Public Disclosure and Private Decisions: Equity Market Execution Quality and Order Routing

Ekkehart Boehmer
Texas A&M University

Robert Jennings
Indiana University

Li Wei
New York Stock Exchange

In 2001, the Securities and Exchange Commission (SEC) required market centers to publish monthly execution-quality reports in an effort to spur competition for order flow between markets. Using samples of stocks trading on several markets, we investigate whether past execution quality affects order-routing decisions and whether the new disclosure requirements influence this relationship. We find that routing decisions are associated with execution quality; markets reporting low execution costs and fast fills subsequently receive more orders. Moreover, the reports themselves appear to provide information that was unavailable previously. Our results are consistent with active competition for order flow that can be influenced by public disclosure. (*JEL* G24, G28, K22)

The U.S. Securities and Exchange Commission (SEC) frequently relies on public disclosure to achieve policy objectives. In defining themselves, the Commission states that, ". . . the SEC is concerned primarily with promoting disclosure of important information, enforcing the security laws, and protecting investors . . ."[1] A theme appearing repeatedly in SEC activities is that well-informed individuals make decisions enhancing security-market efficiency. Recently enacted SEC Rule 11Ac1-5 illustrates this approach. Equity-market trades frequently occur at prices other than those quoted, but brokers/traders find it difficult to anticipate

---

We thank an anonymous referee, Bruno Biais, Beatrice Boehmer, Jay Coughenour, Joel Hasbrouck, Craig Holden, Pamela Moulton, Maureen O'Hara (the editor), Gideon Saar, Erik Sirri, Kumar Venkataraman, Lei Yu, and the seminar participants at the New York Stock Exchange's Roundtable on Execution Quality, the 2003 Financial Management Association meeting, the 2005 Western Finance Association meeting, Texas A&M University, and the University of Texas at Austin for very helpful comments. The opinions expressed in this article do not necessarily reflect the opinions of the members, officers, or employees of the New York Stock Exchange. Address correspondence to: Ekkehart Boehmer, Department of Finance, Mays Business School, Texas A&M University, College Station, TX 77843-4218, or e-mail: eboehmer@mays.tamu.edu.

[1] See the SEC website at http://www.sec.gov/about/whatwedo.shtml.

doi:10.1093/rfs/hhl011 Advance Access publication June 30, 2006

the execution quality they will obtain in different market centers. Historically, the lack of standardized execution quality statistics on submitted orders contributed to this uncertainty. On November 17, 2000, the SEC mandated that U. S. market centers publish a broad set of standardized execution-quality metrics each month. The goal of this action was to "... empower market forces with the means to achieve a more competitive and efficient national market system for public investors."[2] In this article, we examine whether and how brokers/traders appear to use these data when making order-routing decisions. Finding that public reports and the attention they create influence security-market participants' actions would support the SEC's reliance on disclosure as a policy tool, at least in this situation.

A substantial literature compares execution quality and market-maker quoting behavior across U. S. equity markets. Relatively little is known, however, about whether order-routing decisions by traders or by firms with the fiduciary responsibility of executing clients' orders are sensitive to execution quality. Traders and brokers have many options when executing customer orders. During our sample period, orders for actively traded securities listed on the New York Stock Exchange (NYSE) can be submitted to the NYSE, five regional stock exchanges, several Nasdaq Intermarket dealers, and eight Electronic Communications Networks (ECNs). Our analysis investigates how these routing decisions are made. Specifically, we examine how a market center's historical execution quality affects subsequent order-routing decisions. Our results help assess whether these routing decisions are consistent with competition for order flow based on execution quality. This is an important public policy issue, because it affects both investors' trading costs and the operational efficiency of the equity market. More broadly, our analysis addresses whether mandatory public disclosure and the scrutiny it naturally focuses on the disclosed information affects market participants' behavior.

Rule 11Ac1-5 ("Dash-5") reports provide important details on execution quality and increase the visibility of quantitative execution-quality measures. The rule is designed to influence the actions of investors, brokers, and market centers. Most investors delegate order-routing decisions to brokers, who have a fiduciary duty to obtain best execution for their customers' orders.[3] A companion rule, 11Ac1-6, requires brokers to summarize their routing practices and describe preferencing arrangements, payments they receive for order flow, and the extent to which

[2] See SEC Release No. 34-43590.

[3] Comprehensive data on the fraction of orders providing explicit routing instructions do not exist. Market participants indicate that most orders do not contain such instructions. Our review of SEC-mandated Rule 11Ac1-6 reports for several brokers finds that over 90% of orders are not directed to specific market centers. Although institutional orders might be more likely to be directed, industry sources suggest that about 80% leave the routing decision to the broker.

they internalize orders.[4] Customers (and regulators) can use these 11Ac1-6 reports to identify brokers routing orders to high-cost market centers. Thus, we expect brokers to rely on the standardized execution-quality statistics provided in Dash-5 reports, because the reports provide brokers a rationale for routing decisions that they can easily justify to clients and regulators.

Anecdotal evidence suggests that Dash-5 reports play an important role in industry practice, that a market center's reported execution quality affects its future order flow and that market centers respond by adapting their trading protocols. Several investment firms claim to use Dash-5 reports in order-routing decisions. *Traders Magazine* (2002) reports that "Goldman (Sachs & Co.) now makes its order-routing decisions based on new execution quality statistics supplied by market centers under Rule 11Ac1-5." Paul Wigdor of Pershing Trading Company says, "We explain our [11Ac1-5] statistics in a way that illustrates our comparative advantage." In *Traders Magazine* (2003), Mary McDermott-Holland, Franklin Portfolio Associates' head trader, indicates that, although she believes there are limitations, her firm uses Dash-5 statistics to route orders. Moreover, the Chartered Financial Analyst Institute suggests that buy-side firms use the statistics when routing orders. This focus on execution quality might influence market centers to institute changes in trading practices that improve their Dash-5 statistics. In *Traders Magazine* (2005), the largest Nasdaq market maker notes that decimal pricing and the Dash-5 statistics forced traditional market makers to automate their trading processes. The Boston Stock Exchange cites Dash-5 for changes in its competing specialist system (SEC Release No. 34-45791), and the Chicago Stock Exchange granted its specialists more pricing flexibility in an effort to remain competitive in the battle for order flow (*Traders Magazine*, 2001). Finally, the NYSE's March 2003 decision to automate quotations and the associated (ongoing) modifications to its direct + automatic execution system might have been influenced by the relatively slow execution speed documented in the reports.

These arguments suggest that the SEC's focus on quantitative execution-quality measures might pressure the industry to more seriously consider execution quality. Specifically, brokers/traders have economic incentives to consider Dash-5 reports in making routing decisions, and markets offering execution services have incentives to respond competitively by trying to improve execution quality for orders covered in these reports. Therefore, implementing Dash-5 could change the way order-routing

[4] Preferencing is the practice of routinely designating a particular market center as a destination in exchange for either cash payments (payment for order flow) or noncash reasons (e.g., soft dollars). Internalization is the practice of using a market center with a business affiliation to the broker (e.g., many broker-dealers act as specialists on exchanges or market-makers on Nasdaq).

decisions are made. Specifically, we expect market centers with better Dash-5 execution quality to receive more order flow.[5]

In this article, we develop an econometric model of order-routing behavior for marketable orders in NYSE-listed stocks and test whether an execution venue's market share, the outcome of brokers' routing decisions, is related to past execution quality. We first conduct an event study examining how the influence of execution quality on order-routing decisions changed with the imposition of Dash-5. Using effective spreads computed from the NYSE's Trade and Quote (TAQ) data, a standard measure of execution costs that is available both before and after the implementation, we show that the sensitivity of market share to execution costs increases after Dash-5. Although it is difficult to infer a causal relationship, this result suggests that the mandatory disclosure of Dash-5 statistics gives brokers an additional incentive to consider quantitative measures of execution quality when making their routing decisions.

Compared to previously available data on execution quality, Dash-5 reports provide new information. They contain statistics on execution speed and a finer categorization of order flow and execution costs across market centers, order types, and order sizes. Our main analysis suggests that this additional information is useful for the typical routing decision. Market centers tend to lose order flow if their Dash-5 execution cost or time to execution is high relative to the competition. We control for several factors that could affect routing decisions, including TAQ execution-cost measures. These results suggest that U. S. equity market participants are subject to competitive pressure to maintain high-quality executions. Brokers appear to respond to differences in market quality, perhaps because it helps them in the competition for clients. Therefore, market centers can compete for order flow on price and speed. However, poor relative performance on Dash-5 reports does not imply an immediate exit from the industry. At least in the medium term, some order flow does not appear to respond to differences in Dash-5 execution quality. This lack of order-flow sensitivity to effective spreads and execution speed could arise because some markets compete on dimensions that spreads and speed do not capture. For example, some markets might specialize in executing difficult orders or make payments to brokers for order flow.

Finding that Dash-5 reports contain incremental information is not trivial, because the reports have several limitations. They cover only about one-third of total order flow and are published with a one-month

[5] In the remainder of this article, we refer to "brokers' routing decisions" and "routing decisions" with the understanding that most are indeed made by brokers but that some investors make routing decisions themselves. Because these tend to be larger investors, they are likely to have the means to use Dash-5 reports, among other resources, to guide their decisions. Therefore, the arguments and inferences we make regarding brokers' decisions generally also apply to individual traders who make their own routing decisions.

lag. Critics argue that Dash-5 reports are of limited use because they are costly to produce, not audited, and sensitive to alternative ways of aggregating the underlying order data. Two market centers, Instinet, LLC and Inet ATS, Inc., were recently fined by the SEC for alleged inaccuracies in their Dash-5 data (see SEC Release No. 2005-151; Burns and Lucchetti, 2005). This event has important implications for our research. First, it illustrates that some reports might contain inaccurate descriptions of execution quality. We address this concern in two ways. We exclude Instinet and Inet (Island) from our analysis. Moreover, we provide a detailed comparison of Dash-5 execution quality with TAQ-based execution quality in the Appendix. In this analysis, we find no evidence of systematic inaccuracies. The second implication of the SEC enforcement action is that Dash-5 reports are indeed policed by regulators, which increases our confidence in the reports provided by other market centers in our sample.

Our work complements a large body of research on U. S. equity market execution quality and intermarket competition. Among many others, papers by Blume and Goldstein (1992), Lee (1993), Easley, Kiefer, and O'Hara (1996), Huang and Stoll (1996), Bessembinder and Kaufman (1997), U.S. Securities and Exchange Commission (1997), Bessembinder (1999), U.S. Securities and Exchange Commission (2001), Bacidore, Ross, and Sofianos (2003), and Boehmer (2005) provide information about execution quality on the NYSE, Nasdaq, and the regional stock exchanges. Chung, Chuwonganant, and McCormick (2004) examine the relation between directed order flow and execution costs. These papers discuss measures of execution quality, compare differences in execution quality across markets, and speculate about the causes of those differences. Other studies, for example, Bessembinder (2003b) and Huang (2002), investigate the quoting behavior in the market for NYSE-listed and Nasdaq stocks, respectively. Bessembinder (2003b) finds that quoting competitiveness has an important relation to *ex post* execution costs in NYSE-listed stocks. Specifically, he shows that Nasdaq Intermarket execution costs tend to be higher only when its quotes are not competitive. When quotes are competitive, the differences in execution quality are negligible. This finding suggests that execution costs vary among markets and over time and raises the question whether participants consider such variations when making order-routing decisions. Data limitations have made it difficult to answer this question previously. Our article contributes novel evidence on this issue by focusing on how routing decisions respond to variation in execution quality at different market centers.

The remainder of this study is organized as follows. In Section 1, we discuss the sample and provide descriptive statistics from Rule 11Ac1-5 reports. In Section 2, we characterize the Dash-5 reports' execution-quality statistics. In the Appendix, we extend this discussion and provide a comparison of execution-cost measures from TAQ and Dash-5. We outline our research design and econometric model in Section 3. In Section 4, we conduct an event study to investigate how the implementation of Rule 11Ac1-5 affects order-routing behavior. The results of the main analysis of order-routing decisions are presented in Section 5. We discuss the consequences of increased competition in Section 6 and provide conclusions in the final section.

1. Sample and Descriptive Statistics

We begin with 2561 NYSE-listed securities not classified as preferred stocks, warrants, rights, derivatives, or "other" securities in the monthly NYSE master file. We use each monthly file between June 2001 and June 2004 and require that a security remains available for at least 12 consecutive months. To obtain a more homogeneous set of securities, we use the Center for Research in Security Prices (CRSP) share code and delete firms not incorporated in the United States, closed-end funds, units, shares of beneficial interest, American Depository Receipts (ADRs), certificates, and 51 firms that do not appear on CRSP. Finally, because variables such as bid-ask spread, volume, and volatility are sensitive to share price levels, we exclude three stocks trading above $1000 per share: Berkshire Hathaway Class A and B shares and Security Capital Group. For the remaining 1435 stocks, we obtain trade and quote information from the NYSE's TAQ data from June 2000 to June 2004 and Dash-5 reports from June 2001 to June 2004.[6]

To select the most active market centers for the analysis, we begin with all market centers reporting executions of marketable orders in at least one of the 1435 stocks. Figure 1 shows the 19 most active off-NYSE market centers based on the aggregate volume of marketable orders from Dash-5 reports between June 2001 and June 2004 (the NYSE receives 81% of marketable orders and is omitted from the graph to improve readability). The horizontal bars indicate each market's share of marketable Dash-5 order volume and show that market shares are heavily skewed, even without the NYSE. Gray bars indicate market centers that we exclude from the analysis.[7] Panel A of Table 1 provides descriptive statistics for the subset of the 1435 stocks traded on these markets. For

[6] We do not include Nasdaq -listed stocks in our sample. Chung, Chuwonganant, and McCormick (2004) estimate that 60–80% of Nasdaq order flow is subject to preferencing arrangements. In such an environment, Dash-5 is less likely to have an impact than in markets where more order flow is open to competition. In addition, we feel that listed stocks offer brokers a wider range of market-center business models (ranging from traditional exchanges to electronic limit order books) from which to choose, than do Nasdaq stocks.

[7] We eliminate CAES and its successor Super Intermarket. These execution systems are available only to Nasdaq broker-dealers, who must report the orders themselves even if using CAES/SIMT. As a result, all CAES/SIMT Dash-5 reports cover orders reported elsewhere. We also exclude Instinet and Island ECN, because a recently settled SEC enforcement action alleges inaccurate reporting by these two market centers during our sample period (see SEC Release 2005-151). We obtain qualitatively identical results when Island and Instinet remain in the sample.

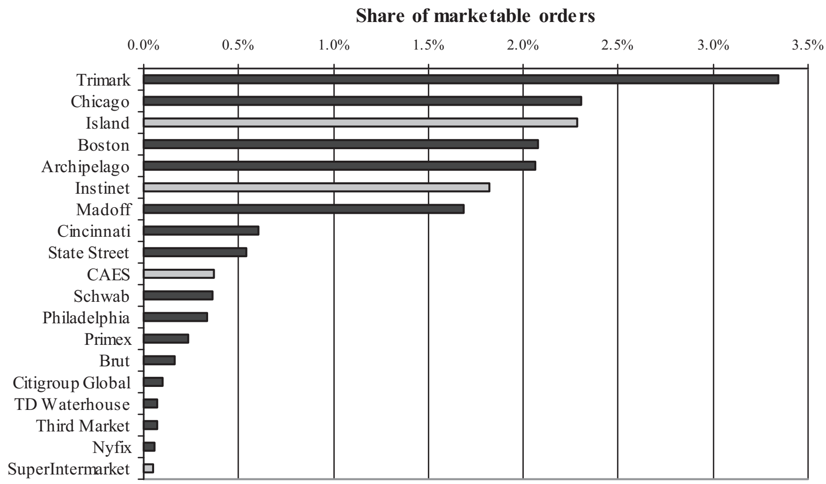

Figure 1
Most active market centers trading NYSE-listed stocks
The sample covers orders between June 2001 and June 2004. For the set of 1435 NYSE-listed stocks that meet our criteria (see Table 1), the figure shows the 20 market centers that receive the largest volume of marketable orders. To improve readability, we omit the NYSE (which receives 81% of such volume) from this graph. Gray bars indicate market centers that we exclude from the sample.

each market center, the table contains the number of stocks for which marketable orders are submitted in at least one month, the daily average closing price, the daily price range relative to the closing price, daily consolidated trading volume in shares and trades, and market capitalization for securities trading at that venue. As in Bessembinder (2003a), we find that most markets (except NYSE and Trimark) concentrate trading on high-volume, high-capitalization stocks.[8]

1.1 Sample selection

To select a sample of stocks from our 1435-stock universe, we take two approaches that balance the number of included market centers and stocks. As our tests focus on routing decisions when brokers have several choices, we believe it is important to select a sample of securities that trade in several markets simultaneously. Figure 1 shows a natural break in market share between the sixth and seventh most active market centers (Madoff and the Cincinnati Stock Exchange). Our first sample ("small sample") uses all 258 securities trading continuously on each of the top six markets over the 35-month sample period. The six markets include a

[8] This is not accidental. For example, Madoff Securities allows only certain brokers to submit orders, limits the securities these brokers can trade via Madoff, and, to a lesser extent, limits the types of orders they can use. This reduces Madoff's exposure to traders who might have private information and makes it less likely that Madoff buys a security just before the price falls and/or sells a security just before a price rises.

Table 1
Descriptive statistics for sample securities

Market center	Number of stocks traded for at least one month	Average daily closing price ($)	Average daily price range (% of closing price)	Average daily volume (shares)	Average daily number of trades	Average market capitalization ($ million)
Panel A: all 1435 sample stocks						
NYSE	1435	27.59	3.40	1,018,840	917	6523
Trimark	1434	27.27	3.39	1,005,603	905	6481
Chicago Stock Exchange	1355	28.19	3.48	1,212,502	1088	7757
Boston Stock Exchange	1012	31.28	3.55	1,843,428	1544	12,024
Archipelago ECN	1427	29.10	3.45	1,196,373	1049	7670
Madoff	528	32.46	3.56	2,503,733	1932	16,935
Cincinnati	926	32.64	4.29	3,202,918	2218	22,033
State Street	1362	31.04	3.32	1,398,371	1224	9583
Schwab	1287	28.63	2.87	1,476,730	1238	6734
Philadelphia	888	32.35	3.62	2,657,477	2014	18,610
Primex	453	32.32	2.65	2,578,256	2193	17,035
Brut	1310	32.30	3.63	2,063,203	1764	12,856
Citigroup Global	851	38.43	3.18	2,306,896	1821	18,760
TD Waterhouse	70	33.70	2.51	7,063,887	4599	55,336
Third Market	549	31.79	5.55	4,089,598	2549	27,289
Nyfix	723	37.72	3.50	3,077,244	2436	22,492
Panel B: final samples						
Small sample:						
258 stocks, six most active market centers	258	35.26	3.6	3,120,660	2342	23,868
Large sample:						
1016 stocks, at least three out of sixteen market centers	1016	29.26	3.4	1,175,315	1066	7921

ECN, Electronic Communications Network; NYSE, New York Stock Exchange.
The table shows sample means for the period from June 2001 to June 2004. The sample is constructed from all 2561 securities classified as common stocks in the NYSE master file. We merge this set of securities with the CRSP header file, resulting in 2510 matches. Based on the CRSP share code, we further delete firms incorporated outside of the United States, closed-end funds, units, shares of beneficial interest, certificates, and ADRs. Finally, we exclude three stocks trading above $1000. This procedure leaves 1435 securities. In Panel B, the first sample includes 258 stocks out of these 1435 that are continuously traded on the six most active market centers. The second sample includes 1016 out of these 1435 that, in each month, are traded on at least three out of the 20 most active market centers. Without affecting our results, we exclude Island and Instinet from the list of the top 20 because of potentially unreliable data. We also exclude CAES and its successor SuperIntermarket from the Top 20, because their executions are reported by other market centers as well.

variety of market structures: the NYSE, two regional exchanges (Boston and Chicago), the Nasdaq Intermarket broker-dealers Madoff and Trimark, and one ECN (Archipelago). As we expect, panel B of Table 1 summarizes that this sample consists of securities with above-average market capitalization and trading activity.[9]

Our second approach is designed to include a larger number of stocks. It relaxes the requirement that stocks trade on the same market centers throughout the entire sample period. Instead, we require that each security trade on at least three of the top 16 sample market centers in a given month. To ensure data availability, we consider only stock market center combinations that survive for at least three consecutive months. The 1016 securities meeting this requirement form our second sample ("large sample"). Although the set of competing market centers changes monthly, Table 1 summarizes that the large sample includes lower capitalization, less actively traded stocks than the small sample and is therefore more representative of the population. Both samples cover the most Dash-5 marketable order volume during our sample period.[10]

Dash-5 reports distinguish order types with varying degrees of marketability, but effective spreads and execution speed are reported only for market orders and marketable limit orders. Market orders instruct the broker to trade immediately at the best available price, while limit orders allow clients to specify a price limit. A limit order is considered marketable if, based on published quotes at the time of order submission, its limit price makes it immediately executable. That is, a buy (sell) marketable limit order has a limit price greater (less) than or equal to the current ask (bid) price.[11] Statistics on these orders are divided into four order-size categories: 100–499 shares, 500–1999 shares, 2000–4999 shares, and 5000–9999 shares.

We impose two additional filters on monthly records. We delete observations where the average monthly price is below \$1 and where the monthly effective spread exceeds one-half the monthly average price. Shares priced below \$1 are subject to delisting and might trade differently than the typical share. The second condition eliminates potential data

[9] In August 2002, Archipelago completed the migration of NYSE-listed securities from Archipelago ECN to Archipelago Exchange, which prints trade reports in the name of the Pacific Stock Exchange. We aggregate the Dash-5 reports from the ECN and the Exchange when both are available on a share-weighted basis and include them with Nasdaq markets when computing statistics based on TAQ.

[10] During the sample period, all market centers together receive market and marketable limit orders for 631 billion shares in the 1435 stocks from which we select. The six markets in the small sample receive orders for 585 billion shares, and the 16 markets in the large sample receive orders for 601 billion shares. The final small sample of 258 stocks represents 341 billion shares, and the final large sample of 1016 stocks represents 532 billion shares.

[11] Orders with special handling instructions, orders that are not submitted electronically, orders for 10,000 or more shares, and limit orders with prices more than 10 cents worse than the relevant quote are excluded from Dash-5. See http://www.sec.gov/rules/final/34-43590.htm for a detailed description of rule 11Ac1-5 and the data contained in the associated monthly reports.

errors in the Dash-5 reports. These two filters together eliminate about 0.7% of the sample observations.

1.2 Measures of execution quality

We focus on two measures of execution quality: round-trip effective spreads and execution speed. Effective spreads measure the noncommission out-of-pocket costs a trader incurs and can be interpreted as the total price impact of a trade. For buy (sell) orders, Dash-5 effective spreads are twice the (negative) difference between the execution price and the National Best Bid and Offer (NBBO) quoted spread midpoint prevailing at the time an order is received at a market center.[12] Effective spreads can be decomposed into a permanent and a temporary component. Realized spreads (the temporary component) exclude the effects of the information content of order flow. For buy (sell) orders, the realized spread is twice the (negative) difference between the execution price and the NBBO quote midpoint five minutes after the trade. Finally, price impact (the permanent component) is defined as the change in the quote midpoint from order receipt to five minutes after execution, or half the difference between effective and realized spreads. It approximates the information component of an order and thus reflects its difficulty. Finally, execution speed is defined as the time between order receipt and execution.

Table 2 provides descriptive statistics for both samples. We report share-weighted execution quality measures for each of our sample market centers, separately for market orders and marketable limit orders. We also report each market's share of all Dash-5 orders for each sample and order type. The NYSE is the dominant market and receives about 91% of Dash-5 marketable orders in both samples (of which 35% are in the form of market orders and 56% in the form of marketable limit orders). Trimark and Madoff have the next largest market shares for the small sample (about 2%, mostly from market orders), while Trimark and Chicago have the next largest for the large sample (about 3%). We observe substantial variation in effective spreads, price impacts, and execution speeds across markets. As in extant studies, we find that NYSE orders generally have more information content than orders arriving at the regional stock exchanges and Nasdaq market makers. In addition, Archipelago also appears to be a destination for informed orders. Moreover, comparing the two samples confirms that stocks in the large sample tend

[12] We construct the NBBO as the best quote (highest bid price and lowest offer price) among all markets quoting the stock. For this computation, we first record all valid quotes for each market center throughout the trading day and then find the highest bid price and lowest offer price at each point in time. A market center's quote is invalid if it has suspended trading in the security, if the market center's bid (offer) price is equal to or greater (less) than the national best offer (bid) price, or if zero depth is posted. The midpoint is one-half the sum of the bid and offer prices.

Table 2
Descriptive statistics on execution quality for marketable orders

Market center	Market orders				Marketable limit orders			
	Market shareof Dash-5 order volume (%)	Dash-5 effective spread ($)	Dash-5 price impact ($)	Dash-5 execution speed (seconds)	Market share of Dash-5 order volume (%)	Dash-5 effective spread ($)	Dash-5 price impact ($)	Dash-5 execution speed (seconds)
Panel A: 258 stocks trading continuously on six market centers								
New York Stock Exchange	35.05	0.038	0.015	18.3	55.70	0.017	0.009	25.2
Boston	1.34	0.027	0.004	18.3	0.33	0.021	0.003	43.8
Chicago	1.08	0.030	0.005	18.4	0.67	0.027	0.010	32.8
Archipelago	0.09	0.045	0.019	44.7	1.63	0.023	0.015	12.3
Madoff	1.76	0.019	0.004	4.7	0.24	0.016	0.004	32.6
Trimark	1.44	0.026	0.006	13.8	0.69	0.016	0.005	44.9
Panel B: 1016 stocks trading on at least three market centers in each month								
New York Stock Exchange	31.46	0.043	0.018	19.0	58.88	0.020	0.011	32.3
Boston	0.83	0.028	0.005	19.3	0.26	0.022	0.004	49.4
Chicago	1.58	0.038	0.008	23.4	1.18	0.030	0.013	52.2
Archipelago	0.07	0.049	0.020	48.0	1.07	0.025	0.016	14.6
Madoff	1.63	0.020	0.004	4.9	0.28	0.016	0.005	37.9
Trimark	2.04	0.032	0.009	15.8	1.24	0.020	0.007	68.1
Cincinnati	0.26	0.023	0.004	19.0	0.14	0.015	0.001	49.5
Philadelphia	0.24	0.031	0.005	29.0	0.10	0.026	0.008	41.9
Brut	0.00	0.030	0.003	16.9	0.14	0.020	0.013	6.0
Nyfix	0.01	0.025	0.002	51.1	0.03	0.022	0.006	110.5
Third Market	0.07	0.053	0.013	31.1	0.04	0.020	0.001	54.9
State Street	0.40	0.037	0.006	39.8	0.21	0.024	0.001	35.4
Schwab	0.54	0.022	0.004	8.6	0.20	0.013	0.002	59.4
Primex	0.00				0.55	0.011	0.001	0.0
Citigroup Global	0.12	0.017	0.000	0.3	0.00			
TD Waterhouse	0.60	0.016	0.002	13.4	0.18	0.014	0.003	34.4

The sample covers orders in NYSE-listed common stocks between June 2001 and June 2004 and is described in Table 1. The table reports share-weighted means for marketable orders based on SEC Rule 11Ac1-5 reports. We impose two filters on monthly records. We delete observations if the mean daily closing price for the month is less than $1 or if the mean monthly effective spread exceeds one-half of the share price.

to be less liquid: orders have larger execution costs, greater price impacts, and take longer to execute than orders in the small sample.

Effective spreads from TAQ are calculated in a slightly different manner, because trade direction and the time of order receipt are not known for these data. For each trade between June 2000 and June 2004, effective spreads are computed as twice the absolute difference between the trade price and the NBBO midpoint at the time of the trade. We exclude trades and quotes outside regular market hours and with irregular settlements and use only trades between 100 and 9999 shares to correspond as closely as possible to the method used for Dash-5 effective spreads. As with Dash-5 reports, we compute monthly share-weighted average effective spreads for each stock.

2. Characteristics of Dash-5 Execution Quality Measures

In this section, we characterize how Dash-5 order-type composition and Dash-5 execution-cost measures vary over time and across market centers. Brokers and traders can respond to differences in execution quality by changing execution venue, order type, and, to some extent, order size. Our empirical tests are designed to capture changes in venue and order size if brokers use Dash-5 eligible orders. If brokers substitute orders not eligible for Dash-5 reports for eligible orders, then we cannot observe the execution quality from Dash-5 reports. Interpretation of changes in market share would be more difficult in such cases, because a venue's overall market share is affected differently than its Dash-5 marketable-order market share. To assess the importance of these issues, we examine the time-series behavior of the order-type mixture for the market centers in the small sample of 258 stocks.

Figure 2 shows each market's composition of incoming order flow. We divide Dash-5 order types into market orders (panel A) and marketable limit orders (panel B). The residual consists of nonmarketable limit orders within 10 cents of the relevant side of the prevailing quote (not shown). We measure incoming order volume as a percentage of each market's total Dash-5 order flow and find substantial variation across markets. For example, market orders constitute around 75% of Dash-5-eligible orders for Madoff, while nonmarketable limit orders represent almost 90% of orders on Archipelago. Although there is some fluctuation over time, we observe a high degree of persistence. These time series are very similar when we examine executed orders (not reported). Although there is a slight trend away from market orders, the results suggest that traders do not dramatically alter their choices of order type during the sample period.

We also examine whether brokers systematically switch from Dash-5 eligible to noneligible orders (not reported). We find that the proportion

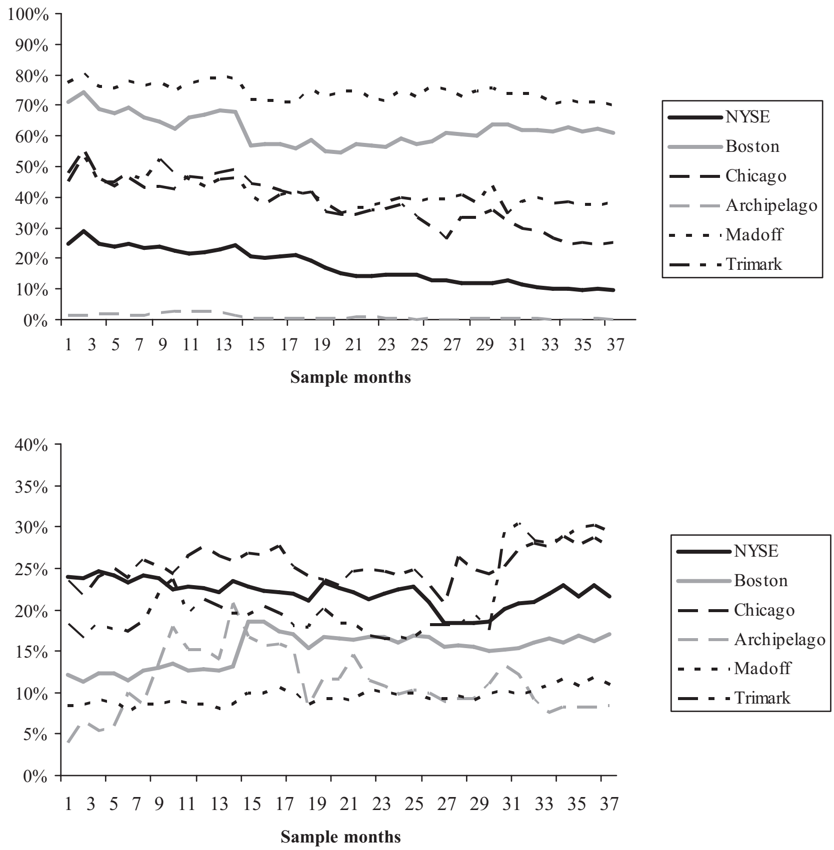

Figure 2
Composition of order flow over time
The sample covers orders in 258 NYSE-listed common stocks between June 2001 and June 2004 and is described in Table 1. Panels A and B show a time series of the percentage of order volume in different order types to total order volume as reported in the Dash-5 statistics for each individual market center. Panel A: Dash-5 market orders as percentage of all Dash-5 eligible orders placed in market center. Panel B: Dash-5 marketable limit orders as percentage of all Dash-5 eligible orders placed in market center.

of Dash-5 volume to total volume increases on all market centers over the sample period; so, it is not likely that brokers attempt to evade the disclosure rules by substituting orders that are not eligible for Dash-5. We consider this helpful in assessing the relation between routing decisions and past execution cost.

To shed more light on the importance of routing decisions, we examine the time-series behavior of each market center's Dash-5 spreads for the small-sample securities. The left half of Figure 3 presents market-specific mean effective spreads, separated into market orders (panel A) and marketable limit orders (panel B). Consistent with Lipson (2003), we

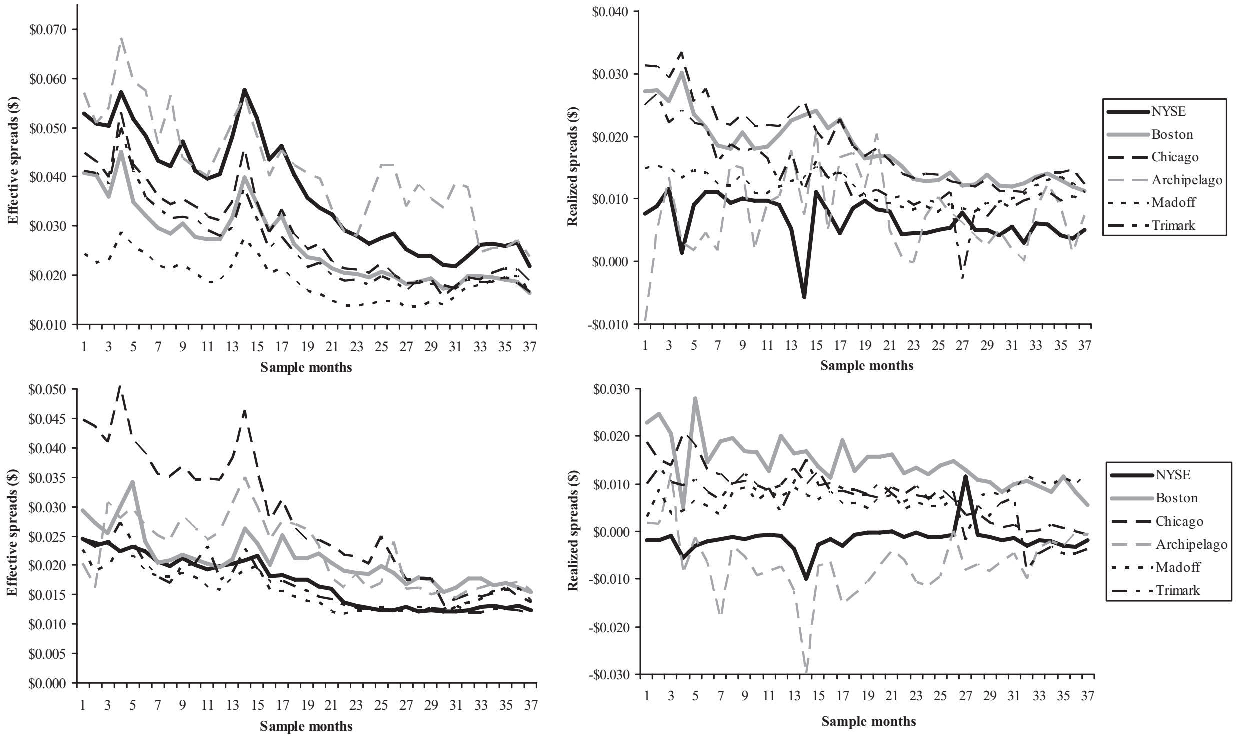

Figure 3
Effective and realized spreads over time
The sample covers orders in 258 NYSE-listed stocks between June 2001 and June 2004 and is described in Table 1. The figures show share-weighted average effective and realized spreads from Dash-5 reports. Panel A: Market orders. Panel B: Marketable limit orders.

find that differences between markets exist and that rankings are fairly persistent over time. It is important to note that a consistent ranking of effective spreads does not trivialize order-routing decisions. For example, suppose that brokers routinely route their most difficult order flow to one market and their easiest to another. Under these circumstances, we expect a consistently greater effective spread in the market to which the more difficult order flow is sent, but this does not make the routing decision irrelevant. Effective spreads reflect the total price impact of a trade, which depends on order difficulty. Order difficulty, in turn, can vary with market conditions and order characteristics. Therefore, effective spreads are useful for routing decisions only conditional on order difficulty.

A comparison of rankings based on effective and realized spreads in Figure 3 is largely consistent with this view. As in Table 2, market-order effective spreads are highest for the NYSE and Archipelago (although Archipelago receives few market orders), while realized spreads are generally the lowest for these two markets. Because realized spreads can be viewed as the portion of execution costs not because of order difficulty, this suggests that the NYSE and Archipelago tend to receive more difficult order flow than the other markets, which explains the high effective spreads. In contrast, Madoff has lower effective spreads and greater realized spreads than the NYSE and Archipelago have. These relations are similar but somewhat less extreme for marketable limit orders. Therefore, conditional on order difficulty, routing decisions seem to matter. An implication is that even markets with consistently high-effective spreads can be useful for brokers with difficult order flow. Similarly, a market with low-effective spreads might not be able to process difficult orders at low cost. Realized spreads do not, however, perfectly measure difficulty-adjusted execution costs. First, Dash-5 measurement of realized spreads focuses on the five-minute period after a trade. Because prices can move for reasons other than the trade in question, realized spreads are noisy estimates of temporary price impacts. Second, the difference between realized and effective spreads depends not only on order difficulty but also on factors unrelated to the order (e.g., the market maker's inventory at the time of execution). For these reasons, and because they represent the actual out-of-pocket execution costs paid by traders, we focus on effective spreads. Instead of attempting to net out order difficulty by using realized spreads, we use the standard approach and control for variation in price, volume, and volatility. We perform robustness tests using realized spreads and obtain qualitatively identical results throughout the article.

3. Econometric Model of Order-Routing Behavior

In this section, we specify an econometric model of routing decisions. We wish to test whether order-routing decisions depend on information published in Rule 11Ac1-5 reports. Unfortunately, we cannot observe individual routing decisions. Rather, we observe each market center's order flow, which represents the aggregate outcome of individual decisions. To make inferences about the determinants of the underlying choices, we follow an extensive marketing literature that addresses the modeling of market shares. Econometric models of market share require specific assumptions to make the estimation logically consistent. For example, predicted values for market shares should lie between zero and one for each market center, and the sum of market shares across market centers should equal one. In addition, brokers' responses to changes in certain explanatory variables can differ across markets, and we must consider possible heterogeneity across different securities. This section addresses each of these issues.

3.1 Model

When making routing decisions, brokers must choose an execution venue. In addition, larger orders can be split, which implies that brokers (or traders) also must select the size of the orders ultimately submitted for execution. As market centers might have comparative advantages in processing certain order sizes (see Lipson 2003, Bessembinder 2003a, or Boehmer 2005), these two choices are dependent. Thus, we assume that brokers choose an optimal combination of market center $m \in \{1, 2, \ldots, M\}$ and order size $s \in \{[100–499$ shares$];[500–1999$ shares$];[2000–4999$ shares$];[5000–9999$ shares$]\}$. Our analysis is limited to orders below 10,000 shares, because Dash-5 reports exclude larger orders. Treating each market center-order size combination as a different choice, which we index by $j \in \{1, 2, \ldots, J\ 4M\}$, allows the effect of execution quality to differ systematically across markets and allows traders to respond to differences in execution quality with changes in submitted order-size categories, changes in venue, or both.

We ask how the choice of j is related to observable characteristics of the different market centers and securities. To model this relationship, we assume that the choice depends on the attraction (utility) A_{ijt} of choice j (a particular order size and market center combination) for stock i at time t. Bell, Keeney, and Little (1975) show that the following relationship between market share S_{ijt} and attraction holds under reasonable assumptions:[13]

[13] Alternative sets of assumptions are sufficient. One possible set is the following: (i) A_j is nonnegative, (ii) the attraction of a subset of all available choices equals the sum of the attractions of the elements in this subset, (3) A_j is finite for all j and nonzero for at least one element, and (4) if two subsets of choices have equal attractions, then their market shares are also equal.

$$S_{ijt} = \frac{A_{ijt}}{\sum_{j=1}^{J} A_{ijt}}. \tag{1}$$

Equation (1) states that the market share of choice j depends on its attraction relative to other contemporaneously available choices for this security. This relationship can vary across securities and over time. Moreover, S_{ij} can be interpreted as the result of individual choices. If individuals choose j according to a multinomial logit model, the aggregation of their choice probabilities is consistent with Equation (1).[14]

Next, we wish to model A_{ijt} in a way that is economically meaningful. We consider the following general model of attraction (see Cooper and Nakanishi, 1988):

$$A_{ijt} = \prod_{k=1}^{k} f_k\left(X_{kijt}\right)^{\beta_k} \tag{2}$$

In this specification, the attraction of the broker's choice for a specific security depends on a set of k variables (the columns of X), representing characteristics of market quality or other considerations important to brokers. The coefficients β_k measure the sensitivity of attraction to these variables. For now, we assume that these coefficients are constant across stocks and choices, but relax this assumption later. The link function f_k is a monotone transformation of X, where $f_k(\bullet) > 0$. For estimation, both the identity and the exponential functions have desirable properties and yield models that are linear in all parameters. We choose an exponential function for our estimation, because it has the additional property that the resulting model is consistent with a multinomial choice model at the (unobservable) broker level. Adding constant and error terms,

$$A_{ijt} = \exp\left(\gamma_j + \sum_{k=1}^{K} \beta_k X_{kijt} + u_{ijt}\right) \tag{3}$$

where the γ_j represent different levels of attraction for different market center-order size combinations. Substituting into Equation (1) and taking logs yields:

[14] See Cooper and Nakanishi 1988, (Section 1.9.3), who present alternative derivations of this result.

$$\ln S_{ijt} = \gamma_j + \sum_{k=1}^{K} \beta_k X_{kijt} + u_{ijt} - \ln \sum_{j=1}^{J} \exp\left(\gamma_j + \sum_{k=1}^{K} \beta_k X_{kijt} + u_{ijt}\right). \quad (4)$$

Summing Equation (4) over all *j*, dividing by J, and subtracting the result from (4) gives:

$$\ln \frac{S_{ijt}}{\tilde{S}_{it}} = (\gamma_j - \bar{\gamma}) + \sum_{k=1}^{K} \beta_k (X_{kijt} - \bar{X}_{kit}) + (u_{ijt} - \bar{u}_{it}) \quad (5)$$

where $\tilde{S}_{it}$ is the geometric mean of market shares in period *t* for security *i* and the bars indicate arithmetic means across choices *j*. This suggests the estimable form:

$$\ln \frac{S_{ijt}}{\tilde{S}_{it}} = \gamma_0 + \sum_{j=1}^{J-1} \gamma_j I_j + \sum_{k=1}^{K} \beta_k (X_{kijt} - \bar{X}_{kit}) + u^*_{ijt} \quad (6)$$

where the I_j are choice fixed effects (representing market centers and order sizes).

Finally, we address two issues related to unobservable heterogeneity across stocks that could make OLS estimates of the coefficients in model (6) inconsistent. In particular, relative market share might be related to factors other than measures of execution quality. To address this concern, we replace the intercept term by a linear combination of security fixed effects and three security-specific control variables—price, trading volume, and volatility. These variables do not vary across choices, but we allow for different coefficients across choices. This approach allows routing decisions to depend on security characteristics as well as measures of execution quality (note that the control variables cannot be included in X, because deviations from choice-specific means would be zero). From CRSP, we compute the log of the average daily closing price (*ClosePrc*) and the log of average daily share volume (*ADV*). Using TAQ, we compute the average of the daily price range standardized by the closing price (*RelRange*). Thus, we replace the intercept term in mode l (6), γ_0, by $\gamma_{ijt} = \sum_{i=1}^{N} \alpha_i I_i + \sum_{j=1}^{J} (\beta_{j,ClosePrc} I_j ClosePrc_{it} + \beta_{j,ADV} I_j ADV_{it} + \beta_{j,RelRange} I_j RelRange_{it})$,where the I_i are security fixed effect and the I_j are choice fixed effects. The resulting model controls for unobservable effects that systematically differ across stocks and choices. By construction, specification (6) produces predicted market shares based on the attraction model (4), which are logically consistent (Cooper and Nakanishi 1988). This specification, which is based on security-specific deviations from the mean across choices at time t, is equivalent to a model of unadjusted variables that includes a set of time-series fixed effects (Nakanishi and Cooper 1982). The deviation from means form is more flexible, however, because we can include controls for security-specific characteristics such as trading volume, volatility, and share price that do not vary across the choice set, but might be related to market centers' attractiveness. In a model with fixed time effects, such variables would be linearly dependent on the effects.

3.2 Variables

3.2.1 Dependent variable. From Dash-5, we compute the market share of orders placed, S_{ijt}, as the share volume of Dash-5-eligible orders in security i in month t sent to the j^{th} market center-order size combination, divided by all marketable order volume in security i and month t across the sample venues. From TAQ, we compute market share as the share volume of trades from 100 to 9999 shares, divided by the aggregate of such volume across the sample venues. For Dash-5 variables, we have individual observations for each market center in our sample. As Nasdaq market centers are not separately identified in TAQ, we aggregate the markets that print trades on Nasdaq for all TAQ measures.

3.2.2 Independent variables. The main independent variables come from two different sources: the Dash-5 reports and TAQ. From Dash-5 reports, we obtain effective spreads, *Dash5ES*, and the time between order arrival and execution in seconds, *Dash5Speed*. These variables are share-weighted averages and computed for each stock, month, order type, order size category, and market center. To specify an economically meaningful model of market shares, we must understand the timing of all variables. We model routing decisions (and thus market shares) during month t, and so, the independent variables must be available to the decision maker at that time. Data on the control variables and the measures constructed from TAQ are available for the previous month, and so, these variables enter the basic model with a one-month lag. In contrast, Dash-5 reports are published by the end of the subsequent month (e.g., January's report must be published by the end of February). Therefore, we lag Dash-5 independent variables by two months.

Regressions also include trade-based effective spreads from TAQ, *TaqES*, for several reasons. First, we wish to understand whether individual routing decisions are associated with TAQ-based measures of execution quality that are publicly available before Dash-5 is imposed. Second, we are interested in the incremental information, in addition to that available in TAQ data, which Dash-5 reports provide. Finally, *TaqES* provides a control for a market's general liquidity because this measure includes all orders below 10,000 shares and not just Dash-5-eligible orders.

We wish to control for the extent to which one security's routing decision is associated with other securities' routing decisions. Although we cannot observe order flow for individual decision makers, we construct an indirect test.[15] In model (6), we transform *TaqES* into the deviation from its mean across market center-order size choices j. For stock i in month t, $TaqES'_{ijt} = TaqES_{ijt} - (1/J)\sum_{j=1}^{J} TaqES_{ijt}$ measures the deviation of TAQ effective spreads for choice j from those for the other available choices, or the relative benefit of sending an order to market center-order size category j. We then compute the mean deviation across all securities except the security of interest. We include the result, $TaqOtherES_{ijt} = \frac{1}{N-1}\sum_{\substack{l=1 \\ l \neq i}}^{N} TaqES'_{ljt}$, as an independent variable to assess how a market's order flow in security i relates to the average relative performance of choice j for all other securities.

4. Order Flow Sensitivity to Execution Cost Before and After Rule 11Ac1-5

In this section, we explore whether Dash-5's implementation affects the relation between quantitative execution-quality measures and order-routing decisions. Before Dash-5, the publicly available sources of execution-quality data were the Consolidated Trade System (CTS) and the Consolidated Quote System (CQS), that is, records of trade prices and quotes. These trade-based data are sold by several providers (e.g., as TAQ by the NYSE) and are available daily after the close of trading, but sophisticated users could capture the data in real time. Dash-5 reports contain monthly averages of order-based spread statistics (and provide statistics on fill rates and execution speed) arranged by order type and size and are available with a one-month lag. Compared to TAQ, Dash-5 data require less manipulation and provide additional and (arguably) more accurate information. The disadvantages include a lower reporting frequency and only partial coverage of order flow.

Whether brokers use trade-based execution quality measures for the order-routing decision before and/or after the implementation of Rule 11Ac1-5 is an empirical question. It is possible that brokers consistently ignore quantitative execution quality. Alternatively, brokers might use trade-based measures before the availability of Dash-5 but place less weight on those as Dash-5 measures become available. Or, brokers might generally raise their focus on quantitative execution measures after seeing the emphasis placed on such measures by the SEC (and potentially by their clients). To

[15] We thank an anonymous referee for this suggestion.

address this question, we conduct an event study comparing order-routing behavior before and after the implementation of Dash-5. We choose June 2000–May 2001 as the pre-Dash-5 period and December 2001–November 2002 as the post-Dash-5 period. We end the pre-event period in May 2001, because Rule 11Ac15 became effective in June 2001. The post-event period starts in December 2001 because some venues provided their first reports only in November 2001 and so that our analysis is not affected by the aftermath of the market closure following September 11, 2001. Specifically, we estimate the following model across stocks i and market centers m:

$$\begin{aligned}\log\frac{S_{imt}}{\tilde{S}_{it}} = &\sum_{i=1}^{N}\alpha_i I_i + \sum_{m=1}^{M-1}\gamma_m I_m \\ &+ \sum_{m=1}^{5}(\beta_{m,ClosePrc} I_m ClosePrc_{i,t-1} \\ &+ \beta_{m,ADV} I_m ADV_{i,t-1} + \beta_{m,RelRange} I_m RelRange_{i,t-1}) \\ &+ \delta After_t + \beta_1 TaqES'_{im,t-1} + \beta_2 After_t \times TaqES'_{im,t-1} \\ &+ \beta_3 TaqOtherES_{im,t-1} \\ &+ \beta_4 After_t \times TaqOtherES_{im,t-1} + \varepsilon_{imt}\end{aligned} \tag{7}$$

with notation as in Equation (6) except that m indexes market centers. This model does not use any Dash-5 data. *After* is a dummy variable that equals one in the post-Dash-5 period and zero otherwise. Using this specification, we can examine (i) whether TAQ information affects routing decisions before Dash-5 by testing whether $\beta_1 = 0$, (ii) whether it is important after Dash-5 by testing whether $\beta_1 + \beta_2 = 0$, and (iii) whether the change is statistically significant by testing $\beta_2 = 0$. We perform similar tests on the importance of execution quality in other securities by examining the coefficients β_3 and β_4.[16]

Table 3 summarizes the regression results for Equation (7), estimated both with and without *TaqOtherES* and the corresponding interaction term. We present the coefficients β_1 through β_4 for both the small sample (panel A) and the large sample (panel B). The first model for the small

[16] All regressions in this article use deviations from means as the dependent variables, and all include security-specific characteristics as control variables. Nevertheless, omitted factors could cause cross-sectional correlation in the errors and affect our inferences. We address this issue explicitly by estimating the correlation across stocks for each regression model (the event study and also the models estimated below). Following Chordia and Subrahmanyam (2004), we sort stocks by name (within market center, order size, and months, where applicable) and compute the correlation coefficients for adjacent residuals. Depending on the specification, the estimated correlation across stocks is about .002 for the small sample of 258 stocks and around –.01 for the large sample of 1016 stocks in each model we estimate. Using the procedure described in Chordia and Subrahmanyam (2004), these correlations are too small to have a measurable effect on nominal significance levels and therefore do not affect inference in this article.

Table 3
Changes in routing behavior around the implementation of Rule 11Ac15

	TaqES	*TaqES* × *After*	*TaqOtherES*	*TaqOther ES* × *After*	N	Adjusted R^2
Panel A: 258 stocks						
Coefficient	0.109	−1.540			24,567	0.92
p-value	.45	.00				
Test that both coefficients are jointly zero:						
Chi-square *p*-value	.00					
Coefficient	−0.380	−0.089	−5.389	−14.991	24,567	0.92
p-value	.01	.72	.00	.00		
Test that both coefficients are jointly zero:						
Chi-square *p*-value	.02		.00			
Panel B: 1016 stocks						
Coefficient	−0.421	−1.562			110,958	0.81
p-value	.00	.00				
Test that both coefficients are jointly zero:						
Chi-square *p*-value	.00					
Coefficient	−0.567	−0.916	−9.901	−32.124	110,958	0.82
p-value	.00	.00	.00	.00		
Test that both coefficients are jointly zero:						
Chi-square *p*-value	.00		.00			

The sample covers trades in NYSE-listed common stocks between June 2001 and June 2004 and is described in Table 1. We use OLS to estimate monthly panel regressions that include fixed effects for each stock and each market center except Nasdaq (all venues reporting to Nasdaq are represented as a single market center). Rule 11Ac1-5 became effective in June 2001, but various market centers began publishing Dash-5 reports between June and September 2001. The estimation period includes 12 months before June 2001 (pre-Dash-5 period) and 12 months after November 2001 (post-Dash-5 period). The dummy variable *After* equals one in the post-Dash-5 period and zero otherwise. The regressions also include the *After* dummy and controls for the average closing price, average daily volume, and the average daily price range scaled by the closing price in month $t-1$, whose coefficients vary across market centers. Coefficients on fixed effects, *After*, and the controls are not reported. The dependent variable is the month t market share of trades between 100 and 9999 shares from TAQ, expressed as the deviation from the geometric mean across market centers. *TaqES* is the share-weighted effective spread from TAQ trades between 100 and 9999 shares. It is expressed as deviations from the arithmetic mean across market centers. *TaqOtherES* is the average deviation of *TaqES* across securities, excluding the security in the current observation, for the market center in the current observation. The "joint tests" refer to tests of the hypothesis that the indicated coefficients sum to zero. We use robust standard errors to compute *p*-values.

sample shows that market share is not significantly related to effective spreads before Dash-5 but has a significantly negative association after its implementation. The latter result suggests that higher effective spreads in month t are associated with a lower market share in month $t+1$. Inference changes when *TaqOtherES* and its interaction are included. In the second model in panel A, the relation between order flow and execution quality is negative both before and after Dash-5, without a significant change around implementation. Thus, both specifications imply that routing decisions depend on trade-based spread after Dash-5, but the relationship before Dash-5 is ambiguous. We also find that the sensitivity to a market's execution costs in other stocks is significantly negative before Dash-5, suggesting that a market center's overall execution quality affects its market share in individual stocks. This sensitivity becomes stronger after Dash-5 is implemented. For the large sample in panel B, where our tests are more powerful, we find significantly negative associations before Dash-5 for both variables, which become more pronounced afterwards. The change in sensitivity is economically large. Using the first model in panel A, which attributes the entire effect to a security's own spread, market share is not sensitive (statistically speaking) to effective spreads before Dash-5, but afterwards a 1 standard deviation (SD) increase in effective spreads decreases market share by 8%. For the large sample, a 1 SD increase in effective spread decreases market share by 3.9% before the implementation of Dash-5 and by 18.2% afterwards. Thus, the implementation of Rule 11Ac1-5 is associated with a significant increase in the importance of trade-based execution costs for routing decisions.

5. Analysis of Order Market Share Sensitivity to Execution Quality Reports

Having shown evidence that public TAQ information appears to become more important in order-routing decisions after Dash-5 is implemented, we now ask whether the Dash-5 reports themselves provide incremental information to brokers. We estimate the following general model based on deviations from means across market center-order size choices j:

$$
\begin{aligned}
\log \frac{S_{ijt}}{\tilde{S}_{it}} = {} & \sum_{i=1}^{N} \alpha_i I_i + \sum_{j=1}^{J-1} \gamma_j I_j \\
& + \sum_{j=1}^{J} (\beta_{j,ClosePrc} I_j ClosePrc_{i,t-1} \\
& + \beta_{j,ADV} I_j ADV_{i,t-1} + \beta_{j,RelRange} I_j RelRange_{i,t-1}) \\
& + \sum_{j=1}^{J} \beta_{j,1} Dash5ES'_{ij,t-2} + \sum_{j=1}^{J} \beta_{j,2} Dash5Speed'_{ij,t-2} \\
& + \sum_{m=1}^{J} \beta_{j,3} TaqES'_{ij,t-1} + \sum_{m=1}^{J} \beta_{j,4} TaqOtherES_{ij,t-1} + \varepsilon_{ijt} \qquad (8)
\end{aligned}
$$

where $Dash5ES'_{ijt} = Dash5ES_{ijt} - (1/J)\sum_{j=1}^{J} Dash5ES_{ijt}$ and $Dash5Speed'_{ijt} = Dash5Speed_{ijt} - (1/J)\sum_{j=1}^{J} Dash5Speed_{ijt}$. The remaining notation is as in Equation (6). In contrast to the trade-based analysis in Section 4, the dependent variable is the market share of orders. We estimate model (8) separately for market orders and marketable limit orders.[17] Using the small sample, we have six market centers, four order-size categories, and 258 securities over 35 months, and so, the maximum number of observations is 216,720. For marketable limit order regressions, we add the order cancellation rate (available from Dash-5) as a control variable. We do not observe the opportunity costs associated with unexecuted orders (see, e.g., Peterson and Sirri, 2002), and the cancellation rate is a readily available proxy to control for these costs.[18]

5.1 Results for all sample market centers combined

Table 4 presents estimation results for Equation (8) with three different sets of restrictions on the slope coefficients. In model 1, we restrict the coefficients on the control variables to be equal across market centers. Model 2 allows their coefficients to vary across market centers, and model 3 allows the coefficients on both the control variables and the TAQ variables to vary across market centers. In each model, coefficients on the Dash-5 variables are restricted to be equal across choices (this assumption will be relaxed in Sections 5.2 and 5.3). For brevity, we do not report the coefficients on control variables or fixed effects.[19]

For the small sample (panel A), we find that the Dash-5 reports provide information that appears useful for routing decisions. Market share declines significantly as past Dash-5 effective spreads increase. Time-to-fill is negatively related to market share for market orders in model 1, but is not significant at the 5% level for market orders in models 2 and 3 or for marketable limit orders. Similar to the results in Table 3, market share also is significantly negatively related to *TaqES*. We obtain similar results for *Dash5ES* and *TaqES* in the large sample (panel B). Past execution

[17] In unreported sensitivity tests, we include the dependent variable lagged by one period (or, alternatively, two periods) as an additional regressor. This addresses the possibility that market-share deviations from the mean may be autocorrelated, which could affect inferences and the causal structure assumed in model (8). We obtain qualitatively identical results for all versions of model (8) that we estimate with and without lagged dependent variables.

[18] For brevity, we refrain from reporting coefficients on the choice-invariant variables that we use to control for order difficulty—price, volatility, and volume. The estimated coefficients are generally significant, but their signs vary across order types and market centers. We obtain qualitatively identical results when we omit the controls. We discuss coefficients on control variables more thoroughly when we estimate market-specific coefficients in Section 5.2, because this is the most interesting setting to examine these effects.

[19] Allowing the coefficients on choice-specific deviations from the mean to vary across market centers is, strictly speaking, not consistent with an attraction model. Rather, it implies a log-linear model where log market share is linearly related to deviations from the arithmetic mean of explanatory variables. We obtain qualitatively identical results using the corresponding attraction-model specification (see Cooper and Nakanishi, 1988, p. 129) and report the coefficients based on the varying-coefficients version of (model 8) to make comparison across models easier.

Table 4
The sensitivity of order-routing decisions to execution quality

		Market orders			Marketable limit orders		
Variable (unit, timing)	Market center	Model 1	Model 2	Model 3	Model 1	Model 2	Model 3
Panel A: 258 stocks							
Dash5ES ($, $t-2$)	All	−0.75***	−0.97***	−0.95***	−0.34***	−0.16***	−0.17***
Dash5Speed (seconds, $t-2$)	All	−0.00007***	−0.00002	−0.00003*	−0.00001	−0.00001	−0.00001
TaqES ($, $t-1$)	All	−1.87***	−1.38***		−2.28***	−1.90***	
TaqES ($, $t-1$)	NYSE			−4.38***			−5.55***
TaqES ($, $t-1$)	Boston			−0.78***			−1.30***
TaqES ($, $t-1$)	Chicago			−1.12***			−1.11***
TaqES ($, $t-1$)	Archipelago			−3.72***			−3.19***
TaqES ($, $t-1$)	Madoff			−1.36***			−1.78***
TaqES ($, $t-1$)	Trimark			−1.51***			−2.14***
TaqOtherES ($, $t-1$)	All	7.52***	5.60***		6.95***	4.40***	
TaqOtherES ($, $t-1$)	NYSE			13.51***			16.27***
TaqOtherES ($, $t-1$)	Boston			−2.67***			−8.59***
TaqOtherES ($, $t-1$)	Chicago			7.38***			20.27***
TaqOtherES ($, $t-1$)	Archipelago			5.68***			−5.20***
TaqOtherES ($, $t-1$)	Madoff			9.25***			6.80***
TaqOtherES ($, $t-1$)	Trimark			0.50			−3.23***
Restrictions on control variable coefficients		All equal	Equal within markets	Equal within markets	All equal	Equal within markets	Equal within markets
Adjusted R^2		0.85	0.87	0.87	0.86	0.88	0.88
Observations		181,351	181,351	181,351	183,598	183,598	183,598
Panel B: 1016 stocks							
Dash5ES ($, $t-2$)	All	−0.59***	−0.50***	−0.48***	−0.32***	−0.26***	−0.25***
Dash5Speed (seconds, $t-2$)	All	0.00005***	0.00003***	0.00002**	−0.00002***	−0.00002***	−0.00002***
TaqES ($, $t-1$)	All	−1.78***	−1.14***		−2.37***	−1.59***	
TaqES ($, $t-1$)	NYSE			−2.91***			−3.32***

Table 4
(continued)

		Market orders			Marketable limit orders		
Variable (unit, timing)	Market center	Model 1	Model 2	Model 3	Model 1	Model 2	Model 3
TaqES (\$, $t-1$)	Boston			−0.40***			−0.83***
TaqES (\$, $t-1$)	Chicago			−1.14***			−1.26***
TaqES (\$, $t-1$)	Archipelago			−1.49***			−2.36***
TaqES (\$, $t-1$)	Madoff			−1.74***			−1.60***
TaqES (\$, $t-1$)	Trimark			−1.08***			−1.68***
TaqES (\$, $t-1$)	Cincinnati			1.69***			−0.24***
TaqES (\$, $t-1$)	Philadelphia			−1.05***			−1.16***
TaqES (\$, $t-1$)	Brut			0.25			−5.02***
TaqES (\$, $t-1$)	Nyfix			1.24***			0.28
TaqES (\$, $t-1$)	Third Market			−1.31***			−4.59***
TaqES (\$, $t-1$)	State Street			−0.52***			−0.86***
TaqES (\$, $t-1$)	Schwab			−2.85***			−2.34***
TaqES (\$, $t-1$)	Primex						−3.43***
TaqES (\$, $t-1$)	Citigroup Global			−4.57***			−14.30***
TaqES (\$, $t-1$)	TD Waterhouse			−8.89***			
TaqOtherES (\$, $t-1$)	All	10.45***	9.65***		6.18***	6.79***	
TaqOtherES (\$, $t-1$)	NYSE			20.04***			24.04***
TaqOtherES (\$, $t-1$)	Boston			3.63***			−3.32***
TaqOtherES (\$, $t-1$)	Chicago			3.54***			10.25***
TaqOtherES (\$, $t-1$)	Archipelago			9.85***			−5.17***
TaqOtherES (\$, $t-1$)	Madoff			4.92***			1.09
TaqOtherES (\$, t-1)	Trimark			5.80***			1.18***
TaqOtherES (\$, $t-1$)	Cincinnati			46.23***			17.40***
TaqOtherES (\$, $t-1$)	Philadelphia			−0.28			−8.28***
TaqOtherES (\$, $t-1$)	Brut			2.71			86.31***
TaqOtherES (\$, $t-1$)	Nyfix			7.42***			17.76***
TaqOtherES (\$, $t-1$)	Third Market			−18.45***			−43.69***
TaqOtherES (\$, $t-1$)	State Street			12.07***			−4.08***
TaqOtherES (\$, $t-1$)	Schwab			−6.03***			−7.42***

TaqOtherES ($, $t-1$)	Primex						32.27***
TaqOtherES ($, $t-1$)	Citigroup Global			6.13***			
TaqOtherES ($, $t-1$)	TD Waterhouse			20.44***			29.24***
Restrictions on control variable coefficients		All equal	Equal within markets	Equal within markets	All equal	Equal within markets	Equal within markets
Adjusted R^2		0.77	0.80	0.80	0.82	0.84	0.84
Observations		678,602	678,602	678,602	660,884	660,885	660,886

The sample covers orders in NYSE-listed common stocks between June 2001 and June 2004 and is described in Table 1. We estimate monthly OLS panel regressions that include fixed effects for each stock and all market center-order size permutations except the largest order size on Trimark. Each regression controls for the average closing price, average daily volume, and the average daily price range scaled by the closing price in month $t-1$. Marketable limit order regressions also include the cancellation rate in month $t-2$ from Dash-5 reports. Coefficients on control variables and fixed effects are not reported. The dependent variable is the month t market share of orders placed, expressed as the deviation from the geometric mean across market center-order size categories. Variables from Dash-5 reports, effective spreads (*Dash5ES*), and execution speed (*Dash5Speed*), are based on share-weighted monthly averages in month $t-2$. They are expressed as deviations from the arithmetic mean across market center-order size categories. All other independent variables are recorded in month $t-1$. *TaqES* is the share-weighted effective spread from TAQ trades between 100 and 9999 shares (this measure does not vary across the market centers reporting to Nasdaq). It is expressed as deviations from the arithmetic mean across market center-order size categories. *TaqOtherES* is the average deviation of *TaqES* across securities, excluding the security in the current observation, for the market center-order size category in the current observation.
*Significance at the 10% level based on robust standard errors.
**Significance at the 5% level based on robust standard errors.
***Significance at the 1% level based on robust standard errors.

costs, both from TAQ and Dash-5, have a significantly negative relationship to future market share. In contrast to the small sample, slower executions significantly decrease future market share of marketable limit orders but are unrelated to the market share of market orders in the predicted manner. A possible reason for the positive coefficient on *Dash5-Speed* is that NYSE orders dominate the large sample. The NYSE is present in every month for each security, while we only require two other markets. The small sample, in contrast, contains the same six markets in every month. Boehmer (2005) shows that the NYSE tends to have slower executions in orders where it offers low spreads. Thus, a positive relationship between time-to-execution and market share could result from brokers trading off slower speed for lower costs. We revisit this issue in our discussion of differences across markets in the next two sections.

To illustrate the economic importance of these estimates, we take model 1 in panel A as an example. Using the unconditional SDs, the coefficients indicate that a 1 SD increase in *Dash5ES* (5.4 cents) decreases, other things constant, the future share of market orders by 4.05%. A 1 SD decrease in *Dash5Speed* (110 seconds) implies a loss of 0.77%. Finally, a 1 SD increase in *TaqES* (3.3 cents) implies a loss of 6.2%. These estimates imply economically significant penalties for poor execution quality. Given an average daily volume of 3.1 million shares (see Table 1), a one-sigma increase in effective spreads reduces daily order flow by 125,550 shares in each stock traded. The decline for a one-sigma decrease in speed is 23,870 shares per stock per day.

The coefficients on *TaqOtherES* are generally positive, and so, market shares do not decline systematically when the relative execution quality of other sample securities at that market center is poor. This suggests that, in a multivariate setting, routing decisions are driven primarily by past execution quality of a specific security, rather than a market center's average performance. This contrasts to the negative coefficients for the trade-based analysis in Table 3. A potential reason for the positive coefficients is that brokers have arrangements promising a minimum order volume to specific markets. These brokers might allocate more order flow in stocks for which a venue offers relatively high execution quality and less order flow in stocks with low execution quality on that venue. Consider an example where a broker allocates order flow in stocks A and B between two venues and has agreed to a minimum volume to both venues in exchange for order-flow payment. If venue one offers low execution costs for stock A and high execution costs for stock B, then the broker routes all the order flow in stock A to venue one and all the order flow in stock B to venue two. Model 3 shows that this association varies across markets, potentially reflecting differences in business models and routing algorithms across market centers and brokerage firms.[20]

Overall, these results suggest that both trade-based and order-based execution cost measures are factors in routing decisions and that Dash-5 statistics provide incremental information over that available from TAQ.[21]

5.2 Results by market center

One implicit assumption in the previous analysis is that brokers use the same market-quality criteria regardless of target venue. However, the varying effects of trade-based execution cost measures suggest that differences in market structure are important. For example, some market centers pay for order flow during the sample period. In addition, markets differ with respect to the dimension of execution quality they emphasize. ECNs tend to provide fast executions, and so, traders might send orders there when speed is important. Exchanges traditionally emphasize the auction process and promise better prices at the expense of execution speed. Nasdaq market makers, such as Trimark and Madoff, operate automatic execution systems that are fast and can provide price improvement for selected orders. Thus, we might expect variation in the sensitivity of routing decisions to execution quality measures across venues. To address this issue, we estimate Equation (8) allowing the Dash-5 coefficient estimates to vary across market centers. To conserve space, we do not report the coefficients on fixed effects or *TaqOtherES*.

For the Dash-5 variables, effective spread and speed, we obtain qualitatively similar results for the small and large sample, and so, we concentrate on the small sample. In Table 5, we provide separate results for market orders and marketable limit orders. Relaxing the equality restriction on Dash-5 coefficients allows several new insights. First, we find that changes in effective spreads have the greatest effect on the NYSE's market share for both order types. In fact, the sensitivity of NYSE market share to *Dash5ES* is about 50% larger than its sensitivity to *TaqES*. This illustrates that the aggregate results in Table 4 understate the economic importance of Dash-5 information, because the NYSE receives over 80% of the orders in our sample. The NYSE also receives more market orders when past execution speed is slow relative to its competitors (i.e., time-to-fill lengthens), which

[20] The coefficients on *TaqOtherEs* appear large in magnitude relative to the coefficients on the security-specific spread variables, *Dash5ES* and *TaqES*. This does not imply, however, that the market-wide effect dominates. In particular, the unconditional SD of *TaqOtherEs* is generally about 80% smaller than the SD of *TaqES*, which in turn is about 40% smaller than the SD of *Dash5ES*. Because changes in the three variables are not independent, it is not meaningful to simply sum the predicted effects on market share. However, comparing the SDs suggests that the market-wide effect does not generally offset the security-specific effects. Moreover, the estimation results are not sensitive to omitting *TaqOtherEs* from the model—there is no discernible effect on the estimated security-specific coefficients or their standard errors.

[21] The finding that Dash-5 reports provide additional information over that contained in TAQ is also supported by a comparison of effective spreads from TAQ and Dash-5, as reported in the Appendix.

Table 5
Differences across market centers in the sensitivity of order-routing decisions

	Market orders			Marketable limit orders		
Market center	*Dash5ES* ($, t – 2)	*Dash5Speed* (seconds, t – 2)	*TaqES* ($, t – 1)	*Dash5ES* ($, t – 2)	*Dash5Speed* (seconds, t – 2)	*TaqES* ($, t – 1)
New York Stock Exchange	−4.52***	0.00076***	−2.95***	−7.50***	−0.00087***	−4.57***
Boston	−0.78***	−0.00015***	−0.79***	−0.13	−0.00001	−1.30***
Chicago	−0.99***	−0.00108***	−1.09***	0.22**	−0.00003	−1.11***
Archipelago	0.05	−0.00001	−3.79***	0.03	0.00006	−3.21***
Madoff	0.26**	−0.00091***	−1.43***	0.30	0.00007***	−1.79***
Trimark	−0.68***	−0.00128***	−1.52***	−0.59***	−0.00001	−2.14***
Adjusted R^2		0.87			0.88	
Observations		181,351			183,598	

The sample covers orders in 258 NYSE-listed common stocks between June 2001 and June 2004 and is described in Table 1. We use OLS to estimate monthly panel regressions that include fixed effects for each stock and all market center-order size permutations except the largest order size on Trimark. Each regression also includes controls for the average closing price, average daily volume, and the average daily price range scaled by the closing price in month t – 1. Their coefficients are allowed to vary across market centers. Marketable limit order regressions also include the cancellation rate in month t – 2 from Dash-5 reports. Coefficients for fixed effects and controls are not reported. The dependent variable as month t market share of orders placed, expressed as the deviation from the geometric mean across market center-order size categories. Variables from Dash-5 reports, effective spreads (*Dash5ES*) and execution speed (*Dash5Speed*), are based on share-weighted monthly averages and recorded in month t – 2. They are expressed as deviations from the arithmetic mean across market center-order size categories. All other independent variables are recorded in month t – 1. *TaqES* is the share-weighted effective spread from TAQ trades between 100 and 9999 shares (this measure does not vary across the three market centers reporting to Nasdaq). It is expressed as deviations from the arithmetic mean across market center-order size categories. *TaqOtherES* is the average deviation of *TaqES* across securities, excluding the security in the current observation, for the market center-order size category in the current observation.
*Significance at the 10% level based on robust standard errors.
**Significance at the 5% level based on robust standard errors.
***Significance at the 1% level based on robust standard errors.

could explain the positive coefficient on speed in panel B of Table 4. Because execution speed and execution costs are inversely related (see Boehmer 2005), this might simply indicate that traders benefit more from reductions in costs than they lose from slower executions. Alternatively, it is consistent with the NYSE being a "market of last resort," which receives more order flow in difficult market conditions. Such difficult conditions may be characterized by slower executions, and our control variables may not fully capture such adverse conditions.

For market orders, most other market centers also have negative coefficients on both Dash-5 variables and all have significantly negative coefficients on *TaqES*. Preferencing of order flow by brokers has been suggested as a practice that is contrary to the best interests of investors. In preferencing agreements, a broker prefers one market center to the others either because of monetary payments (payment for order flow) or because of other considerations. Recent SEC regulation allows us to determine whether the sample market centers systematically enter preferencing agreements. Specifically, SEC Rule 11Ac1-6, promulgated at the same time as Dash-5, requires that brokers disclose payment for order flow and other preferencing arrangements. According to Rule 11Ac1-6 reports published during the sample period, Boston, Chicago, Madoff, and Trimark pay for order flow. We find that poor past execution quality in Boston, Chicago, and Trimark reduces their future order flow, although to a lesser extent than on the NYSE. Thus, despite preferencing, brokers appear to place value on execution quality in these market centers.

For marketable limit orders, we find that Dash-5 execution quality is generally less important than for market orders, but again document a strong effect on NYSE order flow. Both *Dash5ES* and *Dash5Speed* have a significantly negative relationship with NYSE market share and as with market orders, the Dash-5 effect is large relative to the TAQ effect. Other market centers' marketable-limit-order order flow is less sensitive to Dash-5 information. Marketable limit orders, however, are relatively unimportant on these markets (see Figure 1), so that traders may route such orders there for reasons other than past execution quality.

In Figure 4, we provide a visual representation of the coefficient estimates on the control variables based on the small sample (most coefficients on the control variables are statistically significant at the five-percent level). Controlling for execution quality, some market centers appear to attract order flow in certain types of stocks. For market orders (panel A), the NYSE and Archipelago receive a greater share of orders in high-priced stocks, while Chicago, Madoff, and Trimark do better in low-priced stocks. Boston and Trimark obtain more order flow in high-volume stocks. Archipelago and Chicago attract more orders in high-volatility stocks, while Madoff and Boston receive more orders in low-volatility stocks. For marketable limit orders, price and volume have effects similar to those reported above, except that the NYSE receives relatively more orders in low-volume stocks and Archipelago receives more in high-volume stocks. Finally, Madoff, Archipelago, and the NYSE receive a greater relative market share in low-volatility stocks, while Chicago and Boston obtain more orders in high-volatility stocks. Some market centers, such as Madoff and Trimark, take an active role in selecting their client base, but most others do not. Therefore, it is difficult to draw general inferences from the coefficients on the control variables, but we note that the sensitivity of order flow to variables other than execution quality differs substantially across markets.

5.3 Results by market centers and order size

As execution costs increase in one market, traders can either send the order elsewhere or package their trading interest into different order sizes.

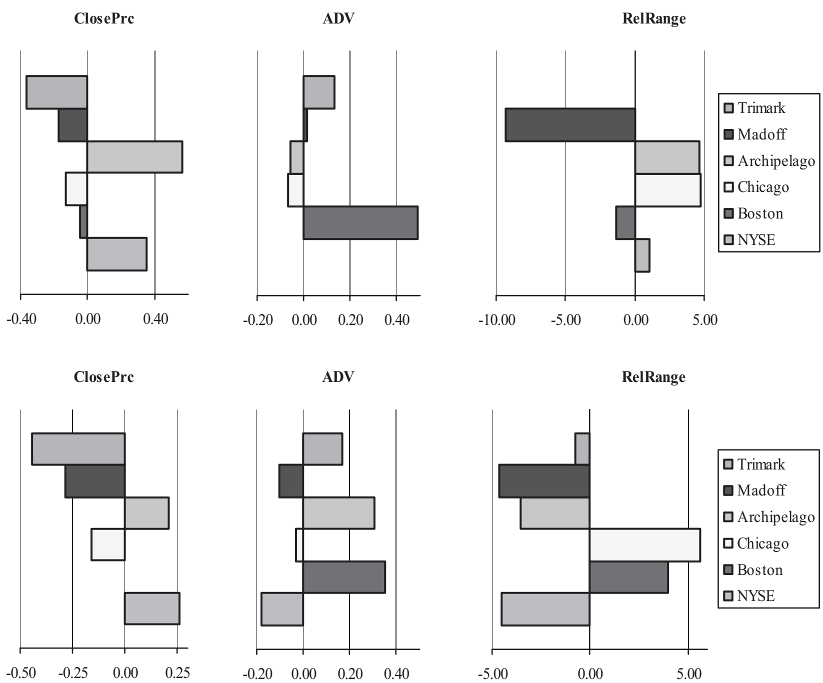

Figure 4
Coefficients on control variables by market center
The figure shows the magnitude of the coefficients on the control variables used in the estimation presented in Table 5, which contains a description of the sample and econometric model. Control variables are monthly averages. ClosePrc is the closing price, ADV is the average daily trading volume in shares, and RelRange is the daily price range divided by the closig price. ClosePrc and ADV are in logarithmic form. Panel A: Market orders. Panel B: Marketable limit orders.

Many Nasdaq broker-dealers, including Trimark and Madoff, execute most small orders automatically and large orders manually. In contrast, during our sample period, the NYSE executes most small orders manually, which slows execution. But on average, small NYSE orders execute at better-than-quoted prices (see Bessembinder, 2003c). Thus, the sensitivity of routing decisions to past execution quality might depend not only on the execution venue but also on order size.

To examine this relationship, we estimate Equation (8) as specified, allowing all Dash-5 execution-quality coefficients to vary across market center-order size choices. One problematical aspect of this approach is that we implicitly assume that each order size represents an available choice for each trader, although some traders may only have, say, 200 shares to trade. In this case, the smallest order-size category is the only available choice. On the contrary, retail investors provide only 4% of the NYSE's order flow (see Boehmer and Kelley, 2005). The remaining investors can probably choose optimally among order sizes, and we would like to assess whether this choice depends on past execution quality.

Table 6 reveals fairly consistent results across venues and order sizes (coefficients on fixed effects and market-center-specific coefficients on TAQ and control variables are estimated but not reported). The sensitivity to past *Dash5ES* frequently is significantly negative, but just over one-third of the market-order coefficients (one-half for marketable limit orders) are not significant. One important exception is the larger size categories submitted to Madoff and the second smallest submitted to Archipelago, which are positive. Traditional exchanges and Trimark generally have the expected negative association between historical cost and current market share. On most markets, the importance of past execution quality is inversely related to order size—variation in execution costs has a greater effect on small orders than on large orders. This contrasts to the NYSE, where the smallest and two largest order sizes are highly sensitive to past costs.

Throughout market centers and order sizes, market share decreases as traders experience slower executions. Notable exceptions, as indicated in Table 5, are NYSE market orders, which are not sensitive to speed in any order-size group. In contrast, the largest marketable limit orders are sensitive to speed only on the NYSE. This could be driven by differences in quoted depth across markets. Orders in the largest size category of 5000 shares or more typically exceed quoted depth, especially on markets other than the NYSE. In these cases, traders probably do not expect immediate execution, and, as a result, execution speed might be less important at these markets.

Although we believe that the pooled regression properly considers the choice of where to send orders of a particular size, we reestimate the

Table 6
Differences across market centers and order size in the sensitivity of order-routing decisions

	Market orders				Marketable limit orders			
Market center	100–499 shares	500–1999 shares	2000–4999 shares	5000–9999 shares	100–499 shares	500–1999 shares	2000–4999 shares	5000–9999 shares
Dash5ES ($, $t-2$)								
New York Stock Exchange	−3.24***	0.62	−6.68***	−4.73***	−5.98***	−1.65*	−3.99***	−12.77***
Boston	−9.96***	−2.86***	−0.17	−0.15	−0.48	−0.10	−0.31	0.26
Chicago	−12.28***	−1.42***	−0.22	−0.29**	0.40	1.34***	0.28	−0.20
Archipelago	−0.19	0.34**	0.25	−1.29***	−0.10	0.68*	−0.04	−0.08
Madoff	−4.51***	1.83***	0.95***	0.52***	−4.35***	−1.29**	0.97***	1.54***
Trimark	−12.40***	−2.80***	−0.48**	0.04	−3.08***	−0.78*	−0.27	−0.20
Dash5Speed (seconds, t-2)								
New York Stock Exchange	0.00067**	0.00058*	0.00090***	0.00066**	−0.00062***	0.00023	−0.00081***	−0.00129***
Boston	−0.00109***	−0.00198***	−0.00019	0.00000	−0.00021**	−0.00011	−0.00004	0.00004
Chicago	−0.00087***	−0.00102***	−0.00120***	−0.00064***	−0.00025**	−0.00031	−0.00016	0.00008**
Archipelago	−0.00001	−0.00007**	0.00004	0.00007	−0.00032*	−0.00014	0.00006	0.00028***
Madoff	−0.00070**	−0.00036	−0.00055*	−0.00185***	0.00001	−0.00002	0.00005	0.00008***
Trimark	−0.00018	−0.00095***	−0.00185***	−0.00097***	−0.00041***	−0.00033***	−0.00004	0.00002
Adjusted R^2		0.87				0.88		
Observations		181,351				183,598		

The sample covers orders in 258 NYSE-listed common stocks between June 2001 and June 2004 and is described in Table 1. We use OLS to estimate monthly panel regression that include fixed effects for each stock and all market center-order size permutations except the largest order size on Trimark. Fixed effect coefficients are not reported. Each regression also includes controls for the average closing price, average daily volume, and the average daily price range scaled by the closing price in month $t-1$. Their coefficients are allowed to vary across market centers. Marketable limit order regressions also include the cancellation rate in month $t-2$ from Dash-5 reports. The dependent variable is the month t market share of orders placed, expressed as the deviation from the geometric mean across market center-order size categories. Variables from Dash-5 reports, effective spreads (*Dash5ES*) and execution speed (*Dash5Speed*), are based on share-weighted monthly averages and recorded in month $t-2$. They are expressed as deviations from the arithmetic mean across market center-order size categories. All other independent variables are recorded in month $t-1$. *TaqES* is the share-weighted effective spread from TAQ trades between 100 and 9999 shares (this measure does not vary across the three market centers reporting to Nasdaq). It is expressed as deviations from the arithmetic mean across market center-order size categories. *TaqOtherES* is the average deviation of *TaqES* across securities, excluding the security in the current observation, for the market center-order size category in the current observation. The table reports only the coefficients for the two Dash 5 variables.
*Significance at the 10% level based on robust standard errors.
**Significance at the 5% level based on robust standard errors.
***Significance at the 1% level based on robust standard errors.

model individually for each order-size category as a robustness check. Because the optimal order-size choice might differ among market centers and depend on past execution quality, the size-specific regressions provide noisier estimates of routing behavior. Nonetheless, our results are quite robust to this alternative specification. In the market order regressions, most coefficient estimates on *Dash5ES* and *Dash5Speed* remain significantly negative in the size-specific models. Moreover, with the exception of the NYSE and Madoff in the smallest size category, no coefficient estimate that is statistically significant in Table 6 becomes statistically significant with the opposite sign. The size-specific marketable limit order regressions compare similarly to the pooled model, although the sensitivity of routing decisions is less sensitive to past execution quality in the smaller size categories.

5.4 Additional robustness tests

One might argue that realized spreads are a more appropriate determinant of routing decisions than effective spreads, because realized spreads do not depend on the information content of an order. However, realized spreads are not a measure of execution costs for the same reason. In particular, traders might not be able to quantify the spread premium associated with their order's difficulty. We nevertheless address this issue empirically and modify the analyses in Tables 3–6. In one approach, we add price impact as an independent variable. Price impact, computed as one-half the difference between effective and realized spreads, approximates an order's information content. In this specification, the coefficient on *Dash5ES* captures variation in effective spreads beyond that caused by information content and can be interpreted as market share sensitivity to changes in realized spreads. In another approach, we replace *Dash5ES* by the corresponding realized-spread variable. In both cases, the resulting estimates (not reported) are qualitatively identical to those for the effective spreads model, although the relationship between market share and Dash-5 spreads becomes stronger in most regressions.

Theory says little about whether and how routing decisions vary across stocks. Our primary interest is the variation across market centers and order sizes, because they represent the choices available to brokers. Thus, we restrict slope coefficients to be equal across securities. In this pooled approach, security-specific fixed effects allow us to focus on within-security variation and to test restrictions across markets and order sizes. Additionally, we use a control variable to explicitly allow routing decisions in one security to depend on the execution quality of other sample securities. To check whether our results are sensitive to this specification, we apply a two-step procedure to the small sample that allows slope coefficients to vary across securities. In the first step, we estimate 258 security-specific regressions where we restrict control variables to be equal across market centers.

In the second step, we aggregate coefficients across securities (using, alternatively, equal and volume weights) and perform tests on the cross-sectional mean and median. This specification (results not reported) produces estimates that are qualitatively similar to the ones presented for the pooled model presented in this article.

5.5 Intelligent order-routing systems as an alternative explanation

Finally, we explore whether the sensitivity of market share to execution quality could be due to factors other than the Dash-5 disclosure requirements. One notable development during our sample period is the increasing popularity of intelligent order-routing systems. For example, since early 2001, Lava Trading Inc. has offered a system that searches for liquidity in different markets simultaneously (see www.lavatrading.com). This system aggregates the order books from several electronic markets and provides functionality to discover non-displayed liquidity. Such systems are designed to find the market that currently offers the highest quality executions, given the size, urgency, and difficulty of a trader's order. Increasing usage of such "smart routers" implies more order flow to markets with low execution costs. If spreads are relatively persistent over time, this could result in significantly negative coefficients on Dash-5 spreads in our market-share regressions. But, in this case, the underlying relationship is driven by the activity of the intelligent routing system and not by the publication of Dash-5 reports.

Although we cannot disentangle statistically whether smart routers or Dash-5 reports explain the relation between market share and execution costs, we expand our analysis to shed some light on the competing explanations. Intelligent routing systems rely, among other inputs, on real-time execution-quality information. Thus, under the smart-router view, the explanatory power of contemporaneous execution-quality information should subsume that of Dash-5 execution-quality metrics, which is at least one month old. Ideally, we would like to use a measure of the router's information as a control variable. Because we do not have a direct measure of this information, we use contemporaneous *TaqES* and contemporaneous *Dash5ES* as instruments. Specifically, we reestimate model (8) with contemporaneous *TaqES* replacing its lagged value as an explanatory variable. In another test, we augment model (8) with contemporaneous *Dash5ES* (leaving the lagged TAQ variable in the model). Neither test has a discernible effect on the regression results. The first test leaves coefficients and their standard errors unchanged. The second test reduces the magnitude of the coefficients on lagged *Dash5ES* by about 20% but leaves their sign and significance unchanged. These results suggest that increased usage of smart routers alone cannot explain our findings.[22]

6. Consequences of Increased Competition

Our results suggest that the influence of execution quality on the competition for order flow increases after Dash-5 reports are available. Figure 2 implies, however, that rankings of market centers on spreads are relatively persistent. Put differently, even markets with persistently high effective spreads appear to survive. We might expect that more vigorous competition on execution quality eventually either eliminates these differences in execution quality or forces high-cost markets out of business. Thus, some routing decisions are apparently unrelated to differences in Dash-5 execution quality. This observation could arise for many reasons. Our model might not fully capture brokers' decision processes. Our execution-quality statistics (spreads and speed) might not encompass all important dimensions of execution quality. It is also possible that order routers evaluate order strategies in a particular security rather than individual orders. That is, they might consider market orders, marketable limit orders, and nonmarketable limit orders as a package. Because we have no comparable data on nonmarketable limit orders, we might misinterpret the determinants of routing decisions. Moreover, some routing decisions might be based on factors unrelated to execution quality, such as payment for order flow. In addition, markets might have a competitive advantage for certain categories of order difficulty that our controls do not capture. If high-cost markets receive more difficult orders, on average, it may make economic sense for brokers to continue using those markets. As we cannot observe the difficulty of individual orders, it is difficult to control for the associated variation in expected execution costs.

To address these concerns more directly, we provide three arguments. First, the general decline in spreads (Figure 2) is consistent with increased competition among market centers, although advances in trading technology, increases in trading volume, and other structural changes undoubtedly contribute to this decline as well.

Second, we posit that spread decreases are greater for venues that initially have higher spreads if competition at least partially explains the

[22] This test still leaves open the possibility that persistence in "true" execution quality may cause the negative relationship between Dash-5 spreads and market share. Effective spreads, whether based on TAQ or Dash-5, are noisy measures of "true" execution quality. If smart routers can somehow base routing decisions on "true" execution quality but we cannot measure it, our tests might still pick up a smart-router effect even when we control for contemporaneous effective spreads. This concern becomes less important the greater the correlation between contemporaneous effective spreads and the information actually used to make routing decisions (and, as we discuss in the introduction, several brokers indeed claim to use Dash-5 data to make routing decisions).

lower average spread. Therefore, spreads across markets should converge over time. Examining Figure 2, this prediction seems consistent with the data. The largest declines in effective (and realized) spreads occur in the highest cost market centers, which causes spreads to converge over time. To address the degree of convergence more rigorously, we compute the SD of effective spreads across market centers for each stock and each month. The average SD declines almost monotonically from 0.04 in August 2001 to 0.02 in June 2004. This suggests that differences in execution costs across markets decrease, consistent with increasing competition for order flow.

Finally, we explore complementary evidence regarding order-routing decisions using quarterly Dash-6 reports obtained from the Transaction Auditing Group (www.tagaudit.com). SEC Rule 11Ac1-6 requires that brokers with the responsibility of routing orders publicly disclose the market centers to which they route, how much of their order flow goes to which markets, and any preferencing arrangements the broker has with the market center. We sample Dash-6 reports beginning in the third quarter of 2001 (the first available) and then annually from the first quarters in 2002–2005 for all brokers where at least three of these five data points are available. We find 164 brokers meeting our data requirement and note the following trends:

- Brokers increasingly use multiple destinations for order flow. The mean number of venues across brokers increases monotonically from 1.60 in 2001 to 2.45 in 2005, and the fraction of brokers using only one market center falls from nearly three-fifths to about one-third.
- Volatility in order-routing relationships appears to increase. For each broker, we compute the coefficient of absolute variation (the absolute value of the change in market share of venue j between year t and $t + 1$ divided by the market share of venue j in year t). We average these across market centers and brokers from 2002 to 2005. The statistic increases from 0.09 in 2001 to 0.10 in 2002 and 0.14 in 2003 and falls to 0.12 in 2004.
- There appears to be binary usage of the NYSE, which declines over time. We find that 9% of brokers never send any orders to the NYSE, and 13% always send more than 95% of orders there. But these extreme order-routing practices become less prevalent over time. In 2001, 23% (33%) of the sample brokers routed all (none) their order flow to the NYSE. Both percentages decline throughout the sample period to 16% (18%) in 2005.

Overall, these observations suggest that brokers increasingly exploit their ability to choose among venues. This implies a more sophisticated approach to routing order flow, consistent with the convergence of spreads across markets and our interpretation that the Dash-5 disclosure requirements foster competition for order flow.

7. Conclusions

We use execution-quality reports required by SEC Rule 11Ac1-5 (Dash-5) to investigate whether order-routing decisions for marketable orders are sensitive to two dimensions of historical execution quality: execution costs and execution speed. Compared to trade-based measures of execution costs that were available previously, Dash-5 reports might be informative, because they report order-based measures computed separately for each market center and order type. First, using execution-cost measures that were available before the enactment of Dash-5, we document that the sensitivity of routing choices to execution costs increases around the Rule's imposition. Second, we show that Dash-5 reports contain information that appears useful in routing decisions. Controlling for other publicly available measures of execution costs, poor Dash-5 execution costs and, to some extent, slow execution speed decrease a market's future share of order flow.

These results suggest that broker-dealers face competitive pressures to route orders to low-cost and/or fast execution venues, and we present anecdotal evidence that their routing behavior changed after the enactment of SEC Rule 11Ac1-5. Moreover, market centers (including those that pay for order flow) can compete on execution quality, because order routers seem to pay attention to past performance. However, we see that differences between market centers persist, and markets with poorer execution quality continue to receive substantial order flow. This could be due to payment for order flow, omitted dimensions of execution quality, or systematic differences in order difficulty across markets that our controls do not capture.

Overall, our analysis suggests that the SEC's emphasis on disclosure to effect public policy can produce beneficial effects. The reports based on SEC Rule 11Ac1-5 appear to have value beyond other publicly available TAQ information and appear to be used in routing decisions. Their disclosure increases competition for order flow based on execution quality, which should make the allocation of resources in the market for equity trading more efficient.

Appendix: Examining the Relation between Dash-5 and TAQ Effective Spreads

Market centers compute Dash-5 reports from large databases of orders and quotes that must be filtered, aggregated, and manipulated regularly. The complexity of this task might affect data quality, because producing the reports is costly to market centers, and the SEC does not require external verification. It also is possible that data quality differs across market centers. There are examples of problems with Dash-5. MSI (www.marketsystems.com) notes that the Archipelago Exchange erroneously calculated the NBBO, on which execution cost measures are based, for its Nasdaq stocks from April 2003 to October 2003. Archipelago's published reports use only Nasdaq quotes rather than the Consolidated BBO as Dash-5 requires. MSI also states that Island reports are erroneous for March and April of 2004. Trimark submitted an inaccurate report for March 2004 that was later replaced by a corrected version. Finally, the SEC enforcement action against Instinet and Island alleges repeated inaccuracies in these markets' reports (SEC Release No. 2005-151).

Therefore, an important question is whether there are sufficient incentives for markets to produce accurate reports. Deliberate, repeated misrepresentation of execution quality would probably prompt legal action by the SEC or clients; the Instinet settlement is the first example that Dash-5 reporting is policed by regulators. Moreover, if brokers use Dash-5 reports for order-routing decisions, repeated relationships should encourage accurate reporting because brokers can verify Dash-5 reports using internal execution cost reports. Although proprietary reports cover only the broker's orders, larger firms can generate meaningful comparisons between Dash-5 reports and their actual experience. Because larger firms are more important to market centers, competition might encourage accurate reporting.

In this appendix, we examine the relation between TAQ and Dash-5 effective spreads. Specifically, we document the frequency with which the two data sources provide similar rankings of the market centers in the small sample (see Section 1.1) and how the rankings change subsequent to disagreements. Unfortunately, comparing Dash-5 to TAQ effective spreads is difficult, because they are not designed to measure identical quantities:

- TAQ is comprehensive, but Dash-5 reports cover only orders below 10,000 shares that have no special execution instructions. Because TAQ trade sizes do not correspond to Dash-5 order sizes and because order instructions are not reported in TAQ, the subset of Dash-5 eligible orders cannot be reconstructed based on TAQ information alone.
- Because order type, order side, and order arrival time are not reported in TAQ, effective spreads in the two databases measure different price impacts. Dash-5 appropriately uses the quote midpoint at the time the order arrives as a benchmark, but TAQ effective spreads rely on the midpoint at the time of the trade (or an arbitrary period before the trade). As a result, effective spreads based on TAQ ignore changes in the quote midpoint between order arrival and execution. Therefore, the costs of difficult executions, which might take longer to execute, are especially likely to be understated in TAQ, as are orders that are rerouted from one market center to another.[23] Furthermore, we must infer order direction when computing TAQ spreads.
- The comparability of Dash-5 and TAQ effective spreads could vary across market centers and depend on market conditions.
- TAQ does not provide separate trade reports for the different Nasdaq market centers.

These issues make it difficult to directly compare the levels or the magnitude of month-to-month changes of Dash-5 and TAQ effective spreads. Instead, we use a nonparametric analysis based on rankings across market centers. For each month and stock, we compute share-weighted effective spreads from TAQ (using trades sizes between 100 and 9999 shares) and Dash-5. We use all marketable orders for the Dash-5 computation, but the results are unchanged when we use only market orders. We create a ranking of the four market centers in our small sample of 258 securities (NYSE, Boston, Chicago, and Nasdaq Intermarket) for both measures, ranging from 1 (lowest effective spread) to 4 (highest effective spread). Based on these rankings, we compute the differences Rank (Dash-5) – Rank (TAQ) for each stock and month. Thus, a positive difference implies that, on average, Dash-5 effective spreads are greater than TAQ effective spreads for a market center. The frequency distribution of these differences is presented in Table A1.

We find that the two measures produce identical rankings for 34% of the stock months and differ by at most one ranking position for 76% of the stock months. The ranking difference is 3 only 5% of the time. The distribution also reveals differences across markets. On the NYSE, Dash-5 tends to produce wider spreads (and a lower ranking) than TAQ. The opposite is true for Nasdaq, and the regional exchanges are relatively balanced. Subject to the qualifications

[23] For a representative sample of 227 NYSE stocks, Boehmer (2005) documents that the average time between order receipt and execution is 22 seconds. Therefore, at least for actively traded stocks, it is feasible that quotes change during this period.

above, these results suggest that the two measures tend to produce similar but not identical rankings. Because we do not know which measure is more accurate, we cannot assess whether the differences are due to data errors. They may result from different sets of included orders, reflect the estimation error of TAQ computations, or inaccuracies in Dash-5 reports.

Another way to compare rankings is to ask what happens in month *t* if the two measures disagree in month *t* – 1. Conditional on the ranking difference in month *t* – 1, Table A2 shows the mean subsequent (one month) change in the respective rankings and the mean change in the difference of rankings. For small ranking differences, neither measure changes much the next month. However, as the disagreement increases, the two spread estimates subsequently tend to converge. Importantly, the more the measures disagree, the stronger is the subsequent change in both metrics. For example, consider the case where the Dash-5 ranking difference in month *t* – 1 is –3, implying a Dash-5 rank of 1 (best) and a TAQ rank of 4 (worst). In the next month, the ranking difference tends to increase by 2.05 ranking positions on average (so the difference in ranking in month *t* is about –1). This convergence is due to changes in both measures: on average, the TAQ rank improves by –0.97 (to about 3) and the Dash-5 rank worsens by 1.08 (to about 2).

Although effective spreads from Dash-5 and TAQ represent different concepts, our results suggest that disagreements between TAQ and Dash-5 are temporary. We also find that the rate of convergence is greater when the disagreement is greater. More importantly, we show that both measures change subsequent to a disagreement. If disagreements were mostly due to Dash-5 errors, one would expect that the TAQ rank remains relatively constant, while the Dash-5 rank adjusts to correct the error. In summary, these results suggest that both TAQ and Dash-5 provide reasonable measures of effective spreads. Moreover, they are consistent with the finding in the main analysis that Dash-5 data provide additional information over that contained in TAQ data.

Table A1

Frequency distribution of differences between effective-spread rankings from Dash-5 and TAQ

Dash-5 rank − TAQ rank	NYSE (%)	Boston (%)	Chicago (%)	Nasdaq (%)	All market centers (%)
−3	0	2	3	10	4
−2	0	5	6	24	9
−1	2	18	18	34	18
0	38	35	42	23	34
1	40	27	23	7	24
2	16	12	8	2	9
3	3	1	0	0	1
	100	100	100	100	100

The sample covers 258 NYSE-listed stocks between June 2001 and April 2004 and is described in Table 1. For each month and stock, we compute share-weighted effective spreads from TAQ and Dash-5. For the Dash-5 calculations, we use all marketable orders. For TAQ, we use only trades between 100 and 9999 shares. Then we create a ranking of the four sample market centers (NYSE, Boston, Chicago, and Nasdaq Intermarket) for both measures, ranging from 1 (lowest effective spread) to 4 (highest effective spread). Based on these rankings, we compute the differences, Rank (Dash-5) − Rank (TAQ), for each stock and month. The table shows the frequency distribution of these differences.

Table A2
Ranking changes after disagreements between Dash-5 and TAQ effective spread rankings

Rank difference in month $t-1$	Ranking change in month t	NYSE	Boston	Chicago	Nasdaq	All market centers
−3	Average change in TAQ rank	−2.00	−1.38	−0.86	−0.90	−0.97
	Average change in Dash-5 rank	1.00	1.68	1.79	0.79	1.08
	Average change in rank difference	3.00	3.06	2.65	1.69	2.05
−2	Average change in TAQ rank	−1.30	−0.79	−0.37	−0.37	−0.43
	Average change in Dash-5 rank	0.90	1.14	1.26	0.45	0.69
	Average change in rank difference	2.20	1.93	1.63	0.81	1.12
−1	Average change in TAQ rank	−0.89	−0.38	−0.28	0.14	−0.12
	Average change in Dash-5 rank	0.71	0.66	0.55	0.11	0.37
	Average change in rank difference	1.60	1.03	0.82	−0.03	0.49
0	Average change in TAQ rank	−0.05	−0.20	−0.26	0.36	−0.09
	Average change in Dash-5 rank	0.49	0.00	−0.22	−0.44	−0.01
	Average change in rank difference	0.55	0.20	0.04	−0.79	0.08
1	Average change in TAQ rank	0.04	0.35	0.46	0.62	0.27
	Average change in Dash-5 rank	−0.14	−0.34	−0.35	−1.06	−0.31
	Average change in rank difference	−0.18	−0.69	−0.82	−1.68	−0.58
2	Average change in TAQ rank	0.09	0.86	1.08	0.85	0.56
	Average change in Dash-5 rank	−0.77	−0.75	−0.68	−1.71	−0.79
	Average change in rank difference	−0.86	−1.61	−1.76	−2.56	−1.36
3	Average change in TAQ rank	0.12	1.41	1.77	1.53	0.66
	Average change in Dash-5 rank	−1.22	−1.29	−1.29	−1.80	−1.27
	Average change in rank difference	−1.34	−2.70	−3.06	−3.33	−1.92

The sample covers orders in 258 NYSE-listed stocks between June 2001 and April 2004 and is described in Table 1. For each month and stock, we compute share-weighted effective spreads from TAQ and Dash-5. For the Dash-5 calculations, we use all marketable orders. For TAQ, we use only trades between 100 and 9999 shares. Then we create a ranking of the four sample market centers (NYSE, Boston, Chicago, and Nasdaq Intermarket) for both measures, ranging from 1 (lowest effective spread) to 4 (highest effective spread). Based on these rankings, we compute the differences, Rank (Dash-5) – Rank (TAQ), for each stock and month. The table shows how rankings change in month t when there is a disagreement between Dash-5 and TAQ rankings in month $t-1$.

References

Bacidore, J., K. Ross, and G. Sofianos, 2003, "Quantifying Best Execution at the NYSE," *Journal of Financial Markets*, 281–307.

Bell, D. E., R. L. Keeney, and J. D. C. Little, 1975, "A Market Share Theorem," *Journal of Marketing Research*, 12, 136–141.

Bessembinder, H., 1999, "Trade Execution Costs on NASDAQ and the NYSE: A Post-reform Comparison," *Journal of Financial and Quantitative Analysis*, 34, 387–407.

Bessembinder, H., 2003a, "Selection Biases and Cross-market Trading Cost Comparisons," working Paper, University of Utah.

Bessembinder, H., 2003b, "Quote-Based Competition and Trade Execution Costs in NYSE-Listed Stocks," *Journal of Financial Economics*, 70, 385–422.

Bessembinder, H., 2003c, "Trade Execution Costs and Market Quality after Decimalization," *Journal of Financial and Quantitative Analysis*, 38, 747–777.

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.

Blume, M., and M. Goldstein, 1992, "Differences in Execution Prices Among the NYSE the Regionals, and the NASD," Working Paper 4–92, Rodney White Center for Financial Research, Wharton School of Business, University of Pennsylvania.

Boehmer, E., 2005, "Dimensions of Execution Quality: Recent Evidence for U.S. Equity Markets," *Journal of Financial Economics*, 78, 463–704.

Boehmer, E., and E. Kelley, 2005, "Institutional Investors and the Informational Efficiency of Prices," working Paper, Texas A&M University.

Burns, J. and A. Lucchetti, 2005, "Instinet Units to Pay $2.2 Million to Settle Inaccurate-data Case," *The Wall Street Journal*, 19 October, C4.

Chordia, T., and A. Subrahmanyam, 2004, "Order Imbalance and Individual Stock Returns: Theory and Evidence," *Journal of Financial Economics*, 72, 485–518.

Chung, K., C. Chuwonganant, and D. T. McCormick, 2004, "Order Preferencing and Market Quality on NASDAQ before and after Decimalization," *Journal of Financial Economics*, 71, 581–612.

Cooper, L. G., and M. Nakanishi, 1988, *Market Share Analysis*, Kluwer Academic Publishers, Boston, MA.

Easley, D., N. Kiefer, and M. O'Hara, 1996, "Cream-Skimming or Profit-Sharing? The Curious Role of Purchased Order Flow," *Journal of Finance*, 51, 811–834.

Huang, R., and H. Stoll, 1996, "Dealer Versus Auction Markets: A Paired Comparison of Execution Costs on NASDAQ and the NYSE," *Journal of Financial Economics*, 41, 313–357.

Huang, R. D., 2002, "The Quality of ECN and Nasdaq market maker quotes," *Journal of Finance*, 57, 1285–1319.

Lee, C., 1993, "Market Integration and Price Execution for NYSE-Listed Securities," *Journal of Finance*, 48, 1009–1038.

Lipson, M., 2003, "Competition Among Market Centers," working paper, University of Virginia.

Nakanishi, M., and L. G. Cooper, 1982, "Simplified Estimation Procedures for MCI Models," *Marketing Science*, 1, 314–322.

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.

Traders Magazine, 2001, "Chicago Exchange move for Price Improvement," 1 January.

Traders Magazine, 2002, “Captive Order Flow is Dead,” 1 February.

Traders Magazine, 2003, “Buy-side Under Fire to Evaluate Brokers,” 1 February.

Traders Magazine, 2005, “A New Vision at Knight,” 21 June.

U.S. Securities and Exchange Commission, 1997, “Report on the practice of preferencing,” 15 April.

U.S. Securities and Exchange Commission, 2001, “Report on the comparison of order execution across equity market structures,” 8 January.