

Competition on the Nasdaq and the growth of electronic communication networks

Jason Fink ^a, Kristin E. Fink ^a, James P. Weston ^{b,*}

^a College of Business MSC 0203, James Madison University, Harrisonburg, VA 22807, United States

^b Jesse H. Jones Graduate School of Management, Rice University, Herring Hall, MS 531, 6100 Main Street, Houston TX 77005-1892, United States

Received 13 May 2005; accepted 10 October 2005

Available online 3 March 2006

Abstract

This paper examines the growth of electronic communication networks (ECNs) and their competitive impact on the Nasdaq. We find that the development of these alternative trading platforms is associated with tighter quoted, effective, and relative bid–ask spreads, greater depths, and less concentrated markets. Further, our results show that an increase in ECN trading may have caused some traditional market makers (wholesaler and national retail dealers) to exit the market for market making. Overall, our results suggest that ECNs provide a source of competition to traditional Nasdaq dealers.

© 2006 Elsevier B.V. All rights reserved.

JEL classification: G18; G10; D43

Keywords: Trading costs; Spreads; Liquidity; ECNs; Competition

0. Introduction

Recent technological advances in computer-automated trading are radically changing the way stock markets operate. One of the most dramatic changes in US markets has been the growth of electronic communication networks (ECNs). ECNs are electronic trading systems that allow investors to execute trades through an open limit order book. Rather

* Corresponding author. Fax: +713 348 6331.

E-mail addresses: finkjd@jmu.edu (J. Fink), finkke@jmu.edu (K.E. Fink), westonj@rice.edu (J.P. Weston).

than place orders with a specialist or traditional dealer, traders on ECNs may anonymously submit orders and trade with each other directly. Competing directly with traditional Nasdaq dealers, ECNs offer a low-cost and anonymous alternative to traditional trading. This paper examines the growth of trading through ECNs on the Nasdaq and its effect on competition in the market for market making.

The competitiveness of the Nasdaq's dealer market system has recently become a contentious debate. Proponents of dealer markets such as the Nasdaq argue that competition for order flow between market makers reduces transaction costs. However, there has been increasing evidence that the Nasdaq system of trading suffered severely from imperfect competition prior to the imposition of recent market reforms. The assertion by [Christie and Schultz \(1994\)](#) of tacit collusion among Nasdaq market makers and the evidence of unusually large spreads relative to other trading systems calls into question the competitiveness of the Nasdaq market.¹ In fact, [Weston \(2000\)](#) finds that as much as 30% of Nasdaq spreads (prior to recent market reforms) may be attributed to economic rents. Finally, evidence of practices such as payment for order flow and preferenced trading on the Nasdaq suggests that traditional dealers may attract order flow through non-price competition which could prevent large spreads from being competed away.² Given these institutional features, the growth in trading via ECNs has the potential to further mitigate the effects of imperfect competition among Nasdaq dealers.

While the force of competition from ECNs has been driven by technological advances, the growth of the ECN market has also benefited from recent SEC regulations. In 1997 the SEC imposed the new Order Handling Rules (OHR). A market maker who receives an inside limit order must now either change its quote to reflect the new order, execute the order, deliver the order to an exchange, send the limit order to another market maker, or deliver the order to an ECN.³ This reform has had the consequence of providing greater access to ECNs for public investors, potentially increasing competition on the Nasdaq.

Theoretically, the effect of ECN competition on market quality is unclear. While it would seem obvious that increased competition from a lower cost provider of liquidity should improve market quality, there are subtle issues that make the net effect of the addition of ECNs ambiguous. The first cost relates to market fragmentation. If orders are drawn away from traditional dealers towards alternative venues, there is the potential that such fragmentation could reduce overall market quality.⁴ The addition of ECNs may be thought of as providing an additional dealer to the market. However, the technological structure of an ECN is different than that of a traditional dealer. It may be that the ECN creates a market effect different than the simple addition of one more dealer. In this case, market quality may be reduced in accordance with the fragmentation hypothesis.

A second potential negative effect of ECNs arises from their anonymity. However, the net effect of anonymity is ambiguous. For example, [Heidle and Huang \(2002\)](#) argue that greater anonymity can lead to larger spreads because intermediaries have a harder time

¹ For example, [Huang and Stoll \(1996\)](#) and [Bessembinder and Kaufman \(1997\)](#) find that trading costs for the NASDAQ are considerably larger than on the NYSE. Similarly, [Barclay \(1997\)](#) and [Barclay et al. \(1998\)](#) find that firms realize a decrease in spreads when they switch exchange listing from the Nasdaq to the NYSE.

² See [Easley et al. \(1996\)](#), [Bloomfield and O'Hara \(1998\)](#), and [Kandel and Marx \(1999\)](#).

³ See Securities and Exchange Commission Release No. 34-38156.

⁴ For example, see [Hamilton \(1979\)](#), [Mendelson \(1987\)](#), [Madhavan \(1995\)](#), [Easley et al. \(1996\)](#), and [Battalio et al. \(1997\)](#) for discussions of fragmentation.

distinguishing informed trades. They note that “if dealers are unable to segregate the informed traders from the liquidity traders, the resulting pooling equilibrium is one where the bid–ask spreads are wider than in an environment where the institutional design permits the segregation of the two parties.” However, in a market consisting of traditional dealers and ECNs, it is unclear that these spreads would increase. Garfinkel and Nimalendran (2003) find that on days in which they can identify corporate insider trading, the NYSE specialist system has larger spreads than the Nasdaq dealer system. But if informed traders prefer the anonymity of trading on ECNs then dealer spreads may decline because of the reduced number of informed traders they face.⁵ In this paper, we simply test whether the net effect of ECNs improves market quality. Our results suggest that the competitive effect of ECNs does outweigh any potentially harmful effects.

Our work is closely related to a number of recent studies on the role of ECNs. For example, Simaan et al. (2003) study 50 stocks over a 10-day period in September 1997. They find that dealers are more likely to post narrow spreads when they can do so anonymously via ECNs. Huang (2002) studies the individual dealer quotations of 30 high volume stocks in July 1997 and November 1998 and argues that ECN investors quickly update their limit orders at the inside and make a significant contribution to price discovery. Biais et al. (2002) examine seven stocks over five trading days in March, 2000 and find that quotes on the Island ECN often undercut market maker quotes. Barclay et al. (2003) use proprietary data to investigate the anonymity and execution quality of ECNs during the month of June 2000. They find that ECNs are associated with smaller bid ask spreads but are also associated with larger price impacts in the cross section. Similarly, Fong et al. (2002) study alternate trading systems on the Australian stock exchange and find that fragmentation across trading venues has no adverse effect on price impact measures. Finally, Conrad et al. (2003) examine a sample of proprietary institutional orders and find that they receive lower execution costs on ECNs relative to traditional dealers, though they are typically smaller orders and have lower fill rates. The recurring theme in all these studies is that ECNs appear to have a positive cross-sectional effect on some measures of market quality. However, each of these studies is limited to a static view of ECN competition and focus on a narrow set of firms or sample period. Thus, it is difficult to infer whether systematic improvements in Nasdaq market quality are driven by the presence of ECNs.

Our paper makes a contribution to this literature. By examining a comprehensive sample of stocks over a long time period (six years), our results address the evolving nature of Nasdaq dealer competition. As a result, we are able to make broader interpretations about the competitive role of ECNs on the Nasdaq market structure. Our focus is not only on the execution quality of particular orders for individual stocks, but also on aggregate measures of liquidity, trading costs, and dealer competition for the Nasdaq as a whole. We find that ECN activity is associated with smaller spreads and, at least prior to decimalization, greater market depth. We therefore conclude that ECN activity provides a significant source of competition to traditional Nasdaq dealers.

As a joint test of the cost-competition hypothesis, we also study the effect that ECN competition has on the entry and exit decisions of traditional Nasdaq market makers. We find that ECNs have not caused, *ceteris paribus*, a net reduction in the average number of NASDAQ market makers. However, when we examine the entry and exit decisions for

⁵ We thank an anonymous referee for this suggestion.

different types of dealers, we find that the success of ECN trading has been at the cost of wholesale and national retail dealers. In fact, we find, *ceteris paribus*, a net reduction in the average number of wholesale and national retail market makers on Nasdaq. At the same time, we find an increase in the number of institutional and regional brokers. Given that wholesale dealers and national retail brokers on Nasdaq (e.g., Bernard L. Madoff Securities) often make payments for order flow, our results suggest that ECNs may have a negative effect on the ability of firms to profitably purchase order flow. These results are also consistent with the hypothesis that ECNs have further reduced economic rents, forcing exit from the industry and improving competition. As a result, the dramatic growth of electronic communication networks has improved the long-run structure, conduct, and performance of the Nasdaq market.

The remainder of the paper proceeds as follows. Section 1 describes our sample selection and data collection. Section 2 documents the growth of ECN trading in the Nasdaq market. In Section 3 we investigate the effect of ECNs on spreads and depths. Section 4 tests the relationship between ECNs and dealer entry and exit. Section 5 presents a series of robustness tests. Section 6 concludes.

1. Data/sample

1.1. Sample construction

Our sample is drawn from the universe of Nasdaq stocks with prices greater than \$5 and an equity market capitalization of greater than \$20 million. We further restrict our sample to firms with available data on TAQ, CRSP, and Nasdaq market maker volume data. Our sample period begins in 1996 and ends in 2002, though most of our analysis is restricted to post-1997. While most of our data series are available at the monthly frequency, many of the series we study (e.g., market value) are highly persistent and/or non-stationary at a high frequency. As a result, we construct quarterly data for most of our analysis, though we also look at annual changes. On average, our sample has roughly 2500 firms each quarter over five years (1998–2002) for a total of 49,307 firm-quarter observations.

Data are collected from three sources. First, data on market maker volumes are provided by the Nasdaq. These data list, monthly, the total number of shares traded for all Nasdaq stocks individually, by each dealer (including ECNs) from January 1996 to December 2002. We form quarterly data from these series by aggregating volume shares for each dealer over the three months in each quarter. Firm characteristics (trading volume, market capitalization, price, volatility of returns) are collected from CRSP. Quoted depths and spreads are collected from the TAQ database provided by the New York Stock Exchange.⁶ Quoted spreads are computed as the average difference between the quoted bid and ask prices over all posted quotes for each stock during the sample period. Relative spreads are computed as the quoted spread divided by the prevailing midpoint of the spread. Since many trades occur within the bid and offer price, the quoted spread may not provide an accurate measure of the cost of trading. To account for this bias, we use Roll's (1984) serial covariance method of constructing the effective spread based on tick

⁶ We employ standard filters on all TAQ data to remove reporting and input errors following Weston (2000).

Table 1
Summary statistics

	No. obs.	Mean	Std. Dev.	25th percentile	Median	75th percentile
<i>Descriptive variables</i>						
Price	49,307	20.24	21.01	9.27	15.00	24.30
Quarterly volume (M Shares)	49,307	9.08	48.73	0.26	1.27	5.14
Market capitalization (M\$)	49,307	1121	10,590	73	173	483
Standard deviation of returns	49,307	0.041	0.023	0.025	0.037	0.052
<i>Liquidity variables</i>						
Quoted spread	49,307	0.289	0.250	0.141	0.223	0.352
Relative spread	49,307	0.019	0.013	0.009	0.016	0.026
Effective spread	49,307	0.151	0.171	0.074	0.114	0.179
Quoted inside depth	49,307	8.66	7.70	5.26	7.38	10.39
Number of market makers	48,325	31.45	35.16	10.00	20.33	40.00
Market concentration (HHI)	49,307	1960	1220	1070	1620	2540
<i>ECN variables</i>						
<i>ECNshare</i>	49,307	0.086	0.077	0.026	0.063	0.129
<i>ECNshare</i> (for Instinet activity only)	49,307	0.026	0.027	0.007	0.021	0.038
<i>ECNshare</i> (for Island activity only)	49,307	0.013	0.024	0.001	0.005	0.015

This table provides summary statistics for our sample. The sample period extends from the first quarter of 1998 to the fourth quarter of 2002. Reported statistics are based on all stock-quarters. Price is the average daily closing price during the quarter. Quarterly volume is computed as total trading volume over the quarter in millions of shares. Market capitalization is the average daily closing price times shares outstanding. Volatility is computed as the standard deviation of daily returns over each stock-quarter. Quoted spreads are computed as the average of the difference between the quoted ask and bid prices over all posted quotes. Relative spreads are computed as the quoted spread divided by the prevailing midpoint of the spread. Effective spreads are computed using the Roll (1984) serial covariance estimator. Quoted inside depth is constructed as the average of the number of shares quoted at the inside bid and ask over all quotations during the quarter, reported in hundreds of shares. Market concentration is computed as the sum over all market makers of the squared market shares for each market in each stock-quarter. *ECNshare* is constructed as the sum of reported ECN volumes divided by total volume for each stock-quarter.

data for all transactions in each stock-quarter to estimate effective spreads. Schultz (2000) demonstrates that this technique works surprisingly well when applied to intra-day data.⁷

Table 1 provides summary statistics for our sample of firms. The sample of firms has considerable variation over market capitalization, volume, and volatility where volatility is measured as the standard deviation of daily returns over the quarter. The average quoted spread is 29 cents (or roughly two eighths) while the relative spread is roughly 1.9%. On average there are 31 Nasdaq dealers making a market in a given stock. The market is also relatively concentrated with a mean Hirshman–Herfindahl index of 1960 but a somewhat lower median value of 1620 which would be considered at least moderately concentrated by the Department of Justice Merger Guidelines.⁸ Only 25% of the stocks in our sample fall below the moderately concentrated threshold. Overall, the market is less

⁷ As Schultz points out, there may be serious errors with matching trades to quotes over the sample period and so relying on a trade-based measure of the effective spread gives a more consistent and reliable metric.

⁸ The Department of Justice Merger Guidelines define a market with a concentration level below 1000 to be less concentrated, a level between 1000 and 1800 to be moderately concentrated and a concentration above 1800 to be highly concentrated.

concentrated in high-volume high-capitalization stocks. For example, many large firms such as Microsoft Corp., Intel Corp. and Cisco Systems have concentration levels well below 500.

In order to test whether ECNs harm or improve competition and market quality, we must identify the extent of ECN activity. To construct this measure, we first compute the market share of each dealer (including all ECNs) each period based on their reported dealer volume. That is, we define a new variable, *Dealershare*, using the Nasdaq reported market share variable for each dealer (j), including ECNs, in each stock (i) at each time (t) as

$$Dealershare_{j,i,t} = \frac{Dealer\ Volume_{j,i,t}}{Total\ Volume_{i,t}}.$$

Among all dealers, we then identify ECNs by their market maker identifier on the Nasdaq market maker volume data. In total, we identify 9 ECNs that operated during our sample period.⁹ The Nasdaq market maker volume data contains identifiers with each market maker that distinguishes ECNs from typical registered Nasdaq dealers. Using these identifiers for our nine ECNs, we construct a proxy for ECN activity for each stock-quarter by aggregating our *Dealershare* variable over all ECNs as in

$$ECNshare_{i,t} = \frac{\sum_{j=1}^9 ECN\ Volume_{j,i,t}}{Total\ Volume_{i,t}}.$$

We use *ECNshare* as our measure of ECN activity throughout the paper. However, this measure is problematic because it underestimates true ECN activity. To directly address any concerns regarding the magnitude or effect of any bias of our *ECNshare* measure, we perform a series of tests in Section 1.2 to demonstrate that our measure is indeed a good instrument for true ECN activity.

The summary statistics for *ECNshare* in Table 1 give an overview of ECN activity in Nasdaq stocks over the sample period. *ECNshare* has a mean of 8.6 percent, although there is considerable time-variation in this measure (detailed in Section 2). A correlation analysis (not reported) demonstrates that *ECNshare* is positively related to measures of liquidity such as volume and turnover, as well as the number of market makers, market capitalization, and the standard deviation of returns. Overall, the correlations suggest that ECNs are more active in large, high-volume stocks.

1.2. Measuring ECN activity

In order to make inferences regarding the effect of competition from ECNs, we need to be sure that our measure, *ECNshare*, is a good measure of ECN activity. The obvious way to construct a robust measure would be to identify trades that are cleared through ECNs versus those cleared through traditional dealers. Unfortunately, this is a difficult statistic to construct. One important problem with such a measure is that the market maker volume data provided by the Nasdaq includes only reported volume by each market maker. As a result, any market maker that reports a trade to Nasdaq has the volume recorded.

⁹ These are: Instinet, Island, Archipelago, RediBook, Brut, Bloomberg Trade Book, Strike, Nex Trade, and Attain.

Instinet (the largest ECN) does not typically report trades to Nasdaq when a registered Nasdaq market maker is on one side of the trade. Instead, the trade is reported by the traditional dealer. Trades that are executed over Instinet are reportedly only printed to Nasdaq by Instinet when parties on both sides of the transaction are not registered market makers. One example of this situation is when two fund managers clear a trade directly over Instinet – bypassing the traditional dealer.

Thus, our *ECNshare* measure for Instinet may be more reflective of institutional trading via Instinet than total volume traded over that ECN. As a result, our estimate of Instinet volume (and, subsequently, of total ECN volume) may be considerably underestimated. For example, a report on electronic trading by the Securities and Exchange Commission estimates that in March 1999, ECNs accounted for roughly 30% of total Nasdaq volume. However, our estimate of ECN activity based on *ECNshare* for that quarter is only 9.7%. This bias is a potential problem in our statistical inference. If our measure is not correlated with true ECN activity (or if it has a cross-sectional or time-series bias), then our inference could be spurious. In this section we test the hypothesis that our *ECNshare* measure is a meaningful instrument for the market share of ECNs.¹⁰

To test whether *ECNshare* represents a meaningful proxy for ECN activity, we turn to three other measures of market maker activity for which we have limited, but more accurate, data. Since the reported Nasdaq volumes are not mis-measured for traditional Nasdaq market makers, our strategy is to find a measure of market maker activity that is highly correlated with the reported monthly volumes for traditional dealers and then test whether this measure is also correlated with our proxy for ECN activity.¹¹

To this end, we use the Nastraq data provided by Nasdaq to compute alternative measures of dealer activity. This database contains all of the dealer quotes for each Nasdaq security along with a time-series of inside quotes (BBO) and transactions records. From this database we construct three alternative measures of dealer activity in a particular stock:

1. **Quote activity.** This statistic is measured as the share of quote revisions for each dealer in a particular stock. That is, we count the number of times dealer i revised their quotes for stock j and then divide by the total number of quote revisions for stock j .
2. **Inside quotes.** For each stock, we match each dealer's quotes with the prevailing inside bid and ask price. We then measure the proportion of the time (during normal trading hours) that the dealer had either the inside bid or offer price. Since an active stock with many dealers decreases the probability that any particular dealer has the best prices, we scale this proportion by subtracting $1/M$, (the equal proportion of time we might expect a given dealer's price to reside at the inside of the spread) where M is the number of dealers in the stock.

¹⁰ In March of 2002, Island switched its reporting from Nasdaq to the Cincinnati stock exchange, creating another potential problem with the *ECNshare* measure. This problem affects only the last three quarterly periods in our data set. We use the full dataset in the generation of our results in all sections but for robustness, we also tested our models excluding the final three quarters of data. Our results are not qualitatively affected.

¹¹ We are indebted to Pete Kyle for useful discussions on the under-reporting problem and for suggesting the instrumental variables approach.

3. **Inside volume.** This measure is constructed in a similar manner to measure 2. However, rather than record the amount of time that each dealer spends at the inside, we match the dealer quotes to both the inside quote data as well as the transaction data. From this database, we measure all of the volume executed while each dealer had either the inside bid or ask price.

The procedure for computing these three measures from the Nasdaq database is computationally intensive. As a result, we limit our consideration to 1000 randomly chosen stocks from our sample that have available data in both July 1998 and October 1999.

With these three alternative measures of dealer activity, we test whether our *ECNshare* statistic is a good measure of ECN activity. In Panel A of Table 2 we present the correlation between our *Dealershare* variable (based on Nasdaq reported volume) and our three

Table 2

Correlations between the proxy for ECN activity and other measures of dealer activity

Panel A: Correlations for <i>Dealershare</i> over all dealers ($N = 16,257$)			
	<i>Dealershare</i>	Inside volume	Inside quotes
Inside volume	0.752	1.000	
Inside quotes	0.605	0.750	1.000
Quote activity	0.615	0.615	0.578
Panel B: Correlations <i>Dealershare</i> using ECNs only ($N = 3549$)			
	<i>Dealershare</i> (ECNs only)	Inside volume	Inside quotes
Inside volume	0.651	1.000	
Inside quotes	0.590	0.659	1.000
Quote activity	0.781	0.696	0.690
Panel C: Correlations for <i>Dealershare</i> using the Instinet ECN only ($N = 950$)			
	<i>Dealershare</i> (Instinet only)	Inside volume	Inside quotes
Inside volume	0.590	1.000	
Inside quotes	0.626	0.785	1.000
Quote activity	0.389	0.562	0.623
Panel D: Correlations for <i>ECNshare</i> ($N = 1000$)			
	<i>ECNshare</i>	Inside volume	Inside quotes
Inside volume	0.748	1.000	
Inside quotes	0.732	0.768	1.000
Quote activity	0.830	0.720	0.674

This table presents correlations between our measure of ECN activity (based on Nasdaq reported volume) with three other measures of ECN activity constructed from the Nasdaq database of individual dealer quotes. The sample is based on 1000 randomly selected firms. The inside volume measure records all of the volume executed during each month that occurred while each dealer had either the inside bid or ask price. Inside quotes is constructed as the proportion of the time during which each dealer was at either the inside ask price or inside bid price less one over the number of market makers in the stock. Quote activity is the number of quote revisions in a stock made by each dealer divided by the total number of quote revisions across all dealers in each stock. *Dealershare* is constructed as reported volume divided by total volume for each stock-quarter for each dealer based on reported Nasdaq volume from the monthly market maker activity data. *ECNshare* is constructed as the sum of reported ECN volumes divided by total volume for each stock-quarter. Panel A reports correlations across all stock-dealer observations. Panel B reports correlations over all stock-dealers but only for ECNs. Panel C reports correlations only for the Instinet ECN and Panel D reports correlations for the summary measures (over all ECNs) across the sample of 1000 stocks. All reported correlations are statistically significant at the one percent level.

measures of dealer activity constructed from the Nastraq database based on all dealers (both ECNs and non-ECN dealers). Panel A shows that, as expected, there is a very strong correlation between *Dealershare* and all three measures of dealer activity from the Nastraq database. What is important to note is that the large positive correlations between all four measures imply that each of our Nastraq market activity measures are good instruments for reported Nasdaq volume shares *when Nasdaq volume shares are accurately measured*.

In Panel B of Table 2 we test whether there is a similar relation between these market activity measures for the sample of ECNs – the sample where the reported volume data is questionable. As we see in Panel B, *Dealershare* (based on reported Nasdaq volumes) is highly correlated with all three measures of Nastraq dealer activity even when we only consider ECNs. While Panel B reports lower correlations of our proxy with the inside volume and inside quote measures, the correlation with quote activity, at 0.781, is actually higher than the correlation computed with the market shares over all dealers. Even the lowest correlation estimate in Panel B is still 0.59. Thus, our *Dealershare* measure using ECNs only is highly correlated with a set of three independent measures of market activity which are in turn highly correlated with accurate measures of Nasdaq reported volume. This exercise demonstrates that, while our proxy for ECN activity may be biased towards zero, the cross-sectional variation in our measure is strong enough to provide a meaningful instrument for ECN activity.

For comparison, Panel C of Table 2 constructs these correlations for *Dealershare* based only on Instinet reported volume. Had we used Instinet volume exclusively as our proxy for ECN activity, we would expect a generally weaker instrument than the proxy we actually do employ, which uses the total volume of nine different ECNs. Panel C demonstrates that this is likely the case. Note that even with the underreporting problem associated with Instinet volume, the Nasdaq reported market share (*Dealershare* for Instinet only) is still a relatively strong instrument with correlations between 0.38 and 0.63.

Finally, Panel D reports the correlations for *ECNshare*, our primary measure of ECN activity. For this aggregate measure across all ECNs, we find a strong positive correlation between *ECNshare* and the three measures of market maker activity constructed from the Nastraq data with correlations between 0.73 and 0.83. Thus, the variable that we use throughout the paper, *ECNshare*, is a strong instrument for ECN activity and we do not expect any bias in our tests from using *ECNshare* as a proxy for ECN activity.¹²

The underreporting problem may cause us to observe a level of ECN activity that is lower than actual. Nevertheless, *ECNshare* appears to be a strong instrument for ECN activity and therefore serves as a meaningful metric in exploring the effect of ECN activity on market quality. Finally, it is important to note that our instrumental variables approach is conservative in the sense that relying on a noisy instrument biases us against finding significant results.

2. The growth of ECN activity

Fig. 1, Panel A presents a time-series plot of our *ECNshare* variable. While the levels of our proxy may not be reflective of true ECN activity, the time-series properties are.

¹² Unreported, we also test for any time-series bias by comparing the magnitude of the correlations between time periods. Overall, we find little evidence that correlations change over time. As a result, the information content of our instrument should be relatively stable over time as well. However, in Section 5 we perform all of our tests quarter-by-quarter and show that our results are not sensitive to any particular time period.

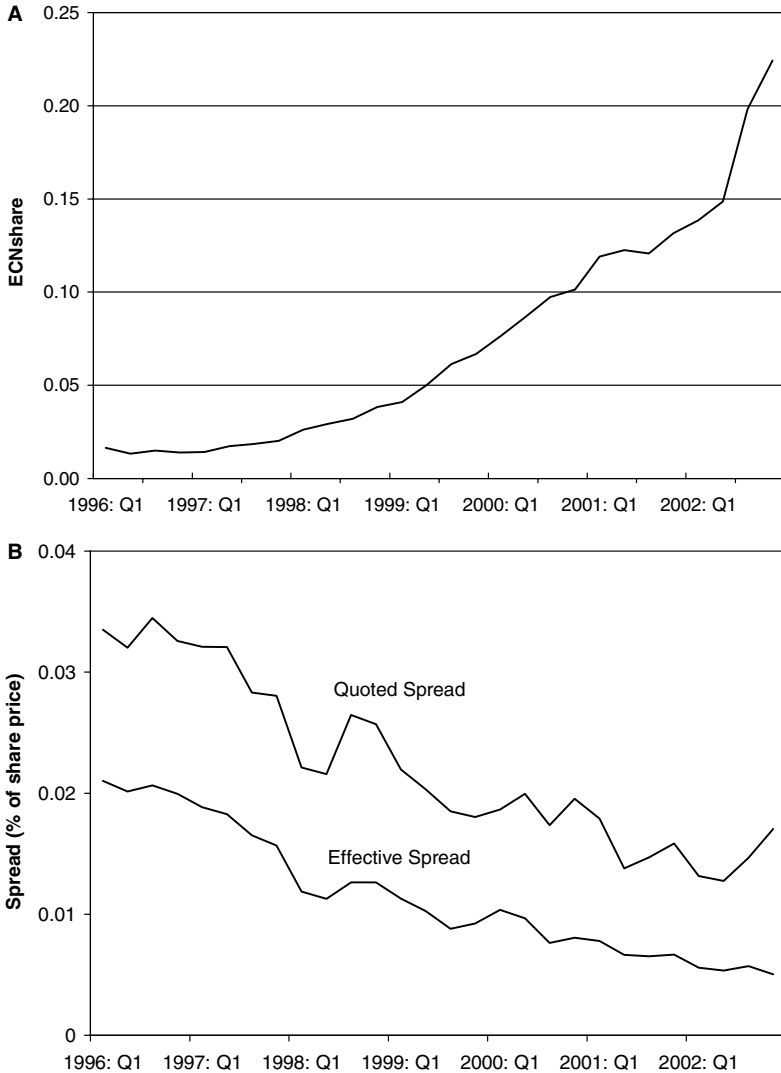


Fig. 1. Quarterly time series of ECN activity proxy (ECNshare) and bid-ask spreads. Panel A. This figure presents a time series of the market share of Electronic Communications Networks. The market share of ECNs is constructed as the sum of reported ECN volume divided by total Nasdaq volume over all market makers for each stock-quarter. Panel B. This figure presents a time series of quarterly average quoted, effective, and relative spreads for all Nasdaq stocks between January 1996 and December 2002. Average quoted spreads are computed for each stock based on all Nasdaq NBBO bid and ask quotes during regular trading hours. Effective spreads are computed using Roll's (1984) procedure based on all intraday transactions. The presented time series are equally weighted cross-sectional means over all stocks in each quarter.

Panel A shows a clear significant growth trend in our measure of ECN activity. At the beginning of the sample (January 1996), *ECNshare* is less than 2%. By the end of the sample (December 2002), our measure increases to over 22% – a 10-fold increase over the sample period.

There are also significant differences in the time series patterns in *ECNshare* by volume decile.¹³ For example, in 2000 *ECNshare* is roughly five times larger in high volume (decile 10) stocks compared to low volume (decile 1) stocks (3.33 compared to 16.37) and the same basic pattern holds in each year. ECN activity has increased significantly over time in all volume deciles. Between 1998 and 2002, the absolute value of *ECNshare* in the largest volume decile increased by over 16 percentage points (from 4.72 to 21.17). Even in the lowest volume stocks, there is still a substantial increase from 1.42 to 5.49. While only a proxy for ECN activity, these values are consistent with other estimates of ECN growth rates and demonstrate an overall regime shift during this period from clearing trades through traditional dealers to electronic trading platforms.

3. The effect of ECN activity on trading costs and liquidity

3.1. The importance of ECN activity

ECNs provide a potentially low-cost alternative to clearing trades through a traditional broker/dealer. Given the dramatic growth of ECN activity from 1998 to 2002, this period provides a natural experiment to test whether the increase in ECN activity has translated into lower trading costs for investors. In this section, we test whether the growth of ECN activity is associated with a reduction in trading costs through smaller spreads and greater quoted depths.

Fig. 1, Panel B presents a time-series of quarterly average quoted and effective bid–ask spreads for our sample of Nasdaq securities from Q1 1996 to Q4 2002. The time-series trend has the inverse shape of the time-series pattern in our ECN activity measure.¹⁴ In fact, the correlation between the time series of quoted spreads and *ECNshare* is -0.92 (significant at the 1% level). This relationship at least suggests that the growth of ECN activity may have had an effect on lowering trading costs.

To test whether the growth of ECNs has had a statistically significant effect on spreads, we perform a multivariate regression analysis. Specifically, we regress the quoted, effective, and relative spreads for stock i in quarter t on a variety of controls that may affect bid–ask spreads in addition to *ECNshare*. Following Stoll (1992) and Wahal (1997), we include trading activity (share turnover), share price, size (market capitalization), and return volatility. Market capitalization is constructed as the quarterly closing price times the number of shares outstanding. Share turnover is measured as total quarterly volume divided by shares outstanding. Volatility is defined as the standard deviation of daily returns for stock i in quarter t . Log transformations of the variables are used to reduce skewness. To avoid the confounding effects on market quality driven by the new order handling rules in 1997, we restrict our analysis to the sample period 1998–2002. To account for unobserved heterogeneity in liquidity for our sample of firms, we estimate a fixed-effects model that allows for within-firm variation in the parameters as well as systematic time effects. The regression equation is given by

¹³ Details and supplemental tables are available from the authors on request.

¹⁴ Much of the decline in spreads during early 1997 may stem from the phase-in of the new order handling rules. Barclay et al. (1998) and Weston (2000) find that the reforms reduced spreads by over 30%. In addition, spreads were also affected by reduction in the minimum tick size from one-eighth to one-sixteenth in June (1997).

$$\ln(\text{Spread})_{i,t} = \alpha_i + \delta_t + \beta_1 \ln(\text{ECNActivity})_{i,t} + \beta_2 \ln(\text{Turnover})_{i,t} + \beta_3 \\ \times \ln(\text{SharePrice})_{i,t} + \beta_4 \ln(\text{Size})_{i,t} + \beta_5 \ln(\text{Volatility})_{i,t} + \varepsilon_{i,t}, \quad (1)$$

where α_i and δ_t refer to the firm-specific and time-specific fixed-effects, respectively. In addition to the within-group estimates based on Eq. (1), we also estimate between-group coefficient estimates for our panel of firms. The between-group estimator computes the time-series averages of the dependent and independent variables in (1) and estimates the coefficients based purely on the cross-sectional dependence across firm-groups. The between-group panel data coefficients are similar to those obtained from a pooled-OLS estimate, but are estimated efficiently and without the size distortion of temporally correlated regressors.

If ECNs are effective low-cost competitors to traditional Nasdaq dealers, we expect a negative sign on both the within-group and between-group estimate of β_1 . In addition, we test the same regression equation for the log of the average quoted depth. If ECNs improve liquidity, *ceteris paribus*, we expect our measure of ECN activity to be positively related to market depth.¹⁵

Table 3 presents the results of our multivariate panel-data regressions on spreads and depths. We present both within-group and between-group coefficient estimates for our four measures of market quality. In all eight specifications we find that our measure of ECN activity has a statistically and economically significant positive effect on market quality. The between-group coefficient estimate of β_1 for the relative spread is -0.132 (Table 3, column 3). Given a one standard deviation change in the log of *ECNshare* (which is approximately equal to one), this coefficient estimate suggests that the relative spread would decrease by 0.132 or roughly 6.8% of the average 1.9% relative spread. For the quoted spread, we find that a one-standard deviation increase in ECN activity lowers the log of the quoted spread by 1.7% (since the coefficient measures the change in the log of a variable, this is roughly equal to a percentage change). There is also an economically large effect based on the within-group fixed effect estimators. For example, column 4 of Table 3 reports that, for the average firm in the sample, a similar increase in ECN activity is associated with a 3.1% decrease in the average firm's relative spread.

However, since our proxy for ECN activity is biased downwards, we must be careful about direct interpretation of the magnitudes of the parameter estimates. If the cross-sectional and time-series properties of true ECN activity are preserved by our instrument for ECN activity, we would expect that the downward bias would yield conservative parameter estimates – or that a one standard deviation increase in ECN activity would lower the relative spread by *at least* 6.8%. At minimum, we find a negative and statistically significant relationship between *ECNshare* and spreads.

¹⁵ Our results are not generally sensitive to the selection of specific control variables. For example, using total assets or total sales as alternative controls for firm size does not change the results. Similarly, using total share volume, dollar volume, or the number of trades instead of (or in addition to) share turnover do not qualitatively change the results. In addition, some of these variables (e.g., spreads and ECN activity) may be jointly determined. As a result, we have also tested specifications using an instrumental variables approach where lagged values of the regressors are used as instruments. The results are not qualitatively changed using this approach. However, these results should be interpreted with caution as many of the variables are highly persistent. As a result, it may be difficult to interpret the direct effects as causations rather than associations.

Table 3
Panel data regression results

	Ln(Quoted spread)		Relative spread		Relative effective spread		Ln(Quoted depth)	
	Between-group (1)	Within-group (2)	Between-group (3)	Within-group (4)	Between-group (5)	Within-group (6)	Between-group (7)	Within-group (8)
Ln(<i>ECNshare</i>)	−0.017 (−2.48)	−0.033 (−16.06)	−0.132 (−7.77)	−0.058 (−12.28)	−0.011 (−3.30)	−0.025 (−2.64)	0.066 (22.66)	0.069 (7.41)
Ln(Average turnover)	−0.318 (−102.64)	−0.271 (−163.84)	−0.661 (−86.44)	−0.505 (−134.59)	−0.130 (−49.12)	−0.174 (−40.49)	0.145 (−62.49)	0.116 (27.77)
Ln(Average share price)	0.968 (162.14)	1.054 (253.42)	0.108 (7.32)	−0.094 (−9.96)	−0.034 (−5.13)	−0.059 (−7.08)	−0.510 (−87.28)	−0.514 (−63.74)
Ln(Market capitalization)	−0.356 (−111.41)	−0.421 (−109.12)	−0.496 (−62.79)	−0.329 (−37.55)	−0.121 (−19.70)	−0.156 (−35.16)	0.109 (20.07)	0.138 (32.00)
Ln(Volatility of returns)	0.462 (64.52)	0.448 (149.69)	1.120 (63.35)	0.935 (137.79)	0.403 (84.29)	0.505 (50.87)	−0.293 (−69.68)	−0.408 (−42.24)
Observations	49,307	49,307	49,307	49,307	49,307	49,307	49,307	49,307
<i>R</i> -squared	0.917	0.825	0.876	0.628	0.427	0.739	0.440	0.702

This table presents the results of panel data regressions on four measures of liquidity and trading costs. The sample includes all Nasdaq stocks each quarter with a price greater than \$5 per share and a market capitalization greater than \$20M. The sample period extends from the first quarter of 1998 to the fourth quarter of 2002. Our proxy for ECN activity is constructed as the sum of reported ECN volumes divided by total volume for each stock-quarter. Average share turnover is constructed as total share volume over the quarter divided by the average number of shares outstanding over the quarter. Average share price is the average daily closing price during the quarter. Market capitalization is the average daily closing price times shares outstanding. Volatility is computed as the standard deviation of daily returns over each stock-quarter. Quoted spreads are computed as the average of the difference between ask and bid prices over all posted quotes for each stock during the sample period. Relative spreads are the quoted spread, scaled by midpoint price. Effective spreads are computed using the Roll (1984) serial covariance estimator. Depth is constructed as the average of the number of shares quoted at the inside bid and ask over all quotations during the quarter. *T*-statistics are reported in parentheses below coefficient estimates.

In addition to reducing spreads, the estimation suggests that ECN activity is also associated with a larger quoted depth. The between-group coefficient on quoted depth suggests that a one standard deviation increase in *ECNshare* is associated with a 6.6% increase in the quoted depth (Table 3, Column 7). Again, caution must be exercised in interpreting the magnitude of the coefficient, but at a minimum we find a positive and statistically significant relationship between quoted depth and ECN activity. Both results support the hypothesis that ECNs may have improved the overall quality of the Nasdaq market. This result is in contrast to Barclay et al. (2003) who find that (for June 2000) ECNs may have a negative effect on depth. This difference may stem from the longer time period and greater variation in depth over our sample.

In all specifications, the control variables generally have the expected signs with an increase in market capitalization and turnover leading to lower spreads and greater depths while an increase in volatility is associated with an increase in spreads and lower depths. Share price is positively associated with quoted spreads but negatively associated with relative spreads and depths. This result makes sense since larger priced stocks should have larger absolute spreads, but lower spreads relative to their share price. Time fixed-effects coefficients (not reported in the table) are jointly significant and capture the systematic negative time-series trend shown in Panel B of Fig. 1.

Overall, these results support the hypothesis that ECNs have a statistically significant effect on lowering transactions costs in dealer markets like the Nasdaq. We find that lower spreads and greater depths are associated with an increase in ECN activity. These findings suggest that the cost-competition effect may outweigh the potentially negative effects of fragmentation and anonymity. It appears ECNs provide a more efficient mechanism for providing competition from limit orders, beyond that provided by traditional dealers. Interpreted in this way, the competitive effect of ECNs may be viewed as a continued strengthening of the limit order competition that resulted from the 1997 reforms.

3.2. The importance of ECN access fees

Our comparisons of market quality based on bid–ask spreads and quoted depths could be somewhat misleading given the different fee structures among Nasdaq market participants. Over the sample period, Nasdaq broker–dealers do not usually charge explicit transactions fees. Instead, traditional brokers report trade prices that have an implicit fee built into them. On the other hand, ECN prices (and, of course, all public limit order prices) are exclusive of fees. That is, an institution pays the ECN price and then pays an explicit access fee on top of that. As Sofianos (2002) points out, this makes it difficult to compare measures based on transaction prices (like effective spreads) since reported prices are a mix of both fee and non-fee adjusted prices.

The access fees for some ECNs are not very economically large. For example, Instinet typically charges institutions two or three cents per share while Island charges 0.5 to 0.25 cents (or lower).¹⁶ Nevertheless, the cost-competition advantage that we document in Section 3.1 must still be considered in light of the access fees typically charged to institutions. To gauge the magnitude of this effect, consider that the average relative spread in our sample is roughly 1.9% on an average 20 dollar stock, or about 38 cents per share. According

¹⁶ As reported in Sofianos (2002) and Island's website (<http://www.isld.com>) at the time of our sample period.

to the point estimates presented in Section 3, a large increase in ECN activity (e.g., two-standard deviations) would reduce the relative spread by 13.6% (6.8% times 2 standard deviations) from 1.9% to roughly 1.67%. This would translate into a reduction in quoted spreads from 39 cents to about 33 cents on a 20 dollar stock for a net decrease of 6 cents. While this is, admittedly, only a rough estimate, it seems clear that ECNs have a net positive effect on market quality as long as the access fees are only fractions of a penny. However, for smaller changes in ECN activity or for coefficient estimates closer to the lower end of our confidence interval, the effect of a 2 penny Instinet fee could be large enough to outweigh the beneficial effect on dealer competition.

4. Market maker entry/exit

In the previous section, we demonstrate that the dramatic growth of ECNs during 1998–2002 is coincident with a decline in trading costs. Our multivariate tests suggest that competition from ECNs is associated with lower spreads and increased depths on the Nasdaq. If ECNs have successfully increased the competitiveness of Nasdaq trading then, *ceteris paribus*, we should expect to see exit from the market for market making. In fact, Wahal (1997) finds that narrower spreads lead to market maker exit for a large sample of Nasdaq stocks while Weston (2000) finds that competition from limit orders (introduced by the 1997 market reforms) causes dealers to exit. In this section, we test whether competition from ECNs has a similar effect on the number of traditional Nasdaq market makers.

In order to test whether ECNs have affected the marginal profitability of traditional Nasdaq dealers, we test whether the number of dealers making a market in a given stock is affected by competition from ECNs. However, Ellis et al. (2002) find that many registered dealers hardly make a market in many stocks they cover. Further, Schultz (2003) finds that many regional dealers often make markets in local firms, regardless of their ability to attract order flow or compete on spreads. Since there is considerable variation in the incentive structure among different types of dealers, we break the total number of dealers into five broad categories based on their self descriptions in the Securities Industry Association Directory.¹⁷ The five categories are:

1. Wholesalers (e.g., Madoff Securities, Knight Securities).
2. National Retail Dealers (e.g., Merrill Lynch, OLDE discount).
3. Institutional Dealers (e.g., Goldman Sachs, Bear Sterns).
4. Regional Dealers.
5. ECNs.

Fig. 2 presents a time series of the market share of total Nasdaq volume for the five different categories of dealers from January 1998 to December 2002.¹⁸ The growth of ECNs (as in Fig. 1) is evident. By the fourth quarter of 2002 our proxy for ECN activity, *ECN-share*, is larger than the market share of both wholesalers and national retailers. This figure highlights where the competition from ECNs may have the greatest effect. Casual inference

¹⁷ We thank Paul Schultz for providing this breakdown of dealer types.

¹⁸ Institutional and regional dealer market shares show little time series variation and were combined for clarity of presentation.

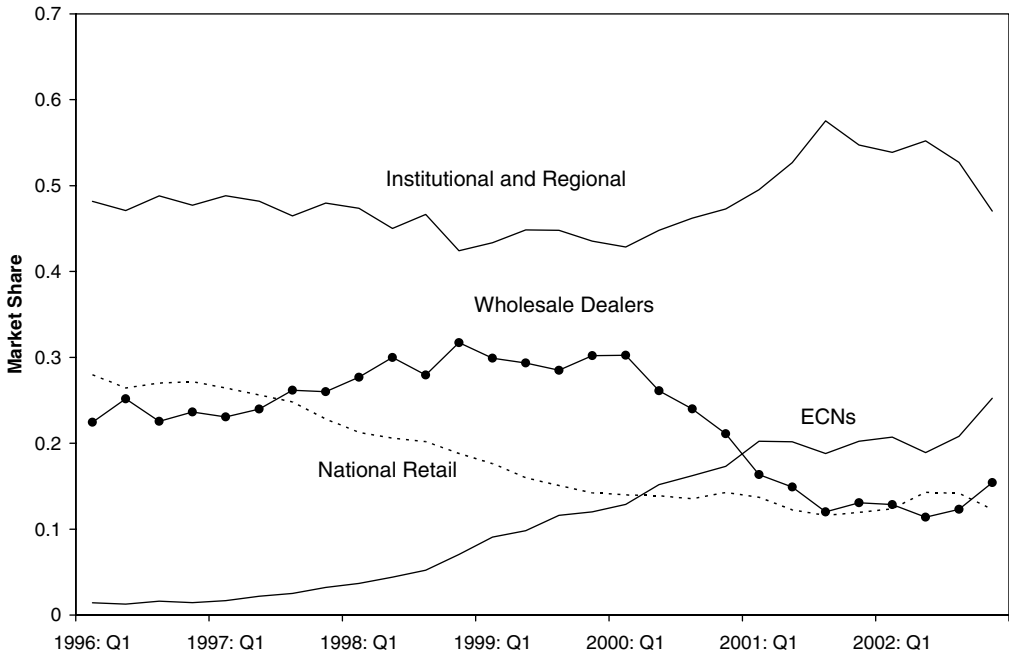


Fig. 2. Time series of market shares by type of dealer. This figure presents a time series of the market share of Nasdaq dealers by dealer type. Dealer types are based on self-descriptions in the Securities Industry Association directory. Market shares are constructed for each dealer type as the ratio of total reported volume over all dealers divided by total Nasdaq volume for each stock-quarter.

suggests that the growth of ECN activity may have a greater effect on national retailers and, at least in the latter half of the sample period, on wholesalers as well.

As in Wahal (1997) we model the number of market makers across all stocks using turnover, share price, market capitalization, return volatility and average quoted spreads as control variables.

However, our statistical analysis of the number of market makers faces some econometric challenges. Consistent with Wahal (1997) and Weston (2000) we find a significant amount of persistence in the quarterly number of market makers in a particular stock. Further, the discrete integer nature of the dependent variable makes standard OLS analysis difficult to interpret. As a result, we compute the percentage change in the number of market makers for each dealer category at the quarterly frequency. This procedure eliminates both the problem of serially correlated errors as well as firm specific fixed effects.¹⁹

Table 4 presents the results of regressions on the change in number of dealers in a given stock-quarter on quarterly changes in the control variables as well as year dummies. Overall, we find that the total number of registered dealers is affected by the presence of ECNs. However, there is considerable variation by the type of Nasdaq dealer. Regional dealers are only marginally affected by the presence of ECNs. There are a slightly higher number

¹⁹ Our results are robust, however, to a non-linear fixed-effects Poisson regression specification. Coefficient estimates from this methodology are difficult to directly interpret. We present the percentage change regressions for simplicity.

Table 4
The effect of ECN activity on market maker entry and exit

	% Change in the number of market makers				
	Total (1)	Regional (2)	Wholesalers (3)	National retail (4)	Institutional (5)
Ln(<i>ECNshare</i>)	0.021 (14.12)	0.028 (13.57)	-0.061 (-5.89)	-0.147 (-7.04)	0.075 (5.40)
Ln(Average turnover)	0.135 (71.92)	0.181 (68.13)	0.077 (6.83)	0.268 (10.78)	0.331 (19.30)
Ln(Average share price)	-0.039 (-10.35)	-0.049 (-9.99)	0.731 (25.31)	2.207 (28.20)	1.048 (20.55)
Ln(Market capitalization)	0.104 (-0.05)	0.136 (-0.07)	-0.478 (-0.10)	-1.770 (-0.25)	-0.894 (-0.18)
Ln(Volatility of returns)	27.488 (-20.48)	27.824 (-19.50)	-16.460 (-7.91)	-22.316 (-8.51)	-17.366 (-8.12)
Relative bid-ask spread	0.033 (15.81)	0.043 (14.34)	0.115 (9.63)	0.062 (2.63)	0.084 (4.83)
Constant	-0.009 (-2.75)	0.083 (18.41)	-1.461 (-61.97)	-3.436 (-69.96)	-3.582 (-76.90)
Observations	38,598	38,598	38,598	38,598	38,598
R-squared	0.397	0.438	0.4380	0.116	0.052

This table reports estimates of regressions relating quarterly changes in the number of market makers to quarterly changes in ECN activity. Each column presents regression results for different subsets of market makers. The sample includes all Nasdaq stocks each quarter with a price greater than \$5 per share and a market capitalization greater than \$20M. The sample period extends from the first quarter of 1998 to the fourth quarter of 2002. Our proxy for ECN activity is constructed as the sum of reported ECN volumes divided by total volume for each stock-quarter. Average share turnover is constructed as total share volume over the quarter divided by the average number of shares outstanding over the quarter. Average share price is the average daily closing price during the quarter. Market capitalization is the average daily closing price times shares outstanding. Volatility is computed as the standard deviation of daily returns over each stock-quarter. Relative spreads are computed as the quoted spread (which is the average difference between the quoted bid and ask prices over all posted quotes) divided by the prevailing midpoint of the spread. *T*-statistics are reported in parentheses below coefficient estimates.

of regional dealers, on average, for stocks with larger ECN activity. This result is consistent with previous findings that regional dealers make markets for reasons other than trading profits. Institutional dealers are also positively affected by ECN activity. However, both wholesale dealers and national retail dealers are negatively associated with ECN activity. Wholesalers and national retail dealers (such as Bernard L. Madoff Securities) often engage in payment for order flow. Since ECN trading (along with the recent regulatory changes) results in tighter bid ask spreads, it may be less profitable for these types of dealers to purchase order flow.

5. Further tests of ECN activity

5.1. Long-run change regressions

Our analysis in the previous section tests whether there is an effect of ECN activity by looking at quarterly data on ECN activity. However, it is always difficult to know whether

such regressions are correctly specified and whether the joint determination of endogenous variables has an effect on our estimates. In this sub-section we take a long-run view of the Nasdaq market and test whether changes in ECN activity between the start of our sample (Q1 1998) and the end of our sample (Q4 2002) are responsible for the improvements in market quality. To perform this analysis, we restrict our sample to 1310 firms for which we have data over the entire 5-year sample period. We then construct the total percentage change in each of the variables over the 19 quarters. To reduce the skewness induced by constructing percentage changes, we winsorize the 1% tails of the data. We use the same control variables as in the previous sections.

Our results (not reported in a table) are generally consistent with the analysis presented above. The parameter estimates suggest that firms with an increase in *ECNshare* one standard deviation above the mean experienced a reduction in quoted spreads of 9.46% (similarly we estimate a reduction of 19% for relative spreads, and 8.5% for effective spreads). Our analysis also finds that a similar increase in *ECNshare* is associated with a 13% decrease in market concentration. In contrast with the findings presented in Table 3, we do not here find statistical significance on the change in quoted depth, although the estimated coefficient on that variable is still positive. Consistent with the analysis in Section 3, we find that firms with a larger increase in our measure of ECN activity had a larger increase in the number of market makers but a significant decrease in the number of wholesale and national retail dealers. Overall, this exercise points out that the effects measured in Sections 3 and 4 are robust in a more conservative estimation based on five-year changes for a restricted sample of stocks.

5.2. Dynamic panel data regressions

In all of the previous analyses, we implicitly assume that the control variables fully capture the amount of systematic variation in the dependent variables. However, given the strong persistence in our variables, there may be dynamic effects that are not captured in our specification. For example, the best estimate of spreads in period t may be some function of spreads in period $t - 1$. If the serial correlation in our dependent variable is not completely attributed to serial dependence in our regressors, then our inference could be based on overstated t -statistics. To explicitly account for such dynamic effects, we re-specify our estimation equation to account for dynamic effects. Since the autocorrelation induces dynamic correlation with the error term, we follow Anderson and Hsiao (1981) and Arellano and Bond (1991) in using a set of instruments based on lagged changes to identify the parameters. Following these estimation methods we continue to find a significant negative association between our measure of ECN activity and our measures of market quality (not reported in a table). However, we must be somewhat cautious in our inference since these estimation procedures place severe restrictions on our sample (because they require many lagged values) and significantly reduce our sample size. Nevertheless, this analysis suggests that our inferences in Section 3 are not driven by spurious statistical inference caused by serial correlation in our dependent variable.

5.3. Vector auto regressions for aggregate measures

The analysis presented in the previous sub-sections analyzes the effect of increasing competition from ECNs only at firm level. In this sub-section we test the hypothesis that

aggregate growth of ECN trading is associated with a systematic improvement in Nasdaq market quality. Here, we construct a measure of aggregate *ECNshare* by computing the percentage of total trading volume associated with ECNs over all stocks in our sample. We also compute average bid–ask spreads, depth, turnover, and return volatility, etc., over the sample of stocks – similar to our analysis in Section 1. To allow for endogeneity of the systematic market measures, we model the joint determination of the time-series based on an unrestricted vector autoregression (VAR). Our estimation method involves specifying a vector, Y_t , of four possibly endogenous variables as

$$Y_t = \begin{bmatrix} ECN \text{ Activity}_t \\ Relative \text{ Spread}_t \\ Share \text{ Turnover}_t \\ Return \text{ Volatility}_t \end{bmatrix}.$$

We then model the evolution of the data as a multivariate VAR as in

$$Y_t = \alpha + \gamma t + Y_{t-1}\Phi_1 + \dots + Y_{t-p}\Phi_p + U_t, \quad U_t \sim IID(0, \Omega). \quad (2)$$

For our specification, we use the AIC criteria to choose p , the number of lags.²⁰ With the estimates from our VAR (not reported in a table), we compute impulse response functions following Hamilton (1995). These results indicate that a positive shock to *ECNshare* has a significant effect on aggregate trading costs for up to three months. We obtain qualitatively similar results for our other measures of market quality. This analysis is consistent with our previous evidence that ECN activity provides a source of competition to traditional dealers and is associated with a systematic improvement in Nasdaq market quality.

5.4. NYSE-Matched sample comparison

In Section 3 we present evidence of both a time-series and cross-section association between ECN activity and liquidity for Nasdaq stocks. However, it is possible that these effects are driven by a secular change in the trading environment, rather than by direct competition from ECNs. For example, our measure of ECN activity could simply proxy for the increasing role of institutional trading over time or for the technological development of automated trading systems which may have improved liquidity. In this section, we employ a matched-sample comparison of NYSE and Nasdaq stocks to disentangle the role of ECN activity from any secular trend in US equity markets over our sample period. Since ECNs have captured a large market share of Nasdaq trading volume but only a small fraction of NYSE trading volume we can test whether the decline in spreads that we document are driven by market wide changes that would affect both NYSE and Nasdaq stocks, or by changes that are specific to the Nasdaq market.

To construct such a test, our strategy is to first form an NYSE-adjusted measure of liquidity for each of our Nasdaq observations and then to repeat the analysis of Section 3 on our NYSE-adjusted measures of liquidity. For example, rather than examine raw Nasdaq spreads for each firm-quarter, we form NYSE-adjusted spreads by taking the difference between the Nasdaq spread and the prevailing spread of a similar NYSE stock. If the results that we document are part of a broad secular trend, then NYSE-adjusting our

²⁰ Our results are generally insensitive to the number of lags chosen.

liquidity measures should remove the effect. Of course, the threat of potential ECN competition could cause improvements in NYSE liquidity provision over our sample period, even if ECNs have only a small market share in listed stocks. However, it seems reasonable that the effect of a threat is lower than that of the actual participation. To the extent that this is true, we expect greater ECN activity in a stock and over time to be associated with better liquidity even after controlling for comparable changes in NYSE liquidity.

Our first challenge in constructing the test is to form a matched-sample of stocks by linking each Nasdaq firm, each quarter, to an NYSE stock with similar characteristics. Following previous studies (e.g., see Weston, 2000; Huang and Stoll, 1996; Bessembinder, 1999), we form our matched sample by finding the NYSE firm with the closest share price, market capitalization, and return volatility over the quarter.²¹ With our sample of matched pairs, we construct NYSE-adjusted liquidity measures by taking the simple difference between the Nasdaq and matched NYSE liquidity measures. With each new NYSE-adjusted measure of liquidity, we replicate the analysis in Section 3 by testing the hypotheses that *ECNshare* has a positive effect on liquidity.

Overall, our results (not reported in a table) are consistent with the analysis presented in Section 3. We continue to find that greater ECN activity reduces bid ask spreads, regardless of how the spreads are measured. For example, we find that, for the average firm in the sample, a one-standard deviation increase in ECN activity is associated with a decrease in NYSE-adjusted Nasdaq relative spreads by roughly 9.5%. Results for quoted and effective spreads are also statistically significant, though somewhat smaller in economic magnitude. We also find that ECN activity continues to have a positive effect on market depth, though the effect for NYSE-adjusted depth is not statistically significant. Overall, the results of our NYSE-matched sample analysis support our previous results. Even when we remove the effect of any secular trend in liquidity over our sample period, we continue to find a positive association between ECN activity and Nasdaq market liquidity.

5.5. Time-series variation in the effect of ECNs on market quality

Given that the Nasdaq market experienced unusually high volatility over our sample period as well as the change to decimal pricing, we want to be sure that our results are robust over time and not driven by any specific sample period. In this sub-section, we test the hypothesis that the effect of competition from ECNs varies over time. To accomplish this, we estimate our multivariate regression quarter-by-quarter for each of the 20 quarters in our sample. The results from this analysis (not reported) suggest that our results are generally robust over time. While the magnitude of the effect varies slightly from quarter to quarter, our results show that ECNs have a negative effect on relative bid–ask spreads in nearly every quarter. The results for the quoted spread and relative effective spread are also generally similar, though there is somewhat less statistical significance quarter to quarter.

²¹ Trading volume is also an important matching characteristic. However, the comparison of volume between the Nasdaq and NYSE presents some problems. In dealer markets, trades are often immediately turned around by the market maker and thus double counted, making it hard to compare with volume in auction markets (see Atkins and Dyl, 1997). Further, the magnitude of this bias may even change over time (see Anderson and Dyl, 2005). However, since volume is such an important determinant of the spread, we replicate our results by also matching based on volume by comparing NYSE volume to a scaled-Nasdaq volume where Nasdaq volume is divided by two. Our results are qualitatively similar.

Interestingly, we find the effect of ECN competition has a consistent positive effect on quoted depths prior to decimalization (third quarter, 2001) but a significantly negative effect after decimalization. This is an important result because it shows that the change to decimalization has had a significant effect on quoted depths, even for ECNs. Thus, our results are consistent with a number of recent studies that examine the effect of decimal prices on various measures of market quality. While we still find that ECNs are associated with improved competition through tighter spreads post-decimalization, we also find that they are associated with lower quoted depths. However, this is not direct evidence that ECNs lead to a deterioration of market quality post-decimalization since stocks with large ECN activity may have greater depth outside of the inside spread. It is therefore hard to interpret the inside depth in a decimal environment since so much liquidity may lie outside of the inside quotes.

6. Conclusion

This paper investigates the growth of trading on electronic communication networks between 1996 and 2002. Consistent with anecdotal evidence and rough measures reported by the SEC and Nasdaq, our measure of ECN activity shows a significant increase over the sample period. We find that ECNs compete directly with traditional Nasdaq dealers and that increases in ECN activity are associated with lower average quoted, effective, and relative spreads. The effect of ECNs on quoted depth is less clear, though we find weak evidence that ECN activity is positively associated with the average quoted depth at the inside (prior to decimalization). Further, the increased competition from these low-cost competitors, *ceteris paribus*, may have forced some traditional dealers out of the market for market making. Our results show that the cost-competition effect appears to outweigh any potentially negative effects on market quality driven by fragmentation or anonymity.

The ability of investors to trade with each other directly is a clear benefit to investors without any noticeable deterioration of market quality along other dimensions. However, it is important to understand the limitations of these results. Our study focuses only on measures of the spread and market quality and not on the total costs that a large trader or institution might face. As a result, our findings relate mostly to competition for small trades. The benefit to large institutional traders is more difficult to gauge, but our results do not suggest any deterioration in market quality from ECN activity. Given this caveat, we find that the growth of electronic communication networks has improved competition on the Nasdaq.

Our results have policy implications for the optimal structure and performance of exchanges. For example, the NYSE's trade-through rule technically forces best-price competition between NYSE specialists and floor brokers. However, ECNs still face a significant barrier to entry because of the slow speed and connectivity of the NYSE's Intermarket Trading System. Our results suggest that more effective competition from ECNs could benefit retail investors in listed stocks. While we provide evidence that ECNs have improved competition on the Nasdaq, it remains to be seen whether ECNs will provide significant competition to NYSE specialists and floor brokers.

Acknowledgements

We would like to thank George Allayannis, Alex Butler, Wake Epps, Gustavo Grullon, Hans Heidle, Yong C. Joe, Pete Kyle, Ingrid Werner, and seminar participants at the

University of Virginia, Rice University, and the Batten Conference at William and Mary for helpful comments. All remaining errors are our own.

References

- Anderson, Anne-Marie, Dyl, Edward, 2005. Market structure and trading volume. *Journal of Financial Research* 28, 115–131.
- Anderson, T.W., Hsiao, C., 1981. Estimation of dynamic models with error components. *Journal of the American Statistical Association* 76, 598–606.
- Arellano, M., Bond, S., 1991. Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Review of Economic Studies* 58, 277–297.
- Atkins, Allan, Dyl, Edward, 1997. Market structure and reported trading volume: Nasdaq vs. the NYSE. *Journal of Financial Research* 20, 291–304.
- Barclay, Michael J., 1997. Bid-ask spreads and the avoidance of odd-eighth quotes on Nasdaq: An examination of exchange listings. *Journal of Financial Economics* 45, 35–58.
- Barclay, Michael J., Kandel, Eugene, Marx, Leslie M., 1998. The effects of transaction costs on stock prices and trading volume. *Journal of Financial Intermediation* 7, 130–150.
- Barclay, Michael J., Hendershott, Terrence, McCormick, D. Timothy, 2003. Competition among trading venues: Information and trading on electronic communications networks. *Journal of Finance* 58, 2637–2666.
- Battalio, Robert, Greene, Jason, Jennings, Robert, 1997. Do competing specialists and preferencing dealers affect market quality? *Review of Financial Studies* 10, 969–993.
- Bessembinder, Hendrik, 1999. Trade execution costs on Nasdaq and the NYSE: A post-reform comparison. *Journal of Financial and Quantitative Analysis* 34, 387–407.
- Bessembinder, Hendrik, Kaufman, Herbert M., 1997. A comparison of trade execution costs for NYSE and Nasdaq-listed stocks. *Journal of Financial and Quantitative Analysis* 32, 287–310.
- Biais, Bruno, Bisiere, Christophe, Spatt, Chester, 2002. Imperfect competition in financial markets: Island vs Nasdaq. Carnegie Mellon University, Working paper.
- Bloomfield, Robert, O'Hara, Maureen, 1998. Does order preferencing matter? *Journal of Financial Economics* 50, 3–37.
- Christie, William, Schultz, Paul, 1994. Why do Nasdaq market makers avoid odd-eighth quotes? *Journal of Finance* 49, 1813–1840.
- Conrad, Jennifer, Johnson, Kevin M., Wahal, Sunil, 2003. Institutional trading costs and alternative trading systems. *Journal of Financial Economics* 70, 99–134.
- Easley, David, Kiefer, Nicholas M., O'Hara, Maureen, 1996. Cream-skimming or profit-sharing? The curious role of purchased order flow. *Journal of Finance* 51, 811–833.
- Ellis, Katrina, Michaely, Roni, O'Hara, Maureen, 2002. The making of a dealer market: From entry to equilibrium in the trading of Nasdaq stocks. *Journal of Finance* 57, 2289–2317.
- Fong, Kingsley, Madhavan, Anath, Swan, Peter L., 2002. Are upstairs markets Pareto improving? University of New South Wales, Working Paper.
- Garfinkel, Jon A., Nimalendran, M., 2003. Market structure and trader anonymity: An analysis of insider trading. *Journal of Financial and Quantitative Analysis* 38, 591–610.
- Hamilton, James L., 1979. Marketplace fragmentation, competition, and the efficiency of the stock exchange. *Journal of Finance* 34, 171–187.
- Hamilton, James L., 1995. *Time Series Analysis*. Princeton University Press.
- Heidle, Hans, Huang, Roger D., 2002. Information-based trading in dealer and auction markets: An analysis of exchange listings. *Journal of Financial and Quantitative Analysis* 37, 391–424.
- Huang, Roger D., 2002. The quality of ECN and Nasdaq market maker quotes. *Journal of Finance* 57, 1285–1319.
- Huang, Roger D., Stoll, Hans, 1996. Dealer vs. auction markets: A paired comparison of execution costs on Nasdaq and the NYSE. *Journal of Financial Economics* 41, 313–357.
- Kandel, Eugene, Marx, Leslie, 1999. Payments for order flow on Nasdaq. *Journal of Finance* 54, 35–66.
- Madhavan, Ananth, 1995. Consolidation, fragmentation, and the disclosure of trading information. *Review of Financial Studies* 8, 579–603.
- Mendelson, Haim, 1987. Consolidation, fragmentation, and market performance. *Journal of Financial and Quantitative Analysis* 22, 189–207.

- Roll, Richard, 1984. A simple implicit measure of the effective bid–ask spread in an efficient market. *Journal of Finance* 39, 1127–1139.
- Schultz, Paul, 2000. Regulatory and legal pressures and the costs of Nasdaq trading. *Review of Financial Studies* 13, 917–957.
- Schultz, Paul, 2003. Who makes markets? *Journal of Financial Markets* 6, 49–72.
- Simaan, Yusif, Weaver, Daniel G., Whitcomb, David K., 2003. Market maker quotation behavior and pretrade transparency. *Journal of Finance* 58, 1247–1268.
- Sofianos, George, 2002. The changing Nasdaq market and institutional transaction fees. Goldman Sachs Trading and Market Structure Research Report.
- Stoll, Hans R., 1992. Principles of trading market structure. *Journal of Financial Services Research* 6, 75–106.
- Wahal, Sunil, 1997. Entry, exit, market makers, and the bid–ask spread. *Review of Financial Studies* 10, 871–901.
- Weston, James P., 2000. Competition on the Nasdaq and the impact of recent market reforms. *Journal of Finance* 55, 2565–2598.