



American Finance Association

Lifting the Veil: An Analysis of Pre-Trade Transparency at the NYSE

Author(s): Ekkehart Boehmer, Gideon Saar, Lei Yu

Source: *The Journal of Finance*, Vol. 60, No. 2 (Apr., 2005), pp. 783-815

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/3694767>

Accessed: 23/11/2009 12:31

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and American Finance Association are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance*.

<http://www.jstor.org>

Lifting the Veil: An Analysis of Pre-trade Transparency at the NYSE

EKKEHART BOEHMER, GIDEON SAAR, and LEI YU*

ABSTRACT

We study pre-trade transparency by looking at the introduction of NYSE's OpenBook service that provides limit-order book information to traders off the exchange floor. We find that traders attempt to manage limit-order exposure: They submit smaller orders and cancel orders faster. Specialists' participation rate and the depth they add to the quote decline. Liquidity increases in that the price impact of orders declines, and we find some improvement in the informational efficiency of prices. These results suggest that an increase in pre-trade transparency affects investors' trading strategies and can improve certain dimensions of market quality.

THE PROLIFERATION OF NEW EXCHANGES and trading platforms in the United States and abroad brings to the forefront many issues in market design. Should a market have at its core an electronic limit-order book? What possible roles can market makers play? What information should market participants observe about order flow and prices? These issues have implications for investor trading strategies, specialist behavior, market liquidity, the informational efficiency of prices, and ultimately for investor welfare. We investigate a key feature of market design: transparency, or the ability of market participants to observe information in the trading process.

Our focus is on a particular form of transparency: the ability of market participants to observe the pending trading interests of other participants, or in other words, the content of the limit-order book. Knowledge about buying and selling interest can be used both to refine one's inference about the value of a

*Ekkehart Boehmer is from the Mays Business School, Texas A&M University. Gideon Saar is from the Stern School of Business, New York University. Lei Yu is from the Mendoza College of Business, University of Notre Dame. We wish to thank Robert Stambaugh (the editor) and an anonymous referee, as well as Yakov Amihud, Cecilia Caglio, Robert Engle, Luca Filippa, Thierry Foucault, Larry Harris, Joel Hasbrouck, Craig Holden, Robert Jennings, Charles Jones, Ronald Jordan, Pete Kyle, Timothy McCormick, Barbara Rindi, Patrik Sandas, James Shapiro, Chester Spatt, Daniel Weaver, and seminar (or conference) participants at Bocconi University, Iowa State University, the New York Stock Exchange, the Securities and Exchange Commission, Southern Methodist University, SUNY at Buffalo, Texas A&M University, Tilburg University, University of Georgia, University of Kentucky, the Fed Atlanta Financial Markets Conference, the NBER Market Microstructure Group meetings, and the Western Finance Association meetings for helpful comments. This research began while Boehmer was a Director of Research at the New York Stock Exchange and Saar was on leave from NYU and held the position of a Visiting Research Economist at the New York Stock Exchange. The opinions expressed in this paper do not necessarily reflect those of the members or directors of the NYSE.

security and to strategically plan the execution of a trading goal to minimize transaction costs. We use the introduction of OpenBook by the NYSE to investigate the impact of an increase in the extent of public information about the content of the limit-order book.

OpenBook, introduced in January 2002, allows traders off the NYSE floor to observe depth in the book in real time at each price level for all securities. Before the introduction of OpenBook, only the best bid and offer (representing orders in the book, floor broker interest, and the specialist's own trading desires) had been disseminated. Our objective in this study is to examine how publicly revealing information about the limit-order book affects investor trading strategies, the way prices evolve in response to order flow, and the resulting state of liquidity in the market.

Our paper is related to other investigations of pre-trade transparency, defined as the availability of information about quotes and trading interest.¹ Most papers look at the influence of quote information in a multiple-dealer market (e.g., Bloomfield and O'Hara (1999) and Flood et al. (1999)) or use the availability of information to characterize different market structures (e.g., Madhavan (1992), Biais (1993), and Pagano and Röell (1996)).² Our work is also part of a growing theoretical and empirical literature about limit-order books.³ Two papers, Madhavan, Porter, and Weaver (2000) and Baruch (2005), specifically construct models to address the question of how revealing more or less information about the content of a limit-order book affects the market. We rely on theoretical predictions from these two papers to guide our investigation into the impact of limit-order book transparency on informational efficiency and liquidity.

In organizing our empirical investigation, we have found it useful to think about the consequences of changes in pre-trade transparency first in terms of direct effects on the trading strategies of market participants, and then in terms of the resulting equilibrium state of informational efficiency and liquidity. Harris (1996) discusses two risks that are associated with the exposure of limit orders: (i) A trader may reveal to the market private information about the value of the security, and (ii) exposed limit orders can be used to construct trading strategies aimed explicitly at taking advantage of these limit orders (e.g.,

¹ For investigations into post-trade transparency, defined as information about executed trades, see Chowdhry and Nanda (1991), Naik, Neuberger, and Viswanathan (1994), Franks and Schaefer (1995), Madhavan (1995), Lyons (1996), and Bloomfield and O'Hara (1999).

² A related literature focuses on the anonymity of traders as a dimension of transparency. For example, Madhavan (1996) investigates the role of information about order flow, but focuses on the availability of information about traders' motives (i.e., whether liquidity traders can be identified). The impact of anonymity on liquidity in an electronic limit-order book is investigated by Foucault, Moinas, and Theissen (2003), while Theissen (2003) examines anonymity in a setting where there is a market maker alongside a limit-order book. Rindi (2002) looks at anonymity in a rational expectations model.

³ The tradeoffs in using limit orders and the nature of equilibrium in limit-order markets were the focus of several theoretical models (e.g., Cohen et al. (1981), Glosten (1994), Seppi (1997), Parlour (1998), and Foucault (1999)). Recent empirical work on limit-order markets includes Biais, Hillion, and Spatt (1995), Handa and Schwartz (1996), Ahn, Bae, and Chan (2001), Sandås (2001), and Hasbrouck and Saar (2004).

“pennying” or front-running the limit orders). Harris details strategies that limit-order traders can use to manage the exposure of their orders: breaking up their orders (submitting smaller limit orders), canceling and resubmitting limit orders more often, and using agents closer to the trading process (floor brokers at the NYSE) to manage order exposure rather than submitting limit orders to the book.

We look at the cancellation of limit orders after OpenBook is introduced and find a higher cancellation rate and shorter time-to-cancellation of limit orders in the book. We also find smaller limit orders after the change in transparency. This evidence is consistent with the idea that traders attempt to manage the exposure of their orders, in line with Harris’s reasoning. However, we do not observe a shift from electronic limit orders to trading with the help of floor brokers. Instead, the volume executed by floor brokers declines compared to that executed against limit orders. Why may that be happening? OpenBook enables traders not just to observe information about demand and supply away from the quote, but also to see how their own actions affect the book. This “visibility effect” may make self-management of orders more appealing to traders, in a manner analogous to the attraction of active traders in Nasdaq stocks to electronic communications networks. Such an effect could dominate the argument that Harris (1996) makes for employing agents, and explain our finding.

We also investigate the trading of one particular type of market professionals—NYSE specialists—who both maintain the limit-order book and trade for their own account (make a market in the stocks). We find that the specialist participation rate in trading declines following the introduction of OpenBook. We also find that specialists reduce the depth they add to the quote (together with floor brokers) beyond what is in the limit-order book. These changes in trading strategies are consistent with an increase in the risk of proprietary trading on the part of specialists due to loss of their information advantage.⁴ Because public limit orders have priority over the specialists’ proprietary trading, these changes are also consistent with a crowding out effect due to more active limit-order strategies employed by investors. Finally, a reduced contribution of the floor to the quoted depth may be due to the shift we observe from floor trades to limit orders that are sent to the book electronically.

Changing strategies of market participants can alter characteristics of the market environment that are important to investors, such as liquidity and informational efficiency. Glosten (1999) presents an informal argument stating that increased transparency should lead to greater commonality of information, implying more efficient prices and narrower spreads. Baruch (2005) asks who benefits from an open limit-order book, and provides a theoretical model showing that opening the book (i) improves liquidity in the sense that the price impact of market orders is smaller, and (ii) improves the informational efficiency of prices. A different view is expressed by Madhavan, Porter, and Weaver (2000).

⁴ On whether the limit-order book provides information about future price movements, see Harris and Panchapagesan (1999), Corwin and Lipson (2000), Irvine, Benston, and Kandel (2000), Kaniel and Liu (2001), and Coppejans and Domowitz (2002).

In their model, greater transparency leads to wider spreads, lower depth, and higher volatility. They also conduct an empirical investigation of the Toronto Stock Exchange's decision in April 1990 to disseminate information about depth at the top four price levels in the book (in addition to the best bid and offer). Since they do not have the detailed order-level data that we have, they are unable to provide evidence about investor strategies or depth in the book, but they do show that spreads are wider after the event and that volatility is higher, both of which are consistent with their theoretical predictions.

Our results contrast with the Toronto Stock Exchange findings and provide support for the view that greater pre-trade transparency is a win-win situation. To examine whether greater pre-trade transparency indeed makes prices more efficient, we use a variance decomposition methodology proposed by Hasbrouck (1993) and document smaller deviations of transaction prices from the efficient (random walk) price. We also find some indication (though weak) of a small reduction in the absolute value of first-order return autocorrelations calculated from quote midpoints. These results are consistent with more efficient prices that are less subject to overshooting and reversal following the introduction of OpenBook.

We then examine measures of liquidity about which we have predictions from the theoretical models: depth in the book and effective spreads (or price impact). We find that depth in the book increases somewhat following the introduction of OpenBook. Most of the increase, however, is in prices away from the current quote. We are unable to examine changes to total depth in the market, because the orders held by floor brokers and the specialist's willingness to provide liquidity behind the quote cannot be measured. Ultimately, however, total depth affects the price impact of trades and we can measure the latter directly using effective spreads. We find that effective spreads of trades decline with the improvement in pre-trade transparency.

The evidence of a decline in effective spreads of trades suggests that the costs incurred by liquidity demanders decrease with the introduction of OpenBook. This evidence may also suggest a decline in investors' compensation for exposing limit orders and supplying liquidity. The decrease in the participation rate of specialists is consistent with such erosion in the profitability of liquidity provision. In an auction market such as the NYSE, an investor can choose whether to be a supplier or a demander of liquidity. Furthermore, acquiring or liquidating a position in a stock can be done with a combination of marketable and limit orders.⁵ Therefore, improved liquidity that is manifested by a smaller price impact of trades does not necessarily mean that the total transaction costs of investors in the new equilibrium are lower than those they experienced under the old regime. We do not have data that allow us to follow the complete trading strategy of an investor and determine its total cost.

However, we can go one step beyond trade effective spreads by focusing on the costs of executing a marketable order at the NYSE. By analyzing orders, as

⁵ By "marketable" we mean both market orders and marketable limit orders. These orders are meant to be executed when they are sent to the market, as opposed to (nonmarketable) limit orders that are meant to stay in the book until prices change and they become executable.

opposed to trades, we can better assess the cost an investor incurs, even when his order is broken down to be executed in multiple trades at different points in time.⁶ We find that effective spreads of orders decline significantly following the introduction of OpenBook. This is true even for small orders, despite a decrease in the depth quoted by the specialist at the bid and ask prices. Our analysis of the change in liquidity around the event uses several econometric models to implement controls and account for potential estimation problems, and the results are robust to the different specifications we use.

While we believe that the effects we document are associated with the increase in transparency that accompanied the introduction of OpenBook, we fully acknowledge that this is an investigation of a single event and that therefore our statistical ability to attribute changes to the event is limited. This issue is a recurring theme in empirical analysis of the financial implications of regulatory changes (see, for example, Schwert (1981)). We believe that the question of how changes in market design affect market quality is important enough to warrant a careful investigation of this particular event. Furthermore, we do address the concern that the changes we document are due to a secular trend in the variables rather than to the introduction of OpenBook. We look at changes in these variables before the event and conclude that the effects we document do not reflect a trend that existed in the market.

With these caveats in mind, our analysis suggests that greater transparency of the limit-order book is beneficial for market quality. This finding is important for several reasons. First, the theoretical literature provides conflicting predictions on how liquidity would change when opening the book, and our results are at odds with those documented when the Toronto Stock Exchange started revealing information about demand in the book. Second, the Securities and Exchange Commission (SEC) has repeatedly emphasized the need for increased pre-trade transparency. Our research is the first empirical study to provide support for such a policy. Third, our results show that market design exerts influence not just on trading strategies, but also on equilibrium liquidity and the informational efficiency of prices. As such, research on market design can help exchanges and regulators improve the functioning of financial markets.

The rest of this paper proceeds as follows. Section I provides details on the OpenBook initiative at the NYSE, describes the event periods, and presents the sample and the data sources used in the investigation. Section II presents the results of our tests concerning the trading strategies of investors, the participation of specialists, informational efficiency, and liquidity. Section III is the conclusion.

I. Research Design

A. OpenBook

Whether or not to make public the content of the limit-order book maintained by specialists at the NYSE has been the subject of discussion for over a decade.

⁶ The analysis of effective spreads of orders follows the conventions set by the SEC for disclosure of execution costs in Rule 11Ac1-5.

In 1991, the NYSE received the approval of the SEC for a program that would have provided snapshots of the book to member firms three times a day. In June of that year, the NYSE announced that it would not implement the system, citing lack of interest among member firms. In 1998, the NYSE announced it was considering providing information about the limit-order book for prices two ticks below and above the best bid and offer. In October 2000, the NYSE again announced intentions to reveal more of the book as part of an initiative called Network NYSE. The implementation was scheduled for the second quarter of 2001, but was postponed. In 2001, the NYSE filed with the SEC for approval of a service called OpenBook that provides information about depth in the book to subscribers, either directly from the NYSE or through data vendors such as Reuters and Bloomberg.

The NYSE's request was approved by the SEC on December 7, 2001, and the OpenBook service was introduced on January 24, 2002, for all NYSE securities simultaneously. OpenBook operates between 7:30 A.M. and 4:30 P.M. It is available for all NYSE-traded securities and shows the aggregate limit-order volume available in the NYSE Display Book system at each price point. The information about depth is updated every 10 seconds throughout the day.⁷ It is important to note that the information disseminated does not include the specialist's proprietary trading interest or floor broker interest. As a result, the information in OpenBook does not reflect total depth in the market, but rather only depth in the limit-order book. Also, OpenBook does not provide any order execution capabilities; it is merely an information dissemination system.

The NYSE charges a fee for the service. Commercial data vendors and large broker-dealers take the raw data directly from the NYSE and pay \$5,000 per month. The NYSE also receives \$50 per month for each subscriber who gets the OpenBook service from a data vendor (or each employee of a broker-dealer with an OpenBook terminal). At the end of 2003, the NYSE was collecting approximately \$870,000 as monthly revenues from the OpenBook service.⁸ Because OpenBook is a paid service, we have a sense of the extent to which this new information is being disseminated. OpenBook had approximately 2,700 subscribers when the service was introduced. This number grew in a steady fashion to about 6,000 during the first 4 months of operation.

B. Event Periods

It is difficult to pinpoint the announcement date for OpenBook. Several times during the decade prior to the actual introduction the idea was announced but did not materialize. Therefore, it is not clear whether the announcement in October 2000, when the NYSE's press release mentioned OpenBook as part of the Network NYSE initiative, had much credibility. Only when the SEC

⁷ While the update is in real time (there is no delay of information), the frequency of updates may be too slow for certain types of automated trading strategies that investors off the exchange floor may want to implement.

⁸ As of October 2003, there were 64 firms getting raw data feeds and over 11,000 subscribers.

approved the service in December 2001 could the NYSE in fact implement the service (though some people might have anticipated it). In contrast, there is no such uncertainty about the implementation date of OpenBook: The service was made available to the public on January 24, 2002. Fortunately, it is the implementation date that matters most for our purpose. While prices may change in anticipation of an event, trading strategies that require information about limit orders in the book cannot be implemented without this information. Therefore, the effects we wish to investigate are best examined around the implementation date.

We are interested in identifying the permanent effects of the change in pre-trade transparency. For that purpose, we need to examine two periods in which the market is in equilibrium with respect to traders' use of order flow information, one before the event and one after the event. We choose 2 weeks (10 trading days) for the length of each period. We believe this choice strikes a balance between our desire to employ more data for the statistical tests on the one hand and both the stability of the estimates and the complexity of handling NYSE order-level data on the other.

Since traders cannot use the information in OpenBook prior to January 24, there is no need to eliminate a long window before the event in order to obtain the steady state of traders' strategies. We choose the full 2 trading weeks prior to the introduction week as the pre-event period (January 7 through January 18). The choice of an appropriate post-event period is more complex. While traders are able to see limit-order book information beginning January 24, learning how to use this information probably takes some time. This is true both for traders who want to use it just to optimize the execution of their orders and for traders who plan to use it to design profitable trading strategies. Furthermore, once such strategies are in place, other traders (e.g., mutual funds' trading desks) may experience poorer execution of their limit orders, prompting more traders to change their strategies until a new equilibrium emerges. Adding to the gradual nature of the process is the fact that the number of subscribers increased in the months following the introduction of OpenBook, which could affect the adjustment of the market to the new pre-trade transparency regime.

To allow for adjustment to an equilibrium state and to examine this adjustment, we use four post-event periods rather than one. As with the pre-event period, we use 2 weeks as the length of a post-event period to capture a reasonably stationary snapshot of the trading environment. More specifically, for each of the first 4 months after the introduction of OpenBook we use the first 2 full weeks of trading: February 4–15, March 4–15, April 1–12, and May 6–17. These four post-event periods enable us to examine how the new equilibrium emerges over time.

C. Sample and Data

The universe of stocks considered for this study includes all common stocks of domestic issuers traded on the NYSE. We eliminate firms that did not trade

continuously between January and May 2002, firms with more than one class of traded shares, closed-end funds, and investment trusts. This results in a population of 1,332 stocks. We then sort by median dollar volume in the last quarter of 2001 and choose a stratified sample of 400 securities that did not experience stock splits or undergo mergers during the sample period. We also divided the sample into four 100-stock groups according to dollar volume in the last quarter of 2001, and conducted the analysis separately for each group. We found that the picture is very similar across groups and therefore we present only the results for the entire sample to simplify the exposition.

Table I provides summary statistics for the entire sample and four trading volume groups. The table testifies to the heterogeneous nature of the sample, which ranges from a median average daily volume of \$59.43 million for the most actively traded group in the pre-event period to a median of \$370,000 for the least actively traded group. All variables—volume, quoted spread, quoted depth, effective spread, and price—change in the expected manner when moving from the most active stocks to the least active stocks. For the most (least) actively traded stocks, median quoted spread is ¢4.4 (¢8.9), and median quoted depth (summing both the bid and ask sides) is 3,445 (1,607) shares. We also observe that prices are higher for the most actively traded stocks in the sample, \$42.74, as compared with \$11.15 for stocks in the least actively traded group.

The data source used for the summary statistics in Table I is the Trade and Quote (TAQ) database distributed by the NYSE. We use these data to analyze effective spreads of trades and informational efficiency.⁹ The rest of our analysis is based on NYSE order-level data provided in the System Order Data (SOD) and Consolidated Equity Audit Trail Data (CAUD) files.¹⁰

The SOD files include detailed information on all orders that arrive at the NYSE via the SuperDot system or that are entered by the specialist into the Display Book system (which powers the limit-order book). SOD files contain about 99% of the orders, representing 75% of NYSE volume, and follow orders from arrival through execution or cancelation. Together with the open limit-order file (LOFOPEN), which describes the exact state of the limit-order book every day before the opening of trading, SOD files allow us to precisely reconstruct the limit-order book on the NYSE at any time. They also enable us to examine how investors change their order submission strategies, to determine how much depth specialists add to the quote beyond what is in the limit-order book, and to compute the effective spreads of orders (as opposed to trades).

⁹ The variables we analyze are calculated using NYSE trades and quotes. We apply various filters to clean the data. We only use trades for which TAQ's CORR field is equal to either zero or one, and for which the COND field is either blank or equal to B, J, K, or S. We eliminate trades with nonpositive prices. We also exclude a trade if its price is greater (less) than 150% (50%) of the price of the previous trade. We eliminate quotes for which TAQ's MODE field is equal to 4, 5, 7, 8, 9, 11, 13, 14, 15, 16, 17, 19, 20, 27, 28, or 29. We exclude quotes with nonpositive ask or bid prices, or where the bid price is higher than the ask. We require that the difference between the bid and the ask be smaller than 25% of the quote midpoint. We also eliminate a quote if the bid or the ask is greater (less) than 150% (50%) of the bid or ask of the previous quote.

¹⁰ A reduced version of these files was the basis for the TORQ database organized by Joel Hasbrouck in 1991. A description appears in Hasbrouck (1992).

Table I
Sample Summary Statistics

The universe of stocks for the study consists of all domestic common stocks listed on the NYSE, excluding firms with multiple classes of shares traded, closed-end funds, and investment trusts. We sort the stocks according to median dollar volume in the last quarter of 2001 and choose a stratified sample of 400 stocks. The table presents summary statistics for the entire sample and for four 100-stock trading volume groups. Five periods are used in the study: the pre-event period (January 7–18) and the four post-event periods: February 4–15, March 4–15, April 1–12, and May 6–17. From the TAQ database, *AvgVol* is the average daily number of shares traded; *QSpread* is the average quoted spread calculated in dollar terms and in percentage terms (the bid–ask spread divided by the quote midpoint); *QDepth* is the average total quoted depth (sum of the depths on the bid and ask sides) measured in dollars and in number of shares; *ESpread* is the average effective spread (twice the distance between the transaction price and the midpoint) in dollars and percentage terms (scaled by the quote midpoint); and *AvgPrc* is the average transaction price of the stock.

		AvgVol (in million \$)		AvgVol (in 100s)		QSpread (in ¢)		QSpread (in %)		QDepth (in \$1,000)		QDepth (in 100s)		ESpread (in ¢)		ESpread (in %)		AvgPrc (in \$)	
		Mean	Median	Mean	Median	Mean	Median	Mean	Median	Mean	Median	Mean	Median	Mean	Median	Mean	Median	Mean	Median
Entire sample	Jan.	31.70	5.69	8498.44	2372.30	7.6	6.2	0.439	0.250	77.69	60.04	30.62	23.34	5.0	4.0	0.296	0.156	30.25	25.49
	Feb.	31.59	5.89	9419.32	2131.70	7.5	6.3	0.458	0.254	76.60	58.76	31.38	23.90	5.0	4.3	0.313	0.160	29.77	24.87
	Mar.	35.16	6.68	9305.33	2505.20	7.1	6.0	0.406	0.228	76.35	57.87	29.37	22.11	4.8	3.9	0.270	0.145	31.64	26.37
	Apr.	30.43	6.11	7675.17	2254.15	6.8	5.7	0.378	0.207	70.91	53.84	26.44	20.35	4.5	3.7	0.248	0.133	32.14	27.04
	May	30.53	5.78	8225.61	2290.75	6.6	5.6	0.372	0.194	73.29	53.80	26.91	20.31	4.4	3.6	0.239	0.126	32.31	27.81
Group 1 (most active stocks)	Jan.	109.80	59.43	26433.18	14941.40	4.9	4.4	0.119	0.108	159.94	131.84	40.56	34.45	3.2	3.0	0.079	0.070	45.86	42.74
	Feb.	108.78	60.61	29888.36	14868.15	5.1	4.6	0.130	0.115	152.89	124.74	41.45	29.80	3.4	3.1	0.087	0.077	44.94	41.25
	Mar.	120.85	67.28	29056.45	17172.45	4.7	4.3	0.114	0.098	156.88	130.14	39.90	33.13	3.2	2.9	0.077	0.066	47.73	43.19
	Apr.	104.20	57.70	23608.54	13764.15	4.4	3.9	0.106	0.090	143.92	123.43	36.86	28.14	2.9	2.6	0.070	0.059	47.82	43.91
	May	102.63	60.53	24815.17	17110.85	4.3	4.0	0.109	0.092	153.54	130.90	40.61	29.49	2.9	2.6	0.073	0.062	47.13	42.23
Group 2	Jan.	13.25	11.03	5345.64	3443.70	6.0	5.6	0.208	0.192	77.85	70.19	28.55	24.37	3.8	3.5	0.133	0.121	32.08	31.21
	Feb.	13.80	11.98	5523.12	3954.65	5.9	5.7	0.214	0.190	78.09	77.67	30.18	23.31	3.8	3.6	0.138	0.122	31.94	31.39
	Mar.	15.18	13.12	5498.91	4198.90	5.7	5.3	0.191	0.169	76.02	70.95	27.07	22.94	3.7	3.4	0.125	0.109	33.87	33.64
	Apr.	13.14	12.10	4604.03	3705.15	5.2	4.9	0.175	0.167	70.77	68.94	24.24	20.48	3.3	3.1	0.111	0.105	34.12	32.80
	May	14.72	11.24	5440.99	3549.60	5.0	4.8	0.173	0.145	70.48	69.80	24.52	20.62	3.2	3.1	0.112	0.094	34.45	33.51
Group 3	Jan.	3.26	2.72	1638.28	1217.05	9.3	7.0	0.355	0.300	49.08	42.68	23.87	18.65	5.9	4.5	0.229	0.194	28.25	23.61
	Feb.	3.18	2.36	1614.53	1165.90	8.8	7.3	0.372	0.306	52.47	44.86	26.89	19.61	5.9	4.6	0.242	0.195	27.61	23.13
	Mar.	3.91	3.19	1937.55	1333.00	8.6	6.6	0.324	0.270	48.35	44.98	22.51	17.42	5.7	4.4	0.208	0.175	29.37	24.81
	Apr.	3.70	2.84	1743.91	1078.30	8.4	6.4	0.302	0.249	45.86	43.16	19.95	15.81	5.5	4.2	0.195	0.156	30.33	25.03
	May	3.78	2.92	1646.25	1212.00	8.2	6.2	0.296	0.236	44.95	39.08	19.51	15.00	5.5	3.9	0.189	0.153	30.70	25.78
Group 4 (least active stocks)	Jan.	0.48	0.37	576.64	326.60	10.3	8.9	1.076	0.833	23.90	19.08	29.49	16.07	7.2	6.3	0.744	0.568	14.82	11.15
	Feb.	0.59	0.43	651.26	361.60	10.0	9.5	1.118	0.849	22.94	19.78	27.01	17.10	7.1	6.4	0.786	0.589	14.59	10.89
	Mar.	0.72	0.54	728.39	445.45	9.6	8.8	0.995	0.751	24.15	20.84	28.00	14.53	6.6	5.9	0.669	0.508	15.57	11.18
	Apr.	0.69	0.58	744.21	432.65	9.4	8.9	0.930	0.689	23.09	19.65	24.70	14.02	6.5	6.1	0.617	0.459	16.30	12.51
	May	1.00	0.61	1000.03	394.40	8.8	8.2	0.909	0.591	24.19	19.57	23.00	14.19	5.9	5.2	0.582	0.382	16.94	13.07

The CAUD files contain detailed execution information on both electronic and manual orders (the latter handled by floor brokers). They enable us to determine the participation rate of specialists in the trading process and the portions of trading volume that originate from either floor brokers or electronic limit orders.

II. Results

Our analysis of the change in pre-trade transparency induced by OpenBook closely follows the exposition of the arguments in the introduction. First, we look at how market participants change their trading strategies as a result of the event. We examine both traders' use of limit orders and specialists' participation in trading and liquidity provision. Second, we examine how these strategies affect the informational efficiency of prices by looking at the deviations of transaction prices from the efficient price and the autocorrelations of quote-midpoint returns. Third, we look at changes in liquidity using depth in the book and effective spreads of both trades and orders. Finally, we examine the question whether our results could be explained by a secular trend in the variables we analyze.

A. Trading Strategies

We use nonparametric univariate tests for the statistical analysis of trading strategies.¹¹ For each period, we compute stock-specific means for all variables. We then report the median across stocks of pairwise differences between each post-event period and the pre-event period, and the p -value from a Wilcoxon test against the two-sided hypothesis that the median is equal to zero. We therefore investigate the total effect of the introduction of OpenBook on strategies, without making an attempt to disentangle which changes represent *direct* effects of the event and which changes are *indirect* effects attributable to changes in other variables. We begin by looking at the conjectures from Harris (1996) that traders react to the risk in order exposure by changing their behavior: canceling and resubmitting limit orders more frequently (shortening the time they are publicly displayed in the book), breaking limit orders into smaller sizes, and making greater use of agents (e.g., floor brokers).

Table II examines the cancelation of limit orders. The first line in Panel A shows an increase in the cancelation rate of limit orders (number of limit orders canceled divided by the number submitted). The median differences between the post- and pre-event periods are positive, and increase monotonically with time. The median change from January to February is 0.68% (though not statistically significant), reaching 4.75% between January and May (and highly statistically significant). The second line in Panel A presents the time-to-cancelation (in seconds) of limit orders that are canceled. It declines following

¹¹ Many of the variables we investigate do not necessarily fit the normality assumption needed for a t -test.

Table II
Analysis of Limit Order Cancellation

This table presents analysis of changes in limit-order cancellation strategies following the introduction of OpenBook. The pre-event period is January 7–18 (Jan.), and the post-event periods are February 4–15 (Feb.), March 4–15 (Mar.), April 1–12 (Apr.), and May 6–17 (May) (each contains 10 trading days). In Panel A, $\Delta\text{CancRate}$ is the change in cancellation rate, defined as the ratio of the number of canceled limit orders to the number of limit orders submitted, and $\Delta\text{TimeCanc}$ is the change in the number of seconds between submission and cancellation of limit orders. We report the cross-sectional median change and the p -value (in parentheses) of a Wilcoxon signed rank test against the hypothesis of a zero median. Panel B presents the results of the duration analysis of time-to-cancellation. For each stock in the sample, we use a parametric approach assuming a Weibull distribution for time-to-cancellation (T):

$$\log T_{it} = \alpha_i + \beta_i I_{it} + \gamma_i \text{Distance from quote}_{it} + \varepsilon_{it}$$
$$\log h_{it}(t) = \alpha_i(t) + \beta_i I_{it} + \gamma_i \text{Distance from quote}_{it} + \varepsilon_{it}$$

In both models, I is a dummy variable that takes the value of one for post-event observations and is zero for pre-event observations. Distance from the relevant quote sides (i.e., bid for sells and ask for buys), standardized by the quote midpoint, is used as a covariate. We report the median of $e^\beta - 1$ (where the coefficients on the dummy variables are estimated separately for each stock), the number of β s significant at the 5% level, and the p -value (in parentheses) of a Wilcoxon signed rank test against the hypothesis of a zero median. **Indicates significance at the 1% level and * indicates significance at the 5% level (both against a two-sided alternative).

Panel A: Analysis of Limit Order Cancellation											
Feb–Jan			Mar–Jan			Apr–Jan			May–Jan		
Variable	Median	(p -Value of Wilcoxon Test)	Median	(p -Value of Wilcoxon Test)	Median	(p -Value of Wilcoxon Test)	Median	(p -Value of Wilcoxon Test)	Median	(p -Value of Wilcoxon Test)	Median
$\Delta\text{CancRate}$	0.0068	(0.1009)	0.0259**	(0.0000)	0.0466**	(0.0000)	0.0475**	(0.0000)			
$\Delta\text{TimeCanc}$	−12.765*	(0.0190)	−25.095**	(0.0000)	−39.902**	(0.0000)	−50.584**	(0.0000)			
Panel B: Duration Analysis											
Feb–Jan			Mar–Jan			Apr–Jan			May–Jan		
$e^\beta - 1$	Median	No. of Significant Out of 400	Median	No. of Significant Out of 400	Median	No. of Significant Out of 400	Median	No. of Significant Out of 400	Median	No. of Significant Out of 400	Median
Weibull	−0.1047**	394	−0.1497**	393	−0.2234**	398	−0.2429**	398	0.1724**	399	
Cox	0.0657**	397	0.1059**	398	0.1700**	397	0.1724**	399			

the event, and declines further with time. Compared to January, time-to-cancellation is 12.77 seconds shorter in February and 50.58 seconds shorter in May. On a pre-event median value of 290 seconds, the decline in time-to-cancellation seems to be quite large (17.4%).

A limitation of the above analysis of time-to-cancellation and the cancellation rate is that it ignores censoring (i.e., limit orders that are executed or expire and therefore cannot be canceled). We use survival (or duration) analysis to estimate two models that take censoring into account (see Lo, MacKinlay, and Zhang (2002) and Hasbrouck and Saar (2004)) on the issue of limit order duration. First, we use an accelerated failure time model that assumes that time-to-cancellation follows a Weibull distribution. The logarithm of time-to-cancellation of limit orders is modeled as a linear function of an intercept, a dummy variable that takes the value 1 after the introduction of OpenBook, and the distance of the limit order from the relevant quote side (the bid for limit buy orders and the ask for limit sell orders) divided by the quote midpoint. The standardized distance from the quote is included as a covariate because it is presumably an important determinant of the probability of both execution and cancellation. The duration model is estimated separately for each stock using all limit orders in the 20-day pre- and post-event periods.

To aid in interpretation of the coefficients, we report in Table II the transformation $e^{\text{coefficient}} - 1$ that provides the percentage change in expected time-to-cancellation between the pre- and post-event periods. The first line in Panel B presents the cross-sectional median of the transformed coefficients on the event dummy variable and the number of the statistically significant coefficients (at the 5% level). In all four post-event periods, the Wilcoxon test is highly significant, and over 394 (of 400) coefficients in the individual stock regressions are statistically significant. For the February post-event period, expected time-to-cancellation of limit orders declines by 10.47%. The decline continues over the sample period and reaches 24.29% by May.

We also report the results of semiparametric Cox regressions (see Cox (1972)), where the logarithm of the hazard rate is modeled as a linear function of an intercept, a dummy variable for the event, and the distance from the quote. While both the Cox model and the Weibull model belong to the class of proportional hazard models, the Cox model does not require that we choose a particular probability distribution for time-to-cancellation. The transformation $e^{\text{coefficient}} - 1$ presented in the second line in Panel B can be interpreted as the percentage change in the estimated cancellation rate of limit orders between the pre- and post-event periods (controlling for the distance from the quote). The results indicate that the cancellation rate increases in a gradual manner: from 6.57% in February to 17.24% in May. The increase in the cancellation rate is highly statistically significant in all four periods.

Panel A of Table III continues our investigation of changes in the trading strategies of investors following the introduction of OpenBook. The first line shows median pairwise differences in the size of limit orders between the post- and pre-event periods. For all four post-event periods, the median changes are negative and statistically different from zero. The magnitude of the changes increases over time after the event. The difference in the size of a typical limit

Table III
Analysis of Trading Strategies

This table presents analysis of changes in trading strategies of investors and specialists following the introduction of OpenBook. The pre-event period is January 7–18 (Jan), and the post-event periods are February 4–15 (Feb), March 4–15 (Mar), April 1–12 (Apr), and May 6–17 (May) (each contains 10 trading days). In Panel A, $\Delta LimitSize$ is the change in the average size of limit orders between the pre- and post-event periods in shares, and $\Delta Floor/Lmt$ is the change in the ratio of the number of shares executed by floor brokers to the number of shares executed using limit orders in the book. Panel B demonstrates the changes in specialist behavior. $\Delta SpecRate$ measures changes in the specialists' participation rate in terms of the number of shares, and $\Delta SpecDepth$ is the change in the specialists' total commitment (in dollars) on the bid and ask sides of the quoted depth. For all variables, the table reports the cross-sectional median and the p -value (in parentheses) of a Wilcoxon signed rank test against the hypothesis of a zero median. ** Indicates significance at the 1% level and * indicates significance at the 5% level (both against a two-sided alternative).

Variable	Feb–Jan		Mar–Jan		Apr–Jan		May–Jan	
	Median	(p -Value of Wilcoxon Test)	Median	(p -Value of Wilcoxon Test)	Median	(p -Value of Wilcoxon Test)	Median	(p -Value of Wilcoxon Test)
Panel A: Differences in Trading Strategies of Investors between Post- and Pre-Event Periods								
$\Delta LimitSize$	–29.513**	(0.0003)	–31.515**	(0.0000)	–55.763**	(0.0000)	–68.423**	(0.0000)
$\Delta Floor/Lmt$	–0.0135	(0.2100)	–0.0323**	(0.0022)	–0.0438**	(0.0001)	–0.0495**	(0.0000)
Panel B: Differences in NYSE Specialists' Behavior between Post- and Pre-Event Periods								
$\Delta SpecRate$	0.0003	(0.5876)	–0.0069**	(0.0001)	–0.0050**	(0.0019)	–0.0088**	(0.0000)
$\Delta SpecDepth$	–1164.76*	(0.0200)	–1320.08	(0.1300)	–2972.83**	(0.0000)	–2599.81**	(0.0000)

order of the same stock between February and January is -29.5 shares, reaching -68.4 in May. On a pre-event median limit-order size of 543 shares, this represents a decline of 12.6%.

The second line in Panel A presents the changes in floor-broker activity relative to electronic limit-order activity. The ratio we compute is the sum of the number of shares bought and sold by floor brokers, divided by the sum of the number of shares bought and sold by limit orders in the book. We document a decline in floor activity relative to limit orders in the book, ranging from -0.014 in February to -0.05 in May (the differences in the last three post-event periods are statistically different from zero).¹² On a pre-event median ratio of 0.52, the magnitude of the decline is almost 10%.

The results are consistent with heightened limit-order exposure management: Smaller limit orders are submitted, limit orders are canceled more often, and limit orders are left for a shorter time in the book. The new ability to see depth in the book seems to make self-management of the trading process more attractive. The shift we document from floor trading to electronic limit orders may indicate that the benefit associated with active trading strategies employed by the traders themselves using OpenBook outweighs the cost of displaying trading interests. The trend in median differences of the variables over the four post-event periods is consistent with the idea that traders learn over time about the new service, learn how to use the information in OpenBook, and adjust their trading strategies accordingly.

The change in pre-trade transparency and the change in the behavior of traders can cause NYSE specialists, who make a market in the stocks, to alter their behavior. We use the CAUD files to examine specialist participation in the trading process. The participation rate is defined as the number of shares bought and sold by the specialist over the total number of shares bought and sold. The first line in Panel B of Table III shows that the specialist participation rate declines in the post-event periods. While the median difference between the first post-event period and the pre-event period is not statistically distinguishable from zero, the median differences for the three other post-event periods are negative and highly statistically significant.¹³

The bid-ask quote disseminated by the NYSE is determined by the specialist. The depth quoted at the bid and ask prices, however, can just reflect the depth available at the best prices in the book. Alternatively, the specialist can add depth to the quote, reflecting the interest of floor brokers or his own interest (in his capacity as a dealer). The second line in Panel B describes the dollar value that specialists (potentially reflecting floor broker trading interest) add to the quoted depth beyond what is in the limit-order book. To create this variable, we use the LOFOPEN and SOD files to reconstruct the book and compare

¹² Separate analysis shows that floor broker activity relative to total volume goes down after the introduction of OpenBook (and the change is statistically significant in three out of the four periods), and electronic limit-order activity relative to total volume increases significantly in all four post-event periods.

¹³ Similar results are obtained when the participation rate is defined in terms of the number of orders rather than the number of shares.

the best prices and depths in the book to the quote disseminated by the specialist every 5 minutes throughout the trading day. We compute the value of the specialist contribution to the quoted depth beyond what is in the limit-order book for each 5-minute snapshot, average over all snapshots, and compute the differences between the post- and pre-event periods for each stock. The specialists' contribution declines monotonically over the four post-event periods, from a median difference of $-\$1,164.76$ to $-\$2,599.81$ (three of the four differences are statistically different from zero).

These results—less participation by the specialists in trading and committing to a smaller quoted depth—are consistent with an increase in the risk associated with the specialists' proprietary trading due to the loss of their information advantage. They are also consistent with a crowding out effect, in that more active management of public limit orders (which have priority over the proprietary trading of specialists) limits the ability of specialists to participate in the trading process. Finally, the reduced depth added by the specialists and floor brokers is also consistent with the shift we observe from floor to electronic limit orders.

B. Information and Prices

Both Glosten (1999) and Baruch (2005) predict that improved transparency would lead to increased informational efficiency of prices. We implement two tests of this hypothesis. The first test is based on the variance decomposition procedure in Hasbrouck (1993). Using information about trade size and execution price for all transactions, Hasbrouck proposes a vector autoregression model to separate the efficient (random walk) price from deviations introduced by the trading process (e.g., short-term fluctuations in prices due to inventory control or order imbalances in the market). More specifically, the variance of log transaction prices, $V(p)$, is decomposed into the variance of the efficient price and the variance of the deviations induced by the trading process, $V(s)$. Because the expected value of the deviations is assumed by the procedure to be zero, the variance is a measure of their magnitude.

The ratio of $V(s)$ to $V(p)$, $VR(s/p)$, reflects the proportion of deviations from the efficient price in the total variability of the transaction price process. If OpenBook allows traders to better time their trading activity to both take advantage of displayed liquidity and provide liquidity in periods of market stress, the proportion of deviations from the efficient price should be smaller after the event. The first line in Table IV shows median changes between the pre- and post-event periods for $VR(s/p)$ (expressing the ratio in percentage terms). While the changes are not significantly different from zero in the February and March post-event periods, they become negative and highly significant in the April and May post-event periods.

Another test of informational efficiency can be formulated by assuming that the quote midpoint is the market's best estimate of the equilibrium value of the stock at every point in time. A more efficient quote-midpoint process would be closer to a random walk and therefore exhibit less autocorrelation (both positive

Table IV
Analysis of Informational Efficiency

This table presents an analysis of informational efficiency around the introduction of OpenBook. The pre-event period is January 7–18 (Jan), and the post-event periods are February 4–15 (Feb), March 4–15 (Mar), April 1–12 (Apr), and May 6–17 (May) (each contains 10 trading days). We use two types of tests to examine changes in the informational efficiency of prices. The first test uses a variable constructed from the variance decomposition procedure in Hasbrouck (1993). ΔVR (s/p) is the change in the ratio (in percentage terms) of the variance of the discrepancies between log transaction prices and the efficient (random walk) price to the variance of log transaction prices. The second test looks at the change in the absolute value of first-order autocorrelations of quote-midpoint returns. We divide the trading day into 30-minute intervals for $\Delta|Corr30|$ and into 60-minute intervals for $\Delta|Corr60|$, and compute the returns from prevailing quote midpoints at the beginning and end of each interval. For all variables, the table reports the cross-sectional median and the p -value (in parentheses) of a Wilcoxon signed rank test against the hypothesis of a zero median. **Indicates significance at the 1% level and *indicates significance at the 5% level (both against a two-sided alternative).

Variable	Feb-Jan		Mar-Jan		Apr-Jan		May-Jan	
	Median	(p -Value of Wilcoxon Test)	Median	(p -Value of Wilcoxon Test)	Median	(p -Value of Wilcoxon Test)	Median	(p -Value of Wilcoxon Test)
ΔVR (s/p)	-0.00047	(0.5294)	0.00043	(0.4390)	-0.00567**	(0.0000)	-0.00691**	(0.0000)
$\Delta Corr30 $	-0.00343	(0.5868)	-0.01181*	(0.0384)	-0.00355	(0.5384)	-0.00100	(0.7996)
$\Delta Corr60 $	-0.00370	(0.2498)	-0.00590	(0.1393)	-0.01424	(0.0906)	-0.01749*	(0.0251)

and negative). The second and third lines in Table IV show changes in the absolute value of the 30-minute and 60-minute first-order quote-midpoint return autocorrelation. For the 30-minute process, we divide the trading day into half-hour intervals and compute the returns from the prevailing quote midpoints at the beginning and end of each interval (a similar construction is used for the 60-minute process). We examine the absolute value of the correlation coefficients because we would like to test how close the return process is to a random walk, which is characterized by zero autocorrelations.

We find that the direction of changes in the autocorrelation is consistent with more efficient prices, but the results are rather weak. While the median changes are negative in all post-event periods, only two of the numbers are statistically different from zero. The results of these two tests together point to some improvement in informational efficiency under the new pre-trade transparency regime. At the very least, the evidence demonstrates that opening the book does not lead to deterioration in the efficiency of prices.

C. Liquidity

What we would like to examine in this section is how the changing strategies of investors and specialists aggregate to create a new state of liquidity provision in the market. This analysis has special significance, since the theoretical arguments we have surveyed disagree on this point—Madhavan, Porter, and Weaver (2000) claim that greater transparency would cause liquidity to deteriorate, while Glosten (1999) and Baruch (2005) claim that it would improve liquidity. In particular, Madhavan, Porter, and Weaver show that depth in the book would decrease and that spreads (or the price impact of trades) would increase when the book is opened. Baruch provides the opposite prediction about spreads, claiming that the price impact of trades would decrease with greater transparency.

What drives the difference in the predictions of these two models? Baruch (2005) features an oligopoly in the market for liquidity provision comprised of fully strategic limit-order traders and a specialist. In equilibrium, these liquidity suppliers have market power and make positive expected profits. Opening the book increases competition among liquidity suppliers, decreasing their profits and improving liquidity. The informed investor in the model trades more aggressively when liquidity is better, which results in more informative prices. In Madhavan, Porter, and Weaver (2000), liquidity provision is competitive, and price impact is derived under the assumption that the expected profit of the limit-order traders is zero. Greater transparency enables traders with private information to make better decisions on order sizes and increase their profits. With fixed information costs, higher profits can translate into more informed trading that would worsen the adverse selection problem and bring about higher spreads (since prices are already determined competitively).

We start by evaluating the prediction on depth in the book. We record snapshots of total depth in the limit-order book for each stock every 5 minutes. We then construct our depth measure by averaging these snapshots for each period.

Because there is much evidence that liquidity is affected by attributes such as volume, we use several parametric approaches to examine the change in liquidity conditional on three control variables. The controls are the average daily dollar volume, intra-day volatility expressed as the average daily range of transaction prices (high minus low), and the average transaction price of the stock (to control for price level effects).

The first econometric specification assumes that the liquidity measure for stock i in period τ (where $\tau \in \{\text{pre}, \text{post}\}$), $L_{i,\tau}$, can be expressed as the sum of a stock-specific mean (μ_i), an event effect (α), a set of control variables, and an error term (η):

$$L_{i\tau} = \mu_i + \alpha\delta_\tau + \beta_1\text{AvgVol}_{i\tau} + \beta_2\text{HiLow}_{i\tau} + \beta_3\text{AvgPrc}_{i\tau} + \eta_{i\tau}, \quad (1)$$

where δ_τ is an indicator variable that takes the value zero in the pre-event period and one in the post-event period, AvgVol represents dollar volume, HiLow is intra-day volatility, and AvgPrc is the price. By assuming that the errors are uncorrelated across securities and over the two periods (although we do not require them to be identically distributed), we can examine differences between the post- and pre-event periods and eliminate the firm-specific mean:

$$\Delta L_i = \alpha + \beta_1\Delta\text{AvgVol}_i + \beta_2\Delta\text{HiLow}_i + \beta_3\Delta\text{AvgPrc}_i + \varepsilon_i, \quad (2)$$

where Δ denotes a difference between the post- and pre-event periods.

We estimate the equation above using OLS and compute test statistics based on White's heteroskedasticity-consistent standard errors. Panel A of Table V presents the intercepts and p -values from the regressions using the change to depth in the book (in round lots) as the liquidity variable. The intercepts for all four post-event periods are positive, and two of the four are statistically significant at the 5% level, indicating some increase in book depth in the post-event period.

Since the event happens to all stocks at the same time, it is possible that the error terms are correlated across stocks. This would cause the standard errors of the intercepts to be biased, but the OLS coefficients would still be consistent. To examine the robustness of our results to this potential problem, we compute the daily values of the variables (e.g., 10 daily averages of book depth indexed by t rather than an average for the entire period), and estimate the following equation pooling all the stocks in our sample:

$$L_{it} = \text{Intercept} + \sum_{k=1}^n (\beta_k \text{Day}_{it}^k) + \gamma_1 \text{Vol}_{it} + \gamma_2 \text{HL}_{it} + \gamma_3 \text{Prc}_{it} + \varepsilon_{it}, \quad (3)$$

where the dummy variables Day_{it}^k ($k = 1, \dots, n$) take the value one for the k^{th} day in the n -day post-event period and zero otherwise, Vol is the daily dollar volume, HL is the daily range of transaction prices, and Prc is the daily average transaction price of the stock. We estimate the model in two ways: (i) for the pre-event period combined with each of the post-event periods (resulting in 10

Table V—Continued

Panel B: Analysis of Book Depth Changes Estimated from Multivariate Regressions at the Daily Frequency														
Feb (8,000 Obs.)			Mar (8,000 Obs.)			Apr (8,000 Obs.)			May (8,000 Obs.)			All Periods (20,000 Obs.)		
Variable	Median β ($n = 10$)	(p -Value of Wilcoxon Test)	Median β ($n = 10$)	(p -Value of Wilcoxon Test)	Median β ($n = 10$)	(p -Value of Wilcoxon Test)	Median β ($n = 10$)	(p -Value of Wilcoxon Test)	Median β ($n = 10$)	(p -Value of Wilcoxon Test)	Median β ($n = 40$)	(p -Value of Wilcoxon Test)		
Depth	13686.06*	(0.0249)	3207.93	(0.6103)	12159.86*	(0.0144)	15118.15**	(0.0080)	12100.86**	(0.0000)				
Panel C: Analysis of Differences in Cumulative Number of Shares Displayed in the Book between Post- and Pre-event Periods														
$\Delta CumShares$	Feb-Jan		Mar-Jan		Apr-Jan		May-Jan							
	Median	(p -Value of Wilcoxon Test)	Median	(p -Value of Wilcoxon Test)	Median	(p -Value of Wilcoxon Test)	Median	(p -Value of Wilcoxon Test)						
0.166% from quote	-1.30	(0.7981)	97.78	(0.1273)	-128.43	(0.4355)	-124.64	(0.0623)						
0.833% from quote	-216.74	(0.4294)	381.35*	(0.0389)	-136.96	(0.3909)	-124.62	(0.2272)						
3.333% from quote	-193.64	(0.7245)	1072.80**	(0.0050)	-130.44	(0.8803)	-309.11	(0.6914)						
16.67% from quote	127.26	(0.5025)	2081.44**	(0.0044)	1059.48*	(0.0293)	485.07	(0.3272)						
Entire Book	1818.60*	(0.0275)	887.52	(0.1304)	366.11	(0.5364)	-838.74	(0.8916)						

coefficients of the daily post-event dummy variables), and (ii) for the entire sample period of 50 pre-event and post-event days (resulting in 40 coefficients of the daily post-event dummy variables).

Panel B of Table V reports the median of the coefficients on the post-event dummy variables and the p -value (in parentheses) of a Wilcoxon signed rank test against the hypothesis of a zero median. The idea behind the test is similar in spirit to the one underlying the Fama and MacBeth (1973) specification. The OLS coefficients on the dummy variables are consistently estimated even with cross-correlated errors, and therefore a test that uses time-series variation in the coefficients is not affected by this potential problem. Three of the four post-event periods show a significant increase in book depth after the introduction of OpenBook, and so does the regression over all periods together.¹⁴

While our analysis so far demonstrates that depth in the book increases following the introduction of OpenBook, the summary statistics reported in Table I suggest that depth quoted by the specialists decreased over the sample period. The median difference between dollar quoted depth in May and January is $-\$3,276$, a statistically significant decrease of about 5.4% from the median quoted depth in January. We know from Section II.A that the specialist adds less depth to the quote following the introduction of OpenBook. The decline in quoted depth suggests that book depth close to current market prices does not increase sufficiently to compensate for the change in the specialist's behavior. It is therefore possible that the increase in book depth is primarily at prices that are further away from the current market price.

To examine this possibility, we compute cumulative depth in the book at five distances from the relevant side of the quote (the ask for limit sell orders and the bid for limit buy orders). To better compare the liquidity of stocks with different share prices, we define the distance from the relevant side of the quote as a percentage of the stock's price: (i) up to 0.166% of the price, (ii) up to 0.833% of the price, (iii) up to 3.333% of the price, (iv) up to 16.667% of the price, and (v) up to the end of the book.¹⁵ These percentage bounds are applied to the price of each stock to create stock-specific dollar bounds for the five categories of cumulative depth in the book that we use.

¹⁴ We implemented yet another procedure to examine robustness to the potential problem of cross-correlated errors by computing cross-sectional daily averages of each variable and using OLS to estimate the following time-series model:

$$CSL_t = \text{Intercept} + \sum_{k=1}^4 (\beta_k D_t^k) + \gamma_1 CSVol_t + \gamma_2 CSHL_t + \gamma_3 CSPrc_t + \varepsilon_t,$$

where the dummy variable D_t^k ($k = 1, 2, 3, 4$) takes the value of one for the k^{th} post-event period and is zero otherwise, $CSVol$ is the cross-sectional average of daily dollar volume, $CSHL$ is the cross-sectional average of intra-day volatility, and $CSPrc$ is the daily mean transaction price averaged across stocks. All four coefficients were positive and statistically significant in the depth regression.

¹⁵ The bounds of the categories were chosen by deciding on several cutoffs in dollar terms ($\$5$, $\$25$, $\$1$, and $\$5$) and dividing them by $\$30$, the average share price of stocks in our sample in the pre-event period. The results are qualitatively similar when we use the dollar bounds.

Panel C of Table V presents the median change in cumulative depth in the book for each category. Of the 20 numbers in the table (five categories, four post-event periods), only five are statistically significant at the 5% level. All of these numbers are positive, indicating an increase in depth (without controlling for changes in volume, volatility, or price). A closer inspection reveals differences that are negative but insignificant in the first two categories (i.e., close to the current market price) in three out of the four periods. On the other hand, the change in cumulative depth in the book up to a distance of 16.67% from the relevant side of the quote is positive in all four periods, and is statistically significant in two out of the four. This may indicate that total depth available for small orders decreases following the introduction of OpenBook, while the total depth available for larger orders increases.

Our tests of depth in the book are motivated by the predication from Madhavan, Porter, and Weaver (2000), but what matters for liquidity in a hybrid market such as the NYSE is the total available depth. At the NYSE, one can think about total depth as the sum of three components: depth in the book, depth on the floor (orders represented by floor brokers), and depth added by the specialist in his capacity as a dealer. While we can examine the change in book depth rather accurately, the other two categories are more difficult to measure. In Section II.A, we show that the contribution of the specialist to the quoted depth beyond what is in the limit-order book decreases following the introduction of OpenBook. However, we are unable to evaluate the change in the specialist's incentives to provide depth at prices behind the best bid and offer (in case the quote is exhausted). Also, we have no way of assessing depth available on the floor, as most floor orders are being worked by floor brokers rather than given to the specialist.

While we cannot judge the change to total depth, the total price impact of an order (or its effective spread) depends on all available sources of depth in the market, and therefore it reflects the change in total depth and provides a good summary measure of the change in liquidity. We also have specific predictions about this variable from the theoretical models: Baruch (2005) states that the price impact should decrease while Madhavan, Porter, and Weaver (2000) claim it should increase. The trade effective spread measure is computed by averaging twice the distance between the transaction price and the prevailing quote midpoint for all transactions in a period (from the TAQ database). The analysis of effective spreads is conducted using the same econometric specifications used for depth with volume, volatility, and price controls.

Panel A of Table VI presents the results of the econometric specification in (2), using differences in trade effective spreads (in cents) as the dependent variable. All coefficients are negative, and they increase in magnitude over time from -0.1046 in February to -0.8001 in May (where the last three post-event periods are statistically significant). Since effective spreads measure the cost of trading and volume measures the quantity of trading, it can be argued that a single-equation specification regressing effective spreads on volume suffers from an endogeneity problem. We examined this potential problem using a simultaneous-equation model of spreads and volume, and the results were

Table VI—Continued

Panel B: Analysis of Trade Effective Spread Changes Estimated from Multivariate Regressions at the Daily Frequency																
Feb (8,000 Obs.)			Mar (8,000 Obs.)			Apr (8,000 Obs.)			May (8,000 Obs.)			All Periods (20,000 Obs.)				
Variable	(p-Value of Wilcoxon Test)		(p-Value of Wilcoxon Test)		Median β (n = 10)	(p-Value of Wilcoxon Test)		Median β (n = 10)	(p-Value of Wilcoxon Test)		Median β (n = 40)	(p-Value of Wilcoxon Test)				
	Median β (n = 10)	(p-Value of Wilcoxon Test)	Median β (n = 10)	(p-Value of Wilcoxon Test)		Median β (n = 10)	(p-Value of Wilcoxon Test)		Median β (n = 40)	(p-Value of Wilcoxon Test)						
$ESpread_t$	-0.0087	(0.9188)	-0.4048**	(0.0059)	-0.8600**	(0.0059)	-1.1311**	(0.0059)	-0.6562**	(0.0000)						
Panel C: Differences in Order Effective Spreads in a Cross-Sectional Multivariate Regression (400 Observations)																
$\Delta ES_{spread,t}$ by Order Size	α	Feb-Jan		Mar-Jan		Apr-Jan		May-Jan		α	(p-Value of t-Statistic)	Adj R^2 (in %)	(p-Value of t-Statistic)	Adj R^2 (in %)	(p-Value of t-Statistic)	Adj R^2 (in %)
		(p-Value of t-Statistic)	Adj R^2 (in %)	(p-Value of t-Statistic)	Adj R^2 (in %)	(p-Value of t-Statistic)	Adj R^2 (in %)	(p-Value of t-Statistic)	Adj R^2 (in %)							
[1, 499] shares	0.0305	(0.7008)	3.31	-0.2610*	(0.0437)	-0.14	-0.4603**	(0.0000)	3.01	-0.5939**	(0.0000)	7.60				
[500, 1,999] shares	-0.3994*	(0.0257)	8.35	-0.4160*	(0.0155)	1.82	-0.8409**	(0.0003)	3.44	-0.7931**	(0.0003)	13.80				
[2,000, 4,999] shares	-1.6877	(0.2155)	4.09	-2.6801	(0.2524)	6.82	-0.1386	(0.6470)	4.48	-1.6680*	(0.0286)	24.42				
[5,000, 9,999] shares	-1.6048	(0.3251)	0.86	-1.5943	(0.1650)	-0.42	-1.2825	(0.3012)	-0.16	-1.8848	(0.0895)	0.35				
[10,000, ∞] shares	-1.5948	(0.2376)	2.87	-0.9467	(0.3780)	0.87	-2.3163*	(0.0109)	2.73	-1.6783	(0.0883)	2.20				
All orders	-1.1688*	(0.0325)	4.53	-0.7978	(0.1386)	2.61	-1.1946	(0.0565)	6.87	-1.7242**	(0.0000)	0.89				
Panel D: Analysis of Order Effective Spread Changes Estimated from Multivariate Regressions at the Daily Frequency																
Feb (8,000 Obs.)			Mar (8,000 Obs.)			Apr (8,000 Obs.)			May (8,000 Obs.)			All Periods (20,000 Obs.)				
$ES_{spread,t}$ by Order Size	(p-Value of Wilcoxon Test)		(p-Value of Wilcoxon Test)		Median β (n = 10)	(p-Value of Wilcoxon Test)		Median β (n = 10)	(p-Value of Wilcoxon Test)		Median β (n = 40)	(p-Value of Wilcoxon Test)				
	Median β (n = 10)	(p-Value of Wilcoxon Test)	Median β (n = 10)	(p-Value of Wilcoxon Test)		Median β (n = 10)	(p-Value of Wilcoxon Test)		Median β (n = 40)	(p-Value of Wilcoxon Test)						
[1, 499] shares	0.1573	(0.1851)	-0.4356**	(0.0059)	-0.6559**	(0.0080)	-0.7718**	(0.0059)	-0.4942**	(0.0000)						
[500, 1,999] shares	-0.0445	(0.3081)	-0.4677**	(0.0059)	-0.9131**	(0.0059)	-1.0761**	(0.0080)	-0.5594**	(0.0000)						
[2,000, 4,999] shares	-0.2889	(0.5408)	-0.5445*	(0.0144)	-0.7452*	(0.0249)	-1.1439**	(0.0059)	-0.5979**	(0.0000)						
[5,000, 9,999] shares	-0.2684	(0.2213)	0.0505	(1.0000)	-0.4224*	(0.0323)	-1.2136**	(0.0059)	-0.3736**	(0.0011)						
[10,000, ∞] shares	-0.9469*	(0.0249)	-0.9690	(0.4148)	-0.9027	(0.0528)	-1.2217**	(0.0059)	-1.0163**	(0.0001)						
All orders	-0.2414*	(0.0191)	-0.3960*	(0.0108)	-0.9472**	(0.0059)	-1.5686**	(0.0059)	-0.6899**	(0.0000)						

similar to those from the single-equation specification.¹⁶ Panel B of Table VI presents the results of the trade effective spreads analysis with the daily time-series specification in (3). Here as well, there is a highly significant decline in effective spreads in three of the four post-event periods and in the regression for the entire sample period.

While we find that effective spreads of trades decrease, Madhavan, Porter, and Weaver (2000) document an increase in spreads following the change in pre-trade transparency implemented by the Toronto Stock Exchange in 1990. What can explain the conflicting results? The answer probably does not lie in differences in market structure between the two exchanges, because Madhavan, Porter, and Weaver document the results for stocks traded on the Toronto Stock Exchange floor that features Registered Traders who are similar to specialists. The conflicting results may be due to developments over the past decade in information processing, order handling, and trading technologies. The ability of buy-side traders to utilize information about the limit-order book to improve their trading strategies is much greater today, which may be the reason why we see emergence of an equilibrium where liquidity improves.¹⁷

It is important to note that the effective spread of a trade does not constitute a perfect measure of transaction costs. A portfolio manager or an individual investor who seeks to change his position may specify a trading strategy that incorporates both marketable and limit orders that can be executed at multiple points in time.¹⁸ If indeed the cost of marketable orders (effective spreads) decreases, it is reasonable to assume that the gain to limit orders that supply liquidity decreases as well. In an auction market, an investor can both demand and supply liquidity, and hence the impact on his overall transaction costs is ambiguous. Furthermore, even market and marketable limit orders that are

¹⁶ The specification of the simultaneous-equation model we used is:

$$\begin{aligned}\Delta ES_{i,t} &= \alpha + \beta_1 \Delta AvgVol_i + \beta_2 \Delta HiLow_i + \beta_3 \Delta AvgPrc_i + \beta_4 \Delta StdInv_i + \varepsilon_i, \\ \Delta AvgVol_i &= \beta_5 + \beta_6 \Delta ES_{i,t} + \beta_7 \Delta HiLow_i + \beta_8 \Delta AvgPrc_i + \beta_9 \Delta SysVol_i + v_i,\end{aligned}$$

where $\Delta StdInv$ is the standard deviation of daily inventory closing positions of specialists (from the NYSE's SPETS file), and $\Delta SysVol$ is the systematic component of dollar volume. The systematic component is obtained from a market model of dollar volume using one year of daily data ending before the beginning of the pre-event period, with an equally weighted portfolio of all common domestic NYSE stocks as a proxy for the market (see Lo and Wang (2000) and Llorente et al. (2002) on the issue of a market model for volume). The simultaneous-equation model was estimated using both two-stage least squares and three-stage least squares. For all post-event periods, the magnitudes of the intercepts from the effective spreads equation and their statistical significance were almost identical to those reported in Table VI, and are therefore omitted for brevity.

¹⁷ In Baruch (2005), the proof of the result that liquidity improves with greater transparency requires a sufficiently large market (i.e., a lower bound on the number of limit-order traders). The average share volume in our sample is almost 20 times greater than the average share volume of the floor-traded securities in Madhavan, Porter, and Weaver's study, and therefore it is possible that Baruch's prediction is applicable to the NYSE, but not to the Toronto Stock Exchange in 1990.

¹⁸ See, for example, the experimental evidence in Bloomfield, O'Hara, and Saar (2003) on the use of market and limit orders by traders.

sent to the exchange can be broken down into several trades. A better measure of the execution cost an investor incurs would benchmark the entire order to the market price when the order is sent to the exchange.

We do not have data that enable us to follow the portfolio changes of specific investors and examine their equilibrium trading strategies and the resulting transaction costs. However, we can go one step beyond looking at the price impact of trades by following the analysis of execution costs proposed by the SEC in Rule 11Ac1-5.¹⁹ This rule, which went into effect in July 2001, dictates that all market centers should provide several measures to help investors compare the quality of execution across trading venues. Unlike the analysis of effective spreads that we provided above, the effective spreads the SEC mandates in Rule 11Ac1-5 are computed for *orders*, rather than for trades.²⁰

To demonstrate the advantage of using the effective spread of orders as a measure of transaction costs, consider the following hypothetical example. Say a marketable buy order for 1,500 shares arrives at the NYSE via the SuperDOT system at 10:15:01 A.M. The quote midpoint at the arrival time is \$30.01. The order is executed in two trades: 500 shares at 10:15:03 A.M. for a price of \$30.03, and 1,000 shares at 10:15:10 A.M. for a price of \$30.05. The quote changes in-between these two trades and the new quote midpoint at 10:15:07 A.M. is \$30.03. The trade effective spread would have two observations: ¢2 per share for the 500-share trade and ¢2 per share for the 1,000-share trade. The effective spread for the order, however, is computed against the quote midpoint at the time of arrival, and is weighted by the shares in the orders, so that the effective spread of the order would be $[500 * (30.03 - 30.01) + 1000 * (30.05 - 30.01)]/1,500 = \text{¢}3.33$.²¹

The data that market centers publicly report to comply with the requirements of the SEC are monthly averages of the measures. We conducted an analysis of differences using December 2001 as the pre-event period and February through May 2002 as post-event periods. The results indicated a reduction in effective spreads of orders in all four post-event periods, with a statistically significant difference in three out of the four periods.²² Our NYSE data, however, enable us to take the SEC definition of effective spreads for orders and compute it ourselves for the intervals we need in the pre- and post-event periods. We therefore use the SOD files to generate effective spreads for market and marketable limit orders for the pre- and post-event periods. In addition, we follow Rule 11Ac1-5 by putting orders into size categories in terms of the number of shares in an order ([1, 499], [500, 1,999], [2,000, 4,999], [5,000, 9,999], and [10,000, ∞]).²³

¹⁹ We thank the referee for suggesting the analysis of Rule 11Ac1-5 data.

²⁰ See Bessembinder (2005), Boehmer (2003), and Lipson (2003) for a discussion of Rule 11Ac1-5 and analysis of data that market centers release to comply with the rule.

²¹ The order effective spread of a stock for a certain period is computed as the share-weighted average of effective spreads across all orders in the period.

²² The results are available from the authors.

²³ Rule 11Ac1-5 specifies the first four categories that we use (up to 9,999 shares). We add the fifth category for completeness.

Panel C of Table VI presents the results of the econometric specification in (2), where differences in order effective spreads between the pre- and post-event periods for each order size category are regressed on an intercept and changes in the control variables. We see that the effective spread of small orders, up to 499 shares, decreases significantly in the last three post-event periods. Similarly, there is a significant decrease in effective spreads for orders between 500 and 1,999 shares in all post-event periods. In fact, the intercepts of all post-event periods and all size categories except one are negative. A decrease in the effective spreads of small orders, despite the decrease in quoted depth, can be explained by the increase in limit-order volume (see footnote 14) coupled with a better ability to make decisions about order types (e.g., taking liquidity when the book is thick, while supplying liquidity when it is thin). In other words, timing market conditions can lead to a reduction in the price impact of marketable orders, irrespective of the change in average depth.

Panel D of Table VI presents the results of the order effective spread analysis using the econometric specification in (3), where daily values of order effective spreads are regressed on dummy variables and the three control variables. Here the evidence of a decrease in order effective spreads is even more pronounced—there is a significant decrease in all size categories by the May post-event period. Similarly, we observe a significant decrease in the pooled regression using all 50 days in the sample period.

The fact that the results are qualitatively similar using different econometric specifications and multiple measures of liquidity indicates that our findings are rather robust. The evidence we present suggests that liquidity improved following the introduction of OpenBook.²⁴

D. Could the Results Just Reflect a Trend?

In the introduction we noted that our study of changes in pre-trade transparency uses a single event: the introduction of OpenBook. As such, we rely on the cross-section of stocks to provide us with statistical significance for the changes we document. In principle, it is possible that these changes reflect a trend that existed in the market even before the introduction of OpenBook.

To examine this issue, we would like to look at the variables prior to the introduction of OpenBook. We therefore proceed to take another 10-day period, August 27–September 10, 2001, and test the differences in the variables from this period to the January pre-event period. Our choice of a “robustness” period is driven by three considerations. First, we want the interval of time between the robustness period and the January pre-event period to be similar to the interval between one of the post-event periods and the pre-event period, so

²⁴ It is interesting to note that the effects of changes in pre-trade transparency may differ from the effects of changes in the degree of anonymity in a market. Foucault, Moinas, and Theissen (2003) look at Euronext Paris, which is organized as a pure electronic limit-order book. They find that decreasing the amount of information by concealing the identities of traders improves liquidity, while we find for the NYSE that an improvement in pre-trade transparency, holding anonymity constant, improves liquidity.

that we can compare the magnitude of the changes. Second, we want to avoid looking at the variables in the immediate aftermath of September 11, 2001, due to the unusual market conditions that prevailed. Third, we do not want to take a period too close to January 2001 in order not to pick up the effects of changes in the minimum tick size. The August 27–September 10 period seems like a reasonable choice that balances these considerations (and the time interval between this period and January 2002 is similar to the one between the pre-event period and the May post-event period).

Table VII shows the changes from the robustness period to the pre-event period for the variables we investigated in Sections II.A–II.C. In general, the magnitude of the changes is much smaller than those we find following the introduction of OpenBook, and their sign is often in the opposite direction. Panel A reports the results on trading strategies. The cancellation rate of limit orders decreased prior to OpenBook, time-to-cancellation increased, and limit-order size increased. These are statistically significant changes, but they go in the opposite direction of the results we document after the introduction of OpenBook. There is no statistically significant change prior to OpenBook in the floor-to-limit ratio, the specialist participation rate, or the contribution of the specialists to the depth of the quote. In contrast, we find significant changes in these variables from January to the May post-event period. It therefore seems unlikely that the changes in trading strategies following the introduction of OpenBook are simply a manifestation of a trend.

Panel B of Table VII presents the results of changes in the informational efficiency variables. None of the changes is statistically significant (and they are all positive, as opposed to the negative changes we document following the introduction of OpenBook).

Panels C and D contain the analysis of liquidity using the econometric specifications in (2) and (3). We observe that depth in the book declines from the robustness period to the pre-event period, and so our finding in Section II.C of an increase in depth does not reflect a trend that existed prior to the introduction of OpenBook. Trade effective spreads seem to decline, but the results are rather weak. The intercept from the regression of changes in trade effective spreads is weakly significant, but much smaller in magnitude than the results we document over a similar time interval from January to May (Table VI, Panel A). The change in trade effective spread is not statistically significant in the multivariate regression at the daily frequency that we implement to mitigate the effect of cross-correlated errors. The effective spread of orders demonstrates no significant change from the robustness period to the January pre-event period.²⁵ It therefore seems that the pronounced change in liquidity that we observe after the introduction of OpenBook cannot be solely attributed to a trend.

²⁵ For the sake of brevity, Table VII omits the analysis of depth and effective spreads by categories. There was no significant increase in depth from the robustness period to January in any of the distance-from-quote categories, and no significant change in order effective spreads was found in any of the size categories. These results are available from the authors.

Table VII
Investigating the Possibility of a Trend

This table presents results of changes in trading strategies, informational efficiency, and liquidity in the months before the event. These results can be used to judge whether there is a secular trend in the variables we analyze in Tables II–VI. We examine the differences in these variables between the 10-day period of August 27–September 10, 2001 (Sep01) and the pre-event period of January 7–18, 2002 (Jan02). In Panel A, $\Delta\text{CancRate}$ is the change in cancellation rate, defined as the ratio of the number of canceled limit orders to the number of limit orders submitted; $\Delta\text{TimeCanc}$ is the change in the number of seconds between submission and cancellation of limit orders; $\Delta\text{LimitSize}$ is the change in the average size of limit orders in shares; $\Delta\text{Floor/Lmt}$ is the change in the ratio of the number of shares executed by floor brokers to the number of shares executed using limit orders in the book; $\Delta\text{SpecRate}$ measures the change in the specialists' participation rate in terms of number of shares; and $\Delta\text{SpecDepth}$ is the change in the specialists' total commitment (in dollars) on the bid and ask sides of the quoted depth. In Panel B, $\Delta\text{VR}(s/p)$ is the change in the ratio (in percentage terms) of the variance of the discrepancies between log transaction prices and the efficient (random walk) price to the variance of log transaction prices; $\Delta|\text{Corr30}|$ and $\Delta|\text{Corr60}|$ measure the changes in the absolute value of first-order autocorrelations of quote-midpoint returns using 30-minute and 60-minute returns, respectively. For all variables in Panels A and B, the table reports the cross-sectional median and the p -value (in parentheses) of a Wilcoxon signed rank test against the hypothesis of a zero median. Panels C and D present the results of the analysis of changes in liquidity variables. A description of the regressions in these panels is detailed in the text of Tables V and VI. **Indicates significance at the 1% level and * indicates significance at the 5% level (both against a two-sided alternative).

Panel A: Differences in Trading Strategies Jan02–Sep01			
Variable	Median	(<i>p</i> -Value of Wilcoxon test)	
$\Delta CancRate$	−0.0229**	(0.0000)	
$\Delta TimeCanc$	20.620**	(0.0000)	
$\Delta LimitSize$	49.902**	(0.0000)	
$\Delta Floor / Lmt$	−0.0058	(0.2321)	
$\Delta SpecRate$	−0.0017	(0.1969)	
$\Delta SpecDepth$	964.30	(0.6400)	
Panel B: Differences in Informational Efficiency Jan02–Sep01			
Variable	Median	(<i>p</i> -Value of Wilcoxon Test)	
$\Delta VR (s/p)$	0.00003	(0.4783)	
$\Delta Corr30 $	0.00031	(0.9891)	
$\Delta Corr60 $	0.01918	(0.0845)	
Panel C: Differences in Liquidity Variables in a Cross-Sectional Multivariate Regression Jan02–Sep01			
Variable	α	(<i>p</i> -Value of <i>t</i> -Statistic)	Adj <i>R</i> ² (in %)
$\Delta Depth$	−1683.49**	(0.0014)	11.75
Trades $\Delta ES_{spread}t$	−0.3238*	(0.0291)	3.43
Orders $\Delta ES_{spread}t$	0.6964	(0.1450)	7.59
Panel D: Analysis of Liquidity Changes Estimated from Multivariate Regressions at the Daily Frequency Jan02–Sep01			
Variable	Median β (<i>n</i> = 10)	(p-Value of Wilcoxon test)	
<i>Depth</i>	−2799.40**	(0.0059)	
Trades <i>ES</i> _{spread} <i>t</i>	−0.2965	(0.1029)	
Orders <i>ES</i> _{spread} <i>t</i>	0.3344	(0.2622)	

III. Conclusions

The structure of securities markets in the United States and around the world is undergoing many changes. Various competing market structures were introduced into U.S. equity trading by alternative trading systems, and the SEC has been instrumental in providing the conditions that helped these new trading platforms to flourish. In particular, the SEC has pushed for greater pre-trade transparency by mandating certain display requirements for limit orders (affecting both market makers and alternative trading systems organized as electronic limit-order books). The insistence of the SEC on pre-trade transparency stems from its conviction that transparency improves not just price discovery, but also the fairness, competitiveness, and attractiveness of U.S. markets (see U.S. SEC (1994)).

So far, however, there has been no consensus in the academic literature on whether greater pre-trade transparency in the sense of disclosing more information about limit orders in the book is beneficial. Our results provide empirical support, for the first time, for the view that improved pre-trade transparency of a limit-order book can be good for investors. Our focus on the largest equity market in the world and the quality of the order-level data we use make our findings even more significant in the debate among academics and policy makers.

We find that investors do change their strategies in response to the change in market design: They submit smaller limit orders and cancel limit orders in the book more quickly and more often. These findings are consistent with a more active management of trading strategies in the face of greater risk of order exposure. Additionally, we find that traders shift activity away from floor brokers toward electronically submitted limit orders. This may indicate that OpenBook enables traders to implement more complex strategies themselves, and therefore reduces their need to delegate that responsibility to floor brokers. We also find that NYSE specialists change their behavior in that they trade less, and together with floor brokers, add less depth to the quote. These changes can reflect the increased risk in proprietary trading without the help of privileged information about the book, a crowding out effect that results from increased competition provided by active limit-order trading, and a shift in investor strategies from using floor brokers to using limit orders submitted electronically via SuperDot.

The equilibrium effects on the state of the market, both in terms of liquidity and informational efficiency, seem to suggest that increased transparency is a win-win situation. We find some improvement in informational efficiency following the introduction of OpenBook, an increase in displayed liquidity in the book, and a decline in the price impact of trades and marketable orders.

The results we document point to two welfare redistributions that are possibly associated with the introduction of OpenBook. The first is from liquidity suppliers to demanders. The decrease in the price impact of trades and marketable orders reduces the compensation for liquidity provision, hurting limit-order suppliers and specialists. The second is from NYSE members to the

exchange itself. We document a decrease in the specialist participation rate, and the evidence of a shift from floor to limit orders is consistent with a decline in the business of floor brokers. At the same time, the NYSE generates revenues from the OpenBook service.

Since we do not observe the complete trading strategy of each investor, we are unable to judge whether the trading costs of investors who utilize both market and limit orders in the new regime are lower than the trading costs when traders did not have information about the book. Nonetheless, there are some reasons to believe that improved liquidity per se can be beneficial for a market. Liquid markets may encourage investor participation. Firms may benefit from liquid markets either because some aspects of liquidity are directly priced (e.g., Amihud and Mendelson (1986) and Easley and O'Hara (2004)) or because greater investor participation lowers the required return in the spirit of Merton (1987). It is important to note, however, that we do not provide evidence for the changes to the total welfare of investors or the cost of capital of firms, but rather document improvements in liquidity and informational efficiency following the introduction of OpenBook.

Beyond providing support for the SEC's beliefs about the importance of transparency to the quality of U.S. markets, our analysis provides evidence that market design indeed has substantial implications for investors. NYSE material stresses that OpenBook was designed to increase transparency in a decimal trading environment. Because the idea of publicly distributing information about the book has been around for many years, implementation at this time indeed seems to have been in response to the change in the tick size. In a sense, one regulatory change in market design (the reduction in tick size) caused the NYSE to implement another change in market design (improving pre-trade transparency). The current securities trading environment is characterized both by frequent regulatory interventions and by competitive pressures. The experience with OpenBook suggests that markets can and should respond to changes in their environments by modifying the design of their trading systems to better meet investor needs.

REFERENCES

- Ahn, Hee-Joon, Kee-Hong Bae, and Kalok Chan, 2001, Limit orders, depth and volatility: Evidence from the Stock Exchange of Hong Kong, *Journal of Finance* 56, 767–788.
- Amihud, Yakov, and Haim Mendelson, 1986, Asset pricing and the bid-ask spread, *Journal of Financial Economics* 17, 223–249.
- Baruch, Shmuel, 2005, Who benefits from an open limit-order book? *Journal of Business* 78, forthcoming.
- Bessembinder, Hendrik, 2003, Selection biases and cross-market trading costs comparisons, Unpublished manuscript, University of Utah.
- Biais, Bruno, 1993, Price formation and equilibrium liquidity in fragmented and centralized markets, *Journal of Finance* 48, 157–185.
- Biais, Bruno, Pierre Hillion, and Chester Spatt, 1995, An empirical analysis of the limit order book and the order flow in the Paris Bourse, *Journal of Finance* 50, 1655–1689.
- Bloomfield, Robert, and Maureen O'Hara, 1999, Market transparency: Who wins and who loses? *Review of Financial Studies* 12, 5–35.

- Bloomfield, Robert, Maureen O'Hara, and Gideon Saar, 2003, The "make or take" decision in an electronic market: Evidence on the evolution of liquidity, *Journal of Financial Economics*, forthcoming.
- Boehmer, Ekkehart, 2005, Dimensions of execution quality: Recent evidence for U.S. equity markets, *Journal of Financial Economics*, forthcoming.
- Chowdhry, Bhagwan, and Vikram Nanda, 1991, Multimarket trading and market liquidity, *Review of Financial Studies* 4, 483–511.
- Cohen, Kalman J., Steven F. Maier, Robert A. Schwartz, and David K. Whitcomb, 1981, Transaction costs, order placement strategy, and existence of the bid–ask spread, *Journal of Political Economy* 89, 287–305.
- Coppejans, Mark, and Ian Domowitz, 2002, An empirical analysis of trades, orders, and cancellations in a limit order market, Unpublished manuscript, Duke University.
- Corwin, Shane A., and Mark L. Lipson, 2000, Order flow and liquidity around NYSE trading halts, *Journal of Finance* 55, 1771–1800.
- Cox, David R., 1972, Regression models and life tables, *Journal of the Royal Statistical Society, Series B*, 34, 187–220.
- Easley, David, and Maureen O'Hara, 2004, Information and the cost of capital, *Journal of Finance* 59, 1553–1583.
- Fama, Eugene F., and James D. MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636.
- Flood, Mark D., Ronald Huisman, Kees G. Koedijk, and Ronald J. Mahieu, 1999, Quote disclosure and price discovery in multiple-dealer financial markets, *Review of Financial Studies* 12, 37–59.
- Foucault, Thierry, 1999, Order flow composition and trading costs in a dynamic limit order market, *Journal of Financial Markets* 2, 99–134.
- Foucault, Thierry, Sophie Moinas, and Erik Theissen, 2003, Does anonymity matter in electronic limit order markets? Unpublished manuscript, HEC and Bonn University.
- Franks, Julian R., and Stephen M. Schaefer, 1995, Equity market transparency on the London Stock Exchange, *Journal of Applied Corporate Finance* 8, 70–77.
- Glosten, Lawrence R., 1994, Is the electronic open limit order book inevitable? *Journal of Finance* 49, 1127–1161.
- Glosten, Lawrence R., 1999, Introductory comments: Bloomfield and O'Hara, and Flood, Huisman, Koedijk, and Mahieu, *Review of Financial Studies* 12, 1–3.
- Handa, Puneet, and Robert A. Schwartz, 1996, Limit order trading, *Journal of Finance* 51, 1835–1861.
- Harris, Lawrence E., 1996, Does a large minimum price variation encourage order exposure? Unpublished manuscript, University of Southern California.
- Harris, Lawrence E., and Venkatesh Panchapagesan, 1999, The information-content of the limit order book: Evidence from NYSE specialist actions, Unpublished manuscript, University of Southern California.
- Hasbrouck, Joel, 1992, Using the TORQ database, Unpublished manuscript, New York University.
- Hasbrouck, Joel, 1993, Assessing the quality of a security market: A new approach to transaction-cost measurement, *Review of Financial Studies* 6, 191–212.
- Hasbrouck, Joel, and Gideon Saar, 2004, Technology and liquidity provision: The blurring of traditional definitions, Unpublished manuscript, New York University.
- Irvine, Paul, George J. Benston, and Eugene Kandel, 2000, Liquidity beyond the inside spread: Measuring and using information in the limit order book, Unpublished manuscript, Emory University.
- Kaniel, Ron, and Hong Liu, 2001, So what orders do informed traders use? Unpublished manuscript, University of Texas at Austin.
- Lipson, Marc L., 2003, Competition among market centers, Unpublished manuscript, University of Georgia.
- Llorente, Guillermo, Roni Michaely, Gideon Saar, and Jiang Wang, 2002, Dynamic volume–return relation of individual stocks, *Review of Financial Studies* 15, 1005–1047.

- Lo, Andrew W., A. Craig MacKinlay, and June Zhang, 2002, Econometric models of limit-order executions, *Journal of Financial Economics* 65, 31–71.
- Lo, Andrew W., and Jiang Wang, 2000, Trading volume: Definitions, data analysis, and implications of portfolio theory, *Review of Financial Studies* 13, 257–300.
- Lyons, Richard K., 1996, Optimal transparency in a dealer market with an application to foreign exchange, *Journal of Financial Intermediation* 5, 225–254.
- Madhavan, Ananth N., 1992, Trading mechanisms in securities markets, *Journal of Finance* 47, 607–642.
- Madhavan, Ananth N., 1995, Consolidation, fragmentation, and the disclosure of trading information, *Review of Financial Studies* 8, 579–603.
- Madhavan, Ananth N., 1996, Security prices and market transparency, *Journal of Financial Intermediation* 5, 255–283.
- Madhavan, Ananth N., David Porter, and Daniel Weaver, 2000, Should securities markets be transparent? Unpublished manuscript, University of Southern California.
- Merton, Robert C., 1987, A simple model of capital market equilibrium with incomplete information, *Journal of Finance* 42, 483–510.
- Naik, Narayan Y., Anthony Neuberger, and S. Viswanathan, 1994, Disclosure regulation in competitive dealership markets: Analysis of the London Stock Exchange, Unpublished manuscript, London Business School.
- Pagano, Marco, and Ailsa Röell, 1996, Transparency and liquidity: A comparison of auction and dealer markets with informed trading, *Journal of Finance* 51, 579–611.
- Parlour, Christine A., 1998, Price dynamics in limit order markets, *Review of Financial Studies* 11, 789–816.
- Rindi, Barbara, 2002, Transparency, liquidity and price formation, Unpublished manuscript, Bocconi University.
- Sandås, Patrik, 2001, Adverse selection and competitive market making: Empirical evidence from a limit order market, *Review of Financial Studies* 14, 705–734.
- Schwert, G. William, 1981, Using financial data to measure effects of regulation, *Journal of Law and Economics* 24, 121–158.
- Seppi, Duane J., 1997, Liquidity provision with limit orders and a strategic specialist, *Review of Financial Studies* 10, 103–150.
- Theissen, Erik, 2003, Trader anonymity, price formation and liquidity, *European Finance Review* 7, 1–26.
- U.S. Securities and Exchange Commission, 1994, Market 2000: An examination of current equity market developments (U.S. Government Printing Office, Washington, DC).