ELSEVIER

Journal of Financial Markets 6 (2003) 309–335

Journal of FINANCIAL MARKETS

www.elsevier.com/locate/econbase

NYSE order flow, spreads, and information☆

Ingrid M. Werner*

Fisher College of Business, The Ohio State University, 2100 Neil Avenue, Columbus, OH 43210, USA

Abstract

This paper uses unique NYSE audit trail data to evaluate spreads and information content for different order types. Actual spreads are positive for liquidity-demanding orders and negative for liquidity-supplying orders after controlling for order direction. However, because a large fraction of liquidity-demanding orders get price improvement, the actual spread for liquidity-demanding orders is up to 50 percent less than the Lee and Ready (1991) algorithm would suggest. Regression results show that the order composition of trades affects traditional measures of spreads and information. They also show that NYSE non-displayed liquidity reduces trading costs facing market orders, and that liquidity-demanding floor broker orders are the most informative order type.

JEL classification: G20; G21; G23; G24

Keywords: New York Stock Exchange; Execution costs; Spreads; Information; Order types

1. Introduction

Execution costs is one of the most intensively researched areas in market microstructure. The focus of recent studies ranges from: comparison of spreads for

☆I thank the NYSE for giving me the opportunity to gain insight into the nature of floor trading during my time as Visiting Research Economist at the NYSE, and for giving me access to the data necessary for this project. I especially thank George Sofianos for his encouragement to pursue this project. I am also grateful for financial support from the Dice Center for Financial Research and for comments and suggestions from Ekkehart Boehmer, Robert Engle, participants at the CAP conference in Denmark, and from seminar participants at Rice University and the New York Stock Exchange. The comments and opinions expressed in this paper are the author's and do not necessarily reflect those of the directors, members or officers of the New York Stock Exchange, Inc. All remaining errors are my own.

*Tel.: +1-614-292-6460; fax: +1-614-292-2418.

E-mail address: werner@cob.osu.edu (I.M. Werner).

PII: S1386-4181(02)00066-6

trades of different size (e.g., Kraus and Stoll, 1972; Holthausen et al., 1987,1990; Jones et al., 1994; Chan and Fong, 2000); comparison of spreads for system market and limit orders (Harris and Hasbrouck, 1996; Peterson and Sirri, 2002); comparison of spreads for upstairs versus downstairs trades (e.g., Keim and Madhavan, 1996; Madhavan and Cheng, 1997); and cross-exchange comparisons of spreads (e.g., Huang and Stoll, 1996; Bessembinder and Kaufman, 1997; Blume and Goldstein, 1997). However, no study has been able to analyze both order executions for New York Stock Exchange (NYSE) system (SuperDOT) and executions based on orders sent to NYSE floor brokers in the same sample.[1] In fact, with the exception of Handa et al.'s (1998b) study of the Amex, there is no empirical work that analyzes spread costs and information content of downstairs floor broker orders. Yet, these orders correspond to over forty percent of NYSE dollar volume.

This paper examines how the composition of NYSE order flow affects traditional measures of spreads and information content. It answers questions like: What spread costs do traders submitting marketable orders through SuperDOT actually pay? How expensive are orders traded actively by NYSE floor brokers? In what way are these costs different from those facing orders that floor brokers leave with the NYSE specialist for execution? How large are the spread gains earned by traders submitting limit orders through SuperDOT? And, how much compensation does the specialist actually get for providing liquidity? The paper also examines theoretical hypotheses regarding what order type informed investors favor.[2] Finally, the paper illustrates how important the order flow composition is for traditional measures of spreads and information content.

The study is based on a broad cross section of NYSE-listed stocks during the periods July 21–August 22, and September 2–October 3, 1997. The data comes from the NYSE audit trail and is different from TAQ and TORQ data in important ways. Specifically, the data identifies the entire order composition on both the buy and the sell side of each trade. Hence, it covers all types of orders: orders that are traded actively by NYSE floor brokers in the downstairs market; orders that are upstairs facilitated; not-held orders that are left by floor brokers with the NYSE specialist (percentage orders); NYSE specialist orders; system market and limit orders that reach the NYSE trading floor via SuperDot; and orders from third markets that reach the floor via the Intermarket Trading System (ITS). It is therefore possible to calculate measures of execution costs such as effective half-spreads, and realized-half spreads for each order type.[3]

[1] Harris and Hasbrouck (1996) compare system market and limit orders, Keim and Madhavan (1996) and Madhavan and Cheng (1997) compare upstairs facilitated and downstairs block trades, Peterson and Sirri (2002) compare system market orders and marketable limit orders, and Bacidore et al. (2002) study system market orders.

[2] Cooney and Sias (2001) study information content of different order types using TORQ data from 1990 to 1991 and conclude that floor orders are most informative.

[3] Peterson and Sirri (2002) study retail system market orders, marketable limit orders, and regular limit orders for NYSE-listed stocks traded on the NYSE and the regional exchanges. Bacidore et al. (2002) study NYSE system market orders. SEC (2001) studies retail market orders, marketable limit orders, and limit orders from the NYSE and Nasdaq. All studies explicitly exclude floor broker orders, percentage orders (not-held orders), and specialist orders from their analyses.

This paper also examines the information content of different order types. An informed trader has a choice of what order type (or types) to submit. Recent models of the NYSE market microstructure assume that impatient informed traders submit market orders, while patient value traders submit limit orders (e.g., Seppi, 1997). An impatient informed trader is assumed to be willing to pay the liquidity premium inherent in submitting a market order or a marketable limit, i.e., buy at the offer or sell at the bid, and in the case of a market order possibly even walk the book. Both order types are thought of as liquidity-demanding orders in the literature, and hence are expected to pay the spread. This spread serves as compensation to liquidity-suppliers, i.e. limit orders and the NYSE specialist, for the adverse selection risk. Since prices are assumed to adjust to information, the models predict that prices will move in the direction of market orders and marketable limits, but in the opposite direction to limit orders and specialist orders.

What about the spreads and information content of orders represented by floor brokers? Floor brokers tend to be modeled as uninformed and as competing with the specialist and limit orders over the liquidity premium paid by market order submitters (e.g., Seppi, 1997; Werner, 1999). If this were the case, floor broker orders would look like other liquidity-suppliers in the data. Hence, the maintained hypothesis is that floor broker orders gain the spread and prices move in the opposite direction to these orders on average. However, this ignores the value added provided by floor brokers as described in Sofianos and Werner (2000). Floor brokers' ability to quickly adjust the trading strategies to current market condition is especially valuable when trading the large and difficult (price-sensitive) orders that are typically entrusted floor brokers. Hence, the alternative hypothesis is that floor broker orders often are information-based. This suggests that these orders pay the spread and that prices tend to move in the direction of the orders on average.

The order-based analysis shows that effective half-spreads for marketable limits and ITS orders are more than twice as high as those for market orders. The difference can be explained by the fact that 45 percent of all market orders get price improvement, while marketable limits by definition do not get any price improvements. Market orders get price improvement in three ways: the specialist stops the order (Ready, 1999); the specialist offers price improvement (Edwards and Harris, 2001); or the order is exposed to non-displayed liquidity through the auction on the NYSE floor. ITS orders do not benefit from the auction process on the NYSE floor, which in part may explain why they pay such high spreads. Interestingly, the effective half-spread for market orders is only slightly higher than for floor broker orders on average despite the fact that floor brokers arguably have an advantage by being on the NYSE floor as explained in Sofianos and Werner (2000). The reasons for this are that: floor brokers often have to offer price improvement to compete with displayed liquidity; executions often occur between the quotes when there is both buying and selling interest on the floor; and floor brokers trading price-sensitive orders often pay the spread. As predicted by theory, specialist orders and limit orders instead have negative effective half-spreads on average. The spread gains for limit orders are considerably larger than those for the specialist. There are two reasons for this. First, the specialist may offer price improvement, and second, the specialist may be willing to pay the spread to manage his inventory or to pursue his own proprietary trading strategy. Percentage orders also enjoy spread gains on average, but these gains are small.

The extent to which share-weighted average mid-quotes move in the direction of the order in the five minute window following the trade is a commonly used measure of information content. The order-based analysis shows that marketable limit orders are the most informative orders, followed by ITS orders, and floor broker orders. In contrast to predictions from theory, market orders have relatively limited information content. Since orders that on average act as liquidity-suppliers face adverse selection risk, we expect mid-quotes to move in the opposite direction of these orders. Indeed, executions based on percentage orders, specialist orders, and limit orders are followed by adverse price moves on average. In other words, these orders get picked off.

The results show that realized half-spreads are generally much smaller in magnitude than effective half-spreads, suggesting that the effective half-spreads tend to compensate for the information content of different order types. However, the offset is not perfect. Realized half-spreads are significantly positive for ITS and market orders. Interestingly, realized half-spreads are also significantly positive for limit orders on average suggesting that the spread gain they earn does not fully compensate for short-term adverse selection risk. By contrast, floor brokers and specialist have significantly negative realized half-spreads. In the case of floor broker orders, this is because the information content of their orders is large compared to the effective half-spread they face. In the case of the specialist, this is because they face relatively limited adverse selection risk compared to the effective half-spreads that they gain on average.

Most researchers do not have access to order-level data from U.S. exchanges. It is therefore important to illustrate how traditional measures of spreads and information content of trades based on trade classification algorithms are affected by order flow composition. Pooled cross-section, time-series regressions are used to test for the effect of order flow on spreads after controlling for stock characteristics, trading day characteristics, and trade size. The results show that order flow composition significantly affects all standard measures of spreads and that ITS orders and market orders pay the highest spreads. Moreover, they show that market orders are significantly better off because they are able to participate in the auction on the NYSE floor. Finally, the results show that floor broker orders by far are the most informative order type, followed by marketable limit orders. By comparison market orders in isolation have virtually no information content. As predicted, prices move in the opposite direction following regular limit orders and specialist orders.

If traders using a particular order type are systematically better informed, they would on average buy at a lower price and sell at a higher price compared to traders favoring other order types. Therefore, how a particular order type fares compared to the volume-weighted average price (VWAP) during the day, or how high their normalized trading profits are on average, provides a complementary view of what order types on average pay for liquidity.[4] Note that this measure uses a more long-term (daily) benchmark price. The results show that limit orders and specialist orders significantly beat the VWAP. So while short-term adverse selection costs wipe out the spread gains, over the course of the day liquidity suppliers gain enough in spreads to compensate for adverse selection. Consistent with the results mentioned already, market orders, marketable limits, and ITS orders significantly underperform the VWAP. By contrast, both percentage orders and floor broker orders execute on average close to the VWAP.

Finally, the data permit an evaluation of the widely used Lee and Ready (1991) algorithm for NYSE data, complementing recent work by, e.g., Ellis et al. (2000), and Odders-White (2000).[5] The misclassification of what we traditionally think of as liquidity-demanding orders is even larger than what previous studies have found. For example, almost 30 percent of all market orders are misclassified using the Lee and Ready algorithm. In other words, market buy (sell) orders frequently execute at or below (above) the mid-quote. The fraction of orders that is misclassified is larger for many other order types. As a result, the Lee and Ready algorithm drastically overestimates the spreads and information content for order types that are usually thought of as demanding liquidity.

The paper proceeds as follows. Hypotheses regarding execution quality measures for different order types are outlined in Section 2. The data are described in Section 3. In Section 4, order executions are compared to the NYSE specialist's quotes. Spread measures are calculated and discussed in Section 5. The effect of order flow on traditional measures of execution costs is discussed in Section 6. The analysis of volume-weighted average prices is in Section 7. Section 8 concludes.

2. The selection of order type, spreads, and information content

The NYSE has a complex microstructure, with many different options for a trader seeking to fill an order. The trader may use SuperDOT to submit a market order; a marketable limit order; or a limit order. Alternatively, the trader may engage a NYSE floor broker to represent the order in the crowd, either actively or passively. Finally, a trader may route an order from a regional exchange or the Third Market to the NYSE via ITS. Of course, the trader can choose any combination of these order types. So, what order type would a particular trader choose? The order selection should mainly be dictated by the trader's demand for immediacy, or his patience. A trader's patience in turn is related to his liquidity needs and the extent to which the trading decision is based on short-lived private information. In other words, a trader with either a liquidity shock and/or short-lived private information is likely to be impatient, and is expected to be willing to pay a premium for immediacy.

[4] Chan and Lakonishok (1993) use this measure in their study of U.S. institutional trades. Choe et al. (2000) use this approach in a recent study of whether or not foreign investors have better information than local investors in Korea.

[5] Similar observations are made by Peterson and Sirri (2002).

By contrast, a patient value trader might be willing to risk waiting for his order to get filled, and possibly risk not getting it filled at all, in return for a more attractive price.

Building on this intuition, recent models of the NYSE market microstructure assume that impatient informed traders submit market orders and patient value traders submit limit orders (e.g., Seppi, 1997). An impatient informed trader is assumed to be willing to pay the liquidity premium inherent in submitting a market order, i.e., buy at the offer or sell at the bid, and possibly even walk the book. A more conservative informed trader may choose to submit a marketable limit order, i.e., a buy (sell) order at the offer or above (bid or below), to eliminate the risk of walking the book. Both order types are thought of as liquidity-demanding orders in the literature, and hence are expected to pay the spread. This spread serves as compensation to liquidity-suppliers, i.e. limit orders and the NYSE specialist, for the adverse selection risk. Since prices are assumed to adjust to information, the models predict that prices will move in the direction of market orders and marketable limits, but in the opposite direction to limit orders and specialist orders.

What about the spreads and information content of orders represented by floor brokers? Sofianos and Werner (2000) show that these orders represent as much as 44 percent of dollar volume on the NYSE so they are clearly an integral part of the trading process. Floor brokers tend to be modeled as uninformed and as competing with the specialist and limit orders for the liquidity premium paid by market order submitters (e.g., Seppi, 1997; Werner, 1999). If this were the case, floor broker orders would look like other liquidity-suppliers in the data. Hence, floor broker orders would tend to gain the spread and prices would move in the opposite direction to these orders on average. This is certainly likely to be the case for percentage orders that floor brokers leave with the NYSE specialist for execution. However, the models ignore the value added provided by floor brokers. Sofianos and Werner (2000) argue that floor brokers are paid a higher commission for their services because they can quickly adjust the trading strategies to current market condition. This is especially valuable when trading large and difficult (price-sensitive) orders as floor brokers do (see Sofianos and Werner, 2000). Hence, the alternative hypothesis is that floor broker orders, particularly those that are traded actively on the NYSE floor, are information-based. Hence, these orders should pay the spread and prices should tend to move in the direction of the orders on average.

3. Data description

The NYSE audit trail data used in this study is very similar to the data described in Sofianos and Werner (2000). They first grouped all NYSE-listed stocks into deciles based on December, 1996, trading volume and then sampled 20 stocks randomly from the most liquid stocks (decile 10), and 10 stocks randomly from each of the other nine deciles. The sample is the 101 stocks that were still traded during the periods: July 21–August 22, and September 2–October 3, 1997.[6] The average market capitalization for stocks in the sample is $3.1 billion, the average stock price is $31.22, the average daily share ($) volume is 275 thousand (12 million), and the average trade size is 1,926 (57,002) shares (dollars). The sample by design has a large cross-sectional dispersion.

Briefly, in addition to the usual information on the execution time stamp (in seconds), size, and the price of trades, the data includes information that permits classifying the underlying order types for all components of the buy *and* the sell side. For example, one trade of 12,000 shares at price $89 13/16 that took place on August 22, 1997, at 9:52:26 may have 10,000 shares executed actively by a floor broker and a 2,000-share market order on the buy side, and an 8,000-share limit order and 4,000 shares executed by the specialist on the sell side.

The focus of this paper is on order executions during regular trading, and therefore opening trades executed through OARS as well as the after hours crossing sessions are excluded. If there are several executions based on orders of a particular type on the same side in a trade, these are aggregated. For example, if there are two limit orders of 500 shares on the buy side and one 1,000-share market order on the sell side, they are aggregated into one 1,000-share limit order on the buy and one 1,000-share market order on the sell side. Finally, 260 trade executions that are more than $0.50 outside the quotes in effect 5 seconds prior to the trade are excluded. The final sample has 613,644 trade executions.

Table 1 provides an overview of the trade data by order execution type. The data are separated into trades executed downstairs, Panel A, and trades that are facilitated in the upstairs market, Panel B. The algorithm developed by Madhavan and Cheng (1997) is used to identify upstairs-facilitated trades. In a nutshell, an upstairs-facilitated trade exceeds 9,999 shares and has the same floor broker represented on both the buy and the sell side. The first column is based on the number of trade reports, counting both buys and sells. System limit orders are the most commonly occurring orders represented in 29.7 percent of all trades, followed by system market orders, 26.0 percent, and specialist orders, 14.8 percent. However, this is largely because executions based on system and specialist orders are smaller than trades based on the orders that are actively traded by NYSE floor brokers and the orders that left by floor brokers with the specialist (percentage orders).[7] The largest fraction of downstairs share volume therefore comes from active floor broker orders −25.4 percent. System limit orders is a close second at 24.4 percent followed by system market orders and percentage orders at 15.3 and 14.5 percent respectively. Specialist orders represent 8.0 percent and ITS orders 1.4 percent of share volume. Note that it is not possible to classify 1.1 percent of twice share volume (not reported).[8]

[6] The attrition was due to mergers and/or de-listings of the following stocks (decile): ACO (1), ALW (7), ATP (6), DFI (5), DNA (6), MFO (2), PDP (8), RTZ (1), and V (8).

[7] NYSE rules dictate that percentage orders (also called not-held orders or CAP orders) left by floor brokers with the specialist be non-discretionary to reduce conflict of interest for NYSE specialists.

[8] The value-weighted distribution of volume (not reported) is very similar to the share-weighted one.

Table 1
Distribution of order execution types[a]

	Fraction of total buys and sells		Execution size Average (shares)
	Trade-weighted	Share-weighted	
A. Downstairs trades			
Market orders	26.0	15.3	1,016
Marketable limit orders	11.2	10.1	1,661
Limit orders	29.7	24.4	1,379
ITS orders	2.6	1.4	961
Floor broker orders (active)	7.2	25.4	5,934
Percentage orders (passive)	7.9	14.5	2,857
Specialist orders	14.8	8.0	828
B. Upstairs facilitated trades			
Market orders	1.1	0.8	5,522
Marketable limit orders	0.0	0.0	2,286
Limit orders	2.6	2.2	9,119
ITS orders	0.0	0.0	2,225
Floor broker orders (active)	92.8	94.0	52,856
Percentage orders (passive)	3.2	2.7	12,761
Specialist orders	0.0	0.0	1,849

[a] Panel A describes the breakdown of trading volume and the execution size distribution for different types of orders for 611,383 downstairs trades in the sample. A marketable limit order to buy (sell) is a limit order that executes above the bid (below the ask) quote posted by the specialist 5 seconds before the trade. Percentage orders are orders that floor brokers leave with the specialist for execution as not-held orders. The columns do not sum to zero due to a small fraction of orders for which the audit trail does not include information on the order type. Panel B describes the breakdown of trading volume and the execution size distribution for 2,261 upstairs-facilitated trades. Upstairs-facilitated trades are defined as in Madhavan and Cheng (1997).

Panel B of Table 1 characterizes the distribution of order executions for upstairs-facilitated trades. As expected, floor brokers represent the vast majority of the volume in these trades, ranging from 92.8 percent for trade-weighted averages to 94.0 percent for share-weighted averages. However, note that percentage orders, limit orders, and market orders, do occasionally participate as the block is broken up when brought to the floor for execution.[9] It is also not surprising that the average execution size for active floor broker orders in Panel B is considerably larger than in Panel A.

The distribution of trading volume is very similar to the one reported for an earlier sample in Sofianos and Werner (2000). The main difference is that the specialist share of volume is smaller while the share of percentage trades is larger in this sample. The reason for these differences is that the more detailed data permit a better

[9] NYSE Rule 76 describes the rules for "Crossing" Orders inside the spread. See also NYSE Rule 127 for crossing blocks outside the prevailing quotes. For a discussion of these and other NYSE rules, see Hasbrouck et al. (1993).

classification of orders executed by the specialist into proprietary trades and those trades that the specialist execute on behalf of floor brokers. It is also possible that the tick size reduction from 1/8 to 1/16 that took place in the month before this sample was taken has affected the distribution of trading volume. Werner (1999) discussed the implications of the tick size reduction for the participation rates of specialists and floor brokers.

Further analysis of the trade composition data shows that 25 percent of trades (57 percent of share volume) do not include any marketable system orders, and 4 percent of trades (6 percent of share volume) include marketable system orders on both sides of the trade. In either of these cases, it is a stretch to designate the buy or the sell side as the liquidity demander simply based on whether or not the trade executes above the mid-quote as is commonly done in the literature (e.g., Lee and Ready, 1991).

Finally, note that the data does not include information on the specific order instructions given to NYSE floor brokers. These instructions typically include a price limit as well as a time frame, and the executions should at least in a volume-weighted average sense respect this price limit. Percentage orders that are left by NYSE floor brokers with the specialist for execution all have explicit price limits.[10] Hence, when interpreting spreads and information content it is important to keep in mind that both active floor broker trades and trades executed by the specialist on behalf of floor brokers are priced and that the execution price often entails a price improvement compared to the original order instructions.

4. Order executions and contemporaneous quotes

To understand the complex nature of NYSE trading and how it affects measures of spreads and information content, it is useful to first characterize how different types of orders execute relative to the contemporaneous NYSE quotes. Recall that liquidity-demanding orders like market orders are expected to execute at the opposite side of the market (e.g., buy at the ask), while liquidity-supplying orders like limit orders and specialist orders are expected to execute at the same side of the market (e.g., buys at the bid). Quotes from TAQ are matched to the NYSE audit trail using a 5 seconds delay as originally proposed by Lee and Ready (1991).[11]

Tables 2 and 3 tabulate the frequency of executions at the quotes on the same side of the market or better (bid for buys and ask for sells); executions inside the quotes, further divided into those that are at the mid-quote and those that are not; and executions at the quotes on the other side of the market or worse (ask for buys and

[10] NYSE Rule 13 and 123A.30 defines a percentage order as: "A limited price order to buy (or sell) 50% of the volume of a specified stock after its entry". During this sample, there were three types of percentage orders: Straight Limit Percentage Orders; Last Sale Percentage Orders; and "Buy Minus"–"Sell Plus" Percentage Orders. Effective March 1, 1999, an additional type was added: "Immediate Execution or Cancel Election".

[11] The results in this paper are qualitatively the same if we use 5, 15, 30, or 60 seconds delay of the quotes. The only major difference is that more executions occur outside the posted quotes the longer the delay.

Table 2
Trade executions for system orders by execution size[a]

	Execution size	Number	Distribution of price relative to quotes						Quoted half-spread (basis points)
			Same side or better	Inside better	At mid-quote	Inside worse	Other side or worse	DI	
Market orders	**All**	**352,703**	**20.3**	**3.6**	**13.4**	**7.9**	**54.8**	**35.4**	**18.7**
	$x<500$	175,726	13.2	4.3	18.2	10.4	53.9	1.5	16.7
	$500\leqslant x<2,000$	129,716	16.3	3.4	14.3	9.2	56.7	14.1	17.6
	$2,000\leqslant x<10,000$	44,058	19.5	3.5	12.6	7.4	57.0	39.0	19.4
	$x\geqslant 10,000$	3,203	27.9	4.4	13.2	6.6	47.8	70.2	21.8
Marketable limits	**All**	**141,957**	**0.0**	**0.0**	**0.0**	**0.0**	**100.0**	**30.8**	**13.6**
	$x<500$	36,554	*0.0*	*0.0*	0.0	0.0	100.0	4.0	12.0
	$500\leqslant x<2,000$	68,855	*0.0*	*0.0*	0.0	0.0	100.0	14.9	12.7
	$2,000\leqslant x<10,000$	33,345	*0.0*	*0.0*	0.0	0.0	100.0	32.1	13.9
	$x\geqslant 10,000$	3,199	*0.0*	*0.0*	0.0	0.0	100.0	44.7	16.1
Limit orders	**All**	**414,220**	**79.9**	**3.8**	**12.4**	**3.9**	**0.0**	**39.9**	**17.0**
	$x<500$	167,635	74.3	5.6	15.1	4.9	*0.0*	1.5	15.4
	$500\leqslant x<2,000$	168,948	72.3	5.4	16.8	5.5	*0.0*	12.9	16.1
	$2,000\leqslant x<10,000$	70,055	79.5	3.6	12.9	4.0	*0.0*	38.3	16.7
	$x\geqslant 10,000$	7,580	86.4	2.5	8.1	2.7	*0.0*	70.9	18.5
ITS orders	**All**	**33,228**	**5.4**	**2.6**	**12.2**	**4.2**	**75.6**	**20.7**	**18.9**
	$x<500$	12,333	4.6	4.7	16.3	9.1	65.3	2.3	17.8
	$500\leqslant x<2,000$	16,858	7.3	3.7	15.3	5.5	68.2	10.0	19.2
	$2,000\leqslant x<10,000$	3,847	4.8	1.6	11.2	3.0	79.3	26.5	18.9
	$x\geqslant 10,000$	190	*2.6*	1.6	6.3	3.7	85.8	44.2	18.4

[a] This table summarizes how different order types are executed relative to the quotes in effect 5 seconds before the order execution. For example, a 3,000 share market buy order execution at $12.0625 when the quoted prices (depths) are $12.000 (2,000)–$12.125 (2,000) will be tabulated as ''Inside'' with a depth improvement (DI). Similarly, a 3,000 share market buy order execution at $12.000 when the quoted prices (depths) are $12.0000 (2,000)–$12.1250 (2,000) will be tabulated as ''At Same Side'' with a depth improvement (DI). Buy and sell orders are lumped together. If more than one buy (sell) order execution of a particular type occurs in one trade, these buy (sell) order executions are aggregated. Quoted half-spread is defined as 10,000 times the difference between the offer and the mid-quote, divided by the mid-quote. Statistics for the All orders category are share-weighted. Figures in italics are not significantly different from zero.

Table 3
Trade executions for floor orders by execution size[a]

	Execution size	Number	Distribution of price relative to quotes						Quoted half-spread (basis points)
			Same side or better	Inside better	At mid-quote	Inside worse	Other side or worse	DI	
FB orders	**All**	**100,062**	**23.9**	**3.9**	**15.2**	**4.7**	**52.3**	**64.4**	**18.3**
	$x<500$	6,629	38.4	10.2	23.7	6.7	21.1	2.2	16.6
	$500\leqslant x<2,000$	28,657	38.9	6.7	19.4	6.1	28.9	15.3	15.6
	$2,000\leqslant x<10,000$	46,628	27.9	4.6	16.9	5.6	44.9	42.6	16.2
	$x\geqslant 10,000$	18,148	20.1	3.7	15.0	4.6	56.6	70.2	18.3
FB orders *Upstairs facilitated*	**All**	**4,517**	**37.4**	**5.0**	**10.7**	**5.1**	**41.9**	**97.3**	**27.6**
Percentage orders	**All**	**118,571**	**44.3**	**5.9**	**16.3**	**4.8**	**28.8**	**59.7**	**18.3**
	$x<500$	26,324	43.2	11.8	24.4	8.0	12.7	2.6	18.2
	$500\leqslant x<2,000$	48,336	42.5	9.0	20.5	7.1	20.9	17.5	17.4
	$2,000\leqslant x<10,000$	35,752	44.5	6.7	17.1	5.2	26.5	46.8	17.4
	$x\geqslant 10,000$	8,159	42.1	4.7	15.5	4.1	33.6	77.2	18.7
Specialist orders	**All**	**224,720**	**41.7**	**14.0**	**21.2**	**4.0**	**19.1**	**25.2**	**19.7**
	$x<500$	111,802	37.5	18.2	30.5	7.7	6.1	1.8	19.4
	$500\leqslant x<2,000$	91,066	44.9	15.7	25.3	5.3	8.9	14.8	19.8
	$2,000\leqslant x<10,000$	20,992	42.9	12.9	18.6	2.7	22.9	27.0	19.9
	$x\geqslant 10,000$	860	23.3	8.7	9.7	1.9	56.5	38.3	20.5

[a]This table summarizes how different order types are executed relative to the quotes in effect 5 seconds before the order execution. For example, a 3,000 share floor broker buy order execution at \$12.0625 when the quoted prices (depths) are \$12.000 (2,000)–\$12.125 (2,000) will be tabulated as "Inside" with a depth improvement (DI). Similarly, a 3,000 share floor broker buy order execution at \$12.000 when the quoted prices (depths) are \$12.0000 (2,000)–\$12.1250 (2,000) will be tabulated as "At the Same Side" with a depth improvement (DI). Buy and sell orders are lumped together. If more than one buy (sell) order execution of a particular type occurs in one trade, these buy (sell) order executions are aggregated. Quoted half-spread is defined as 10,000 times the difference between the offer and the mid-quote, divided by the mid-quote. Statistics for the All orders category are share weighted. Figures in italics are not significantly different from zero.

bid for sells) for system orders and floor orders respectively.[12] The tables also report the percent of trades with depth improvement (where the order exceeds the quoted size). Tables 2 and 3 also break down the measures by order execution size (*execution size* for short): less than or equal to 499 shares; between 500 and 1,999 shares; between 2,000 and 9,999 shares; and those exceeding 9,999 shares. Note that since a trade may have several types of orders represented on the buy and sell sides respectively, the execution size itself may be smaller than or equal to the actual trade size. Numbers for the ''all'' category are share-weighted, while numbers when broken down by execution size are execution-weighted. Executions for non-floor broker orders that are part of upstairs-facilitated transactions are excluded.

Slightly more than half the market order volume pays the spread. Twenty percent of market order volume instead executes at the best quotes on the same side or better, earning the full spread or more instead of paying it. An additional 25 percent of market order volume executes inside the best quotes. These orders pay a liquidity premium, but it is less than the quoted spread. Hence, market orders frequently get price improvement—the specialist may step in and improve the NYSE quotes (Edwards and Harris, 2001); the specialist may stop the order (Ready, 1999); or the orders may be exposed to floor broker represented orders that are not reflected in the NYSE quotes. Interestingly, the fraction of market orders that receives price improvement *increases* strongly in order size. The reason is that large orders are more likely to interact with the crowd at the specialist's post and hence benefit from the added non-displayed liquidity provided by floor brokers. Market orders in this sample get better execution than what is reported by Bacidore et al. (2002): 45 percent of market order volume receives price improvement in the 1997 sample compared to 36 percent in their 1999 sample. The difference is likely attributable to the fact that this data cannot account for multiple fills.[13] On average 35 percent of market order volume gets depth improvement.

A marketable limit order is defined as one where the purchase (sale) price is larger than (smaller than) or equal to the ask (bid) price *at the time of execution*. Hence, by definition all marketable limit orders execute at the other side of the market or worse. When using a marketable limit order, the submitter obtains priority over existing orders in the specialist's Display Book by offering a better price. These orders technically do not pay a spread, but the price improvement they offer compared to existing orders in the book can be thought of as a transaction cost (see, e.g., Peterson and Sirri, 2002). Marketable limit order executions are typically small and on average 31 percent of marketable limit order volume gets depth improvement.

[12] The qualitative results are very similar for liquid (deciles 7–10) and less liquid stocks (deciles 1–6) separately.

[13] If prices move in the direction of the trade, we will compare the execution to a more favorable benchmark than if we were able to pin-point the quotes at the time of order submission. Another difference is that we do not condition on whether or not the order received depth improvement.

As expected, 80 percent of limit order volume executes at the quotes at the same side or better. However, 20 percent of regular limit order volume executes inside the spread. In other words, one-fifth of all limit order volume is executed before it is displayed (reflected in the specialist's quotes). As for market orders, the ability to gain the spread increases strongly with the execution size of regular limit orders. On average 40 percent of trades based on regular limit orders get depth improvement.

Roughly three-quarters of ITS volume pays the spread or gets worse execution than the best quotes. The table also shows that the larger the ITS order, the worse is its execution. On average 21 percent of ITS volume gets depth improvement. Thus, while the execution pattern resembles that of system market orders, ITS orders systematically get worse execution than comparably sized system market orders. One possible explanation is that ITS orders by NYSE rules do not benefit from the auction on the NYSE floor.

In the sample, 52 percent of active floor broker volume pays the spread or more, while 24 percent gains the spread or better. The remaining 24 percent executes inside the best quotes. Thus, contrary to the assumptions in theoretical models (e.g., Seppi, 1997) floor orders do not act exclusively as liquidity suppliers. Instead, they frequently hit the opposite quotes, effectively acting as market orders. Both the fraction of volume that gains the spread and the fraction occurring inside the spread are declining in order size. By contrast, the fraction of executions that pays the spread is increasing in order size. This suggests that it is particularly for large, possibly price-sensitive, orders that floor brokers are willing to hit the opposite quotes. On average 64 percent of floor broker volume receives depth improvement.

For completeness, the table also reports the results for the upstairs-facilitated trades in the sample. Since the bulk of these consist of floor broker orders, only results for the floor broker component is reported. As expected, upstairs-facilitated orders are much larger than the regular actively represented floor broker executions. It is slightly more likely that the upstairs-facilitated volume executes at the other side or worse than at the same side or better. Trades inside the quoted spread represent 21 percent of all upstairs-facilitated volume. Not surprisingly, 97 percent of upstairs-facilitated volume receives depth improvement

For non-discretionary orders left by floor brokers with the NYSE specialist, or percentage orders, 44 percent of volume gains the spread or more, while 29 percent pays the spread on average. The remaining 27 percent of volume executes inside the quoted spread. The fraction of executions that gains the spread is pretty constant across order sizes, while the fraction that pays the spread increases. As a result, executions inside the spread decline as order sizes get larger. Depth improvement is offered to 60 percent of percentage volume.

Finally, the specialist gains the spread or more on 42 percent of executions, and pays the spread or more on 20 percent of his trades. Inside executions amount to 40 percent of specialist volume. The larger is the order, the more likely it is that the specialist actually pays the spread. Executions inside and executions gaining the spread both decrease in order size. On average, 25 percent of specialist volume gets depth improvement.

5. Spreads and information content by order type

The previous section shows that a large fraction of what are traditionally thought of as liquidity-demanding buy (sell) orders execute at or below (above) the mid-quote. Similarly, a large fraction of what are traditionally thought of as liquidity-supplying buy (sell) orders execute at or above (below) the mid-quote. This will tend to reduce the actual spread costs for liquidity-demanding order types, and reduce the actual spread gain for liquidity-supplying order types compared to estimated trading costs based on traditional measures.

To compute spreads, it is necessary to first identify a measure of the true value of the stock. This paper follows the literature in using the mid-quote at the time of the execution as the benchmark.[14] Thus, the effective half-spread is defined as the logarithm of the execution price (mid-quote) minus the logarithm of the mid-quote (execution price) for buy (sell) orders. The second spread measure, the realized half-spread, is defined as the logarithm of the execution price minus the mid-quote five minutes after the trade.[15] The change in the mid-quote for buys is defined as the logarithm of the mid-quote 5-minutes after the trade, minus the logarithm of the mid-quote at the time of the trade. The sign is reversed for sells. This means that the realized half-spread is the effective half-spread *minus* the change in the mid-quote. The change in the mid-quote is a measure of the information content of the trade. Hence, the realized half-spread incorporates the information that might have been conveyed by the order itself.

For reference, applying the standard Lee and Ready (1991) algorithm to trades in this sample generates an estimated effective half-spread of 16.3 basis points, an information content of 14.2 basis points, and a realized half-spread of 2.0 basis points (see Table 6 below). Moreover, the algorithm fails to classify over 100 thousand trades, or more than 1/6th of the sample's trades. This is consistent with the evidence presented earlier of a high incidence of orders executing at the mid-quote, which are difficult to classify with the algorithm. Recall that the Lee and Ready algorithm classifies trades as either buyer- or seller-initiated. By contrast, the methodology used in this paper accounts for each order on the buy *and* the sell side of every trade. Hence, only to the extent that there are systematic differences across order types in how they execute relative to the mid-quotes, and in their information content, are the actual spread and information measures going to be non-zero for individual order types.

Tables 4 and 5 report spread measures and information content for system and floor orders respectively.[16] Table 4 shows that effective half-spreads are highest for

[14] Ideally, the benchmark should be the quotes at the time of order submission. Unfortunately, the audit trail data does not record the time and size of the originally submitted orders.

[15] For example, Huang and Stoll (1996) use 5-minute and 30-minute windows, Bessembinder and Kaufman (1996, 1997) use 30-minute and 24-hour windows, Madhavan and Cheng (1997) use a 20-trade window, Handa et al. (1998a, b) use a 15-trade window, and the recent SEC Report (2001) uses a 5-minute window.

[16] The qualitative results are very similar for liquid (deciles 7–10) and less liquid stocks (deciles 1–6) separately.

Table 4
Spreads and information content for system orders by execution size[a]

	Execution size		Effective half-spread	Information content	Realized half-spread
Market orders	**All**	**352,703**	**6.6**	**3.7**	**3.0**
	$x<500$	175,726	6.9	1.7	5.2
	$500 \leqslant x < 2,000$	129,716	7.6	3.9	3.8
	$2,000 \leqslant x < 10,000$	44,058	7.3	4.1	3.3
	$x \geqslant 10,000$	3,203	4.4	3.4	*0.1*
Marketable limits	**All**	**141,957**	**14.3**	**14.3**	**0.0**
	$x<500$	36,554	12.1	10.0	2.2
	$500 \leqslant x < 2,000$	68,855	12.8	12.5	0.3
	$2,000 \leqslant x < 10,000$	33,349	14.1	13.9	*0.1*
	$x \geqslant 10,000$	3,199	16.3	16.8	*−0.5*
Limit orders	**All**	**414,220**	**−12.6**	**−14.1**	**1.5**
	$x<500$	167,635	−10.2	−5.8	−4.5
	$500 \leqslant x < 2,000$	168,948	−10.3	−8.8	−1.4
	$2,000 \leqslant x < 10,000$	70,055	−12.1	−13.9	1.8
	$x \geqslant 10,000$	7,580	−15.4	−19.3	4.1
ITS orders	**All**	**33,228**	**12.1**	**8.7**	**3.4**
	$x<500$	12,333	9.5	5.9	3.6
	$500 \leqslant x < 2,000$	16,858	10.1	6.0	4.2
	$2,000 \leqslant x < 10,000$	3,847	12.9	9.1	3.9
	$x \geqslant 10,000$	190	14.8	13.7	*1.0*

[a] This table reports effective half-spreads, defined as 10,000 times ln(price(t)/mid-quote(t)) [ln(mid-quote(t)/price(t))] for a buy [sell] order, the information content, defined as 10,000 times ln(mid-quote($t+T$)/mid-quote(t)) [ln(mid-quote(t)/mid-quote($t+T$))] for a buy [sell] order, and realized half-spreads, defined as 10,000 times ln(price(t)/mid-quote($t+T$)) [ln(mid-quote($t+T$)/price(t))] for a buy [sell] order, where T is 5 minutes. Spreads for the All category for each order type are share weighted. Figures in italics are not significantly different from zero.

marketable limit orders at 14.3 basis points followed by ITS orders at 12.1 basis points. By contrast, the effective half-spread for market orders is only half as large, 6.6 basis points which is natural since as much as 45 percent of these orders get price improvement. Spread costs for liquidity-demanding orders are generally increasing in order size. By contrast, limit orders tend to supply liquidity and they gain the spread on average. The effective half-spread is −12.6 basis points, and the spread gain is increasing in order size.

The information content of marketable limits perfectly offsets the spread, so the realized half-spreads are zero on average, and only for small orders are they significantly positive. While quotes move in the direction of ITS orders by an average 8.7 basis points, this is not sufficient to outweigh the initial spread and the realized half-spreads are significantly positive at 3.4 basis points as a result. The information content of marketable limits and ITS orders is strongly increasing in order size. By

Table 5
Spreads and information content for floor orders by execution size[a]

	Execution size		Effective half-spread	Information content	Realized half-spread
FB orders	**All**	**100,062**	**5.4**	**7.1**	**−1.6**
	$x<500$	6,629	−2.4	4.9	−7.3
	$500\leqslant x<2,000$	28,657	−0.7	6.4	−7.0
	$2,000\leqslant x<10,000$	46,628	3.2	6.2	−3.0
	$x\geqslant 10,000$	18,148	6.8	7.6	*−0.7*
FB orders *Upstairs facilitated*	**All**	**4,517**	**1.4**	**1.2**	**0.2**
Percentage orders	**All**	**118,571**	**−2.7**	**−2.2**	**−0.7**
	$x<500$	26,324	−5.1	−0.9	−4.1
	$500\leqslant x<2,000$	48,336	−3.4	*−0.2*	−3.1
	$2,000\leqslant x<10,000$	35,752	−3.1	−1.1	−1.9
	$x\geqslant 10,000$	8,159	−1.8	−2.3	*0.4*
Specialist orders	**All**	**224,720**	**−6.9**	**−2.3**	**−4.6**
	$x<500$	111,802	−6.2	−5.0	−1.1
	$500\leqslant x<2,000$	91,066	−8.0	−5.7	−2.4
	$2,000\leqslant x<10,000$	20,992	−7.3	−1.2	−6.2
	$x\geqslant 10,000$	860	−1.4	9.2	−10.5

[a] This table reports effective half-spreads, defined as 10,000 times ln(price(t)/mid-quote(t)) [ln(mid-quote(t)/price(t))] for a buy [sell] order, the information content, defined as 10,000 times ln(mid-quote($t+T$)/mid-quote(t)) [ln(mid-quote(t)/mid-quote($t+T$))] for a buy [sell] order, and realized half-spreads, defined as 10,000 times ln(price(t)/mid-quote($t+T$)) [ln(mid-quote($t+T$)/price(t))] for a buy [sell] order, where T is 5 minutes. Spreads for the All category for each order type are share weighted. Figures in italics are not significantly different from zero.

comparison, the information content of market orders is only 3.7 basis points on average, which is insufficient to offset the initial spread. As a result, realized half-spreads for market orders are positive at 3.0 basis points. Interestingly, the information content of market orders is non-monotonic in order size and the largest information content is for market orders between 500 and 10,000 shares.[17] Finally, the spread gain that limit orders garner at the time of the trade is insufficient to compensate for the adverse selection risk of the orders they trade with. This is consistent with limit orders getting "picked-off" by traders with better information. The prices move against limit orders by 14.1 basis points on average, and the resulting realized half-spread for these orders is significantly positive at 1.5 basis point. Note, however, that realized half-spreads for limit orders differ significantly across order sizes: small limit orders gain the spread even after accounting for the adverse selection while large limit orders actually face negative realized half-spreads.

[17] This is consistent with Barclay and Warner's (1993) finding that medium sized trades are most informative.

Table 5 repeats the exercise for floor orders. Among floor orders, only orders traded actively by floor brokers face significantly positive, albeit low, effective half-spreads. This is natural as floor broker orders are a mix of liquidity-demanding and liquidity-supplying orders. Indeed, note that small floor broker orders gain the spread while orders exceeding 9,999 shares pay as much as 6.8 basis points on average. Interestingly, spreads for orders traded by floor brokers are lower than for market orders despite their very similar execution patterns relative to the quotes. There are three potential reasons for why spread costs are lower for floor brokers. First, the auction on the NYSE floor often results in trades between floor brokers executed between the displayed quotes. Second, floor brokers are present on the NYSE floor and might therefore be better able to adjust their trading strategies to market conditions (see, Sofianos and Werner, 2000). Third, Sofianos and Werner (2000) show that floor brokers are more active in liquid stocks that have lower spreads. By contrast, orders left by floor brokers with the NYSE specialist earns the effective half-spread. On average, the spread is −2.7 basis points and it is increasing in order size. In other words, it is mainly smaller percentage orders that gain the spread. NYSE specialists gain the effective half-spread on average, but the gain is at 6.9 basis points much smaller than the average gain for limit orders. This is consistent with the specialist frequently offering price improvement. Moreover, the spread gains are non-monotonic in order size. These results are consistent with the specialist's fiduciary obligation to help bridge gaps in liquidity for smaller orders.

Prices move significantly in the direction of active floor broker orders by 7.1 basis points, and the information content is increasing in order size. This suggests that a substantial portion of floor broker orders is information based. As a result, the realized half-spread for floor broker orders is negative on average at −1.6 basis point. By contrast, prices move against passive percentage orders by 2.2 basis points on average, and large percentage orders face significant adverse selection risk. However, floor brokers using percentage orders manage to maintain the spread gain even after accounting for adverse selection risk. While prices move against the NYSE specialists on average, in the relatively few instances when specialists in the sample trade orders in excess of 9,999 shares the prices move in the direction of their orders. Thus, the data suggests that the specialists are skilled in timing their (large) discretionary trades. Interestingly, the adverse selection risk facing specialists tends to decrease in order size. Note also that NYSE specialist orders experience spread-gains even after accounting for adverse selection risk.

One striking feature of Tables 4 and 5 is that effective half-spreads for each order type are very low compared to the ones estimated by applying the Lee and Ready (1991) algorithm to sample *trades*, e.g., an effective half-spread of 16.3 basis points. To understand what is going on, consider Table 6. It reports the estimated execution costs *for each order type* based on the Lee and Ready (1991) algorithm applied the following way. First, assume that the algorithm only knows the order type, not the direction of the order. Second, assume that market orders, marketable limits, ITS orders, floor broker orders, and upstairs orders are liquidity demanding, and that limit orders, percentage orders, and specialist orders are

Table 6
Spreads and information content based on Lee and Ready (1991)[a]

	Percent incorrect	Effective half-spread	Information content	Realized half-spread
A. Trades	**n.a.**	**16.3**	**14.2**	**2.0**
$x<500$	n.a.	12.0	6.3	5.7
$500 \leqslant x < 2,000$	n.a.	13.2	10.7	2.4
$2,000 \leqslant x < 10,000$	n.a.	14.2	13.3	0.9
$x \geqslant 10,000$	n.a.	16.4	15.5	0.9
B. Liquidity demanding orders				
Market order	29.7	15.8	11.4	4.3
Marketable limit order	0.0	14.3	14.3	0.0
ITS order	21.0	16.0	12.9	3.1
Floor broker order	50.4	16.4	15.6	0.8
Upstairs order	54.4	25.2	16.8	8.4
C. Liquidity supplying orders				
Limit order	25.0	−14.4	−15.2	0.8
Percentage order	30.6	−15.9	−12.5	−2.9
Specialist order	16.9	−14.5	−12.0	−2.5

[a] Panel A displays the results of applying the original Lee and Ready (1991) algorithm to sample trades overall and broken down by trade size. In Panels B and C, the Lee and Ready (1991) algorithm is applied separately to each order type. Signs for liquidity supplying orders (Panel C) are the reverse of those assumed by the algorithm, e.g., the effective half-spread for a buy order is 10,000 times the logarithm of the execution price minus the logarithm of the mid-quote at the time of the trade. Percent incorrect refers to the number of orders that are either not classified (mid-quote) or erroneously classified by the algorithm.

liquidity-supplying orders. Third, assume that the algorithm is told what limit orders are marketable.[18]

The column labeled "percent incorrect" in Table 6 reports the fraction of orders that are either not classified (i.e., orders at the mid-quote that do not occur on an up- or a downtick) or incorrectly classified (i.e., a buy (sell) order below (above) the mid-quote gets counted as a seller-initiated (buyer-initiated) order). As much as 30 percent of market orders and 21 percent of ITS orders are classified incorrectly. For order involving floor brokers, the algorithm misclassifies over 50 percent of the orders. Marketable limit orders are correctly classified because of the third assumption. Liquidity supplying orders like limit orders, percentage orders, and specialist orders, would also frequently erroneously be classified as liquidity-demanding orders.

Table 6 also shows that effective half-spreads and information content of orders are dramatically overestimated when estimated based on the Lee and Ready (1991)

[18] If the algorithm instead is asked to treat all limit orders that execute outside of the quotes as marketable, many buy orders below the bid would erroneously be classified as liquidity-demanding sell orders.

algorithm. Effective half-spreads for market orders would be 15.8 basis points on average based on the algorithm, which is more than double the *actual* effective half-spread for market orders—6.6 basis points (Table 4). The algorithm exaggerates the information content by a similar amount, but the problem is less severe for realized spreads. The difference across methodologies between the estimated spreads for ITS orders is less significant, but they are a small part of the sample. Effective half-spreads and information content for floor broker orders are estimated to be more than three times as high as they actually are (e.g., effective half-spreads of 16.4 basis points compared to 5.4 basis points (Table 5)). For limit orders, spread revenues as well as adverse selection (negative of information content) are also overstated, but the magnitude is relatively small. The table also shows that spread gains and adverse selection for specialist orders and percentage orders are overestimated using the Lee and Ready algorithm. The nature of the problem with trade classification algorithms is perhaps clearest for upstairs-facilitated trades. The Lee and Ready algorithm estimates effective (realized) half-spreads to be 25.2 (8.4) basis points when the *actual* spread is 1.4 (0.2) basis points.

It is important to point out that there are problems with trade-classification algorithms across all order sizes and stock liquidity (not reported). The fraction of orders that are either not classified or erroneously classified by the Lee and Ready (1991) algorithm is very similar across order sizes and are almost identical for large capitalization (top 4 deciles) and small capitalization (bottom 6 deciles) stocks. As a result, effective half-spreads are overestimated for liquidity-demanding orders and spread gains are overestimated for liquidity-supplying orders across all order sizes and size-groups. For example, estimated effective half-spreads for market orders in large (small) capitalization stocks are 14.1 (33.7) basis points compared to actual spreads of 5.8 (15.0) basis points. For floor broker orders, estimated effective half-spreads in large (small) capitalization stocks are 15.5 (33.7) basis points compared to actual spreads of 4.8 (15.2) basis points.

6. Order flow characteristics and spreads and information content of trades

Unfortunately, most researchers do not have access to intraday order data from the NYSE. It is therefore useful to consider how the composition of order flow affects spreads and information content that would be estimated by an econometrician using for example TAQ data and relying on the Lee and Ready (1991) algorithm to classify trades as either buyer- or seller-initiated. Implicitly, this analysis controls for the aggressiveness of the orders by classifying orders relative to the mid-quote.

Aside from the order composition of the trade itself, the literature suggests that there are four main sources of cross-sectional and time-series variation in spreads. First, spreads are affected by stock-specific characteristics relating to liquidity and risk. The previous literature has found that spreads decrease in liquidity (dollar volume and market capitalization) and increase in risk (volatility).[19] Second, the

[19] See, e.g., Ho and Stoll (1981).

literature suggests that the tick size may be important, and it has been found that spreads tend to increase in the inverse of the stock price.[20] Third, spreads are affected by trading conditions during the day. Spreads have been found to be higher for orders in the direction of an order-imbalance.[21] That is, spreads for a buy (sell) order would be higher (lower) if there is a buy-imbalance or if the depth at the offer (bid) is lower than the depth at the bid (offer). The literature also suggests that spreads should be tighter if trading volume is unusually high. Further, spreads have been shown to decline during the course of the day, which suggests that a time of day effect may be important.[22] Fourth, the size of the trade itself is likely to affect spreads. Theory predicts, and the literature has generally found, that large trades pay higher effective half-spreads, and that they have larger information content.[23]

The control variables are defined as follows. Market capitalization is defined as the logarithm of market capitalization as of December 31, 1996. Liquidity is further captured by the logarithm of average daily dollar volume. Volatility is defined as 10,000 times the daily average logarithm of (high/low). The inverse of dollar price is taken to be the mid-quote 5 seconds prior to the trade. Depth imbalance is defined as the depth at the opposite side to the trade in effect 5 seconds prior to the trade divided by the sum of the depths at the bid and offer. The second measure of order imbalance is the cumulative signed (in the direction of the trade) dollar volume divided by the average daily dollar volume for the stock.[24] The time of day takes on a value of 1 for trades between 9:30 and 11:30, takes on a value of 2 for trades between 11:30–2:00, and a value of 3 for trades between 2:00 and 4:00. Liquidity during the day is defined as the cumulative dollar volume divided by average daily dollar volume for the stock. Trade size is expressed as the logarithm of the number of shares divided by 1,000. All control variables except trade size are demeaned to facilitate interpretation of the estimated parameters.

To capture the order flow composition, the regressions use variables coded to capture the net orders of a particular type in the direction of the trade. Hence, the regressions test if spreads, $S \in$ {Effective Half-Spreads, Information Contents, and Realized Half-Spreads} depend significantly on the composition of order flow after controlling for other variables that might affect spreads:

$$\begin{aligned} S = \alpha &+ \beta_1 \text{ Mkt.Cap} + \beta_2 \text{ Volatility} + \beta_3 \text{ \$Volume} + \beta_4 (1/\text{Price}) \\ &+ \beta_5 \text{ Depth.Imb} + \beta_5 \text{ Cum.Signed.Vol} + \beta_6 \text{ Cum.Vol} + \beta_7 \text{ Time.of.Day} \\ &+ \beta_8 \text{ Trade.Size} + \sum_{j=1}^{8} \gamma_j N_j + \varepsilon, \end{aligned} \quad (1)$$

[20] Harris (1994) makes this point.

[21] Huang and Stoll (1994) make this point and also Handa et al. (1998a) suggest that this variable is an important determinant of bid-ask spreads in order driven markets. See also Handa et al. (1998b).

[22] E.g., Foster and Viswanathan (1993) and Werner and Kleidon (1996) show that spreads decrease in the time of day and in cumulative volume as information asymmetry is reduced.

[23] See, e.g., the model developed by Easley and O'Hara (1987).

[24] Dollar volume is signed based on the Lee and Ready (1991) algorithm.

where N_j takes on the value of 1 (0) [−1] if there is a (a) [no] buy order and no (a) [a] sell order of type j is present in the trade, $j \in$ {market order, marketable limit order, limit order, ITS order, floor broker order, percentage order, specialist order, other order}. Note that N_j predominantly takes on a non-negative value for what is usually labeled liquidity-demanding orders, and a non-positive value for liquidity-supplying orders. Buyer- and seller-initiated trades, properly signed, are included in the regressions. The regressions are estimated using GMM to obtain robust standard errors.

Table 7 reports the regression results. The stock characteristics are highly significant and, except for market capitalization, have the predicted sign: spreads and information content are increasing in volatility; spreads and information content are decreasing in dollar volume; and spreads are higher for low priced stocks and trades in these stocks move prices more.[25] The variables capturing trading day characteristics are also significant and are of the predicted sign. Spreads and information contents are lower the more depth there is at the opposite side of the trade, are increasing in signed dollar volume preceding the trade, are decreasing in dollar volume preceding the trade, and are declining over the course of the trading day. Finally, as predicted, larger trades are associated with significantly wider spreads and more substantial information content.

The rest of the parameters in the regressions capture the effect of order flow, after controlling for the other factors. With the exception of the residual "other orders", all order flow variables are highly significant both for effective half-spreads and information content. The way to read the regression results is as follows: a trade executing above the mid-quote with a market buy order (1) trading against a limit sell order (−1) has on average a 2.76 (0.904 + 1.851) higher spread. A trade executing above the mid-quote with a market buy order (1) trading against a floor broker order (−1) has on average a 0.14 basis point higher spread (0.904 − 0.768). ITS orders pay even higher effective half-spreads, 3.05 basis points against a limit order and 0.43 basis points against a floor broker order. Similar calculations show that a market order and ITS orders in the direction of the trade trading with a percentage order or with the specialist also pay lower effective half-spreads than a market order trading with a limit order would. Presumably, this comes from the price improvements these market participants have to offer to gain priority over the book. The results thus suggest that market order submitters are able to lower their spread costs significantly because of the presence of the crowd.

The estimated information content is highest for floor broker orders when other factors are controlled for. A floor broker order executing above the mid-quote against a limit order will move mid-quotes as much as 8.81 basis points in the direction of the trade. A marketable limit executing above the mid-quote against a limit order will move mid-quotes considerably less, by 4.72 basis points. Note also that the information content of a trade is drastically reduced if it is a floor broker order or percentage order that supplies liquidity. This further reinforces the

[25] Higher market capitalization is associated with significantly higher spreads after controlling for dollar trading volume. Hence, stocks with lower turnover have higher spreads all else equal in this sample.

Table 7
Effective half-spreads, information content, realized half-spreads and trade composition[a]

	Effective half-spreads	Information content	Realized half-spreads
Intercept	**41.450**	**49.407**	**−8.163**
Stock characteristics			
Log(market capitalization)	**0.549**	0.032	**0.525**
Volatility(high/low)	**0.262**	**0.424**	**−0.162**
Log($ volume)	**−3.223**	**−3.095**	−0.126
Inverse of price	**2.562**	**0.330**	**2.233**
Trading day characteristics			
Depth imbalance	**−0.322**	**−2.333**	**1.984**
Cumulative buy volume	**3.705**	5.245	−1.578
Cumulative volume	**−2.335**	**−2.432**	0.087
Time of day	**−0.457**	**−0.950**	**0.484**
Trade characteristics			
Trade size	**1.016**	**2.158**	**−1.144**
Net order type in direction of trade			
Market orders	**0.904**	**−0.597**	**1.526**
Marketable limit orders	**−0.150**	**3.122**	**−3.235**
Limit orders	**−1.851**	**−1.623**	−0.209
ITS orders	**1.200**	**0.694**	0.535
Floor broker orders	**0.768**	**7.184**	**−6.384**
Percentage orders	**0.448**	**4.409**	**−3.924**
Specialist orders	**−1.139**	**−1.203**	0.060
Other orders	0.218	−0.273	0.559
Adj. R-square	0.437	0.045	0.047
Observations	543,911	535,493	535,493

[a]The table summarizes GMM regressions of downstairs spreads and information content, on a set of control variables, trade size, and dummy variables for the net presence of order-types in the direction of the trade. The Lee and Ready (1991) algorithm is used to determine trade direction. Market capitalization is measured in as of December 31, 1996. Volatility is measured as 10,000 times the daily average ln(high/low) during the sample period for the stock. Dollar volume is average daily dollar volume over the sample period. The inverse of the price is defined as 100 times the inverse of the mid-quote 5 minutes prior to the trade. Depth imbalance is defined as the depth in effect 5 seconds prior to the trade on the opposite side divided by the sum of the depths at the bid and the offer. Cumulative signed volume is the net-buy volume executed prior to the trade in the direction of the trade during the trading day, divided by average daily share volume. Cumulative volume is unsigned share volume prior to the trade execution during the trading day, divided by average daily share volume. Time of day takes on the value of 1 if the trade occurs between 9:30 and 11:30, the value of 2 if the trade occurs between 11:30 and 2:00, and the value of 3 if the trade occurs between 2:00 and 4 : 00 p.m. Trade size is measured in the logarithm of shares divided by one thousand. For each trade, the net orders of each order type in the direction of the trade are recorded, e.g., a market buy and no (a) market sell for a buyer-initiated trade is recorded as a 1 (0). The table pools buyer and seller initiated trades. Boldfaced type indicates that the coefficient is significant at the 1 percent level.

conjecture that informed traders are using floor brokers to execute their trades, and that the most price-sensitive orders are actively represented at the specialist's post. Cooney and Sias (2001) reach the same conclusion based on 1990–1991 TORQ data.

Surprisingly, a market buy order executing above the mid-quote against a limit order will move the mid-quotes only by 1.02 basis points. By contrast, if a market order trades together with a floor broker order against a limit order, the mid-quotes will move by 8.21 $(-0.597 + 7.184 + 1.623)$ basis points. Thus, it is likely that the relatively high information content found for market orders in the order-based analysis results from market orders often trading side by side with floor brokers (or other informative orders, e.g., marketable limits). These results suggest that market orders in and of themselves have relatively limited information content.

7. Value-weighted average trading revenues by order type

The analysis so far uses the contemporaneous mid-quotes and mid-quotes 5 minutes following the trade to evaluate spreads and information content. The final analysis instead uses a more long-term benchmark—the daily volume-weighted average price (VWAP)—to compute trading revenues by order type. The analysis is motivated by the simple observation that if the traders using a particular order type are systematically better informed, they should on average buy at a lower price and sell at a higher price than traders favoring other order types.

The daily, d, volume-weighted average execution price, A_{dj}, is first calculated for all stocks, j, based on all trades, i:

$$A_{dj} = \sum_{i=1}^{T} \frac{P_i V_i}{\sum_{i=1}^{T} V_i}, \tag{2}$$

where P_i is the execution price and V_i is the number of shares. For each order type, k, and each stock, j, the daily volume-weighted average buy (sell) price, $B_{dkj}^{B(S)}$, is then computed as

$$B_{dkj}^{B(S)} = \sum_{i=1}^{T} \frac{P_{ik}^{B(S)} V_{ik}}{\sum_{i=1}^{T} V_{ik}}, \tag{3}$$

where $P_{ik}^{B(S)}$ is the execution price of a buy (sell) order of type k, and V_{ik} is the number of shares in the trade of order type k. The normalized measure of volume-weighted average trading revenue for orders of type k is

$$\mathrm{VWTR}_k = 10{,}000\,\frac{1}{2}\left(\sum_{d=1}^{D} \sum_{j=1}^{N} (B_{dkj}^{S} - B_{dkj}^{B})/A_{dj}\right). \tag{4}$$

To make the implicit spread gains and spread costs are comparable to those reported in earlier tables, trading revenues are divided by two. Note that only stock-days with both buy orders and sell orders by the same order type are included.

The results are in Table 8. Trading profits are highest for specialists, 13.1 basis points, followed by regular limit orders, 13.0 basis points. Thus, despite the fact that

Table 8
Volume-weighted average trading revenues by order execution type[a]

	Market orders	Marketable limit orders	Limit orders	ITS orders	Floor broker orders	Percentage orders	Specialist orders
All trades	**−13.71**	**−8.76**	**12.96**	**−3.83**	**−2.50**	**2.62**	**13.12**
(t-stat: $H_0 = 0$)	(−33.19)	(−17.06)	(32.51)	(−4.84)	(−4.53)	(5.09)	(28.78)
Number of stock-days	4,248	3,456	4,341	2,258	2,475	2,761	3,981
$x \leqslant 500$ *shares*	−12.05	−4.65	11.43	−0.87	−0.08	2.39	9.71
(t-stat: $H_0 = 0$)	(−30.82)	(−6.65)	(26.37)	(−0.72)	(−0.04)	(2.65)	(21.78)
Number of stock-days	3,850	2,282	3,782	1,291	572	1,756	3,547
$500 \leqslant x < 2{,}000$ *shares*	−12.41	−6.71	10.65	−1.27	2.43	1.67	12.15
(t-stat: $H_0 = 0$)	(−27.79)	(−11.32)	(24.13)	(−1.23)	(2.29)	(2.43)	(22.37)
Number of stock-days	3,704	2,791	3,718	1,529	1,414	2,190	3,246
$2{,}000 \leqslant x < 10{,}000$ *shares*	−9.05	−7.01	10.91	−2.26	−1.60	1.31	9.48
(t-stat: $H_0 = 0$)	(−12.52)	(−9.79)	(18.19)	(−0.87)	(−2.29)	(2.17)	(9.30)
Number of stock-days	2,249	1,959	2,422	389	1,970	2,123	1,513
$x \geqslant 10{,}000$ *shares*	−3.45	−8.35	20.65	3.13	−3.95	−0.33	−3.32
(t-stat: $H_0 = 0$)	(−1.45)	(−3.35)	(12.13)	(0.14)	(−5.63)	(−0.34)	(−0.50)
Number of stock-days	385	367	595	8	1,299	1,095	93

[a] Volume-weighted average buy and sell prices are calculated for downstairs executions based on each order type for each stock and day (B, S). In addition, the volume-weighted average trade price based on all trades for each stock and day (A) are computed. For buy orders, the equally-weighted average across stock-days of 10,000 times one minus the ratio of the value-weighted average price for an order type, divided by the volume-weighted average trade price based on all trades (i.e., $10{,}000 * (1 - B/A)$). The sign is reversed for sell orders. The table reports volume-weighted sells minus buys, or trading revenues $10{,}000 * ((S - B)/A)$. For comparability with the other tables in this paper, the figures are divided by two. A positive number means better than average execution. Finally, statistics are broken down by order size. t-Statistics are reported in parentheses.

specialists and limit order traders face substantial short-term adverse selection as illustrated in previous sections, they capture sufficiently large spread gains to more than compensate for this on a daily basis. More substantial losses are faced by marketable limit orders at 8.8 basis points, and particularly market orders at 13.7 basis points. This is consistent with the earlier findings that market orders have limited short-term information content, yet pay significant spreads. By comparison, trading losses/gains are smaller in magnitude for percentage orders and floor brokers. Percentage orders are by definition mostly VWAP trades, so it is not surprising that trading profits for these orders are close to zero. For floor broker orders, the explanation lies in that these orders are a mix of liquidity-supplying and liquidity-demanding orders. Moreover, they are a large proportion of floor broker orders are evaluated against the VWAP.

The table also reports trading profits by order execution size. Note that the benchmark value-weighted execution price for these calculations is based only on trades of the same size as the specific order execution size. Market orders and marketable limit orders lose, and limit orders gain, for all order execution sizes. ITS orders lose for small orders only, while percentage orders and specialist orders generate significantly positive trading revenues for all sizes except the very largest. Floor-broker executions below 2,000 shares are not significantly different from the benchmark. Hence, it is the larger executions that cause overall trading profits to turn slightly negative.

8. Conclusions

This paper provides a comprehensive analysis of NYSE spreads and information content for different order types. While several authors previously have studied market orders, marketable limit orders, and specialist trading, no one has previously studied NYSE ITS orders, floor broker orders and percentage orders in detail.

A large fraction of orders obtain price improvements compared to the displayed liquidity in the NYSE specialist's quotes. For example, almost half the market orders and floor broker orders receive price improvement. Price improvements are often offered by NYSE specialists, and may be obtained by interacting with non-displayed liquidity on the NYSE trading floor. The results show that traditional methods for estimating spreads based on classifying trades into buyer- and seller-initiated, e.g., Lee and Ready (1991), significantly overestimate spreads and information content particularly for market orders but also for other order types. While there is nothing wrong per se with labeling more aggressive trades as liquidity demanding, it is problematic when the estimates are used to proxy for the spread costs facing system market orders (or the spread gains for system limit orders).

The order-based analysis shows that both effective half-spreads and information content of market orders are quite limited. Interestingly, marketable limits face by far the highest effective half-spreads, and also have larger information content. This suggests that informed traders favoring system orders are using marketable limits.[26] By comparison, floor broker orders pay relatively low spreads on average but have substantial information content. The evidence thus supports the notion that floor brokers, or their off-floor clients, have superior information and/or better timing ability than traders using SuperDOT. The results also show that limit orders and specialists get "picked off" on average as suggested by theory. However, note that these order types do not systematically lose money since they earn spread-gains that are sufficient to compensate for the adverse selection over a longer horizon.

This study shows how important the composition of order flow is for commonly used measures of NYSE spreads and information content. The trade-based analysis

[26] Alternatively, marketable limits are more frequently used in less liquid, more volatile stocks.

shows that measures based on trade classification algorithms are significantly affected by order flow composition, even after controlling for stock characteristics, trading day characteristics, and trade size. The regression analysis also shows that market orders benefit from the non-displayed liquidity on the NYSE floor. Moreover, they show that liquidity-demanding floor broker orders carry substantial information, reinforcing the previous results that informed off-floor traders are willing to pay the higher cost of hiring a floor broker to work their orders. Contrary to theoretical predictions, the regressions show that regular market orders do not in and of themselves have significant information content. Only when they trade side by side with other informed orders do prices move significantly in the direction of these orders.

Two caveats are in order. First, the NYSE audit trail data does not reveal the original order submission times, nor does it reveal the original size of orders that receive multiple executions. Therefore, it is not possible to account for systematic differences across order types in the delay between order arrival and execution. Bacidore et al. (2002) show that the delay for marketable system orders is small, but it is likely to be much higher for both limit orders and for orders routed to NYSE floor brokers. The NYSE recently announced that it is moving to electronic capture of floor broker orders. This means that better data will soon be collected, and hopefully be made available also to academics. Second, neither the order-based nor the trade-based analyses control for the fact that the order-type selection is endogenous. Indeed, preliminary work suggests that order-type selection is significantly affected by stock characteristics and trading day characteristics. An avenue for future research is therefore to control for endogeneity as suggested by Madhavan and Cheng (1997).

Aside from characterizing trading on the NYSE, and estimating spread costs and information content for different order types, another important lesson can be drawn from this paper: It is bound to be difficult, probably even impossible, to accurately compare spreads across markets with different order flow composition.

References

Bacidore, J., Ross, K., Sofianos, G., 2002. Quantifying best execution at the New York Stock Exchange: market orders. Journal of Financial Markets, this issue.

Barclay, M.J., Warner, J.B., 1993. Stealth trading and volatility: which trades move prices? Journal of Financial Economics 34, 281–305.

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

Blume, M., Goldstein, M., 1997. Quotes, order flow, and price discovery. Journal of Finance 52, 221–244.

Chan, K., Fong, W.-M., 2000. Trade size, order imbalance, and the volatility-volume relationship. Journal of Financial Economics 57, 247–273.

Chan, L.K.C., Lakonishok, J., 1993. Institutional trades and stock price behavior. Journal of Financial Economics 33, 173–199.

Choe, H., Kho, B.-C., Stulz, R.M., 2000. Do domestic investor have more valuable information about individual stocks than foreign investors? Dice-Center Working Paper No. 21, December.

Cooney, J.W., Sias, R.W., 2001. Informed trading and order type. Working paper, University of Kentucky.
Easley, D., O'Hara, M., 1987. Price, trade-size, and information in securities markets. Journal of Financial Economics 19, 69–90.
Edwards, A.K., Harris, J.H., 2001. Stepping ahead of the book. Working paper, University of Notre Dame, May.
Ellis, K., Michaely, R., O'Hara, M., 2000. The accuracy of trade-classification rules: Evidence from Nasdaq. Journal of Financial and Quantitative Analysis 35, 529–551.
Foster, F.D., Viswanathan, S., 1993. Variations in trading volume, returns volatility and trading costs: evidence on recent price formation models. The Journal of Finance 48, 187–211.
Handa, P., Schwartz, R., Tiwari, A., 1998a. Determinants of the bid–ask spread in an order driven markets. Working paper, University of Iowa.
Handa, P., Schwartz, R., Tiwari, A., 1998b. The value of the AMEX trading floor. Working paper, September.
Harris, L., 1994. Minimum price variations, discrete bid–ask spreads, and quotation sizes. The Review of Financial Studies 1, 148–178.
Harris, L., Hasbrouck, J., 1996. Market vs. limit orders: the SuperDOT evidence on order submission strategy. Journal of Financial and Quantitative Analysis 31, 213–231.
Hasbrouck, J., Sofianos, G., Sosebee, D., 1993. New York Stock Exchange systems and trading procedure. NYSE Working Paper No. 93-01.
Ho, T., Stoll, H., 1981. Optimal dealer pricing under transactions and return uncertainty. Journal of Financial Economics 9, 47–73.
Holthausen, R., Leftwich, R., Mayers, D., 1987. The effect of large block transactions on security prices: a cross-sectional analysis. Journal of Financial Economics 26, 237–267.
Holthausen, R., Leftwich, R., Mayers, D., 1990. Large-block transactions, the speed of response, and temporary and permanent stock-price effects. Journal of Financial Economics 26, 391–398.
Huang, R., Stoll, H., 1996. Dealer versus auction markets: a paired comparison of execution costs on NASDAQ and the NYSE. Journal of Financial Economics 41, 313–358.
Keim, D.B., Madhavan, A., 1996. The upstairs market for large-block transactions: analysis and measurement of price effects. The Review of Financial Studies 9, 1–36.
Kraus, A., Stoll, H., 1972. Price impacts of block trading on the New York Stock Exchange. Journal of Finance 27, 569–588.
Lee, C.M.C., Ready, M.J., 1991. Inferring trade direction from intraday data. Journal of Finance 46, 733–746.
Madhavan, A., Cheng, M., 1997. In search of liquidity: block trades in the upstairs and downstairs markets. The Review of Financial Studies 10, 175–203.
Odders-White, E., 2000. On the occurrence and consequences of inaccurate trade classification. Journal of Financial Markets 3, 259–286.
Peterson, M., Sirri, E., 2002. Limitations of trade and quote data: biases in execution cost estimators. Journal of Financial Markets, this issue.
Ready, M.J., 1999. The specialist's discretion: stopped orders and price improvement. The Review of Financial Studies 12, 1075–1112.
SEC, 2001. Report on the comparison of order executions across equity market structures. Office of Economic Analysis, January.
Seppi, D., 1997. Liquidity provision with limit orders and a strategic specialist. The Review of Financial Studies 10, 103–150.
Sofianos, G., Werner, I.M., 2000. The trades of NYSE floor brokers. Journal of Financial Markets 3, 139–176.
Werner, I.M., 1999. Who gains from steenths. Working paper, Stanford University.
Werner, I.M., Kleidon, A.W., 1996. U.K. and U.S. trading of British cross-listed stocks: an intraday analysis of market integration. The Review of Financial Studies 9, 615–659.