

The Shrinking New York Stock Exchange Floor and the Hybrid Market

Terrence Hendershott
Pamela C. Moulton *

April 11, 2008

* Hendershott is at Haas School of Business, University of California, Berkeley (hender@haas.berkeley.edu) and Moulton is at Fordham Graduate School of Business (pmoulton@fordham.edu). We thank Paul Bennett, Hank Bessembinder, Carole Comerton-Forde, Andrew Ellul, Sonali Hazarika, Pankaj Jain, Charles Jones, Bob Jennings, Albert Menkveld, Michael Pagano, Marios Panayides, Andriy Shkilko, and seminar participants at Wilfrid Laurier University, the NBER Market Microstructure meeting, and the Triple Crown Conference for helpful comments. We thank the New York Stock Exchange for providing data. Hendershott gratefully acknowledges support from the National Science Foundation. Moulton gratefully acknowledges support from a Fordham University Faculty Research Grant.

The Shrinking New York Stock Exchange Floor and the Hybrid Market

Abstract

At the end of 2006 the New York Stock Exchange introduced its Hybrid market, greatly expanding automated electronic trading. We find that as floor trading decreases, cooperation among floor traders declines. This decline in cooperation along with faster electronic trading leads to higher trading costs and adverse selection. Intraday volatility also increases, but the noise in stock prices declines, suggesting that more information is incorporated into prices and prices become more efficient. Together these findings support the existence of a tradeoff in market mechanisms among trading costs and speed and price efficiency. Hybrid moved the New York Stock Exchange to a different position on that tradeoff.

1. Introduction

As electronic trading comes to dominate financial markets (Jain (2005)), are there drawbacks to faster, anonymous electronic trading as opposed to slower, non-anonymous human-intermediated trading? From October 2006 through January 2007 the New York Stock Exchange (NYSE) introduced its ‘Hybrid’ market, which removed frequency and size restrictions for automated electronic execution and consequently reduced floor trading.¹ We use this event to study the impact of the increase in electronic trading speed and the decline in NYSE floor trading.

Execution speed is an important component of execution quality for professional traders but has been little studied in the academic literature. Boehmer, Jennings, and Wei (2007) confirm that execution speed matters to traders by finding that a market center receives more order flow when its reported execution speed increases. Boehmer (2005) discusses several reasons that traders prefer faster executions. Delay induces uncertainty about the probability of execution or the price at which execution occurs. Traders’ risk aversion makes such uncertainty undesirable. Even if traders are risk neutral, many trading strategies are more difficult to implement with slower execution. Strategies contingent on prices, long-short strategies involving simultaneous trades in multiple securities, and strategies which break larger orders into smaller orders all perform worse as execution times increase. The Hybrid market lowered the execution time for market orders from over 10 seconds to less than one second.

Past work comparing different market structures suggests that the overall quality of a market is multidimensional (Battalio, Hatch, and Jennings (2003) and Boehmer (2005)). These papers compare execution quality across different market structures and find that faster markets are associated with higher trading costs. By examining the change within a market, we provide direct evidence on how stock exchanges can move along the speed-cost tradeoff. In addition to the change in speed, automatic execution allows us to explore how speed affects the rate and efficiency of information incorporation into prices.

¹ Lucchetti (2007) discusses how NYSE member firms reduced the number of floor traders with Hybrid’s introduction. The NYSE has closed three of the five rooms on its trading floor (see McGeehan (2007)), and two firms have closed their specialist businesses (see Dowell (2007)).

Our empirical strategy is an event study of the Hybrid rollout. To control for changes in overall market conditions we match NYSE stocks with Nasdaq stocks. We then perform tests using a difference-in-difference approach, comparing the difference between NYSE and Nasdaq stocks before and after the Hybrid change. We find that Hybrid increases NYSE trading costs: From the month prior to the month subsequent to each stock's Hybrid activation date, quoted spreads increase from 7.9 basis points to 8.3 basis points and effective spreads increase from 5.6 basis points to 5.9 basis points. This is despite the fact that trade size falls by 25% with Hybrid's introduction. In contrast, over this same interval the spreads for the matched Nasdaq stocks decline. When we control for changes in Nasdaq spreads, Hybrid results in NYSE trading costs increasing by roughly 10% over their pre-Hybrid levels.

These overall measures of trading costs include changes in trading costs for electronic trading, for floor trading, and for trading that includes both electronic orders and floor participants (specialists and floor brokers).² In addition, the relative magnitudes of the types of trading changes: From the month prior to the month subsequent to each stock's Hybrid activation date, floor participants' share of NYSE trading volume drops from 15 percent to 11 percent. To calculate trading costs conditional on the makeup of the trade, we use NYSE audit trail data that identifies whether a trade is composed of electronic orders, floor participants, or a mix of the two. Conditioning shows that Hybrid increases trading costs for all types of trades, but that the increase is largest for pure floor trades, which prior to Hybrid had lower trading costs than trades involving electronic orders only.

Prior to the Hybrid market introduction, the lower cost of floor trading arose from lower adverse selection. A number of papers argue that reputations formed through repeated interactions, such as those of participants physically collocated on a trading floor, reduce adverse selection and enhance liquidity (see, e.g., Benveniste, Marcus, and Wilhelm (1992) and Chan and Weinstein (1993)).³ Our results are consistent with Battalio, Ellul, and Jennings' (2007) evidence on the reputational benefits of floor

² For a discussion of the multiple ways trades occur on the NYSE, see Sofianos and Werner (2000), Werner (2003), and Moulton (2006).

³ The relative costs of trading under different market structures are examined in research including Lee (1993), Huang and Stoll (1996), and Bessembinder and Kaufman (1997).

participants' repeated interaction.⁴ After Hybrid's introduction adverse selection increases substantially for floor trades, suggesting that cooperation declines and reputations wane as trading shifts away from face-to-face interactions.

Floor participants may use their advantageous role in the trading process to the disadvantage of off-floor traders (Rock (1990), Ready (1999), and (Stoll and Schenzler (2006))).⁵ We also find evidence that floor participants benefit relative to system orders prior to Hybrid, as ex-post measures of trading profitability – realized spreads – favor floor participants over system orders. After the introduction of Hybrid, this difference declines significantly. Thus, we find evidence supporting both the costs and benefits of floor trading. If Hybrid's overall impact is interpreted as being solely due to the decline of floor trading, then our findings suggest that the net effect of floor trading is to reduce trading costs.

The increase in trading speed due to Hybrid leads to an overall increase in adverse selection. In conjunction with the rise in adverse selection, intraday (five-minute) volatility increases. To examine whether the increase in volatility is due to noise or additional information being incorporated into prices, we measure of price efficiency through the ratio of the five-minute variance to 30-minute variance. The variance ratio declines with Hybrid's introduction, suggesting that there is less noise in prices and that the increase in speed allows more information to be incorporated into prices more efficiently.

The remainder of the paper is organized as follows. Section 2 provides an overview of the Hybrid market changes. Section 3 describes our data and sample. Section 4 examines the relation between the Hybrid market introduction and transaction costs. Section 5 investigates how the changes in liquidity surrounding Hybrid are related to changes in the way traders interact on the NYSE. Section 6 analyzes changes in adverse selection with the Hybrid introduction. Section 7 examines how Hybrid influences market efficiency. Section 8 concludes.

⁴ A number of studies empirically examine the trading of specialists and floor brokers. Hasbrouck and Sofianos (1993) and Madhavan and Smidt (1993) and others examine specialist trading. Sofianos and Werner (2000), Werner (2003), and Handa, Schwartz, and Tiwari (2004) study floor broker trading.

⁵ See Ip and Craig (2003) for some details on the NYSE specialist investigations. While not studying a floor-based market, Christie and Schultz (1994), Christie, Harris, and Schultz (1994), and Barclay (1997) provide evidence on how intermediaries collude to use their advantages against outside investors.

2. An Overview of the Hybrid Market

Before we discuss the Hybrid market changes, some background on how continuous trading on the NYSE was conducted before the Hybrid implementation is useful. The traditional auction mechanism on the NYSE requires that a specialist manually execute each trade, allowing the specialist (who is a designated market maker) and floor brokers (who represent customer orders) to provide liquidity and participate in trades at the point of sale. Electronic trading on the NYSE began with the DOT system in 1976, which allowed electronic submission for market orders of 100 shares. Upon reaching the NYSE trading floor, the electronic DOT orders were executed by the specialist in the traditional auction mechanism. The DOT system's capabilities were expanded over time to support limit orders and larger sizes, and the system was renamed SuperDOT in 1984. Harris and Hasbrouck (1996) report that floor trading was 70 percent of total volume for 1990-1991. Werner and Sofianos (2000) find that this fraction declines to 55 percent by 1997. By 1999 electronic and floor trading are roughly equal on the NYSE.

Figure 1 graphs aggregate floor trading as a percentage of aggregate NYSE dollar volume, as well as its breakdown by floor broker and specialist trading, for 1999 through mid-2006. Floor trading activity begins to noticeably decline in 2002. The beginning of this decline appears around the NYSE's January 2002 introduction of OpenBook, which provides limit-order-book information to traders off the exchange floor (Boehmer, Saar, and Yu (2005)). Initially OpenBook data were released every 10 seconds, later reduced to every five seconds, and on May 1, 2006 OpenBook began to be disseminated as continuously as the NYSE systems allow. In addition to the frequency of dissemination, OpenBook is limited in that it does not include floor participants' interest and there are still lags in executions on the floor. Despite these limitations, the substitution of electronic trading for floor trading identified in Boehmer, Saar, and Yu (2005) continues and grows from 2002 onwards. By June 2006 floor trading represents slightly more than 20 percent of NYSE volume.

[Figure 1 Here]

Automatic execution was introduced on the NYSE in 2000. Prior to the Hybrid market, automatic execution was restricted to priced orders (i.e., limit orders) of up to 1,099 shares and subject to a 30-second rule for repeat executions for accounts belonging to the same beneficial owners. Automatic execution orders could be executed only at the inside quote, i.e., the same trader could not “walk the book” by executing at multiple prices, and had to be specifically designated. Furthermore, the default treatment of marketable limit orders and the only option for market orders was execution via the auction mechanism. In 2003 the NYSE began automatically updating best bid and offer quotes to reflect changes in the limit order book; prior to 2003 the best bid and offer were refreshed manually by the specialist.

The NYSE gave three reasons for launching the Hybrid market (NYSE Group (2006b)). First, they believe that customers want a choice of using the existing auction mechanism for the possibility of better prices or accessing the book electronically to achieve faster execution. Second, they expect trading volume to continue to increase, and higher volume can be handled more efficiently in a more automated system. Third, the Securities and Exchange Commission (SEC) Reg NMS Order Protection Rule protects better-priced quotes from being traded through only in markets that are “fast”, defined as markets that offer automatic execution at the posted quotes.

The biggest change in the Hybrid market is the expansion of automatic execution. Orders are no longer limited to 1,099 shares (the new limit is one million shares), the frequency restriction is eliminated, orders may walk the book beyond the best bid and offer, and non-priced (market) orders as well as limit orders are eligible for automatic execution. Market and marketable limit orders are now automatically executed by default, rather than requiring a special code. In Hybrid the NYSE also introduced Liquidity Replenishment Points (LRPs), which are stock-specific price ranges intended to defend against erroneous trades and dampen volatility by converting the market from fast (automatic execution available) to “auction only” (auction mechanism, no automatic execution available) when prices move quickly in either direction. Immediately following Hybrid introduction the market was fast 98.9 percent of the time (NYSE Group (2007)); in February 2007 the NYSE reset the LRPs to less restrictive levels.

Figure 2 shows average execution speeds for market and marketable limit orders – the two types of orders used by those most desiring speed – in the four months surrounding the Hybrid introduction (data from the SEC Dash-5/Rule 605 filings). Time to execution declines by more than 50 percent from the month before to the month after Hybrid’s introduction. Execution time falls for the smallest orders (under 500 shares) as well as overall, evidence that improved execution speed was not strictly a result of the declining average trade size.

[Figure 2 Here]

In addition to reducing execution time, the expansion of automatic execution reduces the opportunities for specialists and floor brokers to participate manually in executions. Another important set of changes in Hybrid gives the specialist and floor brokers ways to participate electronically that correspond to their prior trading capabilities: placing undisplayed as well as displayed orders on the limit order book. In addition, the specialist for each stock can use a proprietary algorithm to interact electronically with customer order flow, subject to a set of rules intended to replicate in an electronic framework what the specialist is allowed to do manually in the auction market (see NYSE Group (2006a)).

The Hybrid market changes apply only during continuous intraday trading: Automatic execution is not available during the opening and closing auctions, which are conducted manually by the specialist as before. Hybrid activation was rolled out gradually between October 6, 2006, and January 24, 2007.⁶ All stocks that trade in 100-share round lots were activated over the four-month period; 43 stocks that trade in round lots of 10 shares (e.g., Berkshire Hathaway Inc., which is priced near \$100,000 per share) were not included in the initial Hybrid rollout.

Figure 3 graphs overall floor trading as well as its breakdown by floor broker and specialist trading for the year surrounding the Hybrid rollout (June 2006 through May 2007). The floor activity at the

⁶ We focus on the expansion of automatic execution under Hybrid, which the NYSE labeled Hybrid Phase 3. Hybrid Phase 1 (rolled out 12/1/05 through 4/5/06) and Phase 2 (rolled out 4/6/06 through 8/21/06) upgraded various NYSE systems to facilitate the Phase 3 expansion of automatic execution. Hybrid Phase 4 (rolled out 1/25/07 through 2/28/07) introduced changes required for the implementation of Reg NMS, such as new order types and new locking and crossing rules.

beginning of Figure 3 does not line up exactly with the end of Figure 1 because the data in Figure 3 exclude certain trades that were unaffected by Hybrid and reflect only the 400 stocks in our sample; see the Data section for details. Using the scale on the right y-axis, Figure 3 also indicates the percentage of stocks for which Hybrid has been introduced. Relatively few stocks went Hybrid in the first two months of the rollout. Almost half of the stocks went Hybrid at the beginning of December 2006, and another 40 percent of stocks went Hybrid over the final few weeks of 2006. Floor activity declines gradually before and after the Hybrid introduction. When the transition to Hybrid is most intense in December 2006, floor activity declines steeply from 15 percent to 11 percent.

[Figure 3 Here]

The fact that many stocks went Hybrid in close proximity to each other requires that our empirical strategy control for contemporaneous changes in market liquidity. We do this by matching NYSE stocks to Nasdaq stocks and following a difference-in-difference approach, examining how the Hybrid event and consequent reduction in floor activity impact the difference between NYSE and Nasdaq stocks.

3. Data and Sample Selection

Our analysis uses data from the NYSE's Trade and Quote (TAQ) database, the Center for Research in Security Pricing (CRSP), the Chicago Board of Options Exchange (CBOE), SEC Rule 11Ac1-5 (Dash-5, now called Rule 605) filings, and the NYSE internal Consolidated Equity Audit Trail (CAUD) database. We collect data from June 1, 2006 through May 31, 2007, which spans the period from roughly four months before to four months after the Hybrid activation interval. This period facilitates the testing of changes both in the window immediately surrounding each stock's Hybrid activation date and over a longer horizon to capture possible delayed adjustments to the changes. We focus on a sample of 400 NYSE-listed stocks that went Hybrid, using a matched sample of 400 Nasdaq-listed stocks to control for market-wide changes in market quality.

3.1 Sample Construction

We construct a sample of 400 NYSE-listed common stocks as follows. We begin by collecting from CRSP the market capitalizations and closing prices of all domestic common stocks listed on the NYSE as of March 31, 2006. From the TAQ Master History file we determine CUSIP numbers that correspond to the symbols in TAQ, to accurately match stocks in CRSP and TAQ. We also use the TAQ Master History file to eliminate stocks that were not listed continuously from March 2006 through May 2007 or changed symbol during the period. We eliminate stocks with prices below \$1 or over \$500, stocks with two or fewer trades per day on average according to TAQ, and stocks that are not included in the Hybrid activation list posted on the NYSE website. Finally, we rank the remaining stocks by market capitalization and randomly select 50 stocks from each of the top eight market capitalization deciles.⁷

We construct a matched sample of 400 Nasdaq-listed stocks as follows. Using one-to-one matching without replacement, we determine a unique Nasdaq match for each stock in our NYSE sample based on CRSP market capitalization and closing price.⁸ We measure the matching criteria at the end of the first quarter of 2006, which precedes our analysis period. We randomize the order of matching by sorting NYSE stocks alphabetically by symbol. We then calculate the following matching error for each NYSE stock i and each remaining Nasdaq stock j :

$$\text{matching error} = \frac{\left| \frac{MCAP_i}{MCAP_j} - 1 \right| + \left| \frac{PRC_i}{PRC_j} - 1 \right|}{2}, \quad (1)$$

where $MCAP$ is the stock's market capitalization and PRC is the stock's closing price. The Nasdaq stock with the lowest matching error is selected as the match for that NYSE stock and removed from the list of potential Nasdaq matches for the remaining NYSE stocks. The mean matching error for the 400-stock sample is 0.76. In earlier analysis we used the same matching procedure for a sample of 160 stocks; the

⁷ We exclude stocks from the two smallest market capitalization deciles because they do not have enough trades to produce valid estimates for several of our market quality measures.

⁸ Davies and Kim (2007) find that one-to-one matching based on market capitalization and share price is the most appropriate method for comparing trade execution costs between NYSE and Nasdaq stocks. They also conclude that eliminating poor matches is not advisable.

160-stock sample has a mean matching error of 0.08. Because all results are qualitatively identical for the 160-stock subsample, we report only the full 400-stock sample results. Table 1 presents descriptive statistics for the NYSE and Nasdaq samples.

[Table 1 Here]

3.2 Data and Measures

The liquidity measures in Table 1 and throughout the paper are calculated from TAQ trade and quote data, as are the intraday volatility and efficiency measures to follow. We determine floor and system trading participation from the CAUD database, which contains detailed information about all trades executed on the NYSE. We obtain execution speed from the SEC Dash-5/Rule 605 data. To measure market-wide volatility we use the daily opening CBOE volatility index (VIX), which is derived from S&P 500 stock index options.

We calculate spreads for NYSE stocks two ways: using trades and quotes from the NYSE only, and using trades and quotes from all markets.⁹ As the results from both samples yield identical inference, we present only the measures and results based on NYSE trades and quotes. Spreads for Nasdaq stocks are calculated using trades and quotes from all markets. We use trades and quotes from regular-hours trading only. Upstairs-arranged trades (see Madhavan and Cheng (1997)), opening trades, and closing trades are excluded because they take place outside of the trading mechanisms that changed under the Hybrid market implementation.

We equally-weight spread measures across trades within the day to calculate measures for each stock each day.¹⁰ The percentage quoted spread is the difference between the best ask price and the best bid

⁹ We apply the following filters to clean the trade and quote data. We use only trades for which TAQ's CORR field is equal to zero, one, or two and for which the COND field is either blank or equal to @, E, F, I, J, or K. We eliminate trades with nonpositive prices or quantities. We eliminate trades with prices more than (less than) 150% (50%) of the previous trade price. We use only quotes for which TAQ's MODE field is equal to 1, 2, 6, 10, 12, 21, 22, 23, 24, 25, or 26. We eliminate quotes with nonpositive price or size or with bid price greater than ask price. We exclude quotes when the quoted spread is greater than 25% of the quote midpoint or when the ask price is more than 150% of the bid price.

¹⁰ Results using measures volume-weighted within the day yield qualitatively similar results, which are available from the authors upon request.

price at the time of a trade, divided by the prevailing midpoint of the bid and ask quotes.¹¹ We calculate quoted depth as the time-weighted average depth at the best bid and ask.

The effective spread for each trade captures the difference between an estimate of the true value of the security (the quote midpoint) and the actual transaction price. The percentage effective spread for stock j at time k on day t is calculated as:

$$ESpread_{j,k,t} = 2 I_{j,k,t} (P_{j,k,t} - M_{j,k,t}) / M_{j,k,t} ,$$

where $I_{j,k,t}$ is an indicator variable that equals one for buyer-initiated trades and negative one for seller-initiated trades, $P_{j,k,t}$ is the trade price, and $M_{j,k,t}$ is the matching quote midpoint. We follow the standard trade-signing approach of Lee and Ready (1991) and use contemporaneous quotes to sign trades—see Bessembinder (2003).

Figure 4 depicts the average effective spread by market capitalization quartile over the window from 20 days before to 20 days after Hybrid activation. Effective spreads generally rise at Hybrid activation (day zero) and remain higher for the next 20 days. Because of the differences across quartiles, we will conduct most of our analysis by quartile as well as for the full sample.

[Figure 4 Here]

Figure 5 follows the same format as Figure 4 and graphs the difference in effective spread between NYSE stocks and their Nasdaq matches. The changes following Hybrid's introduction in the difference between NYSE and Nasdaq spreads in Figure 5 are similar to Figure 4. This suggests that changes in spreads on the NYSE in the months surrounding Hybrid were not due to market-wide changes in liquidity that also affected Nasdaq stocks. If anything the changes in NYSE spreads subsequent to Hybrid appear larger after controlling for the spreads of matching Nasdaq stocks.

[Figure 5 Here]

Figure 6 extends the analysis window from 20 days to four months before and after Hybrid activation and shows effective and quoted spread NYSE – Nasdaq differences for the full 400-stock sample. The

¹¹ Results using dollar spreads yield identical inference, so we present only percentage spreads for brevity. Dollar-spread results are available from the authors upon request.

increases in effective and quoted spread differences appear to be more than just a transitory adjustment to the Hybrid changes.

[Figure 6 Here]

4. Hybrid and Transaction Costs

To examine how Hybrid affects transaction costs we move beyond the simple spread graphs. For each stock we calculate the above-described measures for the 20 days before and the 20 days after Hybrid activation. The first four columns of Table 2 present these results for quoted and effective spreads. Consistent with Figure 3, the quoted and effective spread measures in Table 2 show that NYSE spreads increase with Hybrid's introduction. Controlling for the matched Nasdaq stocks' spreads increases the pre/post difference (as in Figure 5). This implies that while NYSE stock spreads widen at the time of Hybrid, Nasdaq spreads narrow. The increase in spreads is generally greater for smaller stocks.

[Table 2 Here]

The last three columns of Table 2 examine trade size, quoted depth, and trading volume. Trade size shows a clear decline with Hybrid's introduction. Floor trades are larger than system trades (e.g., Moulton (2006)), so the shrinking trade size is likely due to the reduction in floor activity and increase in system activity. Note that smaller trades generally have lower effective spreads, so the declining trade size does not explain the widening of spreads around Hybrid introduction. Quoted depth shows little change with Hybrid, which suggests the wider spreads are not associated with additional liquidity at the quote. Trading volume also shows no significant change.

Table 2 uses standard univariate t -tests to calculate the statistical significance of changes associated with Hybrid. Given that the pre- and post- periods surrounding Hybrid activation overlap for many stocks (Figure 3), the assumption of independence across observations may overstate statistical significance. To properly control for this we adopt a panel data approach in Table 3. For each spread variable we run the following regression:

$$Sprd_{i,t} = \alpha + \beta Hybrid_{i,t} + \gamma Volatility_t + \sum_{k=1}^8 \delta_k ControlVariable_{i,k} + \varepsilon_{i,t}, \quad (2)$$

where $Sprd_{i,t}$ is the average quoted or effective spread for stock i less its Nasdaq match on day t ; $Hybrid_{i,t}$ is an indicator variable taking the value of one if the stock is in Hybrid mode on day t , otherwise zero; $Volatility_t$ is the opening value of CBOE's VIX index on day t ; and $ControlVariable_{i,k}$ are eight stock-level control variables: the log of the market capitalization and the log of the stock price on March 31, 2006, and the dollar turnover and return volatility in the first quarter of 2006, for each NYSE and Nasdaq matched stock.¹² We also run a variation on Equation (2) that includes stock fixed effects, α_i . We conduct inference using double-clustered Thompson (2006) standard errors, which are robust to both cross-sectional correlation and idiosyncratic time-series persistence. Regressions are also run by quartile.¹³

[Table 3 Here]

We conduct our analysis over two periods: the 40-day window surrounding each stock's Hybrid activation (first four columns) and from June 1, 2006 to May 31, 2007 (last four columns), roughly one year surrounding the Hybrid activation period. Hybrid represents a significant change in the trading environment. It is possible that market participants take time to adjust their trading strategies, leading to different effects in the long run than in the month following Hybrid introduction.

The Hybrid coefficients in the quoted and effective spread regressions over the 40-day window (first four columns of Table 3) are of the same magnitude as the average changes in Table 2. The inclusion of volatility in Table 3 demonstrates that the increase in spreads is not due to changes in volatility affecting NYSE and Nasdaq securities differently. The coefficients on Hybrid are all positive and statistically significant, with the smallest t -statistic 2.2 and the largest 6.9. As in Table 2 the Hybrid impact generally increases in the smaller quartiles. For the full sample regression of effective spreads with stock fixed effects, the coefficient on Hybrid is 0.53 basis points. This is almost a 10 percent increase in spreads from

¹² Omission of the control variables from equation (2) does not affect the coefficients on our variable of interest.

¹³ On December 1, 2006, the NYSE eliminated the monthly transaction fee cap, raised the per-transaction fee, and eliminated specialist commissions. Including a dummy variable corresponding to the NYSE's fee structure change does not significantly affect the coefficients of interest.

the 5.6 basis point pre-Hybrid average in Table 2. While the magnitude of the Hybrid coefficient is larger for smaller stocks, the spreads are also wider for smaller stocks. As a percentage of the pre-Hybrid average, the increase due to Hybrid is greatest for the largest stocks: 0.43 basis points on an average of 2.8 basis points (Table 2) for an increase of 15 percent. The results for the full-year regressions (last four columns) show similar coefficients on Hybrid to those for the 40-day window. For example, the effective spread increase attributed to Hybrid is 0.66 basis points in the full-year analysis versus 0.53 basis points in the 40-day window (full sample with fixed effects). Not surprisingly, extending the length of the sample period increases the statistical significance.

The analysis in this section shows that trading costs increase following Hybrid and that these changes are not merely transitory adjustment effects. Next we more fully explore how the changes in execution costs affect floor and system trading. In particular, we test whether or not the increases in trading costs are isolated to floor trades and whether there is evidence of declining cooperation among floor participants.

5. Separating Hybrid's Impact on the Costs of Floor and System Trading

Up to this point we have used the publicly available TAQ data to measure trading costs. The NYSE's CAUD file provides additional information on whether the counterparties in a trade are floor participants (specialists and floor brokers) or electronic orders. The CAUD database matches buyers and sellers for each NYSE trade, providing information about all of the parties on each side of a trade. Note that there can be more than one type of participant on each side of a single trade. For example, a system buy order for 1000 shares of ABC may execute against a system offer of 500 shares, a floor broker offer of 300 shares, and a specialist offer of 200 shares, all at the same price and time. Participation rates are computed by summing the purchases and sales by each type of market participant (system, floor broker, and specialist) and dividing by twice total volume, since the numerator double-counts volume. In contrast,

who trades with whom is determined by identifying all of the types of market participants involved in each trade and then categorizing the trade as follows:¹⁴

Pure Floor = Specialist and Floor Brokers, or Floor Brokers only;

Pure System = System participants only;

Floor and System Interaction = Specialist and System participants, or Floor Brokers and System participants, or Specialist, Floor Brokers, and System participants.

We further decompose floor and system interaction trades into those initiated by floor participants, those initiated by system participants, and those with mixed initiator types (meaning both floor and system participants on one or both sides of the trade).¹⁵ The ABC trade described above would be categorized as a mixed-initiator floor and system interaction trade, because it involves a specialist, a floor broker, and system participants, and both floor and system participants are on one side of the trade. Who-trades-with-whom trade type percentages are calculated by summing volume across trades in each category for each stock each day, then dividing by total traded volume in that stock that day.

Using this categorization of trades, Table 4 shows that the most significant switch in the 40-day window surrounding Hybrid activation is a roughly eight percent change towards pure system trading away from mixed initiator floor-system interaction trades. This likely stems from Hybrid's automatic execution precluding floor participants from joining what would otherwise have been pure system trades all along. The fraction of trading that is pure floor is nearly unchanged in the 40-day window surrounding Hybrid's introduction, remaining under two percent in all quartiles. Floor-initiated interaction trades decrease about one to two percent while system-initiated interaction trades increase a similar amount. The last five columns reveal similar patterns over the one-year period surrounding the introduction of Hybrid.

[Table 4 Here]

Effective spreads are perhaps the most common measure of trading costs. To examine changes in adverse selection and liquidity provider profits it is useful to decompose effective spreads into their

¹⁴ Percentage (CAP) executions are included as floor broker executions. Incoming Intermarket Trading System (ITS) executions are included as system participant executions.

¹⁵ For more detailed decompositions of who trades with whom, see Moulton (2006).

permanent and transitory portions. The percentage price impact for each trade in stock j at time k on day t reflects the permanent effect, a measure of adverse selection, and is calculated as:

$$PImpact_{j,k,t} = 2 I_{j,k,t} (M_{j,k+5,t} - M_{j,k,t}) / M_{j,k,t} ,$$

where $I_{j,k,t}$ is an indicator variable that equals one for buyer-initiated trades and negative one for seller-initiated trades (see Data section for details), $M_{j,k,t}$ is the matching quote midpoint, and $M_{j,k+5,t}$ is the quote midpoint five minutes after the trade. The realized spread reflects the temporary effect, approximating the profit earned by the liquidity provider, and is equal to the difference between the percentage effective spread and the price impact:

$$RSpread_{j,k,t} = 2 I_{j,k,t} (P_{j,k,t} - M_{j,k+5,t}) / M_{j,k,t} ,$$

where $P_{j,k,t}$ is the trade price and other variables are as defined above.

To study how Hybrid affects relationships on the trading floor and the interaction between the floor and system participants, we run regressions of the following form:

$$\begin{aligned} SprdDecomp_{i,k,t} = & \alpha_i + \sum_{k=1}^5 \gamma_k Hybrid_{i,t} \times Type_k + \sum_{k=2}^5 \beta_k Type_k + \delta Volatility_t \\ & + \sum_{m=1}^8 \lambda_m ControlVariable_m + \varepsilon_{i,t}, \end{aligned} \tag{5}$$

where $SprdDecomp_{i,k,t}$ is the effective spread, price impact, or realized spread for trades of type k in stock i on day t less the Nasdaq match; α_i are stock fixed effects; $Type_k$ is a dummy variable indicating the trade type (pure system, pure floor, floor-initiated interaction, system-initiated interaction, or mixed-initiator interaction); $Hybrid_{i,t}$ is an indicator variable taking the value of one if the stock is in Hybrid mode on day t , otherwise zero; and $Volatility_t$ is the opening value of CBOE's VIX index on day t . The type dummy is omitted for pure system trades. Therefore, the coefficients on other type dummies represent differences from pure system trades. The coefficients on the Hybrid dummy variable interacted with the trade type dummy variable measure the change in that type trade following the Hybrid introduction.

Table 5 presents the regression results for effective spreads, five-minute price impacts, and five-minute realized spreads. The first three columns present results for the 40-day window surrounding the Hybrid introduction; the last three columns extend the analysis to the full year surrounding the Hybrid introduction. Effective spreads increase for all trade types with the Hybrid introduction. This shows that the increase in effective spreads occurs for all market participants and is not solely due to floor trading becoming more expensive. The pure system trades increase by an amount similar to the overall increase seen in Table 3. Floor participants appear to be more affected than system participants: Effective spreads for the other type categories increase more than pure system trades, and this difference persists over the full-year analysis as well as in the period immediately surrounding Hybrid introduction.

[Table 5 Here]

Over the 40-day period surrounding the Hybrid introduction, price impact increases by nearly two basis points for pure floor trades. Prior to Hybrid pure floor trades have significantly less price impact than pure system trades, consistent with repeated interaction leading to cooperation among floor participants, attenuating adverse selection. After Hybrid, pure floor trades have more adverse selection than pure system trades, suggesting a breakdown of cooperation on the floor. The price impact of system-initiated interaction trades increases by more than two basis points. This suggests that Hybrid makes it more difficult for floor participants to avoid electronically-arriving informed order flow. Losing this ability could explain the continued decline of floor activity. Meanwhile the price impact of floor-initiated interaction trades declines. This may be due to floor participants' inability to utilize system latency to use off-floor limit orders as free trading options. Alternatively, Hybrid may enable informed traders to get better execution using electronic orders, so they use floor brokers less.

The 40-day window realized spread analysis provides evidence consistent with floor participants' being less able to profit. The over-two-basis-point increase in realized spreads for floor-initiated interaction trades translates into a rise in profitability for the system orders that provide liquidity. In contrast, system-initiated interaction trades become less profitable for floor-based liquidity providers.

Using the realized spread as an *ex post* estimate of profitability suggests that Hybrid shifts the balance between floor participants and system participants of the profitability for liquidity demand and supply.

Extending the analysis to the one-year period surrounding Hybrid's introduction (last three columns of Table 5) produces coefficients that are generally of similar magnitude and statistical significance. The interesting differences are price impacts and realized spreads for pure system trades. In the 40-day sample, the Hybrid coefficients for realized spreads are significantly positive and for price impact are positive but insignificant. These suggest that over the shorter horizon around Hybrid's introduction the increase in effective spreads is due to greater profits for limit orders supplying liquidity. Over the one-year sample, Hybrid introduction leads to an increase in the price impact of pure system trades of 0.76 basis points with a *t*-statistic of 7.1. Hybrid's introduction leads to lower realized spreads for pure system trades, but the decline is not significant. Thus, over the longer term Hybrid leads to greater adverse selection for pure system trades, pure floor trades, and system-initiated floor and system interaction trades. This is consistent with the decline of cooperation on the floor and informed traders preferring anonymous electronic execution (as in Barclay, Hendershott, and McCormick (2003)).

The analysis thus far establishes that Hybrid increases trading costs and that trading cost changes arise both from increases in costs for trades that involve only system orders and from floor participants' losing their advantages and cooperating less.

6. Adverse Selection Changes in Hybrid

The analysis of adverse selection in the previous section is limited to the simple price impact measure because there are too few trades in many of the who-trades-with-whom categories to allow estimation of adverse selection measures based on consecutive trades. In this section we expand our analysis to examine how adverse selection changed in the Hybrid market using the Hasbrouck (1991a, 1991b) impulse response measure, which considers persistence in order flow, as well as the simple five-minute price impact measure described in Section 5. (See appendix for details on the Hasbrouck decomposition.)

We also examine the five-minute realized spread measure, as in Section 6. We use the same panel data approach as in Section 4. For each spread decomposition measure we run the following regression:

$$SprdDecomp_{i,t} = \alpha_i + \beta Hybrid_{i,t} + \gamma Volatility_t + \sum_{k=1}^8 \delta_k ControlVariable_{i,k} + \varepsilon_{i,t}, \quad (6)$$

where $SprdDecomp_{i,t}$ is the average impulse response, price impact, or realized spread measure for stock i less its Nasdaq match on day t ; α_i are stock fixed effects; $Hybrid_{i,t}$ is an indicator variable taking the value of one if the stock is in Hybrid mode on day t , otherwise zero; $Volatility_t$ is the opening value of CBOE's VIX index on day t ; and $ControlVariable_{i,k}$ are eight stock-level control variables: the log of the market capitalization and the log of the stock price on March 31, 2006, and the dollar turnover and return volatility in the first quarter of 2006, for each NYSE and Nasdaq matched stock.

[Table 6 Here]

Table 6 shows the results from estimating the regression over both the 40-day window (first three columns) and the full year surrounding the Hybrid introduction (last three columns). Analyses at both horizons show that adverse selection generally rose with the Hybrid introduction. The Hasbrouck impulse response measure, which reflects order persistence out to ten lags, captures a significant increase in adverse selection over both horizons. The simple five-minute price impact shows no significant change over the 40-day window, but reveals larger increases that are significant in the full-year analysis. This is consistent with the finding in Table 5 that pure system trades (which account for over 70% of trades, see Table 4) experience a significant increase in price impact over the longer period. The 0.61 basis point increase in price impact for the full sample in Table 6 is comparable to the 0.76 basis point increase in price impact for pure system trades in Table 5. Conversely, we find that realized spreads increase in the period immediately surrounding the Hybrid introduction but are insignificant over the long run. Together these findings support the idea that the initial increase in effective spreads is due to higher profits for liquidity suppliers in the short run, although in the longer run the increase is attributable to higher adverse selection.

7. Market Efficiency

The analysis up to this point shows that although Hybrid dramatically increased execution speed on the NYSE, it also led to higher transaction costs as adverse selection rose. The increase in the execution speed and informativeness of trades could also increase market efficiency. We test this by examining the five-minute trading range, five-minute quote return volatility, and five-minute/30-minute variance ratio. The trading range is the five-minute high minus low traded price divided by the last traded price in each non-overlapping five-minute interval, averaged over the trading day for each stock each day. The quote return volatility is the standard deviation of midquote returns in all non-overlapping five-minute periods of the day, calculated for each stock each day. The trading range focuses on the most extreme price movements, perhaps making it better suited for detecting high-frequency transitory volatility, but because it is based on transaction prices it also reflects microstructure noise such as bid-ask bounce. The use of quote midpoints in the quote return volatility provides a measure that is not affected by bid-ask bounce, although it may miss very high frequency volatility. The five-minute/30-minute variance ratio is six times the five-minute variance of midquote returns divided by the 30-minute variance of midquote returns, calculated for each stock each day.¹⁶ The variance ratio evaluates whether short-term price changes are reversed on average. Such reversals, if they exist, would indicate that order flow or other shocks over short horizons push prices away from their longer term equilibrium level. Variance ratios are typically greater than one, indicating some excess volatility over very short horizons, so a decline in the variance ratio would indicate an increase in market efficiency.

Figure 7 presents the three measures for the NYSE stocks minus their matched Nasdaq stocks over the period from four months before to four months after each stock's Hybrid activation. Both the trading range and quote volatility measures increase around the Hybrid activation and remain higher on average thereafter. This increase in intraday volatility could simply reflect increased information flows, consistent with the findings of higher adverse selection following Hybrid activation in Section 6. The variance ratio

¹⁶ To reduce the noisiness of variance ratios measured at the stock/day level, we winsorize stock/day variance ratios at the first and 99th percentiles before calculating NYSE – Nasdaq differences. Using the unwinsorized variance ratio series yields qualitatively similar results, which are available on request.

analysis supports this interpretation, as the ratio between five-minute and 30-minute midquote return variances falls following Hybrid implementation, suggesting that NYSE prices became more efficient following the Hybrid implementation.

[Figure 7 Here]

To more precisely examine the relation between the Hybrid introduction and market efficiency, we run the following regressions, using the same panel data approach as before:

$$MktEff_{i,t} = \alpha_i + \beta Hybrid + \gamma Volatility + \sum_{k=1}^8 \delta_k ControlVariable_{i,k} + \varepsilon_{i,t} \quad (7)$$

where $MktEff_{i,t}$ is the average five-minute trading range, five-minute quote volatility, or five-minute/30-minute variance ratio for stock i less its Nasdaq match on day t ; α_i are stock fixed effects; $Hybrid_{i,t}$ is an indicator variable taking the value of one if the stock is in Hybrid mode on day t , otherwise zero; $Volatility_t$ is the opening value of CBOE's VIX index on day t ; and $ControlVariable_{i,k}$ are eight stock-level control variables: the log of the market capitalization and the log of the stock price on March 31, 2006, and the dollar turnover and return volatility in the first quarter of 2006, for each NYSE and Nasdaq matched stock.

[Table 7 Here]

Table 7 presents the results from estimating Equation (7) over both the 40-day window (first three columns) and the full year surrounding the Hybrid introduction (last three columns). Analyses at both horizons show that the five-minute trading range generally rose with the Hybrid introduction. The change in the trading range is larger than the change in quoted or effective spreads in Table 3 (1.25 basis points for the trading range over the 40-day window, versus 0.80 and 0.53 for quoted and effective spreads in Table 3), so the wider trading range is not solely due to the wider spreads following Hybrid. The five-minute quote volatility does not show robust changes with Hybrid's introduction over the 40-day window, but does show a significant increase in the largest stock quartiles over the full-year window. Similarly, the variance ratio demonstrates no significant change in the 40-day window, but a significant decline of 0.11 relative to Nasdaq, suggesting an improvement in market efficiency, in the full-year window.

8. Conclusion

The NYSE trading floor is one of the most famous and studied markets in the world. In this paper we examine the impact of the NYSE's introduction of its Hybrid market, an event that speeds up electronic trading and represents a major blow to trading on the NYSE floor. We show that the Hybrid market raises trading costs by increasing adverse selection and reduces cooperation on the trading floor. However, the trading process is faster and more information is incorporated into prices more efficiently. The Hybrid market represents a new position for the NYSE on the tradeoff between these different dimensions of market quality. The Hybrid experience suggests that it may be difficult to capture the benefits of both the repeated interaction that in-person trading offers and the speed and efficiency at which electronic systems operate.

For some market participants, e.g., institutional traders working large orders via electronic limit orders or implementing complex strategies, the gains in speed may be more valuable than the wider spreads. For small retail traders speed may be unimportant, in which case they are worse off under the new regime. An aggregate welfare analysis is difficult when there are changes along multiple dimensions of quality and traders have heterogeneous preferences.

The SEC intends to enhance competition between markets with Reg NMS. But by allowing faster markets to ignore slower markets' quotes, Reg NMS's Order Protection (trade-through) Rule effectively precludes traditional floor trading because human interaction is too slow. The goals of Reg NMS, to "give investors, particularly retail investors, greater confidence that they will be treated fairly when they participate in the equity markets" and to "promote deep and stable markets that minimize investor transaction costs" (SEC (2005)) are laudable. However, the Hybrid experience suggests that while investors who favor faster executions benefit, investors who are more concerned with lower execution costs than speed may be worse off in a world without floor trading.

An interesting question is the extent to which the benefits of human interaction can be replicated in fully electronic systems. In the context of the NYSE's introduction of the Hybrid market, could the

benefits that derive from repeated interactions and lack of anonymity on the floor have been preserved in an increasingly electronic setting? Non-financial markets such as eBay demonstrate that reputations can be built in electronic markets. Whether such systems can work in financial markets, where the potential gains from manipulation are substantial, is a question of considerable practical and academic significance.

Appendix: Hasbrouck Decomposition

Hasbrouck (1991a, 1991b) introduces a Vector Autoregression (VAR) based model that makes almost no structural assumptions about the nature of information or order flow, but instead infers the nature of information and trading from the observed sequence of prices and orders. In this framework, all stock price moves end up assigned to one of two categories: They are either associated or unassociated with a recent trade. Although the model does not make any structural assumptions about the nature of information, we usually refer to price moves as private-information-based if they are associated with a recent trade. Price moves that are orthogonal to recent trade arrivals are sometimes considered to be based on public information (examples of this interpretation include Jones, Kaul, and Lipson (1994) and Barclay and Hendershott (2003)).

To separate price moves into trade-related and trade-unrelated components, we construct a VAR with two equations: The first equation describes the trade-by-trade evolution of the quote midpoint, while the second equation describes the persistence of order flow. Define q_{jt} to be the buy-sell indicator for trade t in stock j (+1 for buys, -1 for sells), and define r_{jt} to be the log return based on the quote midpoint of stock j from trade $t-1$ to trade t . The VAR picks up order flow dependence out to 10 lags:

$$r_t = \sum_{i=1}^{10} \alpha_i r_{t-i} + \sum_{i=0}^{10} \beta_i q_{t-i} + \varepsilon_{rt},$$

$$q_t = \sum_{i=1}^{10} \gamma_i r_{t-i} + \sum_{i=1}^{10} \phi_i q_{t-i} + \varepsilon_{qt},$$

where the stock subscripts j are suppressed from here on. The VAR is inverted to get the Vector Moving Average (VMA) representation:

$$y_t = \begin{bmatrix} r_t \\ q_t \end{bmatrix} = \theta(L) \varepsilon_t = \begin{bmatrix} a(L) & b(L) \\ c(L) & d(L) \end{bmatrix} \begin{bmatrix} \varepsilon_{rt} \\ \varepsilon_{qt} \end{bmatrix},$$

where $a(L)$, $b(L)$, $c(L)$, and $d(L)$ are lag polynomial operators. The permanent effect on price of an innovation ε_t is given by $a(L)\varepsilon_{rt} + b(L)\varepsilon_{qt}$, and because we include contemporaneous q_t in the return

equation, $\text{cov}(\varepsilon_{rt}, \varepsilon_{qt}) = 0$ and the variance of this random-walk component can be written as:

$$\sigma_w^2 = \left(\sum_{i=0}^{\infty} a_i \right)^2 \sigma_r^2 + \left(\sum_{i=0}^{\infty} b_i \right)^2 \sigma_q^2,$$

where the second term captures the component of price discovery that is related to trade, and the first term captures price changes that are unrelated to trading (sometimes referred to as public information). As discussed in Hasbrouck (1991a, 1991b), this method is robust to price discreteness, lagged adjustment to information, and lagged adjustment to trades. The VAR and the trade-related and non-trade-related standard deviations are estimated for each stock each day. The trade-related component is our impulse response measure in Table 6.

References

- Barclay, Michael, 1997, Bid-ask spreads and the avoidance of odd-eighth quotes on Nasdaq: An examination of exchange listings, *Journal of Financial Economics* 45, 35–58.
- Barclay, Michael, Terrence Hendershott, and D. Timothy McCormick, 2003, Competition among trading venues: Information and trading on electronic communication networks, *Journal of Finance* 58, 2637-2665.
- Barclay, Michael and Terrence Hendershott, 2003, Price Discovery and Trading After Hours, *Review of Financial Studies* 16, 1041-1073.
- Battalio, Robert, Andrew Ellul, and Robert Jennings, 2007, Reputation effects in trading on the New York Stock Exchange, *Journal of Finance* 62, 1243-1271.
- Battalio, Robert, Brian Hatch, and Robert Jennings, 2003, All else equal? A multi-dimensional analysis of retail market order execution quality, *Journal of Financial Markets* 6, 143–162.
- Benveniste, Lawrence, Alan Marcus, and William Wilhelm, 1992, What’s so special about the specialist? *Journal of Financial Economics* 32, 61-86.
- Bessembinder, Hendrik, 2003, Issues in assessing trade execution costs, *Journal of Financial Markets* 6, 233-257.
- Bessembinder, Hendrik, and Herbert Kaufman, 1997, A cross-exchange comparison of execution costs and information flow for NYSE-listed stocks, *Journal of Financial Economics* 46, 293-319.
- Boehmer, Ekkehart, 2005, Dimensions of execution quality: Recent evidence for U.S. equity markets. *Journal of Financial Economics* 78, 463-704.
- Boehmer, Ekkehart, Robert Jennings, and Li Wei, 2007, Public disclosure and private decisions: Equity market execution quality and order routing, *Review of Financial Studies* 20, 315-358.
- Boehmer, Ekkehart, Gideon Saar, and Lei Yu, 2005, Lifting the veil: An analysis of pre-trade transparency at the NYSE, *Journal of Finance* 60, 783-815.
- Chan, Yuk-Shee and Mark Weinstein, 1993, Bid-ask spread and market structure, *Financial Analysts Journal*, 57-62.
- Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam, 2007, Liquidity and market efficiency, *Journal of Financial Economics*, forthcoming.
- Christie, William, Jeffery Harris, and Paul Schultz, 1994, Why did Nasdaq market makers stop avoiding odd-eighth quotes? *Journal of Finance* 49, 1841–1860.
- Christie, William, and Paul Schultz, 1994, Why do Nasdaq market makers avoid odd-eighth quotes? *Journal of Finance* 49, 1813–1840.
- Davies, Ryan, and Sang Soo Kim, 2007, Using matched samples to test for differences in trade execution costs, *Journal of Financial Markets*, forthcoming.
- Dowell, Andrew, 2007, NYSE’s floor suffers pair of defections, *Wall Street Journal*, November 16, C2.
- Handa, Puneet, Robert Schwartz and Ashish Tiwari, 2004, The economic value of a trading floor: Evidence from the American Stock Exchange, *Journal of Business* 77, 331-355.
- Harris, Lawrence, and Joel Hasbrouck, 1996, Market vs. limit orders: The SuperDOT evidence on order submission strategy, *Journal of Financial and Quantitative Analysis* 31, 213-231.
- Hasbrouck, Joel, 1991a, Measuring the information content of stock trades, *Journal of Finance* 46, 179-207.

- Hasbrouck, Joel, 1991b, The summary informativeness of stock trades: An econometric analysis, *Review of Financial Studies* 4, 571-595.
- Hasbrouck, Joel, and George Sofianos, 1993, The trades of market-makers: An analysis of NYSE specialists, *Journal of Finance* 48, 1565-1594.
- Huang, Roger, and Hans Stoll, 1996, Dealer versus auction markets: a paired comparison of execution costs on Nasdaq and the NYSE, *Journal of Financial Economics* 41, 313-357.
- Ip, Greg, and Susanne Craig, 2003, Trading cases: NYSE's 'Specialist' probe puts precious asset at risk, *Wall Street Journal*, April 18, A1.
- Jain, Pankaj, 2005, Financial market design and the equity premium: Electronic versus floor trading, *Journal of Finance* 60, 2955-2985.
- Jones, Charles M., Gautam Kaul, and Marc L. Lipson. 1994. "Information, Trading, and Volatility." *Journal of Financial Economics* 36, 127-154.
- Lee, Charles, 1993, Market integration and price execution for NYSE-listed securities, *Journal of Finance* 48, 1009-1038.
- Lee, Charles, and Mark Ready, 1991, Inferring trade direction from intraday data, *Journal of Finance* 46, 733-747.
- Lucchetti, Aaron, 2007, The NYSE: Faster (and Lonelier), *Wall Street Journal*, January 24, C1.
- Madhavan, Ananth, and Minder Cheng, 1997, In search of liquidity: Block trades in the upstairs and downstairs markets, *Review of Financial Studies* 10, 175-203.
- Madhavan, Ananth, and Seymour Smidt, 1993, An analysis of daily changes in specialist inventories and quotations, *Journal of Finance* 48, 1595-1628.
- McGeehan, Patrick, 2007, Next to downsize on Wall Street? The Exchange floor, *New York Times*, September 23, 37.
- Moulton, Pamela, 2006, Who trades with whom? Working paper.
- NYSE Group, 2006a, Hybrid market training program, Booklet available online at http://www.nyse.com/pdfs/hm_booklet.pdf.
- NYSE Group, 2006b, NYSE Hybrid FAQ, Document available online at <http://www.nyse.com/productservices/nyseequities/1126821290257.html>.
- NYSE Group, 2007, NYSE completes Hybrid market phase III activation, Press release available online at <http://www.nyse.com/pdfs/HybridPhaseIV1.24.07.pdf>.
- Ready, Mark, 1999, The specialist's discretion: stopped orders and price improvement, *Review of Financial Studies*, 12, 1075-1112.
- Rock, Kevin, 1990, The specialist's order book and price anomalies, Working paper.
- Securities and Exchange Commission, 2005, Regulation NMS, available online at <http://www.sec.gov/rules/final/34-51808.pdf>.
- Sofianos, George, and Ingrid Werner, 2000, The trades of NYSE floor brokers, *Journal of Financial Markets* 3, 139-176.
- Stoll, Hans, and Christoph Schenzler, 2006, Trades outside the quotes: reporting delay, trading option or trade size? *Journal of Financial Economics* 79, 615-653.
- Thompson, Samuel, 2006, Simple formulas for standard errors that cluster by both firm and time, Working paper.

Werner, Ingrid, 2003, NYSE order flow, spreads, and information, *Journal of Financial Markets* 6, 309-335.

Table 1: Sample Descriptive Statistics

Descriptive statistics are presented for the sample of 400 NYSE stocks and 400 Nasdaq stocks matched on market capitalization and price. Market capitalization and closing price are from CRSP as of March 31, 2006. Quoted spread (*QSpread*) and effective spread (*ESpread*) are calculated from TAQ data and averaged for each stock over the period January through March 2006. Mean, median, and standard deviation are calculated across 400 stocks in the full sample, 100 stocks in each quartile, with Quartile 1 comprising the largest stocks.

	NYSE Sample				Nasdaq Sample			
	Market Cap. (\$ mil)	Closing Price (\$)	QSpread (bps)	ESpread (bps)	Market Cap. (\$ mil)	Closing Price (\$)	QSpread (bps)	ESpread (bps)
Full Sample								
Mean	9,258	41.51	9.2	6.4	6,426	37.72	10.7	9.6
Median	3,053	36.84	7.6	5.3	2,008	33.91	8.7	7.8
Std. Deviation	19,522	26.24	6.5	4.3	19,115	27.53	9.9	8.3
Quartile 1								
Mean	29,035	51.41	4.7	3.3	18,895	39.38	6.5	6.7
Median	17,735	49.82	4.3	3.1	8,802	33.75	4.7	5.0
Std. Deviation	31,613	21.93	1.6	1.2	35,433	41.04	4.6	4.3
Quartile 2								
Mean	4,813	44.16	7.5	5.2	3,789	41.84	8.4	7.7
Median	4,473	40.32	6.4	4.4	3,606	37.55	6.8	6.6
Std. Deviation	1,423	28.29	7.2	4.4	1,882	21.29	8.9	7.2
Quartile 3								
Mean	2,132	40.87	10.3	7.2	1,963	40.07	11.0	9.6
Median	2,078	36.19	8.9	6.2	1,822	36.01	9.1	8.1
Std. Deviation	469	32.45	5.6	4.0	478	25.19	7.2	5.3
Quartile 4								
Mean	1,053	29.59	14.2	9.8	1,059	29.58	16.8	14.3
Median	1,007	27.99	12.5	8.7	1,023	28.35	13.0	11.3
Std. Deviation	234	13.78	5.9	4.0	234	13.74	13.7	11.9

Table 2: Market Quality Measures Pre and Post Hybrid

Averages are calculated for each stock over the 20 days immediately preceding hybrid activation (*Pre-Hybrid*) and 20 days immediately following Hybrid activation (*Post-Hybrid*). Cross-sectional means and mean differences between NYSE stocks and their Nasdaq matches (labeled "- Match") are presented for the full sample of 400 stocks and by quartile, with Quartile 1 comprising the largest stocks. *QSpread* is the quoted spread; *ESpread* is the effective spread; *QVolatility* is the 5-minute midquote return volatility; *TRange* is the 5-minute trading range volatility; *Trade size* is the average trade size; *QDepth* is the average quoted depth at the best bid and ask; *Volume* is daily dollar volume. All measures are calculated from TAQ data. Significance levels of mean changes are from univariate t-tests; ** (*) denotes significance at the 1% (5%) level.

	QSpread	QSpread -	ESpread	ESpread -	Trade Size	QDepth	Volume
	(bps)	Match	(bps)	Match	(shares)	(100s)	(\$ mn)
Full Sample							
Pre-Hybrid	7.9	-0.3	5.6	-1.4	396	16.1	37.1
Post-Hybrid	<u>8.3</u>	<u>0.5</u>	<u>5.9</u>	<u>-0.9</u>	<u>303</u>	<u>16.1</u>	<u>38.3</u>
Change	0.4 **	0.8 **	0.3 **	0.5 **	-93 **	0.0	1.3
Quartile 1							
Pre-Hybrid	3.9	-1.6	2.8	-2.7	609	30.7	102.3
Post-Hybrid	<u>4.1</u>	<u>-1.2</u>	<u>3.1</u>	<u>-2.3</u>	<u>414</u>	<u>29.5</u>	<u>106.0</u>
Change	0.2 **	0.4 **	0.3 **	0.4 **	-195 **	-1.2	3.7
Quartile 2							
Pre-Hybrid	6.4	-0.2	4.5	-1.4	386	15.2	26.1
Post-Hybrid	<u>6.6</u>	<u>0.3</u>	<u>4.7</u>	<u>-1.0</u>	<u>302</u>	<u>15.8</u>	<u>26.5</u>
Change	0.2	0.5 *	0.1	0.4 *	-84 **	0.6	0.4
Quartile 3							
Pre-Hybrid	8.8	0.1	6.3	-1.0	305	9.5	12.2
Post-Hybrid	<u>9.3</u>	<u>1.0</u>	<u>6.5</u>	<u>-0.6</u>	<u>256</u>	<u>9.6</u>	<u>12.7</u>
Change	0.5 **	0.9 **	0.3 **	0.5 **	-49 **	0.1	0.4
Quartile 4							
Pre-Hybrid	12.3	0.6	8.7	-0.6	286	9.0	7.7
Post-Hybrid	<u>13.0</u>	<u>2.0</u>	<u>9.1</u>	<u>0.2</u>	<u>241</u>	<u>9.6</u>	<u>8.2</u>
Change	0.7 **	1.4 **	0.5 **	0.8 **	-45 **	0.6	0.5

Table 3: Panel Regressions of Spreads on Hybrid

Analysis periods are the 40-day window surrounding each stock's Hybrid activation (first four columns) and June 1, 2006 to May 31, 2007 (last four columns). Quoted spread (QSpread) and effective spread (ESpread) differences between NYSE stocks and their matched Nasdaq stocks are regressed on a dummy variable set equal to one if the stock has been activated in Hybrid (*Hybrid*), daily market *Volatility* as measured by the VIX index, and the following control variables for each stock: log(market capitalization), log(price), dollar turnover, and return volatility. Coefficients for control variables, constant, and stock fixed effects are not reported. All dependent variables are in basis points. Full sample is 400 stocks; Quartile 1 comprises the largest stocks. T-statistics, reported in parentheses below coefficient estimates, are calculated using standard errors that are double-clustered on date and stock.

Dependent Variable	40-day Window				June 1, 2006 to May 31, 2007				
	QSpread - Match		ESpread - Match		QSpread - Match		ESpread - Match		
Full Sample									
Hybrid	0.85 (5.2)	0.80 (6.5)	0.58 (4.2)	0.53 (5.6)	0.53 (2.6)	0.61 (3.1)	0.60 (3.8)	0.66 (4.2)	
Volatility	-0.11 (-0.7)	-0.02 (-0.3)	-0.11 (-0.8)	-0.03 (-0.6)	0.02 (1.0)	0.02 (1.2)	-0.01 (-0.9)	-0.01 (-0.7)	
Fixed Effects	no	yes	no	yes	no	yes	no	yes	
Observations	16,000	16,000	16,000	16,000	100,400	100,400	100,400	100,400	
Adj. R ²	0.17	0.76	0.15	0.72	0.15	0.63	0.14	0.58	
Quartile 1									
Hybrid	0.61 (3.4)	0.43 (5.3)	0.61 (3.4)	0.43 (6.9)	0.27 (1.0)	0.34 (1.3)	0.67 (2.6)	0.75 (2.8)	
Volatility	-0.22 (-1.3)	0.00 (-0.1)	-0.25 (-1.4)	-0.04 (-0.9)	-0.04 (-1.8)	-0.03 (-1.6)	-0.06 (-3.0)	-0.06 (-2.9)	
Fixed Effects	no	yes	no	yes	no	yes	no	yes	
Observations	4,000	4,000	4,000	4,000	25,100	25,100	25,100	25,100	
Adj. R ²	0.48	0.93	0.51	0.94	0.51	0.74	0.54	0.73	
Quartile 2									
Hybrid	0.51 (2.4)	0.53 (2.2)	0.38 (2.4)	0.39 (2.2)	0.01 (0.0)	0.27 (0.8)	0.36 (2.1)	0.54 (2.4)	
Volatility	-0.45 (-1.1)	-0.19 (-1.4)	-0.37 (-1.1)	-0.17 (-1.5)	0.00 (0.0)	0.01 (0.3)	-0.02 (-0.5)	-0.01 (-0.4)	
Fixed Effects	no	yes	no	yes	no	yes	no	yes	
Observations	4,000	4,000	4,000	4,000	25,100	25,100	25,100	25,100	
Adj. R ²	0.25	0.81	0.19	0.79	0.22	0.71	0.18	0.70	
Quartile 3									
Hybrid	0.87 (5.5)	0.88 (4.9)	0.47 (3.7)	0.47 (3.4)	0.39 (1.5)	0.33 (1.3)	0.40 (2.1)	0.33 (1.8)	
Volatility	0.20 (2.0)	0.15 (1.5)	0.16 (1.9)	0.11 (1.5)	0.04 (0.9)	0.04 (0.9)	0.01 (0.3)	0.01 (0.2)	
Fixed Effects	no	yes	no	yes	no	yes	no	yes	
Observations	4,000	4,000	4,000	4,000	25,100	25,100	25,100	25,100	
Adj. R ²	0.23	0.62	0.14	0.55	0.19	0.51	0.12	0.47	
Quartile 4									
Hybrid	1.38 (4.1)	1.36 (4.2)	0.84 (3.0)	0.82 (3.1)	1.60 (2.8)	1.50 (2.6)	1.10 (2.3)	1.03 (2.1)	
Volatility	0.16 (1.6)	-0.02 (-0.2)	0.12 (1.1)	-0.02 (-0.2)	0.08 (1.6)	0.08 (1.5)	0.01 (0.4)	0.01 (0.3)	
Fixed Effects	no	yes	no	yes	no	yes	no	yes	
Observations	4,000	4,000	4,000	4,000	25,100	25,100	25,100	25,100	
Adj. R ²	0.31	0.74	0.27	0.66	0.24	0.58	0.19	0.49	

Table 4: Who Trades with Whom Pre and Post Hybrid

Averages are calculated for each stock over the days preceding hybrid activation (Pre-Hybrid) and following Hybrid activation (Post-Hybrid), within a 40-day window in the first four columns, within the period from June 1, 2006 to May 31, 2007 in the last four columns. Cross-sectional means are presented for the full sample of 400 stocks and by quartile, with Quartile 1 comprising the largest stocks. *Pure System* trades involve only system participants; *Pure Floor* trades involve only floor brokers and/or the specialist; *Floor & System Interaction* trades involve some combination of floor and system participants. Floor and System Interaction Trades are further categorized by which type of participant initiated the trade: *Floor-Initiated*, *System-Initiated*, and *Mixed Initiator*, which are trades in which floor and system participants are on the same side. Who Trades with Whom percentages are calculated for trade type as share volume divided by total volume. Statistics are calculated from CAUD data. Significance levels of mean changes are from univariate t-tests; ** (*) denotes significance at the 1% (5%) level.

	40-day Window					June 1, 2006 to May 31, 2007				
	Pure System	Pure Floor	Floor & System Interaction			Pure System	Pure Floor	Floor & System Interaction		
			Floor-Initiated	System-Initiated	Mixed Initiator			Floor-Initiated	System-Initiated	Mixed Initiator
Full Sample										
Pre-Hybrid	70.5%	1.4%	4.7%	8.4%	15.0%	67.2%	1.7%	5.1%	9.1%	16.9%
Post-Hybrid	<u>78.2%</u>	<u>1.4%</u>	<u>3.5%</u>	<u>10.5%</u>	<u>6.4%</u>	<u>81.4%</u>	<u>1.0%</u>	<u>2.9%</u>	<u>9.7%</u>	<u>5.1%</u>
Change	7.7% **	0.0%	-1.3% **	2.1% **	-8.6% **	14.2% **	-0.7% **	-2.2% **	0.6% **	-11.9% **
Quartile 1										
Pre-Hybrid	67.0%	1.4%	5.1%	7.0%	19.5%	62.7%	1.8%	5.5%	7.6%	22.4%
Post-Hybrid	<u>75.0%</u>	<u>1.6%</u>	<u>4.1%</u>	<u>11.3%</u>	<u>8.0%</u>	<u>78.7%</u>	<u>1.2%</u>	<u>3.5%</u>	<u>10.4%</u>	<u>6.3%</u>
Change	8.0% **	0.2%	-1.0% **	4.3% **	-11.5% **	16.0% **	-0.6% **	-2.0% **	2.7% **	-16.1% **
Quartile 2										
Pre-Hybrid	70.7%	1.3%	4.8%	8.1%	15.0%	67.0%	1.8%	5.1%	8.9%	17.2%
Post-Hybrid	<u>78.6%</u>	<u>1.4%</u>	<u>3.3%</u>	<u>10.4%</u>	<u>6.3%</u>	<u>81.7%</u>	<u>1.0%</u>	<u>2.8%</u>	<u>9.5%</u>	<u>5.1%</u>
Change	7.9% **	0.0%	-1.6% **	2.3% **	-8.7% **	14.6% **	-0.8% **	-2.3% **	0.6%	-12.1% **
Quartile 3										
Pre-Hybrid	72.3%	1.5%	4.5%	8.8%	12.9%	69.5%	1.7%	4.7%	9.5%	14.6%
Post-Hybrid	<u>79.8%</u>	<u>1.3%</u>	<u>3.2%</u>	<u>9.9%</u>	<u>5.8%</u>	<u>82.3%</u>	<u>1.0%</u>	<u>2.6%</u>	<u>9.4%</u>	<u>4.6%</u>
Change	7.4% **	-0.1%	-1.3% **	1.1% **	-7.1% **	12.8% **	-0.7% **	-2.1% **	-0.1%	-10.0% **
Quartile 4										
Pre-Hybrid	71.9%	1.3%	4.6%	9.7%	12.6%	69.5%	1.7%	5.0%	10.2%	13.6%
Post-Hybrid	<u>79.4%</u>	<u>1.3%</u>	<u>3.4%</u>	<u>10.4%</u>	<u>5.6%</u>	<u>82.8%</u>	<u>0.9%</u>	<u>2.7%</u>	<u>9.4%</u>	<u>4.2%</u>
Change	7.5% **	-0.1%	-1.2% **	0.8%	-7.0% **	13.2% **	-0.7% **	-2.4% **	-0.8% *	-9.3% **

Table 5: Panel Regressions of Spreads on Hybrid by Who Trades with Whom Type

Analysis periods are the 40-day window surrounding each stock's Hybrid activation (first three columns) and June 1, 2006 to May 31, 2007 (last three columns). Effective spread, 5-minute price impact, and 5-minute realized spread differences between NYSE stocks and their matched Nasdaq stocks are regressed on dummy variables set equal to one for each of the five who-trades-with-whom categories (Pure System, Pure Floor, Floor-Initiated Interaction, System-Initiated Interaction, and Mixed Initiator Interaction), who-trades-with-whom category variables times a dummy variable equal to one for stocks that have been activated in Hybrid, volatility as measured by the VIX index, and the following control variables for each stock: log(market capitalization), log(price), dollar turnover, and return volatility. Coefficients for control variables and stock fixed effects are not reported. All dependent variables are in basis points. Full sample is 400 stocks; Quartile 1 comprises the largest stocks. T-statistics, reported in parentheses below coefficient estimates, are calculated using standard errors that are double-clustered on date and stock.

Dependent Variable	40-day Window			June 1, 2006 to May 31, 2007		
	ESpread - Match	PImpact - Match	RSspread - Match	ESpread - Match	PImpact - Match	RSspread - Match
Hybrid x Pure System	0.51 (5.1)	0.10 (1.2)	0.40 (4.3)	0.68 (4.4)	0.76 (7.1)	-0.11 (-0.8)
Hybrid x Pure Floor	1.36 (8.1)	1.95 (3.9)	-0.62 (-1.2)	1.85 (9.5)	2.34 (9.3)	-0.49 (-2.0)
Hybrid x Floor-Initiated Interaction	1.60 (12.4)	-0.70 (-2.5)	2.25 (7.2)	2.05 (11.8)	-0.14 (-0.6)	2.14 (8.9)
Hybrid x System-Initiated Interaction	0.95 (6.5)	2.24 (9.0)	-1.30 (-5.6)	1.17 (6.9)	3.56 (20.2)	-2.43 (-12.3)
Hybrid x Mixed Initiator Interaction	0.74 (5.5)	1.11 (4.7)	-0.40 (-1.7)	0.91 (5.4)	1.12 (6.0)	-0.25 (-1.4)
Pure Floor	-0.13 (-1.2)	-0.88 (-2.5)	0.72 (2.1)	-0.10 (-1.5)	-1.21 (-8.8)	1.05 (7.3)
Floor-Initiated Interaction	-0.62 (-9.1)	-0.85 (-3.9)	0.23 (1.0)	-0.56 (-7.8)	-0.74 (-5.1)	0.17 (1.0)
System-Initiated Interaction	-0.12 (-1.8)	-2.07 (-11.0)	1.95 (11.5)	-0.05 (-0.8)	-2.70 (-22.2)	2.64 (21.9)
Mixed Initiator Interaction	0.93 (12.8)	0.83 (6.7)	0.10 (0.6)	1.00 (15.3)	1.08 (13.3)	-0.08 (-0.8)
Observations	74,546	74,546	74,546	465,283	465,283	465,283
Adj. R ²	0.42	0.04	0.05	0.35	0.03	0.04

Table 6: Panel Regressions of Adverse Selection and Liquidity Provider Revenues

Analysis periods are the 40-day window surrounding each stock's Hybrid activation (first three columns) and June 1, 2006 to May 31, 2007 (last three columns). Hasbrouck impulse response, five-minute price impact, and five-minute realized spread differences between NYSE stocks and their matched Nasdaq stocks are regressed on a dummy variable set equal to one if the stock has been activated in Hybrid (*Hybrid*), daily market *Volatility* as measured by the VIX index, and the following control variables for each stock: log(market capitalization), log(price), dollar turnover, and return volatility. Coefficients for control variables, constant, and stock fixed effects are not reported. All dependent variables are in basis points. Full sample is 400 stocks; Quartile 1 comprises the largest stocks. T-statistics, reported in parentheses below coefficient estimates, are calculated using standard errors that are double-clustered on date and stock.

Dependent Variable	40-day Window			June 1, 2006 to May 31, 2007		
	Impulse	PImpact	RSspread	Impulse	PImpact	RSspread
	- Match	- Match	- Match	- Match	- Match	- Match
Full Sample						
Hybrid	9.48 (10.0)	0.16 (1.8)	0.37 (3.8)	10.04 (6.7)	0.61 (6.1)	0.06 (0.5)
Volatility	-0.23 (-0.4)	0.12 (1.5)	-0.15 (-1.9)	0.26 (0.8)	-0.05 (-2.6)	0.03 (1.7)
Observations	16,000	16,000	16,000	100,400	100,400	100,400
Adj. R ²	0.50	0.27	0.31	0.39	0.18	0.26
Quartile 1						
Hybrid	6.50 (6.0)	0.20 (1.9)	0.23 (2.1)	3.99 (2.3)	0.62 (3.2)	0.13 (0.9)
Volatility	0.37 (0.6)	0.10 (1.8)	-0.14 (-2.0)	1.16 (3.6)	-0.08 (-3.8)	0.02 (1.6)
Observations	4,000	4,000	4,000	25,100	25,100	25,100
Adj. R ²	0.68	0.46	0.37	0.51	0.37	0.16
Quartile 2						
Hybrid	7.13 (5.2)	0.14 (1.0)	0.25 (1.6)	7.70 (4.0)	0.45 (3.5)	0.09 (0.4)
Volatility	-1.83 (-2.0)	0.02 (0.2)	-0.19 (-1.8)	0.21 (0.6)	-0.03 (-1.5)	0.02 (0.5)
Observations	4,000	4,000	4,000	25,100	25,100	25,100
Adj. R ²	0.54	0.31	0.41	0.39	0.18	0.34
Quartile 3						
Hybrid	12.01 (7.5)	0.23 (1.5)	0.24 (1.7)	12.73 (5.6)	0.59 (3.9)	-0.26 (-1.8)
Volatility	-0.28 (-0.4)	0.21 (1.9)	-0.09 (-0.8)	-0.24 (-0.6)	-0.04 (-1.7)	0.05 (1.3)
Observations	4,000	4,000	4,000	25,100	25,100	25,100
Adj. R ²	0.51	0.19	0.26	0.41	0.14	0.28
Quartile 4						
Hybrid	12.37 (6.1)	0.08 (0.4)	0.74 (2.9)	16.06 (6.7)	0.76 (3.4)	0.27 (0.8)
Volatility	0.96 (1.1)	0.14 (0.8)	-0.16 (-1.0)	-0.18 (-0.5)	-0.04 (-1.1)	0.05 (1.3)
Observations	4,000	4,000	4,000	25,100	25,100	25,100
Adj. R ²	0.36	0.20	0.26	0.26	0.13	0.20

Table 7: Panel Regressions of Intraday Volatility and Variance Ratio on Hybrid

Analysis periods are the 40-day window surrounding each stock's Hybrid activation (first three columns) and June 1, 2006 to May 31, 2007 (last three columns). Five-minute trading range (TRange), five-minute quote volatility (QVolatility), and five-minute/30-minute variance ratio (VRatio) differences between NYSE stocks and their matched Nasdaq stocks are regressed on a dummy variable set equal to one if the stock has been activated in Hybrid (*Hybrid*), daily market *Volatility* as measured by the VIX index, and the following control variables for each stock: log(market capitalization), log(price), dollar turnover, and return volatility. Coefficients for control variables, constant, and stock fixed effects are not reported. All dependent variables are in basis points. Full sample is 400 stocks; Quartile 1 comprises the largest stocks. T-statistics, reported in parentheses below coefficient estimates, are calculated using standard errors that are double-clustered on date and stock.

Dependent Variable	40-day Window			June 1, 2006 to May 31, 2007		
	TRange	QVolatility	VRatio	TRange	QVolatility	VRatio
	- Match	- Match	- Match	- Match	- Match	- Match
Full Sample						
Hybrid	1.25 (6.4)	0.28 (0.8)	0.01 (0.3)	2.78 (9.3)	0.86 (3.7)	-0.11 (-6.0)
Volatility	-0.47 (-2.3)	-0.15 (-0.7)	0.04 (1.5)	-0.08 (-1.6)	-0.02 (-0.4)	0.01 (2.2)
Observations	16,000	16,000	16,000	100,400	100,400	100,400
Adj. R ²	0.60	0.41	0.04	0.51	0.37	0.02
Quartile 1						
Hybrid	1.42 (3.9)	0.11 (0.4)	0.02 (0.4)	3.95 (7.0)	1.66 (5.0)	-0.05 (-1.7)
Volatility	-0.69 (-2.3)	-0.46 (-1.8)	0.11 (2.7)	-0.15 (-1.9)	-0.09 (-1.5)	0.00 (-0.5)
Observations	4,000	4,000	4,000	25,100	25,100	25,100
Adj. R ²	0.67	0.47	0.04	0.56	0.43	0.02
Quartile 2						
Hybrid	1.69 (4.4)	0.73 (1.5)	0.07 (1.1)	2.85 (7.0)	0.90 (3.0)	-0.11 (-3.9)
Volatility	-0.41 (-1.5)	-0.03 (-0.1)	0.09 (2.2)	-0.12 (-1.7)	0.01 (0.2)	0.01 (1.0)
Observations	4,000	4,000	4,000	25,100	25,100	25,100
Adj. R ²	0.59	0.41	0.05	0.52	0.37	0.02
Quartile 3						
Hybrid	1.10 (3.1)	0.36 (0.8)	-0.03 (-0.5)	2.38 (3.6)	0.49 (1.0)	-0.13 (-4.2)
Volatility	-0.44 (-2.1)	-0.08 (-0.3)	-0.03 (-0.9)	-0.06 (-0.6)	-0.01 (-0.1)	0.02 (2.9)
Observations	4,000	4,000	4,000	25,100	25,100	25,100
Adj. R ²	0.61	0.38	0.03	0.49	0.36	0.02
Quartile 4						
Hybrid	0.82 (1.9)	-0.08 (-0.1)	-0.03 (-0.5)	1.93 (3.4)	0.36 (0.7)	-0.17 (-4.8)
Volatility	-0.33 (-1.1)	-0.05 (-0.2)	-0.01 (-0.2)	0.02 (0.3)	0.03 (0.5)	0.01 (2.5)
Observations	4,000	4,000	4,000	25,100	25,100	25,100
Adj. R ²	0.54	0.37	0.05	0.44	0.31	0.02

Figure 1: Long-run Floor Activity

This chart graphs the participation of specialists, floor brokers, and the entire floor (specialists plus floor brokers), measured as a percentage of twice total regular-hours trading volume for each stock each day. Daily percentages are equal-weighted averages across all NYSE stocks from January 1999 through May 2006, and the 20-day moving average is presented in the chart. Data in this chart represent all NYSE trading whereas data in the rest of the paper excludes certain types of trades. Data are from the NYSE CAUD file.




Figure 2: Execution Speeds by Order Type and Size

This chart graphs average execution speed in seconds for the NYSE sample over the eight-month window surrounding the hybrid activation date for each stock. The solid lines represent execution speed averaged across all order size categories for market orders (*Mkt*) and marketable limit (*MLimit*) orders. The dashed lines represent execution speed for market orders and marketable limit orders of fewer than 500 shares. Equal-weighted averages are calculated across all 400 stocks. Data are from Dash-5.

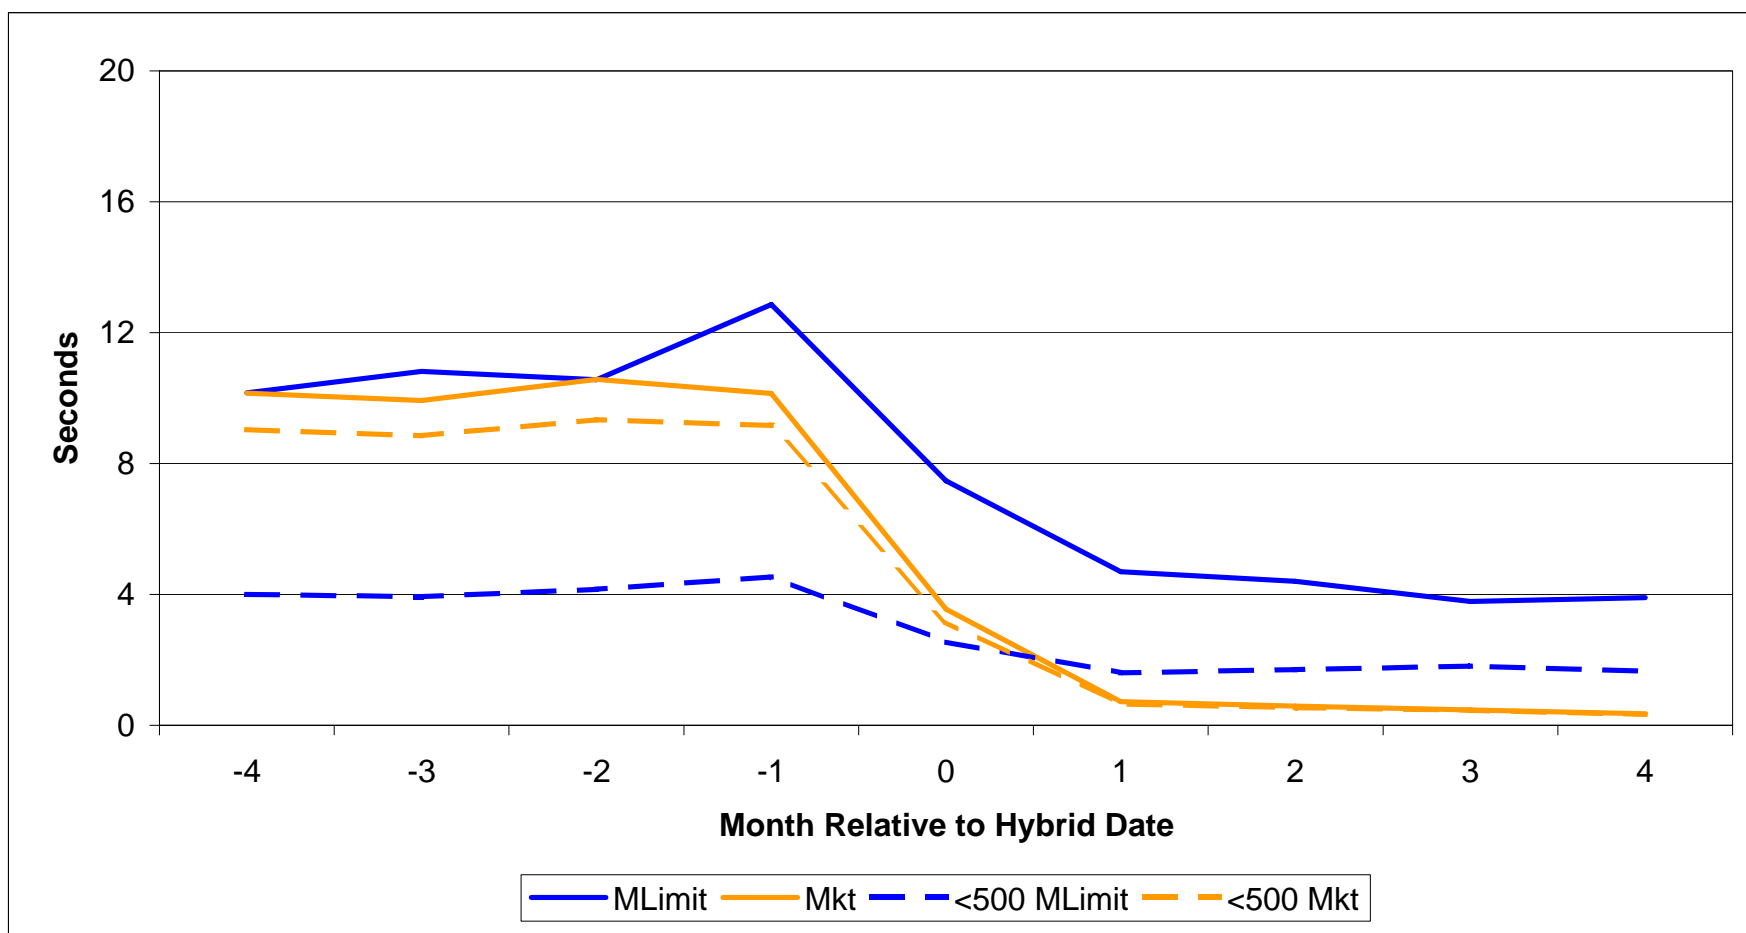


Figure 3: Floor Activity versus Hybrid Activation

This chart graphs the average participation of specialists, floor brokers, and the entire floor (specialists plus floor brokers), measured as a percentage of twice regular-hours trading volume for each stock each day, excluding opening and closing trades. Daily percentages are equal-weighted averages across the sample of 400 NYSE stocks from June 2006 through May 2007. The Hybrid line represents the cumulative percentage of the 400 stocks that have been activated in Hybrid; the box highlights the Hybrid activation period. Data are from the NYSE CAUD file.

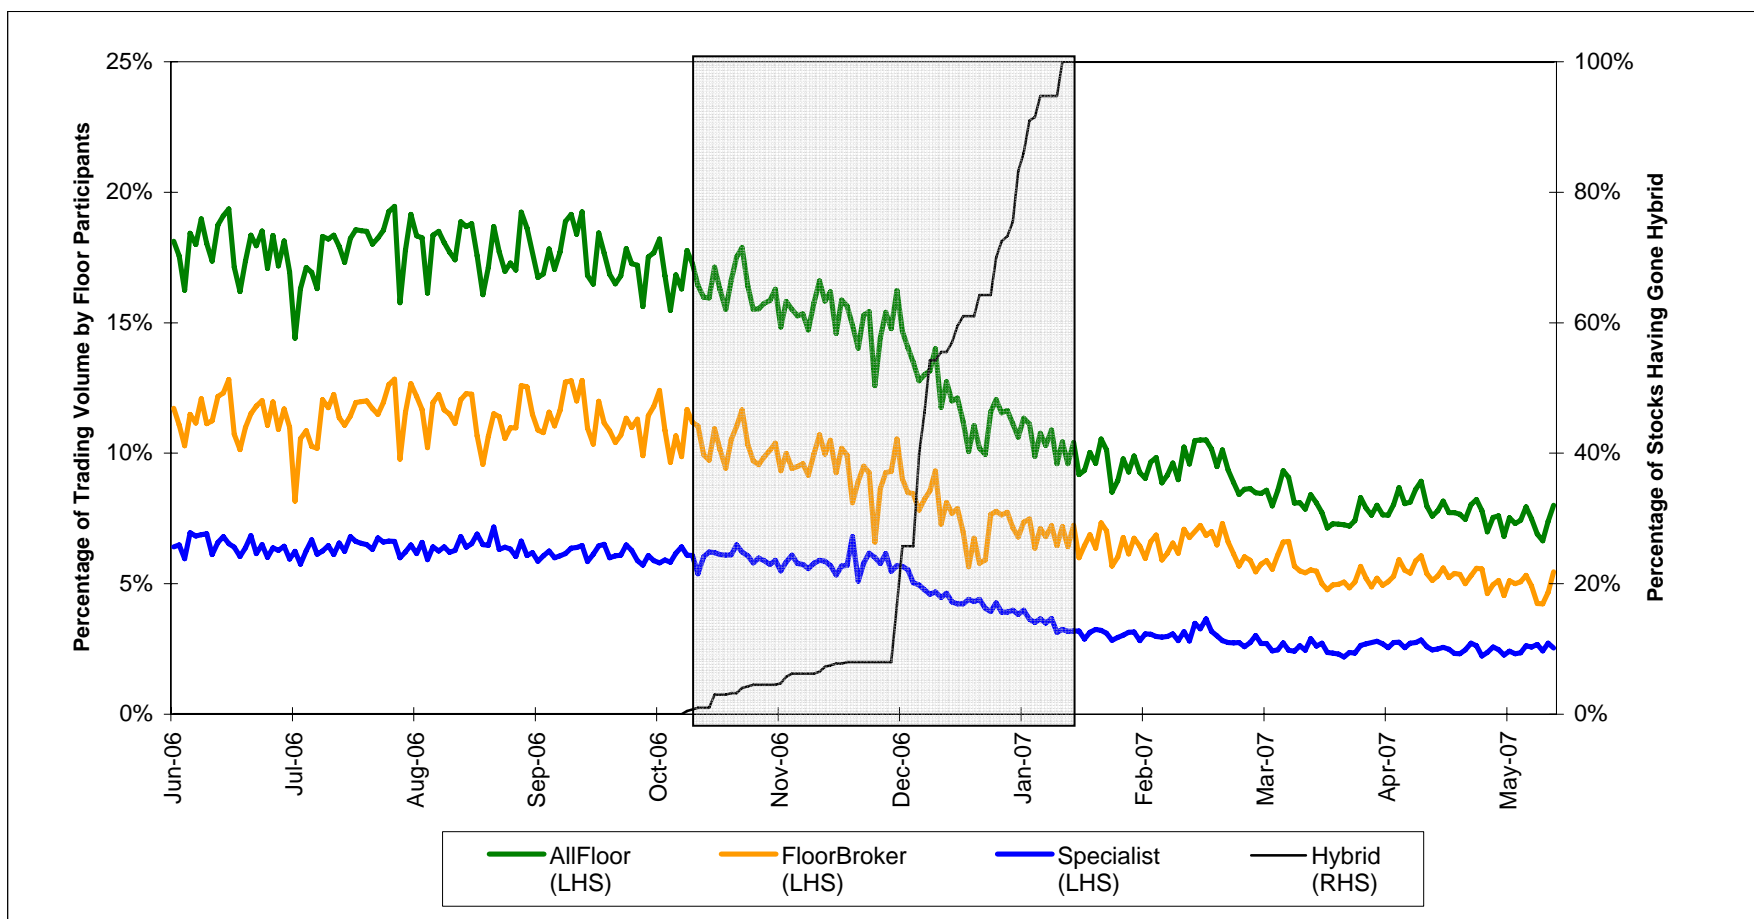


Figure 4: Effective Spreads

This chart graphs average effective spreads in basis points for the NYSE sample over the 40-day window surrounding the hybrid activation date for each stock. Equal-weighted averages across stocks are presented by quartile, where Q1 comprises the 100 largest stocks in the 400-stock sample. Spreads are calculated from TAQ data.

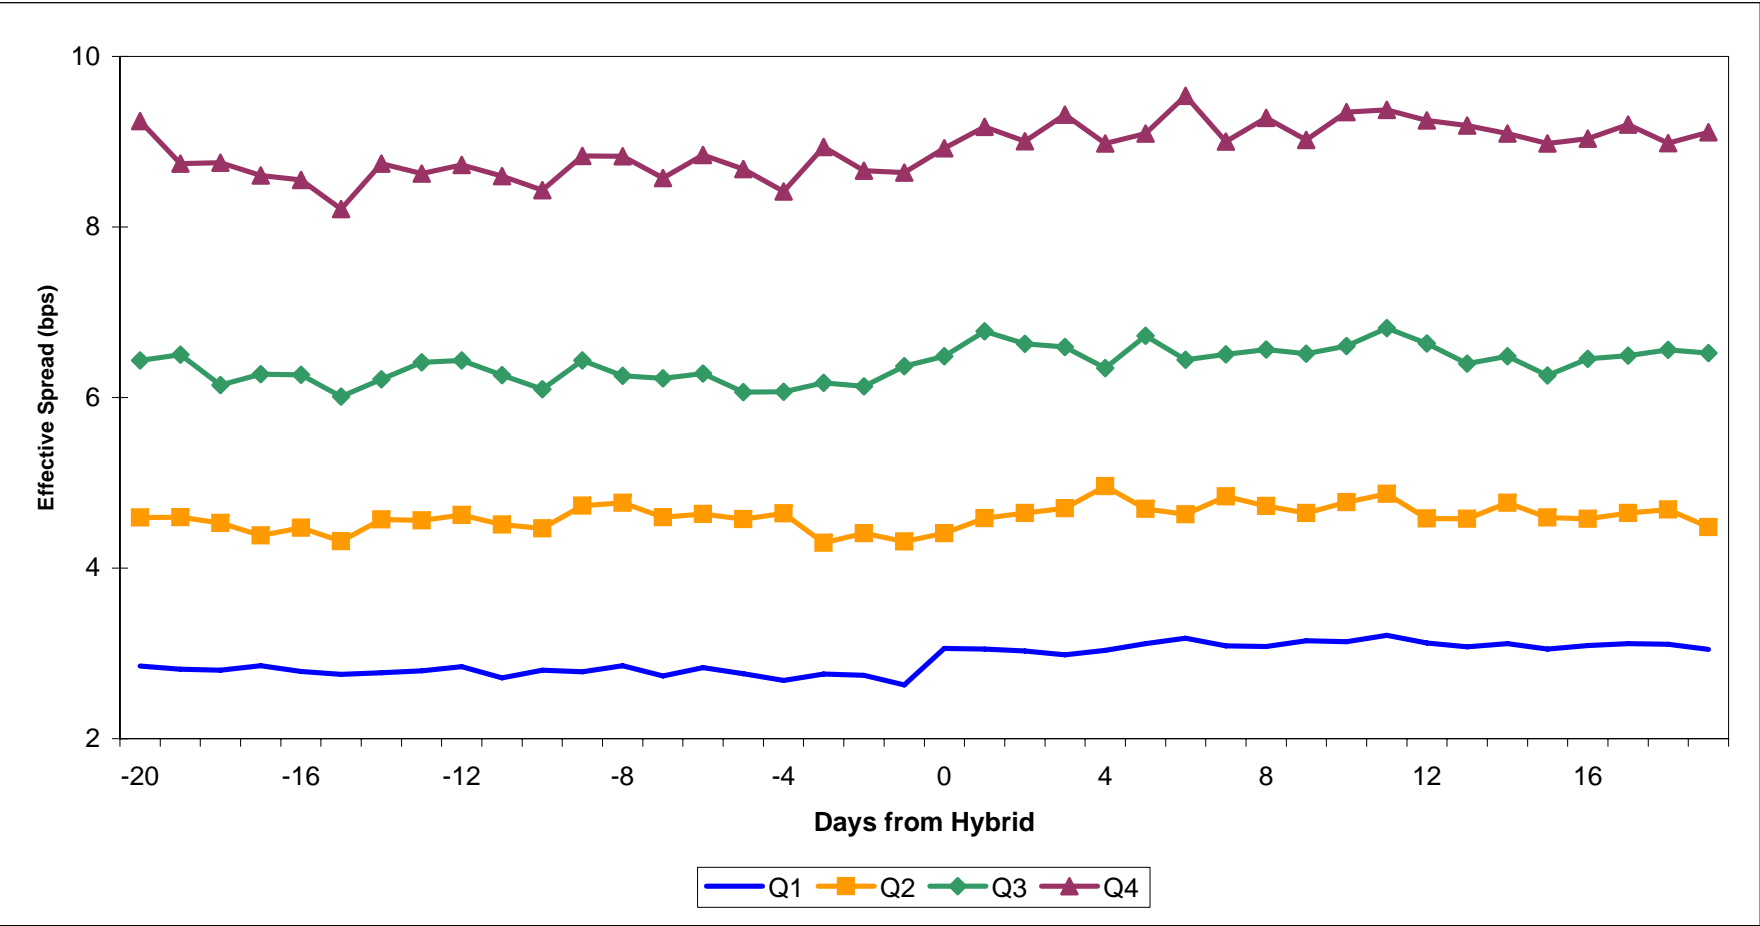


Figure 5: NYSE - Nasdaq Effective Spread Differences

This chart graphs average effective spread difference in basis points for the NYSE stocks minus their Nasdaq matches over the 40-day window surrounding the hybrid activation date for each NYSE stock. The effective spread difference is calculated for each stock each day; equal-weighted averages across stocks are presented by quartile, where Q1 comprises the 100 largest stocks in the 400-stock sample. Spreads are calculated from TAQ data.




Figure 6: Quoted and Effective Spread Differences

This chart graphs the five-day moving average of the average effective and quoted spread differences in basis points for the NYSE stocks minus their Nasdaq matches over the four-month window surrounding the hybrid activation date for each NYSE stock. Spreads are calculated for all 400 NYSE stocks and their Nasdaq matches from TAQ data.




Figure 7: Intraday Volatility and Variance Ratio Differences

This chart graphs the five-day moving average of the average five-minute trading range, five-minute quote volatility, and five-minute/30-minute variance ratio in basis points for the NYSE stocks minus their Nasdaq matches over the four-month window surrounding the hybrid activation date for each NYSE stock. Volatilities and variance ratio are calculated for all 400 NYSE stocks and their Nasdaq matches from TAQ data.

