

Order Consolidation, Price Efficiency, and Extreme Liquidity Shocks

Michael J. Barclay, Terrence Hendershott, and Charles M. Jones*

Abstract

We show that the consolidation of orders is important for producing efficient prices, especially during times of high liquidity demand. The NYSE's centralized opening call market performs better than Nasdaq's decentralized opening process on typical trading days. The NYSE is much better than Nasdaq on witching days when index arbitrage activity subjects S&P 500 stocks to large, predictable, and mostly informationless order flow around quarterly futures contract expirations. Nasdaq opening price efficiency improves to NYSE levels once Nasdaq initiates a consolidated opening call in November 2004, but prices on the decentralized Nasdaq remain less efficient at other times of day.

I. Introduction

The examination of trading under different market structures has greatly expanded our understanding of financial market design. Various characteristics of markets with different structures have been examined, including the efficiency of security prices, competition among liquidity suppliers, the strategies of informed traders, and trading costs. Arguably the most fundamental of these is a market's ability to produce efficient prices.

In this paper, we are particularly interested in a market's ability to absorb order imbalances without unduly moving prices. Such order imbalances can temporarily affect prices in financial markets even if these imbalances reveal little or

*Barclay, formerly of the Simon School of Business, University of Rochester, Box 270100, Rochester, NY 14627, died in a seaplane accident on August 16, 2007; Hendershott, hender@haas.berkeley.edu, Haas School of Business, University of California, Berkeley, 545 Student Services Bldg #1900, Berkeley, CA 94720, and Jones, cj88@columbia.edu, Graduate School of Business, Columbia University, 3022 Broadway, New York, NY 10027. We thank Nasdaq and the NYSE for providing data. We also thank Bill Christie (the referee), Larry Glosten, Frank Hatheway, Dwight Jaffee, Rich Lyons, Tim McCormick, Bill Schwert, Jeff Smith, Bob Whaley, and seminar participants at Columbia University, Emory University, the Nasdaq Economic Advisory Board, the University of Rochester, the University of California, Berkeley, the 2003 Workshop on Information Systems and Economics, the 2004 Western Finance Association Meetings, the 2004 Center for Financial Studies Conference on Market Design, and the 2006 Financial Intermediation Research Society Conference on Banking, Corporate Finance, and Intermediation for helpful comments and suggestions. Hendershott gratefully acknowledges support from the National Science Foundation. Part of this research was conducted while Hendershott and Jones were visiting economists at the NYSE. This paper subsumes an earlier paper titled "Which Witches Better? A Cross-Market Comparison of Extreme Liquidity Shocks."

no information about fundamental values. Because these distortions reduce the efficiency of prices and adversely affect traders in both spot and derivative markets, an important aspect of financial market quality is the ability to minimize these temporary price changes.

The ability of a financial market to absorb liquidity shocks and minimize temporary price changes will be determined by various market characteristics, such as the effectiveness of the competition between liquidity suppliers. Consolidation of orders in one place, either virtually or physically, may also be important (Demsetz (1968)). As decentralized electronic markets increase their share of securities trading,¹ it is natural to ask whether there are advantages to consolidating trading in a single physical location, such as the NYSE trading floor. It is also useful to determine the extent to which these advantages can be replicated in an electronic trading environment.

Unfortunately, it is difficult to assess the link between consolidation and price efficiency using simple cross-market comparisons because it is difficult to control for other important factors that vary across markets, including differences in market conditions and security characteristics. We overcome these difficulties in two ways. First, we identify large, informationless liquidity shocks that strike certain NYSE and Nasdaq stocks at the same time (at the open). Because the NYSE has a centralized opening call market and Nasdaq has a decentralized opening procedure, we use these common shocks to compare the resulting efficiency of prices in centralized and decentralized markets. Second, we study a significant change in the Nasdaq opening mechanism that results in much more order consolidation. In November 2004, Nasdaq changed its opening procedure by instituting its own consolidated call auction. We are able to gain further insights about the value of the consolidation by comparing the efficiency of Nasdaq's prices—both at the open and at other times—before and after this change.

Opening a financial market creates unusual stress because the opening price must be determined with little or no trading at a time when uncertainty about fundamental values is high. This stress is compounded when there are large order imbalances at the open, even if these order imbalances are unrelated to changes in fundamental values. The expiration of index futures contracts provides a sample of large order imbalances caused by largely informationless liquidity shocks. In this paper, we examine the abilities of Nasdaq and the NYSE to produce efficient prices around the quarterly expiration of the Standard & Poor's (S&P) 500 futures contract.

Expiration days for the S&P 500 futures contract are often called triple-witching or quad-witching days because the futures contract, equity index options, individual equity options, and more recently single-stock futures all expire or settle on the same day. S&P 500 futures and options settle based on the opening prices of the constituent stocks. As a result of this settlement procedure, witching days generate large liquidity shocks at the open that are generally attributed to the unwinding of index arbitrage positions.² Arbitrage strategies dictate that initial

¹ See Barclay, Hendershott, and McCormick (2003), Rime (2003), and Barclay, Hendershott, and Kotz (2006) for evidence in the stock, foreign exchange, and government bond markets, respectively.

² Index arbitrage activity is reported to the NYSE because Rule 80A suspends certain types of index arbitrage program trading on volatile trading days. NYSE audit trail data (CAUD) from 2003

long or short cash market trades must be reversed on the expiration day to close out the arbitrage position. Therefore, arbitrageurs are likely to submit large buy or sell orders in many stocks at the open on the expiration day. A temporary mismatch between these orders could send the price up or down in response to the imbalance between supply and demand.

The S&P 500 index contains both NYSE- and Nasdaq-listed stocks. Thus, witching days generate extreme liquidity shocks in both markets. We use these common shocks to determine which market is better able to absorb large liquidity shocks and produce more efficient and less volatile prices. Because witching days produce liquidity shocks in both markets, we are able to use a differences-in-differences approach, where each stock serves as its own control.³ For each stock, we first calculate the differences in price efficiency, temporary volatility, and other market characteristics between witching and non-witching days. We then compare these witching day differences across Nasdaq and the NYSE. By focusing on the differences between witching and non-witching days for each stock, we control for most of the normal cross-sectional variation in trading characteristics. In a regression framework, we also control directly for differences in observable trading characteristics, such as trading volume and volatility.

The NYSE's opening call performs better than Nasdaq on non-witching days and much better than Nasdaq on witching days, which indicates that there are substantial benefits from consolidating trade, particularly during times of extreme market stress. The NYSE's advantage in opening price efficiency disappeared, however, when Nasdaq initiated its own consolidated opening call auction in November 2004. This suggests that consolidating order flow is more important to the trading process than the physical proximity of traders because consolidating orders across Nasdaq's dispersed electronic trading platforms duplicates most of the benefits of the NYSE trading floor.

After the opening each day, Nasdaq returns to a market structure that is much more decentralized than the NYSE. While Nasdaq's opening prices are now on par with the NYSE's, we find that Nasdaq prices are still less efficient during the rest of the trading day. This suggests that Nasdaq's lack of consolidation after 9:30 AM hurts the efficiency of its prices during the day.

The remainder of the paper is organized as follows. Section II reviews related literature and discusses opening procedures on the two markets. Section III provides a general description of our data and sample. Section IV documents the

through Q3 2005 for S&P 500 stocks shows that the opening trade is more than 10 times its normal size on witching days, averaging \$15.05 million per stock on witching days versus \$1.45 million per stock on other days. Index arbitrage activity at the open (buys plus sells) is 50 times its usual level on witching days, averaging \$12.74 million per stock on witching days versus \$0.24 million on other days. The average opening order imbalance caused by index arbitrage trades (the absolute difference between index arbitrage buys and sells) is \$3.21 million per stock on witching days versus \$0.14 million per stock on other days. Thus, the average witching day opening imbalance caused by index arbitrage trades is twice as large as the total normal opening trade. It is possible that the remaining increase in opening volume on witching days is due to other traders seeking to trade at a time of predictable high trading volume or the result of improper reporting of index arbitrage trades. Because Rule 80A has not been in effect at the open, members may be less careful in coding opening index arbitrage activity.

³Using other stocks as controls is difficult because virtually all actively traded stocks are part of some futures contract, e.g., the Russell 1000 and 2000, expiring at the same time and settling in the same manner as the S&P 500.

higher volatility on witching days. Section V investigates the relative efficiency of prices around the open. Section VI examines prices at the open and at other times following Nasdaq's adoption of an opening call auction in November 2004. Section VII concludes the paper.

II. Related Literature

The potential benefits of consolidation have been studied empirically and theoretically with a primary focus on traders' choices of where or when to trade. Admati and Pfleiderer (1988), Foster and Viswanathan (1990), and Barclay and Hendershott (2004) examine how traders can benefit by consolidating their trading in time. Garbade and Silber (1979), Mendelson (1987), Pagano (1989), Amihud, Lauterbach, and Mendelson (2003), and Hendershott and Jones (2005) analyze how liquidity externalities can also arise when traders consolidate across markets and securities. Of the papers on consolidation across markets, only Hendershott and Jones (2005) focus on price efficiency, and their results arise from a regulatory shock to consolidation, not from a comparison of different market structures.

Because we compare the performance of the NYSE opening call market with the decentralized continuous trading at the Nasdaq opening, our results are related to the literature on comparative market structures and price efficiency.⁴ For example, Madhavan (1992) predicts that a periodic call auction yields more efficient prices than a continuous market mechanism. But, Amihud and Mendelson (1987) find greater volatility at the NYSE call market opening than at the NYSE continuous market close.⁵ Pagano and Schwartz (2003) find that the introduction of a closing call in Paris improved price efficiency and Pagano and Schwartz (2005) and Smith (2005) find that the introduction of Nasdaq's closing call also improved price efficiency. Finally, Ellul, Shin, and Tonks (2005) find that the opening call on the London Stock Exchange provides more efficient prices than the off-exchange dealer system. Our results support the view that call markets are better able to absorb extreme liquidity shocks without producing inefficient prices and excess volatility.

A financial exchange's opening mechanism is important because it is used to determine prices when uncertainty about fundamental values is particularly high after an extended non-trading period. Different exchanges have developed different mechanisms to discover the opening price.⁶ Beginning at 9:30 AM, the NYSE

⁴Our results are also related to existing cross-market comparisons of trading costs. Huang and Stoll (1996) and Bessembinder and Kaufman (1997) compare trading costs on the NYSE and Nasdaq in a regression framework. Boehmer (2005) uses recent Rule 11ac1-5 data to compare the two trading venues, while Barclay (1997), Jones and Lipson (1999), and Bennett and Wei (2006) examine the trading costs of firms that switch from Nasdaq to the NYSE. These studies generally find higher trading costs on Nasdaq than on the NYSE. Bessembinder and Rath (2008) examine changes in volatility and trading costs for firms switching from Nasdaq to the NYSE to see if lower trading costs lead to lower volatility and find that they do not. Venkataraman (2001) compares trading costs on the NYSE to the Paris Bourse.

⁵However, Amihud and Mendelson (1991) conclude that the greater volatility at the open is not due to the call mechanism, but instead is due to the greater difficulty of discovering the opening price after 17.5 hours without trading overnight.

⁶See Domowitz and Madhavan (2000) for a survey of these mechanisms.

operates an opening call auction in each listed stock. Prior to this opening auction, investors can submit market and limit orders electronically, along with various other less common order types. Floor traders also indicate their trading interest to the specialist. The specialist consolidates these orders and sets an opening price, absorbing any order imbalance that may remain. Only the specialist and floor traders can observe all of the submitted orders.⁷ Most other market participants off the floor observe only the opening price and the aggregate opening volume in each stock.⁸

Stoll and Whaley (1990b) argue that the NYSE opening process generates noise in the opening price and show that the opening price is significantly noisier than the closing price. In contrast, Madhavan and Panchapagesan (2000) argue that the ability to observe the evolution of the limit order book conveys valuable information to the specialist. Using order-level data, they show that trading by the NYSE specialist produces a more efficient opening price. However, because other market participants condition their actions on the existence of the specialist, it is not possible to know whether NYSE opening prices would be more or less efficient without the specialist.

Before Nasdaq instituted an opening cross in November 2004, its opening process was highly decentralized. Actively traded Nasdaq stocks can have close to 100 market makers that execute trades on multiple trading venues with multiple limit order books.⁹ Unlike the opening call auction on the NYSE, the official opening price for Nasdaq stocks was simply the price of the first trade reported at or after 9:30 AM, regardless of where that trade occurred.

Market makers post quotes before the open, but these quotes are not binding and often result in a locked or crossed market.¹⁰ Cao, Ghysels, and Hatheway (2000) examine the Nasdaq pre-open in 1994 and argue that market makers communicated with each other through non-binding quotes in an attempt to ensure that the opening bid and ask quotes reflected their available information. Today, many Nasdaq stocks begin trading on electronic communications networks (ECNs) before the official 9:30 AM opening. Pre-open trading volume in the most active Nasdaq stocks has increased from several thousand dollars per stock per day in the early 1990s to more than \$2 million per stock per day in 2000. These pre-open trades are informative and improve the efficiency of the opening price (Barclay and Hendershott (2003), (2007)). Nevertheless, because trading costs

⁷Beginning in 2002, subscribers to NYSE's Openbook product can observe each stock's aggregate system limit order book before the open, but they cannot observe system market orders or orders represented on the floor.

⁸The procedure is different if there is a severe order imbalance at the open. In that case, the specialist may disseminate a non-binding, indicative range of opening prices in an attempt to convey information about the order imbalance and solicit stabilizing orders.

⁹In October 2002, Nasdaq attempted to consolidate the various limit order books in its SuperMontage trading system. However, SuperMontage is a voluntary system and several large market participants, including the largest electronic communications networks (ECNs), decided not to participate.

¹⁰On June 5, 2000, Nasdaq instituted a "trade-or-move" rule to reduce the likelihood that the market was locked or crossed at the open. This rule was amended in October 2002 in conjunction with the new SuperMontage opening procedure that resolved locked or crossed markets by matching and executing the locked or crossed quotes and orders at 9:29:30 AM.

are high in the pre-open and volume is low, prices tend to be less efficient during the pre-open than during the trading day.¹¹

Opening a market in the morning is somewhat similar to reopening a market after a crash or trading halt, which has been studied by Lee, Ready, and Seguin (1994), Corwin and Lipson (2000), Christie, Corwin, and Harris (2002), and Goldstein and Kavajecz (2004), among others. However, order imbalances after a trading halt are typically the result of acute shocks to fundamentals or severe information asymmetry. In contrast, we are interested in how market structures handle the large, predictable, but largely informationless liquidity shocks submitted by index arbitrageurs at the open on witching days.

Regulators often express concerns about high volatility in stock markets due to index derivative expirations.¹² To this end, there have been a number of regulatory changes in an attempt to mitigate abnormal stock price movements on expiration days. Most exchanges settle the stock-index futures contracts using opening, closing, or average trade prices of the stocks in the index. For example, the Chicago Mercantile Exchange and the New York Futures Exchange moved the settlement of S&P 500 and NYSE index derivatives from the close to the open of trading in June 1987. The opening quotation for the index (and, thus, the settlement price for the contract) is calculated from the opening price of each component stock.

Several studies examine return volatility on so-called "triple-witching" days.¹³ Stoll and Whaley (1987) investigate the effects of large transactions on prices and find significant price volatility on expiration days. Stoll and Whaley (1990a) find evidence of downward price pressure on expiration days. Stoll and Whaley (1991) and Hancock (1993) find that after settlement is moved to the Friday opening price in 1987, the expiration effects move to the Friday opening for NYSE stocks and the S&P 500 cash index. Sofianos (1994) finds minimal effects of expiration-related trading on NYSE stocks' opening times and price reversals.

Figure 1 graphs the daily open interest in the next-to-expire pit-traded and e-mini S&P 500 futures contracts from 1998 through the third quarter of 2005. The aggregate open interest in these contracts increased steadily from \$100 billion in 1998 to over \$300 billion in 2005. Open interest in the e-mini contracts is less than 40% of the open interest in pit-traded contracts, but this fraction has increased steadily over time. Open interest in the S&P 500 futures contracts is much larger than in any other contract. For example, the 2005 open interest is

¹¹Quoted and effective bid-ask spreads are three to four times larger during the pre-open than during the trading day and there are only 1/20th as many trades per minute during the pre-open as there are during the trading day (Barclay and Hendershott (2004)).

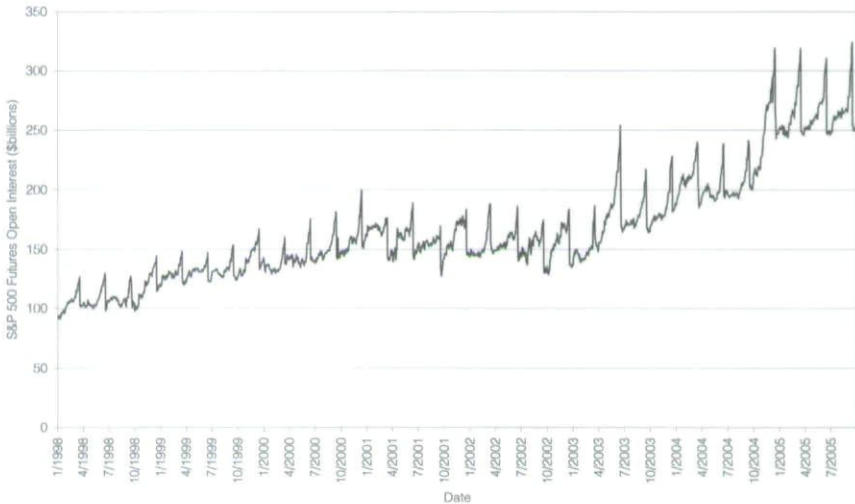
¹²The expiration day effects of stock-index derivative contracts are typically attributed to the unwinding of hedged positions in the cash and derivative markets. Because the derivative settles in cash, arbitrageurs must close out the cash market position at or close to the settlement price to realize the arbitrage profit. At expiration, arbitrageurs submit large buy or sell orders that can create a potential mismatch between supply and demand and generate large price changes that are subsequently reversed. See Roll, Schwartz, and Subrahmanyam (2007) for a general discussion of the relation between liquidity and the cash/futures basis.

¹³On July 24, 2002, the SEC approved single-stock futures, and they began trading later that year. These along with single-stock options, index futures, and index options expire on the same day, and this expiration day is now referred to as quadruple-witching day. We refer to these days as witching days for simplicity.

less than \$10 billion in the Russell 2000 futures contract and less than \$20 billion in the Nasdaq 100 futures contract.

FIGURE 1
Open Interest on S&P 500 Futures

Daily open interest in billions of dollars for all S&P 500 futures contracts from January 1, 1998 to September 30, 2005.



Within each calendar quarter, open interest in the S&P 500 futures contracts increases until the expiration date when, on average, about 20% of the pit-traded contracts and 50% of the e-minis are settled.¹⁴ From 1998 to 2002, \$34 billion of S&P 500 futures contracts settled on an average witching day. Since late 2004 between \$50 and \$75 billion of S&P 500 futures contracts settled on each witching day.

III. Data and Descriptive Statistics

Several data sets are used for our analysis. Closing stock prices, stock splits and distributions, and most descriptive statistics are obtained from CRSP through 2004 and from the TAQ master file for NYSE stocks in 2005. Yahoo Finance and TAQ provide stock splits and distributions for Nasdaq and NYSE stocks, respectively, in 2005. Opening, intraday, and pre-open stock prices are obtained from TAQ. The internal CT file at the NYSE is used in place of the TAQ data in 2005 and also for obtaining official opening prices. Prior to the existence of an official Nasdaq opening price, we use the price of the first non "T" trade after 9:30 AM (as in Barclay and Hendershott (2003), (2004), (2007)).

¹⁴During the last weeks before each quarterly expiration of the S&P 500 futures, the bulk of trading volume begins to shift away from the next-to-expire (nearby or lead) contract toward the second-to-expire (next out) contract. At some point, the CME formally redesignates the next out as the new lead contract, and the next out replaces the nearby in the futures pit location designated for the lead contract.

Table 1 describes the S&P 500 stocks traded on the NYSE (Panel A) and Nasdaq (Panel B) for each year from 1998 through the third quarter of 2005. The number of NYSE-listed S&P 500 stocks declines from 459 in 1998 to 424 in 2001 and remains roughly constant thereafter. The number of Nasdaq-listed stocks increases from 38 to 74 over this same period.¹⁵ The Nasdaq-listed S&P 500 stocks have a larger average market capitalization and a higher average share price from 1998 to 2000. But the Nasdaq and NYSE stocks are similar in market capitalization from 2001 through 2005, and the NYSE stocks have higher average prices during this later period. Daily stock return volatility is significantly higher for Nasdaq stocks than for NYSE stocks during the entire sample period.

TABLE 1
Descriptive Statistics for S&P 500 Stocks by Market and Year

For each year from 1998 to Q3 2005, the daily average market capitalization, price, daily return standard deviation, dollar trading volume, and number of trades is calculated for each S&P 500 stock traded on the NYSE (Panel A) or Nasdaq (Panel B).

Panel A. NYSE Stocks

Year	No. of S&P 500 Stocks	Market Cap (\$billions)	Share Price (\$)	Daily Return Std. Dev. (%)	Trading Day		Open Volume (\$millions)
					Volume (\$millions)	No. of Trades	
1998	459	16.4	50.01	2.40	52.0	599	1.57
1999	458	19.8	46.24	2.54	71.0	840	2.17
2000	439	21.5	41.66	3.11	87.6	1,068	2.66
2001	424	21.1	39.58	2.57	86.1	1,301	2.03
2002	421	18.3	36.36	2.85	83.4	1,729	1.69
2003	424	17.6	35.40	1.94	77.0	2,266	1.44
2004	424	20.8	41.63	1.51	87.8	2,598	1.56
2005	425	24.6	44.30	1.47	99.2	3,286	1.92

Panel B. Nasdaq Stocks

Year	No. of S&P 500 Stocks	Market Cap (\$billions)	Share Price (\$)	Daily Return Std. Dev. (%)	Trading Day		Pre-Open	
					Volume (\$millions)	No. of Trades	Volume (\$millions)	No. of Trades
1998	38	23.4	45.79	3.32	247.0	3,824	1.4	22
1999	41	39.8	55.28	3.57	378.2	6,199	2.1	39
2000	60	44.0	56.32	5.26	530.7	11,336	2.9	88
2001	74	22.1	32.37	4.70	302.7	12,976	1.4	96
2002	77	16.5	25.54	4.17	218.8	13,705	0.9	81
2003	74	18.4	28.70	2.56	202.9	14,777	0.9	88
2004	75	21.5	33.16	2.16	220.0	16,244	0.9	88
2005	74	21.4	34.30	1.91	213.4	18,489	0.8	89

Reported trading volume is higher for Nasdaq stocks than for NYSE stocks, although it is difficult to interpret cross-market comparisons of trading volume because the two markets intermediate trades differently. Because Nasdaq trades are smaller, on average, than NYSE trades, the difference in the number of trades is greater than the difference in dollar trading volume.¹⁶ While there is little observable trading of NYSE stocks in the U.S. before the NYSE opening, pre-

¹⁵The number of Amex-listed stocks in the S&P 500 declines from five to one during our sample period.

¹⁶The number of trades is calculated from TAQ. Trade accounting is difficult on the NYSE because multiple parties may transact at one price. The number of official trades determines tape revenues and the NYSE's overall official number of trades reported to the Consolidated Tape Association is approximately twice the number in TAQ.

open trading volume on ECNs for Nasdaq stocks averages between \$1 million and \$3 million per stock per day.¹⁷

The criteria for inclusion in the S&P 500 ensures that all index stocks are large and actively traded. Table 1 shows, however, that there are systematic differences in the average trading volume, stock return volatility, share price, and market capitalization across the two markets. These cross-market differences motivate our use of an exogenous event—witching days—to examine the effect of liquidity shocks on markets with substantially different structures. The differences-in-differences approach (comparing witching and non-witching days for each stock) controls for the cross-sectional differences in stock characteristics.¹⁸ In Section VI, we supplement this analysis by examining the introduction of Nasdaq's opening cross.

A. Trading Volume and Witching Days

Our analysis begins by establishing that there is a significant shock to trading volume on witching days. We first explore the magnitude and duration of that shock before measuring its effects on volatility and price efficiency.

To measure the abnormal trading volume on witching days, we first calculate the cross-sectional average dollar trading volume in millions of dollars (vol_t^i) for all NYSE or Nasdaq S&P 500 stocks on each day (t) for various intraday time periods (i). Time periods include: the trading day prior to the witching day and the final 30 minutes of trading on the day prior to the witching day; the opening trade for NYSE stocks and the 90 minutes before the open for Nasdaq stocks on the witching day; the first 30 minutes of trading on the witching day; and from 10 AM to the close on the witching day. Because trading volume is serially correlated, we control for lagged trading volume. We also allow the persistence of trading volume to differ on witching and non-witching days because the expiration effect is not persistent. Thus, for each exchange, we regress the cross-sectional average dollar trading volume in period i (vol_t^i) on the average dollar trading volume on each of the previous five trading days (vol_{t-j}), a witching day dummy variable ($Witch_t$), and lagged witching dummies interacted with lagged trading volume:

$$(1) \quad vol_t^i = \alpha + \sum_{j=1}^5 \beta_j vol_{t-j} + \omega Witch_t + \sum_{j=1}^5 \phi_j Witch_{t-j} vol_{t-j} + \varepsilon_t^i.$$

The regression is estimated separately for the NYSE and Nasdaq, which yields separate coefficients for each market. For each market, we report the witching day dummy variable coefficients in Table 2 with Newey-West corrected standard errors for statistical inference. Because there is virtually no NYSE trading

¹⁷Prior to July 12, 2002, TAQ does not capture off-exchange trades in NYSE stocks prior to 9:30 AM. Proprietary data from the National Association of Securities Dealers shows essentially no pre-open trading in NYSE stocks in 1999 and 2000 (Barclay and Hendershott (2003), (2004), (2007)). Beginning on July 12, 2002, trading on the Archipelago ECN appears for NYSE-listed stocks prior to 9:30 AM, but these pre-open trades are much less frequent than in Nasdaq stocks. Starting in September 2005, Nasdaq requires all members to report pre-open trading in NYSE stocks in a timely manner.

¹⁸Another possible control could be non-S&P 500 stocks, but the Russell 1000 and Russell 2000 futures settle at the same time and in the same manner as the S&P 500 futures.

before the open, and because the Nasdaq open represents only one individual trade, we list the NYSE opening trade with Nasdaq pre-open trades.

TABLE 2
Trading Volume Surrounding Witching Days

Data is for 1998 to Q3 2004. For each day (t) and market, the cross-sectional average dollar trading volume (in millions of dollars) for time period i (vol_t^i) is regressed on the average dollar trading volume for the previous five days (vol_{t-j}), a witching day dummy variable ($Witch_t$), and lagged witching dummies interacted with lagged trading volume:

$$vol_t^i = \alpha + \sum_{j=1}^5 \beta_j vol_{t-j} + \omega Witch_t + \sum_{j=1}^5 \phi_j Witch_{t-j} vol_{t-j} + \varepsilon_t^i$$

Time periods include: the trading day prior to the witching day and the final 30 minutes of trading on the day prior to the witching day; the opening trade for NYSE stocks and the 90 minutes before the open for Nasdaq stocks on the witching day; the first 30 minutes of trading on the witching day; and from 10 AM to the close on the witching day. The regression is estimated for each market and time period. The witching day coefficients are given with Newey-West corrected standard errors in parentheses.

Period	NYSE	Nasdaq
Trading Day	1.30	-4.85
Day Prior to Witching Day	(1.90)	(10.01)
Last 30 Minutes of Day Prior to Witching Day	0.49 (3.09)	0.41 (1.47)
Open/Pre-Open	12.68	4.51
Witching Day	(0.22)	(0.38)
Open to 10 AM	15.08	18.62
Witching Day	(0.34)	(1.70)
10 AM to 4 PM	0.57	-7.51
Witching Day	(1.75)	(8.98)

Table 2 presents the average excess witching day trading volume at the open (or pre-open) and for the first half hour of the trading day. The opening trade on the NYSE is approximately seven times larger on witching days than on non-witching days, reflecting an additional \$12.7 million of witching day trading volume per stock at the open. NYSE trading volume is also higher from the open to 10:00 AM on witching days, but most of this increase is due to the larger opening trade. For Nasdaq stocks, pre-open trading volume is nearly five times its normal level on witching days. The abnormal witching day trading volume for the average Nasdaq stock is \$4.5 million in the pre-open and \$18.6 million from the open to 10:00 AM.

Table 2 also shows that the abnormal trading volume is confined to a short period around the open on witching days. There is no significant increase in trading volume on the previous trading day, during the last 30 minutes of trading on the previous trading day, or between 10:00 AM and the close of trade on witching days.


Figure 2 graphs the excess witching day trading volume by minute from 9:00 AM to 9:45 AM. For each witching day, the cross-sectional average dollar trading volume in each minute is divided by the average dollar trading volume in the same minute for all trading days in the quarter. The average of these quarterly ratios is graphed separately for the NYSE and Nasdaq. For the NYSE, the excess trading volume in the opening trade is also reported.

For Nasdaq stocks, trading volume at 9:00 AM is about 50% higher on witching days than on other days. Excess trading volume increases steadily during the

FIGURE 2

Relative Witching Day Trading Volume by Minute from 9:00 AM to 9:45 AM

For each quarter from 1998 to Q3 2004 and for each market, the witching day dollar trading volume is divided by the average daily dollar trading volume in the same quarter and market. The average of these quarterly ratios is graphed by minute.



pre-open, reaching three times the average level by 9:23 AM and more than five times the average level by 9:30 AM. Because the majority of pre-open trading occurs shortly before the open (Barclay and Hendershott (2003), (2004)), total pre-open trading is almost five times higher than average on witching days. After the open, the abnormal trading volume on witching days declines rapidly and returns to normal levels by 9:45 AM.

Excess trading volume on the NYSE follows a similar pattern on witching days. Although there is little trading in NYSE-listed stocks before the open, the NYSE opening trade is 7.4 times larger than normal on witching days. After the open, NYSE trading volume returns to normal levels by 9:45 AM.

Table 2 and Figure 2 document that both NYSE and Nasdaq stocks in the S&P 500 experience significant excess trading volume around the open on witching days. Trading volume is normal on both the day before the witching day and after 9:45 AM on the witching day. Thus, the opening on witching days provides an exogenous, predictable, and short-lived liquidity shock that is common to both markets and allows us to compare how different market structures handle these extreme liquidity events.

IV. Volatility on Witching Days

In this section, we measure the impact of excess witching day trading volume on stock return volatility. Because the excess trading volume is concentrated

at or around the open, we focus our attention on volatility around the open. Because stock return volatility is serially correlated, we control for lagged volatility using computationally tractable absolute return autoregressions (see, for example, Schwert and Seguin (1990)). Similar to the volume regressions, for each day t and market m (NYSE or Nasdaq), we calculate the cross-sectional average absolute return from time i to time j ($|ret_{m,t}^{i,j}|$). We regress these average absolute returns on the absolute close-to-close return in the same market for the previous five days and on dummy variables set equal to one for Nasdaq stocks ($Nasdaq_{m,t}$), witching days ($Witch_t$), and Nasdaq stocks on witching days ($Nasdaq_{m,t}Witch_t$):¹⁹

$$(2) \quad |ret_{m,t}^{i,j}| = \alpha + \sum_{j=1}^5 \delta_j |ret_{m,t-j}^{cc}| + \beta_1 Witch_t + \beta_2 Nasdaq_{m,t} + \beta_3 Nasdaq_{m,t}Witch_t + \varepsilon_{m,t}^{i,j}.$$

In Table 3, we refer to these regressions as the "time-series" regressions because they are estimated with one observation per market per day.

The time-series regressions in Table 3 show that the average absolute close-to-open return is 75 basis points for NYSE stocks. The close-to-open return is about 50% more volatile for Nasdaq stocks, as indicated by the positive and significant coefficient on the Nasdaq dummy variable of 36 basis points. Nasdaq's higher close-to-open volatility on non-witching days could be caused by a variety of factors, such as differences in the characteristics of the stocks traded in each market (which we control for more directly below) or differences in the trading mechanism itself. Thus, although the higher opening volatility on Nasdaq is suggestive, we do not interpret it as a direct test of the relative performance of centralized and decentralized opening procedures.

Having controlled for differences in volatility across markets on non-witching days, we can use the excess opening volatility on witching days to determine whether a centralized or decentralized opening procedure is better able to absorb extreme liquidity shocks. Table 3 indicates that on both Nasdaq and the NYSE, the close-to-open volatility is higher on witching days than on non-witching days. However, the excess volatility on witching days is significantly greater on Nasdaq. Close-to-open volatility on witching days is 28 basis points higher than normal on the NYSE (as indicated by the coefficient on the $Witch_t$ dummy variable) and 64 basis points higher than normal on Nasdaq (as indicated by the sum of the coefficients on the $Witch_t$ and $Nasdaq_{m,t}Witch_t$ dummy variables). The difference in witching day volatility between Nasdaq and the NYSE is statistically significant and suggests that the centralized opening procedure on the NYSE is better able to handle the large liquidity shocks on witching days than the decentralized procedure on Nasdaq.

While the time-series regressions in Table 3 control for differences in the average volatility between Nasdaq and NYSE stocks, they do not control for firm characteristics that could affect excess volatility on witching days. Thus, in "pooled" regressions (one observation per stock per day) we control for the following stock characteristics: daily return standard deviation (σ), log daily dollar

¹⁹Amihud and Mendelson (1987) and others find evidence of excess volatility at the close. However, our results are qualitatively unchanged using the 3:00 PM price instead of the closing price.

TABLE 3
Volatility

Data is for 1998 to Q3 2004. For the time-series regressions, for each market each day the cross-sectional average absolute return (in percent) from the close to time period i ($|ret_{s,t}^{c,i}|$) is calculated. These absolute returns are regressed on a constant and dummy variables set equal to one if it is Nasdaq (*Nasdaq*), one if the day is a witching day (*Witch*), and the Nasdaq and witching dummies interacted (*Witch * Nasdaq*). In the time-series regressions, standard errors are corrected using the Newey-West procedure. For the pooled regressions, absolute returns for each stock are regressed on stock-specific control variables calculated over the previous month: daily return standard deviation (σ), log daily dollar trading volume ($\log(\$vol)$), and log market capitalization ($\log(mkt\ cap)$) expressed as deviations from the overall mean. These control variables are also interacted with the witching dummy. Absolute close-to-close returns for the previous five days are also included in the regression, but not reported. For the pooled regressions, the standard errors correct for contemporaneous correlation and heteroskedasticity.

	Close to Open		Open to 10 AM		Close to 10 AM	
	Time Series	Pooled	Time Series	Pooled	Time Series	Pooled
Constant	0.75 (0.01)	0.38 (0.01)	0.72 (0.01)	0.43 (0.01)	1.09 (0.01)	0.59 (0.01)
<i>Nasdaq</i>	0.36 (0.01)	0.13 (0.01)	0.43 (0.01)	0.19 (0.01)	0.47 (0.01)	0.10 (0.01)
<i>Witch</i>	0.28 (0.11)	0.22 (0.11)	0.01 (0.07)	0.00 (0.04)	0.01 (0.11)	-0.04 (0.08)
<i>Witch * Nasdaq</i>	0.36 (0.12)	0.57 (0.12)	0.21 (0.07)	0.29 (0.06)	-0.01 (0.11)	0.11 (0.07)
σ		0.06 (0.00)		0.01 (0.00)		0.08 (0.00)
$\log(\$vol)$		-0.03 (0.01)		0.10 (0.01)		0.02 (0.01)
$\log(mkt\ cap)$		-0.03 (0.01)		-0.13 (0.01)		-0.11 (0.01)
<i>Witch * σ</i>		0.08 (0.05)		0.01 (0.01)		0.00 (0.03)
<i>Witch * $\log(\\$vol)$</i>		-0.24 (0.08)		-0.04 (0.02)		-0.12 (0.06)
<i>Witch * $\log(mkt\ cap)$</i>		0.16 (0.07)		0.02 (0.02)		0.10 (0.05)
Adj. R^2	0.26	0.03	0.29	0.07	0.31	0.05
Obs.	3,442	849,182	3,442	849,182	3,442	849,182

trading volume ($\log(\$vol)$), and log market capitalization ($\log(mkt\ cap)$). Stock characteristics are calculated for the prior calendar month and expressed as deviations from the mean. To allow these control variables to have different effects on witching and non-witching days, we also interact the stock characteristics with the witching day dummy variable. Defining these control variables for stock s at date t as the vector $X_{s,t}$, we estimate the following regression:²⁰

$$(3) \quad \left| ret_{s,t}^{i,j} \right| = \alpha + \sum_{j=1}^5 \delta_j \left| ret_{s,t-j}^{c,c} \right| + \gamma' X_{s,t} + (\phi' X_{s,t}) Witch_t + \beta_1 Witch_t + \beta_2 Nasdaq_{s,t} + \beta_3 Nasdaq_{s,t} Witch_t + \varepsilon_{s,t}^{i,j}$$

The pooled volatility regressions are presented in Table 3 next to the comparable time-series regressions. For overnight volatility, including firm-specific controls has little effect on the coefficient for the *Witch_t* dummy variable (22

²⁰We follow Rogers (1993) and calculate generalized White (1980) standard errors that allow for both heteroskedasticity and contemporaneous correlation in the error terms of this regression. See Petersen (2008) for further discussion of the Rogers standard errors.

basis points with controls compared with 28 basis points without firm-specific controls). However, if anything, the difference in excess witching day volatility between Nasdaq and the NYSE is greater after including the firm-specific controls. The coefficient on the *NasdaqWitch* dummy variable increases from 36 basis points in the time-series regression to 57 basis points in the pooled regression.

On Nasdaq, the excess close-to-open volatility on witching days persists into the trading day. When we estimate the excess volatility regressions from the open to 10:00 AM, we find that average absolute returns on Nasdaq stocks are 21 to 29 basis points higher on witching days, depending on the specification. In contrast, NYSE stocks are not more volatile from the open to 10:00 AM on witching days. On both exchanges, however, the excess volatility on witching days is temporary. Both the pooled and time-series regressions in Table 3 show that there is no excess volatility from the close to 10:00 AM on witching days. Similar regressions (not reported) also show that there is no excess volatility on the day before the witching day. Together, these regressions support our earlier claim that witching days generate essentially informationless liquidity shocks. Nevertheless, the excess witching day volatility is quite high at the open, especially on Nasdaq, which has important consequences for the futures contracts that settle at these temporarily distorted prices.

Figure 3 shows the excess witching day volatility minute by minute. To construct Figure 3, we estimate equation (2) using the return from the previous close to minute i for each minute from 9:00 AM to 10:00 AM for Nasdaq stocks and from the open to 10:00 AM for NYSE stocks. The figure shows that the excess witching day volatility for Nasdaq stocks increases steadily during the pre-open, reaching a maximum of about 64 basis points at the open. After the open, excess volatility falls sharply and, by 9:50 AM, witching days are no more volatile than non-witching days for Nasdaq stocks. Excess witching day volatility for NYSE stocks also reaches its maximum level of nearly 28 basis points at the open and falls sharply thereafter. By 9:50 AM witching days are also no more volatile than non-witching days for NYSE stocks.

V. Price Efficiency on Witching Days

The analysis in the previous section indicates that the liquidity shock at the open on witching days generates a transitory shock to stock return volatility. This volatility shock is no longer detectable after 10:00 AM on a witching day, which implies that the higher witching day volatility at the open is due to noise in opening prices that is subsequently reversed. To see this, note that if returns are continuously compounded, then for times c , i , and j with c corresponding to the close and $c \leq i \leq j$, the return from the close to time j can be written as:

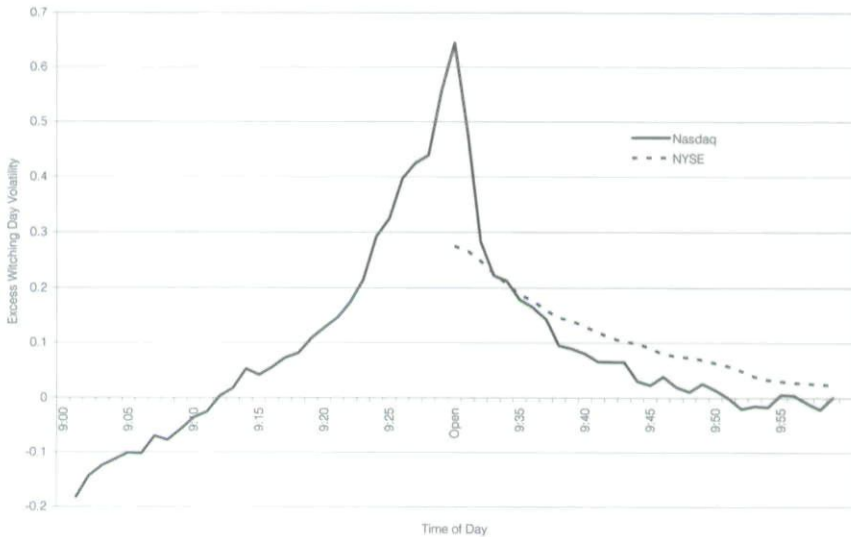
$$(4) \quad \text{ret}_{s,t}^{c,j} = \text{ret}_{s,t}^{c,i} + \text{ret}_{s,t}^{i,j}$$

and the variance of this sum can be written as:

$$(5) \quad \text{var} \left(\text{ret}_{s,t}^{c,j} \right) = \text{var} \left(\text{ret}_{s,t}^{c,i} \right) + \text{var} \left(\text{ret}_{s,t}^{i,j} \right) + 2 \text{cov} \left(\text{ret}_{s,t}^{c,i}, \text{ret}_{s,t}^{i,j} \right).$$

FIGURE 3
Excess Witching Day Volatility

From 1998 to Q3 2004, the average absolute return from the close to minute i for stocks in each market (m) is regressed on an intercept, a witching day dummy, and the witching dummy interacted with a Nasdaq dummy. The witching day dummy (NYSE) is graphed from the open to 10 AM and the witching dummy interacted with a Nasdaq dummy (Nasdaq) is graphed from 9 AM to 10 AM. Returns at time i are only calculated for those stocks that have opened by that time.



Let i and j correspond to the open and 10:00 AM, respectively. The evidence in the previous section indicates that overnight volatility (the first term on the right-hand side) rises on witching days. The second term (volatility from the open to 10:00 AM) also rises for Nasdaq stocks or stays the same for NYSE stocks on witching days. But the volatility of returns from the close to 10:00 AM (the term on the left-hand side) is not higher on witching days. This implies that the covariance term is more negative on witching days and that witching day opening price changes are more strongly reversed.

We use the return covariance to measure the extent of the reversal of overnight stock returns. For each stock s on day t , the reversal from time i to time j ($rev_{s,t}^{i,j}$) is defined as the product of the percent return from the close to time i ($ret_{s,t}^{c,i}$) and the percent return from time i to time j ($ret_{s,t}^{i,j}$).²¹ For each day, we average the reversals across stocks in each market m and regress them on dummy variables set equal to one for Nasdaq stocks ($Nasdaq_{m,t}$), witching days ($Witch_t$), and Nasdaq stocks on witching days ($Witch_t Nasdaq_{m,t}$):

$$(6) \quad rev_{m,t}^{i,j} = \alpha + \beta_1 Witch_t + \beta_2 Nasdaq_{m,t} + \beta_3 Nasdaq_{m,t} Witch_t + \varepsilon_{m,t}^{i,j}$$

Table 4 provides estimates of the magnitude of the reversals from the open to 10:00 AM and from 10:00 AM to 11:00 AM. Price reversals would be expected on

²¹ Our covariance measure of price reversals is proportional to Roll's measure of the bid-ask spread (Roll (1984)). Roll intended this measure, as we do, to capture the deviations from fundamental value that are subsequently reversed. Because intraday returns have means that are close to zero, we can directly measure the covariance term (and thus the size of the reversal) using the uncentered product of the two adjoining returns.

most days because of bid-ask bounce and other market frictions. On non-witching days, NYSE stocks have an average reversal from the open to 10:00 AM of -0.11 (as indicated by the intercept in the time-series regression). Nasdaq stocks have significantly larger average reversals on non-witching days of -0.37 (as indicated by the sum of the intercept and the Nasdaq dummy in the time-series regression).

Reversals from the open to 10:00 AM are significantly larger on witching days than on non-witching days. The incremental witching day reversal from the open to 10:00 AM is -0.59 on the NYSE (as indicated by the $Witch_t$ dummy variable) and -2.36 on Nasdaq (as indicated by the sum of the $Witch_t$ and the $Witch_t Nasdaq_{m,t}$ dummy variables). The individual witching day coefficients and the differences across the markets are both statistically significant. If the reversals are measured from the close to 10:00 AM to 11:00 AM, none of the dummy variable coefficients is statistically significant. This indicates that the noise introduced into prices by the witching day liquidity shock is short-lived and essentially disappears by 10:00 AM.

TABLE 4
Reversals

Data is for 1998 to Q3 2004. The reversal from time i to time j is defined as the percent return from the close to time i times the percent return from time j to time i . For the time-series regressions, the cross-sectional average reversal is calculated for each market on each day. The average reversals are then regressed on a constant and dummy variables set equal to one for Nasdaq stocks (*Nasdaq*), witching days (*Witch*), and Nasdaq stocks on witching days (*Witch * Nasdaq*). Standard errors are corrected using the Newey-West procedure. For the pooled regressions, reversals for each stock are regressed on stock-specific control variables calculated over the previous month: daily return standard deviation (σ), log daily dollar trading volume ($\log(\$vol)$), and log market capitalization ($\log(mkt\ cap)$) expressed as deviations from the overall mean. These control variables are also interacted with the witching dummy. Standard errors in the pooled regressions are corrected for contemporaneous correlation and heteroskedasticity.

	Open to 10 AM		10 AM to 11 AM	
	Time Series	Pooled	Time Series	Pooled
Constant	-0.11 (0.01)	-0.12 (0.02)	0.02 (0.02)	0.02 (0.02)
<i>Nasdaq</i>	-0.26 (0.04)	-0.19 (0.04)	-0.12 (0.16)	-0.08 (0.07)
<i>Witch</i>	-0.59 (0.22)	-0.43 (0.18)	-0.19 (0.37)	-0.18 (0.35)
<i>Witch * Nasdaq</i>	-1.77 (0.62)	-2.28 (0.68)	0.12 (0.33)	0.03 (0.19)
σ		-0.07 (0.01)		0.03 (0.02)
$\log(\$vol)$		0.01 (0.03)		0.00 (0.04)
$\log(mkt\ cap)$		0.00 (0.03)		0.00 (0.03)
<i>Witch * σ</i>		-0.45 (0.18)		6.38 (12.04)
<i>Witch * $\log(\\$vol)$</i>		0.80 (0.25)		0.00 (0.12)
<i>Witch * $\log(mkt\ cap)$</i>		-0.47 (0.18)		0.06 (0.10)
Adj. R^2	0.0417	0.0007	0.0003	0.0001
Obs.	3,442	849,182	3,442	849,182

The time-series regressions in Table 4 control for differences in the average level of reversals on Nasdaq and the NYSE. As with the volatility regressions,

however, they do not control for firm characteristics that could affect the excess price reversals on witching days. Therefore, in these pooled regressions (one observation per stock per day) we control for the same firm characteristics that appeared in the volatility regressions: daily return standard deviation (σ), log daily dollar trading volume ($\log(\$vol)$), and log market capitalization ($\log(mkt\ cap)$). Stock characteristics are calculated for the prior calendar month and expressed as deviations from the mean. To allow these control variables to have different effects on witching and non-witching days, we interact the stock characteristics with the witching day dummy variable. Defining these control variables for stock s at date t as the vector $X_{s,t}$ we run the following regression:

$$(7) \quad rev_{s,t}^{i,j} = \alpha + \gamma' X_{s,t} + (\phi' X_{s,t}) Witch_t + \beta_1 Witch_t + \beta_2 Nasdaq_{s,t} + \beta_3 Nasdaq_{s,t} Witch_t + \varepsilon_{s,t}^{i,j}.$$

Standard errors in these regressions are based on Rogers (1993) and account for contemporaneous correlation and heteroskedasticity.

Controlling for stock characteristics such as stock return volatility, trading volume, and market capitalization has little effect on the differences in witching day price efficiency between Nasdaq and the NYSE. The price efficiency results in Table 4 closely parallel the volatility results in Table 3. On witching days, opening prices are more volatile and less efficient on Nasdaq than on the NYSE. However, on witching days in both markets, stock return volatility and price efficiency return to normal by 10:00 AM.

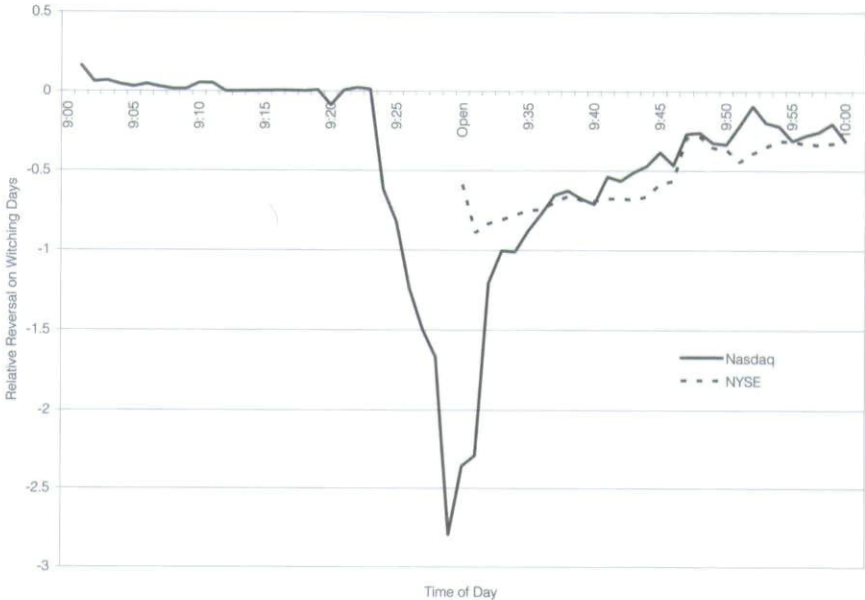
Figure 4 graphs excess witching day price reversals by minute from 9:00 AM to 10:00 AM. Excess witching day price reversals for the NYSE are measured by the coefficient on the witching day dummy variable in the time-series regression (6) and excess witching day price reversals for Nasdaq are measured as the sum of the coefficients on the $Witch_t$ and $Witch_t Nasdaq_{m,t}$ dummy variables. For Nasdaq stocks, pre-open prices become less efficient as the open approaches on witching days. After the open, the noise is quickly reversed and essentially disappears by 9:45 AM. Witching day opening prices for NYSE stocks are less efficient than normal during this period, but the relative inefficiency is smaller than for Nasdaq stocks.

VI. Nasdaq's Opening Cross

Nasdaq's opening procedure has been harshly criticized. In a letter dated May 16, 2000 to Frank Zarb, then chairman and chief executive officer of the National Association of Securities Dealers, Arthur Levitt of the SEC wrote, "I urge the NASD to pursue a unified opening procedure, and in the interim, to press forward with measures to make the opening process more reliable and fair to investors." S&P has also been critical of Nasdaq's opening and closing procedures. In a March 2004 trial, S&P used Amex closing prices for some Nasdaq stocks to calculate its index values, and S&P considered using Amex opening prices to calculate settlement prices for the S&P 500 futures, even though the Amex has only a tiny fraction of trading volume in these Nasdaq issues.

FIGURE 4
Witching Day Reversals

From 1998 to Q3 2004, the average return reversal ($ret_{i-1}ret_i$) in percent from the close to minute i to noon across stocks in each market (m) is regressed on an intercept, a witching day dummy, and the witching dummy interacted with a Nasdaq dummy. The witching day dummy (NYSE) is graphed from the open to 10 AM and the witching dummy interacted with a Nasdaq dummy (Nasdaq) is graphed from 9 AM to 10 AM. Returns at time i are only calculated for those stocks that have opened by that time.



These complaints, along with our results, led to changes in Nasdaq's market structure. On July 28, 2004, Nasdaq announced plans to improve its opening process by adding an electronic opening call auction for each Nasdaq stock. The resulting crossing price would represent the official Nasdaq opening price for each stock and be used for settling S&P 500 futures contracts, among other things.

Following SEC approval, the opening cross was phased in during the fourth quarter of 2004, beginning on November 8, 2004 with 75 actively traded stocks, of which 15 were members of the S&P 500. A second set of 181 Nasdaq stocks implemented the opening cross a week later. The rollout continued until the last 995 Nasdaq stocks were added on Monday, December 13, 2004. Four days later, on Friday December 17, the opening cross handled its first witching day. There have now been a total of four witching days handled by the Nasdaq opening cross.

Under Nasdaq's current procedures, the Nasdaq pre-market trading session begins at 8:00 AM and continues until 9:30 AM.²² Pre-market trading is virtually identical to regular hours trading in terms of functionality. During this time period, participants can also submit orders to two other separate order books: the regular hours order book and the opening cross order book. Orders in these two

²² At 8:00 AM, the Nasdaq system initiates the pre-market trading session with procedures to unlock and uncross the market in each security. Prior to April 18, 2005, this unlocking/uncrossing procedure was implemented at 9:25 AM. For details on the unlocking procedures, see NASD Rule 4704(b).

books are not executed until 9:30 AM, when the opening cross takes place in every Nasdaq stock. The regular hours trading session begins immediately thereafter.

In the opening cross order book, market participants can submit market-on-open orders, limit-on-open orders, and imbalance-only orders. An imbalance-only order is a priced order that, as its name suggests, executes only if it helps to reduce an order imbalance during the opening cross. At 9:30 AM, the opening order book is combined with the regular hours Nasdaq order book, and the opening price is the price that maximizes the number of shares crossed. If there are multiple prices that would cross the same number of shares, then the opening price minimizes the order imbalance, and if there are still multiple prices that are tied on these two criteria, the opening price is chosen to minimize the distance from the midpoint between the post-cross bid and ask on the regular hours order book. All orders that are priced better than the crossing price participate fully in the cross. If rationing is required for orders priced at the crossing price, displayed orders have priority over reserve (undisplayed) orders, followed by time priority.

Hypothetical cross results and net order imbalances are disseminated frequently to market participants beginning at 9:25 AM. There has been some slight variation in the dissemination times, but the following schedule has been in place since March 2005. Beginning at 9:25:30 AM, updates are disseminated every 15 seconds until 9:28:15 AM and every 5 seconds thereafter until the opening cross takes place at exactly 9:30 AM. After 9:28 AM, restrictions go into effect in order to prevent gaming and reduce order imbalances, thereby facilitating orderly price discovery at the open. For example, market-on-open orders cannot be entered or cancelled after 9:28 AM, while imbalance-only orders and limit-on-open orders can be improved but not cancelled. Limit-on-open orders that could create or worsen an order imbalance are not accepted after this time.²³ Finally, regular hours order submissions or cancellations submitted after 9:28 AM are held and are only acted upon after the opening cross.

A. Trading Volume and Price Efficiency Following Nasdaq's Opening Cross


The opening cross had an immediate impact on Nasdaq's opening trade size. Figure 5 graphs the daily average logarithm of the opening trade size for S&P500 stocks on Nasdaq and the NYSE. The size of the Nasdaq opening trade increased by a factor of 10 immediately after the opening cross was initiated, and further tripled during the following quarter. By mid 2005, the opening trade on non-witching days was still smaller on Nasdaq than on the NYSE, perhaps reflecting the continued fragmentation in the trading of Nasdaq stocks. On witching days, however, the sizes of the opening trades are now similar on Nasdaq and the NYSE.

Because Nasdaq trading is fragmented and ECNs continue to trade continuously prior to and during the opening cross, we might expect the Nasdaq opening cross to be less successful than the NYSE opening call in concentrating order flow at the open. There also could be an advantage to having a monopolist specialist in this particular situation or to having traders in close physical proximity on a

²³Exact details on the acceptance criteria for late limit-on-open orders can be found in NASD Rule 4704(a)(3).

FIGURE 5
Opening Trade Size

The logarithm of average daily open trade size (in dollars) averaged across Nasdaq and NYSE S&P 500 stocks, respectively, from January 1, 1998 to September 30, 2005.



trading floor. To study the impact of the new opening procedure, we modify the reversal regressions in equations (6) and (7) to include a dummy variable for the post-cross period ($Cross_t = 1$ on or after November 15, 2004 and 0 otherwise). The resulting time-series regression has one average reversal measure for each market m on each day t ($rev_{m,t}^{i,j}$) regressed on a set of indicator variables:

$$(8) \quad rev_{m,t}^{i,j} = \alpha + \beta_1 Witch_t + \beta_2 Nasdaq_{m,t} + \beta_3 Nasdaq_{m,t} Witch_t + \beta_4 Cross_t + \beta_5 Cross_t Witch_t + \beta_6 Cross_t Nasdaq_{m,t} + \beta_7 Cross_t Nasdaq_{m,t} Witch_t + \varepsilon_{m,t}^{i,j}.$$

The pooled regression is similar to equation (7) and differs only in that it includes a set of dummy variables for the period after the opening cross was implemented ($Cross_t = 1$):

$$(9) \quad rev_{s,t}^{i,j} = \alpha + \gamma' X_{s,t} + (\phi' X_{s,t}) Witch_t + \beta_1 Witch_t + \beta_2 Nasdaq_{s,t} + \beta_3 Nasdaq_{s,t} Witch_t + \beta_4 Cross_t + \beta_5 Cross_t Witch_t + \beta_6 Cross_t Nasdaq_{s,t} + \beta_7 Cross_t Nasdaq_{s,t} Witch_t + \varepsilon_{s,t}^{i,j}.$$

With just four witching days following the initiation of Nasdaq's opening cross, it is possible to examine each witching day individually. On June 17, 2005, a witching day, a Nasdaq member firm experienced a system error and mistakenly submitted badly mispriced orders for many Nasdaq stocks just before the open. The price impact of these erroneous orders was sufficiently large that Nasdaq took

the extreme step of breaking trades and adjusting opening prices for many stocks on that day.²⁴ While no adjustments were made for stocks in the S&P 500, the erroneous orders probably affected the opening prices for some of these stocks. To investigate the influence of these erroneous orders, we also estimate the price efficiency regressions with an additional dummy variable for June 17, 2005.

The time-series regressions in Table 5 are designed for simple differences-in-differences analyses: we are interested in the change in Nasdaq efficiency versus the change in NYSE efficiency after the Nasdaq opening cross is implemented. Relative to NYSE opening prices on non-witching days, Nasdaq opening prices improve significantly once the opening cross is implemented (the coefficient on $Cross_t Nasdaq_{m,t}$ is 0.25, with a t -statistic of around 5). This coefficient on the post-cross Nasdaq dummy variable exactly offsets the coefficient on the full sample Nasdaq dummy variable and indicates that the opening cross has completely eliminated the gap in opening price efficiency between Nasdaq and the NYSE on non-witching days. This result holds across all specifications that we estimate.

On witching days, Nasdaq also does better after the opening cross is in place. The coefficient in the time-series regression on $Cross_t Witch_t Nasdaq_{m,t}$ is 1.03, with a standard error of 0.49. Even though there are only four witching days where Nasdaq opens with a cross and, hence only four independent observations, the improvement on Nasdaq is statistically significant and economically large. However, Nasdaq does not quite catch up to the NYSE, which continues to open more efficiently on witching days. When we sum all of the relevant coefficients to compare post-cross witching day openings on the two markets, we find that Nasdaq reversals are larger than NYSE reversals by 0.74, with a p -value of 0.018.²⁵ The results are unchanged when we add market cap, size, and volatility control variables in a pooled regression. In fact, if anything the post-cross witching day gap between the two markets is slightly bigger under this specification.

After Nasdaq's implementation of its opening cross, most of the difference between the Nasdaq and the NYSE is attributable to June 17, 2005. When we include a dummy variable in the regression for this one day, the difference between Nasdaq and the NYSE on the remaining three witching days is not statistically significant.²⁶

Does Nasdaq still open less efficiently than the NYSE on witching days? The answer to that question depends on whether June 17, 2005 is included in the analysis or thrown out. One could argue that the system problems on that day were one-time events and are not representative of witching day performance on Nasdaq. In this case, it makes sense to exclude that day. However, the NYSE has always argued that human involvement reduces the impact of system errors. Had the erroneous orders been routed to the floor of the NYSE, the specialists would

²⁴See http://www.nasdaqtrader.com/dynamic/newsindex/headtraderalerts_2005.stm for alerts and information that Nasdaq disseminates to its members. June 17, 2005 stands out as the day when Nasdaq experiences its most severe and widespread problems.

²⁵On both the Nasdaq and NYSE, the temporary component in overnight returns continues to disappear rapidly after the markets open; prices are nearly efficient by 9:40 AM on both witching and non-witching days.

²⁶In addition, after dummifying out June 17, 2005, any remaining difference in the point estimates is largely attributable to the first post-cross witching day, December 17, 2004, when traders might still have been learning about how to use the new opening mechanism on witching days.

TABLE 5
Reversals after Nasdaq's Opening Cross

The reversal from the open to 10 AM is the percent return from the previous close to the open times the percent return from the open to 10 AM. The day column is the sum of ten half-hour reversals: the percent return from 10 AM to 10:30 AM times the percent return from 10:30 AM to 11 AM, the return from 10:30 AM to 11 AM times the return from 11 AM to 11:30 AM, and so on up to the return from 2:30 PM to 3 PM times the return from 3 PM to 3:30 PM. For the time-series regressions, the cross-sectional average reversal is calculated for each market and each day. The relevant reversal measure is regressed on a constant and dummy variables set equal to one for Nasdaq stocks (*Nasdaq*), witching days (*Witch*), and Nasdaq stocks on witching days (*Witch * Nasdaq*). Day specifications use data from Nov 15, 2004 to 2005Q3. Overnight reversal specifications use data from 1998 to Q3 2005 and the following additional variables: a dummy equal to one if the date is after 15 Nov 2004 when Nasdaq began its opening cross (*Cross*), a dummy set to one for 17 Jun 2005 (6-17-05), and a dummy set to one if the date is after 14 Apr 2003 when Nasdaq instituted an official closing price (*NOCP*). For the pooled regressions, reversals for each stock are regressed on stock-specific control variables calculated over the previous month: daily return standard deviation, log daily dollar trading volume, and log market capitalization, all of which are expressed as deviations from the overall mean. These control variables are also interacted with the witching dummy (coefficients not reported). Time-series regressions report Newey-West standard errors. Pooled regression standard errors correct for contemporaneous correlation and heteroskedasticity.

	Time Series				Pooled			
	Open to 10 AM			Trading Day	Open to 10 AM			Trading Day
Constant	-0.11 (0.01)	-0.11 (0.01)	-0.14 (0.02)	-0.01 (0.01)	-0.12 (0.02)	-0.12 (0.02)	-0.15 (0.02)	-0.01 (0.01)
<i>Nasdaq</i>	-0.25 (0.04)	-0.25 (0.04)	-0.29 (0.05)	-0.06 (0.01)	-0.20 (0.04)	-0.20 (0.04)	-0.24 (0.06)	-0.07 (0.01)
<i>Witch</i>	-0.59 (0.22)	-0.59 (0.22)	-0.73 (0.23)	-0.01 (0.02)	-0.52 (0.23)	-0.52 (0.23)	-0.61 (0.23)	-0.01 (0.02)
<i>Witch * Nasdaq</i>	-1.77 (0.62)	-1.77 (0.62)	-2.02 (0.67)	-0.03 (0.03)	-2.35 (0.70)	-2.35 (0.70)	-2.79 (0.88)	-0.02 (0.02)
<i>Cross</i>	0.06 (0.03)	0.06 (0.05)	-0.03 (0.03)		0.09 (0.02)	0.09 (0.02)	0.01 (0.01)	
<i>Cross * Nasdaq</i>	0.25 (0.05)	0.25 (0.05)	0.13 (0.04)		0.20 (0.04)	0.20 (0.04)	0.09 (0.03)	
<i>Cross * Witch</i>	0.53 (0.25)	0.53 (0.23)	0.03 (0.11)		0.17 (0.29)	0.40 (0.13)	-0.13 (0.15)	
<i>Cross * Witch * Nasdaq</i>	1.03 (0.49)	1.54 (0.71)	0.16 (0.61)		1.32 (0.64)	1.96 (0.81)	0.09 (0.58)	
6-17-05		-0.43 (0.10)				-0.49 (0.09)		
<i>Nasdaq*6-17-05</i>		-1.14 (0.45)				-1.85 (0.43)		
<i>NOCP</i>			0.10 (0.02)				0.11 (0.06)	
<i>NOCP * Nasdaq</i>			0.16 (0.05)				0.16 (0.04)	
<i>NOCP * Witch</i>			0.64 (0.26)				0.40 (0.22)	
<i>NOCP * Witch * Nasdaq</i>			1.12 (0.62)				1.72 (0.84)	
Adj. R^2	0.0444	0.0452	0.0523	0.0388	0.0005	0.0005	0.0006	0.0013
Obs.	3,876	3,876	3,876	402	938,822	938,822	938,822	98,905
<i>Nasdaq Worse than NYSE Compound Hypotheses Tests</i>								
Post-Cross Average Day		0.00	0.00		0.00		0.01	
Wald p -value		0.945	0.945		0.866		0.979	
Post-Cross Witching Day	-0.74	-0.23	-0.74		-1.03	-0.39	-0.97	
Wald p -value	0.018	0.286	0.018		0.002	0.148	0.005	


likely have recognized the problem and resolved it before the open, rather than breaking trades and adjusting opening prices after the open. If the Nasdaq open remains more vulnerable to system errors, perhaps it makes sense to give this day equal weight in the analysis. Ultimately, we leave it to the reader to draw his or her own conclusions from this evidence.

B. Other Changes at Nasdaq around the Initiation of the Opening Cross

In this section, we examine several other changes that occurred on Nasdaq around the time that Nasdaq initiated its opening cross to separate the effect of the opening cross from these other factors. Figure 6 shows that Nasdaq opening prices are least efficient during the period 2000–2002 and improve steadily thereafter. The timing in the improvement in opening price efficiency corresponds closely with the time-series changes in stock return volatility. Nasdaq stocks were most volatile during 2000–2001, with annualized return volatilities (defined as the standard deviation of log daily returns) for stocks in the Nasdaq 100 Index as high as 80%. By the beginning of 2003, the Nasdaq 100 volatility had fallen below 30%, and it continued a gradual decline to as low as 10% in mid-2005. Our regressions control for time-series variation in volatility, however, and these changes in volatility do not seem to account for Nasdaq's improved opening behavior.

FIGURE 6
Time-Series Patterns in Reversals

Return reversals ($ret_{t-1}ret_t$) in percent from the previous close to the open to the close are averaged across stocks in each market (NYSE or Nasdaq) and then averaged across all days in a calendar quarter. The sample extends from the beginning of 1998 to September 30, 2005.



Nasdaq also made a number of changes to its opening and closing procedures during this time period (as noted in Figure 6). In April 2003, Nasdaq instituted the Nasdaq Official Closing Price (NOCP), which used multiple trades around the 4:00 PM close to establish an official closing price rather than just the last printed trade. In April 2004, Nasdaq began to phase-in a closing cross in order to determine the official closing price. The opening cross discussed above followed in the last quarter of 2004. Because our reversal measures are based on the previous day's closing price, changes in closing procedures could have some effect on the measured reversals. As noted earlier, however, we obtain the same time-series

patterns and regression results as in Tables 3 and 4 when we use, say, 3 PM prices instead of closing prices.

To provide additional information about these competing explanations for Nasdaq's improved opening performance, we decompose the opening price behavior since 2003. To be precise, in addition to dummy variables equal to 1 once the opening cross is implemented, we add a set of dummies that correspond to the implementation of the official closing price ($NOCP_t = 1$ on or after April 14, 2003 and 0 otherwise). The time-series regression is:

$$(10) \quad rev_{m,t}^{i,j} = \alpha + \beta_1 Witch_t + \beta_2 Nasdaq_{m,t} + \beta_3 Nasdaq_{m,t} Witch_t + \beta_4 Cross_t + \beta_5 Cross_t Witch_t + \beta_6 Cross_t Nasdaq_{m,t} + \beta_7 Cross_t Nasdaq_{m,t} Witch_t + \beta_8 NOCP_t + \beta_9 NOCP_t Witch_t + \beta_{10} NOCP_t Nasdaq_{m,t} + \beta_{11} NOCP_t Nasdaq_{m,t} Witch_t + \varepsilon_{m,t}^{i,j}$$

and the pooled regression is:

$$(11) \quad rev_{s,t}^{i,j} = \alpha + \gamma' X_{s,t} + (\phi' X_{s,t}) Witch_t + \beta_1 Witch_t + \beta_2 Nasdaq_{s,t} + \beta_3 Nasdaq_{s,t} Witch_t + \beta_4 Cross_t + \beta_5 Cross_t Witch_t + \beta_6 Cross_t Nasdaq_{s,t} + \beta_7 Cross_t Nasdaq_{s,t} Witch_t + \beta_8 NOCP_t + \beta_9 NOCP_t Witch_t + \beta_{10} NOCP_t Nasdaq_{s,t} + \beta_{11} NOCP_t Nasdaq_{s,t} Witch_t + \varepsilon_{s,t}^{i,j}$$

This implicitly breaks the sample into three periods: January 1998 to April 2003 (before any changes to Nasdaq opening or closing procedures), April 2003 to December 2004 (new closing procedures, unchanged opening), and December 2004 to September 2005 (opening via the cross). The comparison between the last two subsamples is of the most interest.

The results are reported in Table 5. We focus on the pooled regressions that control for variation in stock return volatility, market cap, and trading volume (coefficients not reported). The NOCP dummy confirms that during the 2003–2004 period, Nasdaq improved in relation to the NYSE on both witching and non-witching days (reversals improve by 1.72 and 0.16, respectively). The regressions now use the 2003–2004 subsample as the base period for evaluating the effect of the opening cross. Relative to this base, the NYSE does not show any change in opening price efficiency, while Nasdaq opening prices reliably improve on non-witching days. Interestingly, in this specification we do not find any significant improvement associated with the opening cross on witching days. The relevant point estimate on $Cross_t Witch_t Nasdaq_t$ is positive at 0.09 but statistically insignificant. Here the power problems are compounded because there are only six witching days in the 2003–2004 period and four witching days in the post-cross period.

Overall, the evidence is clear on non-witching days. For S&P 500 stocks, the efficiency of Nasdaq opening prices has improved with the opening cross and now is indistinguishable from the efficiency of NYSE opening prices. Even though Nasdaq trading remains more fragmented overall, the opening cross effectively

consolidates order flow and facilitates price discovery at the open. The story is somewhat different on witching days. If one is willing to discard one anomalous witching day in June 2005, then there is little difference in opening price efficiency between the two markets on witching days. If the June 2005 witching day is included in the analysis, however, then Nasdaq still opens less efficiently than the NYSE on witching days, although this result is not reliably present in all of our specifications.

C. Price Efficiency during the Rest of the Trading Day

One possible explanation for our earlier results is that Nasdaq is improving overall, not just at the open. If true, then Nasdaq price efficiency should be comparable to NYSE price efficiency at other times during the day. Another possible reading of the earlier results is that they are quite general, with consolidation an important contributor to price efficiency. If this is true, we would expect worse price efficiency on Nasdaq during the rest of the trading day because Nasdaq is more fragmented than the NYSE outside of the opening and closing crosses, and there have been no comparable changes to Nasdaq's market structure during the rest of the trading day.

To measure price efficiency during the trading day, we calculate half-hour reversal measures that are analogous to the overnight reversal measures computed above. For stock i on day t , the intraday reversal measure $idrev_{it}$ is defined as:

$$(12) \quad idrev_{it} = r_{it}^{10:00,10:30} r_{it}^{10:30,11:00} + r_{it}^{10:30,11:00} r_{it}^{11:00,11:30} \\ + \dots + r_{it}^{14:30,15:00} r_{it}^{15:00,15:30},$$

where $r_{it}^{s,\tau}$ is the percent return on stock i on day t from time s to time τ .

We continue to compute a cross-sectional average $idrev_{m,t}$ for each market m and, as in (4), regress the two resulting daily time series on a Nasdaq dummy ($Nasdaq_{m,t}$), a dummy for witching days ($Witch_t$), and Nasdaq stocks on witching days ($Witch_t Nasdaq_{m,t}$):

$$(13) \quad idrev_{m,t} = \alpha + \beta_1 Witch_t + \beta_2 Nasdaq_{m,t} + \beta_3 Nasdaq_{m,t} Witch_t + \varepsilon_{m,t}.$$

To ease exposition and focus on the most current market structures, the sample period for this regression begins with the adoption of the opening cross on November 15, 2004 and extends through the third quarter of 2005.

As in (5), we also estimate a pooled version, with one observation per stock s per day t and the same normalized control variables $X_{s,t}$ as before (daily return standard deviation, log daily dollar trading volume, and log market capitalization expressed as deviations from cross-sectional means):

$$(14) \quad idrev_{s,t} = \alpha + \gamma' X_{s,t} + (\phi' X_{s,t}) Witch_t + \beta_1 Witch_t \\ + \beta_2 Nasdaq_{s,t} + \beta_3 Nasdaq_{s,t} Witch_t + \varepsilon_{s,t}.$$

The time-series and pooled results are virtually identical and are reported in the "Day" columns of Table 5. Among S&P 500 stocks on non-witching days, the Nasdaq dummy is -0.06 , indicating that Nasdaq half-hour reversals during

the day are significantly larger than the corresponding NYSE reversals. The t -statistic of about six is surprisingly large given that there only about 200 days in the sample. For completeness, we also look at intraday reversals on witching days. We do not find any evidence that intraday reversals are worse on witching days versus non-witching days, though statistical power may be low with only four such days in this sample period.

These results, taken together with the earlier findings, suggest that consolidation of order flow is important for the efficiency of prices. When Nasdaq's open is fragmented (before the adoption of the opening cross), the quality of opening prices is very poor. Once the opening cross is adopted and opening orders are consolidated in an electronic auction, Nasdaq opening prices improve dramatically and arguably become just as efficient as opening prices on the NYSE. However, this is not true during the middle of the trading day. During the trading day, order flow in Nasdaq stocks is more fragmented than order flow in stocks listed on the NYSE, and Nasdaq prices are not as efficient as NYSE prices. This points to a positive causal relation between order flow consolidation and price efficiency.

We realize that there are other possible ways to explain the cross-market differences in mid-day price behavior. Nasdaq and the NYSE differ in other ways besides consolidation of order flow, and perhaps one of these other differences accounts for the NYSE's superior price efficiency during the trading day. However, the evidence here is tantalizing and suggests that consolidation and price efficiency are closely linked.

VII. Summary

This paper examines whether market structure matters for the efficiency of stock prices and, in particular, whether order consolidation improves market efficiency. We find that prior to November 2004 the NYSE's opening call market performs better than the decentralized Nasdaq open on typical days and the NYSE performs much better than Nasdaq on witching days. The Nasdaq opening price is about 50% more volatile than the NYSE on non-witching days. However, the excess volatility on witching days is more than twice as large on Nasdaq as on the NYSE. These results indicate that the consolidation of orders on the floor of the NYSE improves the opening efficiency on normal trading days and is particularly beneficial when the market is stressed by extreme liquidity shocks.

In November 2004, Nasdaq introduced its own consolidated opening call auction and successfully consolidated a large volume of trade at the open. Nasdaq's opening cross essentially eliminated any difference in opening price volatility between the NYSE and Nasdaq on non-witching days. This result indicates that it is the consolidation of orders, rather than the physical proximity of the traders or the intervention of the specialist, that is important for price efficiency in a typical market opening. The results on witching days are more difficult to interpret. Nasdaq's opening is still more volatile than the NYSE on witching days, but most of this difference is attributable to a single post-cross witching day when a system error caused a firm to submit a large number of mispriced orders. If this particular witching day is ignored, then Nasdaq's post-cross openings on witching days have been almost as efficient as witching day openings on the NYSE.

However, it can also be argued that if these erroneous orders were submitted to the NYSE, the specialist would have intervened before the open.

We also provide evidence that consolidation is important at other times, not just at the open. Once Nasdaq adopts its opening call, Nasdaq and NYSE opening prices are equally efficient. But Nasdaq prices are less efficient during the rest of the trading day when Nasdaq returns to a decentralized, fragmented market structure. This suggests that consolidation, whether virtual or physical, is an important aid to efficient prices.

While our results are in the context of U.S. equity markets, our conclusion that the consolidation of orders improves efficiency should generalize to other settings, including other security markets (such as the dealer-oriented bond and foreign exchange markets) and markets in other geographic locations. Our results also have potentially important implications for derivative traders and the design of settlement procedures in derivative markets. When Nasdaq instituted its opening cross in November of 2004, it became easier for derivative traders to close out their index arbitrage positions. With the initiation of the opening cross, index arbitrage traders benefited from a reduction in the volatility of the opening price on witching days and from the ability to guarantee a trade execution at the price used for the settlement of the derivative contract. Although with only one structural change it is difficult to demonstrate a causal link, the improvement in Nasdaq's opening procedure is followed immediately by a significant upswing in open interest on the S&P 500 futures contract. Figure 1 shows a jump of \$50 to \$75 billion in open interest in November 2004 that is unparalleled in our sample. This jump in open interest was accompanied by a jump in index arbitrage trading at the open on witching days. NYSE audit trail data shows that after the initiation of Nasdaq's opening cross, the size of the NYSE opening trade increased by a factor of three on witching days—from about \$9 million per stock per day in 2003 and 2004 to almost \$27 million per stock per day thereafter, while there was no increase in the size of the opening trade on non-witching days—and more than half of this increase is due to trades identified as index arbitrage. The magnitude of the potential externality of Nasdaq's opening cross on the NYSE and on the S&P 500 futures contract is startling. The results suggest that the ability to close out index arbitrage positions easily and efficiently on witching days may materially impact both the trading and open interest of the futures contract.

References

- Admati, A., and P. Pfleiderer. "A Theory of Intraday Patterns: Volume and Price Variability." *Review of Financial Studies*, 1 (1988), 3–40.
- Amihud, Y.; B. Lauterbach; and H. Mendelson. "The Value of Trading Consolidation: Evidence from the Exercise of Warrants." *Journal of Financial and Quantitative Analysis*, 38 (2003), 829–846.
- Amihud, Y., and H. Mendelson. "Trading Mechanisms and Stock Returns: An Empirical Investigation." *Journal of Finance*, 42 (1987), 533–554.
- Amihud, Y., and H. Mendelson. "Volatility, Efficiency, and Trading—Evidence from the Japanese Stock Market." *Journal of Finance*, 46 (1991), 1765–1789.
- Barclay, M. "Bid-Ask Spreads and the Avoidance of Odd-Eighth Quotes on Nasdaq: An Examination of Exchange Listings." *Journal of Financial Economics*, 45 (1997), 35–58.
- Barclay, M., and T. Hendershott. "Price Discovery and Trading after Hours." *Review of Financial Studies*, 16 (2003), 1041–1073.

- Barclay, M., and T. Hendershott. "Liquidity Externalities and Adverse Selection: Evidence from Trading after Hours." *Journal of Finance*, 59 (2004), 681–710.
- Barclay, M., and T. Hendershott. "A Comparison of Trading and Non-Trading Mechanisms for Price Discovery." Working Paper, University of California, Berkeley (2007).
- Barclay, M.; T. Hendershott; and K. Kotz. "Automation versus Intermediation: Evidence from Treasuries Going Off the Run." *Journal of Finance*, 61 (2006), 2395–2414.
- Barclay, M.; T. Hendershott; and T. McCormick. "Competition among Trading Venues: Information and Trading on Electronic Communications Networks." *Journal of Finance*, 58 (2003), 2637–2666.
- Bennett, P., and L. Wei. "Market Structure, Fragmentation, and Market Quality." *Journal of Financial Markets*, 9 (2006), 49–78.
- Bessembinder, H., and H. Kaufman. "A Comparison of Trade Execution Costs for NYSE and Nasdaq-Listed Stocks." *Journal of Financial and Quantitative Analysis*, 32 (1997), 287–310.
- Bessembinder, H., and S. Rath. "Does Market Structure Matter?" In *Stock Market Liquidity: Implications for Market Microstructure and Asset Pricing*, G. Gregoriou and F.-S. Lhabitant, eds. Hoboken, NJ: John Wiley & Sons (2008).
- Boehmer, E. "Dimensions of Execution Quality: Recent Evidence for U.S. Equity Markets." *Journal of Financial Economics*, 78 (2005), 463–704.
- Cao, C.; E. Ghysels; and F. Hatheway. "Price Discovery without Trading: Evidence from the Nasdaq Pre-Opening." *Journal of Finance*, 55 (2000), 1339–1365.
- Christie, W.; S. Corwin; and J. Harris. "Nasdaq Trading Halts: The Impact of Market Mechanism on Prices, Trading Activity, and Execution Costs." *Journal of Finance*, 57 (2002), 1443–1478.
- Corwin, S., and M. Lipson. "Order Flow and Liquidity around NYSE Trading Halts." *Journal of Finance*, 55 (2000), 1771–1801.
- Demsetz, H. "The Cost of Transacting." *Quarterly Journal of Economics*, 82 (1968), 33–53.
- Domowitz, I., and A. Madhavan. "Open Sesame: Alternative Open Algorithms in Securities Markets." In *Building a Better Stock Market: The Call Market Alternative*, Robert Schwartz, ed. Boston, MA: Kluwer Academic (2000).
- Ellul, A.; H. Shin; and I. Tonks. "Opening and Closing the Market: Evidence from the London Stock Exchange." *Journal of Financial and Quantitative Analysis*, 40 (2005), 779–801.
- Foster, F. D., and S. Viswanathan. "A Theory of Intraday Variations in Volume, Variance, and Trading Costs in Securities Markets." *Review of Financial Studies*, 3 (1990), 593–624.
- Garbade, K., and W. Silber. "Structural Organization of Secondary Markets: Clearing Frequency, Dealer Activity and Liquidity Risk." *Journal of Finance*, 34 (1979), 577–593.
- Goldstein, M., and K. Kavajecz. "Trading Strategies During Circuit Breakers and Extreme Market Movements." *Journal of Financial Markets*, 7 (2004), 301–333.
- Hancock, G. "Whatever Happened to the Triple Witching Hour?" *Financial Analysts Journal*, 49 (1993), 66–72.
- Hendershott, T., and C. M. Jones. "Island Goes Dark: Transparency, Fragmentation, and Regulation." *Review of Financial Studies*, 18 (2005), 743–793.
- Huang, R., and H. Stoll. "Dealer versus Auction Markets: A Paired Comparison of Execution Costs on NASDAQ and the NYSE." *Journal of Financial Economics*, 41 (1996), 313–358.
- Jones, C. M., and M. Lipson. "Execution Costs of Institutional Equity Orders." *Journal of Financial Intermediation*, 8 (1999), 123–140.
- Lee, C.; M. Ready; and P. Seguin. "Volume, Volatility and New York Stock Exchange Trading Halts." *Journal of Finance*, 49 (1994), 183–214.
- Madhavan, A. "Trading Mechanisms in Securities Markets." *Journal of Finance*, 47 (1992), 607–641.
- Madhavan, A., and V. Panchapagesan. "Price Discovery in Auction Markets: A Look Inside the Black Box." *Review of Financial Studies*, 13 (2000), 627–658.
- Mendelson, H. "Consolidation, Fragmentation, and Market Performance." *Journal of Financial and Quantitative Analysis*, 22 (1987), 189–207.
- Pagano, M. "Endogenous Market Thinness and Stock Price Volatility." *Review of Economic Studies*, 56 (1989), 269–287.
- Pagano, M., and R. Schwartz. "A Closing Call's Impact on Market Quality at Euronext Paris." *Journal of Financial Economics*, 68 (2003), 439–484.
- Pagano, M., and R. Schwartz. "Nasdaq's Closing Cross." *Journal of Portfolio Management*, 31 (2005), 100–111.
- Petersen, M. "Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches." *Review of Financial Studies*, forthcoming (2008).
- Rime, D. "New Electronic Trading Systems in the Foreign Exchange Market." In *New Economy Handbook*, Derek Jones, ed. Amsterdam, NL: Academic Press/Elsevier (2003).
- Rogers, W. H. "Regression Standard Errors in Clustered Samples." *Stata Technical Bulletin*, 13 (1993), 19–23. Reprinted in *Stata Technical Bulletin Reprints*, 3, 88–94.

- Roll, R. "A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market." *Journal of Finance*, 39 (1984), 1127-1139.
- Roll, R.; E. Schwartz; and A. Subrahmanyam. "Liquidity and the Law of One Price: The Case of the Futures/Cash Basis." *Journal of Finance*, 62 (2007), 2201-2234.
- Schwert, G. W., and P. Seguin. "Heteroskedasticity in Stock Returns." *Journal of Finance*, 45 (1990), 1129-1155.
- Smith, J. "Nasdaq's Electronic Closing Cross: An Empirical Analysis." Working Paper, Nasdaq Economic Research (2005).
- Sofianos, G. "Expirations and Stock Price Volatility." *Review of Futures Markets*, 13 (1994), 39-108.
- Stoll, H., and R. Whaley. "Program Trading and Expiration-Day Effects." *Financial Analysts Journal*, 43 (1987), 16-28.
- Stoll, H., and R. Whaley. "Program Trading and Individual Stock Returns: Ingredients of the Triple-Witching Brew." *Journal of Business*, 63 (1990a), S165-S192.
- Stoll, H., and R. Whaley. "Stock Market Structure and Volatility." *Review of Financial Studies*, 3 (1990b), 37-71.
- Stoll, H., and R. Whaley. "Expiration-Day Effects: What Has Changed?" *Financial Analysts Journal*, 47 (1991), 58-72.
- Venkataraman, K. "Automated versus Floor Trading: An Analysis of Execution Costs on the Paris and New York Exchanges." *Journal of Finance*, 56 (2001), 1445-1885.
- White, H. "A Heteroscedasticity-Consistent Covariance Matrix Estimator and a Direct Test of Heteroscedasticity." *Econometrica*, 48 (1980), 817-838.