

Order Preferencing and Market Quality on U.S. Equity Exchanges

Mark A. Peterson

Southern Illinois University

Erik R. Sirri

Babson College

We present a detailed view of market quality in the presence of preferencing arrangements. A unique dataset provides the opportunity to measure trading costs of marketable orders and fill rates and ex post costs of limit orders across trading venues. For market orders, we find the primary exchange provides the lowest execution costs. However, the preferencing exchanges are no worse than, and in most cases better than, the nonpreferencing regional exchanges. For limit orders, the regionals execute limit orders more frequently than the primary market and with an ex post execution cost that is not very different from the primary market.

We analyze the execution quality of retail order flow on the primary and regional U.S. equity exchanges. Over the past decade the market share of regional exchanges in retail-size orders has risen dramatically. In 1996 the New York Stock Exchange (NYSE) had a market share of only 47% for transactions of size less than 1,000 shares; lower for orders less than 500 shares. This order flow migration has sparked heated debate among stakeholders in the trading process about both the fairness and the economic rationale for trading away from the primary exchange.¹

The question we seek to answer in this article is the following. If I, as a customer, knew a certain broker was about to execute my order on a particular regional exchange, would I have qualms about doing business with that particular broker? To answer this question we examine a sample of contemporaneous orders, quotations, and trades on the NYSE and the five regional exchanges: Boston (BSE), Chicago (CHX), Cincinnati (CSE), Philadelphia

This work was begun while the authors were visiting economist and chief economist at the Securities and Exchange Commission (SEC) in Washington, DC. Lois Lightfoot and Peter Martin of the SEC contributed substantially to earlier drafts of this work. This article has benefited from seminar participants at the University of Maryland, Arizona Finance Conference, NBER Microstructure Meetings, University of Wisconsin, George Washington University, Dartmouth College, Harvard Business School, University of Delaware, American Finance Association meetings, University of Minnesota, Ohio State University, Notre Dame University, Tulane University, Michigan State University, and from the comments of the referee, George Benston, Hank Bessembinder, David Brown, Dan Deli, and especially Larry Glosten (the editor). The views expressed herein are our own and do not reflect the views of the commissioners or any of our former colleagues at the SEC. Address correspondence to: Mark A. Peterson, Department of Finance, Rehn Hall, Southern Illinois University, Carbondale, IL 62901-4626, or e-mail: map1@cba.siu.edu.

¹ In this article we take the NYSE to be the primary market for all NYSE-listed equity securities. We do not study any American Stock Exchange equities.

(PHLX), and Pacific (PSE). Our study is unique in that the data employed admit a thorough analysis of the question at hand. For example, the data allow us to examine trading costs of marketable orders, something that otherwise can't be done without estimation error. We describe some of the difficulties of estimating trading costs absent the data on the underlying order in more detail below. In addition, our data include information on submitted limit orders, allowing us to analyze market quality for such orders.

The assessment of market quality across trading venues is important because principles of fiduciary duty and federal securities laws state that brokers have an obligation to provide "best execution" for customers' orders. In an effort to attract trading volume and work around the rules of the NYSE, a number of broker-dealers on the regional exchanges have engaged in practices such as payment for order flow and preferencing.² These practices lead to a potential conflict of interest between broker-dealers and their customers. The benefits of routing orders to regional exchanges accrue to the broker-dealer, while the cost of the potentially inferior execution is borne by the customer. These conflicts have been studied in some detail by Petersen and Fialkowski (1994), Easley, Kiefer, and O'Hara (1996), Battalio (1997), Battalio, Greene, and Jennings (1997, 1998), and Bessembinder and Kaufman (1997). We examine to what extent orders routed to various exchanges are consistent with brokers' fiduciary duties.

The BSE and CSE have formally approved preferencing plans in place in which dealers on these exchanges are allowed to direct order flow to themselves, in some cases bypassing time priority. Dealers on the other regional exchanges obtain orders by paying for them, usually at a price of two to three cents per share. Our tests are designed to compare the market quality of the primary market, that is, NYSE, to the preferencing exchanges and the nonpreferencing regional exchanges. We examine quote quality, execution costs of marketable orders, limit order fill rates, and ex post execution costs across exchanges. Thus we analyze preferencing and market quality at the exchange level and not at the broker-dealer level, though broker-dealer practices can have very similar, but not identical economic effects.

Preferencing has the potential to lead to inferior execution. If dealers make contractual or other arrangements for obtaining order flow, they may have little incentive to quote aggressively on a regional exchange, leading to inferior executions. Preferencing dealers will free ride on the quotations of nonpreferencing dealers or of other exchanges trading the same security. Alternatively, limit orders left with preferencing dealers may not fully interact with incoming orders for that exchange and may not execute when they otherwise would have, or might be subject to a heightened adverse selection problem.

² We define *preferencing* to be the dealers' practice of stepping ahead of preexisting customer orders at the same price on a particular exchange. *Payment for order flow* is the brokers' practice of routing an order to a particular venue in exchange for a cash payment. These practices are discussed in more detail below.

Throughout most of our analysis we take the routing decision as exogenous. Battalio et al. (2002) study whether it matters where brokers route customers' limit orders by examining limit order execution quality while attempting to control for market conditions and order submission strategies. The distinction between our study and theirs is that we do not address whether orders sent to one venue might receive better execution elsewhere. Specifically Battalio et al. address whether orders routed to regional exchanges might have enjoyed an even higher fill rate if they had been routed to the NYSE. Their conclusions are that the limit order routing decision may not affect retail limit order traders substantively. However, they do present evidence implying brokers can, in some instances, strategically route limit orders to improve execution quality.

Our results show that the NYSE dominates the regional exchanges for most measures of market quality, including quotation activity and effective spreads. However, among regional exchanges, the preferencing exchanges dominate the nonpreferencing regional exchanges for these same measures. Limit order executions show a somewhat different picture. Limit orders receive higher fill rates and no worse execution quality on the regional exchanges than on the NYSE. We hypothesize that this may be due to the thin limit order books away from the NYSE. There are two possible reasons for regional exchanges to have thin limit order books. First, the dealers who make markets on the regional exchanges have the option of selecting their own customers, who are brokers. Preferencing dealers may select brokers whose retail customers are not big users of limit orders. Second, as will be discussed later, the preferencing dealers may rapidly execute the limit orders as principal to prevent the orders from setting less attractive prices at which other preferencing dealers at that exchange must trade.

The article is organized as follows. Section 1 gives institutional background on the regional exchanges, preferencing, and other forms of internalization of orders. Section 2 describes our data. Section 3 presents our empirical results. A final section summarizes the results and discusses the implications of the article for regulatory policy.

1. Institutional Concerns

1.1 The National Market System and preferencing arrangements

The five regional stock exchanges form part of the National Market System, whose development was encouraged by Congress in the 1975 Securities Act Amendments. Along with the NYSE, these exchanges make continuous markets in NYSE-listed securities through an order-driven auction market. The market centers are linked together via the Intermarket Trading System (ITS), which posts quotes of various exchanges for all market centers. In addition, ITS allows one exchange to send an ITS "commitment" to another exchange

requesting an order be filled on the receiving exchange at the posted price. Through such a mechanism, customers are provided an assurance that upon execution of their order, they received the national best bid and offer (NBBO) price for their order. This execution is either provided by the exchange posting the superior price via ITS, or it may be done on the exchange on which the order was originally sent if that exchange executes the order either at or better than the NBBO. Membership in ITS is compulsory for national security exchanges. Orders on a particular exchange are executed with regard to that exchange's rules regarding price priority and precedence.

In an effort to improve market share and attract trading volume, two regional exchanges, CSE and BSE, initiated programs to obtain more order flow directed to their exchange [Battalio, Greene, and Jennings (1997)]. There are at least two reasons why such arrangements exist. In some cases, the dealer or market maker on the exchange is affiliated with the broker who routes the order to the exchange. Examples of this are Pershing's preferencing unit on the CSE, which is affiliated with Donaldson, Lufkin, and Jenrette, Inc., and Fidelity's unit on the BSE. A second reason is that the preferencing dealer may pay an unaffiliated introducing broker for orders routed to the dealer. Payments typically average two to three cents per share.

On the CSE, there are six preferencing dealers during our sample period, each of which maintains her own limit order book. An individual dealer must give precedence to her limit order book for same-price orders, but not the quotations or limit order books of other preferencing dealers. The CSE also has a central limit order book facility that takes priority over all same-price dealer quotes. However, this central order book is usually devoid of orders.

On the BSE, preferencing is implemented via the Competing Specialist Initiative (CSI) program. An issue in the CSI program will have one or more competing specialists in addition to the regular specialist. Unlike the CSE preferencing dealers, price priority and time precedence are maintained across all specialists, though the regular specialist does have some advantages. However, when the quotes of the BSE are not at the NBBO, time precedence and price priority are not in force among specialists and a broker may route an order to a particular specialist, either because of affiliation or for explicit cash payment. Thus when the BSE is quoting a wide market away from the NBBO, its specialists have the option to engage in preferencing.

1.2 A broker's "best execution" obligation

A broker who is charged with overseeing the execution of a customer's order has a fiduciary responsibility to see the customer receives favorable terms of trade. In the equity markets, this obligation is termed a duty of best execution [Macey and O'Hara (1997)]. In the recent past, this duty has been interpreted by some brokers as being satisfied if they provide their customer with an execution at the NBBO price. With more varied market venues, sophisticated

electronic routing and execution systems, and preferencing arrangements, it is less clear such a standard is appropriate.

It is possible for orders to receive prices better than the NBBO, that is, to receive “price improvement.” Price improvement arises when a market maker elects to pay more than the quoted bid or receive less than the quoted ask for a trade. This might happen if the market maker’s quote was set by a customer’s limit order and the specialist wanted to participate in the trade as principal. To do so he would be forced to offer a better price to step ahead of the customer.³ Alternatively, the specialist may cross two market orders at a spread midpoint or may cross the order with a limit order that is not displayed.

The Securities and Exchange Commission (SEC) has adopted the view that preferencing arrangements, if properly monitored by routing brokers, are not necessarily inconsistent with the broker’s best execution obligations. Best execution does not mean a customer must receive the best possible price across all trading venues on an order-by-order basis. Such a standard would likely be too difficult to implement with current technology. It does require that brokers regularly review their execution data and incorporate the results into future order routing decisions. The SEC maintains that a broker’s failure to provide best execution is in violation of not only fiduciary duty, but also the antifraud provisions of federal securities laws (Section 10(b) and Rule 10b-5), and has prosecuted cases based on such lapses by brokers. With regard to listed securities, the SEC has stated that brokers automating executions at the NBBO have failed to provide best execution because such an arrangement forgoes the opportunity for orders to interact and transact between the bid and the ask price. Moreover, in a recent private class-action case, *Newton v. Merrill*, the Appeals Court found a group of Nasdaq market makers failed to provide best execution by not looking for prices superior to the NBBO for their customers, and thus committed fraud under Rule 10b-5.⁴

Preferencing, internalization, and other order flow inducement practices provide order flow to market centers or participants that otherwise might not have sufficient volume to remain viable. Thus these practices are potentially beneficial to the extent that they may foster competition. In this sense these practices replace competition between dealers on a given exchange for a particular order with competition between competing market centers for large blocks of order flow.

2. Data Description

2.1 General

Order data were obtained from each of the five regional exchanges for the four weeks of October 28 to November 22, 1996, and from the NYSE for the

³ Ross, Shapiro, and Smith (1996) analyze price improvement on the NYSE in detail.

⁴ *Newton vs. Merrill Lynch*, no. 96-5045, Third Circuit Court of Appeals, filed January 30, 1998.

week of October 28 to November 1, 1996.⁵ The analysis is conducted only for NYSE-listed issues. Shares that are listed on the American Exchange or exclusively on a regional exchange are not included in this study. Order data include information such as the order arrival time, the size of the order, a buy/sell indicator, an indicator or means to identify market and limit orders, and limit order prices. This order information allows for precise estimates of trading costs, something not possible using only trade and quote data. The buy/sell information obviates reliance on some form of tick test [e.g., Lee and Ready (1991)] that can be noisy, especially when inferring trade direction in minimum variation markets. In addition, the order arrival times allow us to establish a clear benchmark price (NBBO) to compute a cost for market and marketable limit orders. Without our unique dataset, the combination of estimating trade direction and estimating the order arrival time, and hence the benchmark quote, lead to considerable biases in estimating trading costs.

The limit order analysis is based on day limit orders except for the CHX, whose data did not distinguish day limit orders from other limit orders. However, in evaluating limit order performance statistics such as ex post transactions costs, only executed limit orders were included in the analysis. Day limit orders (those expiring at the end of the trading day if not filled) account for the majority of limit orders submitted.

Orders and executions taking place between 9:30 A.M. and 4:00 P.M. EST are included in the analysis. Opening orders, tick sensitive orders, and market orders with price qualifiers are also excluded. As a result of these screens, the study only considers three types of orders: regular-way market orders, marketable limit orders (limit orders whose limit prices make them immediately executable, such as the limit order to buy at a price equal to the current offer), and nonmarketable limit orders. In the tables that follow, results are partitioned based on order type and, for the limit orders, the limit price relative to the NBBO.

There remains a problem of deciding exactly how to select a sample. Each exchange trades a different set of securities endogenously selected by exchange members. For example, the range of securities traded runs from 338 on the CSE to more than eight times as many names on the NYSE and PHLX. In order to maintain comparability, and to guard against sample selection and endogeneity effects, we restrict our analysis to a set of liquid securities trading on all exchanges. This sample is constructed by taking all regular common stocks that trade on the CSE and have at least 10 trades over the sample period. Though this diminishes our overall sample size, it reduces concerns regarding endogeneity. The sample selection process results in a total of 334 securities. Below we provide additional detail on the databases provided by each exchange.

⁵ Only one week of NYSE data was used because of the large number of orders on the NYSE relative to the other exchanges and because of limitations in data processing.

2.2 NYSE order data

The NYSE SOD file contains all orders entered via the SuperDot system, both market and limit orders, and is distinct from the NYSE Consolidated Audit Trail data. SOD does not contain orders entered into the auction from the floor or by other means. According to the NYSE *Fact Book* for the year 1995, 85% of all orders and 33% of volume went through the SuperDot system. Because this study primarily focuses on retail orders, this limitation should not be severe, as relatively few retail-size orders are entered from the floor.

2.3 BSE order data

BSE data are taken from the BEACON system. This system records the entry and execution of market orders. BEACON does not record information about the entry and disposition of unexecuted limit orders, nor does it have detailed quote records of the interaction between competing specialists in the BSE's CSI program. To obtain data on unexecuted limit orders, BSE hard-copy records of unexecuted day limit orders were entered by hand for one week of the period under study. These hand-entered data were merged with the BSE electronic data.

2.4 CHX order data

The CHX data file combines orders and executions in a single observation. Orders can be entered either electronically or from the floor. Because floor orders are not entered electronically, the order information, such as order arrival time, is unavailable. Order information is incomplete for ITS trades. Therefore floor orders and ITS orders are excluded.

2.5 CSE order data

The CSE is a multiple market-maker exchange. The data in this study are taken from six of the seven CSE preferencing dealers: Prudential, Olde, Pershing, Fidelity, Piper, and Redwood. Data for the seventh preferencing dealer, Dain Bosworth, were unavailable. However, Dain Bosworth accounts for a relatively small proportion of CSE activity. The CSE's preferencing dealers pair orders, subject to exchange rules, at their trading desks and send the paired trades to the CSE's facilities in Chicago for execution. The CSE data include order and trade data for those trades executing on the CSE. Unexecuted limit order data are also available from the CSE.

2.6 PSE order data

The PSE data file contains both order and execution information for orders received on both of the PSE's trading floors. Trades used in our analysis consist primarily of trades processed through the PSE's P\COAST system. Records for floor trades and manually reported trades do not contain an order entry time and were excluded from the analysis. ITS orders sent to

other exchanges from the PSE and ITS orders received on the PSE are both included in the file, but there is no indication on which exchange the ITS trade was printed. Therefore ITS trades were also excluded from the analysis.

2.7 PHLX order data

The PHLX's market surveillance department retains separate trade and order files. The order file consists of all electronically placed orders. This file does not include orders phoned in by brokers. On March 3, 1997, 95% of the orders on PHLX were entered electronically, representing approximately 57% of the volume. The trade file includes all PHLX prints, including ITS trades.

3. Empirical Results

3.1 Sample characteristics

Table 1 describes the characteristics of the entire sample in terms of firm size, trading activity, and order size. It illustrates the differing business models of the NYSE and the five regional exchanges, reflected in the selection of the stocks traded.⁶ The first row of the table reports the number of different equities in the datasets for each exchange. The NYSE, which generates revenue from listing issuers, trades 2,256 different firms in our sample, the most of the six exchanges. The CSE trades only 338 different names, the smallest number of securities among all of the exchanges. For all the regional exchanges, the choice of what shares to trade is up to the exchange and its members. Issuers whose stock is traded on a regional exchange generally do not pay any fees to the regional exchanges for having their shares traded there. Thus the choice of what stocks to trade is highly dependent on the profitability of trading. Of interest is that the two preferencing exchanges, CSE and BSE-CSI (stocks on the BSE with dealers participating in the CSI), chose to trade the smallest number of securities.

The next two blocks of rows in the table, describing market capitalization and daily trading volume, respectively, indicate CSE and BSE-CSI have elected to trade the largest and most active of the names. The NYSE and the nonpreferencing regionals (BSE-non-CSI, CHX, PHLX, and PSE) trade shares with median market capitalization ranging from \$600 million to more than \$1 billion. BSE-CSI and CSE trade stocks with a median market capitalization of approximately \$7 billion. The same pattern is seen in trading volume, where the median trading volume of a preferencing exchange is several times greater than the NYSE and nonpreferencing regionals. The

⁶ The nature of the decision of what stocks to trade differs between the NYSE and the regional exchanges. To trade on the NYSE the stock must generally be listed there, which involves an important certification and auditing role by the NYSE. The regional exchanges do not list many stocks; instead they trade stocks listed on the NYSE pursuant to an unlisted trading privileges plan that allows them to trade shares listed on other exchanges.

Table 1
Stock characteristics by exchange

	Preferencing exchanges			Nonpreferencing regional exchanges			
	NYSE	CSE	BSE CSI	BSE-non-CSI	CHX	PSE	PHLX
Number of stocks	2,256	338	79	1,305	1,899	1,611	1,816
Market capitalization (\$ million)							
25 th percentile	194	2,146	3,292	401	280	359	277
50 th percentile	591	6,265	7,328	1,132	817	1,045	830
75 th percentile	2,079	17,160	28,060	3,488	2,631	3,363	2,772
Daily volume/stock (000s shares)							
25 th percentile	19	273	437	41	28	36	27
50 th percentile	46	431	622	94	66	83	67
75 th percentile	143	706	1,069	211	174	210	186
Percent coverage (for the 334 liquid stocks only)							
Number of trades	55.6%	58.5%	62.5%	59.5%	60.8%	65.1%	71.4%
Trading volume	24.2%	43.4%	44.5%	37.8%	21.8%	47.4%	41.4%

Table reports the quartile points of the distributions of equity market capitalization and trading volume for NYSE-listed stocks with trades in electronic databases on the New York Stock Exchange (NYSE), Boston Stock Exchange (BSE), Chicago Stock Exchange (CHX), Cincinnati Stock Exchange (CSE), Philadelphia Stock Exchange (PHLX), and Pacific Stock Exchange (PSE). BSE CSI refers to stocks trading on Boston and with dealers participating in the competing specialists initiative. BSE-non-CSI refers to stocks trading on BSE with no dealers participating in the competing specialists initiative. Market capitalization is measured as the product of the number of shares outstanding on October 28, 1996 and the share price. Daily trading volume is measured over the previous year. Percent coverage refers to the proportion of trading using market and marketable limit orders represented in the sample dataset to the total trading in the TAQ dataset for the week of October 28 to November 1, 1996. The row beginning with "Number of trades" is the ratio of orders in the sample dataset to the number of trades in the TAQ dataset. The row beginning with "Trading volume" is the ratio of order volume in the sample dataset to the volume of trades in the TAQ dataset.

preferencing exchanges have elected to take orders in large-capitalization, high-volume liquid securities, as these shares are among the most popular for retail investors whose orders may be routed pursuant to preferencing arrangements.

The next block of rows estimates the overlap between our dataset and the NYSE's TAQ dataset between October 28, 1996, to November 1, 1996. In this comparison we include only marketable orders. In terms of the number of trades, our data include more than 50% of the trades on each exchange. However, because the system orders are mostly retail, the percentage of trading volume is less than 50%.

Table 2 presents a summary of the order data used in the tables that follow for each of the exchanges. Because of the differences in the characteristics of stocks traded on each of the exchanges, as shown in Table 1, we restrict our analysis to only the most active and liquid securities. We arrive at this set by finding those stocks trading in common with the CSE and NYSE, and eliminate any of those having fewer than 10 trades over the sample period. This results in a set of 334 equity securities trading on the NYSE and all of the regional exchanges. Consideration of only such securities should ameliorate concerns of cross-sectional sample selection resulting from the different business strategies of the six exchanges.⁷

⁷ We thank the referee and the editor for making this important suggestion.

Table 2
Distribution of order type and order size

Order size	Order type	NYSE	Preferring exchanges		Nonpreferring regional exchanges			
			CSE	BSE CSI	BSE-non-CSI	CHX	PSE	PHLX
Small	Market	99,154	117,201	22,518	17,104	90,546	115,066	62,726
	Marketable limit	19,493	10,925	1,304	1,146	9,675	10,486	5,340
	Other limit	74,577	10,811	2,540	1,633	10,136	10,849	6,317
	Total	193,224	138,937	26,362	19,883	110,357	136,401	74,383
Medium	Market	22,419	14,770	2,719	2,550	8,969	14,955	6,101
	Marketable limit	14,050	3,393	328	368	2,298	2,891	1,189
	Other limit	45,115	4,489	779	526	3,048	4,009	1,573
	Total	81,584	22,652	3,826	3,444	14,315	21,855	8,863
Large	Market	21,419	6,504	1,640	2,049	3,494	5,622	2,203
	Marketable limit	27,963	2,376	252	370	1,196	1,948	706
	Other limit	71,715	2,902	550	497	1,300	2,283	759
	Total	121,097	11,782	2,442	2,916	5,990	9,853	3,668
All	Market	142,992	138,475	26,877	21,703	103,009	135,643	71,030
	Marketable limit	61,506	16,694	1,884	1,884	13,169	15,325	7,235
	Other limit	191,407	18,202	3,869	2,656	14,484	17,141	8,649
Total		395,905	173,371	32,630	26,243	130,662	168,109	86,914
Average order size								
	Market	963	391	429	520	308	362	318
	Marketable limit	2,343	818	734	869	572	716	581
	Other limit	1,961	873	830	1,081	611	729	581

Table reports the number of orders recorded in electronic databases on the New York Stock Exchange (NYSE), Boston Stock Exchange (BSE), Chicago Stock Exchange (CHX), Cincinnati Stock Exchange (CSE), Philadelphia Stock Exchange (PHLX), and the Pacific Stock Exchange (PSE) for 334 stocks that trade on the CSE and NYSE. The table excludes stocks with less than 10 orders on the CSE, tick sensitive orders, and cases where the NBBO does not exist, for example, the preopen. The study period includes one week of data (October 28 to November 1, 1996) for the NYSE and four weeks of data (October 28 to November 22, 1996) for the BSE, CHX, CSE, PHLX, and PSE. Small orders are for 100 to 500 shares. Medium orders are for 501 to 1,000 shares. Large orders are for more than 1,000 shares. Marketable limit orders are limit orders with the limit price greater (less) than or equal to the offer (bid) price at the time the order arrives at the exchange for buy (sell) orders.

We classify orders based on their size into small (100–500 shares), medium (501–1,000 shares), and large (more than 1,000 shares) orders. In addition, we break orders down into three types: market orders, marketable limit orders, and other limit orders. “Market” orders are unpriced orders to buy or sell shares immediately at the best price available in the market when the order arrives. A “marketable limit” order is a priced order to buy (sell) stock where the limit price is greater (less) than or equal to the offer (bid) price at the time the order arrives at the exchange.⁸ Such a limit order is immediately

⁸ Throughout the text, unless otherwise noted, “bid” and “ask” refer to the national best bid and ask, respectively.

executable at the prevailing quotes or better. Limit orders not immediately executable are grouped together in this table as "Other Limit."

The sample consists of approximately one million orders, of which 395,905, or 40%, were sent to the NYSE. Recall the table reports only one week of orders for the NYSE and four weeks of data for the regional exchanges, so the true NYSE market share of all order flow is in fact higher than this. The CSE and PSE receive the most orders among the regional exchanges per unit time, followed by the CHX, PHLX, and BSE. The table shows the relative order mix of market, marketable limit, and other limit orders is roughly constant across the regional exchanges. However, it clearly illustrates the point that the regional exchanges receive a disproportionate share of order flow in the form of small market orders relative to the NYSE. For example, small orders of all types account for 49% (193,224/395,905) of all NYSE orders, but this ratio ranges from 76% to 86% for the regional exchanges. This difference is even more pronounced for small market orders, which account for 25% of NYSE orders, but between 65% and 72% of regional orders.

The difference in the relative size and order type mix reflects the specialization of the regional exchanges into the business of executing retail customer orders. The regional exchanges are used primarily for the execution of smaller orders. This can be seen in the bottom portion of Table 2. The average size of system market orders on the NYSE is 963 shares, whereas for the regional exchanges, the average market order is only about one-third as large. Table 2 shows that this ratio is similar for limit orders as well. The regional exchanges are also used to a certain extent by the upstairs market to execute large-block cross trades that circumvent the priority of other orders on the NYSE floor.⁹

Looking at the mix of orders, it is also clear, even after controlling for order size, that the regional exchanges receive proportionally far more market orders than the primary market, that is, the NYSE. The top portion of Table 2 shows that small market orders comprise 51% of the NYSE's small orders, whereas this percentage is more than half again as large for the regional exchanges. The effect is even stronger as order size increases. For example, 17.7% (21,419/121,097) of the NYSE's large orders are market orders; this percentage is between three and four times greater for the regional exchanges.

Finally, it is interesting to note the relative use of market versus marketable limit orders across exchanges. Recall that marketable limit orders are limit orders priced so that they are immediately executable when they arrive at the exchange. Such orders are useful to an investor who may wish to constrain the actions of the specialist to prevent him from moving quotes disadvantageously after the order arrives at the specialist post, but before it is executed, thus giving a more costly execution to the order. The conclusions from Table 2 are that the regional exchanges receive a higher fraction

⁹ See PHLX Rule 126 and PSE Rule 5.14(b).

of small orders and a higher fraction of market orders (versus limit orders) than the NYSE. In addition, as order size increases, all exchanges, but the NYSE in particular, receive a greater fraction of their order flow in the form of limit orders.

3.2 The quality of market quotations

Before considering trading activity, it is instructive to look at quotations by individual exchanges. There are two basic measures to consider when evaluating the quality of market quotations: the quoted price and the depth, or number of shares, for which the dealer is willing to trade at the quoted price. Other things being equal, a market with a small difference, or spread, between the best price bid and the best price offered is generally more liquid than a market with a larger spread. Similarly, other things being equal, a market whose quotes are for a larger number of shares is generally more liquid and deeper than an identical quote for a smaller number of shares.

Table 3 compares the cross-sectional average time-weighted bid-ask spreads and associated time-weighted quotation depths ($1/2 \times (\text{ask depth} + \text{bid depth})$) for the NYSE and each of the five regional exchanges. This comparison is done only for the 334 liquid NYSE/CSE traded stocks. The NYSE far and away quotes the tightest market in terms of spread. The average bid-ask spread quoted by the NYSE is 15.5 cents, narrower than the next closest exchange, the CSE, by 11.1 cents. It is notable that the CSE has the narrowest average spread of any of the regional exchanges. Given that it is a

Table 3
Quoted bid-ask spreads and depths

	Preferring exchanges			Nonpreferring regional exchanges			
	NYSE	CSE	BSE CSI	BSE- non-CSI	CHX	PSE	PHLX
Time-weighted spread (cents/share)							
Mean	15.5	26.6	39.4	40.4	36.8	37.7	51.1
Median	15.1	24.2	39.9	40.5	37.9	35.8	43.8
Time-weighted depth (shares)							
Mean	18,529	767	316	231	484	734	268
Median	13,233	761	112	102	297	453	102
Percent of time that quotes are equal to							
NBBO bid and NBBO ask	89.1%	29.5%	1.9%	0.7%	17.3%	13.0%	3.1%
NBBO bid or NBBO ask	10.8%	57.8%	9.8%	3.8%	33.7%	42.8%	11.2%
Neither of the above	0.1%	12.7%	88.3%	95.5%	49.0%	44.2%	85.7%
Average depth (shares) when quotes are equal to							
NBBO bid and NBBO ask	18,568	649	1,867	1,467	672	956	1,131
NBBO bid or NBBO ask	19,282	783	1,468	1,030	811	948	1,388
Neither of the above	16,355	864	196	115	346	444	141

The table reports the cross-sectional average time-weighted quoted spread and depth ($0.5 \times (\text{ask depth} + \text{bid depth})$) for stocks on the New York Stock Exchange (NYSE), Boston Stock Exchange (BSE), Chicago Stock Exchange (CHX), Cincinnati Stock Exchange (CSE), Philadelphia Stock Exchange (PHLX), and the Pacific Stock Exchange (PSE) for 334 stocks that trade on the CSE and NYSE. The table is estimated using one week of data (October 28 to November 1, 1996) for the NYSE and four weeks of data (October 28 to November 22, 1996) for the BSE, CHX, CSE, PHLX, and PSE. "NBBO" is the national best bid and offer.

preferencing exchange, one might expect its quotations to be wider. Specialists on that exchange need not depend on their quotation to attract order flow, a topic we will return to at the end of the article. The other regional exchanges post bid-ask spreads in the range of 37 to 51 cents, about three times as wide as the NYSE.

The next lines of Table 3 show the NYSE's average depth is about 20 to 80 times greater than the average depth of the regional exchanges. The ratio is similar whether the comparison is based on means or medians. The average NYSE quoted depth is almost 20,000 shares, whereas for each regional exchange the average depth is less than 1,000 shares. The deeper NYSE quotes are most likely due to the large number of limit orders received by the primary market. The CSE, a preferencing exchange, has the greatest average quoted depth of all the regional exchanges. The low size of the median regional quotes is likely due to a procedure known as "autoquoting," in which an exchange automatically programs its electronic quotation system to quote a market of 100 shares one tick outside the NBBO on each side of the market.¹⁰ This is confirmed by noting that the average spread on most of the regionals is about two ticks, or 25 cents, greater than the average spread of 15 cents quoted on the NYSE.

It is possible that the comparisons above understate the quote quality of the regional exchanges. In particular, because the regional exchanges have much less volume than the NYSE, it may not be profitable for their dealers to quote actively all of the time, or to quote competitively on both sides of the market. If the regionals have fewer limit orders than the NYSE, this could also contribute to the wider quotes. However, a regional exchange may have a competitive quote on one side of the market even though its bid-ask spread is large. In addition, though its average quoted depth may be low, a regional exchange may provide considerable liquidity to the market by quoting a greater depth when it has the best quotation price than it does when its quoted price is less competitive. Alternatively, a regional exchange specialist or dealer may try to control risk by decreasing its quoted depth when it narrows its spread.

The bottom half of Table 3 explores the possibilities noted above. We compute the fraction of the time the regional exchange is at the national best bid and offer on both sides of the quotation, on one side of the NBBO (but not both), and on neither side of the NBBO. This measure of quote quality varies considerably across exchanges. The BSE non-CSI and PHLX were away from the NBBO on both sides of the market 95.5% and 85.7% of the time. In comparison, the CHX and the PSE quoted more aggressively and were away from both sides of the NBBO only about 49% and 44% of the time. The CSE has by far the best quotation performance among the regional exchanges based on the data in Table 3. The CSE is at one or both sides of

¹⁰ Exchanges are, by rule, only allowed to autoquote for a single round lot.

the NBBO over 85% of the time. These results suggest that the CSE provides competitive price quotes a substantial portion of the time.¹¹

To explore quote quality a bit further, the bottom-most portion of Table 3 examines quoted depth when an exchange is on both sides of the NBBO, when it is on only one side of the NBBO, and when it is on neither side of the NBBO.¹² With the exception of the CSE, quotation depth on the regional exchanges is larger when its price quote is at the NBBO than when it is not. This effect is more pronounced for the BSE and PHLX. The PHLX's average quoted depth is 1,131 shares when its quoted prices are at the NBBO on both sides of the market, and 141 when it is on neither side of the NBBO, which is 85.7% of the time. The NYSE also exhibits slightly larger average quoted depth when it is on one or both sides of the NBBO than it does in the infrequent instances when it is not at either side of the NBBO.

In summary, the quotation quality of the regional exchanges, though showing considerable variation, is generally lower than that of the NYSE. On average, the NYSE quotes tighter spreads and deeper markets than the regionals. In addition, two of the regional exchanges, BSE (for CSI and non-CSI stocks) and PHLX, are seldom on one or both sides of the NBBO. While the PSE and CHX are on at least one side of the NBBO about half of the time, only the CSE is on at least one side of the NBBO more than 85% of the time. On the other hand, the regional exchanges do, with the exception of the CSE, provide more quoted depth when they are at the NBBO than when they are away from the market. The CSE's quoted depth when it is at the NBBO is comparable to the depth when it is not at the NBBO.

3.3 Market order execution costs

In this section we analyze the execution quality of marketable orders (i.e., market and marketable limit orders). Execution quality is measured by calculating the effective spread. For a customer buy order, the effective spread is calculated by doubling the difference between the trade price and the midpoint of the bid-ask spread (NBBO) measured at order arrival time. Thus the effective spread, ES , for a marketable buy order can be calculated as follows:

$$ES_{buy} = 2 \times [trade\ price - 0.5 \times (bid\ price + ask\ price)]. \quad (1)$$

For a customer-marketable sell order, the effective spread is calculated by doubling the difference between the midpoint of the bid-ask spread (NBBO)

¹¹ Table 3 reports the composite CSE quote, including both preferencing and nonpreferencing dealers. When looked at separately, preferencing dealers are found to be on both sides of the NBBO 19.6% of the time, on only one side of the NBBO 62.8% of the time, and on neither side 17.6% of the time. The sample used for this calculation differs slightly from the one used in Table 3.

¹² The cross-sectional time-weighted depth may not equal the linear combination of the percent of time and the average depth because it is sometimes the case that a stock does not fall into all three categories on a particular exchange.

measured at order arrival time and the trade price:

$$ES_{sell} = 2 \times [0.5 \times (bid\ price + ask\ price) - trade\ price]. \quad (2)$$

If all trades were executed on the opposite-side quote, then the effective spread would precisely equal the quoted spread. In general, some fraction of trades occurs at prices better than the quotes. The effective spread captures this effect.

In a minimum variation 1/8-point market, effective spreads may be less than the minimum variation of 12.5 cents. This arises because some marketable orders receive price improvement. Some marketable orders to buy are executed at the bid and some marketable orders to sell are executed at the ask, which causes the average effective spread to fall below 12.5 cents. Our data allow us to detect an effective spread of less than a tick because we have order data, not only trade and quote data. Order data allow us to assign the correct buy/sell indicator to each order. Absent such information, an observer looking at data from sources such as TAQ would have no way to distinguish whether an order traded at the bid price was a market sell order or a price-improved buy order. This is an extremely important point. Roughly 85% of the trading activity of these liquid stocks takes place in 1/8-point markets. If the buy/sell assignment cannot be made, then there is virtually no hope of distinguishing market quality among the six exchanges for the majority of trading situations. In working with trade and not order data, Lee (1993:1027) acknowledges that “much of the order flow targeted for purchase is executed when the spread is one eighth” but goes on to argue that “these trades provide little opportunity for price improvement.” In our data, which are from a period almost a decade after the Lee sample, a higher fraction of trades occur in 1/8-point markets and we find price improvement is in fact a significant factor. Battalio, Greene, and Jennings (1997) work around the problem by only considering trades in 1/4-point markets, which permits price improvement to occur at the spread midpoint, making the need to assign the trade direction much less important.

We calculate, for each stock, its average effective spread (along the dimensions of NBBO spread at order submission), order type, order size, and trading venue. Our tests are designed to evaluate the execution quality between the primary market (NYSE), the preferencing exchanges, and the nonpreferencing regional exchanges. Therefore we pool orders from the CSE and the BSE-CSI into the preferencing group and orders sent to the BSE-non-CSI, CHX, PSE, and PHLX into the nonpreferencing regional group. Paired *t*-tests are used to statistically evaluate the differences in means. Note that despite pooling, some stocks do not have representation in all of the categories (i.e., spread, order type, and order size).

Table 4 reports the cross-sectional average effective spread for marketable orders of various sizes in 1/8- and 1/4-point markets. In 1/8-point markets, we

Table 4
Effective spreads (in cents) in 1/8 and 1/4-point markets

Order size	Order type	PREF	NYSE	REGL	NYSE	PREF	REGL
1/8-point market (NBBO)							
Small	Market	11.1	9.6*	11.6	9.6*	11.1	11.6*
	M. limit	11.6	11.3*	12.0	11.4*	11.6	12.0*
Medium	Market	11.7	10.8*	12.1	10.8*	11.7	12.1*
	M. limit	12.1	11.7*	12.2	11.7*	12.1	12.1
Large	Market	12.2	12.1	12.8	12.2*	12.2	12.8*
	M. limit	12.0	12.0	12.2	12.1*	12.1	12.2
All	Market	11.2	10.2*	11.7	10.2*	11.2	11.7*
	M. limit	11.7	11.8	12.1	11.8*	11.7	12.1*
1/4-point market (NBBO)							
Small	Market	9.7	6.9*	14.8	7.0*	9.7	14.7*
	M. limit	11.8	10.0	21.2	10.6*	11.0	20.8*
Medium	Market	11.9	9.5*	16.8	9.7*	11.9	16.3*
	M. limit	14.7	15.7	20.7	14.7*	13.6	19.4*
Large	Market	13.8	13.5	18.7	13.3*	13.9	18.4*
	M. limit	17.0	17.0	20.9	18.1*	17.3	20.5
All	Market	10.4	8.2*	15.1	8.2*	10.4	15.0*
	M. limit	12.8	14.4*	21.2	15.0*	12.4	20.9*

The table reports the cross-sectional average effective spread (in cents) for the sample and time period described in Table 2 for market and marketable limit (M. limit) orders. Cell values in adjacent columns represent the average for stocks present in both columns. For buy (sell) orders, the effective spread is 2 (–2) times the difference between the volume-weighted average execution price and the quote midpoint. Small orders are for 100 to 500 shares. Medium orders are for 501 to 1,000 shares. Large orders are for more than 1,000 shares. An “*” indicates the effective spread in one category is significantly different from the effective spread in the other category at the 5% level using a paired *t*-test. “PREF” indicates preferencing regional exchanges. “REGL” indicates nonpreferencing regional exchanges.

find that the effective spread for market orders ranges from about 10.2 cents per share on the NYSE to 11.7 cents per share for orders routed to the nonpreferencing regionals. Such orders are presumably, mostly retail customer trades. This pattern is somewhat surprising in light of predictions of adverse selection models in the economic literature on market making [Admati and Pfleiderer (1988), Easley and O’Hara (1987)]. Such models predict orders with low information content, such as those of individual investors, should execute at more favorable prices than those of informed trades. Paired *t*-tests indicate both small and medium-size marketable orders on the regional exchanges trade at higher effective spreads than similar-size orders on the NYSE. For large marketable orders, the NYSE has significantly lower costs than the nonpreferencing regionals, but has costs comparable to those on the preferencing exchanges.

The effective spread for orders executed on the NYSE increases with order size. This is consistent with the predictions of traditional adverse selection models. Although the effective spread for orders executed on the regional exchanges increases with size, the increase is less sharp than for the NYSE. This may be due to the routing practices of the brokers who send orders to the regional exchanges. If a high enough portion of the order flow comes from either preferencing or purchasing arrangements, then order size may not pose an adverse selection problem to such dealers and specialists because the preferred order flow is not likely to be obtained from informed traders.

For market orders, the nonpreferencing regional exchanges economically and statistically underperform the preferencing exchanges, and especially the NYSE, in 1/8-point markets. In this regard, for example, the CHX's price improvement rate is adversely affected by the fact that orders in minimum variation markets are not eligible for its SuperMax and Enhanced SuperMax automated price improvement programs. Although not reported in a table, less than 4% of all market orders are price improved on the CHX in 1/8-point markets.

It is interesting that any price improvement occurs on preferencing exchanges, because the purpose of preferencing must be to obtain the right to transact orders at the quotes. Further examination of the subsequent price movement of price-improved orders compared to non-price-improved orders, indicates that price-improved orders on the CSE tend to have the price move against them. For example, a buy order receiving price improvement tends to have a midpoint that decreases following the trade. The average amount of the percentage decrease is about 20 basis points. This is in contrast to a 2 basis point average percentage increase in non-price-improved buy orders. This observation is consistent with dealers speculating when offering price improvement.

The last two columns of Table 4 compare the average effective spreads on the preferencing and nonpreferencing regional exchanges. We can reject equality of effective spreads for all sizes of market orders at the 5% level. Thus the preferencing regional exchanges execute orders at effective spreads that are less than the spreads of their nonpreferencing regional counterparts. Given the decreased competition at the point of trade on a preferencing exchange, this may seem surprising. However, as discussed above, preferencing is just one of several mechanisms whereby broker-dealers are able to capture rents from informationless order flow. Payment for order flow and reciprocal arrangements are alternative arrangements that can achieve similar economic ends. To the extent that the nonpreferencing regional exchanges engage in these activities, their trading costs may be higher than what would prevail in a pure auction market.

The lower panel of Table 4, which reports the results for trades in 1/4-point markets, shows more variation in the estimates of the average effective spread. The cross-sectional average effective spread on the NYSE increases with order size: small market orders have a cross-sectional average effective spread of 7 cents per share and large market orders execute at a cross-sectional average effective spread of 13 cents per share. The average effective spread of the preferencing exchanges, while significantly greater than that for the NYSE in 1/4-point markets for small and medium-size market orders, is nonetheless significantly lower than that of the other regional exchanges. The nonpreferencing regional exchanges show markedly higher effective spreads than do the NYSE and the preferencing exchanges. In fact, the effective

spread of marketable orders on the nonpreferencing regional exchanges easily exceeds that of the preferencing regionals for every size and order type, with the exception of large marketable limit orders.

The effective spreads in Table 4 reveal that marketable limit orders generally receive less price improvement than market orders. These data are consistent with the findings of Harris and Hasbrouck (1996) and SEC (1997). Peterson and Sirri (2002) attribute the difference in costs of market and marketable limit orders to a selection bias. Specifically they find marketable limit orders are used more often for larger orders, especially when the order size is for more than the posted depth, and for orders that are more likely to execute at the next price step. For these reasons the average marketable limit order may have higher effective spreads than market orders.

A comparison of the two panels in Table 4 highlights a curious anomaly pertaining to the size of discrete spreads on U.S. exchanges. In the top panel, which analyzes markets with a quoted spread of 1/8-point, the NYSE's effective spread across all market orders is 10.2 cents. However, in the bottom panel, which analyzes wider 1/4-point markets, the NYSE's effective spread is a narrower 8.2 cents. Intuition suggests effective spreads should increase when quoted spreads increase. A likely explanation for the counterintuitive result is the presence of a midpoint at which to trade in a 1/4-point market that serves to narrow the effective spread. In a minimum variation 1/8-point market, dealers may have difficulty providing price improvement (and narrowing the effective spread) because they are constrained by exchange rules concerning time priority and the priority of the limit order book over a specialist trading for its own account (see, e.g., NYSE Rule 92(b)). Thus customers may receive better trade prices, on average, in markets with a quoted spread of 1/4-point than in a 1/8-point market, but only because the 1/4-point market has a feasible midpoint for a trade price and the 1/8-point market does not. The lower effective spread in 1/4-point markets is not caused by the wider spread per se, but is a consequence of a spread midpoint on which to trade. A smaller minimum tick size would create a spread midpoint for the current minimum variation 1/8-point markets. The results of Table 4 suggest that the presence of such a spread midpoint would lower effective spreads, and thus trading costs.

3.4 Effective spread analysis using ordered probit

The results in Table 4 hold NBBO spread, order size, and order type constant. However, there may be other factors beyond these variables that determine trading costs. To address this issue we estimate an ordered probit regression with explanatory variables to control for time of day, order size, quoted depth, stock liquidity, and trading venue. The independent variables include the following: MIDDLE is a dummy variable set to one if the order arrives at an exchange between 11:00 A.M. and 2:30 P.M.; LATE is a dummy variable set to one if the order arrives at an exchange after 2:30 P.M.; LogSize is

the log of the order size; LogSameSideDepth is the log of the NBBO depth on the same side and LogOppSideDepth is the log of the NBBO depth on the opposite side;¹³ LogVolume is the log of the trading volume from October 1, 1995, to October 1, 1996; PREFERENCING is a dummy variable set to one if the order is sent to a preferencing regional exchange, zero otherwise; REGIONAL is a dummy variable set to one if the order is sent to a nonpreferencing regional exchange, zero otherwise.

We use three levels in the regression for orders submitted in 1/8-point markets: $-\$1/8$, $\$1/8$, and $\$3/8$. An effective spread of $-\$1/8$ occurs when an order receives price improvement, an effective spread of $\$1/8$ occurs when an order executes at the quote, and an effective spread of $\$3/8$ occurs when more of the order trades at prices worse than the NBBO. This may occur for an order for more shares than indicated at the posted depth. For example, suppose a stock is quoted $\$20-\$20\frac{1}{8}$ and a buy order executes at $\$20\frac{1}{4}$, the effective spread is $(2)(\$20\frac{1}{4}-\$20\frac{1}{8})$ or $\$3/8$. In 1/4-point markets we use three levels for the regression: $-\$1/4$, $\$0$, and $\$1/4$.

Table 5 reports the results of the ordered probit. Note that the marginal effects of the regressors on the probabilities are not equal to the coefficients. Therefore, to provide further insight from the model we estimate the effective spread under the following conditions for a 500 share order: the trade occurs in the middle of the day ($MIDDLE = 1$, $LATE = 0$) and the depth is symmetric ($LogSameSideDepth$ and $LogOppSideDepth = 7.824$ (2,500 shares)) for a stock with annual trading volume of 125,000,000 shares in the previous year.¹⁴ The model indicates that such a market order in a 1/8-point market has an effective spread of 10.9 cents, 11.7 cents, and 12.2 cents on the NYSE, preferencing regional exchange, and nonpreferencing regional exchange, respectively. These values are consistent with the values reported in Table 4 and provide additional confidence in our results in Table 4. In 1/4-point markets, the effective spreads are 9.4 cents, 11.9 cents, and 17.3 cents for market orders submitted to the NYSE, preferencing regionals, and nonpreferencing regionals, respectively. Finally, to enable a clean test of whether preferencing exchanges have significantly different effective spreads than nonpreferencing exchanges, we estimate the model with the base case being an order sent to a nonpreferencing regional exchange and include an indicator variable for NYSE orders. The coefficient estimates (not reported in a table) indicate that orders sent to the preferencing exchanges have significantly smaller effective spreads than those same orders sent to the nonpreferencing regional exchanges.

¹³ Here we define the NBBO (consolidated) depth as do Bacidore, Ross, Sofianos (1999), namely, as the highest depth across markets at the best price.

¹⁴ This figure corresponds to roughly 500,000 shares per day, representing a typical stock on the preferencing exchanges (see Table 1).

Table 5
Effective spread analysis: ordered probit regression

	1/8-point market (NBBO)		1/4-point market (NBBO)	
	Market	Marketable limit	Market	Marketable limit
Intercept1	-1.928*	-2.568*	-2.669*	-3.127*
Intercept2	2.321*	3.796*	-0.437*	-1.328*
MIDDLE	-0.003	-0.067*	-0.046*	-0.065
LATE	-0.050*	-0.035	-0.066*	-0.091
LogSize	-0.056*	-0.115*	0.052*	-0.115*
LogSameSideDepth	-0.311*	-0.362*	-0.043*	-0.064*
LogOppSideDepth	0.326*	0.265*	0.132*	0.203*
LogVolume	0.035*	0.132*	-0.015*	0.031
PREFERENCING	-0.291*	-0.288*	-0.230*	0.015
REGIONAL	-0.508*	-0.770*	-0.766*	-0.924*
Pseudo- R^2	0.093	0.044	0.086	0.109
Estimated spreads under specified conditions (in cents)				
NYSE	10.9	11.3	9.4	16.4
Preferencing	11.7	11.8	11.9	16.2
Regional	12.2	12.3	17.3	22.8

The table reports the parameter estimates from a regression of effective spread on several explanatory variables for the sample of orders described in Table 2. MIDDLE is a dummy variable equal to one if the order is placed between 11:00 and 2:30 p.m., zero otherwise. LATE is a dummy variable equal to one if the order is placed after 2:30 p.m., zero otherwise. LogSize is the log of the order size. LogSameSideDepth is the log of the NBBO depth on the same side. LogOppSideDepth is the log of the NBBO depth on the opposite side. LogVolume is the log of the trading volume in the previous year for each stock. PREFERENCING is a dummy variable equal to one for a preferencing exchange, zero otherwise. REGIONAL is a dummy variable equal to 1 for a regional exchange, exclusive of preferencing exchanges, zero otherwise. Results are from ordered probit regressions, assuming three levels: $-\$1/8$, $\$1/8$, $\$3/8$ for NBBO spreads equal to $\$1/8$ at order arrival and $-\$1/4$, $\$0$, and $\$1/4$ for NBBO spreads equal to $\$1/4$ at order arrival. An "*" indicates significance at the 1% level. Estimated spreads are reported for the following order type: MIDDLE = 1, LATE = 0, LogSize = 6.215 (500 shares), LogSameSideDepth = 7.824 (2,500 shares), and LogOppSideDepth = 7.824 for a stock with trading volume of 125,000,000 shares in the previous year.

3.5 Limit order execution analysis

Estimating the quality of limit order submissions on a particular exchange is a more complex problem than that posed by market orders. In particular, because limit orders are priced orders, the execution price cannot be used to assess a limit order's execution quality. The trade is required to occur at the limit price.¹⁵ We measure execution quality of limit orders in two ways: the probability of limit order execution and the ex post transaction cost conditional on execution.

Limit order fill rates are reported in Table 6. As in Table 4, we calculate for each stock its fill rate and report the cross-sectional average fill rate in each category. The results are partitioned into three groups. The first is marketable limit orders, which as discussed above, are priced orders to buy or sell and are immediately executable at the current market quotes. The second group is quote-improving limit orders, which are orders whose price lies inside the quoted bid-ask spread. The third group is at-the-quote limit orders, which would be a buy limit order with a limit price set equal to the prevailing bid or a sell limit order with a limit price set equal to the prevailing ask. All

¹⁵ Technically, for the exchanges, the trade must occur at the limit price or better. As a practical matter, there are few executions of limit orders at prices better than the limit price.

Table 6
Limit order fill rates (%)

Order size	Limit order type	PREF	NYSE	REGL	NYSE	PREF	REGL
1/8-point market (NBBO)							
Small	Marketable	99.6%	96.0%*	97.6%	95.8%*	99.6%	97.6%*
	At-the-quote	90.0	49.4*	72.6	49.2*	80.2	73.6*
Medium	Marketable	99.8	93.7*	98.1	93.4*	99.8	98.6*
	At-the-quote	73.4	48.8*	67.1	48.4*	72.9	68.3*
Large	Marketable	97.8	92.1*	97.7	92.0*	97.7	98.2*
	At-the-quote	71.1	48.7*	63.7	48.9*	71.3	64.9*
All	Marketable	99.5	93.6*	97.5	93.5*	99.5	97.6*
	At-the-quote	76.5	48.4*	69.9	48.3*	76.7	70.2*
1/4-point market (NBBO)							
Small	Marketable	99.5	96.4*	96.8	95.4*	99.5	96.9*
	Qt. improving	96.0	85.6*	90.8	84.8*	96.8	91.4*
	At-the-quote	66.1	29.2*	57.7	28.8*	66.9	62.0
Medium	Marketable	98.6	91.7*	99.5	90.2*	98.7	98.0
	Qt. improving	93.9	80.9*	88.4	81.5*	93.6	89.2*
	At-the-quote	63.3	32.2*	48.3	31.9*	65.9	50.8*
Large	Marketable	100.0	92.3*	99.1	92.8*	100.0	97.6
	Qt. improving	91.8	79.6*	86.2	78.8*	92.3	85.9*
	At-the-quote	59.3	32.3*	51.8	36.6*	59.6	54.5
All	Marketable	99.4	92.3*	97.4	92.3*	99.4	97.5*
	Qt. improving	95.1	81.8*	90.0	81.6*	95.1	90.1*
	At-the-quote	62.7	29.9*	54.3	28.8*	64.7	56.2*

The table reports the cross-sectional average fill rates for the sample and time period described in Table 2. Cell values in adjacent columns represent the average of stocks present in both columns. Orders are defined to be marketable limit orders if the limit price better or equals the NBBO at order arrival time. Orders are quote improving if the order is within the bid-ask spread. Orders are at-the-quote if the order is a buy (sell) order with a limit price equal to the NB bid (NB ask) price. Small orders are for 100 to 500 shares. Medium orders are for 501 to 1,000 shares. Large orders are for more than 1,000 shares. An “*” indicates the fill rate on one market category is significantly different from the fill rate on the other category at the 5% level using a paired *t*-test. “PREF” indicates preferencing regional exchanges. “REGL” indicates nonpreferencing regional exchanges.

other limit orders with limit prices away from the market are not included in any subsequent analysis.

The top panel of Table 6 reports the limit order fill rates in 1/8-point markets for marketable limit orders and at-the-quote limit orders, there being no possibility to submit a quote-improving order for the sample stocks in minimum variation markets. Because marketable limit orders generally convert immediately into market orders, their fill rates are very high, in excess of 93% on the NYSE and in excess of 97% on the regional exchanges.¹⁶ The story is much different for at-the-quote orders. For all order sizes on the NYSE, about half of at-the-quote limit orders are filled, whereas for regional exchanges this number is between 64% and 80%. In general, execution rates on both the preferencing and the nonpreferencing regional exchanges are statistically higher than those of the NYSE using a paired *t*-test.

The difference in fill rates for the NYSE and the regional exchanges is striking and pronounced. There are at least two reasons for this difference.

¹⁶ Possible reasons for the nonexecution of marketable limit orders include (1) orders may exceed the quoted depth; (2) orders may be canceled shortly after submission; and (3) orders may arrive under adverse market conditions.

First, as shown in Table 3, the NYSE provides considerably more depth. Some of these quotes are likely due to the presence of orders on the limit order book, whether entered through SuperDot or from the exchange floor. An incoming at-the-quote limit order must therefore get in line behind a large number of shares on the NYSE. However, the relative shallow quotes of the regional exchanges in Table 3 suggest the limit order book is sparsely populated, allowing the incoming at-the-quote limit order to move to the front of the execution queue sooner. The fill rate difference between the NYSE and the regional exchanges is consistent with thin limit order books being associated with higher fill rates.

The second reason for the higher execution rate likely arises from the preferencing practices of the exchange specialists. As in all U.S. equity markets, customer interest takes priority over market-maker interest at a given price. Thus a customer who enters an at-the-quote buy order at \$20 in a market that is quoted at \$20 to $20\frac{1}{8}$ has priority over the specialist at a price of \$20. For the specialist to preference order flow and buy as principal against the next incoming market sell order, he must either step up and buy at $20\frac{1}{8}$ or else take out the customer buy order at \$20. The buy effectively blocks the specialist from dealing as a preferencing specialist, giving him an incentive to simply remove the limit buy order by trading with it, either immediately as principal or as an agent with the next incoming market sell order. Evidence supporting our second conjecture is that the preferencing regionals have higher fill rates than the nonpreferencing regionals.

Turning to 1/4-point markets, the results are qualitatively similar to the 1/8-point panel for marketable and at-the-quote limit orders. For the preferencing exchanges, the chance of a small or medium at-the-quote limit order remaining unexecuted is approximately 35%, where for the NYSE the probability is approximately 70%. Similar though less dramatic numbers are seen for at-the-quote limit orders on the nonpreferencing regional exchanges. In all cases we can reject the hypothesis that the execution rate of at-the-quote limit orders on the preferencing and nonpreferencing regionals is equal to the rate on the NYSE. As above, the data also show execution rates on preferencing exchanges are typically statistically higher when compared to nonpreferencing regional exchanges.

For quote-improving limit orders, there is variation in the probability of execution among the exchanges. On the preferencing exchanges the cross-sectional average fill rate is more than 95%, while on the NYSE this number is 82%. The nonpreferencing regional exchanges have a higher (lower) fill rate for quote-improving limit orders than the NYSE (preferencing exchanges). This may arise because of the absence of a significant limit order book on the regional exchanges. Again, because public orders take priority over the dealer or specialist, a limit order on a regional exchange is more likely to be binding on the specialist's trading activities than such an

order on the NYSE. Accordingly, the order may be taken out quickly by the dealer to facilitate future proprietary trading activities.

Battalio et al. (2002) report fill rates on the NYSE and some of the regional exchanges. Our fill rates on the regional exchanges are similar to those in Battalio et al.; however, our NYSE fill rates are considerably lower. The difference in NYSE fill rates may be partially explained by canceled limit orders. That is, Battalio et al. exclude canceled orders and we do not because we do not have the ability to distinguish cancelled orders in our dataset. As a reference, Alexander and Peterson (1999) report cancellation rates on the NYSE of 40% and 55% for at-the-quote sell orders in 1/8- and 1/4-point markets and 20% for quote-improving sell orders. Thus the NYSE fill rates may be lower because traders cancel orders sent to the NYSE more frequently than they cancel orders sent to the regional exchanges.

Similar to the analysis of effective spreads, the results in Table 6 hold the NBBO spread, order size, and order type constant. In Table 7 we present the results of a probit regression predicting whether an order was filled. The explanatory variables are the same as those in Table 5. The results, in general, support the results in Table 6. Specifically the fill rates on the regional exchanges, including the preferencing exchanges, are significantly higher than on the NYSE regardless of NBBO spread and limit price relative to the quotes. Consistent with Table 6, the at-the-quote limit orders

Table 7
Fill rate analysis—probit regression

	1/8-point market (NBBO)		1/4-point market (NBBO)		
			Type of limit order		
	Marketable	At-the-quote	Marketable	Quote-imp.	At-the-quote
Intercept	1.945*	-1.722*	1.865	0.468	-1.799*
MIDDLE	0.094*	-0.009	0.097	0.092*	-0.074*
LATE	0.005	-0.155*	0.101	-0.008	-0.220*
LogSize	-0.086*	0.007*	-0.006	-0.037*	0.085*
LogSameSideDepth	0.008	-0.320*	-0.074*	-0.072*	-0.294*
LogOppSideDepth	0.192*	0.161*	0.157*	0.242*	0.122*
LogVolume	-0.074*	0.166*	-0.061	-0.039*	0.125*
PREFERENCING	0.765*	1.047*	1.036*	0.722*	1.226*
REGIONAL	0.622*	0.812*	0.896*	0.348*	0.932*
Pseudo- R^2	0.034	0.168	0.044	0.047	0.124
Estimated fill rates under specified conditions					
NYSE	95.3%	56.3%	92.5%	82.5%	36.0%
Preferencing	99.3	88.6	99.3	95.1	80.7
Regionals	98.9	83.4	99.0	90.0	71.6

The table reports the parameter estimates from a probit regression of whether a limit order fills on several explanatory variables for the sample of orders described in Table 2. MIDDLE is a dummy variable equal to one if the order is placed between 11:00 and 2:30 p.m., zero otherwise. LATE is a dummy variable equal to one if the order is placed after 2:30 p.m., zero otherwise. LogSize is the log of the order size. LogSameSideDepth is the log of the NBBO depth on the same side. LogOppSideDepth is the log of the NBBO depth on the opposite side. LogVolume is the log of the trading volume in the previous year for each stock. PREFERENCING is a dummy variable equal to one for a preferencing exchange, zero otherwise. REGIONAL is a dummy variable equal to one for a regional exchange, exclusive of preferencing exchanges, zero otherwise. An “*” indicates the significance at the 1% level. Estimated fill rates are reported for the following order type: MIDDLE = 1, LATE = 0, LogSize = 6.215 (500 shares), LogSameSideDepth = 7.824 (2,500 shares), and LogOppSideDepth = 7.824 for a stock with trading volume of 125,000,000 shares in the previous year.

have the largest variation in fill rates across trading venues. Of interest is that Table 6 indicates at-the-quote fill rates increase with order size on the NYSE in 1/4-point markets. The coefficient estimate on *LogSize* indicates that although fill rates are positively related with order size, the effect is not very big. To test whether fill rates are statistically different between preferencing and nonpreferencing regional exchanges, we also estimate the probit regressions by considering orders sent to nonpreferencing regionals as the base case and including an indicator variable for NYSE orders. The coefficient estimates indicate that in all cases, except marketable limit orders in 1/4-point markets, the fill rates on preferencing regionals are statistically higher than the fill rates on nonpreferencing regionals.

3.6 Ex post transaction costs of limit orders

The second measure of limit order market quality we examine is the ex post cost of a limit order execution. The price the limit order pays is usually established by the limit price. Typically limit orders are buying low and selling high, all the while competing with the specialist. Conditioned on a limit buy order executing, there is a chance the market price for those shares will continue to fall through the limit price and keep falling for some period of time, as in Rock (1991). Thus having bought shares on the bid with a limit buy order, one might be curious to know what the price is for those shares at some time in the future. If the price falls, then the limit order trader incurs adverse selection costs.

One measure of the adverse selection cost is the difference between the execution price and the price of shares at some time in the future. If the adverse selection problem is severe, a buy (sell) limit order will execute prior to a market decrease (increase). The difference between the future market price of shares and the limit price can be used to develop a ranking of limit order execution quality among exchanges.¹⁷ A similar technique is used in Harris and Hasbrouck (1996).

Table 8 illustrates the adverse selection problem by presenting the cross-sectional average ex post transaction costs. The ex post cost, $Cost_{ex\ post}$, for buy orders is computed as the difference between the volume-weighted execution price of the limit order and the NBBO bid price five minutes after execution, and analogously for sell orders:

$$Cost_{ex\ post, buy} = \text{trade price} - \text{bid price}_{NBBO, t+5\ min}. \quad (3)$$

$$Cost_{ex\ post, sell} = \text{ask price}_{NBBO, t+5\ min} - \text{trade price}. \quad (4)$$

¹⁷ A more thorough analysis would include in the cost measure the opportunity cost of a limit order not executing. Such costs of unexecuted orders have been shown to be large in other instances [Perold and Sirri (1994)]. Because we do not know the nature of the limit order strategy, which could involve sequences of cancellations and resubmissions, no unambiguous measure of opportunity cost can be computed.

Table 8
Average ex post transaction cost (in cents) for limit orders

Order size	Limit order type	PREF	NYSE	REGL	NYSE	PREF	REGL
1/8-point market (NBBO)							
Small	Marketable	12.5	10.2*	12.7	10.3*	12.5	12.6
	At-the-quote	6.7	6.5	6.1	6.6*	6.7	6.0*
Medium	Marketable	12.3	9.5*	12.8	9.7*	12.3	12.7
	At-the-quote	6.4	6.1	7.0	6.0*	6.2	6.8
Large	Marketable	11.4	8.5*	11.2	8.5*	11.3	11.3
	At-the-quote	5.5	5.4	6.3	5.4*	5.8	6.3
All	Marketable	12.0	9.4	12.4	9.4*	12.0	12.4
	At-the-quote	6.7	6.1*	6.3	6.1	6.8	6.2
1/4-point market (NBBO)							
Small	Marketable	14.0	13.0	19.2	12.9*	13.5	19.2*
	Qt. improving	8.3	8.1	9.4	8.1*	8.3	9.2*
	At-the-quote	7.4	6.6	7.5	6.8	7.1	7.7
Medium	Marketable	14.7	8.6*	17.8	9.9*	14.1	16.6
	Qt. improving	8.4	7.4*	9.9	7.4*	8.3	9.9*
	At-the-quote	7.0	6.4	7.0	5.8	6.6	6.7
Large	Marketable	14.9	11.0*	13.9	11.2*	14.7	13.6
	Qt. improving	7.5	6.1*	9.2	6.2*	7.7	9.3
	At-the-quote	8.2	5.0*	5.8	5.5	6.8	4.9
All	Marketable	14.3	11.1*	18.3	11.4*	14.1	18.2*
	Qt. improving	8.2	7.3*	9.3	7.2*	8.2	9.4*
	At-the-quote	7.3	6.1*	6.9	6.1	7.3	7.1

The table reports the cross-sectional average ex post transaction cost for the sample described in Table 2. Cell values in adjacent columns represent the average of stocks present in both columns. The ex post cost is measured (cents per share) for buy orders as the difference between the volume-weighted price and the best bid five minutes after execution. For sell orders, the ex post cost is the difference between the best ask five minutes after execution and the volume-weighted price. Orders are defined to be marketable limit orders if the limit price betters or equals the NBBO at order arrival. Orders are quote improving if the order is within the bid-ask spread. Orders are at-the-quote if the order is a buy (sell) order with a limit price equal to the bid (ask) price. Small orders are for 100 to 500 shares. Medium orders are for 501 to 1,000 shares. Large orders are for more than 1,000 shares. An "*" indicates the ex post transaction cost on one market category is significantly different from the cost on the other category at the 5% level using a paired *t*-test. "PREF" indicates preferencing regional exchanges. "REGL" indicates nonpreferencing regional exchanges.

Turning to the first panel, we see marketable limit orders in 1/8-point markets have ex post costs on the order of one tick for the regional exchanges. Conditional on executing a marketable limit order at the ask price, five minutes later the bid price has not moved by very much and is on average at the same price level as when the marketable limit order was executed. This is consistent with the notion that marketable limit orders carry little information and usually do not lead to revisions in quoted prices. Note, however, for the NYSE, the ex post costs decrease with increases in the order size, consistent with larger marketable limit orders having more information. In 1/4-point markets, small marketable limit orders also have larger costs than large marketable limit orders.

Next consider at-the-quote limit orders. In a 1/8-point market when controlling for order size, we find the ex post costs are about the same on the NYSE and the preferencing exchanges. However, on the nonpreferencing regionals, small at-the-quote limit orders have smaller ex post costs and large at-the-quote limit orders have larger ex post costs than on the NYSE.

At-the-quote limit orders in 1/4-point markets do not show much variation in ex post costs, as the paired *t*-tests do not indicate significance.

The results for quote-improving limit orders indicate the NYSE has lower adverse selection costs than the regionals. For the nonpreferencing regionals, on average, the costs are about 2 cents higher than on the NYSE and the difference becomes more pronounced as order size increases. This is consistent with larger quote-improving orders carrying little or no information, leading to an ex post cost of just under one tick. On the NYSE, such orders may presage a change in the quote, perhaps because they were placed during an interval when the spread was temporarily wide, leading to a lower measured ex post cost.

In aggregate, the results on ex post costs for limit orders are mixed.¹⁸ In 1/8-point markets, economically there is not much variation across exchanges in the ex post costs of at-the-quote limit orders, although the nonpreferencing regional exchanges tend to have the lowest costs. Coupled with the thin limit order books on the nonpreferencing regionals as indicated in Table 3, the lower adverse selection costs are consistent with the model of Glosten (1994). In 1/4-point markets, it is sometimes the case that the NYSE has lower ex post costs, but this may be due to the nature of the order flow arriving at the primary market. In no case do we find the average ex post costs of the preferencing regional exchanges statistically higher than the average ex post costs on the nonpreferencing regional exchanges, and in some cases those costs are statistically and economically lower. It is fair to conclude, however, that there is no evidence the regional exchanges in general, and the preferencing regionals in particular, are systematically imposing relatively high ex post transaction costs on their order flow.

The limit order analysis leaves us with the curious conclusion that it appears the regional exchanges are in many instances the preferred point for routing a customer limit order. This has been noted in Battalio et al. (2002), which concludes that brokers can strategically route limit orders to improve execution quality. A combination of thin limit order books, incentives of specialists to unblock their proprietary operations, and moderate ex post costs make regional exchanges in many cases a more attractive venue than the NYSE. It is somewhat of a mystery why more customer orders are not routed there. The answer may lie in the nature of the business practice itself on the regionals. The preferencing dealers have no desire to compete against the customer limit orders. Order routing regulations, however, prevent them from routing limit orders to the central market and market orders to themselves. Nonetheless, regional market makers have found ways to increase the ratio of market-to-limit orders on their exchanges, as shown in Table 2. Customers are almost certainly not aware of the difference in execution probabilities and

¹⁸ Though the prices used as proxies for market price are taken by measuring the same side quote five minutes after the trade, the results do not change appreciably if prices are observed 15 or 60 minutes after the trade.

costs. Even if they are, customers do not have the ability to select the location to which their orders are routed. In fact, customers do not have the right to know where a given broker routes their order flow. Brokers are required to release only their routing policies, including whether they engage in payment for order flow. They do not have to tell customers where they route market or limit orders. Unless the customer is willing to sample a large number of brokers to learn their routing practices, they will not be able to select the venue for their orders.

3.7 Quote-based competition for order flow

As a final exploration of the effects of preferencing on order executions, this section considers other order inducements on the order routing behavior of brokers. In preferencing and payment for order flow arrangements, orders may arrive at a particular venue not because of the prices or depths that a market is quoting, but because of an institutional order routing arrangement. Specialists at a particular venue have little incentive to quote aggressively since doing so would provide a free option to the market [Copeland and Galai (1983), Glosten and Milgrom (1983)] with little probability of garnering additional order flow.

To empirically test whether such a description characterizes the market for order flow, we compute the probability a particular exchange receives an order and whether that probability depends on the quotations of the market maker. To make this calculation we look at the individual price quotations and depths quoted by each of the exchanges and consider to which exchange a market order is routed. The top row of Table 9 is the unconditional probability the next market order arrives at a particular exchange. This is effectively the

Table 9
Quote-based competition and conditional order arrival probabilities

	Exchange E						
	Preferencing exchange			Nonpreferencing regional exchange			
	NYSE	CSE	BSE CSI	BSE-non-CSI	CHX	PSE	PHLX
Pr{order arrives at E }	55.0%	11.7%	3.5%	3.1%	8.7%	10.6%	7.4%
Pr{order arrives at $E E$ not NBBO}	60.0	11.7	3.5	3.1	8.7	10.7	7.4
Pr{order arrives at $E E$ is NBBO}	54.7*	11.8	14.7*	3.8	7.0*	8.1*	6.5

The table gives the unconditional and the conditional order arrival probabilities on the NYSE and the five regional exchanges. The first row is the unconditional probability of an order arriving at exchange E , where E is one of the six exchanges, and the numbers in the row are expressed in percentages. The second and third rows express the probabilities that an order arrives at exchange E conditional on that exchange either not quoting or quoting, respectively, the NBBO. The NBBO is defined as the lowest offer or the highest bid price with the greatest quoted depth. For example, if the CSE has a low bid quote far away from the NBBO, the probability that the next order sell arrives at the CSE is 11.7%. If the CSE is at the NBBO on the bid side of the market, the probability that the next sell order is sent to the CSE rises to 11.8%. The “*” denotes a p -value of less than 0.01 for a chi-squared test of the equality in the probability of order arrival conditional on whether the exchange E is or is not quoting at the NBBO. Thus a comparison between the second and third row is a measure of the degree of elasticity of the order flow with respect to price; that is, a measure of the increased likelihood of receiving an order based on the aggressiveness of exchange E 's quotations. All numbers are expressed as percentages. The BSE results are reported separately for the competing specialist initiative (CSI) stocks and non-CSI stocks. Sample is market orders for 334 liquid NYSE/CSE traded stocks from October 28 to November 1, 1996.

market share of all market orders in our sample during the first week. The NYSE has the largest share of these orders, followed by CSE, PSE, and the other regional exchanges.

The following two lines then measure how this order arrival probability changes when the exchange is providing the best quote in the marketplace, that is, it is quoting at the NBBO on the side of the incoming market order. For example, in the case of an incoming market order, we consider which exchange has the best price. If there are multiple exchanges at the best price, then the exchange with the greatest depth is deemed to have the best quotation. The second row of the table gives the arrival probability if that particular exchange is not at the NBBO. For example, if the CSE does not have the best bid in the market, then it has an 11.7% chance of receiving the next market sell order. However, the third row shows if the CSE steps up and quotes the best bid price with greatest depth, that probability rises to only 11.8%. In the case of the CSE, this change is not significant. A comparison of rows two and three is a measure of the extent of elasticity of the order flow with respect to quotation. It is clear that, with the exception of the BSE CSI, a preferencing exchange, the order flow shows no propensity to increase when an exchange provides a more aggressive quotation. Thus it appears the routing decision is not made on an order-by-order basis, sending orders to a particular exchange based on their quotations at that instant. This does not mean that order flow is wholly insensitive to quotations. It may be that introducing brokers route orders based on the average quality of a venue's quotation and executions. At any one moment, the broker does not change the routing decision but instead depends on long-run averages to use as a guide in the routing decision. Such a policy is not inconsistent with the statutory guidance provided by the SEC, which does not demand order-by-order comparisons. The results of Table 9 suggest that any model of market quality should not take as given that the market structure resembles a pure auction market. The results are consistent with the notion that preferencing and other order routing practices have decreased minute-by-minute quote-based competition and that the instantaneous elasticity of order flow with respect to price may be slight.

4. Discussion and Conclusions

We analyze market quality on the NYSE (the primary market), the CSE and the BSE (regional exchanges with formal preferencing programs), and on the CHX, PSE, and PHLX (regional exchanges without formal preferencing programs). We compare quotation and execution quality across exchanges for the same time period for the same set of securities. The market quality statistics based on quoted spread and quoted depth indicate better performance for the NYSE than the regional exchanges. The regional exchanges show considerable dispersion in quoted spreads and the quality of quotations

on the preferencing exchanges is not uniform. The CSE quotes the tightest spreads of all regional exchanges, and as measured by the percentage of the time the quotes are at the NBBO, the CSE provides more competitive quotes than the other regional exchanges.

For market order executions, the NYSE has smaller effective spreads than either the preferencing exchanges or the nonpreferencing regional exchanges. Most trades occur in 1/8-point markets where the NYSE has an effective spread advantage for small market orders of about 1.5 cents over the preferencing exchanges and about 2 cents over the nonpreferencing regional exchanges. In 1/4-point markets the NYSE's cost advantage for small market orders widens to more than 2 cents when compared to preferencing exchanges and almost 7 cents for nonpreferencing regional exchanges.

The picture is slightly different for limit orders. The NYSE is the primary market and receives most of the limit order traffic. Limit orders generally are executed subject to price priority and time precedence. The NYSE execution rate of limit orders submitted at the prevailing quote is about 50% in 1/8-point markets and 30% in 1/4-point markets. On the preferencing regional exchanges, the probability of an at-the-quote limit order executing is 77% (63%) in 1/8-point (1/4-point) markets. On the nonpreferencing regional exchanges, the probability of an at-the-quote limit order executing is 70% (54%) in 1/8-point (1/4-point) markets. There is little evidence to suggest that limit orders are subject to greater adverse selection by receiving an execution on a regional exchange rather than on the NYSE.

It is important to point out that the U.S. stock markets have changed considerably since our sample period. One significant change that may have an impact on our findings is the move to decimal pricing. For example, in a decimal trading environment it may be easier for a specialist to step in front of an order. This practice may result in lower spreads for small market orders and a lower fill rate for limit orders. Using the November 2001 monthly disclosure of market quality under Rule 11Ac1-5 for orders of 100 to 499 shares on the 28 NYSE-listed stocks in the Dow Jones Industrial Average, we find that the NYSE's advantage over the regional exchanges in terms of the effective spread has diminished considerably. In fact, the NYSE's cross-sectional average effective spread for the Dow sample is about 0.8 cents higher than the same average on the preferencing regional exchanges (3.0 cents versus 2.2 cents). The NYSE's effective spreads are also larger than the nonpreferencing regional exchanges, but the difference is a smaller 0.4 cents (3.0 cents versus 2.6 cents). With regard to limit order fill rates, the evidence is very similar to our sample data in Table 6. Fill rates for at-the-quote limit orders of 100 to 499 shares in the Dow sample average 48% on the NYSE. Significantly higher fill rates are found for the same stocks on the nonpreferencing regionals (60–70%) and even higher fill rates are found on the preferencing regional exchanges (73–98%).

In summary, those who argue that preferencing harms the market may say it does so by leading to worse executions on all exchanges. Unfortunately we cannot directly observe the effect of preferencing on all exchanges. The data indicate the effective spreads of market orders trading on preferencing regional exchanges tends to be lower than the effective spreads of market orders trading on nonpreferencing regional exchanges. In addition, limit orders have a greater probability of executing on preferencing regional exchanges than on the nonpreferencing regional exchanges. Thus the answer to the question posed in the introduction is that customers should be concerned with the order routing practices of brokers. In most cases the market quality of preferencing regional exchanges is superior to the market quality of nonpreferencing regional exchanges.

References

- Admati, A., and P. Pfleiderer, 1988, "A Theory of Intraday Patterns: Volume and Price Variability," *Review of Financial Studies*, 1, 3–40.
- Alexander, G., and M. Peterson, 1999, "Short Selling on the New York Stock Exchange and the Effects of the Uptick Rule," *Journal of Financial Intermediation*, 8, 80–116.
- Bacidore, J., K. Ross, and G. Sofianos, 1999, "Quantifying Best Execution at the New York Stock Exchange: Market Orders," working paper, New York Stock Exchange.
- Battalio, R., 1997, "Third Market Broker-Dealers: Cost Competitors or Cream Skimmers?," *Journal of Finance*, 52, 341–352.
- Battalio, R., J. Greene, B. Hatch, and R. Jennings, 2002, "Does the Limit Order Routing Decision Matter?," *Review of Financial Studies*, 15, 159–194.
- Battalio, R., J. Greene, and R. Jennings, 1997, "Do Competing Specialists and Preferencing Dealers Affect Market Quality?," *Review of Financial Studies*, 10, 969–993.
- Battalio, R., J. Greene, and R. Jennings, 1998, "Order Flow Distribution, Bid-Ask Spreads, and Liquidity Costs: Merrill Lynch's Decision to Cease Routinely Routing Orders to Regional Stock Exchanges," *Journal of Financial Intermediation*, 7, 338–358.
- Bessembinder, H., and H. Kaufman, 1997, "A Cross-Exchange Comparison of Execution Costs and Information Flow for NYSE-Listed Stocks," *Journal of Financial Economics*, 46, 293–319.
- Copeland, T., and D. Galai, 1983, "Information Effects of the Bid-Ask Spread," *Journal of Finance*, 38, 1457–1469.
- Easley, D., D. Kiefer, and M. O'Hara, 1996, "Cream Skimming or Profit Sharing? The Curious Role of Purchased Order Flow," *Journal of Finance*, 51, 811–833.
- Easley, D., and M. O'Hara, 1987, "Price, Trade Size, and Information in Securities Markets," *Journal of Financial Economics*, 19, 69–90.
- Glosten, L., 1994, "Is the Electronic Limit Order Book Inevitable?," *Journal of Finance*, 49, 1127–1161.
- Glosten, L., and P. Milgrom, 1983, "Bid, Ask, and Transactions Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 85, 71–100.
- Harris, L., and J. Hasbrouck, 1996, "Market vs. Limit Orders: The SuperDOT Evidence on Order Submission Strategy," *Journal of Financial and Quantitative Analysis*, 31, 213–232.
- Lee, C., 1993, "Market Integration and Price Execution for NYSE-Listed Securities," *Journal of Finance*, 48, 1009–1038.

Lee, C., and M. Ready, 1991, "Inferring Trade Direction from Intraday Data," *Journal of Finance*, 46, 733–746.

Macey, J., and M. O'Hara, 1997, "The Law and Economics of Best Execution," *Journal of Financial Intermediation*, 6, 188–223.

Perold, A., and E. Sirri, 1994, "The Cost of Institutional Equity Trading," working paper, Harvard University.

Petersen, M., and D. Fialkowski, 1994, "Posted Versus Effective Spreads: Good Prices or Bad Quotes?," *Journal of Financial Economics*, 35, 269–292.

Peterson, M., and E. Sirri, 2002, "Order Submission Strategy and the Curious Case of Marketable Limit Orders," *Journal of Financial and Quantitative Analysis*, 37, 221–241.

Rock, K., 1991, "The Specialist's Order Book," working paper, Harvard University.

Ross, K., J. Shapiro, and K. Smith, 1996, "Price Improvement of SuperDot Market Orders on the NYSE," working paper, New York Stock Exchange.

SEC, 1997, "Report on the Practice of Preferring; Pursuant to Section 510(c) of the National Securities Markets Improvement Act of 1996," Securities and Exchange Commission, Washington, DC.

Copyright of Review of Financial Studies is the property of Society for Financial Studies and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.