

Competition among Trading Venues: Information and Trading on Electronic Communications Networks

MICHAEL J. BARCLAY, TERRENCE HENDERSHOTT,
and D. TIMOTHY McCORMICK*

ABSTRACT

This paper explores the competition between two trading venues, Electronic Communication Networks (ECNs) and Nasdaq market makers. ECNs offer the advantages of anonymity and speed of execution, which attract informed traders. Thus, trades are more likely to occur on ECNs when information asymmetry is greater and when trading volume and stock-return volatility are high. ECN trades have greater permanent price impacts and more private information is revealed through ECN trades than through market-maker trades. However, ECN trades have higher ex ante trading costs because market makers can preference or internalize the less informed trades and offer them better executions.

TECHNOLOGICAL INNOVATIONS THAT ENABLE HIGH-SPEED, low-cost electronic trading systems are dramatically changing the structure of financial markets. Exchanges and markets around the world are merging or forming alliances to improve liquidity and reduce costs in the face of increased competition from each other and from these computerized trading systems. Trading volume on Electronic Communications Networks (ECNs) has grown rapidly over the past several years. ECNs are now involved in more than a third of Nasdaq trading volume and are attempting to increase their market share in NYSE-listed

* Barclay is at the Simon School of Business, University of Rochester; Hendershott is at the Haas School of Business, University of California at Berkeley; and McCormick is at Nasdaq Economic Research. We thank Rick Green (the editor), an anonymous referee, Robert Battalio, Ken Kavajecz, and seminar participants at the Financial Markets Research Center Conference on Market Quality at Vanderbilt University, Notre Dame University, the University of California at Berkeley, the University of California at Davis, the University of Rochester, the University of Texas at Austin, the 2001 NBER Market Microstructure meeting, the 2001 Workshop on Information Systems and Economics, the 2001 Western Finance Association meetings, and the 2002 American Finance Association meetings for helpful comments and suggestions. An anonymous referee provided numerous excellent ideas and Joel Hasbrouck provided many helpful suggestions regarding the nonstationary time-series analyses. A previous version of this paper titled "Electronic Communications Networks and Market Quality" won the New York Stock Exchange Award for the best paper on equity trading at the 2001 Western Finance Association meetings. The views expressed herein are not intended to represent the views or official policy of The Nasdaq Stock Market, Inc. Hendershott gratefully acknowledges financial support from the National Science Foundation. Any errors or omissions are the responsibility of the authors alone.

issues as well. ECNs offer the promise of greater operational efficiency, lower trading costs, improved limit order exposure, trader anonymity, and faster executions.

ECNs differ from Nasdaq market makers in several important dimensions. ECNs are open limit order books that facilitate customer-to-customer trades. ECNs compete with market makers by providing anonymous and immediate execution and by improving the quotes (sometimes using smaller tick sizes). Market makers, in contrast, enter into long-term agreements with brokers who preference or internalize their customers' orders regardless of the current quotes. The majority of small trades with market makers are preferenced or internalized and market makers contract with brokers to attract the most profitable trades. Market makers also negotiate with institutional investors to intermediate large orders.

Using 1 month of data from June 2000 that identifies all ECN and market-maker trades for a stratified sample of 150 Nasdaq National Market stocks, we explore how these differences between ECNs and market makers affect the trading choices of different types of traders. First, we examine the market conditions under which ECN trades are more likely to occur. We then investigate whether informed trades are more likely to occur with market makers or on ECNs, and study the role of ECN trading in the aggregate price discovery process. Finally, we document the trading costs on these electronic trading systems and compare them with the trading costs for market-maker trades. Together, these results provide new insights about the competition for order flow across trading venues that offer different services to their customers.

Trades on ECNs are smaller than trades with market makers and are more likely to occur during periods of high trading volume and high stock-return volatility. ECN trades are also more likely to occur when spreads are narrow and, not surprisingly, when ECNs are offering better prices. In the cross section, ECN trades are more prevalent in stocks that have high trading volume, large market capitalization, and fewer market makers.

The higher frequency of ECN trades during periods of high trading volume and high stock-return volatility suggests that ECNs attract a higher fraction of informed trades than market makers. There are several reasons why this might be the case. First, market makers are able to preference the less-informed order flow and may force the more-informed trades on to an ECN. In addition, especially in fast moving markets, informed traders may prefer the ECNs' speed of execution and pre- and post-trade anonymity.

To determine where informed trades are more likely to occur and to measure the relative importance of ECN and market-maker trades in the price discovery process, we decompose the stock-price variance into its trade-related and trade-unrelated components (Hasbrouck (1991a)), and show that ECN trading explains about two-thirds more of the stock-price variance than market-maker trades. To control for the contemporaneous correlation between ECN and market-maker trades, we also estimate vector autoregressions (Hasbrouck (1991b)) and measure the permanent price impact of trades in different venues. When we conservatively assume that all contemporaneous correlation between ECN and market-maker

trades is attributable to market-maker trading, we find that the permanent price impact of an ECN trade is 50% larger than the permanent price impact of a market-maker trade.¹

The higher frequency of informed trades on ECNs directly affects both ex ante and ex post trading costs. If market makers compete for the right to preference the less-informed orders, one might expect small trades with market makers to have lower ex ante trading costs than similar trades on ECNs. Indeed, small trades (less than 1,000 shares) that are executed by market makers have slightly lower effective spreads than small trades on ECNs, even after controlling for market conditions at the time of the trade. However, because ECNs trades are more informed, they have significantly smaller ex post trading costs, as measured by the realized spread.

The competition for large trades is more complex. Market makers can negotiate directly with institutional traders and price discriminate based on their inferences about traders' motives and information. Because these negotiations are repeated, market makers can quote competitive prices for large liquidity trades and impose costs on traders who take advantage of their private information. Our results show that market makers are extremely good at identifying large liquidity trades and protecting themselves against large informed trades. Following large trades with market makers, prices actually move in the direction opposite the trade (down following customer purchases and up following customer sales), which makes the realized spreads on these trades larger than the effective spreads.

Traders who demand liquidity for large positions, but who cannot credibly signal that their trades are uninformed, receive inferior prices from market makers. As an alternative to paying a large spread to a market maker, these traders can submit a limit order to an ECN. By posting a limit order on an ECN, the liquidity trader may receive the bid-ask spread rather than paying it, but he also bears the adverse-selection costs of trading with someone who has better information. Liquidity traders will use limit orders when the expected adverse-selection and opportunity costs are lower than the spread charged by a market maker. Large trades will execute on an ECN only when one of these "natural" counterparties is offering sufficient depth. Thus, although there are relatively few medium and large trades on ECNs, these trades have significantly lower effective spreads than similar trades with market makers.

Most ECNs use a smaller tick size than prescribed by Nasdaq.² The ability to submit limit orders on fractional ticks also has a differential effect on informed and uninformed traders. Informed traders need to trade before the market price

¹These results are consistent with the results on trading after hours in Barclay and Hendershott (2003). They find that after the close, information asymmetry is low and trades are large and tend to occur with market makers. Before the open, information asymmetry is high, trades are smaller, more likely to occur on an ECN, and are more informed.

²During our sample period, some ECNs used a tick size of 1/256 compared with 1/16 on Nasdaq. All quotes displayed on Nasdaq, including ECN quotes, are rounded to the nearest 1/16.

reflects their information. Uninformed traders can be more patient. Because they have no reason to believe the price will move against them, uninformed traders can profit by waiting for opportunities to get better executions. Thus, submitting a limit order to an ECN that improves the quote by a fractional tick will have little effect on the arrival of informed trades, but may attract an uninformed trade that is searching for better execution. We show that trades on fractional ticks have lower effective spreads, but larger realized spreads than trades on integer-ticks. Together, these results show that limit orders on fractional ticks do attract the less informed and consequently more patient traders.

Our results are related to the existing literature on multimarket trading. However, our results highlight important differences between ECNs and other trading venues, such as regional exchanges, that have skimmed order flow from primary markets. Generally, when a secondary market skims orders from the primary market, the secondary market skims the least informed and, consequently, most profitable orders. For example, Easley, Kiefer, and O'Hara (1996) show that the Cincinnati Stock Exchange attracts mostly uninformed orders in NYSE stocks and Bessembinder and Kaufman (1997) find this for the regional exchanges in general. In contrast, we find that trades on ECNs are more informed than trades with market makers. On Nasdaq, the market makers' preferencing and internalization agreements allow them to retain the less-informed retail orders. Because ECNs match customer orders without participating in the trades, the more informed orders spill onto the ECNs.

ECNs are open limit order books that create a hybrid dealer/auction market on Nasdaq. Thus, our results are related to the literature on limit order books and automated markets. For example, Domowitz and Wang (1994) study the distributions of price and the dynamics of a limit order book for exogenous order submission strategies. Glosten (1994) examines the optimal prices schedule by competing liquidity suppliers in multiple markets and Parlour and Seppi (2003) study competition between a specialist/limit order market and a pure limit order market. Examples of empirical analyses of limit order books include Coppejans, Domowitz, and Madhavan's (2001) investigation of the dynamics of liquidity provision in an electronic limit order futures market and Hollifield, Miller, and Sandås' (2001) estimation of the asset valuation implied by liquidity suppliers' order placements and executions.

Several previous papers have examined the effect of ECN quotes. For example, Barclay et al. (1999) analyze the impact of Nasdaq's new order-handling rules and show that ECN quotes play an important role in reducing trading costs. Simaan, Weaver, and Whitcomb (2003) show that ECN quotes are more likely to be on an odd tick than market-maker quotes and how quote anonymity affects market makers' quotes. Huang (2002) finds that ECN quote updates are more informative than market-maker quote updates. Because these studies focus on quotes rather than on trades, however, they are unable to examine the role of ECN trading in the price discovery process or determine the venue in which informed traders choose to trade. This paper is the first to use comprehensive data to examine

trading on ECNs.³ By examining ECN trades in addition to ECN quotes, we are able to focus on the demanders of liquidity, rather than on the suppliers of liquidity. We also are able to calculate more direct measures of market performance such as effective and realized spreads, price improvement, and the permanent price impact of trades.

The paper proceeds as follows. Section I provides an overview of ECNs. Section II describes our data. Section III examines investors' choices of trading venues. Section IV investigates the price discovery process and provides evidence that more informed trades occur on ECNs. Section V compares the cost of ECN and market-maker trades. Section VI concludes.

I. An Overview of ECNs

The SEC defines ECNs as "electronic trading systems that automatically match buy and sell orders at specified prices" (<http://www.sec.gov/answers/ecn.htm>) and describes ECNs as having "become integral to the modern securities markets" (U.S. Securities and Exchange Commission (2000), Part II, Section I.A). In August 2002, ECNs accounted for approximately 38% of total share volume and 41% of the dollar volume traded in Nasdaq securities. ECNs accounted for approximately 3% of total share and dollar volume in listed securities. In contrast, in 1993 ECNs accounted for only 13% of share volume in Nasdaq securities and only 1.4% of listed share volume (U.S. Securities and Exchange Commission (2000)).

Competing ECNs offer different fee structures and levels of service and cater to different investor clienteles. However, all ECNs provide the same basic transaction services. ECN subscribers submit limit orders that are posted on the system for other subscribers to view. The ECN then matches contra-side orders for execution. In most cases, the buyer and seller remain anonymous, as the trade execution reports list only the ECN as the contra-side party. Subscribers may use additional features of the ECN, such as negotiation or reserve size, and may have access to the entire ECN order book that contains important real-time information about the depth of trading interest.

When ECNs first developed, they served primarily as private trading vehicles for institutional investors and broker-dealers. The prices posted on ECNs by these professional traders often were better than the prices posted on Nasdaq. Because the ECNs were not integrated into the Nasdaq market, many investors, particularly retail investors, traded at prices inferior to those displayed by market makers and other subscribers on ECNs. Essentially this created a two-tiered

³In the cross section, Hasbrouck and Saar (2001) examine volatility's impact on limit orders on Island. Biais, Bisiere, and Spatt (2002) study quote competition between Nasdaq and Island and the profitability of limit order strategies. Conrad, Johnson, and Wahal (2002) compare trading costs for institutional trades on crossing networks and Instinet to those with traditional brokers. Weston (2001) uses ECNs' monthly activity to examine their evolving competitive role in Nasdaq.

market—the traditional public market, and the new ECN market with better prices and limited access.

In 1996, the SEC adopted new order-handling rules to integrate these markets. Before the adoption of the order-handling rules, market makers could post quotes in private ECNs that were better than the quotes they posted in the public market. This allowed market makers to segment their market, charging higher prices to retail customers and lower prices to more price-sensitive institutional investors. Under the new order-handling rules, market makers and specialists were required to reflect in their public quote any better prices that they placed on an ECN. The new order-handling rules had a large and immediate impact on the securities markets—trading costs fell dramatically, resulting in significant cost savings for investors (Barclay et al. (1999)).

Today, almost any investor can trade through an ECN, including retail investors, institutional investors, market makers, and broker-dealers. The recent proliferation of new electronic markets led the SEC to consider how to incorporate these trading venues into the national market system. In December 1998, the SEC adopted Regulation ATS to establish a regulatory framework for alternative trading systems to more fully integrate them into the national market system. The goals of Regulation ATS were to provide investors with access to the best prices, provide a complete audit trail and surveillance on alternative trading systems, and reduce the potential for market disruption due to system outages.

Once ECNs were fully integrated into the Nasdaq market, investors had to decide whether and when to utilize them. ECNs offer several potential benefits to investors. First, ECNs typically offer an advantage in the speed of execution. Traditional orders are sent first to a broker, either electronically or over the phone, who determines the market where they will be sent for execution. There are no SEC regulations concerning the time required to complete this task. Although trade executions are usually seamless and quick, they do take time. In fast-moving markets, investors using traditional brokers will not always receive the price they see on their computer screen or the price their broker quotes over the phone. By the time their orders reach the market, the price of the stock could be slightly—or very—different. The immediate execution offered by ECNs is one of their major selling points.⁴ Immediate execution and anonymity are closely related and both appeal to those investors who may have short-lived information about future price changes.

Second, ECNs sometimes offer better prices than Nasdaq's National Best Bid and Offer (NBBO) because they use finer tick sizes and because not all ECN quotes are displayed in the Nasdaq NBBO. In the Nasdaq quote montage, ECN quotes are rounded to the nearest Nasdaq tick, but ECN trade prices are not.⁵

⁴ Some market makers, such as Knight Securities, offer automated execution that is as fast as ECNs and the price is guaranteed, whereas the limit order on the ECN may disappear before the trade can be executed.

⁵ Effective June 3, 1997, Nasdaq moved to a tick of 1/16 for stocks with a bid price of \$10 or more, implying that ECN quotes are rounded to the nearest 1/16 of a dollar. Decimalization on Nasdaq will reduce rounding, but will not eliminate it because some ECNs already utilize a tick of less than one cent.

Thus, trades may occur on an ECN at a fraction of a tick below the NBBO. In addition, non-market makers (and ECNs with less than 5% of the volume in a stock) may choose to have their orders displayed only on the ECN and not included in the NBBO. Such limit orders can execute only against orders originating on the same ECN.

The potential benefits of trading on an ECN must be weighed against the costs. ECNs charge fees for their services that are paid directly by subscribers and indirectly by nonsubscribers.⁶ In addition, market makers sometimes execute orders at prices better than the NBBO. This practice, known as price improvement, is one dimension on which market makers compete with each other.⁷ In addition, market makers sometimes execute orders larger than the inside quoted depth, a practice referred to as size enhancement or size improvement. Finally, institutional trading arrangements on Nasdaq, such as preferencing and internalization, affect the ECNs' ability to attract orders and may allow market makers to retain order flow with more desirable characteristics.

Institutional investors and broker-dealers who subscribe to an ECN can place orders directly with the ECN. These traders may make dynamic choices to route an order to an ECN or to a market maker based on current market conditions. For example, best execution (see Macey and O'Hara (1997) for a more detailed discussion of best execution rules) does not require market makers to match the rounded or nondisplayed quotes on an ECN. Thus, customers may go directly to the ECN to get these better prices. If a subscriber sends a marketable order to an ECN and the ECN is not currently posting the best bid or offer, then the ECN may send the order to another market for execution. But brokers choosing to route orders to ECNs, not the ECNs themselves, are responsible for meeting the best execution requirements.

Individual investors can access an ECN through their broker. Some individual investors choose a particular broker expressly because that broker routes their orders to a specific ECN. Other individual investors may or may not be aware of the market to which their broker routes their orders. Market makers can also route customer orders to an ECN. Market orders can be routed to an ECN when the ECN is posting the best bid or offer, and the market maker does not want to match that price. Limit orders submitted to a market maker can also be routed to an ECN in accordance with the new order-handling rules.

⁶ For subscribers these fees include a fixed component, the cost of purchasing the ECN terminal and line feed, and a per-share fee for execution. For nonsubscribers an access fee of 0.25 to 2 cents per share is charged for orders that execute against a standing ECN order. During our sample period, this fee was paid by the intermediary routing the order to the ECN and was not charged directly to investors.

⁷ Because not all ECN orders are displayed in the NBBO, ECN trades sometimes occur at prices better than the NBBO. However, because ECNs simply match orders to buy with orders to sell, generally there is no opportunity for price improvement beyond the unrounded or undisplayed quote on an ECN.

II. Data and Descriptive Statistics

Our data contains all of the trades and quotes for 150 Nasdaq National Market stocks during the normal trading hours from 9:30 a.m. to 4:00 p.m. in June 2000. The 150 stocks in our sample include the 50 highest dollar-volume stocks (high-volume), stocks with dollar-volume ranks from 501 to 550 (medium-volume), and stocks with dollar-volume ranks from 1,001 to 1,050 (low-volume). The data include almost 19 million trades for \$825 billion and a similar number of quote revisions.

The bulk of our analysis uses data obtained from detailed trade, quote, and clearing databases maintained by Nasdaq. The data identify the unrounded trade price and volume, and include a number of unique features that allow us to accurately classify trades by venue. The trading venue is defined as ECN if the trade is executed by an ECN or routed to or from an ECN for execution. All other trades are classified as market-maker trades.⁸ The ECNs included in the analysis are Archipelago, Attain, Bloomberg TradeBook, Brass Utility LLC, Instinet, Island, Market XT, NexTrade, Redibook, and Strike Technologies LLC. A dummy variable, *ecn*, is set equal to one for trades that occurred on an ECN.

We also classify trades by size and construct dummy variables, *size1*, *size2*, and *size3*, that are set equal to one for trades that are small (1,000 shares or less), medium (1,001 to 9,999 shares), or large (10,000 shares or more), respectively. The dummy variables *ecn1*, *ecn2*, and *ecn3* are equal to the trade-size dummies interacted with the ECN dummy. To examine the importance of smaller tick sizes, *frac1*, *frac2*, and *frac3* are the trade-size dummies interacted with a dummy variable that is set equal to one if the trade price is not divisible by 1/16, and *ecnfrac1*, *ecnfrac2*, and *ecnfrac3* are these variables interacted with the ECN dummy.

Using the Lee and Ready (1991) algorithm, trades are classified as buyer initiated if the trade price is greater than the quote midpoint, and seller initiated if the trade price is less than the quote midpoint.⁹ Trades executed at the midpoint are classified with the tick rule: Midpoint trades on an up-tick are classified as buyer initiated, and midpoint trades on a downtick are classified as seller initiated. We represent the buy/sell indicator with the variable $x = +1$ for a buy order and $x = -1$ for a sell order.

Table I provides descriptive statistics for the stocks in our sample by volume category. Not surprisingly, the highest volume category has the highest average market capitalization, trade size, and price per share, as well as a lower average quoted spread and more market makers.¹⁰ Average trading volume in the

⁸ For confidentiality reasons, Nasdaq would not identify trades by individual ECN. Clearing data is used to determine whether the trade involved an ECN and to correct for any missing data in the trade database. Additional details on this process are available from the authors.

⁹ Trades are matched with quotes using execution times and the following algorithm that has been found to perform well for the Nasdaq market. For SelectNet, SOES, and ACES trades, we match the trade with the inside quote 1 second or more before the trade execution time. Because SelectNet, SOES, and ACES are electronic trading systems run by Nasdaq, the execution times are very reliable. For other trades, we match the trade with the inside quote 3 seconds or more before the trade report time.

¹⁰ On a daily basis, we identify market makers in each security from the individual market participant quote data. Any firm that has at least one valid quote during the day is classified as a market maker in that security.

Table I
Sample Descriptive Statistics

Descriptive statistics for 150 Nasdaq National Market Stocks in June 2000. Volume categories are determined by ranking all firms traded on the Nasdaq National Market by total dollar trading volume in June 2000. The categories are defined as high volume (ranks 1 to 50), medium volume (ranks 501 to 550), and low volume (ranks 1,001 to 1,050). Quoted depth is the average of the aggregate size of all inside bid quotes and all inside ask quotes in round lots. Descriptive statistics are calculated for each security and then averaged across securities.

Variable	Volume Categories			
	High	Medium	Low	All
Daily trading volume (\$million)	733.79	12.57	3.54	249.97
Market cap (\$billion)	61.08	1.34	0.53	20.98
Average trade size (\$000)	65.58	50.85	32.76	49.73
Share price	102.09	34.42	20.11	52.21
Daily return std. dev. (%)	6.55	6.21	6.28	6.35
Quoted half-spread (basis points)	7.01	31.37	46.32	28.23
Quoted depth (00)	23.85	8.24	6.36	15.37
Number of marker makers	53.32	20.74	15.14	29.73
Dollar ECN volume (%)	39%	27%	23%	30%

high-volume stocks is 60 times larger than trading volume in the medium-volume stocks, and more than 200 times larger than trading volume in the low-volume stocks. Stock-return volatility is also slightly higher for the high-volume stocks, although volatility has little variation across the volume categories. Finally, 30% of the trading volume in our sample occurs on an ECN. The percentage of trading on an ECN declines from 39% in the high-volume category to 23% in the low-volume category.¹¹

Table II shows the percentage of dollar volume that occurs on ECNs by volume category and trade size. ECN trades are smaller, on average, than market-maker trades. For the high-volume stocks, 51% of the small-trade volume occurs on ECNs. This percentage declines to 21% for medium-size trades, and to only 2% for large trades. The lower-volume stocks exhibit the same inverse relation between trade size and ECN volume. However these lower-volume stocks have fewer ECN trades in all trade-size categories. The percentage of small trades occurring on an ECN is 34% for the medium-volume stocks and 26% for the low-volume stocks. Relatively few medium and large trades occur on ECNs in the medium- and low-volume categories. The percentage of dollar volume occurring on ECNs is generally lower than the percentage of trades, particularly for the large trades where about 2% of the dollar volume occurs on ECNs in all dollar volume categories.

¹¹The positive association between normal trading volume and the fraction of volume on an ECN is evidence of a liquidity externality (Mendelson (1982) and others).

Table II
Percentage Dollar Volume on ECNs by Volume Category and Trade Size

Trades are classified as ECN if either the buyer or seller submitted its order through an ECN. Volume categories are determined by ranking all firms traded on the Nasdaq National Market by total dollar trading volume in June 2000. The categories are defined as high volume (ranks 1 to 50), medium volume (ranks 501 to 550), and low volume (ranks 1,001 to 1,050). Trade-size categories are based on number of shares traded. Percentage dollar volume on ECNs is calculated stock by stock and then averaged across stocks.

Trade-Size Categories	Volume Categories			
	High	Medium	Low	All
1–1,000	51.3%	33.7%	26.3%	37.1%
1,001–9,999	21.4%	11.9%	15.0%	16.1%
10,000 +	2.3%	2.5%	1.6%	2.1%
All	38.7%	26.9%	23.1%	29.6%

III. The Choice of Trading on ECNs or with Market Makers

This paper explores the competition for order flow across venues that have different market designs and that offer different advantages and services to their potential customers. As a first step, this section explores the market conditions under which trades are most likely to occur on an ECN. Estimating a probit regression for the choice of trading venue helps identify the market conditions that cause traders to send their trades to an ECN rather than to a Nasdaq market maker, or the conditions under which traders who predominantly trade on an ECN are most active.

We examine two types of variables describing the market condition at the time of the trade. The first type is related to the quotes and describes the supply of liquidity at the time of the trade. To proxy for the quotes, we include the inside quoted half-spread ($qsp = 10,000 * (ask - bid)/(ask + bid)$), the quoted depth at the inside ($depth = (depth\ at\ the\ ask + depth\ at\ the\ bid)/2$), and the quoted depth imbalance on the side of the trade ($depthimb = buy/sell\ indicator * (depth\ at\ the\ ask - depth\ at\ the\ bid)$). We are also interested in the relative supply of liquidity by ECNs and market makers so, in a second set of regressions, we replace the inside quoted half-spread with the market-maker inside quoted half-spread ($mmqsp = 10,000 * (mmask - mmbid)/(ask + bid)$, where $mmbid$ and $mmask$ are the best bid and ask offered by a market maker) and the difference between the market-maker and ECN inside half-spreads ($qspdif = 10,000 * ((mmask - mmbid) - (ecnask - ecnbid))/(ask + bid)$, where $ecnbid$ and $ecnask$ are the best bid and ask offered by an ECN).

The second group of variables is related to the level of market activity and describes the information environment at the time of the trade. The market activity variables that we include are the total dollar trading volume in the stock during the previous 15 minutes ($dvol15$), the stock-return volatility during the previous 15 minutes ($\sigma15 = the\ absolute\ value\ of\ the\ stock\ return\ over\ the\ previous\ 15\ minutes$), and the stock-price momentum ($mom15 = buy/sell\ indicator * stock\ return\ over\ the\ previous$

15 minutes). Finally, in addition to the variables listed above, we include the size of the trade ($svol = \text{the number of shares traded}$), firm fixed effects, and dummy variables for each half-hour time period during the trading day.¹²

A. Liquidity Variables

Table III reports the estimated coefficients from the probit regression with the corresponding linear probability slopes and chi-square statistics. The negative coefficient on the inside quoted half-spread in all three volume categories suggests that ECNs attract volume by offering better prices and narrowing the quotes. When quotes are narrow, it is more likely that an ECN quote is on the inside, making an ECN trade more likely. This intuition is confirmed when we replace the inside quoted half-spread with the market-maker inside quoted half-spread and the quoted spread difference. In this specification, we find that in all volume categories ECNs attract more orders when they are offering better prices than market makers, as indicated by a positive coefficient on the quoted spread difference (although the coefficients in the medium and low-volume categories are extremely small and round to zero at two decimal places). The effect of a change in the market-maker quoted spread, holding constant the ECN quoted spread, is captured by the sum of the coefficients on the market-maker quoted spread and quoted spread difference variables. As expected, the sum of these coefficients is positive in all volume categories, which indicates that when market makers increase their quoted spreads, trades are more likely to occur on ECNs.

The coefficient on the quoted depth variable is negative (and the chi-square statistic is large) for the high-volume stocks. Because ECNs are usually at the inside for high-volume stocks, greater quoted depth at the inside usually will be associated with more aggressive market-maker quotes, which will reduce the likelihood of an ECN trade. In the medium- and low-volume stocks, the coefficients on the quoted depth variables are positive. For these stocks, ECNs are less likely to be at the inside and there are fewer market makers per stock. In this situation, more quoted depth on the inside generally indicates more liquidity supplied by the ECNs, which increases the likelihood of an ECN trade. However, the chi-squared statistics for these categories are small (not statistically significant for the low-volume category).

For all volume categories, more depth offered on the side of the trade (depth imbalance) is associated with a lower probability that the trade will occur on an ECN. Previous studies have shown that depth imbalance is often associated with subsequent price changes. If this is true, then the negative coefficient for the depth imbalance coefficient suggests that ECN trades are more likely to occur in

¹²A number of other variables were also tested: the number of trades, dollar trading volume, stock-return volatility, and stock-price momentum during the previous 30 minutes, 60 minutes, and the previous trading day; the depth measured one, two, or three ticks from each stock's time-weighted quoted spread; the spreads necessary to generate several different levels of quoted depth; and the time-weighted spread and depth during the previous 5, 15, or 30 minutes. These variables add little to the explanatory power of the regression and do not affect the coefficients of interest.

Table III
Probit Regressions for the Choice of Trading Venue

The dependent variable is equal to one for ECN trades and zero for market-maker trades. Trades are classified as ECN if either the buyer or seller submitted its order through an ECN. All other trades are classified as market-maker trades. The volume categories are determined by ranking all Nasdaq National Market firms by total dollar trading volume in June 2000. The categories are defined as high volume (ranks 1 to 50), medium volume (ranks 501 to 550), and low volume (ranks 1,001 to 1,050). The number of shares traded is *svol*. The inside quoted half-spread is *qsp*. The market-maker inside quoted half-spread is *mmqsp*. The difference between the market-maker and ECN inside half-spreads is *qspdif*. The quoted depth at the inside (average of the sum of all bid and ask quotes) is *dep*. The quoted depth imbalance on the side of the trade (buy/sell indicator * (depth at the ask—depth at the bid)) is *depimb*. The total dollar trading volume in the stock during the previous 15 minutes is *dvol15*. The absolute value of the stock return over the previous 15 minutes is σ_{15} . The buy/sell indicator times the stock return over the previous 15 minutes is *mom15*. Firm fixed effects and dummy variables for each half-hour time period are not reported.

Variable	Volume Categories					
	High		Medium		Low	
<i>svol</i> * 1,000	-0.28	-0.28	-0.40	-0.40	-0.56	-0.56
Probability slope	-0.11	-0.11	-0.15	-0.15	-0.18	-0.18
Chi-square statistic	15,975	15,569	1,582	1,582	852	851
<i>qsp</i> * 100	-0.05		-0.09		-0.15	
Probability slope	-0.02		-0.03		-0.05	
Chi-square statistic	270		558		541	
<i>mmqsp</i> * 100		-0.29		0.05		0.07
Probability slope		-0.11		0.02		0.02
Chi-square statistic		13,504		839		894
<i>qspdif</i> * 100		0.49		0.00		0.00
Probability slope		0.19		0.00		0.00
Chi-square statistic		65,907		236		175
<i>dep</i> * 1,000	-5.53	-4.52	4.45	3.88	1.16	0.94
Probability slope	-2.19	-1.79	1.62	1.41	0.38	0.31
Chi-square statistic	8,469	5,598	64	49	3	2
<i>depimb</i> * 1,000	-4.79	-5.22	-5.73	-6.02	-3.29	-3.31
Probability slope	-1.89	-2.06	-2.09	-2.20	-1.08	-1.08
Chi-square statistic	30,022	35,166	490	540	116	117
<i>dvol15</i> * 100	9.04	8.77	2.58	3.43	2.26	3.80
Probability slope	3.58	3.47	0.94	1.25	0.74	1.25
Chi-square statistic	45,716	41,890	361	630	172	482
σ_{15} * 100	0.63	1.29	1.78	1.54	1.68	1.28
Probability slope	0.25	0.51	0.65	0.56	0.55	0.42
Chi-square statistic	396	1,652	626	462	325	186
<i>mom15</i> * 100	-0.62	-0.12	0.17	0.17	0.29	0.22
Probability slope	-0.02	-0.05	0.06	0.06	0.10	0.07
Chi-square statistic	9	34	12	11	20	12
Observations	17,692,539	17,692,539	752,575	752,575	372,689	372,689

the same direction as the subsequent price change. The question of whether ECN or market makers trades are more informed is explored in depth in the subsequent sections.

B. Market Activity Variables

The coefficients on the market-activity variables indicate that ECN trading is more likely when there is higher trading volume and when there is new information as measured by stock-return volatility. However, ECN trading is not significantly related to stock-price momentum, suggesting that short-term momentum traders are not concentrated in a single trading venue.

Consistent with the summary statistics in Table II, the negative coefficients on the size of the trade (*svol*) indicate that large trades are less likely to occur on ECNs. The coefficients on the time-period dummy variables are not reported, but contain several patterns. The likelihood of an ECN trade is almost 10% lower in the first half-hour of the trading day than during any other time period. This probably reflects the large number of orders that accumulate overnight and execute with market makers at the open. The probability of trading on an ECN increases slightly through the morning and early afternoon, and is highest in the last hour of the day.

As an alternative to the specification using firm fixed effects, we also estimate the probit regressions including firm characteristics (not reported). These regressions indicate that ECN trades are more prevalent in stocks that have high trading volume, large market capitalization, and fewer market makers. The coefficients on the remaining variables are largely unaffected by the inclusion of firm fixed effects or firm characteristics.

IV. Information and Price Discovery from ECN and Market-Maker Trades

The higher frequency of ECN trades during periods of high trading volume and high stock-return volatility suggests that ECNs attract a higher fraction of informed trades than market makers. To determine whether informed trades are more likely to occur on an ECN, and to measure the relative importance of ECN and market-maker trades in the price discovery process, we conduct several analyses. First, we decompose the stock-price variance into its trade-related and trade-unrelated components (Hasbrouck (1991a)), and show that ECN trading explains about two-thirds more of the stock-price variance than does market-maker trading. Then, to control for the contemporaneous correlation between ECN and market-maker trades, we estimate vector autoregressions (Hasbrouck (1991b)) and measure the permanent price impact of trades in different venues. When we conservatively assume that all contemporaneous correlation between ECN and market-maker trades is attributable to market-maker trading, we find that the permanent price impact of an ECN trade is 50% larger than the permanent price impact of a market-maker trade.

Generally, when a secondary market skims orders from a primary market, there is a concern that the secondary market will skim the most profitable, that is, least informed, orders. This cream skinning will increase the adverse selection in the primary market, resulting in higher trading costs and lower depth in the primary market. There is some evidence, for example, that the regional

exchanges skim less-informed order flow from the NYSE (Easley, Kiefer, and O'Hara (1996) and Bessembinder and Kaufman (1997)). On Nasdaq, the opposite seems to occur. Nasdaq market makers apparently have been able to retain the most profitable orders while the more informed and consequently less profitable orders are executed on the secondary (ECN) market.

This suggests that market makers, through preferencing and internalization agreements, are able to capture and retain the less informed orders. In addition, market makers are skilled at protecting themselves against large informed trades. Trading large orders with market makers is a repeated game that is not anonymous. Therefore, informed traders are unlikely to place their most profitable trades with a market maker because the market maker will use this information in future negotiations. Instead, informed traders will break up their large orders in an attempt to camouflage their trades with the uninformed order flow (Barclay and Warner (1993)). Because trading on an ECN is anonymous, informed traders can take out whatever depth is available there without concern about how this will affect subsequent negotiations with market makers.¹³

A. Price Discovery from Market Makers and ECNs: Variance Decomposition

To determine both the relative and absolute amounts of price discovery from trades with market makers and trades on ECNs, we use the variance decomposition technique from Hasbrouck (1991a). We decompose the variance of the efficient stock price into a component that is correlated with prior trades and a component that is uncorrelated with prior trades. The component of the stock-price variance that is correlated with prior trading is interpreted as the revelation of private information through trading. This decomposition is performed separately for market-maker trades and for ECN trades. Thus, in addition to controlling for serial correlation in the signed trades, the variance decomposition provides a measure of the total contribution to price discovery from ECN trades.

Following Hasbrouck, the time scale (t) is defined as the transaction sequence. We represent a trade at time t by the variable $x_t = +1$ for a buy order and $x_t = -1$ for a sell order. The percentage change (log return) in the quote midpoint subsequent to that trade (but prior to the next trade at $t + 1$) is denoted r_t . We then estimate the following VAR of trades and quote changes:¹⁴

$$r_t = \sum_{i=1}^p \alpha_i r_{t-i} + \sum_{i=0}^p \beta_i x_{t-i} + \varepsilon_{1,t} \quad (1)$$

¹³ As discussed subsequently, we find that the average price impact, measured as the signed difference between the bid-ask midpoint at the time of the trade and the bid-ask midpoint 5 minutes after the trade, is negative after large trades with market makers. A negative price impact implies that the realized spread is larger than the effective spread on these trades (see Table VI) and occurs when the stock price declines after investor purchases and increases after investor sales.

¹⁴ Identification also requires the following restrictions on the innovations (as in Hasbrouck (1991a, 1991b)): $E\varepsilon_{1,t} = E\varepsilon_{2,t} = 0$ and $E\varepsilon_{1,t}\varepsilon_{1,s} = E\varepsilon_{2,t}\varepsilon_{2,s} = E\varepsilon_{1,t}\varepsilon_{2,s} = 0$, for $s < t$.

and

$$x_t = \sum_{i=1}^p \gamma_i r_{t-i} + \sum_{i=1}^p \delta_i x_{t-i} + \varepsilon_{2,t}. \tag{2}$$

The trading process is assumed to restart at the beginning of each day, at which time all lagged values of x_t and r_t are set to zero. Because the average number of trades per unit time is sensitive to normal trading volume, we use $p = 100$ lags for high-volume stocks, 20 lags for medium-volume stocks, and 10 lags for low-volume stocks.¹⁵

Once estimated, the VAR representation can be inverted to generate the following vector moving average (VMA) model:

$$\begin{pmatrix} r_t \\ x_t \end{pmatrix} = \begin{pmatrix} a(L) & b(L) \\ c(L) & d(L) \end{pmatrix} \begin{pmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{pmatrix}, \tag{3}$$

where $a(L)$, $b(L)$, $c(L)$, and $d(L)$ are lag polynomial operators. The coefficients of the lag polynomials in this moving average representation are the impulse response functions implied by the VAR. Within the VAR framework, calculating the fraction of total price discovery that is due to private information revealed through trades is a straightforward variance decomposition. Following Hasbrouck, we decompose the (logarithm) of the spread midpoint, denoted p_t , into a random-walk component m_t and a stationary component s_t :

$$p_t = m_t + s_t, \tag{4}$$

where $m_t = m_{t-1} + v_t$ and $v_t \sim N(0, \sigma_v^2)$ with $E v_t v_s = 0$ for $t \neq s$. We refer to the random-walk component (m_t) as the permanent component of the price, and we refer to the stationary component (s_t) as the transitory component of the price. Defining $\sigma_{\varepsilon_1}^2 = E \varepsilon_{1,t}^2$ and $\sigma_{\varepsilon_2}^2 = E \varepsilon_{2,t}^2$, we can further decompose the variance of the permanent (or random walk) component of the price changes, σ_v^2 , into the price changes caused by the arrival of public information and price changes caused by the arrival of private information through trades:

$$\sigma_v^2 = \left(\sum_{i=0}^{\infty} a_i \right)^2 \sigma_{\varepsilon_1}^2 + \left(\sum_{i=0}^{\infty} b_i \right)^2 \sigma_{\varepsilon_2}^2, \tag{5}$$

where the second term in this equation, $\sigma_x^2 = \left(\sum_{i=0}^{\infty} b_i \right)^2 \sigma_{\varepsilon_2}^2$, represents the fraction of the total price discovery that is attributed to private information revealed through trades.

This variance decomposition is performed separately for market-maker and ECN trades by using only the market-maker or ECN trades to estimate the VAR. Because the variance of the efficient price (σ_v^2) is the same in both venues,

¹⁵We also estimate, but do not present, a model in which x_t is a vector containing signed trade, signed trade volume, and signed trade volume squared (as in Hasbrouck (1991a, 1991b)). Adding signed trade volume and signed trade volume squared provides little additional explanatory power, primarily because large trades on Nasdaq do not appear to contain more information than small trades (see the subsequent related discussion of the effective and realized spreads by trade size). We also estimate the system using varying numbers of lagged trades and quote changes. Our results are not sensitive to the choice of the number of lags.

Table IV
Decomposition of the Efficient Stock-Price Variance by Volume Category and Trading Venue

Volume categories are determined by ranking all firms traded on the Nasdaq National Market by total dollar trading volume in June 2000. The categories are defined as high volume (ranks 1 to 50), medium volume (ranks 501 to 550), and low volume (ranks 1,001 to 1,050). The fraction of the efficient stock-price variance explained by trades in a given trading venue (σ_x^2/σ_v^2) is based on the VAR/VMA model described in Section IV.A. For each volume category and venue, the cross-sectional mean and standard deviation are provided. For each firm in each venue, the standard deviation of 1,000 Monte Carlo simulations is calculated to test the statistical significance of the differences across venues.

	Volume Categories and Venues					
	High		Medium		Low	
	ECN	MM	ECN	MM	ECN	MM
Mean σ_x^2/σ_v^2	0.28	0.12	0.23	0.14	0.21	0.13
Std. Dev. (σ_x^2/σ_v^2)	(0.10)	(0.07)	(0.09)	(0.06)	(0.08)	(0.05)
Fraction of firms with						
ECN > MM	0.98		0.88		0.80	
ECN > MM at 0.10 level	0.96		0.40		0.38	
ECN > MM at 0.05 level	0.96		0.36		0.34	
ECN > MM at 0.01 level	0.96		0.32		0.22	

comparing the ratio of private information to total information (σ_x^2/σ_v^2) for ECN and market-maker trades indicates which trading venue contributes more to price discovery. This approach is similar to the information share employed by Hasbrouck (1995) for the NYSE and regional exchanges and by Huang (2002) for ECNs and several categories of market makers. However, the information share estimates in these papers are based solely on quotes and do not directly investigate where informed traders trade.

The results of the variance decompositions described above are reported in Table IV. For the high-volume stocks, ECN trades explain more than twice the amount of the efficient price variance than market-maker trades (28% for ECN trades and 12% for market-maker trades), even though less than half of all trades in this volume category are executed on ECNs. For the medium- and low-volume stocks, ECN trades explain about 60% more of the efficient price variance than market-maker trades, even though less than one-third of all trades in these volume categories are executed on ECNs. The difference in means between σ_x^2/σ_v^2 for ECN and market-maker trades is significant at the 0.01 level for all three volume categories with *t*-statistics of 16, 8, and 7 for the high-, medium-, and low-volume categories, respectively.¹⁶

¹⁶The variance decomposition was also performed in calendar time using time intervals. This approach yields very similar results for the high-volume stocks. In the medium- and low-volume categories, ECN trades continue to explain more of the efficient price variance than do market-maker trades, but the differences are smaller.

For each stock and trading venue, the standard deviation of σ_x^2/σ_v^2 is calculated using 1,000 Monte Carlo simulations with random draws from the estimated VAR coefficients. We use these standard deviations to determine the number of firms for which the amount of private information due to ECN trading is greater than the amount due to market-maker trading at the 0.10, 0.05, and 0.01 levels. For high-volume stocks, ECN trades explain more of the efficient price variance for 49 of the 50 stocks, and the difference is significant at the 0.01 level for 48 stocks. For medium- and low-volume stocks, ECN trades explain more of the efficient price variance for 44 and 40 of the 50 stocks, respectively, and the difference is significant at the 0.01 level for 16 and 11.

The variance decomposition technique controls for serial correlation in trades by including lagged trades in the VAR. However, this technique does not control for contemporaneous correlation between ECN and market-maker trades. In our data, the contemporaneous correlation between ECN and market-maker trades is high enough to warrant further exploration.¹⁷

B. The Impulse Response Function

Several previous papers examine the price discovery process for stocks trading in multiple venues. These papers typically examine quote changes in the different venues and attempt to determine which market “moves first.” These papers must also deal with the issue of contemporaneous correlation between trades in the different venues. Hasbrouck (1995) and Huang (2002) deal with this problem using Cholesky factorization to bound the information share. They alternately assume that any contemporaneous quote changes are caused entirely by quote changes in one venue or the other. We use this type of assumption to calculate upper and lower bounds for the price impact of ECN and market-maker trades. If we assume that ECN trades cause the contemporary market-maker trades and quote changes, then we can calculate the upper bound for the price impact of ECN trades and the lower bound for the price impact of market-maker trades. If we assume that market-maker trades cause the contemporary ECN trades and quote changes, then we can calculate the upper bound for the price impact of market-maker trades and the lower bound for the price impact of ECN trades. We report results aggregating trades over short (5-second) intervals to reduce the contemporaneous feedback from quote changes to trades. However, results for 15- and 30-second intervals are qualitatively similar.

Specifically, we use trades and quote changes aggregated over 5-second intervals to construct a VAR with a quote equation, an ECN trade equation, and a market-maker trade equation. Define x_t^e to be the sum of the signed ECN trades (+1 buys, -1 sells) during the 5-second interval t . Similarly, define x_t^m to be the sum of the signed market-maker trades during interval t , and define r_t to be the percentage change (log return) in the quote midpoint during interval t . The VAR

¹⁷The contemporaneous correlation between ECN and market-maker trades over 5-second intervals is 0.19 for our full sample and 0.28, 0.15, and 0.14 for the high-, medium-, and low-volume categories, respectively.

using 10 lags consists of three equations:

$$r_t = \sum_{i=1}^{10} \alpha_i r_{t-i} + \sum_{i=1}^{10} \beta_i x_{t-i}^e + \sum_{i=0}^{10} \zeta_i x_{t-i}^m + \varepsilon_{1,t}, \quad (6)$$

$$x_t^m = \sum_{i=1}^{10} \delta_i r_{t-i} + \sum_{i=1}^{10} \phi_i x_{t-i}^e + \sum_{i=1}^{10} \gamma_i x_{t-i}^m + \varepsilon_{2,t}, \quad (7)$$

and

$$x_t^e = \sum_{i=1}^{10} \varphi_i r_{t-i} + \sum_{i=1}^{10} \kappa_i x_{t-i}^e + \sum_{i=0}^{10} \lambda_i x_{t-i}^m + \varepsilon_{3,t}. \quad (8)$$

In this specification, the contemporaneous market-maker trading variable, x_t^m appears in the quote and ECN-trade equations. Thus, market-maker trades are assumed to cause contemporaneous quote changes and ECN trades. We can reverse this assumption by removing the contemporaneous market-maker trades and inserting the contemporaneous ECN trades. Because the clock is based on calendar time rather than transaction time, the VAR uses the same number of lags, 10, for all volume categories.

As in the previous section, the VAR is inverted to get the VMA representation. However, we now focus on the impulse response functions that measure the permanent price impacts from shocks to the trade equations. The estimated bounds for the permanent price impacts of ECN and market-maker trades are reported in Table V. As discussed in Hasbrouck (1991a, 1991b), this method is robust to price discreteness, lagged adjustment to information, and lagged adjustment to trades.

Table V shows that the permanent price impact of ECN trades is at least 50% larger than the permanent price impact of market-maker trades. For high-volume stocks, the lower bound for the ECN price impact is almost 50% larger than the upper bound for the market-maker price impact. The relative magnitudes of the ECN lower bound and the market-maker upper bound are similar in the medium- and low-volume categories. The difference between the lower bound of the ECN price impact and the upper bound of the market-maker price impact is significant at the .01 level for all three volume categories, with t -statistics of 8, 7, and 6, respectively.

As in the previous section, we calculate the standard deviation of the permanent price impacts from 1,000 Monte Carlo simulations with random draws from the estimated VAR coefficients. We use these standard deviations to determine the number of firms for which the lower bound for the ECN price impact is greater than the upper bound for the market-maker price impact at the 0.10, 0.05, and 0.01 levels. For high-volume stocks, the lower bound of the ECN price impact is greater than the upper bound for the market-maker price impact for 46 of the 50 stocks, and the difference is statistically significant at the 0.01 level for 40 stocks. For medium- and low-volume stocks, the lower bound of the ECN price impact is greater than the upper bound of the market-maker price impact for 42 of the 50 stocks, and the difference is statistically significant at the 0.01 level for 24 and 22 stocks, respectively.

Table V
Impulse Response Functions by Volume Category and Trading Venue

Volume categories are determined by ranking all firms traded on the Nasdaq National Market by total dollar trading volume in June 2000. The categories are defined as high volume (ranks 1 to 50), medium volume (ranks 501 to 550), and low volume (ranks 1,001 to 1,050). The impulse response of a trade innovation (in basis points) is based on the VAR/VMA model described in Section IV.B. For each volume category and venue, the cross-sectional mean and standard deviation are provided. For each firm in each venue the standard deviation of 1,000 Monte Carlo simulations is calculated to test the statistical significance of the differences across venues.

	Volume Categories and Venues					
	High		Medium		Low	
	ECN	MM	ECN	MM	ECN	MM
ECN upper/MM lower						
Mean impulse response	1.73	0.74	11.18	6.82	16.44	10.05
Std. Dev.	(0.64)	(0.45)	(3.70)	(2.80)	(6.77)	(4.72)
ECN lower/MM upper						
Mean impulse response	1.31	0.89	10.47	7.54	15.47	10.67
Std. Dev.	(0.51)	(0.52)	(3.59)	(3.02)	(6.55)	(5.02)
Fraction of firms with						
ECN lower > MM upper	0.92		0.84		0.84	
ECN lower > MM upper at 0.10 level	0.82		0.60		0.52	
ECN lower > MM upper at 0.05 level	0.80		0.56		0.50	
ECN lower > MM upper at 0.01 level	0.80		0.48		0.44	

V. Trading Costs: Market Makers versus ECNs

In the sections above, we show that the different types of services offered by ECNs and Nasdaq market makers attract different investor clienteles. In particular, a higher fraction of informed trades occur on ECNs and consequently a disproportionate share of the price discovery occurs there. In this section, we explore the relationship of this sorting of informed and uninformed traders to the trading costs in the two markets. In particular, we examine the nature of the equilibrium in which neither demanders nor suppliers of liquidity have an incentive to switch trading venues. The choice of whether to send orders to a market maker or to an ECN will depend on the total expected trading costs, including the implicit, explicit, and opportunity costs. Because we do not have order submission data or a complete breakdown of commissions, we focus on the implicit trading costs as indicated by various measures of the bid-ask spread, which do not capture the benefits of speed and anonymity offered by ECNs.

A. Effective Spreads

Table VI reports average percentage effective half-spreads (measured as the absolute difference between the transaction price and the bid-ask midpoint) by volume category, trade size, and trading venue. The effective half-spreads for high-volume Nasdaq stocks are quite small, averaging between 7 and 16 basis

Table VI
**Percentage Effective and Realized Half-Spreads by Volume Category,
 Trade Size, and Trading Venue**

Trades are classified as ECN if either the buyer or seller submitted its order through an ECN. All other trades are classified as market-maker trades. Volume categories are determined by ranking all Nasdaq National Market firms by total dollar trading volume in June 2000. The categories are defined as high volume (ranks 1 to 50), medium volume (ranks 501 to 550), and low volume (ranks 1,001 to 1,050). Trade-size categories are based on the number of shares traded. All values are expressed in basis points. Panel A reports the effective half-spread. Panel B reports the realized half-spread.

Trade-Size Categories	Volume Categories and Venues								
	High			Medium			Low		
	ECN	MM	All	ECN	MM	All	ECN	MM	All
Panel A: Effective Half-Spread									
≤1,000	7.92	7.22	7.56	31.56	33.41	32.85	55.26	49.14	50.68
1,001–9,999	8.01	12.07	10.57	31.69	36.37	35.43	54.67	53.95	54.10
10,000 +	9.62	15.59	15.35	29.55	48.57	48.03	48.71	58.72	58.39
Panel B: Realized Half-Spread									
≤1,000	0.57	3.04	1.85	2.12	12.25	9.14	9.16	16.57	14.70
1,001–9,999	0.91	15.21	9.89	6.00	33.09	27.67	12.01	38.46	33.01
10,000 +	4.47	19.65	19.04	3.82	55.68	54.24	12.76	59.34	57.84

points depending on trade size and trading venue. Consequently, there is not much room for large absolute differences in effective spreads between ECN and market-maker trades for these stocks. Nevertheless, small trades on ECNs have slightly larger effective spreads than small trades with market makers. For medium- and large-size trades in the high-volume category, however, the market-maker effective spreads are more than 50% larger than the ECN effective spreads.¹⁸

Spreads are wider in the medium- and low-volume categories than in the high-volume category, and the absolute differences between ECN and market-maker effective spreads generally are also larger. ECN trades have lower effective spreads than market-maker trades for all trade sizes in the medium-volume category. In the low-volume category, market-maker trades have lower effective spreads for small- and medium-size trades, and ECN trades have lower effective spreads for large trades. It is interesting to note that for the medium- and low-volume categories, effective spreads increase with trade size for market-maker trades, but decline with trade size for ECN trades.

¹⁸ Unlike ECNs, market makers typically add a three to six cents per share charge, in lieu of commissions, on institutional trades. This may account for some of the cost differences we find for large trades. This could also cause prices subsequent to market-maker trades to move in the opposite direction of the trade by this amount. These effects are small compared to the differences between large trades on ECNs and with market makers.

Table VI also shows that for each volume category, the relative cost advantage for ECN trades increases with trade size. Due to both higher inventory and adverse selection costs, most microstructure models predict that the cost of trading with an intermediary increases with trade size. Therefore, finding a natural counterparty for large trades on an ECN generates the greatest cost savings.

The probit regressions in Table III demonstrate that market conditions at the time of the trade have a significant impact on the likelihood that a trade is executed by a market maker or on an ECN. These market conditions are also likely to affect the effective spreads. Therefore, in Table VII, we control for the effect of these market conditions by regressing the percentage effective half-spread for each trade on trade-size dummies, an ECN dummy variable interacted with the trade-size dummies, and control variables describing the market conditions at the time of the trade.¹⁹ As in the probit regressions, the control variables include liquidity variables (quoted inside spread, quoted depth at the inside, and quoted depth imbalance) and market activity variables (dollar trading volume, stock-return volatility, and stock-price momentum during the 15 minutes prior to the trade).²⁰ To control for time period and stock characteristics, these regressions also include firm fixed effects and dummy variables for each half-hour during the trading day.

The effective-spread regressions are reported separately for each volume category. The variables of interest in these regressions are the trade-size dummies interacted with the ECN dummy (*ecn1*, *ecn2*, and *ecn3*). The coefficients on these variables indicate the average difference in percentage effective half-spreads (measured in basis points) between ECN and market-maker trades after controlling for trade size and other variables that have been shown to affect trading costs.

The results in Table VII generally are consistent with the univariate comparisons in Table VI. However, controlling for market conditions at the time of the trade generally reduces the measured cost of ECN trades in relation to market-maker trades. The fact that ECN trades are more likely to occur when stock-return volatility is high, and when there is less depth on the side of the trade, is consistent with the notion that ECN trades occur when trading is more difficult or more costly. Therefore, controlling for these market conditions improves the relative performance of the ECN trades.

As noted above, small trades in high-volume stocks are slightly more expensive on ECNs than when they are executed by market makers. The coefficient on the dummy variable *ecn1* is only 0.28 basis points, however, which is small by any

¹⁹ Regressions run separately by trade size or using the probit results from Section III in the Heckman (1979) two-stage procedure to control for selectivity provide similar results.

²⁰ Adding the quoted spread as a control variable in the effective-spread regressions measures the quality of execution in relation to the quoted spread, which is commonly called price improvement. If we exclude the quoted spread from these regressions, the coefficients on the ECN dummy variables in the high- and medium-volume categories are largely unaffected. However, for the low-volume category, excluding the quoted spread from the regression reduces the relative cost of ECN trades.

Table VII
**Regressions of the Percentage Effective Half-Spread on ECN Indicators,
 Trade-Size Dummy Variables, and Market Conditions at the Time
 of the Trade**

Volume categories are determined by ranking all firms traded on the Nasdaq National Market by total dollar trading volume in June 2000. The categories are defined as high volume (ranks 1 to 50), medium volume (ranks 501 to 550), and low volume (ranks 1,001 to 1,050). Trade-size dummy variables that take the value one for trade sizes 100 to 1,000 shares, 1,001 to 9,999 shares, and greater than 10,000 shares, respectively, are *size1*, *size2*, and *size3* (*size1* is omitted). The trade-size dummies interacted with an ECN dummy are *ecn1*, *ecn2*, and *ecn3*. The trade-size dummies interacted with a dummy variable that is set equal to one if the trade price is not divisible by 1/16 are *frac1*, *frac2*, and *frac3* and these variables interacted with an ECN dummy are *ecnfrac1*, *ecnfrac2*, and *ecnfrac3*. The inside quoted half-spread is *qsp*. The quoted depth at the inside (average of the sum of all bid and ask quotes) is *dep*. The quoted depth imbalance on the side of the trade (buy/sell indicator * (depth at the ask – depth at the bid)) is *depimb*. The total dollar trading volume in the stock during the previous 15 minutes is *dvol15*. The absolute value of the stock return over the previous 15 minutes is σ_{15} . The buy/sell indicator times the stock return over the previous 15 minutes is *mom15*. Firm fixed effects and dummy variables for each half-hour time period are not reported. *T*-statistics are in parentheses.

Variable	Volume Categories					
	High		Medium		Low	
<i>ecn1</i>	0.28 (48.08)	0.74 (119.28)	-0.66 (-8.59)	0.11 (1.31)	4.86 (28.08)	5.49 (29.90)
<i>ecn2</i>	-3.03 (-119.16)	-2.21 (-82.50)	-3.60 (-15.80)	-3.11 (-12.82)	2.94 (6.08)	2.93 (5.66)
<i>ecn3</i>	-7.71 (-57.19)	-6.63 (-46.84)	-21.39 (-14.39)	-18.35 (-11.63)	-10.37 (-2.57)	-9.50 (-2.26)
<i>ecnfrac1</i>		0.49 (19.20)		-4.10 (-16.72)		-4.80 (-9.07)
<i>ecnfrac2</i>		-3.18 (-23.76)		-3.08 (-4.41)		1.39 (0.97)
<i>ecnfrac3</i>		-27.07 (-50.31)		-37.37 (-8.14)		-12.03 (-0.81)
<i>frac1</i>		-3.50 (-152.94)		-7.33 (-46.05)		-6.73 (-23.59)
<i>frac2</i>		-1.81 (-14.62)		-8.62 (-21.18)		-9.48 (-12.31)
<i>frac3</i>		19.37 (57.65)		15.04 (9.88)		6.51 (1.89)
<i>size2</i>	3.64 (212.52)	3.57 (207.02)	3.35 (24.82)	3.62 (25.42)	1.93 (7.31)	2.27 (8.09)
<i>size3</i>	9.76 (244.78)	9.38 (234.38)	18.21 (48.76)	16.72 (43.53)	13.40 (14.00)	12.16 (12.19)
<i>qsp</i>	-0.15 (-570.52)	-0.15 (-576.16)	0.26 (315.90)	0.26 (315.14)	0.77 (412.89)	0.77 (413.09)
<i>dep</i>	0.01 (329.15)	0.01 (329.47)	0.05 (46.87)	0.06 (57.60)	0.02 (13.66)	0.03 (16.08)
<i>depimb</i>	0.00 (-223.36)	0.00 (-216.58)	-0.02 (-30.15)	-0.02 (-39.55)	-0.02 (-25.62)	-0.03 (-28.45)
<i>dvol15</i>	-2.06 (-717.36)	-2.05 (-716.68)	-1.03 (-34.87)	-1.07 (-36.27)	0.01 (0.23)	-0.03 (-0.50)
σ_{15}	2.05 (802.73)	2.05 (806.26)	1.90 (124.37)	2.00 (130.93)	2.06 (71.25)	2.14 (73.87)
<i>mom15</i>	-0.19 (-96.97)	-0.19 (-98.75)	0.33 (29.20)	0.30 (26.96)	0.46 (21.29)	0.43 (20.20)
Observations	17,692,539	17,692,539	752,575	752,575	372,689	372,683

metric. Small ECN trades have lower effective spreads for medium-volume stocks and higher effective spreads for low-volume stocks, although these differences are also small.

For medium and large trades, the lower average effective spreads for ECN trades in the medium- and high-volume categories are economically significant. The three basis point reduction in effective spreads for medium-size ECN trades in the high-volume stocks represents a discount of 25% from the market-maker effective spreads. Large ECN trades have significantly lower effective spreads than market-maker trades for all volume categories with the percentage cost difference ranging from 50% for high-volume stocks to 20% for low-volume stocks. This is consistent with Conrad, Johnson, and Wahal (2002) who examine order-level data for institutions and find that ECNs are cheaper than traditional mechanisms for executing large orders.

Although medium and large trades get better prices when they are executed on an ECN, it is clear that these better prices on ECNs are often unavailable. Market makers generally offer more depth than ECNs. For high-volume stocks, for example, Table II shows that ECNs capture 51% of the volume in small trades, but only 21% of the volume in medium trades, and only 2% of the volume from large trades. When investors want medium or large trades, they generally go to a market maker. Medium and large trades go to an ECN only when there is a natural counterparty supplying liquidity on the ECN.

During our sample period, ECNs used price increments that were smaller than the minimum tick size of 1/16 on Nasdaq. In our sample, almost 19% of the small ECNs trades occurred on noninteger ticks, while only 3% of small market-maker trades occurred on noninteger ticks. The percentage of medium- and large-size trades on noninteger ticks is more similar across the trading venues. ECN executed 12% of medium- and large-size trades on noninteger ticks while market makers executed 9% of medium-size trades and 12% of large trades on noninteger ticks. To study the effect of noninteger ticks, we create three dummy variables, *frac1*, *frac2*, *frac3*, that are set equal to one for small, medium, and large trades on noninteger ticks, and three more dummy variables, *ecfrac1*, *ecfrac2*, and *ecfrac3*, that are set equal to one for small, medium, and large trades on noninteger ticks that occur on ECNs.

The ECN trades on noninteger ticks have significantly lower effective spreads than ECN trades on integer ticks (to calculate the impact of an ECN trade on a fractional tick, one adds the appropriate *frac* and *ecfrac* coefficients). For high-volume stocks, the effective spread for small, medium, and large trades on noninteger ticks is approximately 3, 5, and 8 basis points lower than ECN trades on integer ticks. Although small in absolute value, this represents a savings of 30% to 40% on the spread-related cost of the trade. For the medium- and low-volume stocks, the average effective spread for ECN trades on noninteger ticks is between 6 and 22 basis points lower than the effective spread for ECN trades on integer ticks. The lower effective spreads for ECN trades on noninteger ticks suggest that ECNs may be attracting orders by shaving the bid-ask spread by a fraction. In the next section, we investigate whether the trades attracted by quotes on fractional ticks are more or less informed than average. Small- and

medium-size trades on noninteger ticks that are executed by market makers also have lower effective spreads, which suggests that market makers may be providing better execution, including price improvement, to attract particular types of trades.

B. Realized Spreads

Realized spreads—the signed transaction price minus the quote midpoint five minutes after the trade—provide an ex post measure of trading cost to the demander of liquidity and an ex post measure of profitability for the provider of liquidity. The difference between an ex ante measure of cost (such as the effective spread) and an ex post measure of cost (the realized spread) is the price impact of the trade. If trades contain information, then prices will move in the direction of the trade (up following purchases and down following sales), reducing the ex post trading costs. If ECN trades are more informed than market-maker trades, then ECN trades should have lower realized spreads than market-maker trades.

Table VI shows that realized spreads are lower for ECN trades than for market-maker trades in all volume categories and for all trade sizes. The difference between the ECN and market-maker realized spreads is larger for larger trades and for lower-volume stocks. To control for trade size and other variables that have been shown to affect trading costs, Table VIII reports regressions of the percentage realized half-spread for each trade on trade-size dummies, an ECN dummy variable interacted with the trade-size dummies, and control variables describing the market conditions at the time of the trade. As in the previous regressions, the control variables include liquidity variables (quoted inside spread, quoted depth at the inside, and quoted depth imbalance) and market activity variables (dollar trading volume, stock-return volatility, and stock-price momentum during the 15 minutes prior to the trade). To control for time period and stock characteristics, these regressions also include firm fixed effects and dummy variables for each half-hour during the trading day. Regressions are reported separately for each volume category. The variables of interest in these regressions are the ECN dummy interacted with the trade-size dummies (*ecn1*, *ecn2*, and *ecn3*). The coefficients on these variables indicate the average difference in percentage realized half-spreads between ECN and market-maker trades.

Consistent with the results in Table VI, after controlling for the market conditions at the time of the trade, realized spreads are significantly lower for ECN trades than for market-maker trades in all volume categories and for all trade sizes. In addition, the coefficients on the ECN dummy variables in Table VIII are two to five times larger than the corresponding coefficients in the effective spread regressions in Table VII.

Table VII shows that the use of smaller minimum price increments generally results in lower ex ante trading costs. Above, we argued that the liquidity traders would be more patient and more price sensitive than informed traders. Thus, by shaving the quote by a fractional tick, it may be possible to attract market (or marketable-limit) orders from uninformed traders. A comparison of Tables VII

Table VIII
Regressions of the Percentage Realized Half-Spread on ECN Indicators, Trade-Size Dummy Variables, and Market Conditions at the Time of the Trade

Volume categories are determined by ranking all firms traded on the Nasdaq National Market by total dollar trading volume in June 2000. The categories are defined as high volume (ranks 1 to 50), medium volume (ranks 501 to 550), and low volume (ranks 1,001 to 1,050). Trade-size dummy variables that take the value one for trade sizes 100 to 1,000 shares, 1,001 to 9,999 shares, and greater than 10,000 shares, respectively, are *size1*, *size2*, and *size3* (*size1* is omitted). The trade-size dummies interacted with an ECN dummy are *ecn1*, *ecn2*, and *ecn3*. The trade-size dummies interacted with a dummy variable that is set equal to one if the trade price is not divisible by 1/16 are *frac1*, *frac2*, and *frac3* and these variables interacted with an ECN dummy are *ecnfrac1*, *ecnfrac2*, and *ecnfrac3*. The inside quoted half-spread is *qsp*. The quoted depth at the inside (average of the sum of all bid and ask quotes) is *dep*. The quoted depth imbalance on the side of the trade (buy/sell indicator * (depth at the ask—depth at the bid)) is *depimb*. The total dollar trading volume in the stock during the previous 15 minutes is *dvol15*. The absolute value of the stock return over the previous 15 minutes is σ_{15} . The buy/sell indicator times the stock return over the previous 15 minutes is *mom15*. Firm fixed effects and dummy variables for each half-hour time period are not reported. *T*-statistics are in parentheses.

Variable	Volume Categories					
	High		Medium		Low	
<i>ecn1</i>	-1.78 (-48.11)	-1.86 (-47.09)	-5.91 (-15.28)	-5.44 (-13.28)	-7.88 (-10.23)	-9.19 (-11.21)
<i>ecn2</i>	-9.21 (-57.38)	-9.10 (-53.51)	-19.08 (-16.60)	-19.43 (-15.88)	-14.77 (-6.85)	-16.49 (-7.13)
<i>ecn3</i>	-18.45 (-21.66)	-18.79 (-20.96)	-62.01 (-8.29)	-64.83 (-8.13)	-67.64 (-3.76)	-65.17 (-3.48)
<i>ecnfrac1</i>		-5.28 (-32.85)		-10.14 (-8.18)		9.83 (4.16)
<i>ecnfrac2</i>		-4.32 (-5.10)		2.26 (0.64)		10.29 (1.61)
<i>ecnfrac3</i>		-10.45 (-3.07)		25.05 (1.08)		-27.09 (-0.41)
<i>frac1</i>		6.73 (46.46)		21.60 (26.85)		13.47 (10.60)
<i>frac2</i>		4.06 (5.18)		6.77 (3.29)		16.89 (4.92)
<i>frac3</i>		15.85 (7.49)		17.60 (2.30)		25.26 (1.66)
<i>size2</i>	7.62 (70.55)	7.75 (71.07)	13.13 (19.33)	14.37 (19.99)	10.68 (9.07)	10.24 (8.19)
<i>size3</i>	17.16 (68.49)	17.14 (67.92)	45.24 (24.15)	45.84 (23.70)	48.43 (11.40)	47.69 (10.75)
<i>qsp</i>	0.05 (30.89)	0.05 (31.60)	0.13 (32.37)	0.14 (32.97)	0.37 (44.20)	0.37 (44.37)
<i>dep</i>	-0.01 (-53.72)	-0.01 (-54.31)	0.04 (7.69)	0.02 (3.02)	0.01 (1.27)	0.00 (0.16)
<i>depimb</i>	0.03 (254.78)	0.03 (253.70)	0.10 (35.74)	0.11 (39.39)	0.09 (21.87)	0.10 (23.13)
<i>dvol15</i>	-0.39 (-21.65)	-0.37 (-20.26)	-1.21 (-8.13)	-1.11 (-7.46)	-1.31 (-5.54)	-1.22 (-5.18)
σ_{15}	-0.20 (-12.23)	-0.20 (-12.44)	0.23 (3.03)	0.05 (0.69)	0.78 (6.03)	0.62 (4.78)
<i>mom15</i>	-0.82 (-67.08)	-0.81 (-66.42)	-2.85 (-49.97)	-2.80 (-48.96)	0.88 (9.19)	0.93 (9.69)
Observations	17,695,956	17,695,956	752,784	752,784	372,726	372,726

and VIII shows that trades on noninteger ticks generally have larger realized spreads than effective spreads, which confirms the hypothesis that these trades are less informed.

C. Size- and Time-Matched Sample

Tables VII and VIII document that after controlling for market conditions, small trades on ECNs have larger effective spreads and smaller realized spreads than do small trades with market makers, while medium and large trades on ECNs have both lower effective and realized spreads than do medium and large trades with market makers. Although we attempted to control for market conditions at the time of the trade, it is possible that we omitted some important variables from these regressions. Thus, to test the robustness of these results, we calculated effective and realized spreads for a matched sample of ECN and market-maker trades. For each stock, we collected all pairs of ECN and market-maker trades in the same trade-size category that occurred within 15 seconds of each other.²¹ Matching on trade-size category and time of day (within 15 seconds) may better control for the contemporaneous market liquidity and information environment. The effective and realized spreads of these matched pairs are reported in Table IX.

The results of the matched pairs are consistent with the regression results. Small trades on ECNs have larger effective spreads but smaller realized spreads than do small trades with market makers. Medium and large trades on ECNs have both lower effective and realized spreads than do trades with market makers. In addition, ECN trades of all sizes generally are more informed, as measured by the difference between the effective and realized spreads, than are trades with market makers. All of these results continue to hold when the sample is split into two equal time periods.

VI. Conclusion

This paper provides insights about the competition for order flow across trading venues that offer different services to their customers and provides the first systematic study of trading on the most prolific alternative trading systems, ECNs. To examine why some traders use ECNs, we study how the differences in services between ECNs and Nasdaq market makers affect the choice of trading venue and the cost of ECN and market-maker trades. We also investigate whether informed trades are more likely to occur with market makers or on ECNs, and study the importance of ECN trades for the aggregate price discovery process.

We show that ECNs attract more informed orders and that these trades occur during periods of high volume and stock-return volatility. ECNs' smaller tick size

²¹ Matching on finer trade-size categories (100, 101–499, 500, 501–999, 1,000, 1,001–5,000, 5,001–9,999, and 10,000–25,000 shares) and a narrower time interval (5 seconds) does not affect these results.

Table IX
Percentage Effective and Realized Half-Spreads by Volume Category and Trade Size for a Matched Sample

Trades are classified as ECN if either the buyer or seller submitted its order through an ECN. All other trades are classified as market-maker trades. Volume categories are determined by ranking all Nasdaq National Market firms by total dollar trading volume in June 2000. The categories are defined as high volume (ranks 1 to 50), medium volume (ranks 501 to 550), and low volume (ranks 1,001 to 1,050). Trade-size categories are based on the number of shares traded. Trades in the same size categories that occur on different venues that occur within 15 seconds of each other are included in the sample. All values are expressed in basis points. ECN values that differ from market maker value at a 0.01 level are denoted with a †. Panel A reports the effective half-spread. Panel B reports the realized half-spread.

Trade-Size Categories	Volume Categories and Venues					
	High		Medium		Low	
	ECN	MM	ECN	MM	ECN	MM
Panel A: Effective Half-Spread						
≤1,000	8.11	7.19†	34.02	33.07†	55.81	51.27†
1,001–9,999	8.25	10.28†	35.79	36.58†	62.88	59.72
10,000 +	9.52	14.10†	32.41	58.41†	41.75	47.95
Panel B: Realized Half-Spread						
≤1,000	0.12	2.62†	– 2.51	3.33†	2.17	4.19†
1,001–9,999	– 0.65	8.38†	– 2.89	13.33†	1.13	4.06†
10,000 +	– 2.93	13.08†	11.39	32.70†	– 86.17	– 26.63

also attracts less informed orders seeking (and receiving) lower ex ante execution costs. The ECNs’ customer-to-customer interactions generate significantly lower trading costs for large trades when there is a natural counterparty offering significant liquidity on an ECN. Because this much depth is seldom available on an ECN, however, traders seeking to trade large quantities typically require the services of an intermediary.

There is widespread concern about the detrimental effects of market fragmentation resulting from the proliferation of alternative trading venues (Mendelson (1987), Grossman (1992), Madhavan (1995), Hendershott and Mendelson (2000), and others). A major concern about multimarket trading is that a secondary market will skim the most profitable (i.e., least informed) orders. This cream skimming would increase the adverse selection in the primary market and result in higher trading costs and lower depth in the primary market. There is some evidence, for example, that the regional exchanges skim less-informed order flow from the NYSE (Easley et al. (1996) and Bessembinder and Kaufman (1997)). On Nasdaq, the opposite seems to occur. Nasdaq market makers apparently have been able to retain the most profitable orders while the more informed and consequently less profitable orders are executed on the secondary (ECN) market. Decomposing the stock-price variance into its trade-related and trade-unrelated

components, we show that ECN trading explains from 60% to 100% more of the efficient stock-price variance than market-maker trading. We find that the lower bound for the permanent price impact of ECN trades is about 50% larger than the upper bound for the permanent price impact of market-maker trades. These results show that ECNs have the unusual feature of attracting the more informed orders, something not found in any theoretical model of multimarket trading. In addition, these results show that ECN trading is an important component of the overall price discovery process.

Several widely used practices of Nasdaq trading, such as preferencing, internalization, and payment for order flow have been widely criticized in the academic literature. The fact that a majority of the aggregate price discovery occurs on ECNs provides a potential explanation for the high quality and efficiency of the Nasdaq stock prices despite these practices. Because ECNs facilitate trades that are more difficult in some dimensions, for example, more informed trades in active and volatile markets, while market makers facilitate trades that are more difficult in other dimensions, for example, larger trades, ECNs provide an important complement, but not a complete substitute, for the traditional trading mechanism of market makers. Thus, our results contribute to an integrated picture of quoting, trading, and price discovery in the hybrid Nasdaq market place.

REFERENCES

- Barclay, Michael, William Christie, Jeffrey Harris, Eugene Kandel, and Paul Schultz, 1999, The effects of market reform on the trading costs and depths of Nasdaq stocks, *Journal of Finance* 54, 1–34.
- Barclay, Michael, and Terrence Hendershott, 2003, Price discovery and trading after hours, *Review of Financial Studies*, forthcoming.
- Barclay, Michael, and Jerold Warner, 1993, Stealth trading and volatility: Which trades move prices? *Journal of Financial Economics* 34, 281–305.
- Bessembinder, Hendrik, and Herbert Kaufman, 1997, A cross-exchange comparison of execution costs and information flow for NYSE-listed stocks, *Journal of Financial Economics* 46, 293–319.
- Biais, Bruno, Christophe Bisiere, and Chester Spatt, 2002, Imperfect competition in financial markets: Island vs. Nasdaq, Working paper, Carnegie Mellon University.
- Conrad, Jennifer, Kevin Johnson, and Sunil Wahal, 2002, Institutional trading and alternative trading systems, *Journal of Financial Economics*, forthcoming.
- Coppejans, Mark, Ian Domowitz, and Ananth Madhavan, 2001, Liquidity in an automated auction, Working paper, Penn State University.
- Domowitz, Ian, and Jianxin Wang, 1994, Auctions as algorithms, *Journal of Economic Dynamics and Control* 18, 29–60.
- Easley, David, Nicholas Kiefer, and Maureen O'Hara, 1996, Cream-skimming or profit-sharing? The curious role of purchased order flow, *Journal of Finance* 51, 811–833.
- Glosten, Lawrence, 1994, Is the electronic open limit order book inevitable?, *Journal of Finance* 49, 1127–1161.
- Grossman, Sanford, 1992, The information role of upstairs and downstairs markets, *Journal of Business* 65, 509–529.
- Hasbrouck, Joel, 1991a, The summary informativeness of stock trades: An econometric analysis, *Review of Financial Studies* 4, 571–595.

- Hasbrouck, Joel, 1991b, Measuring the information content of stock trades, *Journal of Finance* 46, 179–207.
- Hasbrouck, Joel, 1995, One security, many markets: Determining the contributions to price discovery, *Journal of Finance* 50, 1175–1199.
- Hasbrouck, Joel, and Gideon Saar, 2001, Limit orders and volatility in a hybrid market: The Island ECN, Working paper, New York University.
- Heckman, James, 1979, Sample selection bias as a specification error, *Econometrica* 47, 153–162.
- Hendershott, Terrence, and Haim Mendelson, 2000, Crossing networks and dealer markets: Competition and performance, *Journal of Finance* 55, 2071–2115.
- Hollifield, Burton, Robert Miller, and Patrik Sandås, 2001, Empirical analysis of limit order markets, Working paper, University of Pennsylvania.
- Huang, Roger, 2002, The quality of ECN and Nasdaq market-maker quotes, *Journal of Finance* 57, 1285–1319.
- Lee, Charles, and Mark Ready, 1991, Inferring trade direction from intraday data, *Journal of Finance* 46, 733–747.
- Macey, Jonathan, and Maureen O'Hara, 1997, The law and economics of best execution, *Journal of Financial Intermediation* 6, 188–223.
- Madhavan, Ananth, 1995, Consolidation, fragmentation, and the disclosure of trading information, *Review of Financial Studies* 8, 579–603.
- Mendelson, Haim, 1982, Market behavior in a clearing house, *Econometrica* 50, 1505–1524.
- Mendelson, Haim, 1987, Consolidation, fragmentation, and market performance, *Journal of Financial and Quantitative Analysis* 22, 187–207.
- Parlour, Christine, and Duane Seppi, 2003, Liquidity-based competition for order flow, *Review of Financial Studies* 16, 301–343.
- Simaan, Yusif, Daniel Weaver, and David Whitcomb, 2003, Market-maker quotation behavior and pre-trade transparency, *Journal of Finance*, 58, 1247–1267.
- U.S. Securities and Exchange Commission, 2000, Electronic communications networks and after-hours trading.
- Weston, James, 2001, Competition on the Nasdaq and the growth of electronic communication networks, Working paper, Rice University.