Are Retail Orders Different?

Charles M. Jones
Graduate School of Business
Columbia University

and

Marc L. Lipson
Department of Banking and Finance
Terry College of Business
University of Georgia

First draft: April, 2003
This draft: February, 2005

We thank Ekkehart Boehmer, Rich Lyons, Gideon Saar, Patrik Sandas, Kumar Venkataraman, and participants in the 2003 NBER Summer Institute for helpful comments. An earlier version of this paper was titled "Why NYSE retail orders get better (yes, better) executions." Part of this research was undertaken while Jones was the visiting economist at the New York Stock Exchange, and the NYSE provided the data as well as financial support for this study. The comments, opinions, and errors are those of the authors only. In particular, the views expressed here do not necessarily reflect those of the directors, members, or officers of the New York Stock Exchange, Inc. Contact Charles Jones at Columbia Business School, Columbia University, New York, NY, 10027, or at cj88@columbia.edu, or contact Marc Lipson at the Terry College of Business, University of Georgia, Athens, GA, 30602, or at mlipson@terry.uga.edu.

Are Retail Orders Different?

Abstract

We use proprietary order-level data on New York Stock Exchange (NYSE) trading to examine how trading costs and price discovery evolve and may be related in a market where heterogeneous sources of order flow are present. We find that retail and non-retail orders do not have the same average execution costs. Effective spreads for retail orders are smaller than effective spreads for similar orders originating from institutions or program trades. The principal explanation is that non-retail order flow appears to be less correlated with information flows. We find that some of the initial price response to retail order flow is reversed following an execution as a result of an inflow of institutional orders in the opposite direction. Finally, institutional and program order flow appears to take advantage of liquidity changes, jumping in when spreads narrow, while retail order flow does not. Our results suggest that differences in the timing of order initiation across order types leads to remarkably different average execution results and relations between order flows.

The dynamics of trading costs and price discovery
in markets with heterogeneous order flow

1. Introduction

Not all order flow is created equal. Order flow originating from retail traders is sufficiently desirable that some market centers pay for it. Institutions often execute order over time and employ sophisticated strategies to mask their activities. Program trades are generated electronically from predetermined algorithms that respond almost immediately to price changes. Evidence on these order flow characteristics has two sources. Some studies document order flow characteristics from a single source without making comparisons. Others document differences in characteristics between market centers and attribute these to differences in the source of order flow. Despite the importance of these differences, there are few direct comparisons and no studies have explored how trading costs and price discovery evolve and may be related in a market where these heterogeneous sources of order flow are all present.

Given inherent differences in the nature of order flow, the timing of order flow is likely to vary across order sources. This difference in timing may lead to differences in execution returns, differences in price movements over time, and interrelations between order flows. We use a proprietary data set of retail, institutional and program order flow on the NYSE to investigate the following questions. First, are there differences in execution results for various order sources when these sources are all present? This question is of particular importance for retail order flow. A number of market centers specialize in retail order flow and these have been shown to provider lower average executions costs. Whether these average differences are relevant to retail order routing decisions depends on whether retail order flow obtains better executions than average on venues that do not specialize.[1] Second, do subsequent prices distinguish between retail and institutional order flow? Given that retail order flow generally contains less price

[1] Securities and Exchange Commission (SEC) Rule 11Ac1-5 specifically requires the public dissemination of average execution statistics in order to facilitate routing decisions: "One of the primary objectives of the Rule is to generate statistical measures of execution quality that provide a fair and useful basis for

relevant information than institutional orders, at some point prices must reflect the actual difference in the informativeness of these order sources.[2] Differences in order flow characteristics may allow markets to make this distinction in a timely manner. Third, do some orders take advantage of the liquidity provision available from retail order flow? If retail order flow provides a pool of liquidity, other sources of order flow may be able to detect and capitalize on the presence of retail orders. Fourth, to what extent do various sources of liquidity appear to time short-duration changes in liquidity? Given differences in trade initiation and routing decisions across order sources, some sources may be better able to adjust to very short duration changes in liquidity.

We find that substantial differences between retail order flow and both institutional and program order flow. On average, we find that retail order flow obtains substantially more favorable executions than other order flow in our sample.[3] For example, effective spreads for retail orders in our sample are about 2.60 pennies and are a half a penny lower, on average, than effective spreads for comparable institutional orders. Retail orders also obtain better executions than orders associated with program trading and all other orders. These results are more pronounced for market orders than marketable limit orders and for smaller order sizes. Retail orders have a higher realized spread (a measure of gross trading profits to liquidity providers), which makes it clear why market centers prefer to execute these orders (see Bessembinder and Kaufman (1997) and Huang and Stoll (1996)).

---

comparisons among different market centers." The comparison of averages is legitimate only if there are no differences in execution results for various sources of order flow.

[2] See Lipson (2003), Huang (2002), and Barclay, Hendershott and McCormick (2002). A relation between execution costs and the information content of order flow has been suggested by Demsetz (1968), Glosten and Milgrom (1985), Easley and O'Hara (1987), among others. See O'Hara (1997) and our discussion below for additional details. Easley, Keifer, and O'Hara (1996) and Battalio (1997) point out that order routing agreements can be used by market centers to draw more profitable uninformed order flow (cream skimming). Chordia and Subrahmanyam (1995) and Battalio, Greene and Jennings (1997), describe the arrangements and agreements that route order flow to various market centers. Related evidence and discussions can be found in Battalio, Greene and Jennings (1997), Bloomfield and O'Hara (1998), Dutta and Madhavan (1997), Bessembinder (1999), Bessembinder (2002a).

[3] We examine a random sample of 60 stocks chosen from the most active 1,000 symbols in November of 2002. The order-level data we obtain provide particularly accurate measures of execution results since the quality measures can acknowledge the time of order submission and can, therefore, incorporate price movements that affect execution results (Harris and Hasbrouck (1996) and Bessembinder (2002b) discuss the advantages of order-level data relative to transaction data). Most importantly, our data allow us to identify the type of account associated with an order and we distinguish between retail, institutional, program and other orders.

Given the difference in execution results for retail orders, we explore the underlying causes of these differences. First, we verify that the results are not due to retail orders being treated differently. Second, we verify that the differences are not driven by variation in order flow during the day. We also find that retail orders are only modestly correlated with institutional, program and other order flows while these other order flows are much more highly correlated with each other. The explanation for generally lower effective spreads must relate to the timing of order flows.

We examine quoted spreads and price movements immediately around order execution. Non-retail order flow seems better able to time changes in liquidity. Spreads narrow markedly before a non-retail order arrives, while spreads narrow less before a retail order arrives. Clearly, liquidity timing would seem to indicate more favorable execution for non-retail orders, so it cannot explain the narrow effective spreads for retail orders.

We find no evidence that retail or institutional orders are chasing price trends. Interestingly, we find that program trades do tend to follow recent trends (with buys following price rises and sells following price declines). Most importantly, prices move dramatically during and immediately following execution. As expected, prices move on average against the order (up for buys, down for sells) and, consistent with retail orders being less informed, prices move less for retail orders. For example, between order arrival and order execution, prices move about 0.13 pennies more for institutional than retail orders. Just after execution, the difference is even larger – about 0.88 pennies. These price movement differences more than offset the slightly larger spreads at the time of order arrival for retail orders, resulting in lower effective spreads.

We examine one possible factor that would contribute to differential price response. Since more active markets are an indicator of more information flows, we look at trading volumes around order arrival and execution. Both before and after order arrival, aggregate order flow is smaller around retail orders. For example, the average share volume for system orders (electronic orders) before a retail order arrives for execution is about 3,263, which is about 458 fewer shares than for institutional orders. Thus, differences in price response may be related to the intensity of trading around execution.

We estimate Hasbrouck (1991) vector autoregressions of quote returns and net order flow by source to document the relations between order flows. As expected, we find that a unit of retail order flow has a small permanent price impact relative to non-retail order flow. We find that non-retail order flow is strongly persistent, and the steady stream of orders in one direction continues to move the price. There is little such persistence in retail order flow, so prices do not continue to move. More important, much of the initial price response to retail orders dissipates (on average) during the following ten minutes. The cause of this reversal appears to be an influx of institutional orders in the opposite direction for the first few minutes after a retail execution.

Our results show that orders originating from different sources vary in their execution results and are interrelated. In particular, even though retail orders are not distinguished as such, they obtain a reduced cost of execution. Furthermore, prices quickly adjust to the lower information content of these orders as a result of order flow from institutional traders. In general, the reduced execution cost appears to be driven by the timing of retail orders, which typically arrive at calmer times and regardless of short-term price momentum.

The remainder of the paper is organized as follows. Section 2 provides a discussion of background issues including the type of data used. Section 3 discusses our sample. Section 4 presents basic results, and Section 5 presents results in a vector autoregression framework. Section 6 concludes the paper.

2. Background

The purpose of this introduction is to provide a brief theoretical and empirical background for discussing statistical evaluation of execution quality. The first section discusses the typical spread measures employed when analyzing trade and quote data. The second section discusses the unique issues and measures associated with order level data.

2. 1 The Measurement and Determinants of Spreads

Spreads are a simple and intuitive measure of trading costs. They reflect the difference between the price at which one sells a security and the price at which one buys. From an investor's point of view, the spread quantifies the round-trip cost of acquiring and then liquidating an investment. Two spread measures are commonly used: the quoted spread and the effective spread.

The quoted spread is equal to the difference between quoted bid and ask prices, expressed either in dollars or as a percentage of the quote midpoint. Quoted spreads reflect a market center's posted willingness to trade.

In contrast, effective spreads are based on actual transaction prices. The effective spread is defined as twice the distance between the price at which an order is executed and the midpoint of a benchmark quote. The benchmark mid-quote should represent the price that would be obtained in the absence of transaction costs. In most studies that look at transaction data, the benchmark quote is the quote prevailing at the time of execution. Here, we take advantage of our order level data and use as our benchmark the quote in effect at the time of order arrival. Effective spreads measure realized execution costs and differ from quoted spreads due to price or depth improvement. Effective spreads also vary with characteristics of the order, such as order size. This variation cannot be easily reflected in a single quoted spread number.

Both effective and quoted spreads vary over time and across securities and depend on market conditions and stock characteristics at the time an order arrives for execution. For example, the spread may reflect the inventory risk faced by liquidity providers from holding the security at that time.[4] As mentioned, the effective spread also reflects characteristics of the order. Liquidity providers incur less risk when trading with a small order, for example, and thus spreads should vary with order size.

It should be stressed that spreads are not a perfect measure of trading costs for many reasons. For example, many orders are worked over time, and spreads cannot capture the price impact of working an order. Furthermore, spreads ignore commissions and any other market center fees or costs.[5] However, spreads are simple to measure, readily available, and are usually reasonable indicators of actual trading costs for small orders.

[4] For NYSE stocks, there are many providers of liquidity other than the specialist. In fact, the floor of the exchange encourages competition for liquidity provision. When we refer to the specialist as a liquidity provider, we mean to include all providers of liquidity.

[5] The conclusions drawn from examining spreads may actually differ from the conclusions reached with more extensive data. For example, almost all studies find that spreads decline with a reduction in tick size,

Theoretical and empirical studies tend to divide the effective spread into two spread components: the information component and the realized spread. These components are important to drawing inferences about execution quality from spread numbers.

The realized spread is the gross trading revenue to liquidity providers. The realized spread is defined as twice the signed difference between an execution price and the mid-quote five minutes after execution. This mid-quote is designed to measure the post-trade value of the security, and therefore the realized spread reflects the gross trading profit to a liquidity provider from taking the other side of an order.

The difference between the effective spread and the realized spread reflects the five-minute price impact of the order. The price impact is often referred to as the information component or adverse selection cost, as it presumably reflects the information content of the order (see, for example, Huang and Stoll (1996)). To put it another way, the liquidity provider initially receives the effective spread, loses the information component as prices move against her, and thus earns only the realized spread as gross trading revenue.

These spread components are important to understanding the characteristics of particular order flows. If an order is perceived to be more informed (whether through characteristics of the order or the time of order arrival), then the order will move prices relatively more than another order. Along the same lines, if a trading venue is earning economic rents by successfully cream-skimming uninformed order flow, realized spreads should be relatively large.

Effective spreads and realized spreads are some of the quantities mandated by SEC Rule 11Ac1-5 (Dash5). Dash5 has become a standard for evaluating execution costs at various market centers. Thus, the Dash5 approach seems particularly suited to an investigation of retail order flow, and we follow many of the conventions established by the Dash5 regulations. For example, as mentioned above, we use order arrival times to benchmark effective spreads. We also examine the set of orders for which Dash5 statistics are required. Most importantly, our data allow us to identify the type of account

but studies of order level data find little if any change (see Jones and Lipson (2002) and Goldstein and Kavajecz (2002)).

associated with an order, and this allows us to compare retail, institutional, program and other orders.

2. 2 Order Level Data

In this study, the order level data are data captured by the NYSE SuperDOT system for orders submitted electronically. Order level data have two main advantages. First, it is possible to identify many of the characteristics of executed orders, such as the account type and order type. Second, order level data allow a more accurate measure of the full cost of execution since the data reflect order arrival times, not just execution times.

Execution costs should be evaluated as much as possible conditioning on characteristics of an order. We follow the Dash5 rules and partition orders across two dimensions:

- **Order Size.** Orders are classified into four order size groups. These are indicated below along with the designation we use to describe the order size category. As with Dash5 statistics, this study does not examine orders of 10,000 shares or more.

Designation	*Order Size*
Very Small	100-499 shares
Small	500-1,999 shares
Medium	2,000-4,999 shares
Large	5,000-9,999 shares

- **Order Type.** Among other things, the order type reflects a customer's degree of urgency. In general, the more patient a customer, the lower the expected cost of execution (and the longer the expected time to execution). Dash5 distinguishes between the following order types. The definitions below apply to buy orders; sell orders are defined analogously. The applicable quote is the quote prevailing at the time of order arrival.

Order Type	*Description*
Market	No limiting price
Marketable Limit	Limit price equals or exceeds the ask
Non-Marketable Limit	Limit price is below the ask

Throughout the paper, we refer to combinations of order size and order type as a "category". In general, we report average share-weighted execution results within each category. We do not examine non-marketable limit orders. Spread measures are problematic for these orders, and Dash5 regulations do not require their publication.

Dash5 guidelines contain many provisions designed to prevent the statistics from being distorted by unusual orders. For example, orders that require special handling or have unusual restrictions are excluded. Also excluded is any portion of an order executed on a day different from when the order was placed. Orders that meet all the requirements for inclusion in the statistics are referred to as "eligible orders". We follow the NYSE implementation of Dash5 rules to identify eligible orders, and we limit our analysis to these orders.

The system data include an indicator of the account type originating the order. We partition the indicators into four groups: retail, institution, program, and other. The orders in the "other" category are generally of less interest but are included for completeness. The account type partitions are:

Account Type Designation	*Description*
Retail	Agency orders that originate from individuals
Institution	Agency orders that do not originate with individuals
Program	Orders associated with program trades.
Other	Mostly orders where NYSE members are trading as principal.

Account types are coded by the submitting broker-dealer based on a set of regulations issued by the NYSE. While they are generally unaudited, these classifications are important to the NYSE and to broker-dealers because they are required for a number of compliance issues. For example, NYSE Rule 80A suspends certain types of index arbitrage program trading on volatile trading days, and account type classifications are important for enforcing this ban. The specialist and traders on the floor do not, however, observe this account type indicator for an incoming system order. In general, these market participants observe only the type, size, and limit price (if applicable) of an order. It is possible for the specialist to research a particular order in real-time and obtain the account type as well as information about the submitting broker.

However, this takes a number of keystrokes and requires a certain amount of time, and given the pace of trading on the exchange and our conversations with specialists, we conclude that the account type indicator is seldom if ever observed before execution.

We believe we are the first academic researchers to study execution quality and order timing for these different groups. Using proprietary Nasdaq data, Griffin, Harris, and Topaloglu (2003) classify trades as either individual or institutional, but they focus instead on momentum trading at the daily horizon for each of these groups. Battalio, Hatch and Jennings (2003) examine compare retail order flow sent to a third-market dealer with similar order flow sent to the New York Stock Exchange.

3. Sample and Summary Statistics

This study examines a sample of 60 symbols for which NYSE system order data were gathered. The sample was chosen as follows. First, NYSE executed share volume for all NYSE listed common equity symbols trading above $5.00 a share was gathered for November of 2002. From this sample, the 1000 most active symbols were identified and were divided into trading volume quintiles. From the most active quintile, we chose 20 symbols at random. From each of the remaining four quintiles, we choose 10 symbols at random. Appendix A lists the symbols studied along with their November consolidated trading volume. Order level data for this sample were collected for every order in the month of November 2002 (twenty trading days).

Table 1 presents summary statistics for the sample. The statistics are given for the full sample and then separately for the 20 symbols from the most active quintile and the remaining symbols. The first part of the table describes firm and share characteristics. Note that the active symbols have a higher share price, greater market capitalization (over $34 billion on average), and by construction a much higher trading volume – over ten times more active than symbols in the less-active subsample. Note that daily trading volume is based on the consolidated tape and includes all trades at all market centers.

The second part of Table 1 describes all NYSE system orders in our sample stocks. It gives the executed share volume for all orders and for relevant partitions.[6] Note that these executed order data count buy and sell orders separately. Hence, overall volume figures should be compared to twice the consolidated volume from the first part of the table. Overall, about 36% of (twice) consolidated volume involves NYSE system orders.

The last part of Table 1 describes the Dash5 eligible orders that make up our sample. Compared to twice the consolidated volume from the first part of the table, our sample covers about 17% of total volume. These numbers are much lower because we follow the Dash5 selection criteria and limit the analysis to system market and marketable limit orders below 10,000 shares. The sample excludes large institutional orders and orders sent to floor brokers. Since the focus of the paper is retail orders, and our methodology seeks similar institutional orders as a basis for comparison, excluding these large or difficult orders should not affect the results.

About 55% of the executed shares in the sample are market orders. The remaining 45% are marketable limit orders. In addition, retail order flow represents only 4% of the executed shares in the sample. There are several reasons this percentage is so low. First, retail orders tend to be relatively small. Second, while most institutional orders and program trades are routed to the NYSE, a substantial amount of retail order flow is either internalized or channeled to alternative venues. Unfortunately, we do not have order level data on retail orders executed elsewhere. Thus, we do not know whether NYSE retail orders are similar to retail order flow that is internalized or sent to other venues. Finally, the account type codes are imperfect. Based on conversations with exchange officials, we are confident that nearly all orders marked as retail are in fact submitted by individual investors. However, some orders submitted by individual investors are not recorded as retail orders, particularly if they are executed by an NYSE member firm on behalf of another broker-dealer.

---

[6] We could also have provided results on orders rather than executions. For market orders, order volume and executed volume will be almost identical. However, for marketable limit orders, order volume will exceed executed volume since the market may move away from a marketable limit order before it is executed. Lipson (2003) provides more detailed results on system order disposition.

It is typically argued that retail order flow is less informed than other order flow. To take this to the extreme, if retail order flow arrives randomly over time and is uncorrelated with contemporaneous informed order flow, then it must be uninformed. Table 2 assesses this null hypothesis by calculating the autocorrelation of and the correlation between the net order flow of different account types. For the 60 stocks in our sample during November 2002, we aggregate all orders of a given account type that execute in the same minute and measure net order flow as the excess of buys over sells during that minute. Net order flow is measured in shares as well as orders executed. The resulting time series has 7,800 observations for each account type (390 minutes per trading day × 20 trading days).

Table 2 contains the relevant correlations and autocorrelations, and the evidence rejects the extreme null. Like other account types, retail order flow is positively autocorrelated, with a one-minute autocorrelation of 0.10. Retail order flow is also positively correlated with order flow from other account types. If measured in shares, retail order flow has a contemporaneous correlation of 0.05 with institutional order flow, and 0.06 with program trades. However, all of these correlations are extremely small, and they are only marginally statistically different from zero. Economically, retail order flow is quite close to being random over time.

Though the absolute correlation levels are different from zero, we might expect relative differences if retail order flow is less informed than other types of order flow. More precisely, we would expect non-retail order flow to be more highly correlated if the different classes of non-retail order flow are motivated by the same information flows. Table 2 shows that, indeed, retail order flow is much less correlated with other order flow. This is particularly true if we consider correlation in the number of orders rather than the number of shares. For example, different types of non-retail orders have correlations that range between 0.30 and 0.55, while the correlation of retail order flow with other account types is between 0.03 and 0.06. In addition, we find that retail orders are the least autocorrelated, and institutional orders the most, with a one-minute autocorrelation coefficient of 0.34.

Similar evidence emerges from the cross-autocorrelation of retail and non-retail order flow. Institutional, program, and other non-retail order flows have similar characteristics, while retail order flow is very different. Retail order flow has almost no predictive power for non-retail order flow in the next minute, with cross-autocorrelations between 0.027 and 0.041. Retail orders seem to lag other orders slightly, as the cross-autocorrelations between non-retail order flows and lagged retail order flow are a bit higher, ranging from 0.062 to 0.079. Of course, the correlation evidence is only suggestive and needs to be confirmed by a closer look at the execution of retail orders.

4. A Detailed Look at Retail Order Execution

4. 1 Execution Quality Measures

Table 3 presents a summary of standard execution quality statistics for our sample by account type. These are simple share-weighted averages across the whole sample. Results are presented for the whole sample, by order type, and by order size. We also indicate the total shares executed in each category.[7] Finally, we include tests of the hypothesis that the given value differs from the corresponding value for retail order flow. Throughout the paper, we conduct statistical inference by aggregating all observations on a single day and base statistical tests on the variation in the weighted time series of daily observations, thus assuming independence across days but not across orders.

For the whole sample, the average effective spread for retail orders is 2.60 cents. This compares to 3.07, 3.05 and 2.46 for institution, program, and other order types. The retail orders have reliably lower spreads than institutional orders and program trades. The differences are substantial – almost half a penny separates institutional and retail spreads. Generally, the results for realized spreads and information component are similar to those in Lipson (2003) – realized spreads are small and the information component is large. The notable difference here is that realized spreads are substantial for retail order flow. The realized spread is over a penny whereas, for example, it is negative (on average) for institution orders. This illustrates the trading revenue that might be available to a market center that can attract retail order flow. From narrow effective spreads and high realized spreads, it follows directly that retail orders have little price impact. Average price impacts are 1.38 cents for retail orders, compared to 3.22

[7] This differs from Table 1, which presents daily averages by symbol. To obtain the totals in Table 3, multiply Table 1 values by 20 (days) × 60 (symbols).

cents for institutional orders and 2.66 cents for program trades. We often refer to the price impact as the information component, because all else equal, a smaller price impact implies that retail orders are relatively more "uninformed". However, it is worth noting that these are simple averages and make no attempt to set all else equal. For example, perhaps retail orders pay smaller spreads because they are simply smaller than other orders on average.

The quoted spread at the time of order execution is reliably smaller for retail than institution orders, though reliably larger than for program and other orders. As we shall see later, these results change considerably once we apply appropriate control variables.

To begin to control for differences in order flow characteristics, we calculate execution quality measures for various partitions of the data. When we partition by order type, the results are weaker for marketable limit orders (see Peterson and Sirri (2002) for issues related to the execution costs of marketable limit orders). For example, the effective spread difference between retail and institutional order flow is about 1.20 cents for market orders, but only about 0.30 cents for marketable limit orders. It should be noted that individuals submit proportionally far fewer marketable limit orders than do the other account types – the market and marketable limit breakdown is more than 80/20 for retail orders and roughly 50/50 for other account types.

A more important control is order size. For smaller order sizes, retail effective spreads are statistically narrower. For the smallest orders of less than 500 shares, retail effective spreads average 1.69 cents, while institution orders' effective spreads average 2.57 cents. For the large orders in our sample (over 5,000 shares), there is no reliable difference in effective spreads between retail and either institution or program trades. As expected, effective spreads are increasing with order size (consistent with Easley and O'Hara (1997)).

These simple controls may not be enough. One possibility is that retail investors trade more in liquid stocks. For example, if retail orders are proportionally more likely in symbols with lower spreads, then effective spreads would be smaller. Table 4 contains the analysis with a full set of control variables. The reported numbers focus on retail orders relative to institutional orders; results for other account types are generally similar.

Table 4 presents a comparison of retail and institution orders using four control variables. Specifically, all orders are aggregated (using a share-weighted average) if they are on the same date in the same stock with the same order size category, same order type, and same account type. Pairs are formed when there are both retail and institutional orders that match along all four other dimensions, and the table reports equal-weighted averages across these pairs. Again, statistical inference is performed using the 20-day time series of these average pair-wise differences. It should be noted that we do not necessarily have observations for every category, so we also report the number of pairs in our analysis.[8]

Across all such pairs, the average effective spread for retail orders is 2.81 pennies. This is 0.50 cents less than the average for institutional orders.[9] We find that effective spreads are reliably smaller than effective spreads for institutions in every case except for the largest order size, where the differences are not statistically reliable. Once again we see that realized spreads are much larger and the information component much smaller for retail orders.[10] Finally, after controlling for stock, trading day, order type, and order size category, it appears that retail orders are submitted when the spread is relatively wide, while institutional orders are submitted when the quoted spread is 0.23 cents narrower. This could indicate that institutions are closely monitoring liquidity as it varies through time, and they pounce when the market is relatively liquid. We return to this issue later in greater detail.

4. 2 Are Retail Orders Treated Differently?

Among other things, the previous section establishes that cheaper retail executions are not an artifact of individuals trading more liquid stocks or submitting smaller orders. In this section, we address another possibility – that retail orders sent to the NYSE are actually treated differently by the specialist or other intermediaries. For

---

[8] The maximum number of pairs would be equal to 20 (days) × 60 (symbols) × 2 (order types) × 4 (order sizes) = 9,600. Thus, for all orders, we only have pairs for about half the possible categories.

[9] The magnitude of the spreads is much larger in Table 4 than Table 3 because we are equally weighting across symbols rather than share weighting. Thus, Table 4 reflects to a greater degree the conditions for smaller and less active symbols.

[10] Interpreting the magnitude of values in Tables 3 and 4 is somewhat complicated. In Table 3, the results are those that would be expected for a trader whose orders are distributed across symbols and days in line

example, Benveniste, Marcus, and Wilhelm (1992) argue that the lack of anonymity in the NYSE's floor-based market structure allows the specialist to separate relatively informed and uninformed order flow, thereby reducing adverse selection risk. Their model implies that uninformed orders should have lower trading costs, which is consistent with the results found here.

However, in the case of retail order flow, differential treatment seems unlikely, since these orders arrive at the trading post electronically, and the specialist cannot easily observe the account type indicator, though he may be able to draw some inference from, say, the size and timing of the order. However, to rule out differential treatment, we construct matched pairs of retail vs. non-retail orders that occur within 5 seconds of each other. These matched pairs are in the same symbol and are also the same order type (market or marketable limit), same direction (buy or sell), and also in the same order size category.

Results of the matched order analysis are given in Table 5. There are 3,306 order pairs that match retail and institution orders, and fewer retail orders that match the other account types. We report equal-weighted averages across all relevant pairs. The execution quality measures for retail orders are generally indistinguishable from the spreads for other account types. Retail orders have slightly lower effective spreads than matched program orders, but this difference is only marginally significant at the 10% level, and the result may be due to imperfect controls (e.g., matched orders need not be exactly the same size or arrive at exactly the same time). Overall, the evidence indicates that orders that arrive around the same time receive the same execution. Thus, it must be the case that retail orders execute at tighter spreads because they arrive at different times than other orders. Our goal in the rest of the paper is to explore market conditions before, during, and after retail order arrival.

4. 3 Time-of-day Differences

One simple possibility is that retail orders tend to trade at different times during the trading day. In general, spreads follow a U-shaped pattern during the trading day. They are higher at the start of trading, decline over the next few hours, and rise again

---

with aggregate volume for that trader type. The results in Table 4 are what a trader might expect for a near the close. If retail orders are predominantly executed in the middle of the day, then this might explain the results. Figure 1 presents the distribution of trading volume over the course of the day. Share volume is aggregated by 5-minute intervals, and the plot records the proportion of total volume in the sample that occurs during that 5-minute interval for that account type. All account types have very similar trading patterns. Retail order flow closely tracks the intraday regularities in other order flows. There are no discernible time-of-day differences in order flow.

4. 4 Quoted spreads before and after execution

Next we explore a number of possible determinants of execution quality differences. In this section we examine quoted spreads and in the next section we examine price changes.

We begin by examining conditions immediately surrounding the time of order arrival and execution. Figure 2 presents the quoted spread at 15 one-minute intervals prior to and at order arrival time, and at 15 one-minute intervals at and subsequent to order execution. The time between order arrival and execution (denoted in the graph by a gap) varies from order to order. All one-minute intervals are calculated relative to the order arrival time (for pre-arrival) and order execution time (for post-execution). The graph only includes orders that arrive later than 15 minutes after the start of trading and are executed at least 15 minutes before the close of trading.

Other than this filter, we apply control variables and aggregate orders following a procedure identical to that used for Table 4. That is, all orders are aggregated (using a share-weighted average) if they are on the same date in the same stock with the same order size category, same order type, and same account type. Pairs are formed when there are both retail and non-retail orders that match along all four other dimensions, and Figure 2 reports equal-weighted averages across these pairs. Statistical inference is performed using the daily time series of these average pair-wise differences.

Figure 2 shows that market conditions are similar 15 minutes before the order arrives. There is little difference in quoted spreads fifteen minutes before a retail vs. non-retail order. The notable feature of this graph is what happens just before retail order

---

randomly chosen symbol and trading day.

arrival. For the non-retail account types, the quoted spread declines markedly in the minutes just before order submission. In contrast, there is relatively little change in quoted spreads in the minutes before a retail order. Thus, it would appear that non-retail orders are timing their order arrivals to take advantage of changes in quoted spreads. For example, these orders may be picking off a limit order that has just arrived to narrow the spread. Retail orders, on the other hand, exhibit less liquidity timing.

At the time of order execution, quoted spreads are narrower for institutional orders than they are for similar retail orders. This matches the evidence in Table 4.

In all cases, quotes widen subsequent to order execution. For retail orders, the quotes narrow back down within a few minutes, whereas spreads do not narrow as much for non-retail orders. Once again, this is consistent with the timing of order flow to take advantage of temporary improvements in spreads. The slow decline may reflect the amount of time it takes for the book to fill back in.

Are non-retail orders simply quicker at pouncing on improved liquidity? To address this question, Table 6 looks at the time between the most recent liquidity improvement and the arrival of the market or marketable limit order for different account types. We look at the time between the last quote change and order arrival, the time since the last limit order arrival that improves the existing quote, and the time since the last quote narrowing. The general empirical strategy is the same as for Table 4. That is, all orders are aggregated (using a share-weighted average) if they are on the same date in the same stock with the same order size category, same order type, and same account type. Pairs are formed when there are both retail and non-retail orders that match along all four other dimensions, and Table 6 reports equal-weighted average times or price changes across these pairs. Statistical inference is performed using the daily time series of these average pair-wise differences.

Table 6 shows no evidence that institution or program trades are quicker at taking advantage of liquidity improvements. For example, the most recent improving limit order arrives an average of 94 seconds before a retail market order arrival, while the corresponding figure for institutional orders is almost identical at 93 seconds. There is some evidence that other (non-retail, non-institution, non-program) orders are quicker, at 83 seconds since the last improving limit order vs. 91 seconds for the matched sample of retail orders. These are mostly proprietary trades by member firms, so it makes sense that these entities would be the quickest on the trigger following an improvement in liquidity.

Overall, there is no evidence that institutions or program trades are faster at taking advantage of improved liquidity. Instead, the evidence suggests that institutions are waiting for substantial improvements in price before submitting a market order.

4. 5 Price changes before, during, and after execution

In addition to timing liquidity, perhaps some order submitters are responding to recent price changes in an effort to time the market. Also, market movements may affect the willingness of market participants to provide liquidity. For example, price movements might affect inventory holdings. We explore this possibility in Table 7 and Figure 3, where we examine price changes before order arrival, between order arrival and execution, and after execution. Table 6 breaks the order execution process into three parts that are analyzed separately:

- *Pre-Arrival* This is the five-minute period before an order arrives at the NYSE.

- *Execution* This period begins when the order arrives at the exchange and ends when the order is reported as executed. This takes an average of about 20 seconds. This interval matches the period used to calculate the effective spread.

- *Post-Execution* This is the five-minute period after an order is executed. This interval matches the period used to determine the realized spread.

We are most interested in the movement of prices around order arrival and execution. We measure this using momentum, which is defined as the average signed change in the midquote return (measured in cents) over the relevant time period.[11] Returns are signed by multiplying by 1 for a buy order and –1 for a sell order. That is, if prices are moving up during a buy order execution or down during a sell, momentum is positive. When positive momentum occurs before order execution, it reflects an adverse move in prices for the order submitter. However, when positive momentum occurs after

[11] We also examined the volatility of returns around order arrival and execution. Results are not reported, because there were no discernible patterns in volatility before, during, or after order execution.

order execution, the price move favors the order submitter. There are several possible sources of momentum during and after an order executes. The momentum could be the result of the executed order itself (reflecting prevailing market conditions), it could be due to other orders arriving at the same time, it could be due to price changes in other stocks, or it could be any other new information that causes the specialist to change the quotes.

The basic idea is to see whether some classes of traders are responding to price trends, to see whether some traders are better able to anticipate short-term price moves, and to document the extent of price responses to orders. On average, program trades in our sample are short-term trend chasers, with prices moving a statistically significant 1.26 cents in the five minutes before order arrival.[12] Institutions also trade in the direction of previous price moves, while retail buy (sell) orders tend to arrive after modest and statistically insignificant price declines (increases) averaging 0.35 cents.

To compare momentum across account types, we again use the Table 4 approach to control for the symbol traded, trade date, order type, and order size category. In terms of five-minute pre-arrival momentum, program trades are statistically distinct from retail orders. However, pre-arrival momentum for retail is not significantly different from that of institutional or other order flow.

Table 7 also reveals that the most interesting quote changes happen during execution. Between order arrival and execution, quoted prices all move in the same direction as the order (up for buys, down for sells). But the price changes are the smallest for retail orders. After controlling for stock, trading day, and order characteristics, average momentum during retail order execution is always statistically lower than average momentum for other account types. Retail vs. institutional momentum is 0.34 vs. 0.61 cents, retail vs. program momentum is 0.31 vs. 0.70 cents, and retail vs. other momentum is 0.32 vs. 0.54 cents.

These differences in price moves during execution account could be the explanation for the difference in the effective spread paid by market order and marketable limit order submitters. To see this, consider again the retail vs. institutional comparison.

The momentum numbers during execution (0.34 cents retail vs. 0.61 cents institutional) imply that this slippage contributes 0.68 cents to the (round-trip) cost of a retail trade and 1.21 cents to the cost of an institutional trade. The difference between the two is 0.53 cents, which is about the same as the 0.50 cent difference in effective spreads for these two account types from Table 4. This is also consistent with the large information component we observe for non-retail orders; interestingly, some of this information is already being incorporated into price prior to execution.

One might worry that momentum during execution might depend on the time required to execute the order. But this does not seem to explain the differences between retail and non-retail momentum. The bigger price moves in non-retail orders are not the result of large systematic differences in the time to execution. The average time from order arrival to order execution is about 20 seconds for all account types.

In the first minute after execution, Table 7 shows that prices move more for non-retail orders. For example, retail vs. institutional price moves are 1.81 vs. 2.56 cents, a statistically reliable difference. Over the next four minutes, the contrast between retail and non-retail orders becomes especially stark. Following a retail order, prices revert by 0.49 cents during this interval. In contrast, comparable institutional orders show a continued average price move of 0.53 cents in the direction of the original order. The net result over the 5-minute post-execution period is not surprising; it is simply another manifestation of the greater information component for non-retail orders found in Table 4. But the pattern of adjustment is very different, with reversion in prices only after retail order executions.

Figure 3 tells the same general story graphically. It presents the cumulative price impact (cumulative momentum) around order arrival and execution. The graph begins fifteen minutes prior to order arrival, extends fifteen minutes subsequent to order execution, and documents the price change each minute. Orders are aggregated as in Table 4; to make comparisons across types, we control for symbol, trade date, order type, and order size category. Also included is a single point that captures quote changes

---

[12] Share-weighted average momentum is calculated for all orders in the same stock on the same day with the same order type, order size category, and account type. The table reports equal-weighted averages for all non-empty classifications of a given account type.

between order arrival and execution, regardless of the elapsed time between arrival and execution.

Figure 3 shows that, in aggregate, neither retail nor institutional orders are chasing trends. The figure confirms that program trades chase recent trends, though it also indicates that these trends have been short-lived, beginning on average 10 minutes prior to the order. Figure 3 shows that institutional orders have a bigger price impact than retail orders. Figure 3 also confirms the Table 7 evidence of mean reversion in prices. While the price impact for institutional orders is permanent at least 15 minutes out, after retail orders prices tend to partially revert to their earlier levels.

Overall, Table 7 and Figure 3 tell a very interesting story. Program trades tend to be short-term trend chasers, while retail and institutional orders do not exhibit any strong trend-chasing or trend-reversing behavior on average. However, during execution, prices start to move in the direction of trade, and they move much more for institutional orders. After an institutional order, the mini-trend continues, as prices continue to move in the same direction. After a retail order, however, prices move less initially, and they tend to revert significantly over the next 10 minutes. This price reversion is an important part of the high realized spreads on retail orders at a five-minute horizon, and the evidence indicates that realized spreads on retail orders are even higher at a horizon of ten minutes post-trade.

These results indicate that, for whatever reason, retail orders tend to arrive when prices respond less dramatically to order flow. What might contribute to a differential price response? It is possible that non-retail orders arrive in more active markets. These active markets might be associated with greater information flows. Active markets might also increase the amount of inventory risk borne by the specialist or other liquidity suppliers.

To investigate this, we look at trading volumes around order arrival and execution. As in the rest of the paper, we compare similar retail vs. non-retail orders, controlling for order type, order size category, symbol, and trade date. We look at net signed trading volume (buyer-initiated less seller-initiated volume) as well as unsigned trading volume.

The results are in Table 8. Non-retail orders tend to execute at relatively active times. Both before and after order arrival, aggregate system volume is smaller around retail orders. For example, average system order volume (electronic orders) is about 3,263 shares in the minute before a retail order arrives, which is about 458 fewer shares than for institutional orders. There is a similar differential during the minute after order execution. The difference in signed volume is even more dramatic. In the minute before a retail order arrives, net signed volume averages 169 shares in the same direction, compared to 702 shares in the minute before an institutional order. This differential persists during order execution and in the minute after order execution. This confirms the evidence in Table 2. Marketable retail orders are close to random over time and are largely uncorrelated with order flow from other account types, while institutional orders tend to cluster in the same direction over short intervals of time. In addition, the unsigned volume evidence indicates that retail orders tend to arrive in calmer times. Thus, it is not surprising that prices do not adjust as strongly in response to a given retail order.

5. Vector autoregressions

In most of the previous section, we take a typical market order or marketable limit order and examine the nearby behavior of prices, spreads, and volume. Table 2 gives some hints about how order flow is related to nearby order flow but does not consider order flow and prices at the same time. In order to model the evolution of order flow and prices over time in an integrated framework, we turn in this section to a vector autoregression of trades and quotes.

Based on Hasbrouck (1991, 1996), we construct a vector autoregression that distinguishes between different types of order flow (see also, for example, Hendershott and Jones (2003)). This involves separate equations for the order flow of each account type, yielding five equations in total: a quote midpoint equation, an equation that describes the evolution of retail signed order flow, and so on for institutional, program, and other order flow. Specifically, for a given stock define x_t^I to be the sum of the signed order flow in shares (positive for market and marketable limit orders to buy and negative for sells) during the one-minute interval t for retail account types. Similarly, define x_t^I for institutional account types, x_t^P for program trades, x_t^O for other order flow, and define r_t to be the percentage change (log return) in the quote midpoint during interval t. The following VAR with five lags is estimated for each stock for each trading day:[13]

$$\begin{bmatrix} x_t^R \\ x_t^I \\ x_t^P \\ x_t^O \\ r_t \end{bmatrix} = \sum_{j=1}^{5} \Phi_j \begin{bmatrix} x_{t-j}^R \\ x_{t-j}^I \\ x_{t-j}^P \\ x_{t-j}^O \\ r_{t-j} \end{bmatrix} + \varepsilon_t \qquad (1)$$

where Φ_j is a 5 x 5 autoregressive matrix and ε_t is a 5 x 1 vector of innovations with covariance matrix Ω.

The VAR is inverted to get the vector moving average representation in order to focus on the impulse response functions to shocks in various types of order flow. Among other things, this allows us to measure the permanent price impact from a shock to each trade equation, as well as the effect of an order flow shock on later order flow of the same or different account type. As discussed in Hasbrouck (1991), this method is robust to price discreteness, lagged adjustment to information, and lagged adjustment to trades. Note that in this case, the order flow variables include only system market orders and marketable limit orders of 10,000 shares or fewer. Executed floor orders are excluded because we lack account types and order types for these executions.

We calculate the response of each variable to a unit shock in net order flow of a certain account type, assuming that all other types of order flow are zero. The unit shock is normalized to 1,000 shares, and contemporaneous quote midpoint changes are included.[14] There is a separate VAR for each trading day, so we average the impulse response curves across the 20 trading days in our sample and report the average impulse response. Estimated impulse responses are assumed independent across trading days, and

---

[13] For actively-traded stocks, the results are insensitive to the length of the interval and the number of lags estimated in the VAR. VAR estimation is limited to the most active stocks, because the lack of order flow in other stocks makes it very difficult to pin down their transition matrices.

[14] We accomplish this by working with orthogonalized residuals, where the order flow type being shocked is the penultimate variable, and the midpoint return is the last variable.

95% confidence intervals are constructed using the variability in the impulse response across days. Impulse reponses are calculated for a total of twenty minutes following the initial shock.

Figure 4 reports results for a single large stock, ExxonMobil. This is the third-largest American company by market capitalization, a member of the Dow Jones Industrial Average, and the third most-active stock by share volume during November 2002. Its VAR results are also representative of the broader sample of active stocks.

Figures 4a, 4b, 4c, and 4d give impulse response functions for shocks to retail order flow, institutional order flow, program order flow, and other order flow, respectively. Non-retail order flow is qualitatively similar. The strongest finding is that own order flow shocks persist over time. For example, a 1,000 share institutional buy tends to be followed by institutional purchases totaling an additional 642 shares over the next 20 minutes. Effects across order types tend to be much weaker. For example, a shock to institutional order flow alone does not tend to be followed by order flow in the same direction from other account types.

The same size trade has very different permanent price impacts for different account types. The permanent price impact is 0.13 basis points for a retail order flow shock of 1,000 shares, 0.64 basis points for an institutional order, and 0.44 basis points for a program order flow shock. The retail price impact is statistically distinct from the other two.

To help us understand why retail price impacts are so low, Figure 4a shows the response to a unit shock in retail order flow. Unlike institutional and program order flow, there is much less persistence in retail order flow. On average, a 1,000-share buy order is followed by only about 40 additional retail shares in the same direction over the next 20 minutes. The cumulative price response shows an initial price move of about one-half basis point in the direction of the trade. Only a little order flow follows in the same direction, so it is not surprising that prices do not continue to adjust in the same direction. In fact, the initial price move reverses quickly, with more than half of the initial move reversed over the next three minutes.

Why does this reversal take place? The answer lies in institutional order flow. In the first five minutes following a retail order execution, institutional order flow arrives in the opposite direction. This institutional order flow is fairly substantial: an unexpected retail order of 1,000 shares is followed by more than 400 institutional shares in the opposite direction. This countervailing order flow is significant and continues in the same direction for the entire twenty minute period studied. We cannot, of course, be sure why institutions are trading in the opposite direction, but this institutional order flow appears to explain the strong temporary component in the cumulative price response.

It is also worth noting that there is also a small reversal following program order flow (Figure 4c). This too appears to be driven by institutional order flow in the other direction, though the magnitudes are smaller. A shock of 1,000 shares in program order flow tends to be followed by 81 institutional shares in the opposite direction in the next two minutes, when the reversal occurs, and 227 institutional shares in the opposite direction over the next 20 minutes. However, it is important to note that program order flow is positively autocorrelated, with the unit shock of 1,000 shares followed by an average of 625 more program shares in the same direction over the next twenty minutes. This is likely to limit the effect of institutional trades in the opposite direction. In any case, program trades have substantial permanent price impacts, so they are qualitatively very different from retail orders.

Next, we report impulse response functions that are aggregated across stocks. For each of the twenty most active stocks in the sample, impulse response functions are calculated for each stock for each trading day, standardized to reflect the impact of an order flow innovation of 1,000 shares. An equal-weighted cross-sectional average impulse response function is calculated for each trading day, and these are then averaged across trading days. Time-series independence of the daily cross-sectional averages is used to conduct statistical inference.

The results are in Figures 5a, 5b, 5c, and 5d for retail, institutional, program, and other order shocks, respectively. The results are qualitatively similar to the single stock counterparts in Figure 4. For non-retail order flow, there is strong own order flow persistence, and modest positive cross-persistence in various types of non-retail order flow. Only retail order flow engenders order flow in the opposite direction. On average across these twenty stocks, an unexpected marketable order of 1,000 shares results in about 130 institutional shares in the opposite direction over the next five minutes, though there is only marginal statistical evidence that the institutional response is different from zero.

Permanent price impacts continue to differ across account type. The pooled average permanent price response to a unit shock of 1,000 shares is lowest for retail orders, at 1.33 basis points. Corresponding figures for institutional orders are 1.82 basis points and 2.37 basis points for program orders.

Figure 5a shows that the price reversal following retail orders is not unique to ExxonMobil. For the twenty active stocks, the response in quote midpoints maxes out at 1.91 basis points after one minute, and about one-third of this initial price response reverses in the next twenty minutes. Only retail order flow engenders such a price reversal.

Overall, the VAR evidence confirms that retail orders have smaller price impacts, and it confirms that the permanent price impact is much lower than the price impact one or two minutes after the order is executed. It also reveals at least part of the mechanism behind this quote reversion: market orders and marketable limit orders in the opposite direction are being sent by institutional accounts.

6. Conclusions

In this paper, we use proprietary system order data from the NYSE to examine the execution quality of NYSE retail order flow. It turns out that retail orders get better executions, on average, than similar non-retail orders. Effective spreads for retail orders are smaller than effective spreads for comparable orders originating from institutions, program trades, or other sources. Nevertheless, retail orders have larger realized spreads, which explains why other market centers are trying to siphon off these orders. This also implies that retail orders have a smaller price impact, which we confirm using impulse response evidence from vector autoregressions.

We rule out a number of explanations for these results. Retail orders are not treated any differently; comparable retail and non-retail orders that arrive at nearly the same time obtain similar executions. Retail and non-retail orders are distributed similarly throughout the day. The results are not driven by differences in quoted spreads at the time of execution, which are actually slightly larger, on average, when retail order flow arrives. In fact, we find that non-retail orders are able to time liquidity, jumping in when quoted spreads narrow substantially. But this effect goes the wrong way, so it cannot explain lower effective spreads for retail orders. Finally, neither institutional nor retail orders are chasing price trends, on average (though program trades do tend to chase them).

The explanation appears to be related to two important differences between retail and institutional orders. First, prices tend to rise (fall) immediately after any kind of buy (sell) order is executed, but the price reaction is smaller for retail orders. There is also a temporary component. For ten minutes after a retail execution, prices tend to partially revert toward their earlier levels. Vector autoregressions reveal that this reversion is at least partially due to institutional order flow in the opposite direction in the first few minutes following a retail order arrival. Second, retail orders seem to arrive at relatively calm times. There is more volume both before and after a non-retail order execution.

Most of this paper focuses on the search for what makes retail order flow different. But the stark differences in retail vs. non-retail order execution quality have important policy implications. Most importantly, Dash5 statistics may not provide sufficient information for routing retail order flow. For example, it is misleading to compare aggregate NYSE execution quality to that of market centers that execute predominantly retail order flow. Unfortunately, only aggregate statistics are required under Dash5 rules, and this promotes "apples-to-bicycles" comparisons. Among other things, our results suggest the New York Stock Exchange should voluntarily publish Dash5 statistics on its retail order flow so order-routers and others can draw meaningful comparisons between the NYSE and retail-oriented market centers.

References

Barclay, Michael, Terrence Hendershott, and Timothy McCormick, 2003, Electronic communications networks and market quality, working paper.

Battalio, Robert, 1997, Third market broker-dealers: cost competitors or cream skimmers?, *Journal of Finance* 52, 341-352.

Battalio, Robert, Jason Greene, and Robert Jennings, 1997, Do competing specialists and preferencing dealers affect market quality?, *Review of Financial Studies* 10, 969-993.

Battalio, Robert, Brian Hatch, and Robert Jennings, 2003, All else equal? A multi-dimensional analysis of retail, market order execution quality, Journal of Financial Markets 6, 337-362.

Benveniste, Lawrence M., Alan J. Marcus, and William J. Wilhelm, 1992, What's special about the specialist?, *Journal of Financial Economics* 32, 61-86.

Bessembinder, Hendrik, 1999, Trade execution costs on NASDAQ and the NYSE: A post-reform comparison, *Journal of Financial and Quantitative Analysis*, 34, 387-407.

Bessembinder, Hendrik, 2003a, Quote-based competition and trade execution costs in NYSE-listed stocks, *Journal of Financial Economics*, forthcoming.

Bessembinder, Hendrik, 2003b, Issues in assessing trade execution costs, *Journal of Financial Markets*, 6, 233-257.

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

Chordia, Tarun, and Avinidhar Subramanyam, 1995, Market making, tick size, and payment-for-order flow: Theory and evidence, *Journal of Business* 68, 543-575.

Demsetz, Harold, 1968, The cost of transacting, *Quarterly Journal of Economics* 82(1), 33-53

Dutta, Prajit, and Ananth Madhavan, 1997, Competition and collusion in dealer markets, *Journal of Finance* 52, 245-276.

Easley, David, N. Kiefer, and Maureen O'Hara, 1996, Cream skimming or profit sharing? The curious role of purchased order flow, *Journal of Finance* 51, 811-834.

Easley, David, and Maureen O'Hara, 1987, Price, trade size and information in securities markets, *Journal of Financial Economics* 19, 69-90.

Glosten, Lawrence, and P. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 13, 71-100.

Goldstein, Michael, and Kavajecz, 2000, Eighths, sixteenths and market depth: Changes in tick size and liquidity provision on the NYSE, *Journal of Financial Economics* 56, 125-149.

Griffin, John M., Jeffrey H. Harris, and Selim Topaloglu, 2003, The dynamics of institutional and individual trading, *Journal of Finance*, forthcoming.

Harris, Lawrence and Joel Hasbrouck, 1996, Market vs. limit orders: The SuperDOT evidence of order submission strategy, *Journal of Financial and Quantitative Analysis* 31, 213-231.

Hasbrouck, Joel, 1991, Measuring the information content of stock trades, *Journal of Finance* 46, 179-207.

Hasbrouck, Joel, 1995, One security, many markets: Determining the contributions to price discovery, *Journal of Finance* 50, 1175-1199.

Hasbrouck, Joel, 1996, Order characteristics and stock price evolution: an application to program trading, *Journal of Financial Economics* 41:129-149.

Hendershott, Terrence and Charles M. Jones, 2003, Island goes dark: transparency, fragmentation, and liquidity externalities, Columbia University working paper.

Huang, Roger, 2002, The quality of ECN and Nasdaq market maker quotes, *Journal of Finance*, 57, 1285-1319.

Huang, Roger, and Hans Stoll, 1996, Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE, *Journal of Financial Economics* 41, 313-358.

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

Lipson, Marc L., 2003, Competition among market centers, NYSE working paper.

O'Hara, Maureen, 1997, "Market Microstructure Theory", Blackwell Publishers, Cambridge, Mass.

Peterson, Mark, and Erik Sirri, 2002, Order submission strategy and the curious case of marketable limit orders, *Journal of Financial and Quantitative Analysis* 37.

Appendix A
List of symbols studied.

Symbol	November 2002 Trading Volume	Name
AMD	291,517,400	ADVANCED MICRO DEVICES INC
HI	271,039,900	HOUSEHOLD INTERNATIONAL INC
XOM	224,264,100	EXXON MOBIL CORP
CD	102,219,600	CENDANT CORP
FNM	86,419,500	FEDERAL NATIONAL MORTGAGE ASSN
UNH	85,477,300	UNITEDHEALTH GROUP INC
SWY	80,895,100	SAFEWAY INC
ABT	77,328,900	ABBOTT LABS
WM	69,010,300	WASHINGTON MUTUAL INC
G	63,427,700	GILLETTE CO
ABC	51,175,400	AMERISOURCEBERGEN CORP
TJX	50,004,900	T J X COMPANIES INC NEW
DAL	45,435,800	DELTA AIR LINES INC
SLE	44,582,300	SARA LEE CORP
PRU	44,418,900	PRUDENTIAL FINANCIAL INC
ACS	40,492,300	AFFILIATED COMPUTER SERVICES INC
KFT	39,333,500	KRAFT FOODS INC
CAT	35,640,200	CATERPILLAR INC
OHP	33,217,100	OXFORD HEALTH PLANS INC
COX	32,347,000	COX COMMUNICATIONS INC NEW
Z	29,385,700	FOOT LOCKER INC
CMS	28,927,100	C M S ENERGY CORP
PFG	26,335,800	PRINCIPAL FINANCIAL GROUP INC
ETR	20,167,300	ENTERGY CORP NEW
BRO	15,833,800	BROWN & BROWN INC
CTL	15,831,000	CENTURYTEL INC
ROK	14,982,500	ROCKWELL INTERNATIONAL CORP NEW
SHW	14,758,300	SHERWIN WILLIAMS CO
PTV	14,656,800	PACTIV CORP
TXT	12,728,800	TEXTRON INC
GTK	12,108,600	GTECH HOLDINGS CORP
AW	11,699,400	ALLIED WASTE INDUSTRIES INC
TCB	11,618,500	T C F FINANCIAL CORP
PPD	10,734,000	PRE PAID LEGAL SERVICES INC
DST	9,431,600	D S T SYSTEMS INC DEL
NCF	7,972,800	NATIONAL COMMERCE FINANCIAL CORP
TEX	7,855,700	TEREX CORP NEW
ATI	6,688,700	ALLEGHENY TECHNOLOGIES
ION	6,155,900	IONICS INC
MW	6,136,200	MENS WAREHOUSE INC
PER	4,355,000	PEROT SYSTEMS CORP
HGR	4,280,400	HANGER ORTHOPEDIC GROUP INC
GVA	4,274,900	GRANITE CONSTRUCTION INC
EV	3,999,200	EATON VANCE CORP
GAS	3,899,900	NICOR INC
CXR	3,808,300	COX RADIO INC
NUI	3,496,600	N U I CORP NEW
BTU	3,266,300	PEABODY ENERGY CORP
PNM	3,188,800	P N M RESOURCES INC
GPN	3,017,500	GLOBAL PAYMENTS INC
BBX	2,808,700	BANKATLANTIC BANCORP INC
BKH	2,449,400	BLACK HILLS CORP
CBM	2,433,200	CAMBREX CORP
HAE	2,167,700	HAEMONETICS CORP MASS
BWS	2,155,800	BROWN SHOE INC NEW
KCP	2,095,800	COLE KENNETH PRODUCTIONS INC
KFY	1,636,300	KORN FERRY INTERNATIONAL
MHO	1,449,400	M I SCHOTTENSTEIN HOMES INC NEW
BKI	1,227,000	BUCKEYE TECHNOLOGIES INC
AIT	1,178,300	APPLIED INDUSTRIAL TECHS INC

Table 1
Summary statistics

The sample combines the 20 most active symbols for the month of November 2002 (measured by consolidated trading volume), plus a stratified random sample of 40 additional symbols. All symbols are common equity with a trade-weighted price of at least $5.00 during November. Dash-5 eligible trades represent SuperDot executions of market and marketable limit orders of 9,999 shares or fewer.

Symbol Characteristics	**Full Sample**	**Active 20**	**Remaining 40**
Price (dollars)	26.64	35.38	22.27
Shares Outstanding (thousands)	344,074	869,426	81,399
Market Value (thousands of dollars)	12,921,236	34,389,762	2,186,973
Consolidated Daily Volume (shares)	1,757,870	4,420,618	426,496
All NYSE System Trading Activity (daily average shares executed)			
All Orders	1,262,357	3,095,893	345,590
Market Orders	479,732	1,229,213	104,992
Marketable Limit Orders	404,086	959,080	126,590
Retail Orders	38,501	97,831	8,835
Orders from Institutions	691,991	1,720,103	177,935
Program Trades	356,980	842,435	114,252
Other Orders	174,886	435,523	44,568
Dash-5 Eligible NYSE System Trading Activity (daily average shares executed)			
All Orders	599,952	1,466,543	166,657
Market Orders	330,388	849,474	70,845
Marketable Limit Orders	269,565	617,069	95,812
Retail Orders	22,081	54,802	5,721
Orders from Institutions	328,285	812,025	86,414
Program Trades	190,549	451,533	60,056
Other Orders	59,038	148,183	14,466

Table 2
Correlation of signed order flow for one-minute intervals

Signed order flow is the net of buys minus sells (measured using the number of shares or the number of orders) for all market orders and marketable limit orders of less than 10,000 shares, aggregated across all stocks in the sample over one-minute intervals. Inference assumes time-series independence.

	Retail	Institution	Program	Other
Autocorrelation (Shares)	0.1026***	0.3443***	0.3638***	0.2810***
Contemporaneous Correlation (shares)				
Institution	0.0537**			
Program	0.0635*	0.5512***		
Other	0.0371**	0.3516***	0.2935***	
Contemporaneous Correlation (orders)				
Institution	0.0347**			
Program	0.0522*	0.5130***		
Other	0.0388**	0.3043***	0.2665***	
Cross-Autocorrelation (shares)				
Lagged Retail	0.1026***	0.0351**	0.0273***	0.0410***
Lagged Institution	0.0668***	0.3443***	0.2421***	0.2488***
Lagged Program	0.0787***	0.2559***	0.3638***	0.2187***
Lagged Other	0.0621***	0.1736***	0.1123***	0.2810***

*, **, and *** indicate significance at the 10%, 5% and 1% levels, respectively.

Table 3
Transaction cost measures by account type

Standard trading cost measures for the entire sample and selected partitions. Values are in pennies and are share-weighted across all observations. For each account type, we test whether the given value differs from the corresponding value for retail orders. Statistical tests are based on the daily time series of share-weighted averages.

			Spread Decomposition			
		Shares (1,000)	Effective Spread	Realized Spread	Information Component	Quoted Spread
All Orders						
	Retail	26,497	2.60	1.22	1.38	3.04
	Institution	393,941	3.07***	-0.15***	3.22***	3.19***
	Program	228,658	3.05**	0.39***	2.66***	2.78***
	Other	70,846	2.46	0.11***	2.34**	2.93***
By Order Type						
Market Orders	*Retail*	21,908	2.82	1.13	1.69	3.12
	Institution	217,028	4.09***	-0.06***	4.15***	3.66***
	Program	121,339	4.38***	0.95	3.44***	3.31**
	Other	36,190	3.38**	0.11***	3.27***	3.48***
Marketable Limit Orders	*Retail*	4,589	1.53	1.63	-0.10	2.66
	Institution	176,913	1.83**	-0.25***	2.07***	2.62
	Program	107,319	1.55	-0.24***	1.79***	2.18***
	Other	34,656	1.49	0.12**	1.37**	2.35***
By Order Size						
Very Small (100 – 499 shs)	*Retail*	5,927	1.69	1.10	0.59	3.24
	Institution	85,411	2.57***	-0.32***	2.89***	3.36**
	Program	77,997	2.93***	-0.26***	3.20***	3.06***
	Other	12,719	2.38***	0.10***	2.28***	3.36*
Small (500 – 1,999)	*Retail*	10,448	2.39	1.09	1.30	3.09
	Institution	165,176	3.11***	-0.58***	3.69***	3.28**
	Program	100,532	2.80**	0.15***	2.65***	2.61***
	Other	29,760	2.52	-0.11***	2.63***	3.11
Medium (2,000 – 4,999)	*Retail*	6,265	2.99	1.08	1.91	2.87
	Institution	86,251	3.17	0.31*	2.85**	3.08
	Program	38,207	3.71**	1.82	1.90	2.69**
	Other	16,104	2.47	-0.20***	2.68	2.60***
Large (5,000 – 9,999)	*Retail*	3,857	3.92	1.96	1.96	2.88
	Institution	57,104	3.56	0.67**	2.90	2.87
	Program	11,922	3.87	2.11	1.76	2.66
	Other	12,263	2.36***	1.08	1.28	2.48***

*, **, and *** indicate significance at the 10%, 5% and 1% levels, respectively.

Table 4
Differences Between Retail and Institutional Orders

All orders are aggregated (share-weighted average) if they are on the same date in the same stock with the same order size category, same order type, and same account type. Pairs are formed when there are both retail and institutional orders that match along all four other dimensions, and the table reports equal-weighted averages or average differences across these pairs. The reported difference is the retail value minus the institutional value. Statistical tests are based on the time series of daily averages.

			Spread Decomposition			
		Category Pairs	Effective Spread	Realized Spread	Information Component	Quoted Spread
All Orders						
	Retail		2.72	0.96	1.76	3.58
	Difference	4,388	-0.50***	1.57***	-2.06***	0.23***
By Order Type						
Market Orders	*Retail*		3.33	1.18	2.15	3.92
	Difference	2,819	-0.61***	1.66***	-2.27***	0.12*
Marketable Limit Orders	*Retail*		1.63	0.57	1.06	2.97
	Difference	1,569	-0.30**	1.40***	-1.70***	0.41***
By Order Size						
Very Small (100 – 499 shs)	*Retail*		1.95	0.92	1.03	3.71
	Difference	1,619	-0.65***	1.42***	-2.07***	0.29***
Small (500 – 1,999)	*Retail*		3.04	0.79	2.26	3.73
	Difference	1,548	-0.27***	1.72***	-2.00***	0.22**
Medium (2,000 – 4,999)	*Retail*		2.98	0.94	2.05	3.37
	Difference	801	-0.70**	1.11**	-1.81***	0.24**
Large (5,000 – 9,999)	*Retail*		4.02	1.84	2.18	2.95
	Difference	420	-0.37	2.40***	-2.77***	-0.03

*, **, and *** indicate significance at the 10%, 5% and 1% levels, respectively.

Table 5
Analysis of Matched Orders

Standard execution cost measures for matched pairs of orders arriving within five seconds of each other. Matches must be the same order type (market or marketable limit), and order direction (buy or sell). The table reports averages across all matched pairs. Inference is conducted using the time series of daily average paired differences.

		Spread Decomposition			
	Matched Pairs	Effective Spread	Realized Spread	Information Component	Quoted Spread
Retail		3.269	0.103	3.167	3.528
Institution		3.288	0.149	3.139	3.518
Difference	9,705	-0.018	-0.047	0.028	0.010
Retail		3.497	0.435	3.062	3.186
Program		3.681	0.564	3.117	3.167
Difference	4,935	-0.184***	-0.129**	-0.055	0.019
Retail		3.377	0.686	2.691	3.637
Other		3.356	0.697	2.659	3.626
Difference	2,070	0.021	-0.011	0.032	0.012

*, **, and *** indicate significance at the 10%, 5% and 1% levels, respectively.

Table 6
Analysis of Order Timing

This table describes the timing of order flow relative to recent changes in quotes and the magnitude of the quote change. For a marketable buy (sell) order, the change in the relevant side is the last change in the ask (bid) price, and a negative number indicates that the terms of trade are improving. Price changes are in cents.

	Time (in seconds)			Last Quote Change	
	Since Last Quote Change	Since Last Improving Limit Order	Since Last Spread Decrease	Change in Relevant Side	Change in Spread
Retail	66.52	94.46	78.95	-0.33	-0.62
Institution	62.97	93.47	78.84	-0.26	-0.74
Difference	3.55	0.99	0.11	-0.07*	0.12**
Retail	65.38	92.96	78.00	-0.34	-0.63
Program	65.24	93.71	78.68	-0.17	-0.71
Difference	0.14	-0.75	-0.68	-0.17***	0.08
Retail	64.39	91.33	77.88	-0.34	-0.62
Other	56.54	82.55	69.91	-0.18	-0.61
Difference	7.85***	8.78***	7.97*	-0.16***	-0.01

*, **, and *** indicate significance at the 10%, 5% and 1% levels, respectively.

Table 7
Momentum Analysis

Price momentum around execution and duration of order executions. Momentum is the price change in the same stock over the specified interval, signed by the direction of the order. Momentum is measured in cents using quote midpoints and is positive if price is moving up around a buy or down around a sell. Comparisons across account types use the approach described in Table 4, which controls for stock, trading day, order type, and order size category. Statistical tests are based on the time series of daily averages.

	Pre Arrival	**Execution**	**Post Execution**	
	5 Minutes Before Arrival	*Arrival to Execution*	*1 Minute After Execution*	*Next 4 Minutes After Execution*
<u>ALL ORDERS</u>				
Momentum (tests are against null of zero)				
Retail	0.104	0.348***	2.150***	-0.546***
Institution	0.409**	0.632***	3.459***	0.373***
Program	1.839***	0.574***	2.711***	0.079
Other	-0.257	0.508***	2.622***	0.010
Average time from arrival to execution (in seconds)				
Retail		22.49		
Institution		21.46		
Program		18.04		
Other		23.66		
<u>COMPARABLE ORDERS ONLY</u>				
Retail vs Institutional				
Retail Momentum	0.162	0.341	1.811	-0.491
Institution Momentum	0.519	0.605	2.556	0.527
Difference	-0.357	-0.264***	-0.745***	-1.018***
Retail vs Program				
Retail Momentum	0.103	0.311	1.413	-0.444
Program Momentum	1.854	0.699	2.294	0.065
Difference	-1.751***	-0.388***	-0.881***	-0.510***
Retail vs Other				
Retail Momentum	0.262	0.320	1.392	-0.520
Other Momentum	-0.215	0.538	2.027	-0.069
Difference	0.477	-0.218***	-0.635***	-0.451***

Table 8
Volume Analysis

NYSE system order volume around retail and institutional order execution, in shares. Comparisons across account types use the approach described in Table 4, which controls for stock, trading day, order type, and order size category. Statistical tests are based on the time series of daily averages.

	Number of Categories	**Pre Arrival** *1 Minute Before Arrival*	**Execution** *Arrival to Execution*	**Post Execution** *1 Minute After Execution*
Retail vs Institutional				
Volume Retail		3,263	1,314	3,152
Volume Institutional		3,721	1,187	3,614
Difference	4,099	-458***	127	-462***
Net Signed Volume Retail		169	20	94
Net Signed Institutional		702	231	597
Difference	4,099	-532***	-212***	-503***
Retail vs Program				
Volume Retail		3,422	1,387	3,301
Volume Program		3,898	1,271	3,748
Difference	3,670	-476***	117	-447***
Net Signed Volume Retail		160	15	65
Net Signed Volume Program		1,108	371	783
Difference	3,670	-948***	-356***	-718***
Retail vs Other				
Volume Retail		3,537	1,440	3,422
Volume Other		3,932	1,319	3,788
Difference	3,548	-395***	121	-366***
Net Signed Volume Retail		177	23	73
Net Signed Volume Other		699	219	494
Difference	3,548	-522***	-196***	-421***

Figure 1
Trading Volume by Time of Day

Distribution of trading volume, by account type, over the course of the trading day. Chart excludes first and last 15 minutes and each point represents a five minute block.

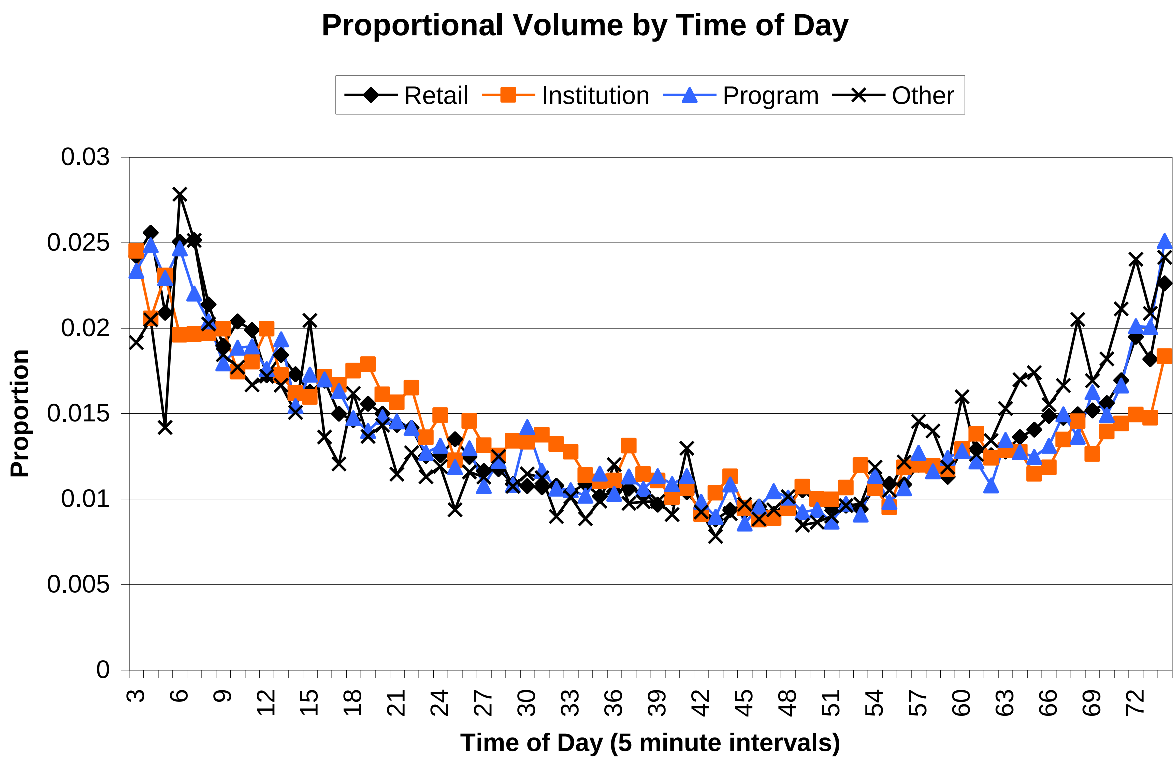

Figure 2
Quoted Spread Around Orders

Share-weighted average quoted spreads in pennies at various times before order arrival (negative numbers) and after order execution (positive numbers) Orders are aggregated and weighted using the approach in Table 4, which controls for stock, trading day, order type, and order size category.

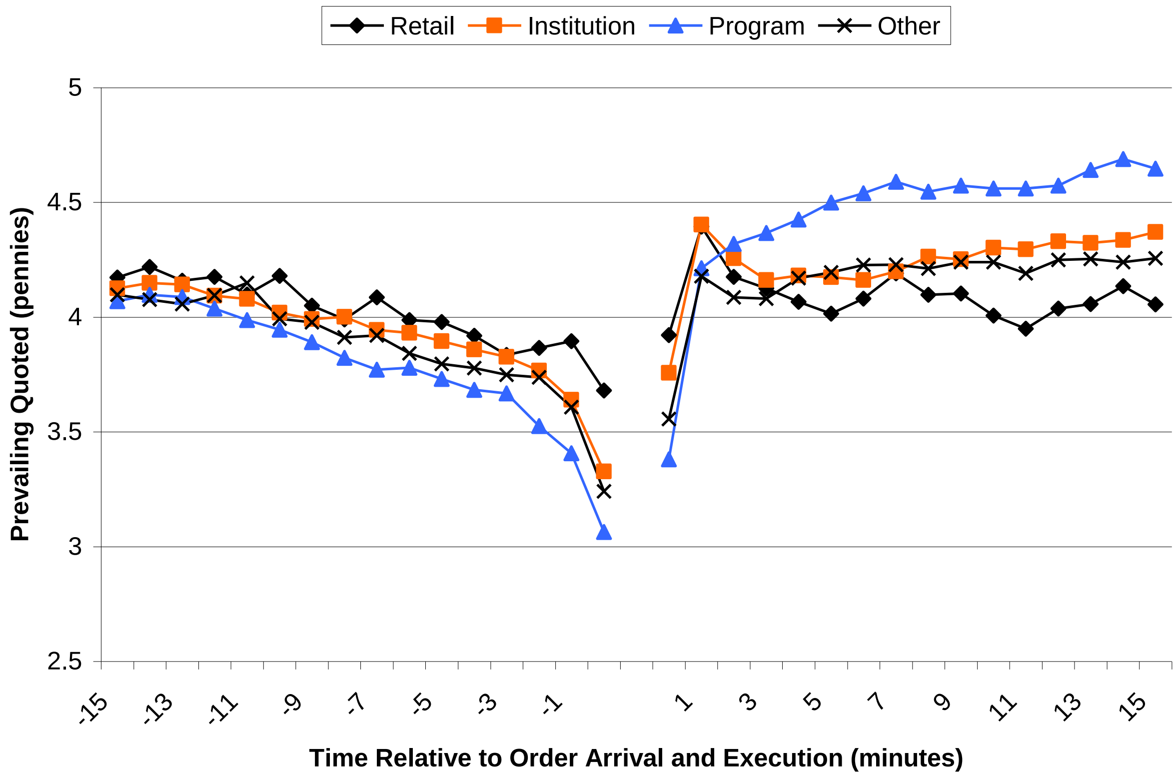

Figure 3
Cumulative Momentum

Cumulative price change over the specified interval, signed by the direction of the order. Price change or momentum is measured using quote midpoints and is positive if price is moving up around a buy or down around a sell. Single points at time zero include the earlier cumulative price changes plus the price change between order arrival and order execution. Orders are aggregated and weighted using the approach in Table 4, which controls for stock, trading day, order type, and order size category.

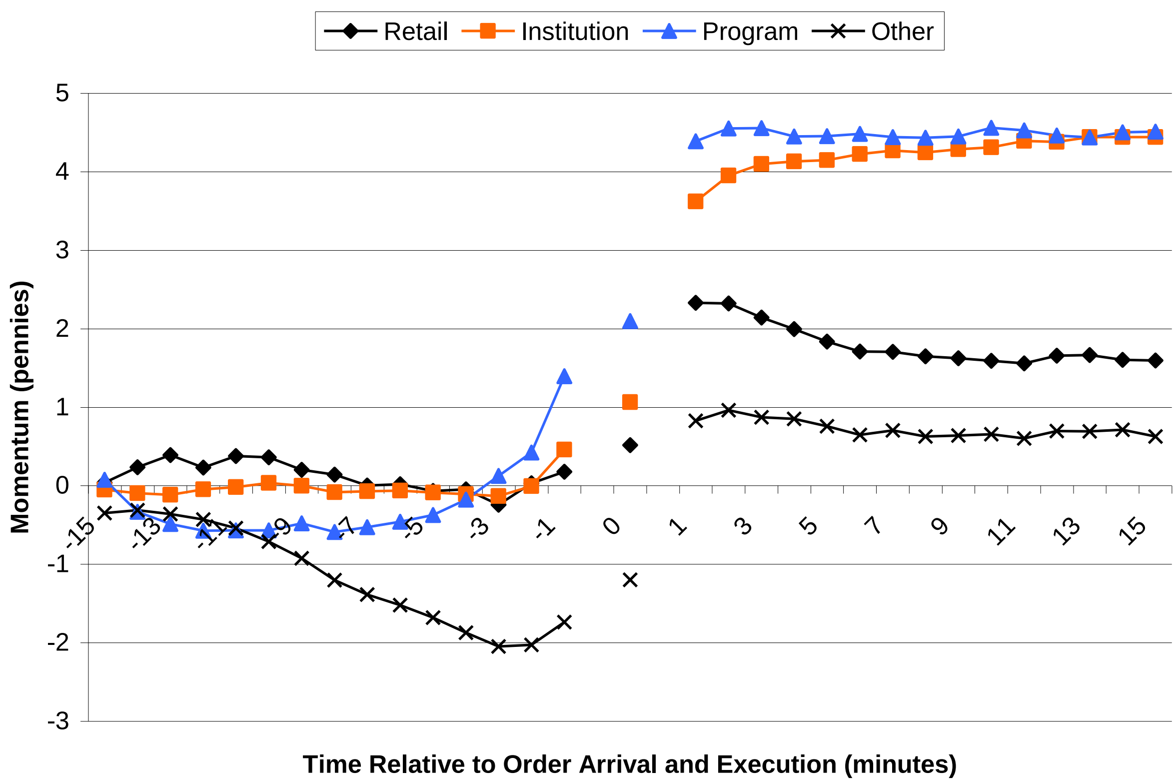

Figure 4a. Unit shock (1,000 shares) to retail net order flow in XOM

Impulse response functions for a vector autoregression in one-minute quote returns and net order flow of various account types. Confidence intervals are constructed by estimating a separate VAR and impulse response function for each trading day and then assuming independence over time. Dashed lines are two standard errors away from the average estimated impulse response and reflect approximate 95% confidence intervals.

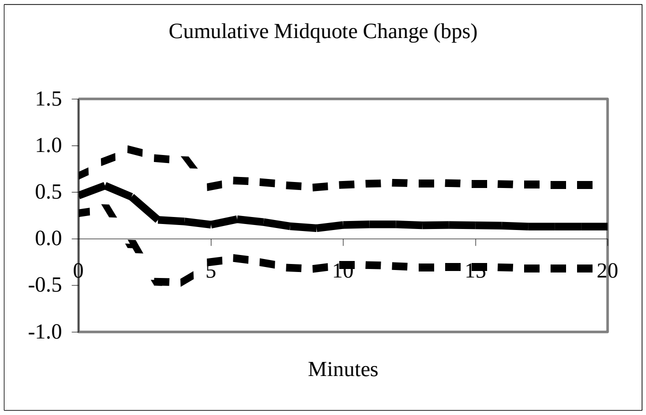

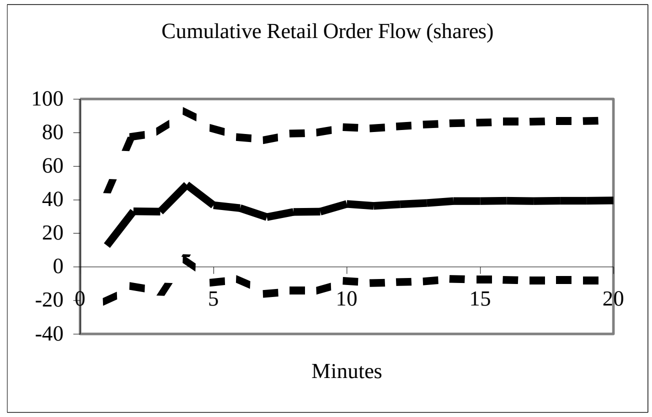

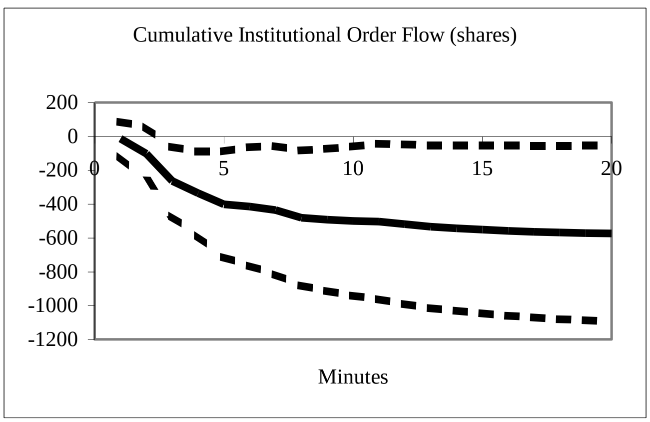

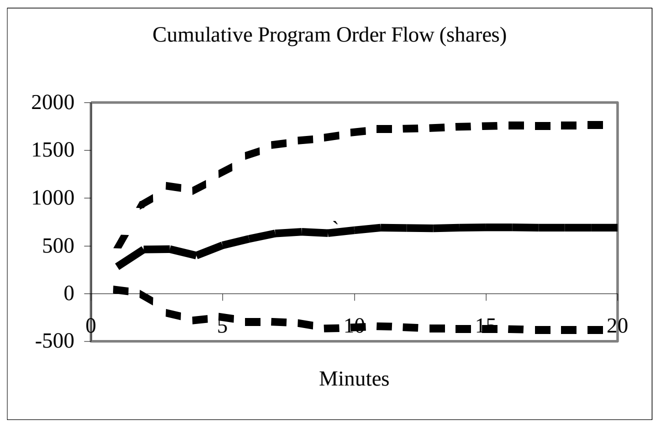

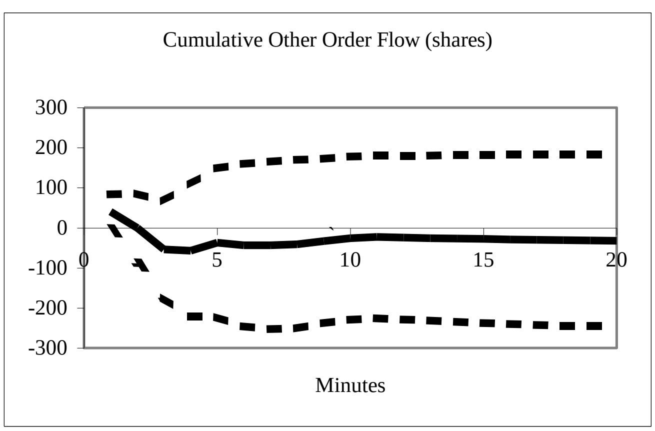

Figure 4b. Unit shock (1,000 shares) to institutional net order flow in XOM

Impulse response functions for a vector autoregression in one-minute quote returns and net order flow of various account types. Confidence intervals are constructed by estimating a separate VAR and impulse response function for each trading day and then assuming independence over time. Dashed lines are two standard errors away from the average estimated impulse response and reflect approximate 95% confidence intervals.

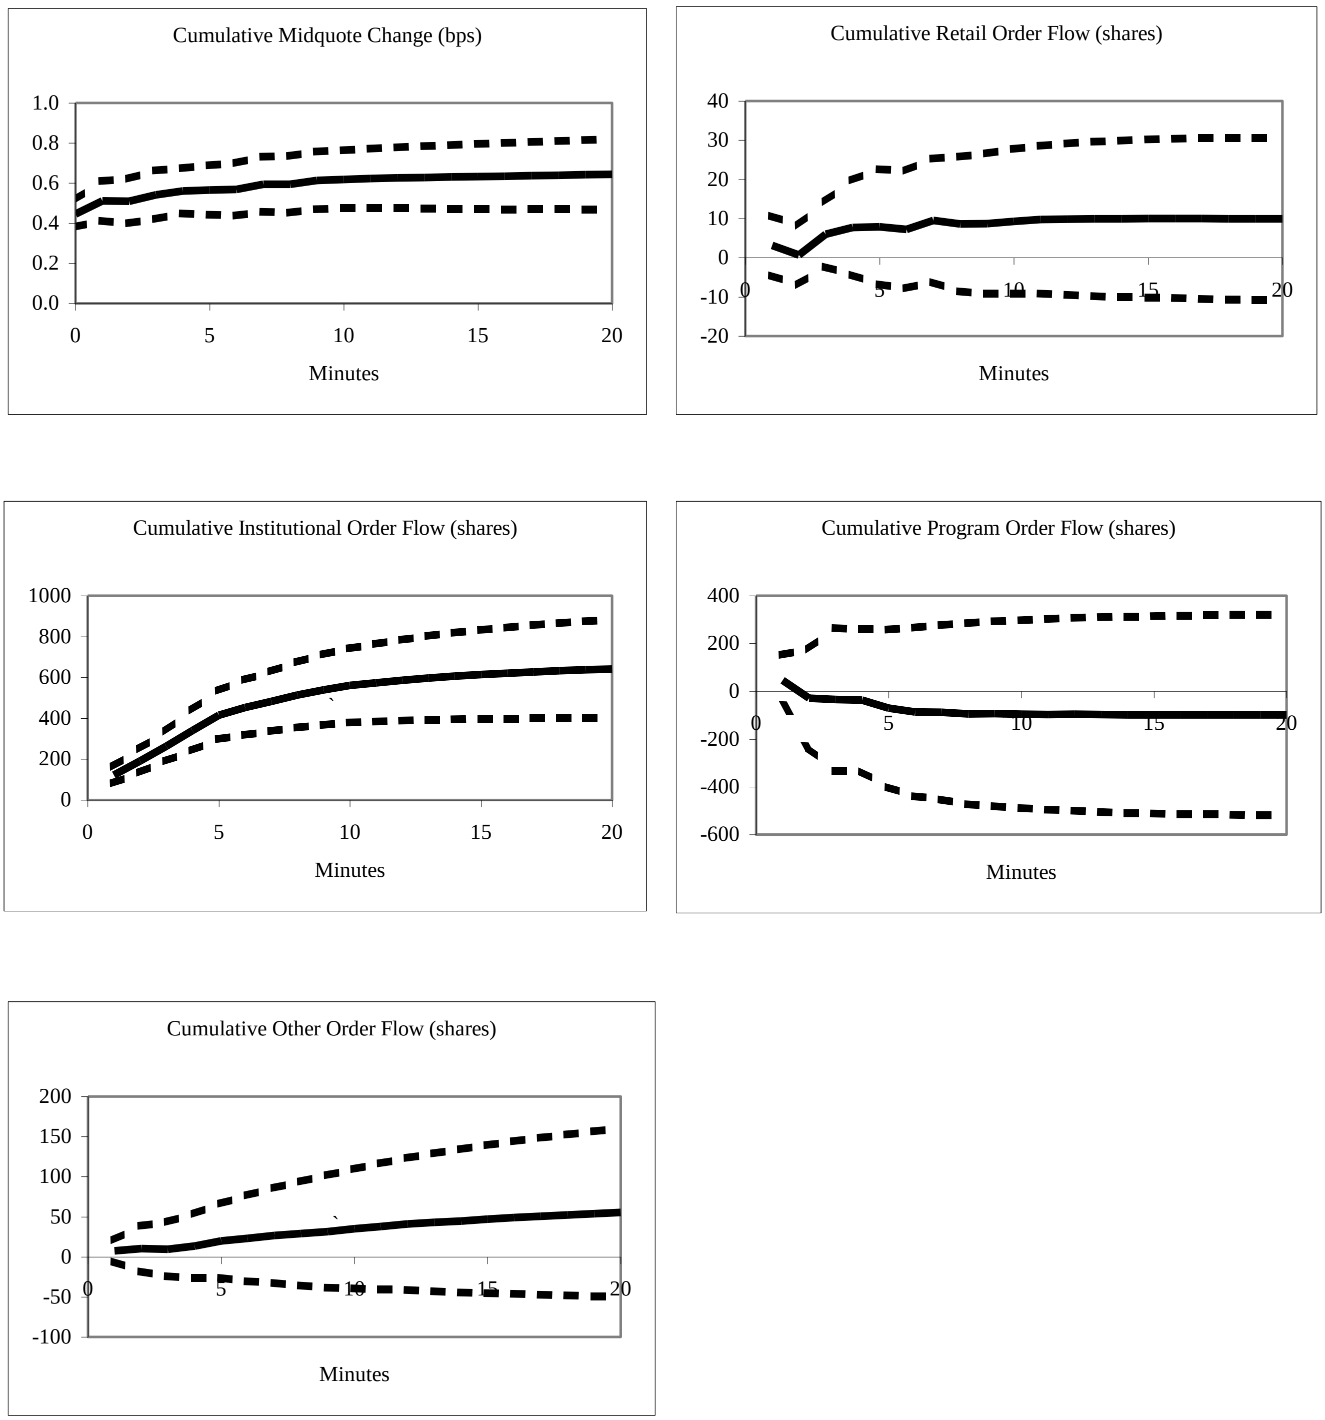

Figure 4c. Unit shock (1,000 shares) to program net order flow in XOM

Impulse response functions for a vector autoregression in one-minute quote returns and net order flow of various account types. Confidence intervals are constructed by estimating a separate VAR and impulse response function for each trading day and then assuming independence over time. Dashed lines are two standard errors away from the average estimated impulse response and reflect approximate 95% confidence intervals.

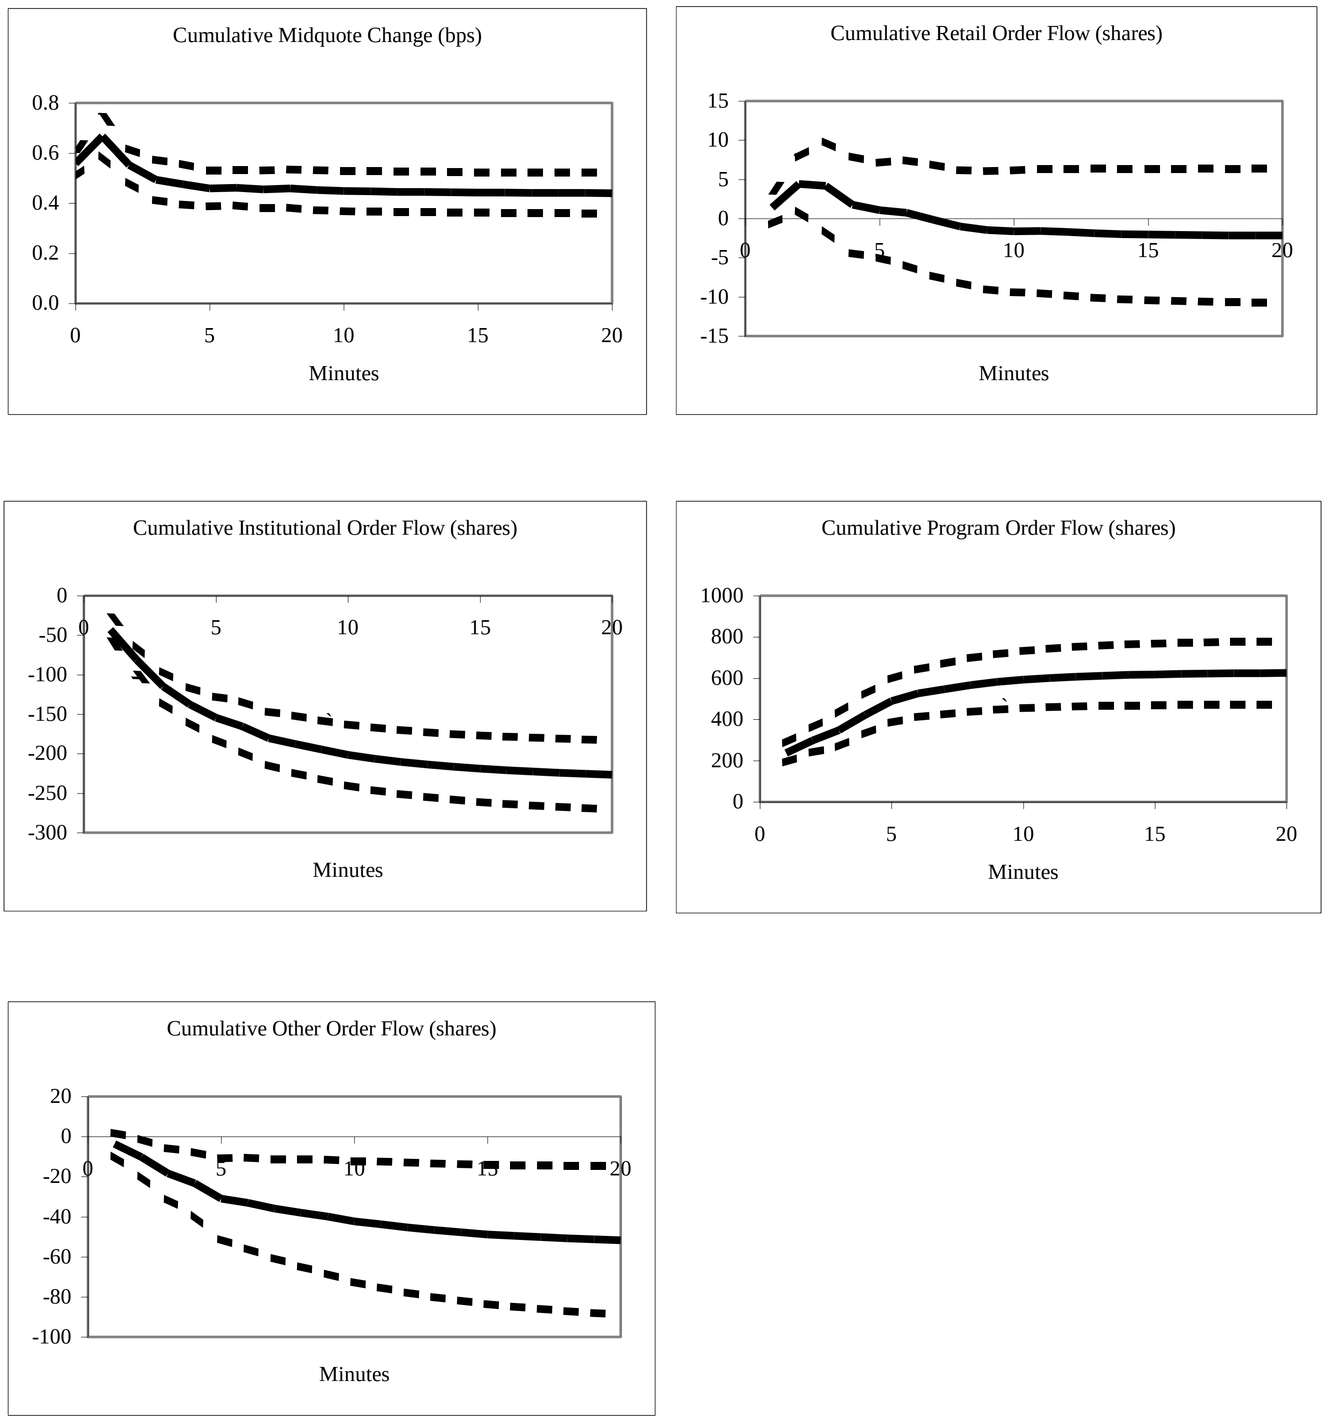

Figure 4d. Unit shock (1,000 shares) to other net order flow in XOM

Impulse response functions for a vector autoregression in one-minute quote returns and net order flow of various account types. Confidence intervals are constructed by estimating a separate VAR and impulse response function for each trading day and then assuming independence over time. Dashed lines are two standard errors away from the average estimated impulse response and reflect approximate 95% confidence intervals.

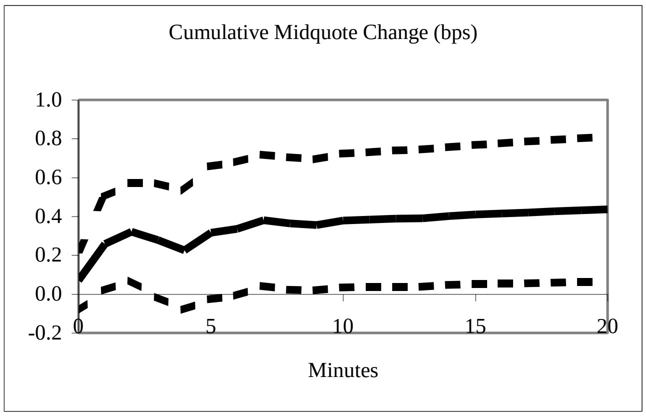

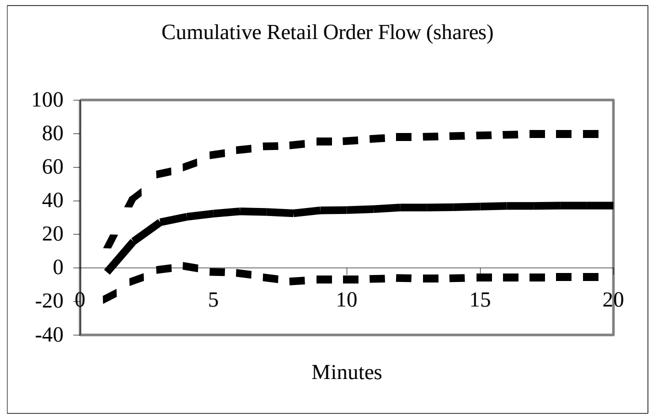

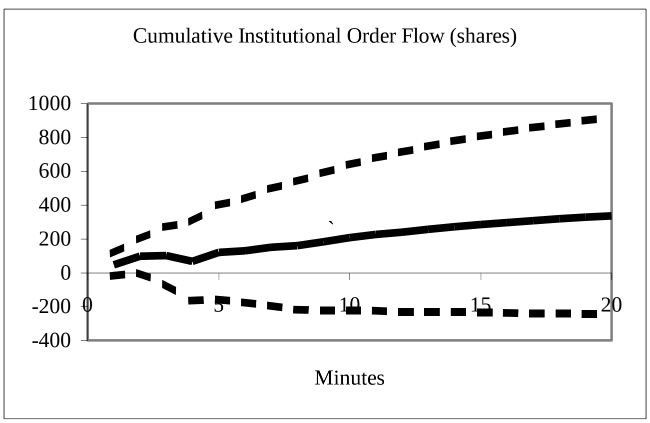

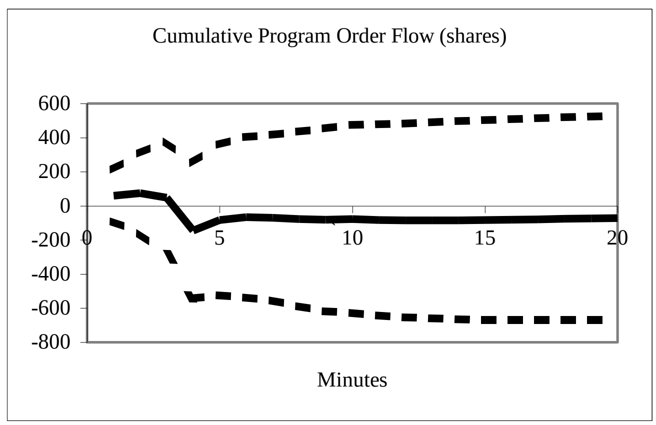

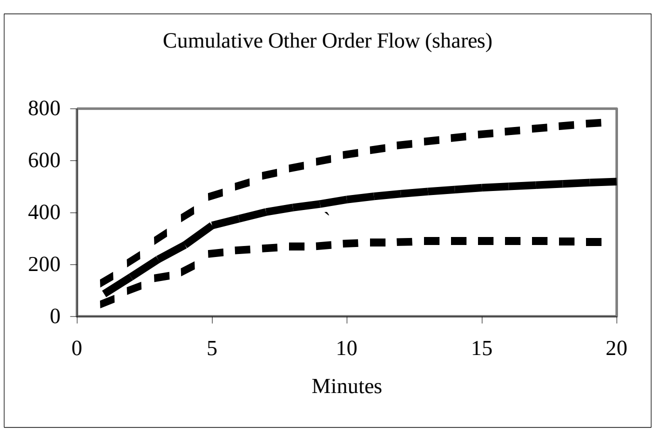

Figure 5a. Unit shock (1,000 shares) to retail net order flow, average of 20 most-active stocks

Results of a vector autoregression in one-minute quote returns and net order flow of various account types. A separate VAR is estimated for each trading day for each stock. The figures report impulse responses averaged across stocks and over time. Confidence intervals assume independence over time. Dashed lines are two standard errors away from the average estimated impulse response and reflect approximate 95% confidence intervals.

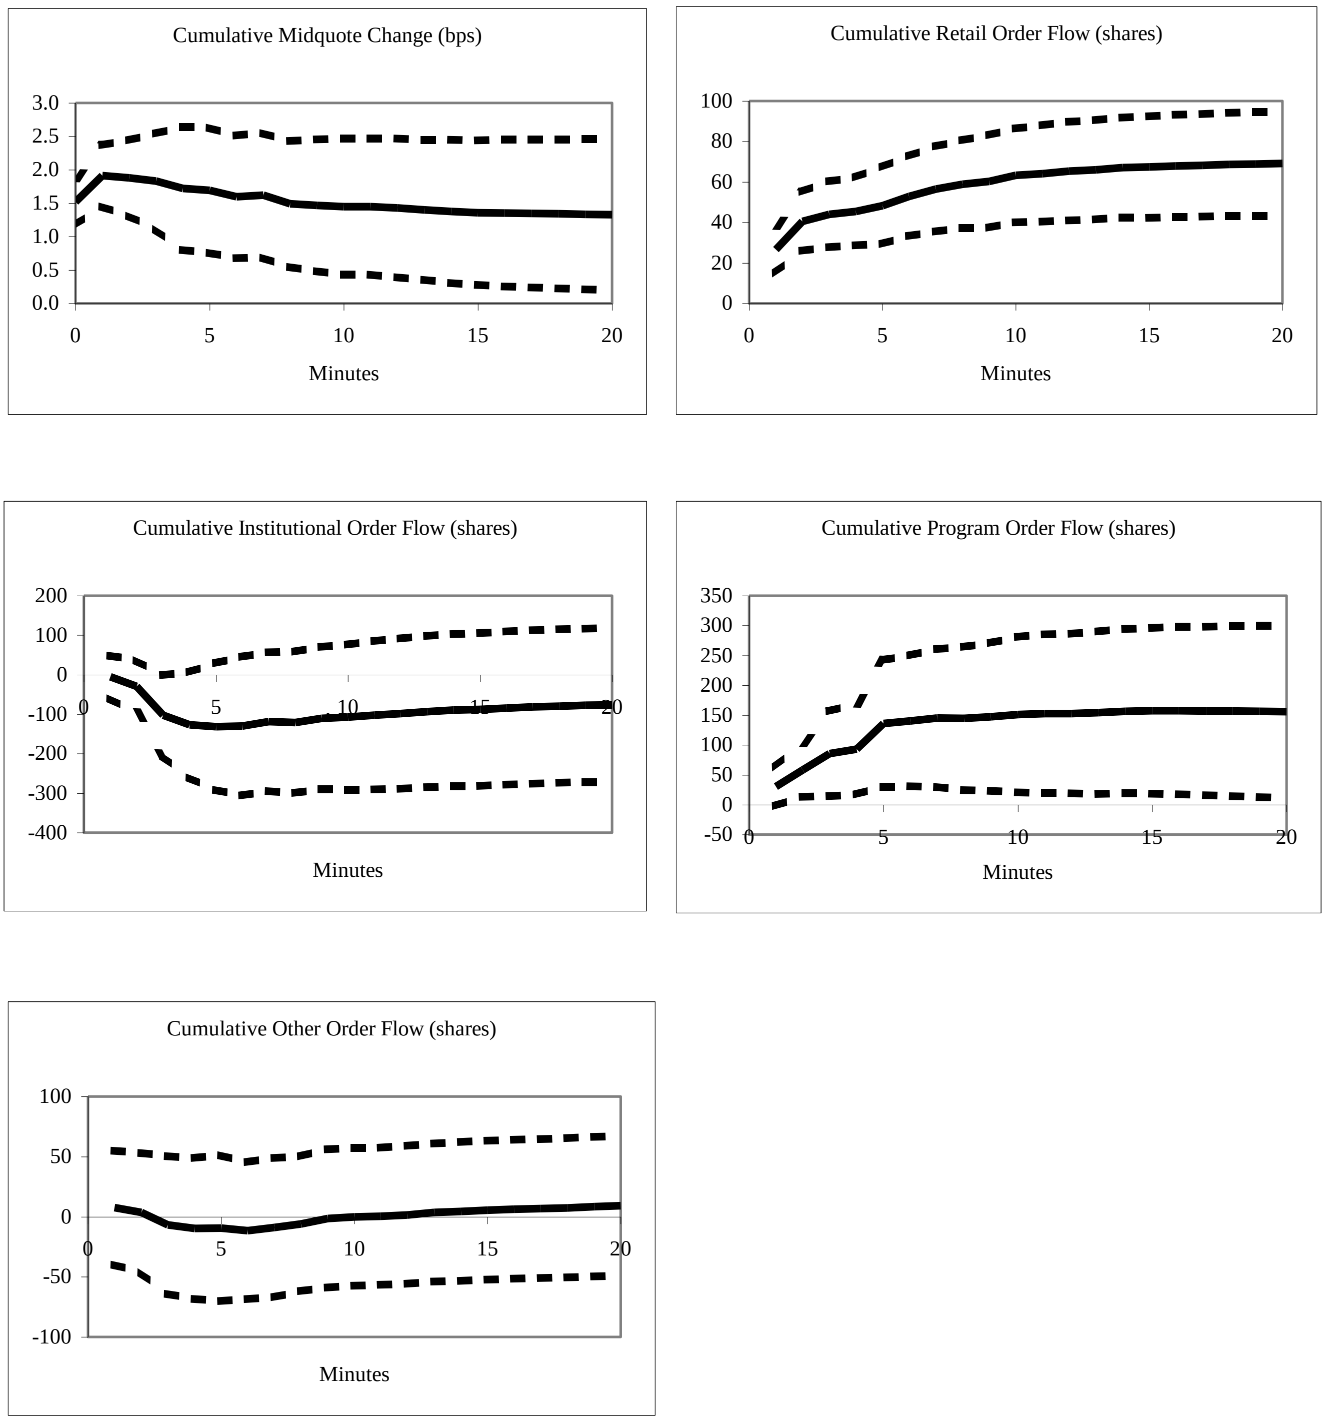

Figure 5b. Unit shock (1,000 shares) to institutional net order flow, avg of 20 most-active stocks

Results of a vector autoregression in one-minute quote returns and net order flow of various account types. A separate VAR is estimated for each trading day for each stock. The figures report impulse responses averaged across stocks and over time. Confidence intervals assume independence over time. Dashed lines are two standard errors away from the average estimated impulse response and reflect approximate 95% confidence intervals.

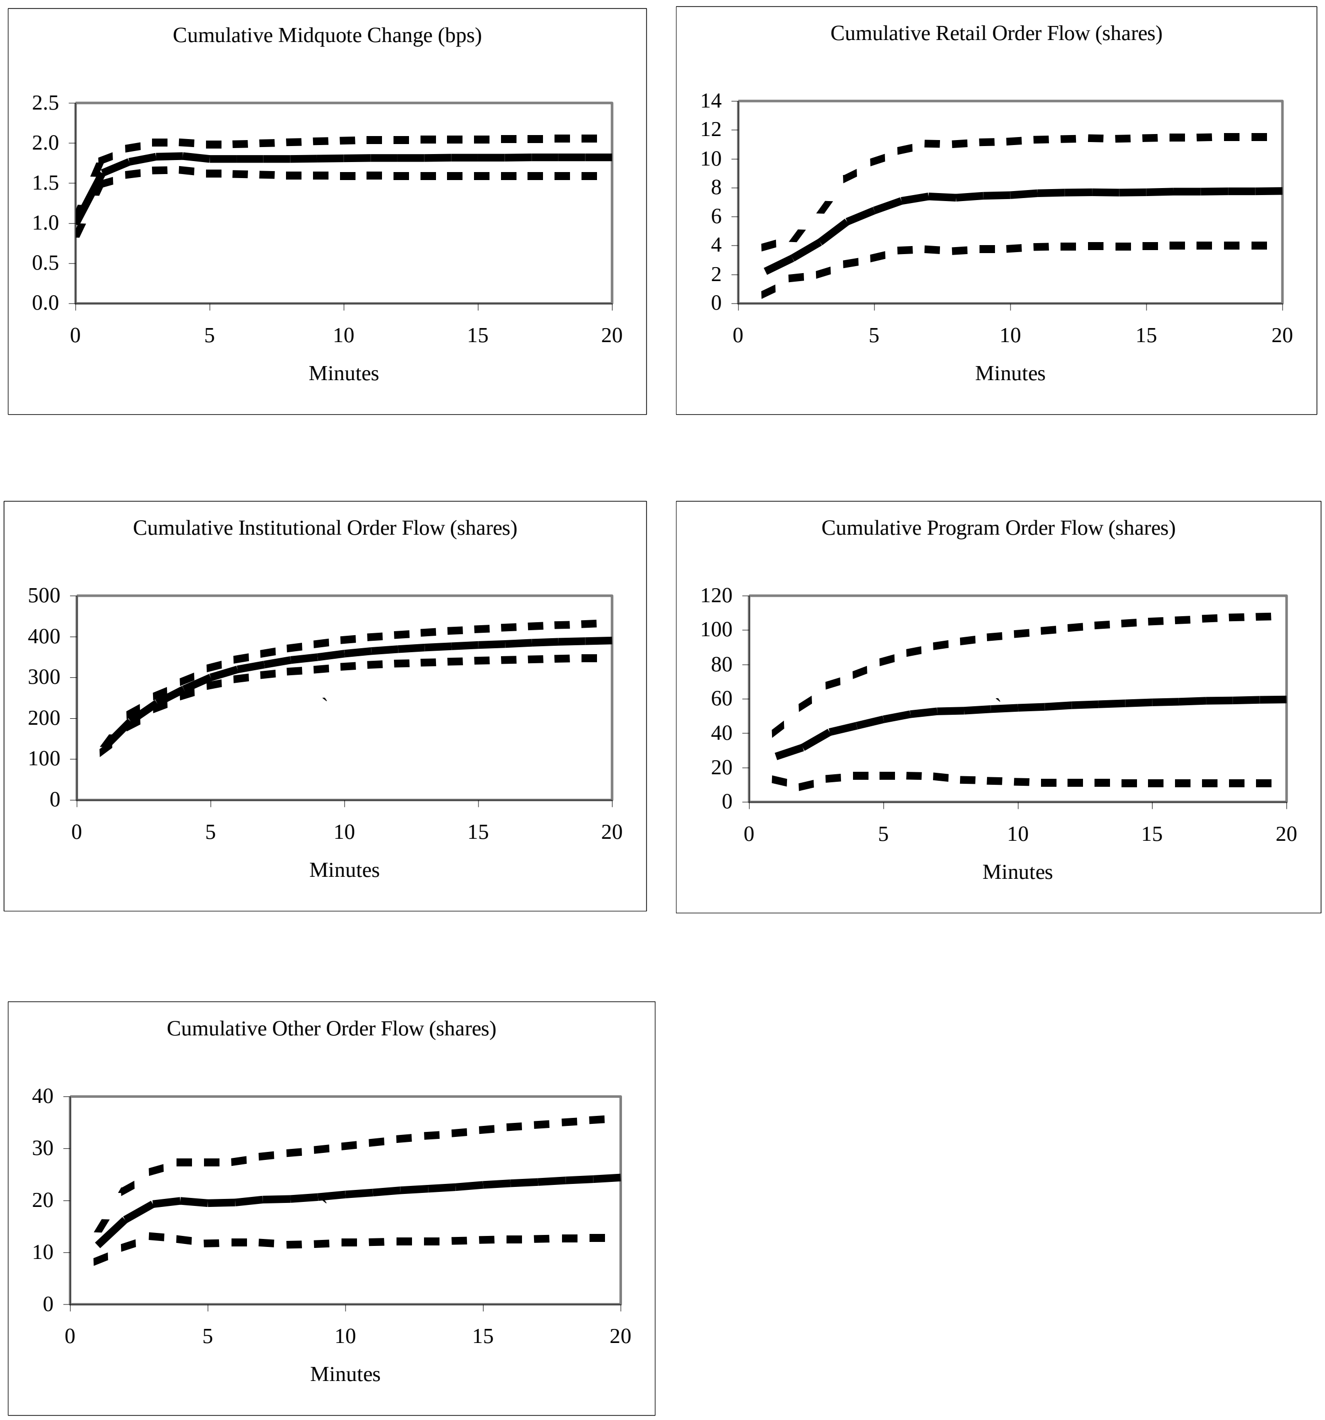

Figure 5c. Unit shock (1,000 shares) to program net order flow, avg. of 20 most-active stocks

Results of a vector autoregression in one-minute quote returns and net order flow of various account types. A separate VAR is estimated for each trading day for each stock. The figures report impulse responses averaged across stocks and over time. Confidence intervals assume independence over time. Dashed lines are two standard errors away from the average estimated impulse response and reflect approximate 95% confidence intervals.

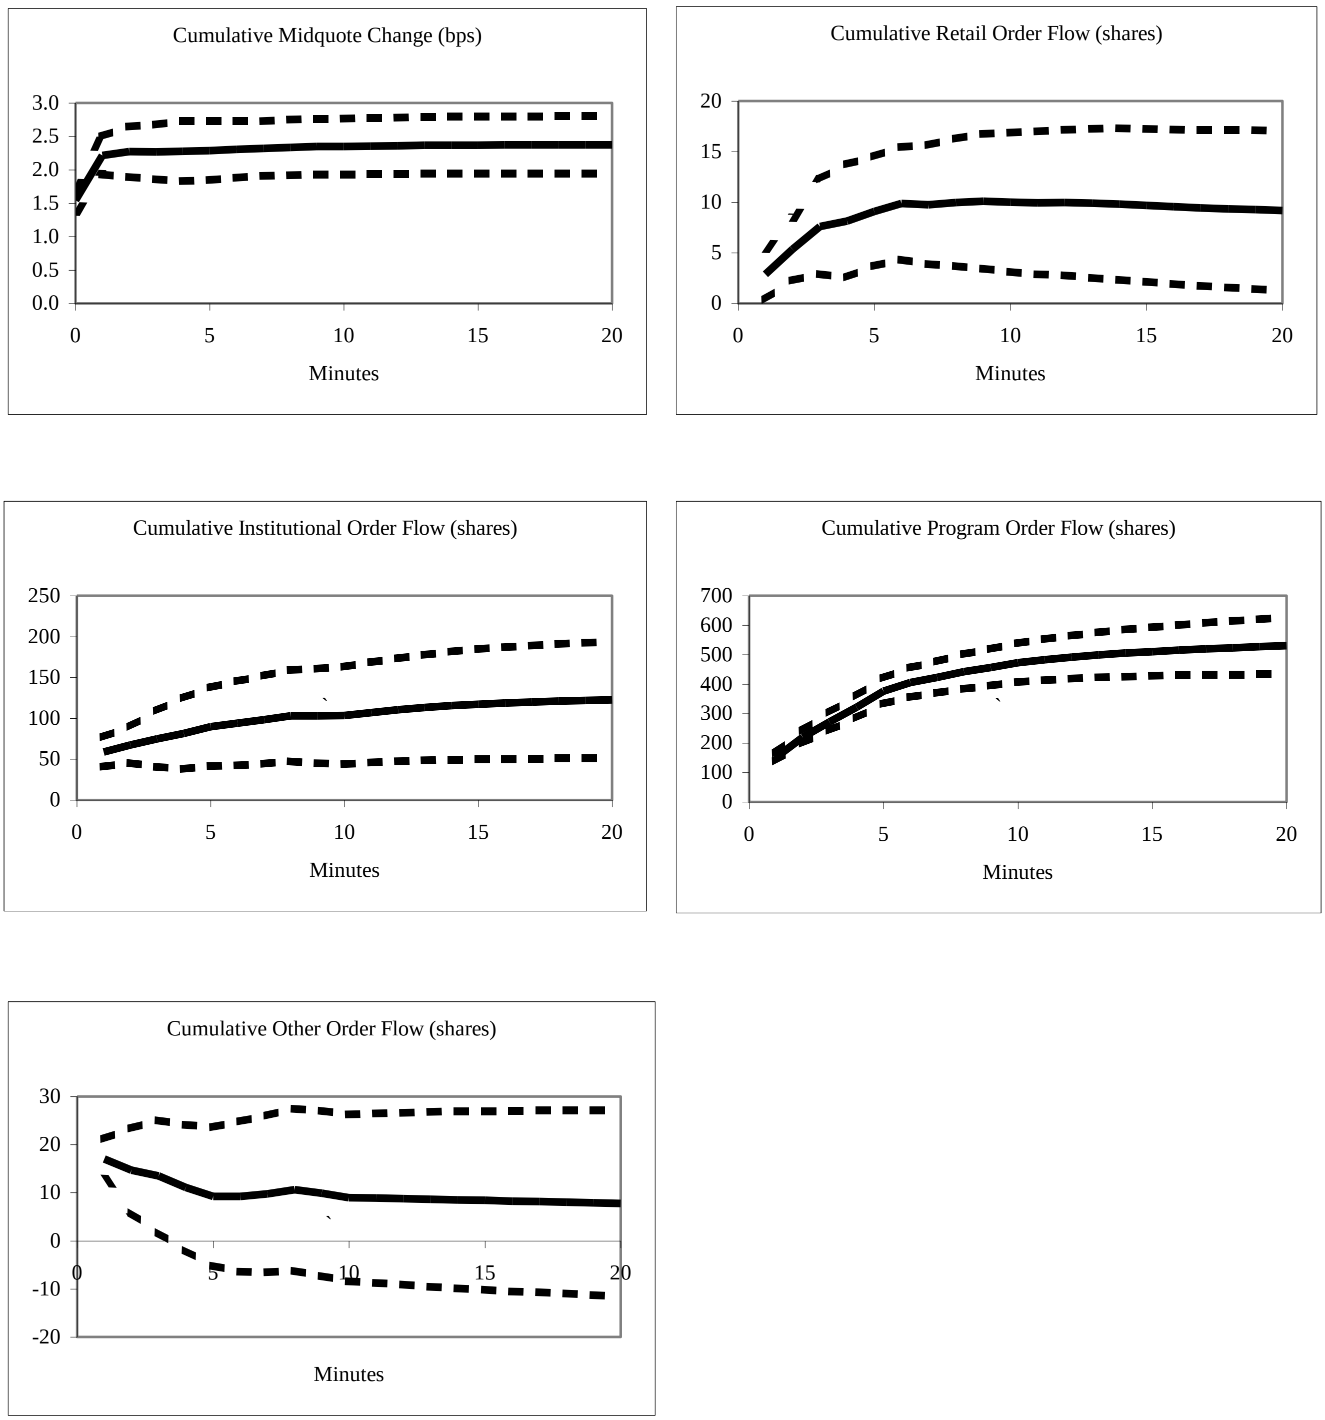

Figure 5d. Unit shock (1,000 shares) to other net order flow, average of 20 most-active stocks

Results of a vector autoregression in one-minute quote returns and net order flow of various account types. A separate VAR is estimated for each trading day for each stock. The figures report impulse responses averaged across stocks and over time. Confidence intervals assume independence over time. Dashed lines are two standard errors away from the average estimated impulse response and reflect approximate 95% confidence intervals.