5-2009

Institutional Trading Frictions

Chiraphol New CHIYACHANTANA
Singapore Management University, chiraphol@smu.edu.sg

Pankaj K. JAIN
University of Memphis

Institutional Trading Frictions

Chiraphol N. Chiyachantana and Pankaj K. Jain

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 or order splitting, and trading commissions. Our empirical analysis reveals that various hidden components of institutional trading frictions such as adverse selection and clean-up costs are persistent and could add significantly to previously measured directly observable components of transaction costs. Our simultaneous system of equations accounts for the endogeniety in institutional order aggressiveness based on potentially superior information as well as order splitting strategies in the implementation stage to reduce transaction costs. Order aggressiveness, market conditions and other stock characteristics are associated with significant variations in trading frictions.

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. Suggestions and comments from participants at 2008 meeting of National Bureau of Economic Research (NBER) market microstructure group, 2009 Mid-Atlantic Research Conference, CFA Institute conferences, and finance seminars at the Indian School of Business, University of Central Florida and University of Memphis are highly appreciated. 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 College of Business and can be reached at *pankaj.jain@memphis.edu*. 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 or order splitting, and trading commissions. Our empirical analysis reveals that various hidden components of institutional trading frictions such as adverse selection and clean-up costs are persistent and could add significantly to previously measured directly observable components of transaction costs. Our simultaneous system of equations accounts for the endogeniety in institutional order aggressiveness based on potentially superior information as well as order splitting strategies in the implementation stage to reduce transaction costs. Order aggressiveness, market conditions and other stock characteristics are associated with significant variations in trading frictions.

Introduction

In his AFA presidential address, Stoll (2000) highlights the importance of trading frictions and presents various measures of observable retail trading costs. He distinguishes between real and informational 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 order and execution, return dynamics, speed or aggressiveness of execution, and trading commissions (acronym PQRST). The issues are even more important to understand in the institutional trading context because institutional players now hold 74% of all outstanding stocks according to Bogle (2008, page 35).

We analyze the determinants of PQRST, its persistence, and its randomness. Whereas the vast majority of academic microstructure studies focus on directly observable average trading costs, 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 can smartly route their orders and optimize their trading strategies to improve alpha capture, reduce price slippage, and prevent orders from being gamed (See for example, Bacidore, Otero and Vasa (2010). Carrie (2008) attributes the rapid proliferation of dark pools of liquidity in alternative trading systems to the increasing importance of hidden trading costs and timing risk, higher costs of traditional block desk worked orders, potential information leaks, and backlog effects from slowdown in execution because of transaction costs. We provide a comprehensive empirical analysis of several of these issues in the context of institutional trading.

Consideration of both observable costs and hidden opportunity costs is a fundamental theoretical concept spanning many areas of financial decision making including the field of portfolio management. 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. He defines 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 types of components. The first component, observable execution cost, is widely studied as it relates to the transactions that institutions actually execute and arises from price impact (P), commission, and other transaction fees and taxes (T).[1] Wagner and Edwards (1993) suggest that directly observable costs are only the tip of the iceberg. The second component, hidden opportunity cost, is only sparsely studied as it relates to opportunity loss on any transactions that institutions fail to execute as well as adverse selection risks of both executed and unexecuted transactions.

Our paper advances an understanding of the quantity aspect of execution (Q), which is important because an order that a fund manager does not fulfill to his satisfaction may contain a gain that investors had to forgo by failing to achieve an ideal investment. The magnitude of the opportunity loss will depend on the direction and the magnitude of returns (R) in the stock. We show that the interaction between the quantity dimension and the return dimension generates

[1] This component of institutional trading friction is widely studied. For instance, Berkowitz, Logue and Noser (1988), Holthausen, Leftwich, and Mayers (1987, 1990), Keim and Madhavan (1996, 1997), Chan and Lakonishok (1993, 1995, 1997), Jones and Lipson (2001), Conrad, Johnson, and Wahal (2001), Chakravarty, Panchapagesan and Wood (2004) analyze the relationship between investment styles, trade motivations, exchange listing, soft dollar arrangements and price impact of trades within the U.S. Chiyachantana, Jain, Jiang and Wood (2004) provide international evidence on determinants of price impact in bull and bear markets in 37 countries. Goldstein, Irvine, Kandel, and Weiner (2009) examine the impact of commission costs on institutional trading patterns.

hidden adverse selection costs and hidden clean up costs.[2] An increased probability of filling orders with adverse short term price movements or not filling orders with favorable short term price movements makes these hidden components of trading friction more acute. In Section I, we provide formal definitions and additional background information about the various components of institutional trading frictions and their potential determinants.

Our order level data from Abel/Noser, described in Section II, allows us to form some immensely powerful research designs that are not possible with trade only datasets such as CRSP or TAQ. We are able to directly examine the speed and aggressiveness dimension (S) of execution quality. For portfolio managers lost time is lost performance in terms of their assets under management and cash drag. Furthermore, 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 clean-up costs. Our simultaneous system of equations framework explicitly recognizes that speed or aggressiveness of execution is an endogenous institutional choice. While, informed institutions are expected to be more aggressive than others, they would trade-off the information advantage with the potentially higher price impact of aggressive trading strategies. We study order size, order splitting over time, order splitting across brokers, and order fill rates to investigate these issues. These are difficult questions to answer so we are very careful in conservatively reporting our empirical results. For example, although no academic researcher can read the minds of all institutional managers to know their dynamically evolving portfolio and

[2] In our framework, adverse selection cost (cost recovery) is loss (gain) from filled orders while a clean-up cost is gain or loss resulting from the failure to fully execute an order.

trading decisions, our judicious use of client supplied information gets us as close to clear-cut measurement of hidden institutional trading frictions as is practically possible.[3]

Our empirical results contribute to the literature in several ways. In section III, we focus on institutional aggressiveness in the order submission stage versus order splitting in the trade implementation stage. When institutions possess better than average information, as proxied by favorable long term future excess return, they submit larger orders, execute them more quickly, avoid using multiple brokers, and finish with higher order fill rates. When institutions are trading under difficult conditions that can potentially cause greater frictions, they adopt conservative implementation strategies. Order splitting over time, order splitting across brokers, and lower order fill rates are common especially for large order sizes, adverse market conditions, small market capitalization stocks, volatile stocks, and Nasdaq stocks.

Section IV focuses on friction. First, we show that hidden clean-up costs are persistent and comparable in magnitude to the directly observable price impact, 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 sheds light on the informativeness of institutional trades by establishing that the 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 trading costs as a result of their informational advantage.

The magnitude of total friction, defined as the sum of price impact, commissions, adverse selection cost or cost recovery, and clean-up costs, in our sample period is 67 basis points or 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

[3] Institutional clients report order volume through the Order Delivery System (ODS), which includes complete history of all orders and transactions of portfolio managers. We exclude orders from other clients for whom order volume is algorithmically generated or inferred by Abel/Noser.

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 of 67 basis points in our sample imply a round-trip cost of 134 basis points. It follows from a 106% annual turnover computed from the CRSP mutual fund dataset from 1999-2005 that the round-trip transaction costs eat into nearly 31% of the value weighted CRSP return of 4.60% in that particular period or any other similar low return environments.

As our final major contribution, we characterize the determinants of friction and its various components. Strength of institutional information, order execution strategies, market conditions, and firm-characteristics stand out as the main types of explanatory variables. The two order splitting mechanisms discussed above have opposite effects on overall friction and its components. Longer duration or splitting orders over time reduces clean-up cost, price impact, and overall frictions. In contrast, splitting orders across multiple brokers increases clean-up cost, price impact, and overall frictions. Our innovative finding about the impact of longer order execution duration on the price impact component is different from prior literature because we explicitly control for the joint effect of order complexity, duration, and number of brokers in two stage regressions.[4] Furthermore, adverse market conditions such as buying (selling) stocks with positive (negative) recent return accelerate cost recovery from filled orders but such conditions are also associated with higher clean-up cost, price impact, and total friction. Overall friction is higher for small and volatile stocks and lower for S&P 500 index components. Consistent with Bessembinder (2003) and Chakravarty, Panchapagesan and Wood (2004), we find that Nasdaq stocks face a bigger price impact. Even though institutions have better cost recovery opportunities in Nasdaq stocks, the effect of price impact dominates and generates a higher

[4] This resolves the puzzle in past literature that longer duration is so common despite being associated with higher costs, unconditionally. After conditioning on order volume, we find that longer duration actually helps lower costs.

overall friction for Nasdaq stocks. Section IV also contains robustness checks and characterizes the variance and persistence of friction components.

Our findings have several practical and academic implications as outlined in the concluding Section V. Institutional investors can use our cost estimates as benchmarks to analyze their own implementation shortfall. The numbers can also provide guidance on whether or not it pays to be aggressive in completely filling large institutional orders. More importantly, the implementation policy can be customized to address the effect of order-dynamics, market conditions and firm-specific characteristics. From the academic perspective, these measures of transaction costs provide evidence of limits to arbitrage and also implore asset pricing models to include these 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.

I. Background on Observable and Hidden Frictions and their Determinants

Our paper contributes to the literature by presenting a comprehensive empirical analysis of various dimensions of the institutional trading costs – observable price impact, hidden clean-up cost, hidden adverse selection cost, and explicit commissions – while bearing in mind endogenous order size aggressiveness, order splitting over time, order splitting across brokers, and partial fill rates.

The execution quantity dimension, which arises from the fact that only some orders are completely filled as desired whereas others are only partially filled, plays an important role in determining the overall friction and its various components. Institutions, in an attempt to minimize the price impact cost of more aggressive strategies, may use a less aggressive strategy and thus willingly face a higher risk or uncertainty about the quantity of execution. Unfilled volume must be filled later at a potentially worse price or could possibly result in completely lost returns from an otherwise good investment decision. We analyze the cost of unfilled volume by computing a clean-up cost for a hypothetical trade execution at a fixed interval after the order submission.

The return dynamics are analogously applicable in the computation of adverse selection cost of filled volume. Adverse selection cost occurs if orders with potentially favorable returns fill at a slower pace whereas orders with potentially adverse returns fill rapidly. However, if an institution is better informed than its counterparts the opposite can happen, i.e., a cost recovery can result from the filled volume instead of a loss from adverse selection costs. If the portfolio decisions have positive short term alphas, filled volume will be associated with at least partial cost recovery. Another form of hidden cost is adverse selection cost. The fill rate and return movement interact to generate adverse selection or cost recovery as well, with two computational differences relative to clean-up cost computation. 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.

These hidden and other observable components of trading frictions can be summarized as follows:

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

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

where P_{t+x} is the closing price x days after the last trade completing an institutional order and P_{d-1} is the closing price on the day before the order arrival, 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.

We present a simple numerical example to elaborate the computation of these costs. Suppose, a portfolio manager identifies an undervalued stock priced at \$100 (P_{d-1}) at time 0 and submits a buy order for 10,000 shares at time d. Subsequently at time t, a transaction of 4,500 shares (w_e) takes place at \$100.90 (*WTP*) for which a commission (C) of \$0.05 per share is paid. [5] The stock price rises to \$101.50 in the short term at time $t+x$ and to \$110 in the long run at time L. In this example, without the knowledge of order fill rates, what would a researcher observe? Based on the initial price of \$100 a 10,000 share order would appear to have a CRSP return of 1.50% or \$15,000 with the price rising to \$101.50 in the short term ($t+x$), 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 because the large transaction size of 4,500 shares moves the trade price to \$100.90 instead of \$100. An amount of \$225 is paid in commissions. If we close the books at $t+x$ by buying the remaining unfilled quantity of 5,500 shares (1-w_e) at \$101.50 instead of \$100, then the clean-up costs amount to \$8,250. This leaves the institutional trader with a short term return of only \$2,475. Assuming that institutions incur similar cost at the time of selling, the long term net return by time L after deducting round trip cost is \$74,950, well short of the CRSP return of \$100,000.

---

[5] For expositional convenience, we assume that an institutional order consists of only one trade in this illustration. In our empirical analysis, an institutional order may consist of multiple trades. We sum up all transaction shares within an order to calculate total transaction shares (w_e) and use principle weighted transaction price (*WTP*).

The goal of our paper is to quantify such deviations 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 at $t+x$ 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 positive clean-up costs or positive adverse selection cost, it is useful to offset these amounts to arrive at the net hidden costs.

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

$$Market-Adjusted\ Clean-up\ Cost = \left\{ \left(\frac{P_{t+x}}{P_{d-1}} - \frac{MI_{t+x}}{MI_{d-1}} \right) * (1-w_e) \right\} * \text{Order Direction} \qquad (2)$$

where MI_{d-1} is the level of CRSP value weighted 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. As a robustness check, we also repeat our analysis using beta-adjusted abnormal returns from the market model as the benchmark return.[6] The results are qualitatively similar.

We now briefly discuss the potential determinants of institutional trading friction. For ease of exposition, we group the variables into three categories although some variables may represent more than one category. The three categories are order execution strategy, market condition, and firm characteristics.

---

[6] Beta adjusted clean-up costs are $\left\{ \left(\frac{P_{t+x}}{P_{d-1}} - \beta \frac{MI_{t+x}}{MI_{d-1}} \right) * (1-w_e) \right\} * \text{Order Direction}$

A. *Information versus Friction and Aggressiveness of Institutional Orders*

We examine the impact of information based trades on institutional trading frictions by classifying an order as informed order if it is associated with a favorable one year future return, i.e., positive (negative) market-adjusted return after a buy (sell). When institutions possess better than average time sensitive information, we expect them to submit larger orders and fill them more aggressively to maximize their payoff.

However, when institutions try to execute a complex order that is several times the size of average daily volume, filling it completely is naturally challenging. Such voluminous trading activity is also likely to reveal information more quickly to the market and results in a greater amount of clean-up costs compared to smaller orders which can be camouflaged more easily. Aggressive trading with such large complex orders creates huge price impact costs. Thus institutions would rationally adopt conservative implementation and endogenously split the orders across brokers and over time.

Bertsimas and Lo (1998) develop a theoretical model of institutional order splitting over time, aimed at trading cost reduction. Typically, an order splitting strategy reduces price impact by lowering trade size but also increases the clean-up costs that arise from unfavorable price movements during the longer order execution period. This trade-off calls for optimization of order size, order duration, and number of brokers used to execute the order. Our study also includes order splitting across brokers. In contrast to duration splitting, the use of multiple brokers saves time and increases fill rates simultaneously. However, the disadvantage of this approach is that the increased probability of information leakage, front running, or other manipulations can exacerbate the lost returns component.

Handa and Schwartz (1996) and Wald and Horrigan (2005) develop 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 the context of institutional trading as well. We account for these issues in a simultaneous system of equations to capture institutional strategy, implementation, and performance in the empirical results section.

B. Liquidity, market condition and order direction affect fill rates and transaction costs

We now discuss the role of market conditions in determining transaction costs. Institutions demand (supply) liquidity when trading on the same (opposite) side with other market participants. Thus, we categorize orders as those demanding (supplying) liquidity when institutions submit buy (sell) order when the stock price is rising (falling). This approach follows Wagner and Edwards (1993) who argue that the liquidity characteristic of an institutional order is one of the most important factors affecting the total transaction cost of the trade. Liquidity demanding orders pay a higher price impact and are indicative of an institution's aggressiveness. Liquidity supplying orders can get a lower price impact and are also more likely to be filled in the institutional trading framework. Since institutions are more likely to demand liquidity and pay a higher price when they possess better than average information, the lost returns component of any clean-up cost is likely to be higher for the unfilled portion of such liquidity demanding orders. A liquidity supplying order indicates an institution's patience and may be associated with only marginal clean-up costs, if any. Liquidity supplying orders also earn the spread and thus face very low or even negative price impact. However, the adverse selection risk for those orders is high and thus the cost-recovery for filled liquidity supplying orders will not be as good as for liquidity demanding orders.

Although the discussion above focuses on the asymmetric friction in terms of order direction and individual stock returns, the same logic can apply equally to market-wide returns. Thus, adverse selection costs of filled orders are expected to be lower or negative (i.e., cost recovery) whereas clean-up costs of non-filled volume are expected to be bigger for buys than for sells in bull markets. The opposite would happen in bear markets with costs being higher for sells than for buys. As for the explicit costs, Chiyachantana et al. (2004) show that market conditions create asymmetric price impacts with buys (sells) costing more in bull (bear) markets.

Another important market condition variable is the listing exchange. NYSE and Nasdaq listed stocks have distinct microstructure conditions and trading mechanisms, which can affect the various dimensions of institutional trading frictions. A large body of empirical studies has compared the performance of different exchanges. Huang and Stoll (1996) and Bessembinder and Kaufman (1997a, 1997b) report higher trading costs on the Nasdaq, a dealer market, than on the NYSE, a hybrid market though the cross-market differential has decreased steadily over time due to changes in order handling rule and reduction in tick sizes (Bessembinder (2003)). Chakravarty, Panchapagesan and Wood (2004) confirm that institutional trading cost, measured by price impact cost, is higher on Nasdaq than on NYSE after decimalization.

C. Friction varies with Firm Characteristics

The problem of clean-up costs of unfilled orders is expected to be more severe for smaller firms due to lack of liquidity. On the other hand, smaller firms offer more research opportunities for finding good bargains, which can help informed institutions generate a cost recovery instead of facing adverse selection cost for the filled volume. Big firms are more heavily researched by the entire market leaving little room for information advantage for any particular institution.

Stock volatility generates mixed predictions. While higher volatility reduces the probability of the non-execution of limit orders according to 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 larger 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.

The next three variables relate to the potential information asymmetries in a given stock. Idiosyncratic risk is one such variable that captures the information asymmetry between market participants. Dierkens (1991) and Moeller, Schlingemann and Stultz (2007) suggest that idiosyncratic risk can serve as a good proxy for the level of information asymmetry. We expect the firms with higher idiosyncratic risk to increase both clean up costs as well as cost recovery opportunities for institutions. Wider analyst coverage has the opposite effect by making more information public and transparent. Similarly, if a stock is an S&P500 index constituent, its greater liquidity can help reduce frictions.

II. Data Sources and Research Design

We obtain institutional trading data from the Abel/Noser. The company offers goal-oriented trading strategies and trading cost measurement to help their 776 domestic institutional clients improve their investment performance and help with legal compliance related to best execution. The dataset includes details about the institutional investment orders and related purchase and sale transactions during the 1999-2005 period. We conduct our analysis with both the full sample of orders worth \$15 trillion and a subset of approximately four million orders worth \$4.59 trillion, for which 107 institutional clients provide complete history of actual order volume and transaction data through the Order Delivery System (ODS).[7] We obtain qualitatively similar results with both datasets. Whereas the full sample is more representative of overall institutional trading, the subset relates 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. Actual order size in the latter specification enables us to work with actual fill rates instead of estimated fill rates.[8] The analysis of filled and unfilled volumes provides us with excellent inroads into the hidden 'quantity dimension' of execution quality.

The data provide comprehensive information on institutional trading orders and actual transactions resulting from each order. 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, price at the time of order release, number of shares in the released order, transaction execution price, quantity of shares traded, and commissions charged. 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. In addition, following the approach of Keim and Madhavan (1995, 1997), we exclude orders that took longer than 21 calendar days to complete.

[7] Order volume can be accurately identified by checking whether the order identification field is populated. Most of the results reported in the paper are based on this subset of observations. The remaining observations without any order identification will still have order volume which is either aggregated or algorithmically generated by Abel/Noser.

[8] Fill rate is defined as actual transaction volume divided by the desired volume or submitted order size.

Our research design and the unique features of the dataset offer several advantages by tracking each institutional buy or sell order down to order placement and trade implementation. Through the analysis of quantity dimension of frictions, we show that there is a significant mismatch between actual portfolio performance and ideal performance that academic papers usually measure using price datasets, such as CRSP or TAQ, because a significant portion of institutional orders don't result in an execution. More importantly, we conduct trading cost analysis using information of both institutional orders and actual transactions while a large body of market microstructure research examining the impact of transaction costs focuses exclusively on observed costs (e.g. spread or price impact).[9] Although limit order book datasets such as NYSE OpenBook take us a step further than TAQ by making the submitted orders transparent, they do not completely resolve the situation. Institutions often use order splitting over time and across different venues and brokers, thus fragmenting the information contained in the limit order book. Datasets containing institutional trades such as those by Elkins/McSherry have been used sporadically by Domowitz, Glen and Madhavan (2001) but those studies do not compare order size with actual transaction size. Moreover, these datasets are proprietary in nature and limited access to such data have restricted the proliferation of empirical research leaving ample scope for further research like ours to enhance our understanding about institutional trading costs.

We merge institutional trading data with CRSP to obtain stock specific information and a value-weighted market index. This index 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

[9] The widely used database, such as NYSE Trade and Quote (TAQ), only provides information on actual execution of trades. Furthermore, 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.

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 in 5,688 securities from 107 institutions with explicit 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, within those partly filled orders, 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 submitted (and 48% of the dollar volume submitted) in the full sample is filled and the remaining 54% shares represent the non-filled volume. The total number of brokers used by the institutions in our final sample is 451. On average each order is split into 2.26 trades.

[Insert Table I about here]

III. Information versus Friction and Aggressiveness of Institutional Orders

The framework of our analysis captures the three stages of institutional trading – strategy, implementation, and performance. We conjecture that order size strategy is determined exogenously based on the strength of institutional portfolio research and information. In contrast, trade implementation strategies – degree of order splitting and fill rates – depends on trading cost considerations. Eventually the degree of aggressiveness at the implementation stage is determined endogenously to trade off the gains of trading based on information advantage against an increase in trading costs associated with aggressiveness. We find that future returns are positively and statistically significantly correlated with order size. The following preliminary order size regression result with calendar year fixed effects and institutional client fixed effects verifies the first conjecture:

$$\mathit{Order\ size} = \underset{}{1{,}060{,}679} + \underset{(11{,}883)}{46{,}001} \ {*}\mathit{Future\ return} - \underset{(0.06)}{2.95} \ {*}\mathit{Firm\ size} + \lambda{*}\mathit{Fixed\ Effects} + \mathit{Error} \quad (3)$$

where order size is the monetary amount of institutional order in dollars. Long term future excess returns is calculated by first dividing the individual stock price appreciation or depreciation in one calendar year following the institutional trade by the stock price on the order submission date and then subtracting the analogous CRSP value weighted return for the corresponding period. Finally, to obtain signed future return, we multiply the excess return with order direction, which is positive one for buys and negative one for sells. Firm size is in millions of dollars. λ is the matrix of coefficients for 5 calendar year fixed effects and 106 institutional client fixed effects. Standard errors for the coefficient estimates are in parentheses. The regression is based on 3,976,387 observations and has an adjusted R-squared of 0.06%. Long term future excess return is a proxy for the strength of institutional information. It's coefficient of 46,001 has a standard error of 11,883 which implies a t-value of 3.87 and statistical significance at the 1% level. Thus, we infer that institutions appear to be submitting larger orders more aggressively when they possess better than average information about a stock's long term future potential.

Aggressive order submission does not automatically translate into aggressive trading because of the offsetting trade-off presented by higher friction. Order splitting over time or across brokers is the mechanism to convert the large orders into smaller trades to reduce the magnitude of friction, particularly its observable components. In this section, we characterize the choice of order splitting for various types of orders. In the next section, duration and number of brokers are used as endogenously determined variables which in turn affect transaction costs in a simultaneous system of equations.

Figure 1, Panel A plots order duration on the right vertical axis and order volume on the left vertical axis for firm and order specific categories on the horizontal axis. For the overall sample, the 25th percentile of order duration is 1 day, the mean is 1.19 days, the 90th percentile is 1 day and the 99th percentile is 6 days. We truncate a miniscule amount of orders from the sample to contain orders completed within 21 days. So, the 99th percentile is still meaningful to analyze as an indicator of maximum duration. Panel B plots the corresponding information for number of brokers.

[Insert Figure 1 about here]

First, we partition the sample into three groups (Informed, Neutral, Liquidity orders) based on terciles of long term future excess returns. Informed orders account for slightly higher order volume than liquidity orders and are executed more aggressively with shorter average duration. 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 the 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 splitting is order complexity, which divides the order specific volume by the average daily market wide volume for a given stock in the last 5 days. Although there are an 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 a 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 99^{th} percentile of 11 days in Panel A. Similarly, Panel B shows that easy orders are filled with a single broker but the 99^{th} percentile of difficult orders use 3 brokers to fill the trades. Given the importance of order complexity, it is used as a key variable in the duration and broker regression in the next section.

In the next partition, liquidity demanding orders represent both a higher volume and a higher mean duration of 1.25 days relative to liquidity supplying orders. NYSE stocks account for a higher proportion of institutional volume but Nasdaq stocks take longer to execute. For the partition based on stock volatility, high volatility stocks account for a larger proportion of order volume and also a longer order duration. Number of brokers is similar across liquidity demander, listing exchange, and volatility partitions.

Next we set out to understand the distribution and determinants of order fill rates. In Table II, we capture the incremental effect of explanatory variables capturing order execution strategies, market conditions, and firm-specific characteristics on our dependent variables, namely, the fill probability and the fill rates in multivariate regression settings. We conduct the analysis using two approaches – Probit and OLS regressions. The probit regression is estimated to understand the fill probability while the OLS regression is estimated to understand the fill rates. In the Probit regression, we estimate the likelihood of an order being filled, following Wald and Horrigan (2005). The dependent variable is equal to one for completely filled orders and zero for orders that are not completely filled. The goal of the Probit regression 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. In the first OLS regression using the full sample of all institutional orders, the dependent variable is the actual fill rate for each order, which is 100% for fully filled order and ranges from 0.01% to 99.99% for party filled orders. In the second OLS regression, we examine the determinants of fill rates for the sub-sample of 771,861 partly filled orders only. This measure can help us understand the strategies that help institutions increase the fill rate of the orders, conditioned on the fact that they are using execution methods that do not result in a complete fill.

Explanatory variables fall in the four broad categories of information, order execution strategy, market condition, and firm characteristics. The indicator variable for an informed order is assigned the numerical value of 1 if the one year future excess return associated with a buy (sell) order is positive (negative), and the value of 0 otherwise. Following Chiyachantana et al. (2004) we define order complexity as the number of shares in an order divided by the average daily trading volume in the given stock over the prior five trading days. Duration is defined as the number of days elapsed from the date of order submission to the date of the final trade for that order package. The next variable is the number of brokers that it took to execute the trades in an order. Following Wagner and Edwards (1993), if an order to buy (sell) is made when the stock return on the order date is positive (negative), we classify it as a liquidity demanding order. An order with the opposite return scenario is classified as a liquidity supplying order. Adverse market condition is an indicator variable that takes a value of 1 if an order to buy (sell) is made in the calendar month when the CRSP value weighted index is positive (negative). Nasdaq listing is an indicator variable which takes the value of 1 if the stock is listed on Nasdaq and 0 if it is on the NYSE. Firm size is the natural logarithm of the market capitalization of the firm in dollars. Stock volatility is calculated as the percentage difference between the highest and the lowest trading price in the past 30 calendar days prior to institutional trading order. All regressions include 106 institution-specific and 5 calendar year-specific fixed effects variables.

The first column of Table II reports the estimated coefficients from Probit regression based on approximately 4 million institutional orders. This regression has a pseudo R-squared of 28.03%. The positive coefficient for informed orders implies that institutions aggressively try to fill orders when they have above average information. The probability of completely filling an order decreases with relative order size as measured by order complexity. The next two variables represent transaction execution methods related to order splitting over time or across brokers. The probability of filling the order decreases (increases) with the number of days (brokers) used to execute the order. Liquidity demanding order, adverse market condition and stock volatility all adversely affect the probability of filling an order while bigger firm size favorably affects the probability of filling an order.

The second and third columns of Table II report the estimated coefficients from OLS fill rate regressions. The adjusted R-squared for the regressions are 26.11% and 8.95% for all orders and partially filled order regressions, respectively. The direction and the significance of coefficients in the OLS regressions are identical to those in the Probit regression with the exception of the Nasdaq listing indicator variable, which now takes a statistically significant negative sign.

[Insert Table II about here]

The overall conclusions from the analysis in this section are consistent with the notion that order implementation is a profit maximizing endogenous institutional choice where they balance their aggressive use of information and large order sizes in the order submission stage against conservative trading strategies of order splitting in the implementation stage to reduce their transaction costs.

IV. Composite transaction costs and their determinants

The end result of various institutional strategies discussed above can be seen in eventual trading cost performance, which is the final step of our study. Here we present the magnitude of institutional trading friction, its components, its persistence, and its determinants.

A. Summary statistics of transaction costs

Table III provides market-adjusted estimates for overall friction and its various components.[10] Filled orders do not have any clean-up costs by definition. Partially filled orders give rise to clean-up costs of 64 basis points. Average clean-up costs for all orders, including both fully and partly filled 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 minimal. 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 significantly alter the gross returns on a stock from order arrival to final execution. 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 that are incurred when institutions are either 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.

[10] Raw costs not adjusted for market returns are very similar to the market-adjusted results.

Gross clean-up costs at 43 basis points are actually higher than price impact but institutions are able to recover part of these hidden costs through the informational advantage on their filled orders. Thus, the net-clean up costs are 19 basis points which account for 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. Friction is related to order execution strategy, market condition, and firm characteristics

In Table IV, we assess the magnitude of total transaction costs and its four components in a univariate setting. The first four columns present the average trade level cost of institutional activity with no fill rate adjustment. The fifth column shows the average fill rate. The last six columns show the effective costs at the order level after adjusting (weighting) trade level costs 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 costs times the fill rate (or one minus fill rate) equal the order level weighted costs.[11] The differences in unadjusted costs and weighted costs highlight the importance of carrying out the analysis using complete order information. The past literature is mostly based on unadjusted costs and assumes that all institutional orders are completely filled. The additional knowledge of simply the average fill rate is not sufficient to learn about the order level frictions due to the significant variations in the fill rates of individual orders. The most appropriate approach is to use the weighted costs. Weighted costs are reported throughout the paper and used for majority of our analysis as they provide the best available measure of institutional trading frictions at the order level. The costs

[11] Because of differences and uniqueness in order volume, fill rates, and transaction volumes for each order, the reported summary statistics at the aggregate sample level presented in the table do not yield themselves to exact equality. The minor departure of aggregated statistics from the order level equation represents Jensen's inequality.

for the overall sample discussed in last column of Table III are shown again in the first row of Table IV but there is significant cross-sectional variation in these costs.

We analyze the determinants of variations in institutional trading frictions in three broadly defined groups: order execution strategy, market condition, and firm characteristics. The first partition within the order execution strategy group is based on terciles of the order complexity variable, which are used to divide the sample into three groups (Easy, Moderate, Difficult). Order complexity appears to be a key driver of order fill rate as well as institutional frictions. Easy orders have the highest fill rate of 92% which is more than twice as much as the 44% fill rate of difficult orders. Cost recovery is low while clean-up and price impact costs are highest for difficult orders. Weighted clean-up cost for difficult orders at 45 basis points is a striking 232 times the clean-up cost for easy orders, which have negligible clean-up costs. The total execution cost for difficult orders is three times larger than that of easy orders. In untabulated results, we calculate the dollar cost of overall friction. The difficult orders face the bulk of total costs accounting for over 99% of the total institutional trading frictions of $30.68 billion in our sample.

Our next two partitions represent order splitting strategies over time or across brokers. We divide the orders into two groups based on whether or not the orders were split over multiple days or using multiple brokers. With respect to fill rate, 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 multi-day orders forcing the institutions to break up their orders. The execution of multi-day 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, a multi-day execution ends up with the bulk of the frictions. Of course, if institutions had tried to execute those orders on a single day, perhaps they would 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 for direct comparison to the easy single day executions. The other type of order splitting is across multiple brokers. Order fill rates of 86% for multiple broker orders is almost twice that of single broker executions. Despite the higher percentage of total costs, multiple broker orders account for only one third of the dollar transaction costs because the vast majority of the orders are completed using a single broker. Given that the execution duration and number of brokers can be endogenously determined by the institutions, we later estimate a simultaneous system of equation where these two methods of order splitting are modeled as institutional choice variables.

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. If 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 at a later time. 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.

Consistent with prior NYSE versus Nasdaq comparative studies of retail spreads (e.g., Bessembinder (2003)), Nasdaq stocks face a larger price impact, but also offer institutions even greater cost recovery opportunities. Overall, higher clean-up cost and price impact cost lead to higher overall frictions for Nasdaq stocks.

Next, we allocate stocks into three firm-size groups with an equal number of stocks based on their market capitalization terciles. 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%. Institutional traders in small stocks could be more informationally advantaged and may be trying to lower their unfilled rates, but the lack of available liquidity might balance their aggressiveness, effectively putting a cap on empirically observed fill rates. 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 the lion's shares of filled and unfilled orders, large capitalization stocks account for the bulk of total dollar transaction costs at $22.33 billion. 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.

The last three variables relate to the potential information asymmetries in a given stock. Analyst coverage indicates the number of analysts following the firm and S&P index-stock is an indicator variable on whether stock is an S&P500 index component. Idiosyncratic volatility is estimated as follows. For each calendar month, excess daily returns of each individual stock are regressed on the daily Fama-French three factors: *Rm-Rf, SMB,* and *HML*. The monthly idiosyncratic volatility of the stock is the product of the standard deviation of the regression residuals and the square root of the number of observations in the month. Wider analyst coverage, inclusion in S&P 500 index and lower idiosyncratic risk reduce the opportunities for institutions to profit from research and are thus associated with lower cost recovery opportunities. As expected, overall frictions are lower for index stocks and higher for stocks with greater idiosyncratic risks.

[Insert Table IV about here]

C. Regression analysis with simultaneous system of equations

In Table V, we present multivariate regression results based on a simultaneous system of equations. The first stage regression estimates in Panel A show how institutions split the execution of their orders over time or across brokers. Order execution duration decreases with the strength of institutional information at the 5% statistical significance level. Every 1% increase in future returns in the favorable direction is associated with a reduction in order duration of 0.21 days. Order duration increases with relative order size (order complexity) and during adverse market conditions (liquidity demanding). The intuition for this result is straightforward as difficult tasks are likely to take longer to complete. Duration is also higher for stocks with a Nasdaq listing, smaller firm size, or higher volatility. Aggressively executing orders with these characteristics could clearly cause potentially greater frictions. Overall, this regression reinforces the notion that splitting orders over time is an endogenous decision.

Multiple broker engagement is another form of order aggressiveness available to institutions. A common characteristic of duration and multiple brokers is that they split order size to help reduce price impact. However, the two forms of splitting have some diagonally opposite implications as well. Multiple brokers can shrink order duration and also increase order fill rates as seen previously. Order splitting across brokers is used more frequently or less frequently for some of the same reasons for which duration splitting was adopted above, but multiple brokers are avoided in situations when institutions possess better than average information. This strategy of limiting the number of brokers will be optimal if brokers' actions in shopping around for counterparties reveal the institutions' information about future returns to a larger number of market participants.

Given the similarities and differences in the two forms of order splitting, we expect them to uniquely affect the different components of friction. The second stage regression reported in Panel B of Table V has overall friction and its various components as dependent variables in separate regressions. Explanatory variables include order splitting variables from the first stage, market conditions, and several stock specific characteristics. Order complexity continues to increase overall friction and its various components. Longer duration increases adverse selection cost, but helps reduce clean-up cost and price impact cost. Our duration result is insightful and it differs from prior research because we explicitly control for order complexity in the two stage regressions. The result is also consistent with the fact that a higher proportion of overall volume is executed over a multiple day horizon.

The empirical results presented earlier show that multiple brokers help increase fill rates and yet most of the order volume is associated with the use of a single broker. The friction results easily help explain that dilemma. Multiple brokers are associated with significantly higher overall friction and higher friction components. Since the final friction is more important than the intermediate fill rate, institutions use multiple brokers for only a relatively small portion of the overall volume.

The negative coefficient for liquidity demanding strategy of buying (selling) stocks with positive (negative) recent return 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 higher overall frictions for such aggressive orders compared to liquidity supplying orders.

Consistent with prior NYSE versus Nasdaq comparative studies of retail spreads (e.g., Bessembinder (2003)), Nasdaq stocks face a larger price impact but also offer institutions greater cost recovery opportunities than NYSE stocks. From a cost-recovery or returns perspective, Nasdaq listing appears to be more beneficial. Since institutional players who now hold 74% of all outstanding stocks according to Bogle (2008, page 35), companies listing their stock could give a substantial weight to institutional perspective. Thus, they may not necessarily see a direct need to list on NYSE for better retail liquidity alone if important components of institutional frictions worsen with such a move. Eventually, higher clean-up costs lead to higher overall frictions for Nasdaq stocks. Thus, our final comparative net cost-benefit analysis of NYSE versus Nasdaq for institutional friction is the same as retail friction, with total transaction cost being higher on Nasdaq.

As expected, overall friction is lower for large stocks and index stocks while it is higher for more volatile stocks. Wider analyst coverage brings information in public domain and eliminates cost recovery opportunities for any specific institution. This leads to a positive adverse selection cost and thus, higher overall costs. Price impact is lower because counterparties are less worried about any particular institution's informational advantage for very widely covered stocks, but hidden friction components still dominate to create higher overall costs.

As hypothesized higher idiosyncratic risk leads to larger clean-up costs, price impact, and commissions. However, institutions are net beneficiaries of increased idiosyncratic risk because of increased level of cost recovery from information based trading in those stocks. The cost recovery dominates other components and leads to negative overall costs for stocks with greater idiosyncratic volatility.

[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 the persistence of these costs as well as variance in the level of friction to understand transaction cost risk.

D. Persistence of hidden costs and transaction costs risk

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 so far for $t+x$ in equation (1), we now consider periods of 1 to 60 days after the completion of the last transaction in an order package. For brevity, results are shown only for a few sub-samples based on complexity and firm size but the patterns are persistent in other categories as well.

For the overall sample, clean-up cost begins at 55 basis points one day after transaction period. Thereafter, clean-up costs remain persistent and range bound between 29 and 55 basis points with an average of 39 basis points that compares well with the 20^{th} day reported value of 43 basis points used and presented throughout the paper. We observe this 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]

The average hidden opportunity costs statistics provide good benchmarks for overall institutional performance in a repeated trading setting. However, any single order carries a 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 the risk was higher for small and highly volatile stocks. We now provide additional insight into the transaction cost risk issue by separately examining the hidden components of friction in Figure 4. Clean-up cost charts are to the left and cost recovery charts are to the right. In each Panel, the horizontal axis captures the transaction cost variation by forming cost range categories with one percent intervals. We consolidate the extreme categories by clubbing together the orders that have less than -20% or more than 20% cost. The vertical axis plots the proportion of all orders with a given characteristic (e.g. large market capitalization) that fall in the transaction cost range on x-axis. For example, in Panel A which is based on firm size, 15% of all large stock orders have clean-up 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 the variance in hidden opportunity costs conditional on liquidity provision. Liquidity supplier and liquidity demander orders have similar variance. However, we can see that the costs are slightly asymmetric. The 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 Table IV that multiple day orders have higher clean-up costs and lower cost recovery. Now we show in Panel C of figure 4 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. The observable and explicit components of implementation shortfall are fairly well understood. They relate to the portion of orders that actually executes and arise 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 order volume and clean-up costs of unfilled order volume and find that the hidden opportunity cost of inaction is very high for institutional orders. The clean-up cost of 43 basis points is the biggest component of institutional frictions followed closely by price impact, both of which dwarf the explicit commissions of 5 basis points. Overall frictions in the full sample are 67 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 equations to account for endogenous order duration and number of brokers and the effects of such order splitting on frictions.

When institutions possess better than average information, as proxied by favorable long term future excess return, they submit larger orders, execute them more quickly, avoid using multiple brokers, and finish with higher order fill rates. When institutions are trading under difficult conditions that cause greater frictions, they adopt conservative trade implementation strategies. Order splitting over time, order splitting across brokers, and lower fill rates are common especially for large order sizes, adverse market conditions, small market capitalization stocks, volatile stocks, and Nasdaq stocks.

In our analysis of friction, we show that hidden clean-up costs of unfilled volume are persistent and comparable in magnitude to the directly observable price impact, which doubles the estimate of total costs in relation to previous studies. The adverse selection cost component of filled volume is negative. This leads us to infer that institutions in our sample are actually more informed than their counterparties and are able to recover part of their trading costs as a result of their informational advantage.

Institutional trading friction varies with order execution strategies, market conditions, and firm characteristics. Friction and its major components are higher for complex, high volume orders, but spreading the execution of such orders over multiple days can help mitigate costs. Institutional decisions to split orders over time must be based on careful consideration of the trade-off between various components of trading costs, some of which increase while others decrease with duration. Order splitting across brokers significantly increases overall friction and its various components, which explains why a large proportion of overall institutional volume is executed with a single broker, despite higher fill rates with multiple brokers. Cost recovery opportunities are larger with liquidity demanding strategies in small, volatile, idiosyncratic, S&P 500 constituent, or Nasdaq stocks. Clean-up cost increases with liquidity demanding strategy, for Nasdaq stocks, and stocks with higher idiosyncratic risk. The total friction increases with order complexity, multiple brokers, liquidity demanding trading strategy, volatility, wider analyst coverage, or Nasdaq listing and it decreases with and use of longer duration, firm size, idiosyncratic risk, and inclusion in S&P 500 index. Often the reduced friction is a result of cost recovery possibly through institutional informational advantage.

The results have several practical and academic implications. Institutional investors can use our cost estimates 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 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 situations and asset classes. From the academic perspective, the concept of total transaction cost 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 highlight the importance of asset pricing models that explicitly include transaction costs because such costs can lead to significant deviations from the ideal performance of a paper portfolio that we so often see in theoretical and empirical papers. With institutional investors playing a bigger and bigger role in our investment portfolios, optimization of institutional trading frictions can have a direct and meaningful impact on our investment return performance.

References

Ahn, Hee-Joon Kee-Hong Bae and Kalok 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.

Bacidore, Jeff, Hernan Otero, and Alak Vasa, 2010, Does Smart Routing Matter? Working Paper available at http://ssrn.com/abstract=1681449.

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.

Bessembinder, H., 2003, Trade Execution costs and market quality after decimalization, *Journal of Financial and Quantitative Analysis* 4, 747-778.

Bessembinder, H., and H. Kaufman, 1997a, A comparison of trade execution costs for NYSE and NASDAQ-listed stocks, *Journal of Financial and Quantitative Analysis* 32, 287-310.

Bessembinder, H., and H. Kaufman, 1997b, A cross-exchange comparison of execution costs and information flow for NYSE-listed stocks, *Journal of Financial Economics* 46, 293-319.

Bogle, John C., 2008, Black Monday and Black Swans, *Financial Analyst Journal* 64, 30-40.

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, 1995, The behavior of stock prices around institutional trades, *Journal of Finance* 50, 1147-1174.

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.

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.

Dierkens, Nathalie, 1991, Information Asymmetry and Equity Issues, *Journal of Financial and Quantitative Analysis 26(2),* 181.

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.

Huang, R. and H. Stoll, 1996, Dealer versus auction markets: a paired comparison of execution costs on NASDAQ and the NYSE, *Journal of Financial Economics* 41, 313-357.

Holthausen, Robert W., Richard W. Leftwich, and David Mayers, 1987, The effect of large block transactions on security prices: A cross-sectional analysis, *Journal of Financial Economics* 19, 237-268.

Holthausen, Robert W., Richard W. Leftwich, and David Mayers, 1990, Large-block transactions, the speed of response, and temporary and permanent stock-price effects, *Journal of Financial Economics* 26, 71-96.

Jerias, Davis and George Sofianos, 2008, Equity execution strategies, *Street Smart* (A Goldman Sacs publication) 36, 1-8.

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, 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.

Moeller, S., F. Schlingemann and R. Stulz, 2007, How do Diversity of Opinion and Information Asymmetry Affect Acquirer Returns, *Review of Financial Studies* 20, 2047-2078.

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.*

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

This table reports sample characteristics of institutional trading data compiled by Abel/Noser for 1999-2005. We eliminate the observations where order level information is merely aggregated or algorithmically generated by Abel/Noser and not directly provided by the client. We provide statistics for all orders, both fully filled orders and partially filled orders. Fill rate is defined as actual transaction volume divided by the submitted order size of customer-supplied order volume.

Sample Characteristics:	
Number of institutions	107
Number of securities	5,688
Average Market Cap (billion $)	3.59
Number of brokers	451
Number of transaction per order	2.26
All Orders:	
Number of orders with client-supplied volume statistics	3,976,387
Order share volume (Billion shares)	147.99
Total 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
Total 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
Total 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 orders and zero for incomplete orders. In the second column, we present the results of an ordinary least square (OLS) regression where the dependent variable is the actual fill rate of an order ranging from 0.01% to 100%. Fill rate is defined as shares traded divided by shares submitted. In the last column a similar OLS regression is estimated within the sub-sample of partially filled orders only. Explanatory variables include an indicator variable for informed order which takes a value of 1 if the one year future excess return associated with a buy (sell) order is positive (negative), and 0 otherwise. Following Chiyachantana et al. (2004) we define order complexity as the number of shares in an order divided by the average daily trading volume in the given stock over the prior five trading days. Duration is the number of days elapsed from the date of the order submission to the date of the final trade for that order package. The next variable is the number of brokers that it took to execute all of the trades in an order. Following Wagner and Edwards (1993), if an order to buy (sell) is made when the stock return on the order date is positive (negative), we classify it as a liquidity demanding order. An order with the opposite return scenario is classified as a liquidity supplying order. Adverse market condition is an indicator variable that takes a value of 1 if an order to buy (sell) is made in the calendar month when CRSP value weighted index is positive (negative). Nasdaq listing is an indicator variable which takes the value of 1 if the stock is listed on Nasdaq and 0 if it is on the NYSE. Firm size is the natural logarithm of the market capitalization of the firm in dollars. Stock volatility is calculated as the percentage difference between the highest and the lowest trading price in the past 30 calendar days prior to institutional trading order. The Probit and OLS regressions include institution-specific and calendar year-specific fixed effects variables. Statistical significance is indicated by ** for the one percent level.

	Probit Regression	*OLS Regression*	
Order Characteristics	**Probability of Fully Filling a Order**	**All Orders**	**Partly Filled Orders**
Intercept	2.580 **	0.958 **	0.130
Informed Order	1.880 **	0.412 **	0.466 **
Order Complexity	-1.400 **	-0.820 **	-0.271 **
Duration (days)	-0.408 **	-0.078 **	-0.033 **
Number of Brokers	0.479 **	0.038 **	0.041 **
Liquidity Demanding Order	-0.061 **	-0.011 **	-0.013 **
Adverse Market Condition	-0.012 **	-0.003 **	-0.008 **
Nasdaq Listing	0.002	-0.004 **	-0.015 **
Firm Size	0.094 **	0.034 **	0.070 **
Stock Volatility	-0.210 **	-0.038 **	-0.033 **
Year Fixed Effects Included	*Yes*	*Yes*	*Yes*
Institution Fixed Effects Included	*Yes*	*Yes*	*Yes*
Number of Observations	3,976,387	3,976,387	771,861
Pseudo R-square	28.03%		
Adjusted R-square		26.11%	8.95%

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

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

where P_{t+x} is the closing price 20 days after the last trade completing an institutional order and P_{d-1} is the closing price on the day before the order submission. 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 the market index return from the 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 an 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 Tables I and II. We use medians or tercile values as cut-off points to create the sub-samples within order complexity, firm size, stock volatility, analyst coverage, and idiosyncratic risk categories. The first four columns present the trade-weighted (or unfilled-volume-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 within the relevant category.

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
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
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
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
Exchange Listing	NYSE	-0.13	0.78	0.53	0.11	43%	-0.06	0.42	0.24	0.05	0.66	0.07
	Nasdaq	-0.99	1.11	1.20	0.11	45%	-0.49	0.56	0.61	0.05	0.73	0.11
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
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
Analyst Coverage	High	-0.28	0.91	0.72	0.10	47%	-0.13	0.48	0.34	0.05	0.74	0.08
	Low	-0.44	0.75	0.64	0.13	44%	-0.20	0.40	0.30	0.06	0.55	0.09
S&P Index Stocks	Yes	-0.18	0.83	0.60	0.10	48%	-0.08	0.45	0.28	0.04	0.70	0.07
	No	-0.65	0.90	0.88	0.13	44%	-0.31	0.46	0.43	0.06	0.64	0.10
Idiosyncratic Risk	High	-0.53	1.19	1.02	0.12	46%	-0.25	0.62	0.50	0.06	0.93	0.11
	Low	-0.07	0.43	0.23	0.09	44%	-0.03	0.24	0.10	0.04	0.35	0.06

Table V. Regression Analysis: Determinants of Institutional Trading Friction

The following system of equations is estimated with duration and number of brokers as endogenous variables in the first stage regression and various firm (*i*) and order (*t*) characteristics as instruments:

$\text{Duration}_t = \alpha_0 + \alpha_1 \text{ complexity}_{it} + \alpha_2 \text{ liquidity-demander}_{it} + \alpha_3 \text{ mcap}_i + \alpha_4 \text{ volatility}_i + \alpha_5 \text{ Informed}_i + \nu_t$

$\text{Broker}_t = \alpha_0 + \alpha_1 \text{ complexity}_{it} + \alpha_2 \text{ liquidity-demander}_{it} + \alpha_3 \text{ volatility}_i + \alpha_4 \text{ Informed}_i + \alpha_5 \text{ Nasdaq-listed}_i + \nu_t$

$\text{friction}_t = \beta_0 + \beta_1 \text{ complexity}_{it} + \beta_2 \text{ duration}_t + \beta_3 \text{ brokers}_t + \beta_4 \text{ liquidity-demander}_{it} + \beta_5 \text{ mcap}_i + \beta_6 \text{ volatility}_i + \beta_7 \text{ idiosyncrastic risk} + \beta_8 \text{ analyst coverage}_i + \beta_9 \text{ index-stock}_i + \beta_{10} \text{ Nasdaq-listed}_i + \varepsilon_t$

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; *broker*$_t$ is the number of brokers used to execute an order, *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 *complexity*$_{it}$ calculated as the ratio of order shares relative to average daily trading volume over the prior five trading days; *liquidity-demander*$_{it}$ takes a value of 1 if an order to buy (sell) is made when the stock return on the order date is positive (negative); *mcap*$_i$ which is the natural logarithm of the market capitalization of firm *i* in dollars; *volatility*$_{it}$ is calculated as the percentage difference between the highest and the lowest trading price in the past 30 calendar days prior to institutional trading order; *informed*$_{it}$ takes value of 1 if the one year excess returns of buy (sell) order is positive (negative). *index-stock*$_i$ is an indicator variable with a value of 1 if the stock is an S&P500 index component; *brokers*$_t$ is the number of brokers engaged in the trades pertaining to the particular order; *Nasdaq-listing*$_i$ takes a value of 1 if the stock is listed on Nasdaq and 0 if it is on the NYSE; *Idiosyncratic volatility* is estimated as follows. For each calendar month, excess daily returns of each individual stock are regressed on the daily Fama-French three factors: *Rm-Rf, SMB,* and *HML*. The monthly idiosyncratic volatility of the stock is the product of the standard deviation of the regression residuals and the square root of the number of observations in the month. *analyst coverage*$_i$ indicates the number of analysts formally following the firm; β_0 and α_0 are intercepts; and ε_t and ν_t are error terms. The regressions include institution-specific and calendar year-specific variables. The analysis is based on 3.97 million observations between 1999 and 2005. Statistical significance is indicated by **, * for one and five percent levels, respectively.

Panel A: Order duration and Number of Broker first stage regression

Order Characteristics	Duration	
Intercept	0.109	**
Informed Order	-0.211	*
Complexity of Order	2.656	**
Liquidity Demander	0.101	**
Nasdaq dummy	0.068	**
Market Capitalization	-0.937	**
Stock Volatility	0.285	**
Fixed Effects Included	Yes	
Adjusted R-Square	0.04%	

Order Characteristics	Number of Broker	
Intercept	1.031	**
Informed Order	-0.061	*
Complexity of Order	6.964	**
Liquidity Demander	0.012	**
Nasdaq dummy	0.001	**
Market Capitalization	-0.147	**
Stock Volatility	0.020	**
Fixed Effects Included	Yes	
Adjusted R-Square	0.02%	

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	0.629	**	-0.182	**	-0.597	**	0.052	**	-0.098	**
Complexity of Order	1.482	**	0.668	**	0.616	**	-0.034		2.733	**
Duration	0.595	**	-0.184	**	-0.586	**	0.050	**	-0.124	**
Number of Broker	4.351	**	1.646	**	1.227	**	-0.054	**	7.171	**
Liquidity Demander	-0.659	**	0.375	**	1.211	**	-0.002	**	0.925	**
Nasdaq dummy	-0.046	**	0.059	**	0.050	**	0.013	**	0.076	**
Market Capitalization	1.190	**	-0.185	**	0.066	*	0.010	**	-1.298	**
Stock Volatility	-0.198	**	-0.034	**	0.171	**	0.065	**	0.400	**
Analyst Coverage	0.004	**	0.001	*	-0.001	**	-0.001	**	0.003	**
Index Stock	-0.036	**	-0.013	**	-0.034	**	-0.015	**	-0.098	**
Idiosyncratic Risk	-0.100	**	0.009	**	0.005	**	0.006	**	-0.080	**
Fixed Effects Included	Yes		Yes		Yes		Yes		Yes	
Adjusted R-Square	0.50%		0.35%		11.26%		10.74%		0.44%	

Figure 1. Order Volume, Duration and Number of Brokers

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

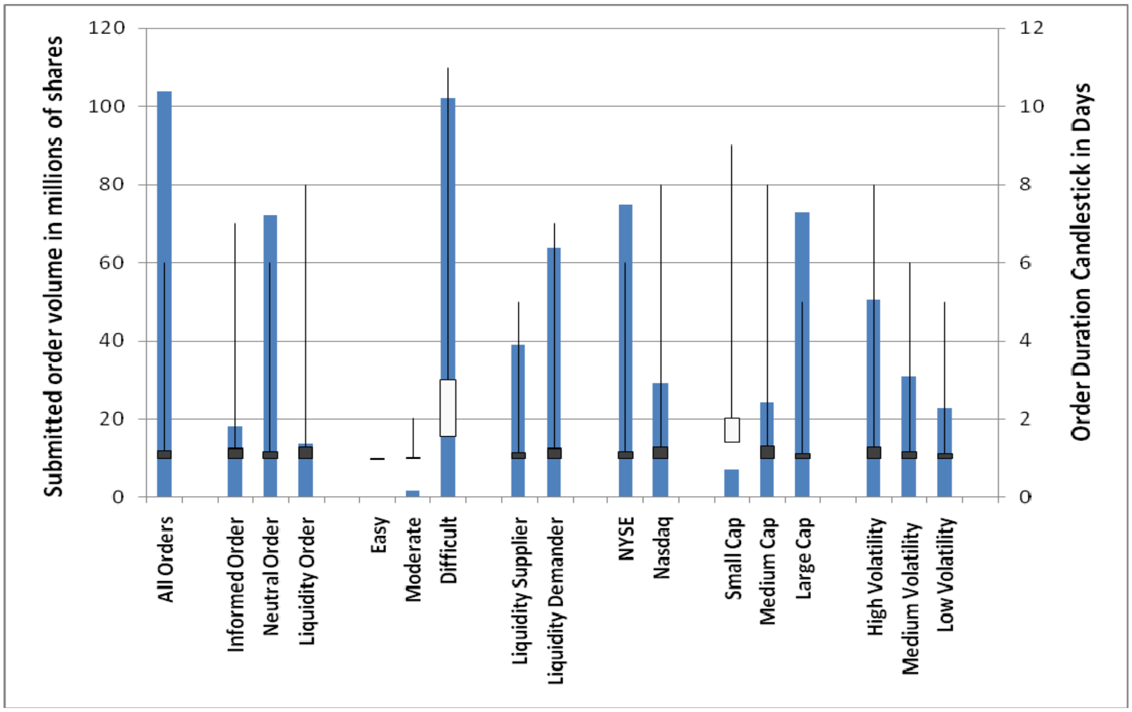

Panel B. Number of Brokers

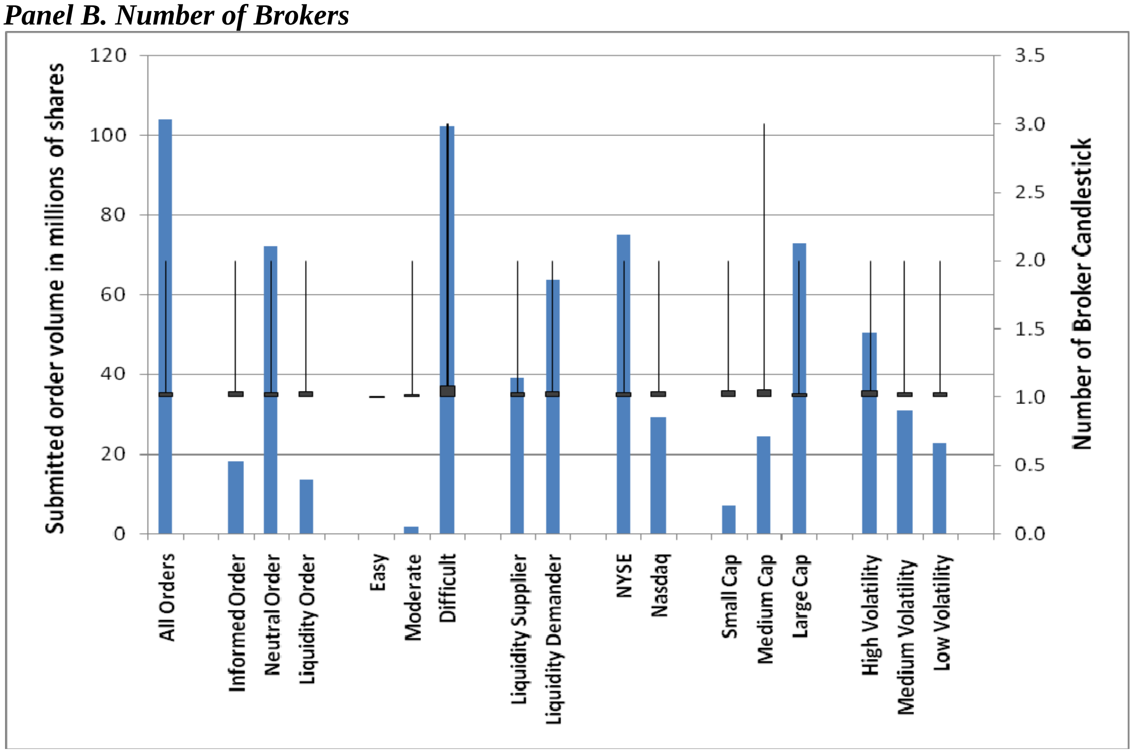

Figure 2. Components of Institutional Trading Frictions

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

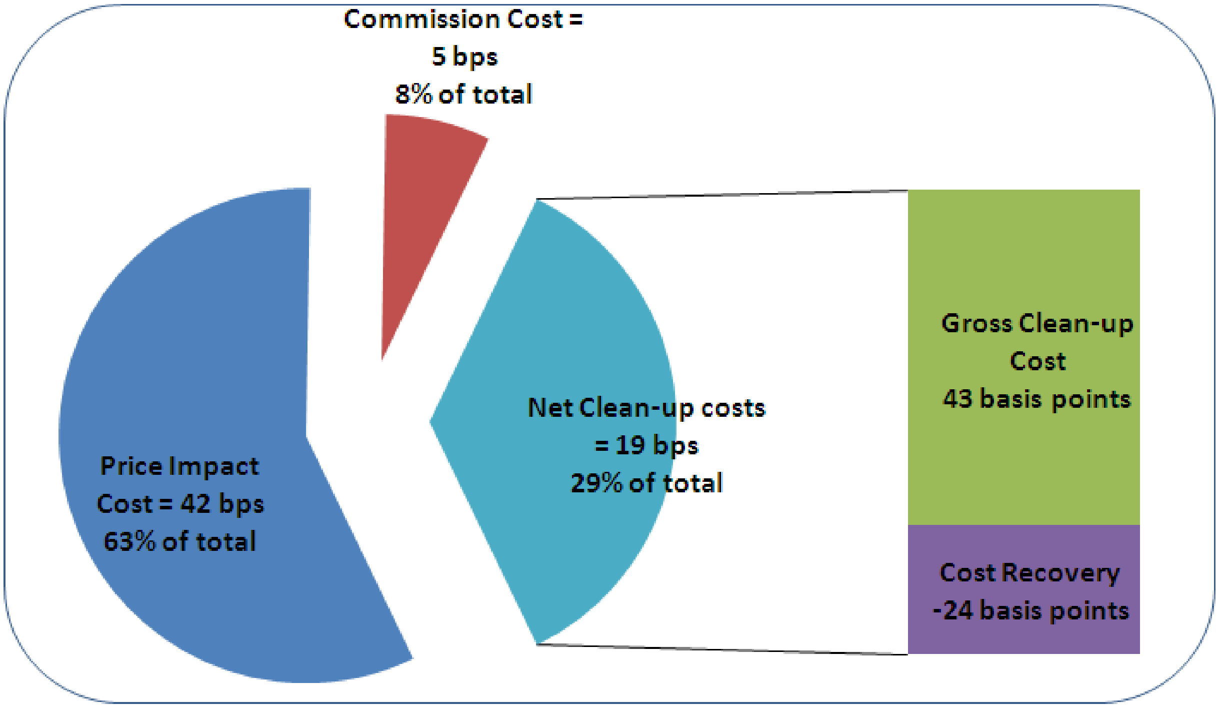

Figure 3. Persistence in clean-up costs over time

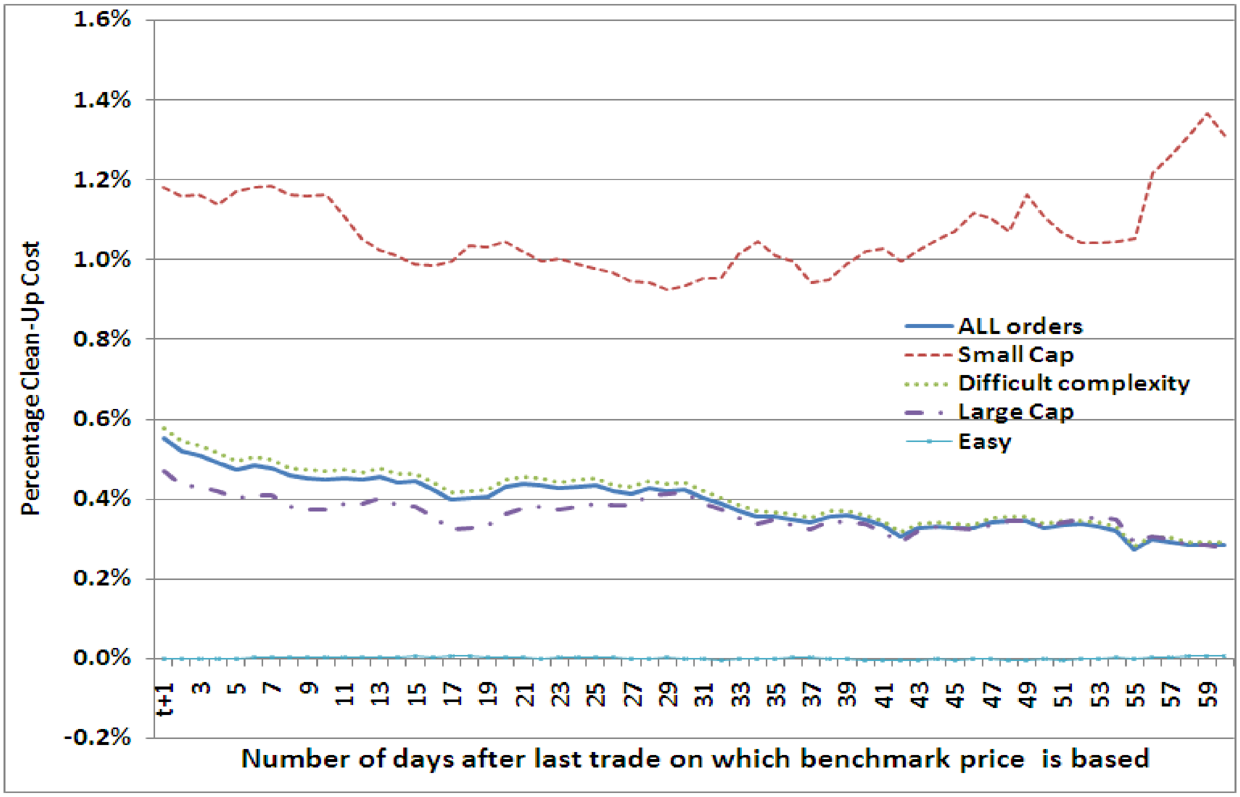

Figure 4. Clean-up cost and cost-recovery distribution

In each picture, the horizontal axis shows transaction cost ranges in 1% intervals. We consolidate the extreme categories by clubbing the orders that have less than -20% or more than 20% cost. The vertical axis plots the proportion of all orders with a specific characteristic that fall in that transaction cost range. Clean-up cost graphs are on the left and cost recovery on the right.

Panel A: Firm size categories

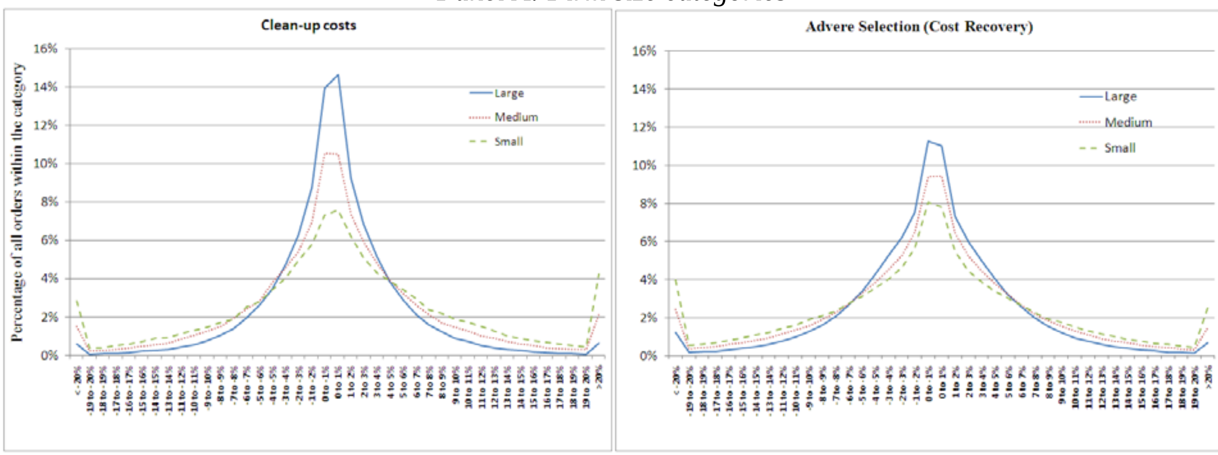

Panel B: Liquidity supplier versus liqudity demander

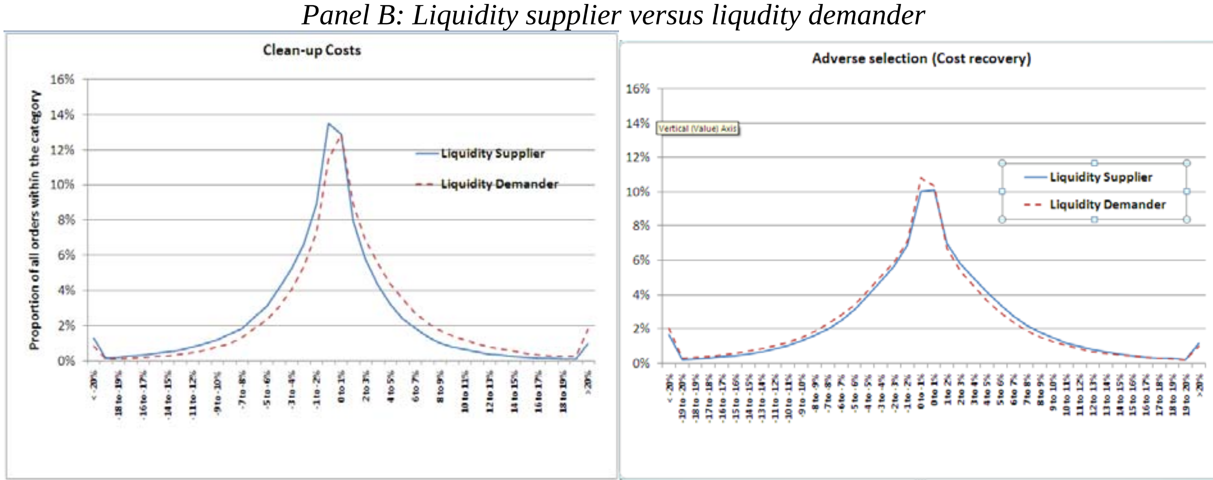

Panel C: Duration