ELSEVIER Journal of Financial Economics 46 (1997) 293–319

JOURNAL OF Financial ECONOMICS

A cross-exchange comparison of execution costs and information flow for NYSE-listed stocks[1]

Hendrik Bessembinder, Herbert M. Kaufman*

Department of Finance, Arizona State University, Tempe, AZ 85287, USA

Received 4 September 1996; received in revised form 6 May 1997

Abstract

We examine execution costs for trades in NYSE issues completed on the NYSE, the NASD dealer market, and the regional stock exchanges during 1994. We find that effective bid–ask spreads are only slightly smaller at the NYSE. However, realized bid–ask spreads, which measure market-making revenue net of losses to better-informed traders but gross of inventory or order-processing costs, are lower on the NYSE by a factor of two to three. This differential is attributable to the successful 'cream skimming' of uninformed trades by off-NYSE market makers. These findings reinforce existing concerns about whether orders are routed so as to receive the best possible execution.

Keywords: Trade execution costs; Cream skimming; Regional stock exchanges; NASD dealer market

JEL classification: G10; G14

1. Introduction

Stocks listed on the NYSE are also traded on the five domestic regional stock exchanges, on the National Association of Securities Dealers (NASD) dealer market, and on electronic networks such as Instinet. While multiple trading locations potentially benefit traders by increasing competition, it is unresolved

* Corresponding author. Tel.: 602/965-3131; fax: 602/965-8539.

[1] Portions of this paper are extracted from our earlier study titled 'Quotations and Trading Costs on the Domestic Equity Exchanges'.

PII S0304-405X(97)00032-9

whether investors actually benefit from the segmentation of trading as it has evolved. This paper presents empirical evidence by exchange on quotations, trade execution costs, and information revelation for NYSE-listed securities in order to shed additional light on this issue.

In the standard economic model of a competitive industry, customers choose among vendors based on price and quality. In contrast, the trade location for equity transactions is typically selected by brokers rather than retail customers. Blume and Goldstein (1997) document that most trades in NYSE issues that are executed off the NYSE occur when the trading exchange is posting inferior quotes, indicating that the trade location is often determined for reasons other than the location of the best quote.

Two controversial practices that affect the execution location are 'payment for order flow' and 'dealer preferencing' programs. In order flow payment programs, dealers executing trades make cash or nonpecuniary payments to brokers in exchange for the routing of retail orders. It is unclear whether order flow payments are effectively passed through to customers in the form of lower commissions. In preferencing programs, dealers with integrated brokerage operations trade directly against their own retail order flow, with the resulting transactions printed by affiliated specialist operations at regional exchanges.[2]

Under law, brokers owe their customers a duty of 'best execution', which the United States Securities and Exchange Commission (SEC, 1996, p. 141) describes as requiring a broker to 'seek the most favorable terms reasonably available under the circumstances for a customer's transaction'. Macey and O'Hara (1996) provide an overview of legal and economic issues related to the concept of best execution. They note the absence of specific definitions of best execution or of an explicit best execution rule. Chordia and Subrahmanyam (1995) and Harris (1993), among others, observe that payments for order flow and dealer preferencing arrangements can create a conflict between brokers' interests and the duty of best execution.

Order flow agreements and Intermarket Trading System (ITS) rules generally provide that trades are to be executed at prices no worse than the 'inside' intermarket quotes. Therefore, if posted quotes always represent the most favorable terms available, there would be little concern as to whether investors are harmed by brokers' order routing decisions. In practice, however, trades are often executed at prices superior to the posted quotes (see, e.g., Peterson and Fialkowski, 1994; Bessembinder and Kaufman, 1996; Blume and Goldstein, 1997), which raises the possibility that a trade executed at the inside bid or ask

[2] Battalio et al. (1995) provide a description of the pilot program for clearing preferenced trades at the Cincinnati exchange (recently made permanent by the SEC) and of a separate pilot program at the Boston exchange. See also 'Merrill shifts on small investors' orders', *The Wall Street Journal*, October 29, 1995 and 'Member firms compete with the big board', *The Wall Street Journal*, May 31, 1995.

quote might have received a better execution at an alternate market. The SEC (1996), (p. 143) has stated that brokers should consider opportunities for price improvement as part of the duty of best execution:

> The Commission believes that broker-dealers deciding where to route or execute small customer orders in listed or OTC securities must carefully evaluate the extent to which this order would be afforded better terms if executed in a market or with a market maker offering price improvement opportunities.

Concerns as to whether orders are routed so as to receive the best possible execution are enhanced if, as Easley et al. (1996) suggest, order flow payment agreements target those orders originating with relatively uninformed traders. Uninformed orders are likely candidates for execution within the quotes in a competitive market, since they do not impose on market makers the adverse selection costs described by Glosten and Milgrom (1985). The immediate concern regarding the possible 'cream skimming' of uninformed orders is whether the diverted orders receive the best possible trade executions.

A second concern is that cream skimming increases the proportion of NYSE trades that originate with informed traders, potentially leading to a widening of NYSE bid–ask spreads. Wider NYSE spreads can increase execution costs off the NYSE as well, since the NYSE quotes often constitute the best intermarket quotes that order flow agreements promise to match.

Further, cream skimming potentially affects the overall integrity of the equity markets. Glosten and Milgrom show that an equity market can fail entirely if privately informed traders become too prevalent. This concern is potentially reinforced by the increasing use of electronic crossing networks such as Instinet and Posit, which may also divert uninformed order flow from the NYSE. Segmentation of the equity markets such that most price discovery takes place on the NYSE, while most liquidity trades are executed elsewhere, is unlikely to constitute a viable long-run equilibrium.

Unlike brokerage commissions, which are charged explicitly, execution costs are implicitly charged to traders through the prices at which trades are completed, and are not separately reported. Given the volatility of security prices, it can be difficult for customers to determine whether their brokers have succeeded in securing the best possible executions for their trades. We provide empirical evidence comparing average execution costs for trades executed at the NYSE to trades in NYSE issues executed on the regional exchanges and on the NASD dealer market. We examine several measures of trade execution costs, including quoted bid–ask spreads, effective bid–ask spreads (which account for executions inside the quotes), and realized bid–ask spreads (which measure trades' temporary price impacts).

Lee (1993) provides related empirical evidence. He compares average quoted and effective bid–ask spreads for trades executed off the NYSE to matched trades in the same issues that are executed on the NYSE within 10 min. However, the need for a matching trade forces Lee to exclude many trades from his analysis. In contrast, we include all trades for our sample of firms. Huang and Stoll (1994) estimate quoted, effective, and realized spreads by exchange for a sample of large capitalization NYSE issues. We extend the analysis to include small and medium capitalization issues as well, and to several issues that Huang and Stoll do not examine, including average quote size and the cream skimming question. We also assess whether variation in economic attributes such as average trade size, firm size, or return volatility can explain variation in average execution costs across exchanges. Further, we compare execution costs for securities eligible for trading under SEC rule 19c-3 (which allows NYSE member firms to compete with the NYSE specialist by executing trades in the OTC market) to costs for non-19c-3 firms.

The cream skimming issue has been empirically investigated by Easley et al. (1996) and Battalio (1997). However, their evidence is indirect and leads to differing conclusions. Easley et al. examine patterns in signed order flow for 30 firms and conclude that cream skimming does occur on the Cincinnati exchange. Battalio examines bid–ask spreads around points in time when a major firm commences to purchase order flow, and concludes from the absence of an observed widening of spreads that cream skimming does not occur. In contrast, we directly measure the information content of trades by assessing the impact on subsequent market prices of executions on each exchange for a broad sample of 300 firms. We find that trades executed off the NYSE contain significantly less information than trades executed on the NYSE. This constitutes strong direct evidence that cream skimming of uninformed order flow occurs.

Our analysis indicates that traders pay higher average execution costs when trades are executed off the NYSE. The differentials in effective bid–ask spreads are moderate, averaging 0.05% of trade value for large-firm trades, 0.08% for medium-firm trades, and 0.18% for small-firm trades. However, given that we also document that non-NYSE market makers successfully cream skim trades that convey relatively little information, it might be expected that the non-NYSE executions would provide lower execution costs. This concern is reinforced by examination of average realized bid–ask spreads, which measure trades' temporary price impacts, or equivalently, market-making revenue net of the adverse selection costs attributable to trades' information content. We find that realized spreads are typically two to three times larger for trades executed off the NYSE than for NYSE trades. In general, average execution costs are somewhat larger and the cream skimming of uninformed order flow is more pronounced for securities eligible for off-exchange trading by NYSE member firms under SEC rule 19c-3.

In the absence of systematic variation in the order processing or inventory costs of market making, it might be expected that competition would tend to equalize bid–ask spreads net of information-related costs. The finding that realized bid–ask spreads for trades executed off the NYSE are substantially larger than for NYSE trades raises the question of whether trades receive the best possible executions and whether the regulations that currently govern the routing of orders in NYSE securities effectively foster competition.

This paper is organized as follows. Section 2 describes the sample and presents some descriptive statistics. Section 3 describes how we measure trade execution costs and the information content of trades, while Section 4 reports averages of these measures by exchange. Section 5 assesses whether variation in average trade execution costs can be explained by systematic differences in trade size or firm characteristics, and Section 6 provides a summary and our conclusions.

2. Sample and descriptive statistics

Transaction price and quotation data for calendar year 1994 are from the Trade and Quote (TAQ) database, which is made available by the NYSE. We include trades and quotes observed on the NYSE, the NASD dealer market, and the Boston, Cincinnati, Chicago (formerly Midwest), Pacific, and Philadelphia stock exchanges. With the exception of the Cincinnati exchange, which we report separately to highlight its relative competitiveness, empirical results for the regional exchanges are aggregated for reporting purposes.

We focus on a set of 300 NYSE-listed common stocks that are described in detail by Bessembinder and Kaufman (1996). The sample is constructed to include 100 each of large, medium, and small capitalization issues. Table 1 reports some summary statistics for these firms. We include most quotes and trades reported in the TAQ database during 1994. However, some trades and quotes are filtered because of a high likelihood that they reflect errors or because they differ from more common trades and quotes in dimensions that are likely to be related to trading costs.[3] In aggregate, these filters eliminate 3.8% of the raw

[3] Trades are omitted if they are indicated (by correction indicator values in the TAQ Consolidated Trade file of 2 or greater) in the TAQ database to be coded out of time sequence, or coded as involving an error or a correction. Trades indicated to be exchange acquisitions or distributions, trades that involve nonstandard settlement, and trades that are not preceded by a valid same-day quote (as indicated by Sale Condition codes A, C, D, N, O, R, and Z in the TAQ Consolidated Trade file) are omitted. We also omit trades that involve price changes (since the prior trade) of 25% or more if the prior price is over $2 per share, and trades that involve a reported fractional price that is not a round multiple of 1/64. Quotes are omitted if (i) either the ask or bid price is nonpositive, (ii) the differential between the ask and bid prices exceeds $4 or is nonpositive, (iii) the indicated bid or ask quote size is nonpositive, and (iv) the fractional portion of either the ask or bid price is not a round multiple of 1/64. We also omit quotes that are reported in the TAQ Consolidated Quote file with quote condition codes 4, 7, 9, 11, 13, 14, 15, 19, and 20, which are quotes associated with trading halts or designated order imbalances, or that are not firm.

Table 1
Mean market capitalization, share price, trading volume and return volatility for the 300 NYSE-listed sample firms

Reported are simple means for each indicated group. Share price is measured in dollars as of December 31, 1993. Market capitalization reflects the product of share price and the number of shares outstanding on December 31, 1993. Trading volume reflects activity during calendar year 1993. The return standard deviation is computed from daily close-to-close returns during calendar year 1994.

	Large-firm sample	Medium-firm sample	Small-firm sample
Market capitalization ($000)	3,009,004	221,192	39,695
Trading volume (shares/day)	241,078	50,040	22,679
Share price ($)	36.62	18.39	7.91
Daily return standard deviation (%)	0.361	0.472	0.726

trade observations and 0.9% of the raw quote observations. We also exclude trades and quotes that are time stamped outside regular NYSE trading hours.

Table 2 reports some detail regarding the sample. Panel A reports on the sample of quotes while Panel B reports on the sample of trades. Each 'market share' reports the proportion of quotes originating at or trades occurring at the indicated exchange. Trades are reported for the overall sample, by firm-size group, and by three trade-size categories: 'small' trades are less than $10 000, 'medium' trades are $10 000–$199 999, while 'large' trades are $200 000 or more. It should be noted that trades executed on Instinet are reported on the TAQ database as NASD trades and cannot be separately identified. If the institutional orders crossed on Instinet typically receive favorable executions, then the average execution costs for NASD trades that we report are biased downward.

Only 31.3% of the quotes for our sample of NYSE-listed firms are entered at the NYSE. The regional exchanges enter a majority, 55.6%, of the quotes. The count of regional quotes is inflated by the presence of 'autoquotes', which are noncompetitive quotes with sizes of 100 shares that are automatically updated in response to quote revisions at the NYSE. Cincinnati and NASD market makers appear to target large NYSE-listed firms in their quoting activity, accounting for 11.6% and 4.8%, respectively, of quotes for large NYSE-listed firms, but only 0.1% and 1.5%, respectively, of quotes for small NYSE-listed firms.

During 1994, 65.4% of the trades in the 300 sample firms occurred on the NYSE. Consistent with the results that Chordia and Subrahmanyam (1995) report for the 1988–1991 interval, non-NYSE market makers tend to attract small trades: the NYSE share of large trades exceeds 89% for each firm-size category, while the NYSE market share declines monotonically across trade-size

Table 2
Trade and quote sample sizes and market shares

Sample sizes reflect the number of quotes and trades for the 300 NYSE-listed firms described in Table 1 that were time-stamped between 9:30 a.m. and 4:00 p.m. eastern time during calendar year 1994 and that passed the error filters described in the text. Shares reflect the percentage of trades and quotes that originate on the indicated exchange. The large, medium, and small firm-size groups consist of firms with mean market capitalization of $3.0 billion, $220 million, and $39.7 million, respectively. The large trade category includes all trades of over $200,000, the medium trade category includes trades ranging from $10,000 to $200,000, while the small trade category includes all trades of less than $10,000. The exchanges are the New York Stock Exchange (NYSE), the National Association of Securities Dealers (NASD) over-the-counter market, the regional exchanges (excluding Cincinnati) and the Cincinnati Stock Exchange. All trade and quote data are obtained from the Trade and Quote (TAQ) database made available by the NYSE.

Firm group	Trade group	Sample size	Market shares by exchange			
			NYSE	NASD	Regionals	Cincinnati
Panel A: Quotes						
All	NA	10,071,167	31.31	4.09	55.58	9.02
Large	NA	7,409,486	29.52	4.80	54.04	11.64
Medium	NA	1,866,642	35.05	2.38	60.16	2.41
Small	NA	795,039	39.18	1.49	59.22	0.11
Panel B: Trades						
All	All	3,896,108	65.42	10.15	21.43	3.00
Large	All	2,962,954	63.97	10.93	21.31	3.79
Medium	All	641,901	69.90	7.54	21.86	0.70
Small	All	291,253	70.34	7.96	21.67	0.03
Large	Large	193,246	93.34	2.57	3.47	0.62
Large	Medium	1,512,654	72.74	8.49	15.66	3.11
Large	Small	1,257,054	48.90	15.15	30.85	5.10
Medium	Large	16,229	90.41	5.13	4.16	0.30
Medium	Medium	251,918	79.74	6.64	13.01	0.62
Medium	Small	373,754	62.38	8.25	28.60	0.78
Small	Large	1617	89.12	8.53	2.35	0.00
Small	Medium	64,715	84.36	6.33	9.30	0.01
Small	Small	224,921	66.17	8.42	25.37	0.04

groups within each firm-size group. The NYSE market share is smallest (48.9%) for small trades in large firms. These market share statistics underscore the fact that the NYSE faces substantial competition for the right to execute trades in NYSE-listed securities. The fact that the NASD dealer market and the regional exchanges attract a larger proportion of small trades is generally consistent with the cream skimming hypothesis if, as seems plausible, small retail orders are less likely to reflect private information. The Cincinnati exchange and the NASD market appear to have further specialized in attracting small trades in the shares of large NYSE-listed firms. An implication of this specialization is that the trades executed on the various exchanges are not random draws from the pool of all trades, but rather vary systematically in dimensions that are potentially important to the determination of trading costs.

2.1. Quote and trade sizes

Table 3 reports average trade and quote sizes by exchange. Quote sizes refer to the maximum trade size for which a posted quote is firm. The largest quote size that can be accommodated by the quotation system used to construct the

Table 3
Quote and trade sizes

Reported are average quotation and trade sizes in shares and dollars for 300 NYSE-listed firms. Each quote size has been averaged across bid and ask sizes.

Quoted:	NYSE	NASD	Regionals	Cincinnati
Panel A: Quote Sizes in Shares				
Firm size				
All	7007.33	459.08	172.53	617.27
Large	7216.32	483.92	166.31	637.60
Medium	5943.54	309.90	172.36	237.68
Small	7773.63	272.42	225.86	127.77
Panel B: Quote Sizes in Dollars				
Firm size				
All	177,318	14,186	5108	25,665
Large	224,076	15,567	6087	26,751
Medium	85,702	5864	3139	5369
Small	41,261	3925	1474	462
Panel C: Trade size in shares				
Firm size				
All	2413.87	990.19	673.43	646.14
Large	2570.49	874.40	648.75	634.45
Medium	2025.41	1458.22	715.90	934.18
Small	1815.54	1631.23	825.95	763.86
Panel D: Trade size in dollars				
Firm size				
All	71,672	24,731	17,506	19,610
Large	86,169	25,949	20,191	19,651
Medium	36,814	23,290	11,152	18,939
Small	13,880	10,715	4765	2840

TAQ database is 99,900 shares. To the extent that market makers sometimes quote larger sizes, the average quote size computed from the TAQ database is biased downward. Also, the quotation data do not indicate whether a quote reflects the interest of a specialist, a dealer, or a limit order. As a consequence, quote size comparisons compare depth available on exchanges rather than quote sizes for specialists versus dealers. The reported quote sizes have been averaged across bid and ask sizes.

NYSE quotations are for much larger sizes than quotes issued on other exchanges. The average NYSE quote is for 7007 shares, or equivalently for $177,000. In contrast, the average NASD quote for NYSE issues is for 459 shares, the average regional quote is for 173 shares, and the average Cincinnati quote is for 617 shares. The average quote size in shares is fairly uniform across firm-size groups on each exchange, while the average quote size in dollars declines across size groups, reflecting the lower average share prices of smaller firms.

NYSE trades are typically larger than trades in NYSE issues executed elsewhere. The mean trade size for NYSE executions for the full sample is approximately 2400 shares, compared to 990 shares for NYSE issues on NASD, 670 shares on the regional exchanges, and 650 shares on the Cincinnati exchange. Though trade sizes on the NYSE exceed those on other exchanges, the differential is far less dramatic than for quote sizes. The average NYSE trade size is only 34% as large as the average NYSE quote size. In contrast, the average Cincinnati trade size is slightly larger than the average Cincinnati quote, the average NASD trade in an NYSE issue is more than twice as large as the average NASD quote, and the average regional trade size is almost four times as large as the average regional quote size.

By posting small average quote sizes, regional and NASD market makers avoid making binding commitments to execute large trades. They do, however, retain the option to selectively execute larger trades. The observation that average trade sizes exceed average quote sizes off the NYSE suggests that the option to selectively execute trades may be used strategically.

3. Empirical measures of trading costs

3.1. Measuring the cost of individual trades

A comprehensive measure of an investor's costs of executing trades is the difference between the actual post-trade value of an investor's portfolio and the value of the portfolio if the investor had been able, instantaneously upon making the trading decision, to transact the desired quantity at a net price equal to the fair value of the asset. Perold (1988) defines such a comprehensive measure, which includes costs related to delays in executing orders, bid–ask spreads, market impact, and commissions or fees. Unfortunately, constructing such a comprehensive measure of trading costs requires access to information that is not generally available, such as the timing of trading decisions. However, several imperfect measures of trading costs can be constructed from the available trade and quote data.

The simplest measure of trading costs is the quoted bid–ask spread. In this study, we report percentage quoted half-spreads, computed as

$$\text{Quoted half-spread}_{it} = 100(A_{it} - B_{it})/(2M_{it}),$$

where A_{it} is the posted ask price for security i at time t, B_{it} is the posted bid price for security i at time t, and M_{it} is the quote midpoint or mean of A_{it} and B_{it}. The quoted spread measures round-trip trade execution costs in percent if trades are executed at the quotes. For comparability to subsequent one-way trading cost measures we report half of the quoted percentage spread. Results for the NASD dealer market are based on the 'inside' quotes posted by NASD market makers.

It should be noted that this measure of trading costs, as well as alternate measures of bid–ask spreads developed below, excludes commission costs, which can also vary across exchanges, but are reported only to traders. Keim and Madhavan (1996) note that some NASD trades are reported on a net basis, with commissions built into the reported trade price. For these trades, the effective half-spread overstates execution costs, though not total trading costs. Unfortunately, these net trades are not separately identified in the TAQ database. Because the frequency with which these trades occur has not been documented, the magnitude of the resulting bias is difficult to gauge.

Many equity trades actually occur at prices inside the posted bid and ask quotes. Let V_{it} denote the true economic value of security i, or the price at which the asset would trade in a frictionless world, just prior to the time t trade. Of course, an observable proxy for V_{it} must be used for estimation. A measure of trading costs that reflects savings due to trading inside the quotes is the percentage effective half-spread:

$$\text{Effective half-spread}_{it} = 100D_{it}(P_{it} - V_{it})/V_{it},$$

where P_{it} is the transaction price for security i at time t and D_{it} is a binary variable that equals one for customer buy orders and negative one for customer sell orders. The effective half-spread measures the closeness of the trade price to the asset's underlying value and thereby estimates the percentage execution cost actually paid by the trader, and gross market-making revenue. In the absence of data on the orders that underlie trades, trade direction must be inferred from the trade and quote data. We categorize trades using the algorithm recommended by Lee and Ready (1991). They also note that trades are often reported with a delay, and recommend comparing trades to quotes in effect five seconds prior to the reported trade time. Hasbrouck et al. (1993) report a median trade reporting delay of 14 s. To be conservative, we compare trade prices to quotes in effect 20 s prior to the reported trade time.

We evaluate the possibility that non-NYSE market makers effectively cream skim uninformed order flow by developing a measure of trades' information content. To do so, we decompose the effective bid–ask spread on each trade into two additive components. Let V_{it+n} denote the asset's post-trade economic value. The two components of the effective spread are defined as

$$\text{Price impact}_{it} = 100D_{it}(V_{it+n} - V_{it})/V_{it},$$

$$\text{Realized half-spread}_{it} = 100D_{it}(P_{it} - V_{it+n})/V_{it}$$

$$= \text{Effective spread}_{it} - \text{Price impact}_{it}.$$

Price impact measures decreases in asset value following customer sells and increases in asset value following customer buys, which reflect the market's assessment of the private information the trades convey. Such price moves constitute a cost to market makers, who buy prior to price decreases and sell prior to price increases. We follow Huang and Stoll (1994) in referring to the second component defined above as the realized half-spread. This component measures the reversal from the trade price to post-trade economic value. As such it measures market-making revenue net of losses to better-informed traders, but before deducting order-processing, inventory, or other costs. Some authors (e.g., Madhavan and Cheng, 1997), refer to 'permanent' and 'temporary' price impacts of a trade. These terms correspond at a conceptual level to the price impact and realized spread measures, respectively, that we discuss. However, Madhavan and Cheng implement the measures by an alternate procedure, using earlier trade prices to proxy for pre-trade economic value. We discuss this implementation issue further in Section 4.3 below.

To operationalize these measures requires the use of observable proxies for economic values. The results we report are based on use of the midpoint of the most recently posted bid and ask quotes as the proxy for pre-trade value, and the first trade price reported at least 24 h after the trade as a proxy for post-trade economic value. We also investigate results obtained when quote midpoints are used instead of transaction prices as proxies for post-trade economic value, and when post-trade prices are observed 30 min rather than 24 h subsequent to the trade. All empirical results, some of which are reported in Bessembinder and Kaufman (1996), are very similar to those reported here. Huang and Stoll (1994) use 5-min and 30-min horizons to estimate post-trade economic value, and report that results are quite uniform across these horizons as well.

A consideration in interpreting these measures of execution costs is that quotes entered at the NYSE and regional exchanges can reflect specialist interest or can reflect limit orders. Trades that 'pick off' stale limit orders incur small or even negative execution costs. We consider it unlikely that executions against stale limit orders affect our empirical results to a substantial degree.[4] Further, any reductions in average trade execution costs that result from the presence of limit orders reflect real cost savings to liquidity demanders. However, since limit order executions cannot be distinguished using the TAQ data, average realized spreads by exchange should be understood to measure market-making revenue net of information costs averaged across all suppliers of liquidity at the exchange, and not necessarily as reflecting specialist revenue.

3.2. Computing average trading costs

The sample includes 3.9 million trades and 10.1 million quotes. The large number of data observations creates some difficulties arising from limitations on data processing and storage capacity. We circumvent these limitations by using a two-stage computational procedure similar to that employed by Huang and Stoll (1996). In the first stage, we analyze trades and quotes on a calendar month basis. The output of this first-stage analysis is a collection of means by month by stock by exchange for variables (quoted half-spread, effective half-spread, realized half-spread, etc.) of interest. The number of trades or quotes used to compute each mean is also recorded.

The second-stage analysis consists of a set of weighted least squares (WLS) regressions of the form

$$C_{ijt} = \alpha_{\text{nys}} I_{\text{nys}} + \alpha_{\text{nas}} I_{\text{nas}} + \alpha_{\text{reg}} I_{\text{reg}} + \alpha_{\text{cin}} I_{\text{cin}} + \varepsilon_{ijt}, \tag{1}$$

where C_{ijt} is the mean observed for security i on exchange j during month t for a trading cost variable of interest (such as the effective half-spread) and the I's are indicator variables that identify exchanges: I_{nys} equals one for observations from the NYSE and zero otherwise, I_{nas} equals one for observations from the NASD market and zero otherwise, etc. The weighting variable for the WLS regression is the number of first-stage observations used to compute each observation on the dependent variable. The resulting α estimates reproduce the sample means by exchange that would have been obtained had it been possible to examine the full dataset in a single stage.

Since the hypothesis that trading costs equal zero is of little practical interest, we do not report standard errors for the estimated means. However, we do report results of testing hypotheses that average trading costs are equal across the NYSE and other exchanges. Statistical inference in this context is complicated by the fact that the regression residuals in Eq. (1) are not independent.

[4] Our conversations with NYSE specialists indicate that limit order traders increasingly recognize the danger of leaving limit orders on the book for substantial periods of time. In particular, limit orders of size posted by institutions are typically monitored in real time and updated as necessary.

In particular, residuals across months for the same firm contain a common component.

To assess statistical significance in this context, we adopt a bootstrap procedure. Efron and Tibshirani (1993) emphasize that valid inference using bootstrap procedures requires the development of a bootstrap probability distribution for the parameters of interest that (i) reflects the properties (such as nonindependence) of the actual regression errors, and (ii) is generated under the null hypothesis. Eq. (1) is estimated by weighted least squares and the residuals are saved. We then create bootstrap samples for each exchange by sampling at random (with replacement) from Eq. (1) residuals. The bootstrap sample for each exchange is drawn from residuals for the same exchange, and bootstrap sample sizes by exchange are constructed to equal the original sample sizes. Eq. (1) is then estimated using the bootstrap sample as the dependent variable, and the resulting coefficient estimates are saved. This bootstrap procedure is repeated 500 times, giving 500 sets of bootstrap coefficient estimates. The bootstrap p-value on the hypothesis that average trading costs are equal across the NYSE and a competing exchange is obtained as the proportion of the 500 bootstrap simulations in which the absolute difference between bootstrap coefficient estimates is as large or larger than the absolute difference in the actual coefficient estimates.

4. Average trading costs by exchange

4.1. Quoted bid–ask spreads

Table 4 reports average quoted half-spreads for the 1994 sample by quoting exchange. Consistent with the large-firm results from the 1987–1991 period that Huang and Stoll (1994) report, the NYSE quotes substantially narrower spreads. For the full sample, quoted half-spreads on the NYSE average 0.575%. Quoted inside NASD half-spreads for NYSE-listed firms average 0.740%. The regional exchanges' quoted half-spreads are greatest, averaging 1.315%. Quoted half-spreads on the Cincinnati exchange average 0.582%, only marginally greater than on the NYSE.

Comparisons of overall means are affected by the distribution of quotes across firm-size categories. As noted above, NASD and Cincinnati quotes for NYSE issues are skewed toward large firms, and quoted spreads are smallest for large firms. Quoted spreads on the NYSE are in fact significantly smaller than on other exchanges, including Cincinnati, within firm-size groups. For example, in the large-firm group, quoted spreads for NYSE-listed firms average 0.321% on the NYSE, compared to 0.533% on the Cincinnati exchange and 0.601% on NASD. The differential becomes more dramatic for small capitalization firms. In the small-firm group, quoted NYSE half-spreads for NYSE-listed firms

Table 4
Quoted half-spreads (%)

Reported are average quoted bid–ask spreads for the 300 NYSE-listed firms described in Table 1. The spreads are stated in percent and, for comparability to subsequent spread measures, have been divided by two. Individual half-spreads are calculated as 100(ask–bid)/(2Mid), where Mid is the average of the ask and the bid. Reported spreads are simple means of all quoted spreads in the indicated category. Bootstrap *p*-values reported in brackets are for the hypothesis that the mean quoted spread does not differ from the mean quoted NYSE spread.

Firm size	Quoted: NYSE	NASD	Regionals	Cincinnati
All	0.575	0.740 [0.000]	1.315 [0.000]	0.582 [0.712]
Large	0.321	0.601 [0.000]	0.750 [0.000]	0.533 [0.000]
Medium	0.796	1.157 [0.005]	1.757 [0.000]	1.427 [0.000]
Small	1.900	3.351 [0.000]	5.071 [0.000]	5.605 [0.000]

average 1.900%, compared to 3.351%, 5.605%, and 5.071% for NASD, Cincinnati, and the other regional exchanges, respectively.

4.2. Effective bid–ask spreads

Many trades are executed at prices superior to the posted bid and ask quotes. As a consequence, quoted spreads overstate the average cost of trade execution. Effective bid–ask spreads provide a more accurate measure of actual execution costs since they account for executions within the quotes.[5] Table 5 reports average effective half-spreads by exchange for the full sample and for firm-size and trade-size subsamples. However, we do not, in this or subsequent tables, report results for firm/trade-size combinations for which the sample contains fewer than 100 trades.

Averaged across all sample trades, the lowest effective half-spread, 0.277%, is observed on the Cincinnati exchange. By comparison, average effective half-spreads are 0.394% for trades executed on the NYSE, 0.413% for trades executed on NASD, and 0.454% for trades executed on the other regional

[5] Effective spreads that are estimated in the absence of data on underlying orders can still provide upward-biased estimates of average execution costs, since these estimates fail to accommodate the possibility that trades are executed at the 'opposite quotes', i.e., sell orders are executed at the ask price or buy orders are executed at the bid price. The magnitude of the bias for a particular exchange is related to the unobserved frequency with which trades at that exchange are executed at the opposite quote. Knez and Ready (1996) use the TORQ (Trades, Orders, Reports, and Quotes) database, which covers 144 NYSE-listed firms for a three-month period during 1991, and report that, among those NYSE trades occurring when the NYSE bid–ask spread is 1/8 of a dollar, 12% are executed at the opposite quote. Similar data for other exchanges is not publicly available.

Table 5
Effective half-spreads (%)

Reported are average effective bid–ask (half) spreads for the 300 NYSE-listed firms described in Table 1. For each trade the effective half-spread is calculated as 100|Price − Mid|/Mid, where Price is the transaction price and Mid is the average of the most recently posted bid and ask prices. Reported spreads are simple means of the effective spreads observed on all trades in the indicated category. Each bootstrap *p*-value reported in brackets is for the hypothesis that the mean effective spread does not differ from that on the NYSE. An asterisk denotes a firm/trade size pair for which the sample contains fewer than 100 observations.

Firm size	Trade size	Traded: NYSE	NASD	Regionals	Cincinnati
All	All	0.394	0.413 [0.018]	0.454 [0.000]	0.277 [0.000]
Large	All	0.239	0.298 [0.000]	0.289 [0.000]	0.264 [0.000]
Medium	All	0.586	0.643 [0.006]	0.683 [0.000]	0.574 [0.806]
Small	All	1.403	1.550 [0.008]	1.597 [0.000]	* *
Large	Large	0.229	0.207 [0.250]	0.237 [0.232]	0.233 [0.720]
Large	Medium	0.226	0.267 [0.000]	0.248 [0.000]	0.235 [0.244]
Large	Small	0.266	0.321 [0.000]	0.315 [0.000]	0.285 [0.012]
Medium	Large	0.419	0.420 [0.920]	0.405 [0.612]	* *
Medium	Medium	0.517	0.481 [0.098]	0.472 [0.014]	0.401 [0.000]
Medium	Small	0.657	0.737 [0.012]	0.750 [0.002]	0.668 [0.888]
Small	Large	0.657	1.150 [0.000]	* *	* *
Small	Medium	1.035	0.850 [0.000]	0.845 [0.000]	* *
Small	Small	1.545	1.704 [0.022]	1.676 [0.036]	* *

exchanges. However, trades are not homogeneous. Some effects of trade heterogeneity are apparent in comparing effective spreads across the New York and Cincinnati exchanges. Although effective spreads averaged across all trades are lower on the Cincinnati exchange, effective spreads are generally similar across the New York and Cincinnati exchanges within firm-size and trade-size subsamples. The apparent Cincinnati advantage is attributable to its degree of specialization in executing trades in large firms.

Within firm-size and trade-size subsamples, effective half-spreads in the NASD dealer market and on the regional exchanges are generally larger than on the NYSE, and most differentials are statistically significant. However, cross-exchange differences in effective spreads are relatively small. For large-firm trades, effective half-spreads average 0.298% on the NASD market, 0.289% on the regional exchanges, and 0.264% at the Cincinnati exchange, as compared to 0.239% on the NYSE. For trades in medium-sized firms, effective half-spreads on the NYSE average 0.586%, compared to 0.643% on the NASD market, 0.574% on the Cincinnati exchange, and 0.683% at the other regionals. For small-firm trades, the average effective half-spread is 1.403% on the NYSE, an advantage of 14.7 basis points over the NASD market and 19.4 basis points over the regional exchanges.

Average effective bid–ask spreads are much more similar across exchanges than are average quoted bid–ask spreads. In part, this reflects the effects of ITS rules requiring that trades be executed at a price no worse than the best quote available on any exchange.[6] It also reflects that trades executed off the NYSE sometimes receive price improvement. Bessembinder and Kaufman (1996) report that the overall percentage of trades executed at a price better than the standing NYSE quote is only slightly higher for trades executed at the NYSE than for trades executed on the NASD market and the regional exchanges.

The observation that average trade execution costs as measured by effective bid-ask spreads are only moderately higher when trades are executed off the NYSE does not necessarily alleviate concerns as to whether trades are routed so as to receive the 'best execution'. If the trades executed off the NYSE involve lower market-making costs, then it might be reasonable to expect that they receive superior rather than similar executions as compared to trades executed on the NYSE. We next investigate whether trades in NYSE issues that are executed off the NYSE involve lower information costs.

4.3. Trades' average price impact: is there evidence of cream skimming?

Easley et al. (1996), among others, suggest that order flow payments target orders originating with relatively uninformed traders. Trades that convey less private information lead to smaller post-trade price movements and, ceteris paribus, are more profitable to execute. The possibility that trades diverted from the NYSE reflect the cream skimming of uninformed order flow can be evaluated by measuring the average information content, as reflected in subsequent price changes, of trades executed on each exchange.

Panel A of Table 6 reports measures of average price impact for trades in NYSE issues executed at each exchange. A striking regularity can be observed: the average price impact of trades in NYSE issues that are executed off the NYSE is uniformly and significantly less than for trades executed on the NYSE. On average, each trade on the NYSE moves subsequent prices by 0.297%. By comparison, trades of NYSE issues on the NASD market move prices by

[6] The Intermarket Trading System is an electronic system linking the NYSE, the American Stock Exchange (AMEX), the five regional exchanges, the Chicago Board Options Exchange, and the NASD dealer market that facilitates trading at the best quote. According to ITS rules, when a market or marketable limit order in an ITS-eligible issue arrives at a market with an inferior quote, the market maker must execute the order at the best ITS quote or send a 'commitment to trade' to the best-quote market where it may be executed. Though the ITS was limited to 11 stocks when it was inaugurated in 1978, by the end of 1994 there were 3293 issues (2817 NYSE-listed and 476 AMEX-listed) in the ITS system, representing most of the exchange-listed stocks eligible for trading on multiple exchanges. See (Hasbrouck et al. 1993, especially pages 27–32) for a more complete discussion of the ITS and ITS trading rules.

Table 6
Average price impact and realized half-spreads

Panel A reports trades' average price impact, while Panel B reports average realized half-spreads. For individual trades price impact is measured as 100indicator(next price − mid)/mid, where indicator is a binary variable that equals one for customer-initiated buy orders and negative one for customer-initiated sell orders, mid is the midpoint of the bid and ask prices prevailing at the time of the trade, and next price is the first transaction price observed at least 24 h subsequent to the trade. The realized spread for each trade is computed as 100indicator(price − next price), where price indicates the trade price, or equivalently as the effective half-spread less computed price impact. Each entry is the simple mean computed across all trades in the indicated category. Each bootstrap p-value reported in brackets is for the hypothesis that the mean does not differ from that for NYSE trades. An asterisk denotes a firm/trade size pair for which the sample contains fewer than 100 observations.

Firm size	Trade size	Traded: NYSE	NASD	Regionals	Cincinnati
Panel A: Price impact					
All	All	0.297	0.166 [0.000]	0.232 [0.000]	0.123 [0.000]
Large	All	0.185	0.119 [0.000]	0.143 [0.000]	0.118 [0.000]
Medium	All	0.486	0.291 [0.000]	0.380 [0.000]	0.247 [0.000]
Small	All	0.928	0.569 [0.000]	0.798 [0.008]	* *
Large	Large	0.238	0.037 [0.000]	−0.002 [0.000]	−0.012 [0.000]
Large	Medium	0.197	0.131 [0.000]	0.146 [0.000]	0.146 [0.000]
Large	Small	0.147	0.113 [0.036]	0.143 [0.760]	0.100 [0.012]
Medium	Large	0.501	0.216 [0.016]	0.298 [0.388]	* *
Medium	Medium	0.529	0.244 [0.000]	0.319 [0.000]	0.324 [0.000]
Medium	Small	0.449	0.318 [0.000]	0.399 [0.052]	0.205 [0.004]
Small	Large	0.561	0.211 [0.300]	* *	* *
Small	Medium	0.964	0.282 [0.000]	0.605 [0.008]	* *
Small	Small	0.918	0.632 [0.000]	0.818 [0.068]	* *
Panel B: Realized half-spreads					
All	All	0.096	0.247 [0.000]	0.222 [0.000]	0.154 [0.000]
Large	All	0.055	0.179 [0.000]	0.146 [0.000]	0.146 [0.000]
Medium	All	0.100	0.352 [0.000]	0.303 [0.000]	0.327 [0.000]
Small	All	0.475	0.983 [0.000]	0.799 [0.000]	* *
Large	Large	−0.009	0.169 [0.000]	0.238 [0.000]	0.241 [0.000]
Large	Medium	0.029	0.136 [0.000]	0.102 [0.000]	0.089 [0.000]
Large	Small	0.119	0.208 [0.000]	0.171 [0.000]	0.185 [0.000]
Medium	Large	−0.083	0.184 [0.016]	0.099 [0.408]	* *
Medium	Medium	−0.012	0.236 [0.000]	0.152 [0.000]	0.076 [0.212]
Medium	Small	0.208	0.419 [0.000]	0.351 [0.000]	0.462 [0.008]
Small	Large	0.075	0.924 [0.016]	* *	* *
Small	Medium	0.070	0.559 [0.000]	0.230 [0.088]	* *
Small	Small	0.627	1.074 [0.000]	0.858 [0.000]	* *

0.166%, trades on the Cincinnati exchange move prices by 0.123%, and trades on the other regional exchanges move prices by 0.232%. The general pattern of smaller price impacts for trades executed on non-NYSE markets is repeated without exception in every firm-size and trade-size subsample. It is important to the interpretation of these results to recognize that trades executed off the NYSE move *all* market prices less than trades executed on the NYSE: these empirical results are invariant when price impact is computed from subsequent trade prices (as reported), from quote midpoints on the same exchange that executes the trade, or from intermarket inside quote midpoints. These observed differences in average price impact across exchanges indicate that the non-NYSE markets attract a relatively high proportion of the order flow originating from uninformed traders.

Keim and Madhavan (1996) note that information about upcoming trades may leak to the market prior to trade execution. If so, some information may be reflected in quote revisions prior to trade execution, and the measure of price impact that we report (which uses the quote midpoint at the time of the trade as the proxy for pre-trade value) will tend to understate the true information flow. They argue that such leakage will tend to be greater for more difficult trades, including trades in smaller stocks and large block trades. In our sample, a larger proportion of large trades and of trades in smaller stocks are executed on the NYSE. Since the NYSE executes a larger portion of the difficult trades, the Keim and Madhavan insights suggest that our measures are likely to understate the actual differential across exchanges in trades' average information content.

Order flow inducements typically target small trades. Our empirical results indicating smaller price impacts for trades executed off the NYSE hold for the (relatively few) large trades executed off the NYSE as well. One possibility is that the large trades executed off the NYSE tend to represent institutional trades directed to upstairs brokers, who then cross the trades at exchanges other than the NYSE. These block trades would tend to be associated with smaller measured price impacts, either because of pre-trade information leakage or because large liquidity traders are more likely than information-based traders to use the upstairs market.

To the extent that the upstairs block market attracts liquidity-motivated trades, the NYSE specialist and trading floor also in effect face cream skimming from those who choose to use the upstairs market. Although Madhavan and Cheng (1997) report that 83% of block volume is still crossed at the NYSE, the specialist and trading floor do not necessarily have the opportunity to act as intermediaries in block trades. However, since the trades completed in the upstairs market are directed there by institutional customers, no obvious concerns regarding best execution are raised.

Easley et al. (1996) also address the issue of trades' information content. They examine signed order arrival frequencies for 30 securities traded at the New York and Cincinnati exchanges, and infer from the order flow data that the proportion of orders originating with privately informed traders is higher on the NYSE. The empirical results reported here are consistent with their results, and because we use transactions prices, provide direct evidence on the magnitude of the difference in trades' average information content across exchanges.

4.4. Realized half-spreads

An examination of realized bid–ask spreads, or equivalently, trades' temporary price impacts, allows for the assessment of market-making revenue net of losses to better-informed traders. Panel B of Table 6 reports average realized half-spreads, which measure price reversals after trades, by exchange. Consistent with the results that Huang and Stoll (1994) report for large firms during the earlier 1987–1991 period, realized bid–ask spreads during 1994 were consistently lower on the NYSE than on other exchanges. We note that this general observation extends to each firm-size and trade-size subgroup. The observed differentials in realized spreads are quite large: realized half-spreads for large-firm and medium-firm trades are on average three times as large for trades executed off the NYSE, while realized spreads for small-firm trades are on average about twice as large for trades executed off the NYSE.

Realized spreads constitute a measure of market-making revenue net of the adverse selection costs that arise from the presence of privately informed traders but before deducting other (inventory, order processing, etc.) costs of market making. The observation that realized spreads are greater for trades executed off the NYSE therefore implies either higher levels of other market-making costs or higher rates of market-making profit off the NYSE. In either case, the question arises whether the trades executed off the NYSE might have received better execution had they not been diverted.

5. Cross-sectional variation in trade execution costs and information content

The empirical results reported in Section 4 indicate that execution costs are greater, while information content is less, for trades executed off the NYSE. We next assess whether cross-sectional variation in trade execution costs and information content is related to observable differences in economic characteristics or to a security's eligibility for off-exchange trading by NYSE member firms.

5.1. Can variation in trades' economic characteristics explain variation in trading costs?

The trades that occur at a given exchange do not generally represent random draws from the pool of all trades. The variation in average trading costs and information content across exchanges documented in Section 4 could reflect variation in the type of trades executed. We assess whether average trading costs vary across exchanges after controlling for variation in a set of economic characteristics. To do so, we augment the WLS regressions described in Section 3.2 to include a set of economic variables, and run a series of regressions of the following form:

$$C_{ijt} = \alpha_{nys}I_{nys} + \alpha_{nas}I_{nas} + \alpha_{reg}I_{reg} + \alpha_{cin}I_{cin} + \gamma X_{ijt} + \varepsilon_{ijt}, \tag{2}$$

where X_{ijt} denotes a vector of economic variables associated with observation C_{ijt}. The economic variables we employ include the beginning-of-sample market capitalization of the traded firm, the inverse of the average transaction price for firm i during month t, the standard deviation of the daily returns for the firm i during month t, average daily trading volume during 1993 for the firm, and the average trade size in dollars for firm i during month t at exchange j. Harris (1994) and others have employed similar economic variables to explain cross-sectional variation in trading costs. Trade execution costs should decrease with trading volume due to economies of scale and with firm size due to enhanced liquidity and reduced information asymmetries. Trading costs should increase with return volatility because market makers are risk averse. Trading costs are known to be related to price levels (in part because minimum allowable price variations constrain quoted bid–ask spreads for lower-priced stocks). Harris has documented that the inverse of share price improves explanatory power as compared to the use of share price. Trade size is expected to be related to trading costs due to correlation with information content or to the inventory holding costs imposed by larger trades. We transform each of the economic variables by deducting the variable's mean computed across all sample observations, and estimate the WLS regression using the transformed variables. We also estimate a strictly cross-sectional version of Eq. (2) on a monthly basis. The resulting intercept estimates are quite stable across months and are quite similar to those obtained using the full year as reported in Table 7. (These monthly estimates are available from the authors on request.)

The intercept coefficients estimated on the indicator variables in this regression allow assessment of whether average trading costs vary across exchanges after adjusting for heterogeneity in economic attributes. Regression intercepts in general estimate the mean of the regression's dependent variable, conditional on each explanatory variable equaling zero. Since we use mean-adjusted economic variables, the estimated intercept coefficients equal average trading costs on each exchange for trades of average size in a firm with economic characteristics equal to sample means. Intercepts from estimating Eq. (2) are reported in Panel A of Table 7, while slope coefficients are reported in Panel B. Panel C of Table 7 reproduces for comparison observed sample means, or the intercepts from estimating Eq. (1), for each trading cost measure. Statistical significance is assessed by bootstrap procedures analogous to those previously discussed.

Estimated slope coefficients are generally of the anticipated signs and most are statistically significant. Measures of trade execution costs decrease with average trading volume and firm size and increase with the inverse of share price and with return volatility. Average price impact increases with mean trade size, while effective and realized spreads decrease with mean trade size.

Table 7
Trading costs and variation in economic attributes

Reported are coefficients obtained in WLS regressions of average monthly trading costs by firm on indicator variables that identify trading exchange and a set of economic variables, including 1993 trading volume (in millions of shares per day), beginning-of-sample market capitalization (in $billions), the inverse of the average trade price during the month, the standard deviation of daily returns during the month, and the monthly average trade size (in units of $100,000). Each economic variable has been transformed by deducting its sample mean. Panel A reports intercept estimates, for which the hypothesis subject to test is that the coefficient does not differ from that estimated for the NYSE. Panel B reports estimated slope coefficients, for which the hypothesis subject to test is that the associated coefficient equals zero. Panel C reproduces observed sample means for comparison.

Dependent variable:	Quoted Half-Spreads (%)	Effective Half-Spreads (%)	Price Impact (%)	Realized Half-Spreads
Panel A: Intercepts				
NYSE	0.475	0.399	0.298	0.101
NASD	0.939**	0.413	0.168**	0.245**
Regionals	1.298**	0.423	0.218**	0.205**
Cincinnati	0.943**	0.388*	0.197**	0.191**
Panel B: Slope coefficients				
Trading Volume	− 0.274**	− 0.080**	− 0.088**	0.008
Market Capitalization	− 0.076**	− 0.055**	− 0.030**	− 0.025**
Inverse Price	10.871**	3.705**	1.559**	2.144**
Return Standard Deviation	0.606**	0.207**	0.332**	− 0.125**
Mean Trade Size		− 0.005**	0.006**	− 0.012**
Panel C: Observed means				
NYSE	0.575	0.394	0.297	0.096
NASD	0.740**	0.413*	0.166**	0.247**
Regionals	1.315**	0.454**	0.232**	0.222**
Cincinnati	0.582	0.277**	0.123**	0.154**

*Bootstrap p-value less than 0.10.
**Bootstrap p-value less than 0.01.

Focusing on the intercept estimates reported in Panel A, quoted half-spreads after controlling for variation in economic variables are smallest on the NYSE, by large margins. The estimated intercept for quoted half-spreads on the NYSE is 0.475%, compared to estimates of 0.939% on NASD, 0.943% on the Cincinnati exchange, and 1.298% on the other regionals. Effective half-spreads for NYSE-listed firms are more uniform across exchanges after controlling for economic variables, equaling 0.399% on the NYSE, 0.413% on the NASD market, 0.388% on the Cincinnati exchange, and 0.423% on the other regionals. Allowing for variation in economic attributes, including mean trade size, does not overturn the conclusion that trades executed on the NYSE contain more information than trades executed off the NYSE. The differentials in price impact are substantial, indicating that the information content of NYSE trades is 37% larger than for regional trades, 51% larger than for Cincinnati trades, and 77% larger than for NASD trades. As a consequence, realized spreads adjusted for variation in trades' economic attributes are smallest on the NYSE. The difference in realized spreads is substantial: coefficients from the augmented second-stage regression for realized spreads are 0.101% on the NYSE, 0.245% for NYSE firms traded on the NASD market, 0.205% for NYSE firms traded on the regionals, and 0.191% at the Cincinnati exchange. On balance, these results support the conclusion that differences in average execution costs and price impact across exchanges are not explained by variation across exchanges in the economic characteristics of trades.

5.2. Execution costs and cream-skimming for 19c-3 stocks

New York Stock Exchange Rule 390 prohibits NYSE member firms from executing trades in NYSE-listed securities off of organized exchanges. The SEC partially abrogated Rule 390 when it adopted Rule 19c-3, which provides that stocks initially listed after April 26, 1979 are not subject to off-exchange trading restrictions, in an apparent effort to foster competition by allowing NYSE member firms to effectively compete with the NYSE specialist.

Our 300-firm sample contains 162 companies that are subject to Rule 19c-3.[7] In Table 8 we report average effective half-spreads, average price impact, and average realized half-spreads by exchange for 19c-3 and non-19c-3 stocks. These results indicate that average effective and realized half-spreads, as well as average price impacts, are substantially larger for 19c-3 stocks than for non-19c-3 stocks on each exchange. However, simple comparisons of means (not reported) indicate that the 19c-3 stocks in the sample are characterized by smaller average market capitalization, lower mean share prices and trading volume, and more volatile returns than non-19c-3 stocks. These differences in economic attributes contribute to larger trade execution costs for 19c-3 stocks and suggest that meaningful comparisons across groups requires controlling for these differences.

To obtain more useful comparisons of trade execution costs across 19c-3 and non-19c-3 stocks, we expand specification of Eq. (2) to include eight indicator variables that identify both trading location and 19c-3 status. The resulting intercept estimates reported in Table 8 indicate that average effective spreads are slightly larger for 19c-3 than for non-19c-3 stocks on the NYSE, the NASD,

[7] We thank George Sofianos of the NYSE for identifying the securities in our sample that are eligible for Rule 19c-3. Hasbrouck et al. (1993) provide a more detailed discussion of SEC Rule 19c-3 and NYSE Rule 390.

Table 8
Comparisons of average trade execution costs for rule 19c-3 stocks

Reported are average trade execution cost estimates for NYSE stocks that are eligible and ineligible for off-exchange trading under the provisions of SEC Rule 19c-3. Sample means reflect simple averages across all trades executed on the indicated exchange. Regression intercepts reflect coefficients obtained in WLS regressions of average trading costs by firm by exchange by month on indicator variables that identify trading location and a set of mean-adjusted economic variables described on Table 7. Hypothesis tests are for the hypothesis that coefficient estimates are equal across 19c-3 and non-19c-3 stocks.

	Sample means				Regression intercepts			
	NYSE	NASD	Regionals	Cincinnati	NYSE	NASD	Regionals	Cincinnati
Panel A: Effective half-spreads								
19c-3 Stocks	0.495	0.503	0.608	0.351	0.420	0.429	0.458	0.375
Non 19c-3 Stocks	0.322**	0.357**	0.367**	0.253*	0.384**	0.401**	0.404**	0.387
Panel B: Price impact								
19c-3 Stocks	0.384	0.192	0.309	0.158	0.327	0.138	0.213	0.156
Non 19c-3 Stocks	0.236**	0.150*	0.189**	0.111	0.278**	0.186**	0.221	0.207*
Panel C: Realized half-spreads								
19c-3 Stocks	0.111	0.310	0.300	0.192	0.093	0.291	0.245	0.218
Non 19c-3 Stocks	0.086**	0.206**	0.178**	0.141	0.106*	0.215**	0.183**	0.180

*Bootstrap p-value less than 0.10.
**Bootstrap p-value less than 0.01.

and the regional exchanges except for Cincinnati. On the NYSE, the price impact of trades in 19c-3 stocks exceeds the price impact of trades in non-19c-3 stocks. In contrast, at the NASD dealer market and at the regional exchanges the price impact of trades in 19c-3 stocks is significantly less than that of trades in non-19c-3 stocks. Realized bid-ask spreads for both 19c-3 stocks and non-19c-3 stocks are greater when trades are executed off the NYSE, but the differential is uniformly greater for 19c-3 stocks than for non-19c-3 stocks. In general, these differentials are most notable when comparing NYSE executions to executions on the NASD market, where many NYSE member firms also maintain dealer operations.

These empirical results indicate somewhat larger average trade execution costs for securities that are eligible for rule 19c-3. The observed differentials in average price impact are consistent with the conclusion that the diversion of uninformed order flow from the NYSE occurs to a greater extent for 19c-3 securities than for non-19c-3 securities. These findings are also consistent with those of Madhavan and Sofianos (1994), who show that specialist participation rates are lower for 19c-3 stocks.

In contrast to the apparent intent of Rule 19c-3, these results provide no support for the rationale that allowing NYSE member firms to execute trades off the NYSE effectively increases competition or reduces trade execution costs. Allowing NYSE member firms to provide an alternative to the routing of orders to the NYSE floor would be more likely to intensify competition if retail customers rather than brokers select the trade location.

6. Summary and conclusions

This paper presents a comprehensive analysis of quotations, execution costs, and information content for trades in 300 NYSE-listed stocks executed during 1994. We find that quotations differ markedly across exchanges. NYSE quotes (including both specialist interest and limit orders) are for relatively large sizes at relatively narrow bid–ask spreads. In contrast, although NASD market makers and those at the regional exchanges post large numbers of quotes, the quotes are on average substantially wider and for relatively small numbers of shares. In posting quotes with smaller sizes and less competitive spreads, the non-NYSE market makers take on less *obligation* to make trades, while retaining the *option* to do so.

We also find that effective bid–ask spreads, which reflect the reductions in trading costs attributable to trades executed within the quotes, are only modestly larger for trades executed off the NYSE. However, the trades in NYSE issues that are transacted off the NYSE contain less information, as measured by their impact on subsequent market prices, than trades executed on the NYSE, implying that non-NYSE market makers are able to cream skim a subset of trades that can be processed at lower cost. As a consequence, realized bid–ask spreads, which measure price reversals after trades and market-making revenue net of information-related costs, are lower by a factor of two to three for trades executed on the NYSE. Controlling for variation in economic characteristics, such as firm size, share price, return volatility, trading volume, and average trade size does not fully explain these cross-exchange differences in trade execution costs or price impact. Interestingly, trade execution costs are on average slightly larger, and there is stronger evidence of cream skimming for securities eligible for rule 19c-3 (which allows NYSE member firms to compete with the NYSE specialist by executing trades in the OTC market) than for non-19c-3 securities.

This set of empirical findings raises two main concerns regarding the trading of NYSE-listed securities off the NYSE. The first is immediate. Trade location is typically selected by brokers rather than customers, giving rise to potential agency costs. The observation that realized bid-ask spreads, which measure market-making revenue net of information-related costs, are substantially greater for trades executed off the NYSE implies either higher non-information market-making costs or higher rates of market-making profit off the NYSE. In either case, the questions arise as to whether the trades currently being diverted from the NYSE might receive better execution if they were not diverted and whether existing rules governing order flow effectively foster competition.

A second, longer-term, concern stems from the observation that market makers on the NASD market and the regional exchanges are able to cream skim a subset of trades that contain relatively little information. Such cream skimming can worsen the adverse selection problem faced by NYSE market makers, potentially increasing bid–ask spreads on the NYSE. If so, execution costs could be increased off the NYSE as well, since ITS rules and order flow agreements generally require only that orders be executed at prices no worse than the best intermarket quote, which is often the NYSE quote. The cream skimming of uninformed trades away from the NYSE could have the salutary effect of facilitating price discovery by allowing NYSE market makers to learn about market clearing prices faster. However, in a segmented market in which price discovery occurs primarily on the NYSE while liquidity trades are diverted elsewhere, NYSE market makers may be unable to capture the benefits of improved price discovery. Glosten and Milgrom (1985) and Milgrom and Stokey (1982) have shown that if the concentration of informed trades becomes too large, the bid–ask spread may have to be widened to the point where the market fails entirely. It is unclear what degree of segmentation of informed versus liquidity trades can be sustained in equilibrium.

Some limitations of our empirical analysis should be noted. We cannot observe the executions that individual trades would have received at alternate locations, nor can we compare the current commission structure to that which would have been obtained in the absence of order flow agreements or preferencing arrangements. However, the results we report are consistent with the reasoning that orders are not always routed so as to receive the best possible executions and that the cream skimming of uninformed order flow occurs. If so, average trade execution costs may be greater than they need be, and the potential for future market instability exists.

Acknowledgements

The authors thank the New York Stock Exchange for partial financial support and for the provision of data for our earlier study (Bessembinder and Kaufman, 1996). We also thank Robert Battalio, Ed O'Malley, Bill Schwert (the editor), George Sofianos, Hans Stoll, and an anonymous referee, as well as participants at the NYSE conference 'The Search for the Best Price' and a University of Oklahoma research seminar, for valuable comments.

References

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

Battalio, R., Greene, J., Jennings, R., 1995. Do competing specialists and preferencing dealers affect market quality? Unpublished working paper. Department of Finance, Indiana University, Bloomington.

Bessembinder, H., Kaufman, H., 1996. Quotations and trading costs on the domestic equity exchanges. Unpublished working paper. Department of Finance, Arizona State University, Tempe.

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

Chordia, T., Subrahmanyam, A., 1995. Market making, the tick size, and payment-for-order flow: theory and evidence. Journal of Business 68, 543–575.

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

Efron, B., Tibshirani, R, 1993. An Introduction to the Bootstrap. Chapman & Hall, New York.

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

Harris, L., 1993. Consolidation, fragmentation, segmentation and regulation. Financial Markets, Institutions, and Instruments 2, 1–14.

Harris, L., 1994. Minimum price variations, discrete bid–ask spreads, and quotation sizes. The Review of Financial Studies 1, 148–178.

Hasbrouck, J., Sofianos, G., Sosebee, D., 1993. New York Stock Exchange systems and trading procedures. Unpublished Working Paper 93-01. New York Stock Exchange.

Huang, R., Stoll, H., 1994. Competitive trading of NYSE-listed stocks: measurement and interpretation of trading costs. Unpublished Working Paper 94-13. Owen Graduate School of Management, Vanderbilt University, Nashville.

Huang, R., Stoll, H., 1996. Dealer versus auction markets: a paired comparison of execution costs on NASDAQ and the NYSE. Journal of Financial Economics 41, 313–358.

Keim, D., Madhavan, A., 1996. Execution costs and investment performance: an empirical analysis of institutional equity trades. Unpublished working paper. Rodney L. White Center, The Wharton School, Philadelphia.

Knez, P., Ready, M., 1996. Estimating the profits from trading strategies. Review of Financial Studies 9, 1121–1164.

Lee, C., 1993. Market integration and price execution for NYSE-listed securities. Journal of Finance 48, 1009–1038.

Lee, C., Ready, M., 1991. Inferring trade direction from intraday data. Journal of Finance 46, 733–746.

Macey, J., O'Hara, M., 1996. The law and economics of best execution. Unpublished working paper. Johnson Graduate School of Management, Cornell University, Ithaca.

Madhavan, A., Cheng, M., 1997. In search of liquidity: block trades in the upstairs and downstairs markets. Review of Financial Studies 10, 175–204.

Madhavan, A., Sofianos, G., 1994. Auction and dealer Markets: an empirical analysis of NYSE specialist trading. Unpublished Working Paper #94-01, New York Stock Exchange.

Milgrom, P., Stokey, N., 1982. Information, trade, and common knowledge. Journal of Economic Theory 26, 17–27.

Perold, A., 1988. The implementation shortfall: paper versus reality. Journal of Portfolio Management 14, 4–9.

Petersen, M., Fialkowski, D., 1994. Posted versus effective spreads. Journal of Financial Economics 35, 269–292.

United States Securities and Exchange Commission, 1996. Order Execution Obligations, Release No. 34-37619A, Washington, DC.