and the clean-up costs of unfilled orders. In Section II, we describe our data sources, research design, empirical methods, and related literature. Section III empirically documents the quantity dimension of friction and its determinants. Section IV computes the adverse selection cost (or cost recovery) of filled orders and clean-up costs resulting from the failure to fully execute an order. Apart from the magnitude, we characterize the variance and persistence of both explicit and hidden costs. This section also discusses the determinants of the various dimensions of friction using a two-stage regression framework. Section V contains concluding remarks and potential directions for future research.

I. Background on Explicit and Hidden Frictions

Consideration of both explicit costs and hidden opportunity cost is a fundamental theoretical concept spanning many areas of financial decision making. In the field of portfolio management, Perold (1988) highlights the importance of both explicit and hidden costs in accurately computing the implementation shortfall, which is the difference between paper performance and actual performance of a portfolio manager or an investment strategy. Wald and Horrigan (2005) and Handa and Schwartz (1996) develop a theoretical models of clean-up costs and bagging costs of trading with limit orders based on a joint distribution for the subsequent return and the order execution probability. Although they state that their model is applicable primarily to small orders, similar issues are present in institutional trading as well. Bertsimas and Lo (1998) develop a theoretical model of institutional order splitting aimed at trading cost reduction. Yet, the empirical estimation of hidden clean-up costs has been a challenge due to data limitations. Even though the standard microstructure datasets such as TAQ provide detailed time-stamped ex-ante information about quotes and ex-post information about trades, they do not entirely reveal the intention of the traders in terms of their desired order size. Limit order book datasets such as NYSE OpenBook take us a step further by making the submitted orders transparent but do not completely resolve the situation. Institutions often use order splitting over time and across venues and brokers, thus fragmenting the information contained in the limit order book. Datasets containing order information about original institutional intention to trade such as those by Elkins/McSherry, Plexus and Abel/Noser have been used sporadically by Perold and Sirri (1993) and Domowitz, Glen and Madhavan (2001) but these datasets are proprietary in nature. Limited access to such data have restricted the proliferation of empirical research leaving ample scope for further research to enhance our understanding about institutional trading costs.

Our paper contributes to the literature by presenting a comprehensive empirical analysis of the various dimensions of the institutional trading costs – particularly its quantity and time dimensions. Quantity dimension arises from the fact that only some orders are executed and their costs can be observed directly. Other orders are not filled as desired if a passive trading strategy is used. Passive trading strategy is necessary to balance the higher price impact cost of aggressive strategies but the former result in higher risk or uncertainty of order execution. Unfilled volume must be filled later using a potentially worse price or could even result in lost returns from an otherwise good investment decision. In either case, there is a cost to clean out the unfilled volume. Distinguishing between the two outcomes of unfilled volume - future execution or lost returns - is empirically difficult, if not impossible because of multiplicity of trading venues and dynamic order duration. Nevertheless, the similarity of the implications of the two outcomes makes it possible to compute a clean-up cost for a hypothetical trade execution for the unfilled volume at a fixed interval after the order submission.

We now present a simple numerical example these explicit and hidden institutional trading frictions. Suppose, a portfolio manager identifies an undervalued stock priced at $100 at time 0 and submits a buy order for 10,000 shares at time 1. At time 2, a transaction of 4,500 shares takes place at $100.90 for which a commission of $0.05 per share is paid. The stock price rises to $101.50 in the short term at time 3 and $110 in the long run at time 365. In this example, without the knowledge of order fill rates, a researcher could observe that for a 10,000 share decision the CRSP return is 1.50% or $15,000 in the short term, which is the focus of our paper. However, the institutional trader does not capture this full CRSP return. An amount of $4,050 is lost in price impact costs and $225 in commissions. If we close the books at time 3 by buying the remaining unfilled quantity of 5,500, then the clean-up costs amount to $8,250. This leaves the institutional trader with a short term return of only $2,475. The goal of our paper is to quantify deviation from CRSP return caused by various explicit and hidden trading frictions in the institutional trading setting. Since the short term return is positive in this example, the adverse selection cost is negative. Had the price on time 3 fallen to 99 instead of rising, then the filled quantity would represent an adverse selection cost of $4,500 but the clean-up costs would be negative. Since an order can only have either clean-up costs or adverse selection cost, it is useful to offset these amounts to arrive at the net hidden costs. These components of trading frictions can be summarized as follows:

Friction={Clean-up cost + Adverse selection cost + Price impact}* Order direction + Commission
OR Cost Recovery

$$PQRST = \left\{ \left(\frac{P_{t+x}}{P_{d-1}} - 1 \right) * (1 - w_e) + \left(1 - \frac{P_{t+x}}{WTP} \right) * (w_e) + \left(\frac{WTP}{P_{d-1}} - 1 \right) * w_e \right\} * \mathrm{OD} + \left(\frac{C_t}{P_{d-1}} - 1 \right) * w_e \tag{1}$$

where P_{t+x} is the closing price x days after the last trade implementing an institutional order and P_{d-1} is the closing price on the day before the order, w_e is the proportion of order shares that actually execute, (1- w_e) is the proportion of unfilled shares, WTP is the volume-weighted trade price of the component trades, order direction is +1 for buys and -1 for sells, and C_t is volume-weighted commission per share.

Clean-up cost of an order has two aspects. It is the product of proportion of the order size that is unfilled and the quantum of return in the stock that is sacrificed by not doing the trade. When an order is completely filled, the clean-up cost of non-execution is zero because the unfilled rate is zero. Similarly, an unfilled order also can also have zero opportunity cost if the stock return after trade execution is zero. Therefore, we simplify our research problem by looking at the fill rates separately before studying the interaction between the two items. The first term in the equation, the return on stock, depends on quality of information and research skills of the equity analyst or portfolio manager. The second term in the equation, the unfilled order rate, depends on the quality of microstructure and trading skills of trading desk manager.

The return dynamics are analogously applicable in the computation of adverse selection cost of filled volume. However, if the institution is better informed than its counterparts, a cost recovery can result. If the portfolio decisions have positive short term alpha on average, filled volume will be associated with at least partial cost recovery. Similar to clean-up cost, the fill rate and return movement interact to generate cost recovery as well, with two computational differences. The first difference is that we use filled volume for cost recovery instead of unfilled volume which is used for clean-up cost computations. The second difference is in the direction of returns. For example, a positive return is a favorable situation for a filled buy in cost recovery calculations but an unfavorable outcome for an unfilled buy in the clean-up cost computation.

Since market-wide returns can influence the measurement of costs, market adjusted costs are computed by deducting market index returns from raw return. For example, for market-adjusted clean-up costs are:

$$Market\ Adjusted\ Clean-up\ Cost = \left\{ \left(\frac{P_{t+x}}{P_{d-1}} - \frac{MI_{t+x}}{MI_{d-1}} \right) * (1 - w_e) \right\} \ * \text{Order Direction}$$

where MI_{d-1} is the level of that index on the day before the order is submitted, MI_t is the index on the day of the last trade of institutional order. The concept is analogously applicable to adverse selection cost and price impact cost but does not apply to commissions.

We analyze the quantity dimension (w_e) of order execution with two alternative measures. The first measure divides the sample into two categories based on whether a given order is completely filled or not. The focus is on the proportion of all orders that are not completely filled, which is obtained simply by dividing the unfilled number of orders by total number of orders submitted. The goal here is to identify any trading strategies or features that are associated with full versus partial execution. This method treats all partially filled orders as equivalent to each other whether the order is filled only 20% or 95% even though the latter is not much different from a completely filled order. We fine-tune the concept in the alternative measure, which is calculated by taking the average unfilled rate of partially filled orders. This measure can help us understand the strategies that help institutions increase the fill rate of the orders, conditional on the fact that they are using execution methods that do not result in complete fill.

The return component of adverse selection cost of filled orders or clean-up cost on unfilled orders is represented by short term price movement. We compute this component with two alternative benchmarks as well. The first measure computes the stock return from one day before the date of the order to twenty days after the last trade implementing that institutional order[2]. This formula represents the transaction cost opportunity loss relative to ideal portfolio assuming that it could be created on the order date itself. From a portfolio perspective, this measure fully captures the implementation shortfall by showing the difference between a paper portfolio and the actual portfolio. This method represents a useful measure of implementation costs, especially for academic studies that form portfolios based on firm characteristics or corporate events and then compute the returns generated by those portfolios. The second alternative measure is based on a similar formulae but it computes the return from the beginning of the last trading day in an order package to twenty days after that date. This alternative method offers a useful measure from a trading perspective. It separates the component of clean-up cost from price impact. In some sense this measure represents money left on the table after considering other major costs that institutions face. This measure assumes that institutions cannot improve their trading performance or lower the price impact that they faced in our sample. Since each method offers its own unique benefit, the combination of the two measures can serve as powerful tool for in-depth analysis of implementation shortfall.

The product of order unfilled rate and lost return represents the clean-up cost of not trading as decided. The proportion of unfilled orders or the unfilled rates is expected to be affected by a variety of microstructure and trading strategy variables. Lost returns on unfilled orders are likely to be affected by quality of research, information environment, and adverse selection problem of trading with other more informed institutions. We also expect that both components of opportunity costs will be affected by firm-specific and order-specific characteristics. We first present the results of our analysis of fill rate (w_e) and then characterize the overall friction which includes all components outlined above.

[2] Later, we show that costs are persistent by varying the interval from 1 day to 60 days after the decision.

A. *Market condition and order direction affect fill rates and transaction costs*

In conjunction with market conditions, order direction can directly affect the fill rates if less aggressive strategies are used. Buys will have higher unfilled rates in bull markets whereas sells will be more difficult to fill in bear markets. The extant literature, especially Chan and Lakonishok (1993) and Keim and Madhavan (1996), also suggests that buys are more informed because institutions choose a few stocks buys out of thousands available after a lot of research. In contrast, sell may be more mechanical based on existing portfolio return goals. Most institutions only sell what they have and do not exploit information by short selling stock that they do not possess. Thus, adverse selection costs of filled orders are expected to be lower or negative (cost recovery) whereas clean-up costs of non-filled volume are expected to be bigger for buys than for sells especially in up markets. As for the explicit costs, Chiyachanta et al. (2004) show that market conditions create asymmetric price impacts with buys (sells) costing more in bull (bear) markets.

B. *Firm size*

Our next variable of interest is firm size. Bigger firms have more trading activity than smaller firms. So, the problem of unfilled order is expected to be more acute for smaller firms due to lack of liquidity. Similarly, big firms are more heavily researched by the entire market leaving little room for information advantage for any particular institution. Smaller firms offer more research opportunities for finding bargains and therefore adverse selection and clean-up costs would be a bigger concern particularly when trading smaller stocks. Likewise, cost recovery for informationally advantaged institutions should be higher in smaller stocks.

C. *Order aggressiveness and trade-off between clean-up cost and price impact*

Order complexity and liquidity issues come to mind as variables that can influence opportunity costs. When an institution is trying to execute an order that is several times the size of average daily volume, filling it completely is naturally going to be difficult. Such voluminous trading activity is also likely to reveal information more quickly to the market resulting in greater amount of clean-up costs compared to smaller orders that can be camouflaged more easily. Aggressive trading with such large orders creates huge price impact costs. Thus institutions would endogenously split the orders across brokers and over time. This passive order splitting strategy reduces price impact but enhances the clean-up costs which are expected to increase with passage of time. This trade-off calls for optimization of order duration. We account for these issues in a simultaneous system of equations described in more details later.

We also categorize momentum orders as those demanding liquidity versus contrarian orders as those that are supplying liquidity. This approach follows Wagner and Edwards (1993) who argue that liquidity characteristics of the order is one of the most important factors affecting transaction cost of institutional order. Liquidity demanding orders pay a higher price impact and are indicative of institutions aggressiveness. Contrarian orders can get a lower price impact and are also more likely to be filled in the institutional trading framework. When everyone is buying, it is much easier to sell resulting in 100% fill rate and vice versa. Since institutions are more likely to demand liquidity and pay a higher price when they possess information, lost returns component of clean-up costs is likely to be higher for the unfilled portion of liquidity demanding orders. Liquidity supplying order indicates institutions patience which comes with a lack of any special information. Therefore, such order may be associated with only marginal clean-up costs, if any. Liquidity supplying order will also earn spread and thus face very low or negative price impact. However, they face adverse selection risk and thus the cost-recovery for filled liquidity supplying orders will not be as good as other orders.

D. Volatility ameliorates fill rates but exacerbates adverse return risk

Stock volatility generates mixed predictions. While higher volatility reduces the probability of the non-execution of limit orders (Ahn, Bae and Chan (2001) and Ellul, Holden, Jain, and Jennings (2007)), it also implies larger return losses for unexecuted orders. Wald and Horrigan (2005) show that the second effect dominates in the context of retail limit orders. If institutional trading philosophy resembles limit order strategy then higher volatility would be associated with bigger clean-up costs on non-filled volume. Volatility may or may not directionally impact adverse selection cost or cost recovery from filled volume but it will enhance all components of transaction costs risk.

E. Order Splitting across brokers and over time; execution venue

The last three variables represent transaction execution methods and properties. Institutions can use floor brokers or market makers who can work the orders slowly to exploit the evolving liquidity conditions. Orders executed over multiple days or using multiple brokers are defined as worked orders. However, the use of more intermediaries also opens the possibility of front running and manipulations. The issue relates to order splitting over time versus across brokers. By continuing to work on an order for several days, the institutions can increase its fill rate while keeping the price impact minimal. But such a strategy can aggravate the lost returns component with the passage of more time. In contrast use of multiple brokers saves time and increases fill rate simultaneously. However, the disadvantage of this approach is that the probability of information leakage increases which can again aggravate the lost returns component. Finally, we distinguish between NYSE and Nasdaq listed stocks. Execution venues differ in their trading mechanisms, which can affect the various dimensions of institutional trading frictions.

II. Data Sources and Research Design

We obtain institutional trading data from the Abel Noser Corporation (hereafter, Abel Noser). The company offers institutional investors goal-oriented trading strategies and trading cost measurement to help the institutions improve their investment performance. The dataset includes details about the investment orders and related purchase and sale transactions by Abel Noser's institutional clients. Abel Noser provides consulting services to 776 domestic clients who collectively transact nearly $20 trillion over period of 1999-2005.

The data provide comprehensive information on institutional trading orders and actual transactions. The variables provided in the dataset include scrambled institutional client code, scrambled institutional manager or trader code, scrambled broker code, scrambled order identifier number, stock ticker symbol, order direction (buy or sell), quantity of shares desired, order placement date, transaction execution date, value-weighted average stock prices (VWAP) on order entry date, VWAP on 1 day prior to order entry date, price at the time of order release, number of shares in the released order, transaction execution price, quantity of shares traded, commissions charged, and type of institution (Mutual Fund or Pension Fund) executing orders. The data are provided to us after removal of the actual names of the managers involved to maintain client anonymity and privacy. To ensure the integrity of the data and filter out possible errors, we eliminate observations with missing prices or order quantities. Many of Abel Noser's institutional clients provide order volume and transaction data for analysis by Abel Noser. We focus on 4 million orders for which order volume can be accurately identified by checking whether the order identification field is populated. The remaining observations without any order identification will still have order volume which is either aggregated or algorithmically generated by Abel Noser. For the most part, we exclude such orders from our analysis, except to carry out of sample robustness of our results. In addition, following the approach of Keim and Madhavan (1995, 1997), we exclude orders that took longer than 21 calendar days to complete.

We merge institutional trading data with CRSP to obtain stock specific information and value-weighted market index. These indices help us to control for market-wide returns. For instance, if the market index rose significantly on a given day, then all purchases, whether institutional or retail, may have more positive price impact for purchases and perhaps negative price impact for sells. Therefore, we conduct the analysis of both the raw transaction costs and market-adjusted costs.

As shown in Table I, our final sample contains about 4 million institutional orders with order identification and client-supplied order volume. The aggregate share volume of these orders is 148 billion and the dollar volume is 4.59 trillion dollars. Of these 3.2 million orders with 44 billion shares worth $1.49 trillion are completely filled. However, the remaining 771,861 orders with 104 million submitted shares are partially filled. Orders that are partly filled represent much larger order size and aggregate order volume. However, only 24 million shares representing 23% of submitted shares are filled and the remaining 80 million shares are unfilled. In aggregate, 46% of the total volume of shares volume submitted (and 48% of the dollar volume) in the full sample is filled and the remaining 54% shares represent the non-filled volume.

[Insert Table I about here]

III. Quantity dimension: Order fill rate and its determinants

A. Order fill-rate variations related to firm and trade-specific characteristics

Now we set out to understand the distribution and determinants of order fill rates. Figure I reports the proportion of filled and unfilled volume in various categories based on firm-specific and trade-specific characteristics. Fill rates are number of shares traded divided by number of shares submitted. As discussed above, the overall fill rate is 46%. The figure also decomposes the filled volume into two portions. The portion of filled volume coming from fully filled orders is 30% and is shown in green and that coming from partly filled orders is 16% as displayed in blue. Unfilled volume of 54% is shown in brown.

In terms of order direction, the fill rates for purchases are marginally higher at 47% than for sells at 44%. Since buyers may be more informed than sells (Chan and Lakonishok (1993) and Keim and Madhavan (1996)), they may be trading more aggressively leading to slightly higher fill rates.

Next, we allocate stocks into 3 groups with equal number of stocks based on their market capitalization. Firm size partitions reveal that trading activity in terms of absolute number of orders and share volume of both filled and unfilled orders is highest for large market capitalization stocks. However, the proportion of filled orders is marginally higher for small stocks at 50% followed by large and medium capitalization stocks at slightly below 50%. Small stock orders also appear to be smaller and fully filled orders contribute a much bigger proportion of traded volume. Institutional traders in small stocks could be more informationally advantaged and may be trying lower their unfilled rates, the lack of available liquidity might balance their aggressiveness effectively putting a cap on empirically observed fill rates.

Following Chiyachantana et al (2004) we compute the complexity of order, defined as number of shares in an institutional order divided by the average daily share volume over the previous 5 days obtained from CRSP. Then we divide the total number of orders into terciles based on order complexity. As expected, easy orders have the highest fill rate of 92% which is more than twice as much as the 44% fill rate of difficult orders. Majority of easy orders are fully filled. The contribution of partly filled orders in trading volume is negligible for easy orders but substantial for difficult orders. The difference between fill rates of easy and difficult orders is even more dramatic when we look at dollar value, which make it clear that difficult orders are the dominant source of unfilled volume and clean up costs.

Next we split the sample based on liquidity characteristics of the order. We define an order to buy (sell) on the day when stock return is positive (negative) as a liquidity demanding order. Conversely, purchase (sell) orders submitted at a time when the prices are falling (rising) are defined as liquidity supplying orders. Our approach is similar to Wagner and Edwards (1993). Fill rates for liquidity supplying orders at 49% is marginally higher than 44% but the difference between their clean-up costs is much more dramatic, as we discuss in next section.

Stock volatility is another important firm characteristic that can affect fill rates. We measure volatility as percentage difference in highest and lowest trading price in the one month preceding the institutional trading order. For the three categories of high, medium, and low stock volatility, the cut-off points were determined such that each category has equal number of orders. Fill rates are comparable for high, medium, or low volatility stocks.

The last two variables represent transaction execution methods related to order splitting. We divide the orders into two groups based on whether or not the orders were split over multiple days or using multiple brokers or both. A vast majority of low volume orders are executed wihin a day with one broker, many high volume and high value order are split across brokers or over time. As a result, a smaller number of split orders command a bigger share of filled and unfilled order volume and value. Orders split over multiple days have a dramatically lower fill rate of 31% compared to that of 65% for single day executions. This difference is likely to reflect the more difficult nature of implementation of such orders forcing the institutions to break up their orders. We account for this potential endogeneity in our subsequent analysis by estimating a simultaneous system of equations where order duration is a function of order complexity among other control variables.

Finally, we split the sample between orders using single versus multiple brokers. Vast majority of orders are implemented using a single broker and such orders dominate share volume and value as well, both for filled and unfilled orders. Nevertheless, many big orders are split across multiple brokers. Order fill rates of 85% for multiple broker orders are almost twice that of single broker executions. Thus, multiple broker order splitting and multiple days order splitting have opposite effects. We analyze the difference between the two types of order splitting in greater details in our subsequent analysis.

[Insert Figure I about here]

B. Regression analysis of determinants of fill rates

In Table II, we capture the incremental effect of each firm-specific and order-specific characteristic on fill probability and fill rates in a multivariate regression setting. Analysis is conducted at two levels -- all institutional orders and partly filled orders. For all orders, probit regression is estimated to understand the fill probability and OLS regression is estimated to understand the fill rates. In the probit regression, we estimate the likelihood of an order being filled, following the suggestion in Wald and Horrigan (2005). The dependent variable is equal to one for completely filled order and zero for orders that are not completely filled. This regression is based on 4 million observations and has a pseudo R-squared of 7.20%. In the OLS regression with the full sample of the same 4 million, the dependent variable is the actual fill rate for each order. Fully filled order have 100% fill rate, partly filled order can have the fill rate from 0.01% to 99.99%, and orders without any trading activity have 0% fill rate. R-squared for this regression is 10.97%. The last regression in the table is based on 771,861 partly filled orders only; the dependent variable in this OLS regression is again the actual fill rate. R-squared is 5.46% in this regression. Large sample sizes lead to powerful statistical tests and highly significant coefficients.

The estimated coefficients from the first regression show that, after controlling for other factors, the probability that an order will be implemented fully decreases with order complexity, liquidity demanding trading strategy, order duration, and bullish market sentiment. The probability of a complete fill increases with market capitalization, higher stock volatility, usage of more brokers, positive market returns, and Nasdaq listing. The direction and significance of the coefficients in the fill rate OLS version in the second regression are identical to those in the probit regression.

Coefficients from the third regression based only on partly filled orders again have similar direction and significance of coefficients with the exception of Nasdaq listing indicator variable which takes a negative sign and firm size and market returns which are not significant. All other coefficients in the three regressions are statistically significant.

Of course, the complete interpretation of these regression results in terms of effectiveness of institutional trading needs to await the adverse selection (or cost recovery) and clean-up cost analysis presented in the next section.

[Insert Table II about here]

IV. Composite transaction costs and their determinants

A. Summary statistics of transaction costs

Table III provides the market-adjusted estimates for overall friction and its various components.[3] Filled orders are complete and do not have any clean-up costs. Partially filled orders give rise to clean-up costs of 64 basis points based on 20 days after trade benchmark. Average clean-up costs for all orders are 43 basis points or $19.88 billion in our sample. Adverse selection costs from fully filled orders are -74 basis points. Negative costs imply that institutions in our sample are better informed than their counterparties. Thus, they are able to recover some of the costs through their informational advantage. Partly filled orders also do not have any positive adverse selection costs but the cost recovery from such orders is nominal. On average, cost recovery of 24 basis points implies that institutions are able to recoup $11 billion dollar of their costs through short term returns. Price impact costs consume significantly into the gross returns on a stock from order time to execution time. For fully filled order, price impact is 97 basis points and for partly filled orders it is 16 basis points. Overall price impact in the full sample is 42 basis points or $19.48 billion. Finally, commissions are 5 basis points or $2.33 billion. The total all-in execution costs are 67 basis points summing up to $30.68 billion in our sample. All these are one way costs while buying or selling shares.

[Insert Table III about here]

Figure 2 shows the respective shares of various components of institutional trading friction. Price impact represents the biggest component and accounts for 63% of total costs. Gross clean-up costs at 43 basis points are actually higher than price impact but institutions are

[3] Raw numbers not adjusted for market are very similar to the market-adjusted numbers.

able to recover part of these hidden costs through the information advantage on their filled orders. Thus the net-clean up costs remain 19 basis points which is still a significant 29% of total trading costs. Price impact and clean-up costs are several times larger than the explicit commission costs that account for only 8% of the total friction.

[Insert Figure 2 about here]

B. Transaction cost variations related to firm and order-specific characteristics

In Table IV, we assess the magnitude of total transaction costs and its four components in the previously defined seven firm-specific characteristics and order-specific categories. The first four columns present the trade-weighted (or non-trade-weighted) average cost based on individual trades. The fifth column shows the average fill rate. The last six columns show the effective costs at the order level after adjusting (weighting) them with the fill rates (or one minus fill rates as applicable) as shown previously in equation 1. At the level of each individual order, the unadjusted trade level cost times the fill rates (or one minus fill rates) equals the order level weighted costs. However, because of differences and uniqueness in order volume, fill rates, and transaction volumes for each order, the summary statistics at the aggregate sample level presented in the table do not yield themselves to that equality. This difference highlights the importance of carrying out the analysis using complete order information. The past literature is mostly based on unadjusted trade level costs. The additional knowledge of simply the average fill rate is not sufficient to learn about the order level frictions because of the significant variations in fill rates of individual orders. Weighted costs are reported throughout the paper and used for majority of our analysis for they provide the best available measure of institutional trading frictions at order level.

The average adverse selection costs are negative for small, medium, and large capitalization stocks and all subsequent partitions at both trade and order levels. Thus, institutions in our sample appear to be better informed than their counterparties for all firm sizes. They use this informational advantage to recover part of the costs arising from trading frictions.

Clean-up costs of non-filled volume as well as price impact and commissions are positive in all market capitalization categories and significantly higher for small stocks than for large stocks. Total execution costs are 111 basis points for small stocks and 61 basis points for large stocks. Small stocks also have the highest transaction cost risk as measured by its standard deviation. However, with a lion's shares of filled and unfilled orders, large capitalization stocks account for the bulk of total dollar transaction costs at $22.33 billion.

Complexity of order continues to be a key driver of institutional frictions just like it was the major driver of order fill rate. Cost recovery is low, and clean-up and price impact are highest for difficult orders. Weighted clean-up cost for difficult orders at 45 basis points is a striking 232 times of the clean-up cost for easy orders, which have negligible clean-up costs. The difficult orders face $31 billion in total costs which account for over 99% of the total institutional trading frictions.

In our next partition, liquidity demanding orders emerge as a major source of clean-up costs even though the proportions of unfilled orders were fairly similar between liquidity demanding and liquidity supplying orders. Of course, the definition of order type for this partition can directly lead us to the observed result. Liquidity supplying orders, in fact, have a negative clean-up cost. The reason is fairly obvious. Liquidity supplying orders are defined as sell orders in up markets and buy orders in down markets. Is such orders are not executed and the market continues its move in the same direction then there is no opportunity loss because one would be able to sell higher or buy lower later on. However during market reversals, liquidity supplying orders would have higher clean-up costs than liquidity demanding orders. The overall transaction costs are negative for liquidity supplying trades. Thus, a contrarian trading strategy appears to earn net positive rents of $19 billion from the business of supplying liquidity. In contrast, liquidity demanders face $50 billion worth of trading frictions.

In the stock volatility partition, high volatility stocks have clean-up costs and total frictions nearly three times higher than that for low volatility stocks. More important, volatility increases transaction cost risk as seen in the standard deviation column.

Our next partitions represent transaction execution duration. Multi day execution of orders is associated with clean-up costs that are nearly six times that for single day orders. After accounting for cost recovery, single day orders end up with net negative costs. Thus, multi-day execution ends up with bulk of the frictions. Of course, if institutions had tried to execute those orders on a single day, perhaps they could have faced astronomical costs. So the appropriate interpretation of our result is that order splitting may help reduce transaction costs but yet it may not be sufficient to make them comparable to the easy single day executions. Despite the higher percentage total costs, multiple broker orders account for only one third of the dollar transaction costs because vast majority of the orders are completed using a single broker. Given that the execution duration can be endogenously determined by the institutions, we later estimate a simultaneous system of equation where duration is a choice variable.

[Insert Table IV about here]

C. Persistence of opportunity costs

In Figure 3, we conduct a robustness test of opportunity cost results by altering the measurement period. In addition to the 20 day period used in equation (1), we now consider periods of 1 to 60 days after the completion of last transaction in an order package. For brevity results are shown only for a few sub-samples based on complexity and firm size but results are similar in other categories as well.

For the overall sample, clean-up costs begin with 55 basis points one day after transaction period. Thereafter, it remains persistent and range bound between 29 and 55 basis points with the average of 39 basis points that compares well with the 20th day reported value of 43 basis points used and presented throughout the paper. We observe the persistence and even tighter ranges of clean-up costs within the order-complexity sub-categories of difficult and easy orders over time. Similarly, large capitalization and small capitalization stocks have persistent differences throughout the 60 day robustness analysis period. Overall, the analysis demonstrates that opportunity costs are highly persistent across time and the differences across the various sub-samples are stable irrespective of the measurement period.

[Insert Figure 3 about here]

D. Order Duration

The aggressiveness with which the institutions complete their transaction reflects in the duration of an order. Duration, in turn, can have important bearings on the various components of institutional trading costs. In this section we characterize the choice of duration for various types of orders. In the next section duration is used as an endogenously determined variable which in turn affects the transaction costs in a simultaneous system of equations.

Figure 4 plots order duration and order volume for firm and order specific categories. For the overall sample, 25th percentile of order duration is 1 day, mean is 1.19 days, 90th percentile is 1 day and 99th percentile is 6 days. We truncate the sample to contain orders completed 21 days. So the 99th percentile is more meaningful to analyze as an indicator of maximum duration. The mean duration of 1.21 days for purchases is just marginally higher than 1.18 days for sells.

Market capitalization categories present an interesting fact about order duration. Small (large) stocks have the lowest (highest) aggregate volume of 7 million (73 million) shares but highest (lowest) mean order duration of 1.40 days (1.11 days). This pattern arises easily once we account for low market-wide volume in small stocks. Thus the appropriate yardstick for understanding order duration is order complexity which divides the order volume with average daily volume in the last 5 days for a given stock. Although there are equal number of orders (1.33 million each) in easy, medium, and difficult categories, bulk of the volume comes from difficult orders. Easy orders have maximum duration of 1 day but account for negligible volume. Difficult orders with 98% of the volume have mean duration of 1.55 days and 99th percentile of 11 days. Given the importance of order complexity, it is used as a key variable in the duration regression in the next section.

In the next partition, liquidity demanding orders represent both a higher volume and higher mean duration of 1.25 days relative to liquidity supplying orders. For the partition based on stock volatility, high volatility stocks account for a larger proportion of order volume and also a longer order duration. The main of the next partition based on single or multiple days is to show that long duration orders account for bulk of the share volume. Finally, multiple broker orders are also associated with durations that are almost twice that of single broker orders. However, bulk of the volume utilizes a single broker.

E. Regression analysis

Several important factors have emerged in the discussion above as potential determinants of opportunity costs. In Table V, we capture the incremental effect of each firm-specific and order-specific characteristic on opportunity costs in a multivariate regression setting based on a simultaneous system of equation. The first stage regression estimates show that institutions split the execution of their orders more i.e. increase their order duration with small stocks, order complexity, stock volatility, bullish market conditions, and non-index stocks. This implies that orders in big stocks, easy orders, orders in less volatile stocks, and those in index-components are executed quickly.

We now discuss the determinants of aggregate institutional trading friction and its components. The coefficient on market-capitalization is negative in each regression. Thus, institutional trading friction and each one of its components is higher in small stocks and reduces as the size of stock increases. Order complexity or higher relative volume is associated with increased hidden components of adverse selection and clean-up costs. Although institutions appear to be controlling and reducing the observable components of price impact and commissions by increasing the duration of complex orders, the overall friction remains slightly higher for complex orders after controlling for endogenously determined duration. Next we focus on stock volatility. Although higher stock volatility leads to bigger clean-up costs, price impact, and commissions, institutions are net beneficiaries of increased volatility because of increased level of cost recovery from information based trading in volatile stocks. The cost recovery dominates other components and leads to negative impact on overall costs as well, for highly volatile stocks.

The negative coefficient for liquidity demander in the adverse selection regression is consistent with the notion that when institutions have an informational advantage then they trade more aggressively and demand liquidity and finish the trade with substantial cost recovery. Whereas filling such orders is rewarding for the institutions, the clean-up cost for missing the trades is also very high. Liquidity demanding orders also have a huge price impact, which eventually dominates and creates high overall frictions for such aggressive orders.

Order splitting over time has diagonally opposite effects of various components of trading costs. Theoretically, the purpose of increasing duration is to reduce price impact but this strategy can also increase the clean-up costs. The importance of endogenously determined order duration and the contribution of our study in that respect are evident from positive coefficient for both price impact and hidden clean-up cost, which present a trade-off for the institutional traders. Orders with higher values of predicted duration are associated with negative adverse selection or cost recovery. Previous studies have focused on direct costs such as price impact and commissions and found positive association between duration and those cost, like we do too. After controlling for everything else in the model, the overall friction is still increasing with duration. Institutional decisions to split orders over time must be based on careful consideration of the trade-off between various components of trading costs.

Another dimension of order splitting is the use of multiple brokers. The insignificant coefficient on number of brokers in the adverse selection regressions points to the possibility that institutions cannot maximize their informational advantage and cost recovery simply by employing multiple brokers. However, the negative coefficient on clean-up costs points to the benefit of using multiple brokers, who help increase the fill rates and lower the lost returns from unexecuted orders. Yet, the positive and dominant coefficient on price impact implies that use of multiple brokers fails to reduce the price impact of large orders and as a result the overall friction increases with the use of multiple brokers. No wonder that multiple brokers are used for a very small portion of overall order volume as seen previously in figure 4.

Next, we have the comparison between Nasdaq and NYSE. Institutions appear to have a higher informational advantage in Nasdaq stocks relative to NYSE stocks as implied by the higher cost recovery (negative coefficient of Nasdaq-listing in the adverse selection regression). Conversely, missed trades in Nasdaq stocks represent a bigger opportunity loss of clean-up cost. Price impact and commissions are higher in Nasdaq stocks. The informational advantage dominates and leads to an overall friction which is lower on Nasdaq than on NYSE.

Finally, we control for the amount of analyst coverage. Any institution's informational advantage could diminish if a greater number of analysts are covering a stock. Positive coefficient for analyst coverage in adverse selection and clean-up cost regressions reflect this notion and imply that institutions do not receive any cost recovery in widely covered stocks. Price impact is lower because counterparties are less worried about any institutions informational advantage for very widely covered stocks.

[Insert Table VI about here]

The regression analysis helps us identify situations that can result in higher institutional trading frictions. However, the focus thus far has been on average costs. In the next section we present additional analysis focusing on variance in the level of friction to understand transaction cost risk.

F. Transaction cost risk

The average opportunity costs statistics provide very good benchmarks for overall institutional performance in a repeated trading setting. However, any single order carries risk because its transaction cost can be very different from the average. For the overall sample, we presented the standard deviation earlier in Table IV of 8.6 basis points and also concluded that risk was higher for small and highly volatile stocks. We now provide additional insight into transaction cost risk issue in separately examining the hidden components of friction in Figure 5. In each Panel, the x-axis captures the transaction cost variation by forming cost categories with one percent intervals. We consolidate the extreme categories by clubbing the orders that have less than -10% or more than 10% cost. The y-axis plots the proportion of all orders with the category plotted that fall in the transaction cost category. For example, in Panel A which is based on firm size, 15% of all large stock orders have opportunity costs ranging between 0% and 1%. In contrast, only 7% of all small stock orders have such low opportunity costs. Less than 1% of large stock orders have clean-up costs exceeding 20% whereas more than 4% of small stock orders have those exorbitant costs. Patterns are similar for adverse selection cost (or cost recovery). Thus, the narrower bell shaped curve for large stocks and a flatter curve for small stocks indicate that small stocks carry a more severe transaction cost risk.

In Panel B, we focus on variance in opportunity cost conditional on liquidity provision. Liquidity supplier and liquidity demander orders have similar variance. However, we can see that the costs are slightly asymmetric. Liquidity supplier curve tilts to the left demonstrating that a higher proportion of liquidity supplier orders lower clean-up costs. But the same liquidity supplier category is skewed to the right for adverse selection costs. From both graphs, we can infer that liquidity demanders are more informed than liquidity suppliers.

Finally, we analyze the implications of order splitting over time.. Previously we established in Tables IV and VI that multiple day orders have higher clean-up costs and lower cost recovery. Now we show in Panel C of figure 5 that long duration orders also have higher clean-up cost risks and lower cost recovery variance.

V. Conclusions

Institutional money managers such as mutual funds typically transact large volumes of shares to implement their portfolio investment strategies. The nature of their activity often results in large transaction costs that can undermine their performance by creating significant implementation shortfall. One of shortfall's components, the execution cost, is fairly well understood. It relates to the transactions you actually execute and arises from bid-ask spread, price impact, commission, and other transaction fees and taxes. We analyze additional hidden components of institutional friction, namely, adverse selection costs of filled volume and clean-up costs of unfilled order volume and find that the opportunity cost of inaction is very high for institutional orders.

We present a comprehensive analysis of these various components of friction including the often ignored component of implementation shortfall -- the clean-up cost of unexecuted institutional orders. Based on client supplied information about institutional order volume, only 46% of the submitted volume results in trade and the remaining 54% of the submitted volume is unfilled. Both the probability that an order will be implemented fully and the fill rate decrease with buy order direction, order complexity, liquidity demanding trading strategy, and order duration; the fill probability and fill rate increase with market capitalization, higher stock volatility, usage of more brokers, and positive market returns.

We integrate the fill rates with post-trade short term returns to fully understand the nature of institutional trading frictions. Adverse short term returns create hidden clean-up costs for non-filled volume and observable adverse selection costs for filled orders. In contrast favorable returns can potentially result in a cost recovery for successfully filled orders. The institutions in our sample appear to have informational advantage and generate a negative adverse selection or cost recovery of 24 basis points on filled volume and face clean up costs or lost returns of 43 basis points resulting from the unfilled volume. The clean-up cost is the biggest component of institutional frictions followed closely by price impact of 42 basis points which dwarf the explicit commissions of 5 basis points. Overall frictions in the full sample are 66 basis points with a standard deviation of 8.6 basis points. To analyze the determinants of institutional friction and its components, we estimate a simultaneous system of equation to account for endogenous order duration choice and the effects of such order splitting over time on frictions. Institutions seem to increase their order duration with order complexity, stock volatility, bullish market conditions, and non-index stocks.

Institutional trading friction and each one of its components are higher for small stocks, and complex high volume orders. Wider analyst coverage reduces the informational advantage of institutions and increases both adverse selection and clean-up costs. The remaining determinants have opposite effects on adverse selection and clean-up costs. For example, adverse selection components increases (or cost recovery decreases) with use of multiple brokers whereas clean-up costs increase with volatility, liquidity demanding trading strategy, order duration and Nasdaq listing. The total friction increases with order complexity, liquidity demanding trading strategy, and use of multiple brokers and it decreases with firm size, volatility, order duration, and

Nasdaq-listing. Often the reduced friction is a result of cost recovery possibly through institutional informational advantage.

We demonstrate that the hidden adverse selection and clean-up costs are persistent by analyzing the benchmark periods ranging from ten days after the unexecuted order to sixty days after the last transaction in the order. We also assess the transaction cost risks by computing standard deviation of 8.6 basis points for total transaction costs. Small stocks and high volatility stocks stand out as high transaction cost risk categories. We also present the overall distribution of adverse selection costs and clean-up costs for various sub-samples.

The results have several practical and academic implications. Institutional investors can use them as benchmarks to analyze their own implementation shortfall. The numbers can also provide guidance of whether or not it pays to be aggressive in completing a large institutional order. More importantly, the implementation policy can be customized to address the affect of market conditions, firm-specific characteristics, and order-dynamics. Although we focus on institutional trading in equities, future studies can examine if the results can be generalized to other trading situation and asset classes. From the academic perspective, the concept of total transaction should include not only the explicitly observed spreads, price impact and commissions but also the hidden opportunity costs such as adverse selection and clean-up costs. These expanded measures of total transaction costs provide the limits to arbitrage and also implore models to include the transaction costs because they can lead to significant deviations from the ideal performance of a paper portfolio that we so often see in theoretical and empirical papers.

References

Ahn, H.-J.; K. Bae; and K. Chan, 2001, Limit orders, depth and volatility: Evidence from the Stock Exchange of Hong Kong, *Journal of Finance* 56, 769-790.

Almgren, Robert and Neil Chriss, 2000, Optimal Execution of Portfolio Transactions, *Journal of Risk* 3, 5-39.

Berkowitz, Stephen A., Dennis E. Logue, and Eugene A. Noser, 1988, The total cost of transaction on the NYSE, *Journal of Finance* 97-112.

Bertsimas, Dimitris, and Andrew W. Lo, 1998, Optimal Control of Execution Costs, *Journal of Financial Markets* 1, 1-50.

Carrie, Carl, 2008, Illuminating the new dark influence on trading and U.S. market structure, *Journal of Trading* 3, 40-55.

Chan, Louis K.C., and Josef Lakonishok, 1993, Institutional trades and intra-day stock price behavior, *Journal of Financial Economics* 33, 173–200.

Chan, Louis K.C., and Josef Lakonishok, 1997, Institutional equity trading costs: NYSE versus Nasdaq, *Journal of Finance* 52, 713-735.

Chakravarty, Sugato, Venkatesh Panchapagesan, and Robert Wood, 2004, Has Decimalization Hurt Institutional Investors?, *Journal of Financial Markets* 8, 400–420.

Chemmanur, Thomas J. and Gang Hu, 2007, Institutional Trading, Allocation Sales, and Private Information in IPOs, EFA 2006 Zurich Meetings Paper Available at SSRN: http://ssrn.com/abstract=959645

Chiyachantana, C., P. Jain, C. Jiang, and R. Wood, 2004, International evidence on institutional trading behavior and determinants of price impact, *Journal of Finance* 59 (2), 865-894.

Conrad, Jennifer S., Kevin M. Johnson, and Sunil Wahal, 2001, Institutional trading and soft dollars, *Journal of Finance* 56, 397-417.

Conrad, Jennifer S. and Sunil Wahal, 2003, Institutional trading and alternative trading systems, *Journal of Financial Economics* 70, 99-134.

Domowitz, Ian, Jack Glen, and Ananth Madhavan, 2001, Liquidity, volatility, and equity trading costs across countries and over time, *International Finance,* 4, 221-255.

Ellul, Andrew, Craig Holden, Pankaj K. Jain, and Robert Jennings, 2007, Order Dynamics: Recent Evidence from the NYSE, *Journal of Empirical Finance* 14, 636–661.

Engle, Robert , and Robert Ferstenberg. 2007, Execution Risk, Journal of Portfolio Management 33, 34-45

Goldstein, Michael, Paul Irvine, Eugene Kandel, and Zvi Wiener, 2009, Brokerage Commissions and Institutional Trading Patterns, Forthcoming in *Review of Financial Studies*.

Handa, Puneet and Robert A. Schwartz, 1996, Limit Order Trading, Journal of Finance 51, 1835-1861.

Jones, Charles M., and Marc L. Lipson, 2001, Sixteenths: direct evidence on institutional execution costs, *Journal of Financial Economics*, 59, 253-278.

Keim, Donald B., and Ananth Madhavan, 1995, Anatomy of the trading process: Empirical evidence on the behavior of institutional traders, *Journal of Financial Economics* 37, 371-398.

Keim, Donald B., and Ananth Madhavan, 1996, The upstairs market for large-block transactions: Analysis and measurement of price effects, *Review of Financial Studies* 9, 1-36.

Keim, Donald B., and Ananth Madhavan, 1997, Transaction costs and investment style: An inter-exchange analysis of institutional equity trades, *Journal of Financial Economics* 46, 265-292.

Kyle, Albert S., 1985, Continuous auctions and insider trading, 1985, *Econometrica* 53, 1315-1335.

Lipson, Marc L., and Andy Puckett, 2007, "Institutional Trading During Extreme Market Movements" (July 2007). Available at SSRN: http://ssrn.com/abstract=687414

Perold, Andre F., 1988, "The implementation shortfall: Paper versus reality," *Journal of Portfolio Management* 14, 4-9.

Perold, Andre F., and Erik R. Sirri, 1997, The cost of international equity trading, *Working Paper 97-012, Harvard Business School- Research Division*.

Puckett, Andy and Xuemin Sterling Yan, 2008, Short-term Institutional Herding and Its Impact on Stock Prices, Available at SSRN: http://ssrn.com/abstract=972254

Stoll, H.R. "Friction." *Journal of Finance*, 55 (2000), pp. 1479-1514.

Wagner, Wayne H., and Mark Edwards, 1993, Best execution, *Financial Analysts Journal* 49, 65-71.

Wald, John K., and H.T. Horrigan, 2005, Optimal Limit Order Choice, Journal of Business 78, 597-619.

Table I. Sample Description of Institutional Orders and Trades

Data are from Abel Noser and represent 776 institutional clients who collectively transact nearly $20 trillion over period of 1999-2005. Many of these institutions provide order volume and transaction data for analysis by Abel Noser. We use only 4 million such observations representing nearly $5 trillion in trades. It is essential to restrict the sample to these orders with customer-supplied order volume to accurately calculate the filled and unfilled volume. These good orders for our purpose can be identified by checking whether the order identification field in the database is populated. The remaining orders which we do not consider good for our purpose have algorithmically computer order volume but their order identification field is blank.

Overall Sample:	
Number of orders with client-supplied volume statistics	3,976,387
Order share volume (Billion shares)	147.99
Order amount (Trillion $)	4.59
Overall percentage of volume filled (traded)	45.69%
Overall percentage of dollar amount filled (traded)	48.17%
Fully Filled Orders:	
Number of orders completely filled	3,204,526
Order share volume (Billion shares)	43.88
Order amount (Trillion $)	1.49
Fill rate (based on share volume)	100%
Partly Filled Orders:	
Number of orders partly unfilled	771,861
Order share volume (Billion shares)	104.11
Order amount (Trillion $)	3.09
Number of shares traded (Billion)	23.74
Traded amount (Trillion $)	0.72
Fill rate within partly filled orders	23%

Table II. Probit and OLS Regressions: Determinants of Order Completion

In the first column, we present estimates of Probit regressions with the full sample where the dependent variable is equal to one for completely filled order and zero for incomplete order. In the second column, we present the results of an ordinary least square (OLS) regression where the dependent variable is fill rate. Fill rate is shares traded divided by shares submitted. In the last column a similar OLS regression is estimated within the sub-sample of partly filled orders only. Explanatory variables include the natural logarithm of market capitalization of the firm in dollars, complexity of order which is calculated as the ratio of order shares relative to average daily trading volume over the prior five trading days, liquidity demander which takes a value of 1 if order to buy (sell) is made when stock return on order date is positive (negative), stock volatility which is calculated as percentage difference in highest and lowest trading price in the past 30 calendar days prior to institutional trading order, duration in numbers of days elapsed from the date of order to the date of final trade for that order package, number of brokers involved in the trades pertaining to the particular order, market return based on the CRSP value weighted index from the day of institutional trading order to the last trading day, and dummy for Nasdaq listing which takes the value of 1 if stock is listed on Nasdaq and 0 if it is on the NYSE. Statistical significance is indicated by ** for one percent level and * for five percent level.

	Probit Regression		*OLS Regressions*			
Order Characteristics	**Probability of Fully Filling a Order**		**All Orders**		**Partly Filled Orders**	
Intercept	0.196	**	0.826	**	0.318	**
Market Capitalization	0.011	*	0.017	**	0.039	
Complexity of Order	-0.011	**	-0.008	**	-0.086	**
Liquidity Demander	-0.061	**	-0.013	**	-0.013	**
Stock Volatility	0.056	**	0.012	**	0.034	**
Duration	-0.401	**	-0.086	**	-0.030	**
Number of Brokers	0.793	**	0.070	**	0.050	**
Market Returns	-3.392	**	-0.573	**	-0.012	
Nasdaq Listing	0.030	**	0.002	**	-0.013	**
Number of Observations	3,976,387		3,976,387		771,861	
Pseudo R-square	7.20%					
Adjusted R-square			10.97%		5.46%	

Table III. Institutional Transaction Cost and Its Components

Total transaction costs of implementing institutional portfolio orders comprise of:
Friction={Clean-up cost + Adverse selection cost + Price impact}* Order direction + Commission
OR Cost Recovery

$$PQRST = \left\{ \left(\frac{P_{t+x}}{P_{d-1}} - 1 \right) * (1 - w_e) + \left(1 - \frac{P_{t+x}}{WTP} \right) * (w_e) + \left(\frac{WTP}{P_{d-1}} - 1 \right) * w_e \right\} * \mathrm{OD} \qquad + \left(\frac{C_t}{P_{d-1}} - 1 \right) * w_e \tag{1}$$

where P_{t+x} is the closing price 20 days after the last trade implementing an institutional order and P_{d-1} is the closing price on the day before the order. w_e is the proportion of order shares that actually execute, (1- w_e) is the proportion of unfilled shares, WTP is the volume-weighted trade price of the component trades, order direction is +1 for buys and -1 for sells, and C_t is volume-weighted commissions per share.

Market adjusted costs are computed by deducting market index returns from raw return. For example, for market-adjusted clean-up costs are:

$$Market\ Adjusted\ Clean-up\ Cost = \left\{ \left(\frac{P_{t+x}}{P_{d-1}} - \frac{MI_{t+x}}{MI_{d-1}} \right) * (1 - w_e) \right\} \quad * \mathrm{Order\,Direction}$$

where MI_{d-1} is the level of that index on the day before the order is submitted, MI_t is the index on the day of the last trade of institutional order. The concept is analogously applicable to adverse selection cost and price impact cost but does not apply to commissions.

Dollar trading costs are obtained by multiplying each component of trading cost to the dollar value of institutional order.

Trading Cost Components	Market-Adjusted Institutional Trading Cost (%)			
	Fully Filled Orders	**Partially Filled Orders**	**All Orders**	**Institutional Trading Cost (Billion $)**
Clean-up Cost	0.00	0.64	0.43	$19.88
Adverse Selection Cost (Cost Recovery)	-0.74	-0.01	-0.24	-$11.01
Clean-up Cost Net of Cost Recovery	-0.74	0.63	0.19	$8.87
Price Impact Cost	0.97	0.16	0.42	$19.48
Commission Cost	0.09	0.03	0.05	$2.33
Total Execution Cost	0.33	0.83	0.67	$30.68

Table IV. Variations in Transaction Costs based on Order Characteristics and Implementation

Market-adjusted institutional transaction costs are presented in percent. Partitioning factors for samples are the same as defined previously in Figure 1. The first four columns present the trade-weighted (or non-trade-weighted) average component costs based on individual trades. The fifth column shows the average fill rate. The next four columns show the effective costs at the order level after adjusting (weighting) them with the fill rates (or one minus fill rates as applicable) as shown previously in equation 1. Total execution costs are the sum of adverse selection, clean-up, price impact and commissions. The last column is the standard deviation of costs across orders.

Order Characteristics		**Adverse Selection Cost (Cost Recovery)**	**Clean-up cost of non-filled volume**	**Price impact of filled volume**	**Commission on filled volume**	**Fill rate**	**Weighted Adverse Selection Cost (Cost Recovery)**	**Weighted Clean-up cost of non-filled volume**	**Weighted price impact**	**Weighted commission**	**Total Execution Cost**	**Standard Deviation of Total Execution Costs**
All Orders		-0.49	0.86	0.85	0.11	46%	-0.24	0.43	0.42	0.05	0.669	0.086
Order Direction	*Purchase*	-0.17	0.33	0.81	0.10	47%	-0.08	0.19	0.33	0.05	0.487	0.085
	Sell	-0.78	1.29	0.89	0.11	45%	-0.38	0.64	0.51	0.05	0.824	0.088
Market Capitalization	*Small*	-1.02	2.28	0.81	0.22	50%	-0.53	1.06	0.47	0.11	1.108	0.114
	Medium	-0.55	1.30	0.82	0.14	43%	-0.26	0.67	0.41	0.06	0.888	0.095
	Large	-0.46	0.71	0.86	0.09	46%	-0.22	0.36	0.43	0.05	0.606	0.076
Complexity of Order	*Easy*	-0.53	0.02	0.02	0.06	92%	-0.49	0.00	0.02	0.06	-0.410	0.078
	Moderate	-0.43	0.14	0.04	0.07	80%	-0.35	0.03	0.03	0.06	-0.233	0.087
	Difficult	-0.50	0.87	0.91	0.11	44%	-0.23	0.45	0.44	0.05	0.709	0.093
Liquidity	*Supplier*	-0.17	-1.11	-0.93	0.10	48%	-0.08	-0.56	-0.47	0.05	-1.054	0.085
	Demander	-0.73	2.10	2.14	0.11	44%	-0.35	1.10	1.02	0.05	1.825	0.087
Stock Volatility	*High*	-0.81	1.22	1.40	0.12	46%	-0.44	0.56	0.76	0.06	0.948	0.113
	Medium	-0.39	0.88	0.57	0.10	45%	-0.18	0.47	0.27	0.05	0.614	0.077
	Low	-0.06	0.37	0.23	0.09	45%	-0.03	0.20	0.11	0.04	0.321	0.061
Duration of Execution (Days)	*Single*	-0.63	0.37	0.20	0.10	65%	-0.41	0.12	0.14	0.06	-0.074	0.085
	Multiple	-0.28	1.08	1.88	0.11	31%	-0.10	0.70	0.66	0.04	1.297	0.097
Order Splitting across brokers	*Single*	-0.49	0.88	0.73	0.11	42%	-0.22	0.48	0.34	0.05	0.640	0.086
	Multiple	-0.50	-0.21	1.53	0.10	86%	-0.43	-0.03	1.35	0.09	0.975	0.099

Table V. Regression Analysis: Determinants of Institutional Trading Friction

The following system of equations is estimated with institutional trading frictions and duration as endogenous variables and various firm (*i*) and order (*t*) characteristics as instruments:

$$\text{duration}_t = \alpha_0 + \alpha_1\text{mcap}_i + \alpha_2\text{complexity}_{it} + \alpha_3\text{volatility}_i + \alpha_4\text{market-return}_i + \alpha_5\text{index-stock}_i + \nu_t$$

$$\text{friction}_t = \beta_0 + \beta_1\,\text{mcap}_i + \beta_2\,\text{complexity}_{it} + \beta_3\,\text{volatility}_i + \beta_4\,\text{liquidity-demander}_{it} + \\ + \beta_5\,\text{duration}_t + \beta_5\,\text{brokers}_t + \beta_6\,\text{Nasdaq-listed}_i + \beta_6\,\text{analyst coverage}_i + \varepsilon_t$$

where $duration_t$ of an order is the actual number of transaction days in the common first stage regression and the predicted value of duration from those regression estimates is used as an explanatory variable in the second stage regressions; $friction_t$ represents adverse selection (cost recovery), clean-up, price-impact, commission, or total friction on order *t* in 5 separate second stage regressions. Explanatory variables and instruments include $mcap_i$ which is the natural logarithm of market capitalization of firm *i* in dollars; $complexity_{it}$ calculated as the ratio of order shares relative to average daily trading volume over the prior five trading days; $volatility_{it}$ calculated as percentage difference in highest and lowest trading price in the past 30 calendar days prior to institutional trading order; $market\text{-}return_t$ is the CRSP value weighted return in the calendar month; $index\text{-}stock_i$ an indicator variable with value of 1 if the stock is a S&P500 index component; $liquidity\text{-}demander_{it}$ takes a value of 1 if order to buy (sell) is made when stock return on order date is positive (negative); $brokers_t$ is the number of brokers engaged in the trades pertaining to the particular order; $Nasdaq\text{-}listing_i$ takes value of 1 if stock is listed on Nasdaq and 0 if it is on the NYSE; $analyst\ coverage_i$ indicates the number of analysts formally following the form; β_0 and α_0 are intercepts; and ε_t and ν_t are error terms. The analysis is based on 3.97 million observations between 1999 and 2005. Statistical significance is indicated by *** for one percent level and ** for five percent level and * for 10% level.

Panel A: Order duration first stage regression.

Order Characteristics	Duration	
Intercept	0.269	***
Market Capitalization	-0.549	***
Complexity of Decision	0.850	***
Stock Volatility	0.265	***
Market Returns	0.168	***
S&P 500 Index Component	-0.165	***
R-Square	1.01%	

Panel B: Institutional friction second stage regressions

Order Characteristics	Adverse selection Or Cost recovery		Clean-up Cost		Price Impact Cost		Commission Cost		Total Execution Cost	
Intercept	0.282	***	-0.150	***	-0.798	***	0.048	***	-0.618	***
Market Capitalization	-0.331	***	-0.257	***	0.027	*	-0.039	***	-0.600	***
Complexity of Decision	0.024		0.352	***	-0.162	***	-0.094	***	0.332	
Stock Volatility	-1.488	***	0.072	***	0.214	***	0.148	***	-1.054	***
Liquidity Demander	-0.667	***	0.371	***	1.205	***	-0.001	***	0.908	***
Duration	-0.306	***	0.167	***	0.178	***	0.106	***	0.146	**
Number of Broker	0.001		-0.025	***	0.175	***	0.007	***	0.157	***
Nasdaq Dummy	-0.237	***	0.068	***	0.063	***	0.023	***	-0.081	***
Analyst Coverage	0.002	***	0.001	***	-0.001	***	-0.001	***	0.001	*
R-Square	0.31%		0.35%		11.30%		2.57%		0.34%	

Figure 1. Institutional Transactions as Proportion of Submitted Order Volume

All orders are based on 148 million shares submitted in 3.97 million orders with client-supplied order volume statistics. Partitions are based on order direction (purchase, sell), market capitalization terciles (small, medium, large), order complexity terciles defined as order volume divided by average daily volume in last 5 days (easy, medium, difficult), liquidity demand (buying in up market or selling in down) versus liquidity supply (doing the opposite), volatility tercile (high, medium, low), execution duration (single or multiple days), and order splitting (single or multiple brokers). The green and the blue bars taken together show the proportion of order shares that result in a trade execution and the brown bar shows the unexecuted or failed or unfilled volume. Green bar is the contribution of fully filled orders and blue bar is the trade contribution of partly filled orders.

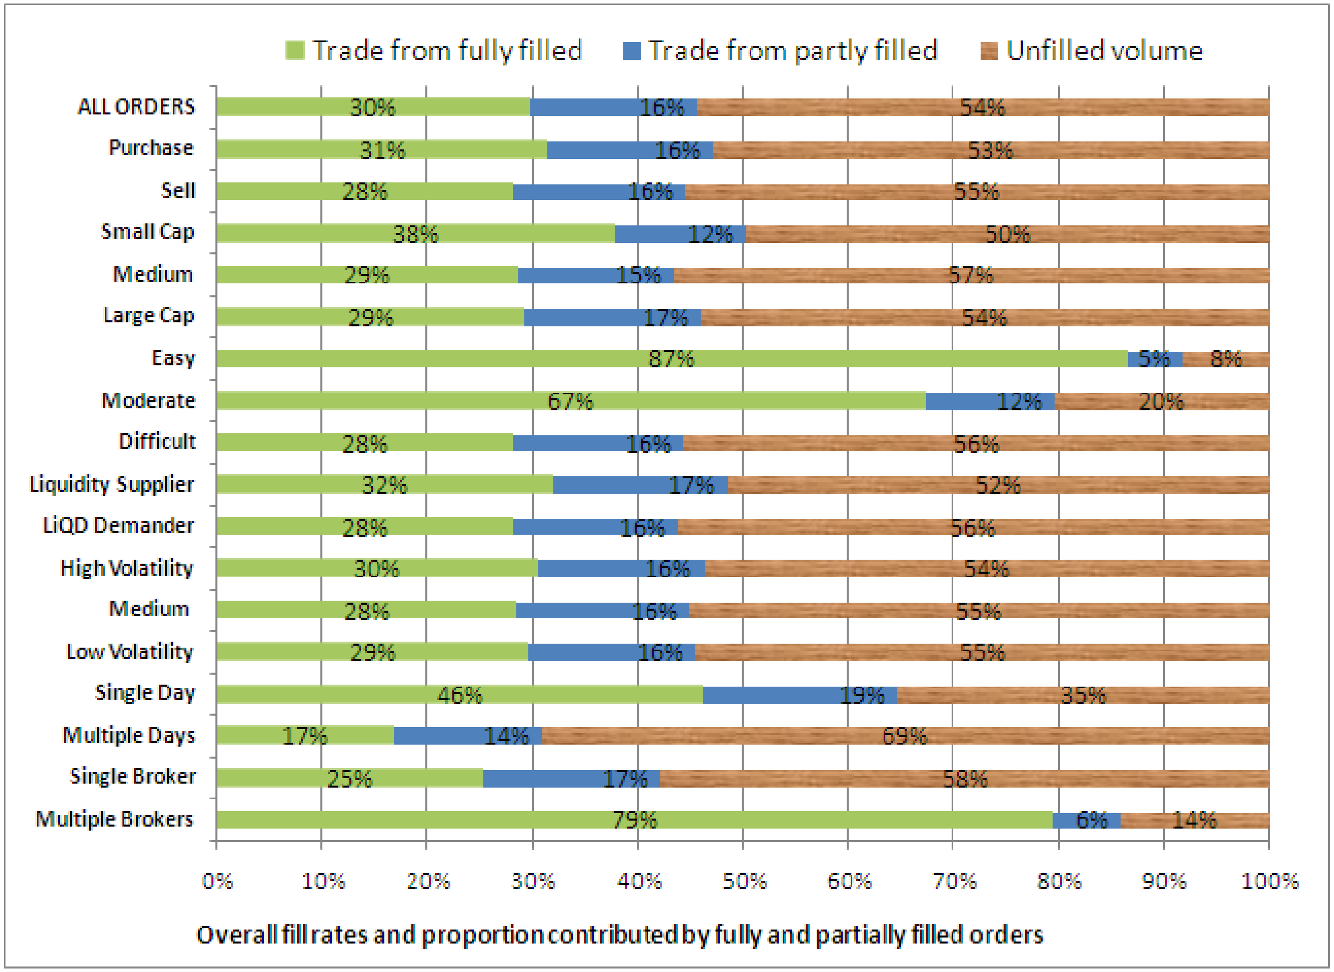

Figure 2. Components of Institutional Trading Frictions

Based on all 3.97 million institutional orders with client-supplied volume for 1999-2005 sample period

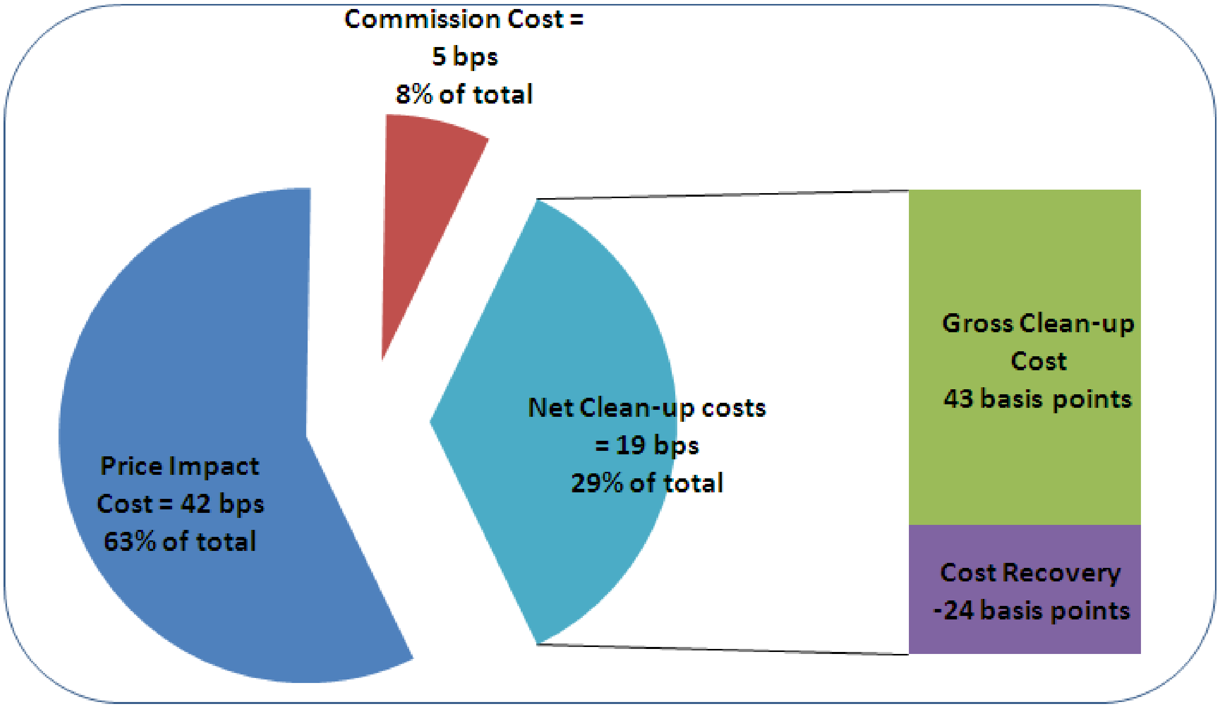

Figure 3. Persistence in clean-up costs over time

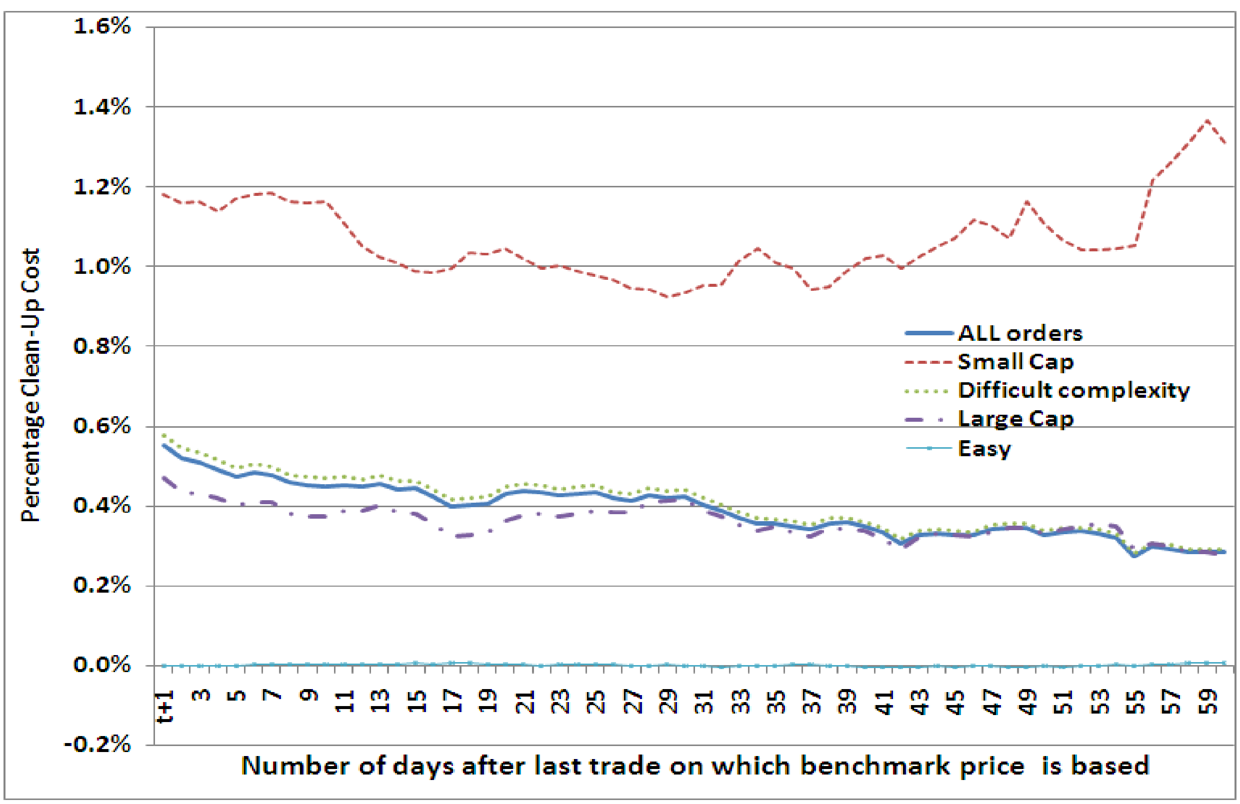

Figure 4. Order Duration and Order Volume

Submitted order volume for each category is shown with the blue bars on the left vertical scale. Order duration is shown with candlesticks. The bottom of the stick shows 25th percentile of order duration within the category. The bottom of the candlestick shows mean order duration. The top of the candlestick shows 90th percentile and the top of the stick shows the 99th percentile of order duration.

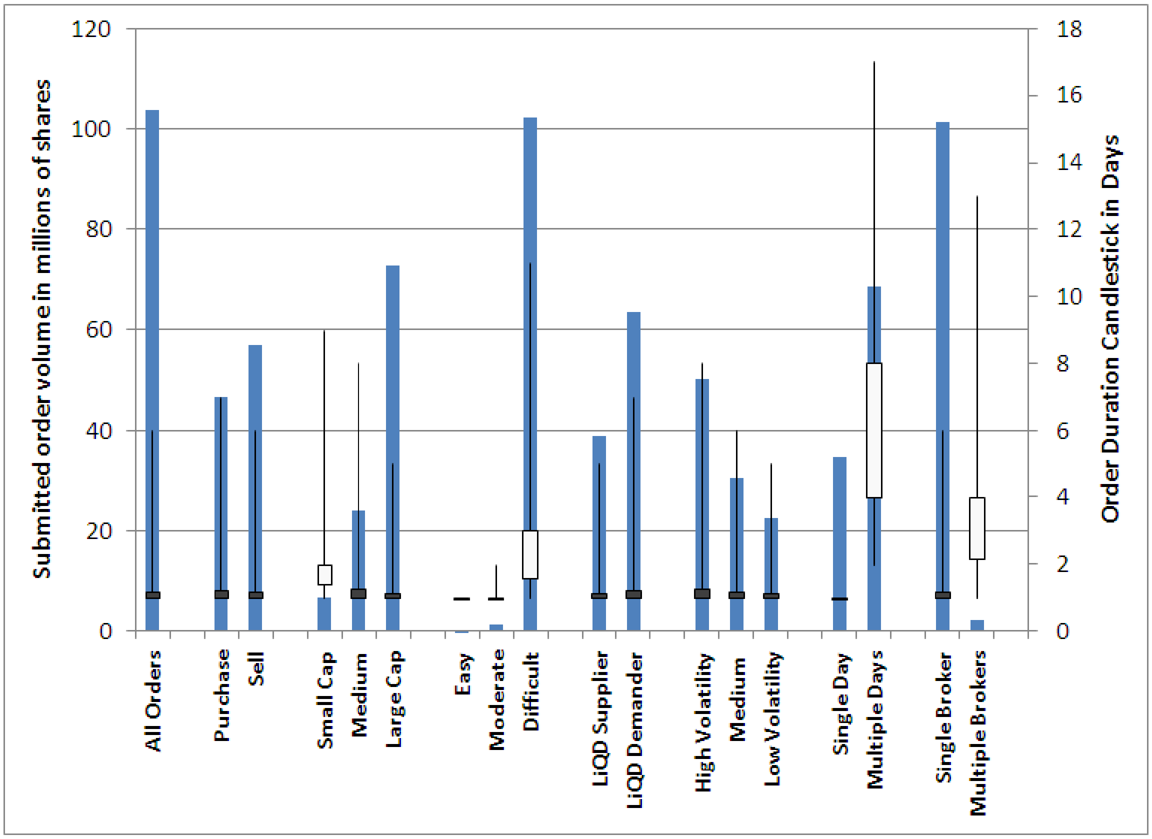

Figure 5. Clean-up cost and cost-recovery variance

Panel A: Firm size categories

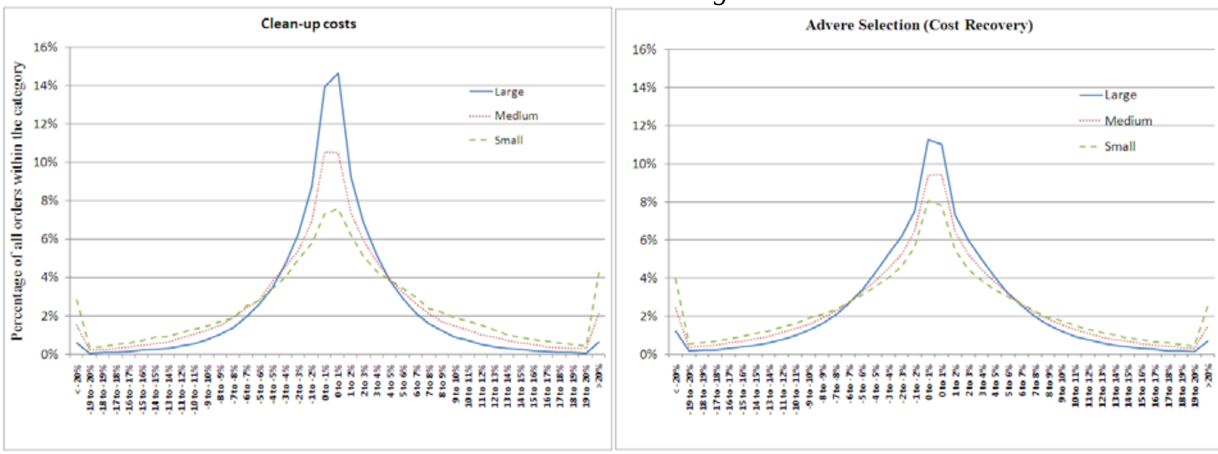

Panel B: Liquidity supplier versus liqudity demander

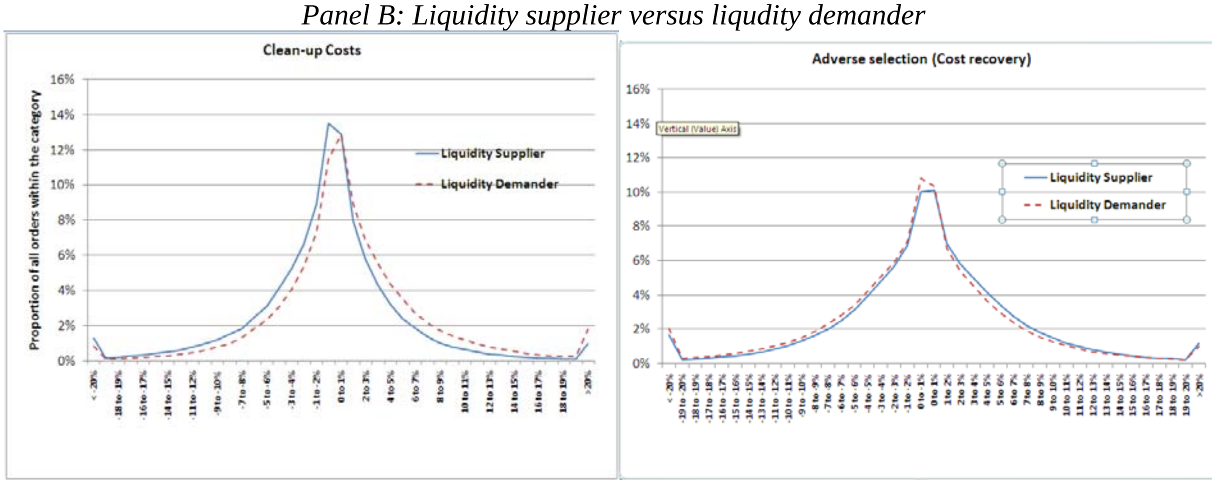

Panel C: Duration