Available online at www.sciencedirect.com
ScienceDirect
Journal of Financial Markets 14 (2011) 568–604
Journal of FINANCIAL MARKETS
www.elsevier.com/locate/finmar

Automation, speed, and stock market quality: The NYSE's Hybrid ☆

Terrence Hendershott[a], Pamela C. Moulton[b,*]

[a] Haas School of Business, University of California, 545 Student Services Building #1900, Berkeley, CA 94720, USA
[b] School of Hotel Administration, Cornell University, 435A Statler Hall, Ithaca, NY 14853, USA

Available online 24 February 2011

Abstract

Automation and trading speed are increasingly important aspects of competition among financial markets. Yet we know little about how changing a market's automation and speed affects the cost of immediacy and price discovery, two key dimensions of market quality. At the end of 2006 the New York Stock Exchange introduced its Hybrid Market, increasing automation and reducing the execution time for market orders from 10 seconds to less than one second. We find that the change raises the cost of immediacy (bid-ask spreads) because of increased adverse selection and reduces the noise in prices, making prices more efficient.

JEL classification: G12; G14

Keywords: Automation; Liquidity; Speed; Transparency; Latency; Adverse selection

☆ We thank the editor, Bruce Lehmann, an anonymous referee, Paul Bennett, Hank Bessembinder, Ekkehart Boehmer, Carole Comerton-Forde, Andrew Ellul, Sonali Hazarika, Pankaj Jain, Charles Jones, Bob Jennings, Albert Menkveld, Michael Pagano, Marios Panayides, Tavy Ronen, Andriy Shkilko, Matt Spiegel, and seminar participants at Villanova University, Wilfrid Laurier University, the 2008 Financial Management Association Conference, the 2007 NBER Market Microstructure meeting, the 2007 Triple Crown Conference, and the 2008 Western Finance Association Conference for helpful comments. We also 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.

*Corresponding author. Tel.: +1 607 255 9882.
E-mail addresses: hender@haas.berkeley.edu (T. Hendershott), pmoulton@cornell.edu (P.C. Moulton).

1386-4181/$ - see front matter
doi:10.1016/j.finmar.2011.02.003

1. Introduction

The automation and speed of the trading process have long been important dimensions of financial market design, and the growth of electronic trading in recent years has intensified the emphasis on these dimensions.[1] If automation and speed reduce transaction costs, they enable more efficient allocation of securities among heterogeneous investors, improve risk-sharing and consumption smoothing, and can raise asset prices (Pastor and Stambaugh, 2003; Acharya and Pedersen, 2005). Automation and speed may also enhance price discovery, or how efficiently new information is incorporated into prices (Chordia, Roll, and Subrahmanyam, 2008; Boehmer and Kelley, 2009). More efficient price discovery contributes to better informed financing and investment decisions, benefiting shareholders by facilitating better corporate decisions.[2] However, theoretical models of limit order books and liquidity provision offer ambiguous predictions regarding the impact of automation and speed. Existing literature compares speed across market structures (Battalio, Hatch, and Jennings, 2003; Boehmer, 2005) and levels of automation across market structures (e.g., Venkataraman, 2001). However, it is difficult to control for all differences across markets. The effect of changing automation and speed within a market is an important and understudied area.[3]

We use the New York Stock Exchange's (NYSE) introduction of its 'Hybrid Market' to study how increasing automation and speed within a market affects market quality. The Hybrid Market was designed to increase the speed of the NYSE's trading so that it would qualify as a fast market under Reg NMS' trade-through rules, which apply only to fast markets (SEC, 2005); see Section 1 for details of the Hybrid Market introduction. The Hybrid Market expands automated electronic execution and lowers the execution time for market orders from over 10 seconds to less than one second. Because the NYSE's pre-existing automated electronic execution was fully anonymous and did not allow traders on the NYSE's trading floor a last-mover advantage, the Hybrid introduction also expands those features.[4] Our analysis examines theoretical predictions of how speed and automation may affect the cost of immediacy (transaction costs for market orders) and price discovery in limit order books and in markets with dedicated market makers (specialists at the NYSE). Our empirical strategy is an event study of the Hybrid Market introduction. We match NYSE stocks with NASDAQ stocks to control for changes in

[1]See Jain (2005a, b) for evidence on the increase in automated markets and its effects on asset prices and liquidity. See The Economist (2007a, b, 2008) and Bunge (2009) for discussions of the importance of speed. The speed of trading in a market is important because delay induces uncertainty about the probability of execution and the price at which execution may occur. 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, strategies involving simultaneous trades in multiple securities, and strategies which break larger orders into smaller orders all perform worse as execution times increase. Boehmer, Jennings, and Wei (2007) find that a market center receives more order flow when its reported execution speed increases.

[2]Feedback from market prices to the firms that issue securities has long been noted; see, for example, Keynes (1936) and Tobin (1969), whose q-theory incorporates the market value of securities.

[3]See Cardella, Hao, and Kalcheva (2010) for a survey of floor trading versus automation as well as citations of studies on derivative markets switching from floor to automated trading.

[4]Automated trading need not be anonymous. The NYSE Hybrid Market implementation of fully automated trading is anonymous for immediately executable orders. Because the significant changes were implemented simultaneously, our results include changes to the speed of execution, the anonymity/transparency of the identities of marketable orders, and the last mover advantage of floor traders.

overall market conditions. We then perform tests using a difference-in-difference approach, comparing the difference between NYSE and NASDAQ stocks before and after the Hybrid introduction.

We find that Hybrid increases standard bid-ask spread measures of the cost of immediacy. From the month prior to the month subsequent to each stock's Hybrid activation date, NYSE 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. When we control for changes in NASDAQ spreads, Hybrid results in the cost of immediacy for NYSE trades increasing by more than 0.5 basis points, which is an increase of about 10% relative to their pre-Hybrid levels.[5]

We next examine whether the higher cost of immediacy after Hybrid reflects an increase in adverse selection. Spreads measure trading costs from the perspective of traders using market orders or marketable limit orders to receive immediate execution. The spreads paid by marketable orders are largely received by floor traders and non-marketable limit orders providing liquidity. Decomposing the spread into the liquidity provider revenues – the realized spread – and adverse selection, we find that the increase in the cost of immediacy is due to an increase in adverse selection. With Hybrid's introduction, the spread widens just enough to compensate liquidity suppliers for the higher adverse selection. While the spread is a zero-sum transfer between the liquidity demander and the liquidity supplier, changes in adverse selection likely result in a transfer to informed traders from uniformed traders.[6]

To explore the causes for our empirical findings, Table 1 lays out how Hybrid's increase in speed and automation may affect the cost of immediacy for floor traders and non-floor traders. Hybrid's introduction reduces the floor's advantages by increasing the anonymity of non-floor orders and decreasing the floor's last-mover advantage. Hybrid's expansion of automated trading also provides faster feedback in terms of more up-to-date trades and quotes. These direct changes could also introduce indirect changes such as an increase in the working of orders and changes in non-floor traders' patience, arrival rate, and incentives to acquire information. The theory and intuition behind these various impacts are described below.

Immediate automated execution eliminates floor traders' (specialists and floor brokers on the NYSE) last-mover advantage via their ability to condition their actions on incoming orders (e.g., to observe the identity of broker submitting the order, the price and size of the order, and the state of the limit order book before deciding whether to trade with the incoming order). Anonymity and faster execution also facilitates the breaking of large orders into smaller pieces for execution, a strategy referred to as working an order. Back and Baruch (2007) show theoretically that large traders work their orders to hide their true demand by pooling with smaller traders. When all orders are worked, floor traders have no information advantage over traders in the limit order book, because floor traders cannot condition their trading decisions on the (unobserved) true demand underlying a market order. The increase in anonymity and speed in the Hybrid Market should increase the working of orders, reducing the information advantage of floor traders. This should result in smaller trades, lower floor participation, and less favorable trades for

[5]Spreads for electronic-only trades, i.e., trades with no floor traders involved, increase by an amount similar to the overall increase in spreads.

[6]This holds if the informed traders use marketable orders or if the uninformed traders face higher than average adverse selection when their limit orders provide liquidity.

Table 1
Possible effects of increasing automation and speed.
Spreads(*PImpact*) refers to a change in the cost of immediacy due to a change in permanent price impact. (*R*)*Spread* refers to a change in the cost of immediacy due to a change in liquidity provider revenues as measured by the realized spread.

	Floor	Non-floor
Floor's advantages ↓		
Non-floor anonymity ↑	*Spreads*(*PImpact*) ↑	*Spreads*(*PImpact*) ↑
Last mover advantage ↓	*Spreads*(*PImpact*) ↑	*Spreads*(*PImpact*) ↑
Automated working of orders ↑		
Floor cooperation breaks down	*Spreads*(*PImpact*) ↑	
Faster feedback for non-floor		
Demanders observe book		
More informed about common value	*Spreads*(*PImpact*) ↑	*Spreads*(*PImpact*) ↑
Suppliers observe book		
Floor/non-floor competition ↑	(*R*)*Spreads* ↓	
Patience of non-floor traders ↑		(*R*)*Spreads* ↑
Order arrival rates ↑		(*R*)*Spreads* ↓

floor traders. We find that 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% to 11%. Also consistent with Back and Baruch (2007), trade size falls and the floor traders' advantageous executions decline as they are exposed to greater adverse selection, which should reduce their profitability.

Floor trading allows reputational benefits to arise from repeated human interaction on the floor (Benveniste, Marcus, and Wilhelm, 1992; Chan and Weinstein, 1993; Battalio, Ellul, and Jennings, 2007). The reduction in floor trading brought on by Hybrid makes reputations harder to sustain, which could lead to a breakdown in cooperation among floor traders. Floor traders would then impose greater adverse selection costs upon each other. We find that pure floor transactions begin to exhibit as much adverse selection as non-floor transactions in the Hybrid Market, suggesting that cooperation among floor traders disappears.

Faster anonymous trading mechanisms attract more informed trading (Barclay, Hendershott, and McCormick, 2003). Faster execution in Hybrid increases transparency by providing off-floor traders more up-to-date information about the state of the market. This could increase the ability of off-floor liquidity demanders to more closely monitor the market for temporary mispricings or stale quotes and pick them off (Foucault, Roell, and Sandas, 2003). The resulting higher adverse selection, or picking-off risk, may raise the cost of immediacy for liquidity demanders as they impose higher adverse selection costs on liquidity suppliers. As discussed in Hasbrouck (1991a, b), empirical measures of adverse selection reflect how informed liquidity demanders are relative to liquidity suppliers. The measures do not distinguish whether this is due to a liquidity demander becoming more informed about "public" or soon-to-be-public information or becoming more privately informed. Given that Hybrid increases information flow on a horizon similar to the increase in the execution time, the former seems more likely.

Greater transparency can also lead to more competition in liquidity provision and lower costs of immediacy (Baruch, 2005; Boehmer, Saar, and Yu, 2005). Therefore, Hybrid's introduction can provide insight into the trade-off between greater speed and automation (i) increasing adverse selection in liquidity demand, and (ii) increasing competition in liquidity supply. We find that spreads and adverse selection increase with Hybrid and that liquidity supplier profits do not decline, which are consistent with Hybrid having a more significant effect on liquidity demand than liquidity supply.

Models of the limit order book often focus on traders' equilibrium order placement decisions and the resulting cost of immediacy (e.g., Foucault, Kadan, and Kandel, 2005; Rosu, 2009).[7] These models do not contain adverse selection, so changes in the cost of immediacy are due to changes in the realized spread. Limit order models provide intuition for how Hybrid could affect competition between the floor traders, particularly the specialist, and (nonmarketable) limit orders. The last-mover advantage of the floor forces limit orders to be more aggressive to receive execution. When the floor loses its advantages, limit orders face less competition and therefore should be less aggressive. The reduction in aggressiveness is similar to an increase in patience by the limit order submitters, which widens the spread and increases the cost of immediacy. The initial increase in realized spread post-Hybrid is consistent with an increase in limit order patience. However, the increase in realized spread is not robust to the longer (eight-month) sample period and we do not find evidence of an increase in patience: The ratio of marketable to non-marketable electronic orders does not change with Hybrid's introduction.

We find mixed evidence that order arrival rates increase with Hybrid's introduction. Foucault, Kadan, and Kandel (2005) identify two competing effects of an increase in order arrival rates. First, with higher arrival rates when limit orders are placed they improve the best price less, which would lead to spreads increasing in arrival rates. Second, the faster arrival rates reduce the expected waiting time for orders in the queue so that limit orders require less compensation for delayed execution. Foucault, Kadan, and Kandel's (2005) Table 3 and discussion suggest that the lower waiting cost effect dominates and spreads should decline with higher order arrival rates. Proposition 2 and point (*i*) in the abstract of Rosu (2009) reach the same conclusion. Therefore, the increase in order arrival rates should not cause the increase in the cost of immediacy associated with Hybrid's introduction.

Predictions of how automation and faster execution should affect the efficiency of price discovery are mixed. In a limit order book, speed can decrease the noise in prices if traders with more extreme private values are more likely to obtain information about the current common value of the asset (Table 9 in Goettler, Parlour, and Rajan, 2009). However, theoretical models motivated by the NYSE [e.g., Glosten, 1989; Leach and Madhavan, 1993], show that the specialist's monopoly position can facilitate price discovery, particularly by reducing market failure. While complete market failure rarely occurs on the NYSE, partial market failure in terms of higher short-run transitory volatility has historically been of great concern to the NYSE. Whether the traders possibly becoming better informed about the current value of the asset outweighs the benefits of a quasi-monopolist liquidity supplier for producing less noisy prices is an empirical question, which the Hybrid introduction allows us to examine. We find that intraday (five-minute quote) volatility increases and the volatility of the efficient price increases with the introduction of the Hybrid Market, consistent with an increase in information production. To measure price efficiency, we calculate the autocorrelation of five-minute midquote

[7]See Kandel and Tkatch (2008) for empirical evidence on the cost of waiting.

returns, the five-minute/30-minute variance ratio, and the pricing error (motivated by Hasbrouck, 1993). Price efficiency improves according to all three measures, suggesting that there is less noise in prices and that the increase in speed allows information to be incorporated into prices more efficiently, which can facilitate better investment and financing decisions.

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 the cost of immediacy. Section 5 analyzes changes in adverse selection with the Hybrid introduction. Section 6 examines how Hybrid influences price efficiency. Section 7 investigates how the Hybrid Market introduction affects the NYSE floor. Section 8 studies Hybrid's effect on off-floor order submission strategies and discusses their impact. Section 9 concludes.

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.[8] 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 is 70% of total volume for 1990–1991. Sofianos and Werner (2000) find that this fraction declines to 55% by 1997. By 1999, electronic and floor trading are roughly equal on the NYSE.

Fig. 1 shows 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 May 2006, floor trading represents slightly more than 20% of NYSE volume.

Automatic execution was introduced on the NYSE in 2000 and allowed orders to execute directly against orders in the limit order book, eliminating the opportunity for the specialist or floor brokers to provide liquidity or participate in trades beyond their orders

[8] A number of studies empirically examine the trading of specialists and floor brokers. Hasbrouck and Sofianos (1993), 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.

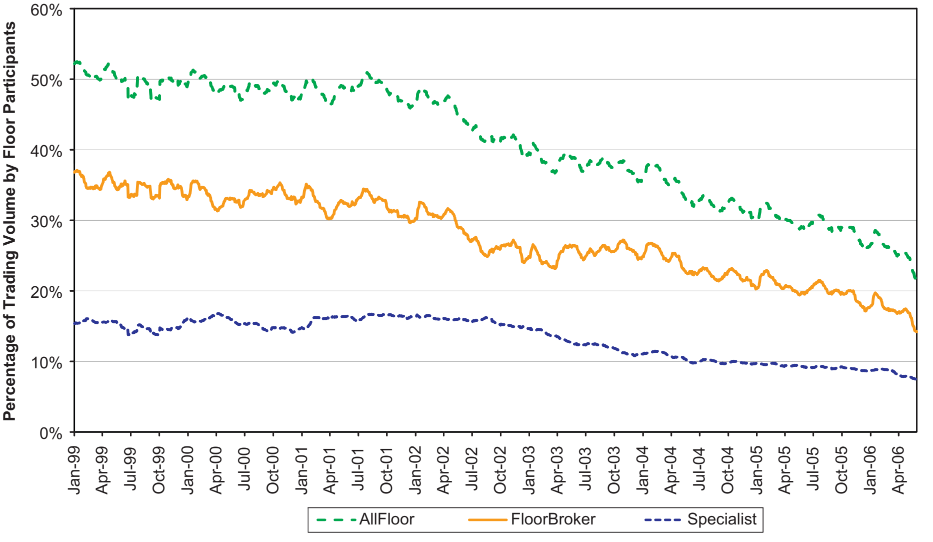

Fig. 1. Long-run floor activity. This chart graphs the participation of specialists, floor brokers, and the entire floor (specialists plus floor brokers), measured as a percentage of twice total regular-hours trading volume for each stock each day. Daily percentages are equal-weighted averages across 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 analysis in the rest of the paper excludes certain types of trades. Data are from the NYSE CAUD file.

already in the limit order book. In the auction market, the specialist could view the identity of marketable order submitters prior to execution. As a result, orders designated for automatic execution have greater anonymity than orders handled in the traditional auction mechanism. 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 order 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

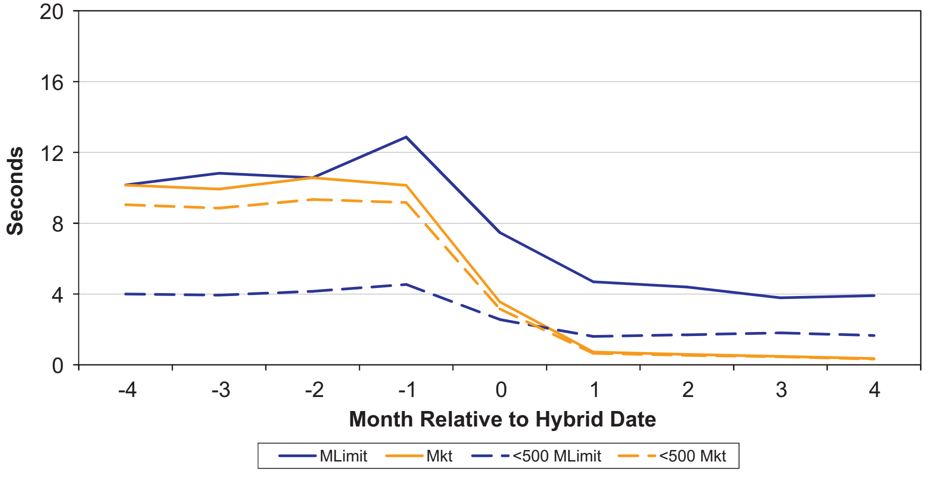

Fig. 2. Execution speeds by order type and size. This chart graphs average execution speed in seconds for the NYSE sample over the eight-month window surrounding the hybrid activation date for each stock. The solid lines represent execution speed averaged across all order size categories for market orders (*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.

execution. Market and marketable limit orders are now automatically executed by default, rather than requiring a special code. As a result of these changes, more electronic orders enjoy the enhanced speed and anonymity of automatic execution in Hybrid.[9]

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% of the time (NYSE Group, 2007); in February 2007, the NYSE reset the LRPs to less restrictive levels.

Fig. 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. Time to execution declines by more than 50% 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 a declining average trade size.

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

[9]Because the expansion of automatic execution increases both the execution speed and the anonymity of electronic orders simultaneously, it is not possible to completely disentangle the effects of the two factors.

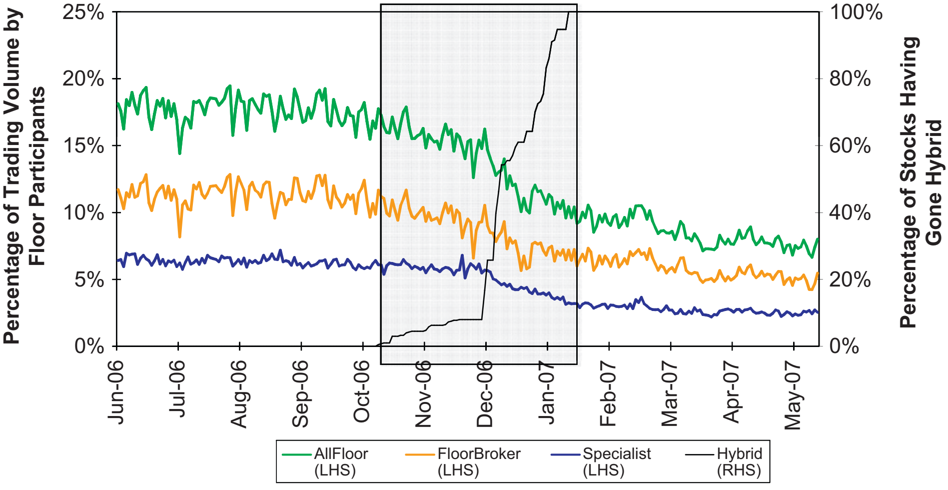

Fig. 3. Floor activity versus Hybrid activation. This chart graphs the average participation of specialists, floor brokers, and the entire floor (specialists plus floor brokers), measured as a percentage of twice regular-hours trading volume for each stock each day, excluding opening and closing trades. 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.

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 (NYSE Group, 2006a).[10]

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.[11] 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.

Fig. 3 depicts 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 beginning of Fig. 3 does not line up exactly with the end of Fig. 1 because the data in Fig. 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, Fig. 3 also indicates the percentage of stocks for which Hybrid has been introduced. Relatively few stocks go Hybrid in the first two months of the rollout. Almost half of the stocks go Hybrid at the beginning of December 2006, and another

[10] Chakrabarty and Moulton (2011) find that Hybrid reduces the effect of specialists' attention constraints on the liquidity of the stocks handled.

[11] 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, but did not change automation and speed (which changed in Phase 3). 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. Analysis of the subsample of stocks whose post-Hybrid (Phase 3) period does not overlap the implementation of Phase 4 shows that our results are not due to Phase 4 changes.

40% of stocks go Hybrid over the final few weeks of 2006. Floor activity declines gradually before and after the Hybrid rollout. When the transition to Hybrid is most intense in December 2006, floor activity declines steeply from 15% to 11%.

The fact that many stocks roll out in Hybrid in close proximity to each other requires that our empirical strategy control for contemporaneous changes in the market-wide cost of immediacy. We do this by matching NYSE stocks to NASDAQ stocks and following a difference-in-difference approach, examining how the Hybrid event impacts 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) and System Order Data (SOD) databases. 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 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.[12]

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.[13] 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

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

[13]Davies and Kim (2009) find that one-to-one matching without replacement 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.

alphabetically by symbol. We then calculate the following matching error for each NYSE stock i and each remaining NASDAQ stock j:

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

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 160-stock sample has a mean matching error of 0.08. We also construct a matched sample of NASDAQ stocks for the 160-stock sample using propensity score matching. Because all the results are qualitatively similar for the 400-stock sample, the 160-stock subsample, and the 160-stock subsample with propensity-score matching, we report only the full 400-stock sample results. Table 2 presents descriptive statistics for the NYSE and NASDAQ 400-stock samples. The NYSE stocks generally have larger market capitalization and higher prices than their NASDAQ matches, reflecting the differences between the firms that tend to list on each exchange. To ensure that our results are not due to poor matches, we include controls for the differences in price and market capitalization for each matched pair.

3.2. Data and measures

The spread measures in Table 2 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 measure the arrival rate of electronic orders, the fraction of shares placed by marketable and nonmarketable orders, and the fraction of orders that are cancellations using data from the SOD database. We obtain execution speeds and spreads on a monthly basis from the SEC Dash-5 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.[14] As the results from both samples yield the same 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) for a description of upstairs-arranged trades and how they are identified in CAUD], opening trades, and closing trades are excluded because they take place outside of the trading mechanisms that changed under the Hybrid Market implementation.

[14]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 also 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.

Table 2
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 100 stocks.

	NYSE sample				NASDAQ sample			
	Market capitalization. ($ mil)	Closing price ($)	*QSpread* (bps)	*ESpread* (bps)	Market capitalization. ($ 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

We use equally-weighted spread measures across trades within the day to calculate measures for each stock each day.[15] The percentage quoted spread is the difference between the best ask price and the best bid price at the time of a trade, divided by the prevailing midpoint of the bid and ask quotes.[16] 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

$$ESpread_{j,k,t} = 2q_{j,k,t}(p_{j,k,t} - m_{j,k,t})/m_{j,k,t}, \quad (1)$$

where $q_{j,k,t}$ 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).

Fig. 4 depicts the average difference in effective spread between NYSE stocks and their NASDAQ matches by market capitalization quartile over the window from 20 days before to 20 days after Hybrid activation. Effective spreads rise by roughly 10% at Hybrid activation (day zero) and remain higher for the next 20 days. Calculating the difference between NYSE and NASDAQ spreads controls for market-wide changes in the cost of immediacy, but does not distinguish between whether the changes in spreads are due to change on the NYSE or on NASDAQ. To study this, we examine the NYSE and NASDAQ spreads separately using both spreads reported monthly by the markets to the SEC under Rule 605 (historically referred to as Dash-5 data) and our daily spread measures calculated from TAQ.

Fig. 5 uses Dash-5 data to examine the longer-term trends in NYSE and NASDAQ effective spreads for the full 400-stock sample over the eight months surrounding the Hybrid introduction. Spreads fall on both markets in the months before and after Hybrid. During the month of and the month after Hybrid's introduction, months zero and one, spreads rise on the NYSE and decline on NASDAQ.[17] While Dash-5's monthly reporting frequency limits our ability to detect the precise impact of the Hybrid introduction, the evidence in Fig. 5 shows that Hybrid's introduction led to an increase in NYSE spreads relative to NASDAQ.

Fig. 6 depicts the backward-looking five-day moving average of daily effective spreads for the NYSE and NASDAQ matched stocks, calculated from TAQ data.[18] Fig. 6 reveals

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

[16]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.

[17]Because Dash-5 data are produced by calendar month and the Hybrid introductions occur within the month, month zero can include only a few days or many days under the Hybrid system. Examining the Hybrid rollout by calendar month, Sofianos and Abrokwah (2007) document a similar deviation from the general decline in NYSE spreads through the Hybrid rollout period. Sofianos and Abrokwah (2007) also show that the pattern is consistent across trade size categories.

[18]Differences between spreads reported under Dash-5 and those computed from TAQ data are due to several factors, which are mainly related to the accuracy with which trades are matched with quotes (Dash-5 should be perfectly accurate as opposed to the approximate matching when using TAQ), how executions are allocated to market centers, and which trades are included. Dash-5 statistics include all orders sent to a market center, regardless of where the resulting trades are eventually executed, while TAQ identifies only the market center where trades are executed. Dash-5 excludes orders exceeding 9,999 shares, even if fewer than 9,999 shares are executed, and limit order prices must be less than 10 cents from the quote, while TAQ includes trades of all sizes and irrespective of any limit price of the original order. These result in differences in the spread levels, but the changes

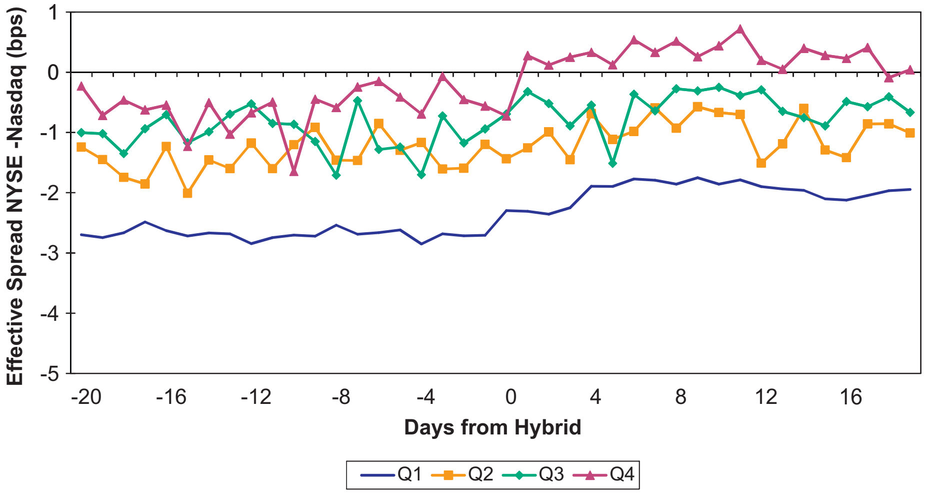

Fig. 4. NYSE–NASDAQ effective spreads. This chart graphs average effective spread difference in basis points for the NYSE stocks minus their NASDAQ matches over the 40-day window surrounding the Hybrid activation date for each NYSE stock. The effective spread difference is calculated for each stock each day; equal-weighted averages across stocks are presented by quartile, where Q1 comprises the largest 100 stocks in the 400-stock sample. Spreads are calculated from TAQ data.

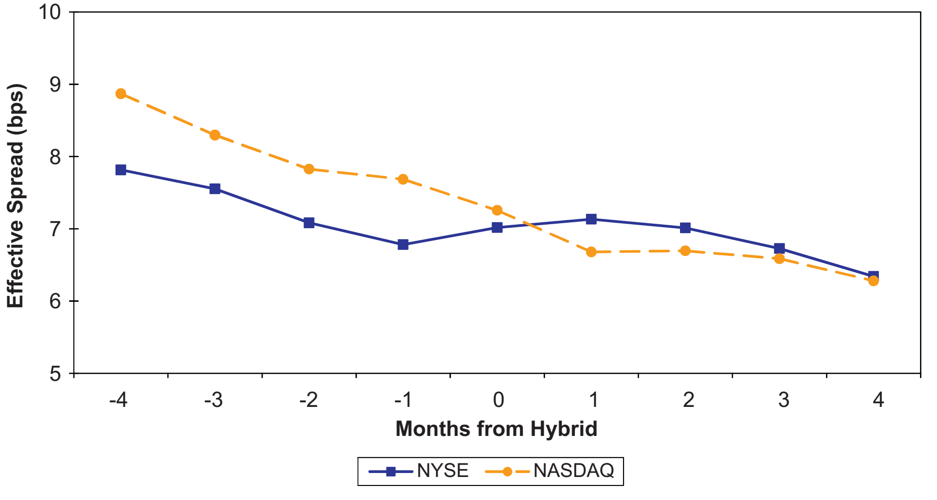

Fig. 5. Effective spreads NYSE versus NASDAQ from Dash-5. This chart graphs the average effective spread in basis points for the NYSE stocks and their NASDAQ matches over the eight-month window surrounding the Hybrid activation date for each NYSE stock. Average spreads are calculated for market and marketable limit orders across all 400 NYSE stocks and their Nasdaq matches from Dash-5 data.

(*footnote continued*)
in NYSE and NASDAQ spreads over time, particularly around the Hybrid introduction, are similar using Dash-5 and TAQ data.

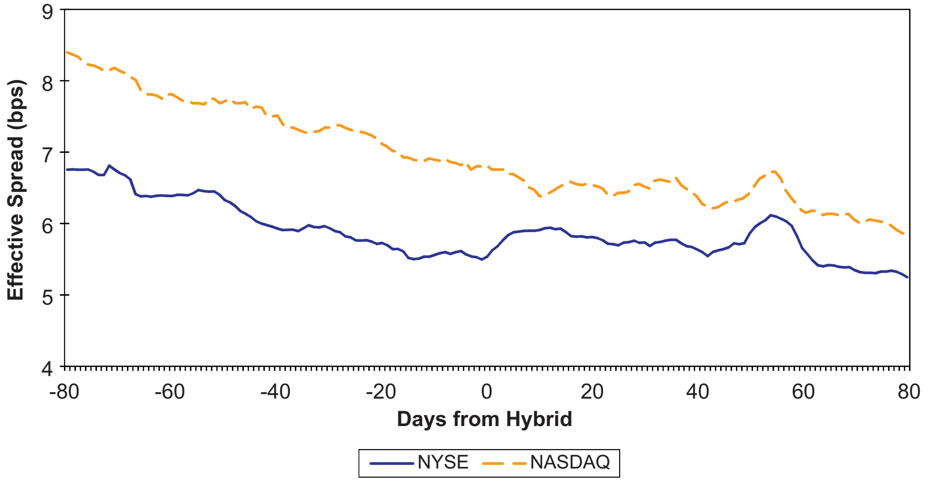

Fig. 6. Effective spreads NYSE versus NASDAQ from TAQ. This chart graphs the backward-looking five-day moving average of the average effective spread in basis points for the NYSE stocks and their Nasdaq matches over the eight-month window surrounding the Hybrid activation date for each NYSE stock. Average spreads are calculated across all 400 NYSE stocks and their NASDAQ matches from TAQ data.

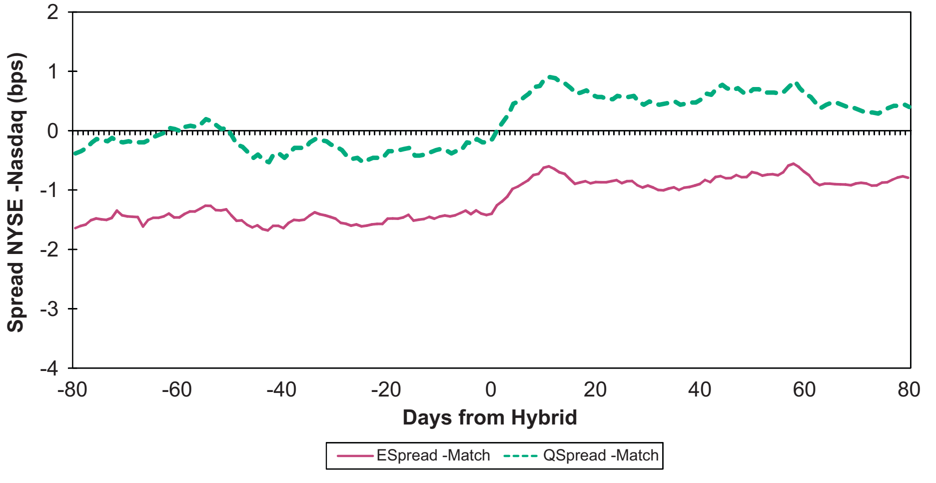

Fig. 7. NYSE–NASDAQ Quoted and Effective Spreads. This chart graphs the backward-looking five-day moving average of the average effective and quoted spread differences in basis points for the NYSE stocks minus their NASDAQ matches over the four-month window surrounding the Hybrid activation date for each NYSE stock. Average spread differences are calculated across all 400 NYSE stocks less their NASDAQ matches from TAQ data.

a marked increase in NYSE spreads on the Hybrid implementation date, preceded and followed by a general downward trend in effective spreads. The notable increase in effective spreads on both markets around day +50 reflects the events of February 27, 2007, when the Shanghai stock market crashed, former Federal Reserve Chairman Greenspan warned of a recession, and Dow Jones reported erroneous values for the Dow Jones Industrial Average: a perfect storm of record volume and market dislocation that led to high spreads on both the NYSE and NASDAQ.

The general decline in spreads depicted in Figs. 5 and 6 suggests that we should control for market-wide spread changes when examining the effect of Hybrid on the cost of immediacy on the NYSE. Fig. 7 shows the backward-looking five-day moving average of effective and quoted spread NYSE–NASDAQ differences for the full 400-stock sample. The increases in effective and quoted spread differences occur upon Hybrid's introduction, and the increase persists in the months following.

4. Hybrid and the cost of immediacy

We move beyond the simple spread graphs to examine how Hybrid affects the cost of immediacy. We calculate the above-described measures for the 20 days before and the 20 days after Hybrid activation for each stock. The first four columns of Table 3 present these results for quoted and effective spreads. Consistent with Fig. 6, the quoted and effective spread measures in Table 3 show that NYSE spreads increase with Hybrid's introduction. Controlling for the matched NASDAQ stocks' spreads increases the pre/post difference (as in Figs. 4 and 7). This confirms our earlier graphical finding that while NYSE stock spreads widen at the time of Hybrid, NASDAQ spreads narrow. The increase in spreads is generally greater for smaller stocks. We conduct most of our analysis by quartile as well as for the full sample to control for heterogeneity across quartiles.

The last three columns of Table 3 examine trade size, quoted depth, and trading volume. Trade size shows a clear decline with Hybrid's introduction, consistent with an increase in worked orders, as in Back and Baruch (2007). 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, suggesting that the wider spreads are not associated with additional liquidity at the quote. NYSE trading volume also shows no significant change.

Table 3 uses standard univariate *t*-tests to calculate the statistical significance of the changes associated with Hybrid. Given that the pre- and post-periods surrounding Hybrid activation overlap for many stocks (Fig. 3), the assumption of independence across observations may overstate statistical significance. We adopt a panel data approach in Table 4 to properly control for this. We run the following regression for each spread variable:

$$Sprd_{i,t} = \alpha + \beta Hybrid_{i,t} + \gamma Volatility_t + \sum_{k=1}^{4} \delta_k\, ControlVariable_{i,t,k} + \varepsilon_{i,t}, \tag{2}$$

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 $ControlVariable_{i,t,k}$ are four stock-level control variables: the log of the absolute daily price difference, the daily turnover difference, the daily stock volatility difference (calculated as Alizadeh, Brandt, and Diebold, 2002), and the log of the absolute difference in market capitalization on March 31, 2006, for each NYSE and NASDAQ matched stock pair.[19] We also run a variation in Eq. (2) that includes stock fixed effects, α_i, and excludes the control

[19]Omission of the control variables from Eq. (2) does not affect the coefficients on our variable of interest. 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 Hybrid coefficient.

Table 3

Summary statistics 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 100 stocks. *QSpread* is the quoted spread; *ESpread* is the effective spread; *QVolatility* is the five-minute midquote return volatility; *TRange* is the five-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 (bps)	*QSpread* -*Match* (bps)	*ESpread* (bps)	*ESpread* - *Match* (bps)	*Trade Size* (shares)	*QDepth* (100s)	*Volume* ($ mn)
Full sample							
Pre-Hybrid	7.9	−0.3	5.6	−1.4	396	16.1	37.1
Post-Hybrid	8.3	0.5	5.9	−0.9	303	16.1	38.3
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	4.1	−1.2	3.1	−2.3	414	29.5	106.0
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	6.6	0.3	4.7	−1.0	302	15.8	26.5
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	9.3	1.0	6.5	−0.6	256	9.6	12.7
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	13.0	2.0	9.1	0.2	241	9.6	8.2
Change	0.7**	1.4**	0.5**	0.8**	−45**	0.6	0.5

Table 4
Panel regressions of spreads on Hybrid.

Analysis periods are the 40-day window (first four columns) and 8-month window (last four columns) surrounding each stock's Hybrid activation date. 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 matched pair of stocks: the daily difference in price, turnover, and stock volatility. Specifications without fixed effects also include the difference in market capitalization between the NYSE and Nasdaq matched pair of stocks. Coefficients for volatility, 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 100 stocks. *t*-Statistics, reported in parentheses below coefficient estimates, are robust to time series and cross-sectional correlation.

Dependent variable	40-day window				8-month window			
	QSpread - Match		*ESpread - Match*		*QSpread - Match*		*ESpread - Match*	
Fixed effects	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*
Full sample								
Hybrid	0.84	0.78	0.13	0.51	0.82	0.69	0.73	0.61
	(5.5)	(6.5)	(4.2)	(5.5)	(4.2)	(4.3)	(4.5)	(4.5)
Observations	16,000	16,000	16,000	16,000	64,000	64,000	64,000	64,000
Adj. R^2	0.06	0.77	0.05	0.74	0.05	0.68	0.05	0.61
Quartile 1								
Hybrid	0.69	0.44	0.25	0.44	0.79	0.61	0.99	0.81
	(2.9)	(5.5)	(3.4)	(7.0)	(2.5)	(2.3)	(3.1)	(3.1)
Quartile 2								
Hybrid	0.50	0.50	0.23	0.36	0.32	0.29	0.46	0.41
	(1.6)	(2.2)	(2.4)	(2.2)	(0.7)	(1.1)	(1.6)	(2.5)
Quartile 3								
Hybrid	0.68	0.82	0.14	0.43	0.46	0.34	0.33	0.23
	(3.4)	(4.9)	(3.7)	(3.3)	(1.5)	(1.7)	(1.4)	(1.6)
Quartile 4								
Hybrid	1.40	1.33	0.30	0.79	1.54	1.47	0.97	0.93
	(3.8)	(4.2)	(3.0)	(3.1)	(3.5)	(3.3)	(2.6)	(2.4)

variable for the market capitalization difference. We conduct inference in this and all subsequent regressions using double-clustered Thompson (2011) standard errors, which are robust to both cross-sectional correlation and idiosyncratic time-series persistence. Regressions are run for the full 400-stock sample and by quartile.[20] For brevity, *R*-squareds and number of observations are reported for only the full sample. *R*-squareds for the quartiles

[20] An alternative specification in which the spreads of all 800 stocks (NYSE and NASDAQ) are regressed on the Hybrid indicator, the Hybrid indicator interacted with an indicator equal to one for NYSE stocks, calendar-day time dummies, volatility, and control variables as described above provides similar inference: The sign, magnitude, and statistical significance of the coefficient on the $Hybrid \times NYSE$ variable in this regression are similar to those of the Hybrid coefficient reported in the paper.

are similar, although typically somewhat higher than the full sample for quartile 1 and lower for quartile 4.

We conduct our analysis over two periods: the 40-day window (first four columns) and the eight-month window (last four columns) surrounding each stock's Hybrid activation. 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 4) are of the same magnitude as the average changes in Table 3. The inclusion of volatility in the regressions 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 full-sample *t*-statistics ranging from 4.2 to 6.5. As in Table 3, the Hybrid impact generally increases in the smaller quartiles. The coefficient on Hybrid is 0.51 basis points for the full sample regression of effective spreads with stock fixed effects over the 40-day window. This is almost a 10% increase in spreads from the 5.6 basis point pre-Hybrid average in Table 3. 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: 0.44 basis points on an average of 2.8 basis points (Table 3) for an increase of 15%. The results for the eight-month 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 0.61 basis points in the eight-month analysis versus 0.51 basis points in the 40-day window (full sample with fixed effects).

The analysis in this section shows that the cost of immediacy increases following Hybrid and that these changes are not transitory adjustment effects. In the following section we examine to what extent the increase in spreads is attributable to an increase in adverse selection, as in Barclay, Hendershott, and McCormick (2003) and Foucault, Roell, and Sandas (2003). In Section 8, we analyze how possible changes in off-floor traders' order submissions strategies were affected by Hybrid and the impact those changes could have on the non-informational costs of immediacy.

5. Adverse selection changes in Hybrid

We now consider how the Hybrid Market affects adverse selection for trading on the NYSE and off-NYSE trading of NYSE-listed stocks.

5.1. Adverse selection on the NYSE

Effective spreads are perhaps the most common measure of the cost of immediacy. It is useful to decompose effective spreads into their permanent and transitory portions to examine changes in adverse selection and liquidity provider profits. 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

$$PImpact_{j,k,t} = 2q_{j,k,t}(m_{j,k+5,t} - m_{j,k,t})/m_{j,k,t}, \tag{3}$$

where $q_{j,k,t}$ is an indicator variable that equals one for buyer-initiated trades and negative one for seller-initiated trades (see the 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

$$RSpread_{j,k,t} = 2q_{j,k,t}(p_{j,k,t} - m_{j,k+5,t})/m_{j,k,t}, \tag{4}$$

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

We also calculate the Hasbrouck (1991a) impulse response measure, which measures adverse selection allowing for potential lagged adjustments to the information in trades and quotes. (See Appendix for details on the Hasbrouck decomposition.) We use the same panel data approach as in Table 4. We run the following regression for each spread decomposition measure:

$$SprdDecomp_{i,t} = \alpha_i + \beta\ Hybrid_{i,t} + \gamma\ Volatility_t + \sum_{k=1}^{3} \delta_k ControlVariable_{i,t,k} + \varepsilon_{i,t}, \tag{5}$$

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 $ControlVariable_{i,t,k}$ represents the three stock-level control variables: the log of the absolute daily price difference, the daily turnover difference, and the daily stock volatility difference for each NYSE and NASDAQ matched stock pair.

Table 5 shows the results from estimating the regression over both the 40-day window (first three columns) and the eight-month window surrounding the Hybrid introduction (last three columns). The impulse response measure, which is robust to price discreteness and lagged price adjustment, shows a significant increase in adverse selection over both horizons. The simple five-minute price impact is insignificant over the 40-day window, but has larger increases that are significant in the eight-month analysis. Overall, the evidence supports adverse selection rising with the Hybrid introduction.

We find that realized spreads increase in the period immediately surrounding the Hybrid introduction but are insignificant over the long run. Together these findings suggest 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. This is consistent with liquidity demanders learning over time how to best utilize the new systems (as in Foucault, Roell, and Sandas, 2003) and liquidity suppliers having short-lived market power after Hybrid's introduction.

5.2. Trading of NYSE-listed stocks on other markets

The analysis so far has examined how the Hybrid Market introduction affects trading of NYSE-listed stocks on the NYSE. These changes may be due to changes in the trading of existing market participants, entry of new market participants, or movement of market participants across markets. Furthermore, NYSE-listed stocks also trade on other markets, and the Hybrid changes may have caused participants to switch from the NYSE to those other markets or to switch from the other markets to the NYSE. While the lack of data linking trading activity to the final investor precludes definitive answers, some analysis on

Table 5
Panel regressions of adverse selection and liquidity provider revenues.

Analysis periods are the 40-day window (first three columns) and 8-month window (last three columns) surrounding each stock's Hybrid activation date. Five-minute price impact (*PImpact*), five-minute realized spread (*RSpread*), and Hasbrouck impulse response (*Impulse*) 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: the daily difference in price, turnover, and stock volatility. Coefficients for volatility, 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 100 stocks. *t*-Statistics, reported in parentheses below coefficient estimates, are robust to time series and cross-sectional correlation.

Dependent variable	40-day window			8-month window		
	PImpact - Match	*RSpread - Match*	*Impulse - Match*	*PImpact - Match*	*RSpread - Match*	*Impulse - Match*
Full sample						
Hybrid	0.13	0.38	9.37	0.59	0.02	12.56
	(1.4)	(3.9)	(10.1)	(6.4)	(0.2)	(11.8)
Observations	16,000	16,000	16,000	64,000	64,000	64,000
Adj. R^2	0.27	0.31	0.50	0.19	0.29	0.46
Quartile 1						
Hybrid	0.17	0.27	6.49	0.64	0.16	7.57
	(1.7)	(2.4)	(6.0)	(3.7)	(1.1)	(5.2)
Quartile 2						
Hybrid	0.07	0.29	6.97	0.36	0.06	10.27
	(0.5)	(1.7)	(5.3)	(2.7)	(0.3)	(6.4)
Quartile 3						
Hybrid	0.19	0.24	11.52	0.56	−0.33	15.24
	(1.3)	(1.7)	(7.2)	(4.3)	(−2.2)	(9.2)
Quartile 4						
Hybrid	0.05	0.74	12.25	0.75	0.18	17.19
	(0.2)	(2.9)	(6.1)	(3.8)	(0.7)	(9.0)

how trading changed off the NYSE may provide insight into possible interpretations of the results.

Although NYSE trading volume in NYSE-listed stocks does not change significantly around the Hybrid introduction (Table 3), trading volume in NYSE-listed stocks in other markets rises significantly throughout our sample period.[21] Fig. 8 shows the average daily trading volume for NYSE-listed stocks on and off the NYSE, equally-weighted across the 400-stock sample with each stock's volume normalized each day relative to trading volume in its matched NASDAQ stock. The growth in off-NYSE trading volume appears roughly constant throughout the period, with the off-NYSE trading market share increasing from 24% to 35% over the eight months. Off-NYSE growth may be due to a long-term trend or due to market participants' anticipating Reg NMS and the introduction of Hybrid.

[21]Off-NYSE trading in NYSE-listed stocks is dominated by NASDAQ and Archipelago, which report about 65% and 33%, respectively, of the off-NYSE trading volume in our sample.

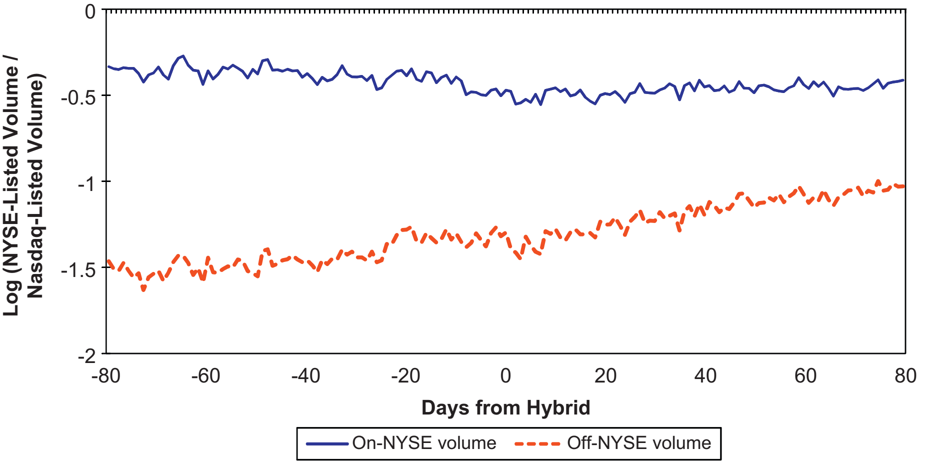

Fig. 8. On-NYSE versus off-NYSE trading volume in NYSE-listed stocks. This chart graphs the average log ratio of share volume in NYSE-listed stocks on and off the NYSE relative to share volume in their matched Nasdaq-listed stocks over the four-month window surrounding the Hybrid activation date for each NYSE-listed stock. Average log ratios are calculated across all 400 NYSE stocks versus their NASDAQ matches from TAQ data.

We run regressions analogous to those previously run for spread measures (Table 4) and decompositions (Table 5) based on off-NYSE trading to examine how trading in NYSE-listed stocks off the NYSE is affected by the Hybrid introduction:

$$DepVar_{i,t}^{off\text{-}NYSE} = \alpha_i + \beta\ Hybrid_{i,t} + \gamma\ Volatility_t + \sum_{k=1}^{3} \delta_k ControlVariable_{i,t,k} + \varepsilon_{i,t}, \tag{6}$$

where $DepVar_{i,t}^{off\text{-}NYSE}$ is the average quoted spread, effective spread, price impact, or realized spread for stock i trading off the NYSE 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 $ControlVariable_{i,t,k}$ represents three stock-level control variables: the log of the absolute daily price difference, the daily turnover difference, and the daily stock volatility difference for each NYSE and NASDAQ matched stock pair.

Table 6 presents the regression results for quoted spreads, effective spreads, five-minute price impacts, and five-minute realized spreads. The first four columns present results for the 40-day window surrounding the Hybrid introduction; the last four columns extend the analysis to the eight months surrounding the Hybrid introduction. Over the 40-day window, the results for off-NYSE trading are similar to those for on-NYSE trading (in Tables 4 and 5): quoted and effective spreads rise significantly with Hybrid, price impact shows little change, and realized spreads rise significantly. Over the longer horizon, off-NYSE trades exhibit large increases in price impact, but insignificant changes in quoted and effective spreads because the higher price impacts are offset by declines in realized spreads. This contrasts with the continuation of higher quoted and effective spreads and unchanged long-horizon realized spreads on the NYSE (Tables 4 and 5).

Table 6
Panel regressions of spreads, adverse selection, and liquidity provider revenues for off-NYSE trading.

Analysis periods are the 40-day window (first four columns) and 8-month window (last four columns) surrounding each stock's Hybrid activation date. Quoted Spread (*QSpread*), effective spread (*ESpread*), five-minute price impact (*PImpact*), and five-minute realized spread (*RSpread*) differences between NYSE stock trades off the NYSE 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: the daily difference in price, turnover, and stock volatility. Coefficients for volatility, 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 100 stocks. *t*-Statistics, reported in parentheses below coefficient estimates, are robust to time series and cross-sectional correlation.

Dependent variable	40-day window				8-month window			
	QSpread - *Match*	*ESpread* - *Match*	*PImpact* - *Match*	*RSpread* - *Match*	*QSpread* - *Match*	*ESpread* - *Match*	*PImpact* - *Match*	*RSpread* - *Match*
Full sample								
Hybrid	0.61	0.23	−0.02	0.25	0.04	−0.05	0.86	−0.91
	(5.0)	(2.4)	(−0.2)	(2.1)	(0.2)	(−0.3)	(7.4)	(−4.9)
Observations	16,000	16,000	16,000	16,000	64,000	64,000	64,000	64,000
Adj. R^2	0.73	0.70	0.21	0.26	0.62	0.55	0.15	0.22
Quartile 1								
Hybrid	0.39	0.25	0.21	0.03	0.31	0.58	0.88	−0.30
	(5.7)	(4.2)	(1.9)	(0.3)	(1.2)	(2.1)	(4.7)	(−1.5)
Quartile 2								
Hybrid	0.32	0.08	−0.05	0.13	−0.31	−0.23	0.59	−0.81
	(1.2)	(0.5)	(−0.3)	(0.6)	(−1.0)	(−0.9)	(3.6)	(−2.5)
Quartile 3								
Hybrid	0.58	0.12	−0.02	0.14	−0.39	−0.56	0.82	−1.38
	(3.6)	(0.9)	(−0.1)	(0.7)	(−1.7)	(−2.8)	(5.2)	(−6.2)
Quartile 4								
Hybrid	1.12	0.44	−0.26	0.70	0.49	−0.11	1.06	−1.18
	(3.5)	(1.7)	(−1.1)	(2.4)	(1.1)	(−0.3)	(4.1)	(−3.8)

The increase in adverse selection (price impact) for both on-NYSE and off-NYSE trading over the longer horizon is consistent with liquidity demanders on and off the NYSE becoming better informed. The increase in speed on the NYSE may encourage new or existing market participants to become more informed about the current common value of the asset or to possibly uncover more information. Over the longer horizon liquidity provider profits (realized spreads) decrease off-NYSE while realized spreads remain unchanged on-NYSE. This indicates increased competition by liquidity providers off-NYSE, possibly due to increasing liquidity externalities on the non-NYSE markets, which increase their trading volume. It is possible that the increase in speed on the NYSE allows better integration of trading across markets, enabling existing liquidity suppliers on-NYSE to face less risk of double execution when also posting limit orders in off-NYSE markets. However, the data cannot rule out other possibilities such as new liquidity providers entering.

6. Volatility and price efficiency

The analysis up to this point shows that Hybrid increases execution speed on the NYSE and leads to higher costs of immediacy as adverse selection rises. The increase in the execution speed and informativeness of trades could also affect the efficiency with which information is incorporated into prices, depending on the balance between increased competition in liquidity provision (as in Goettler, Parlour, and Rajan, 2009) and the diminished role of a quasi-monopolist liquidity supplier (as in Glosten, 1989; Leach and Madhavan, 1993). We analyze this issue by studying the volatility of price changes and whether these changes reflect information or noise.

We first examine three measures of intraday volatility: the five-minute trading range, the five-minute quote return volatility, and the five-minute volatility of the efficient price. 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, incorporating both high-frequency transitory volatility and 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 volatility of the efficient price is measured for each stock each day as in Hasbrouck (1993) (see Appendix). The volatility of the efficient price provides a measure of whether volatility increases reflect a change in the amount of information being incorporated into prices or an increase in noise. We run the following regressions, using the same panel data approach as before:

$$IntradayVol_{i,t} = \alpha_i + \beta Hybrid_{i,t} + \gamma Volatility_t + \sum_{k=1}^{3} \delta_k ControlVariable_{i,t,k} + \varepsilon_{i,t}, \quad (7)$$

where $IntradayVol_{i,t}$ is the average five-minute trading range, five-minute quote volatility, or five-minute volatility of the efficient price 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 $ControlVariable_{i,t,k}$ represents three stock-level control variables: the log of the absolute daily price difference, the daily turnover difference, and the daily stock volatility difference for each NYSE and NASDAQ matched stock pair.

Table 7 presents the results from estimating Eq. (7) over both the 40-day window (first three columns) and the eight months surrounding the Hybrid introduction (last three columns). Analyses at both horizons show that the five-minute trading range generally rises with the Hybrid introduction. The five-minute quote volatility does not show robust changes with Hybrid's introduction over the 40-day window, but it does show a significant increase in the largest stock quartile over the eight-month window. The increase in the efficient price volatility over the same period indicates that the higher volatility for the largest stocks is not simply noise, but reflects information being incorporated faster into prices.

We next examine three measures of price efficiency: the absolute value of the autocorrelation of five-minute midquote returns, the five-minute/30-minute variance ratio,

Table 7
Panel regressions of intraday volatility on Hybrid.

Analysis periods are the 40-day window (first three columns) and 8-month window (last three columns) surrounding each stock's Hybrid activation date. Five-minute trading range (*TRange*), quote volatility (*QVolatility*), and efficient price volatility (*EPVolatility*) 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: the daily difference in price, turnover, and stock volatility. Coefficients for volatility, 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 100 stocks. *t*-Statistics, reported in parentheses below coefficient estimates, are robust to time series and cross-sectional correlation.

Dependent variable	40-day window			8-month window		
	TRange - Match	*QVolatility - Match*	*EPVolatility - Match*	*TRange - Match*	*QVolatility - Match*	*EPVolatility - Match*
Full sample						
Hybrid	1.19	0.20	0.10	2.21	0.36	0.23
	(6.6)	(0.6)	(0.4)	(9.4)	(1.8)	(1.2)
Observations	16,000	16,000	16,000	64,000	64,000	64,000
Adj. R^2	0.76	0.57	0.27	0.72	0.57	0.23
Quartile 1						
Hybrid	1.56	0.21	1.00	3.37	1.11	1.56
	(4.8)	(0.9)	(1.9)	(7.1)	(3.9)	(4.3)
Quartile 2						
Hybrid	1.49	0.55	0.43	1.99	0.32	0.29
	(4.9)	(1.2)	(0.9)	(6.8)	(1.2)	(0.9)
Quartile 3						
Hybrid	1.06	0.25	−0.85	1.73	−0.03	−0.95
	(3.3)	(0.5)	(−1.4)	(4.2)	(−0.1)	(−2.1)
Quartile 4						
Hybrid	0.61	−0.29	−0.17	1.51	−0.14	−0.24
	(1.7)	(−0.5)	(−0.4)	(3.2)	(−0.4)	(−0.8)

and the pricing error. The return autocorrelation is based on five-minute midquote returns, calculated for each stock each day. If prices follow a random walk, the return autocorrelation should be equal to zero. The absolute value of the autocorrelation measures the extent to which quote returns diverge from a random walk in either direction, so a decline in the absolute value of the quote return autocorrelation would indicate an increase in price efficiency.[22] 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.[23] The variance ratio evaluates whether

[22]The five-minute quote autocorrelations are on average negative. If we use the signed autocorrelation rather than the absolute value of the autocorrelation, the qualitative inference in unchanged, i.e., with Hybrid's introduction the signed autocorrelations move closer to zero by becoming less negative.

[23]Calculating the variance ratios daily is noisy because the 30-minute volatility is sometimes close to zero. Therefore, we winsorize the stock/day variance ratios at the first and 99th percentiles before computing NYSE minus NASDAQ differences. Calculating pooled variance ratios by the pre and post periods yields qualitatively similar results, but makes statistical inference problematic because of the common factor in returns across stocks.

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 price efficiency. The pricing error is measured for each stock each day and is motivated by Hasbrouck (1993) (see Appendix). A Vector Autoregression (VAR) of quote returns and trade directions is estimated using ten lags, from which we identify the random walk component (interpreted as the efficient price) and the residual stationary component (deviations of midquote prices from efficient prices) of the stock price process. The residual stationary component has a mean of zero, so the pricing error is defined as the standard deviation of deviations of midquote prices from efficient prices over time. This approach produces a measure that is a lower bound on the pricing error. A decline in the pricing error indicates an increase in price efficiency.

Fig. 9 presents the three efficiency 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. All three measures fall following the Hybrid implementation, suggesting that NYSE prices become more efficient following the Hybrid implementation.

We run the following regressions to more precisely examine the relation between the Hybrid introduction and price efficiency, using the same panel data approach as before:

$$Eff_{i,t} = \alpha_i + \beta\ Hybrid_{i,t} + \gamma\ Volatility_t + \sum_{k=1}^{3} \delta_k ControlVariable_{i,t,k} + \varepsilon_{i,t}, \tag{8}$$

where $Eff_{i,t}$ is the absolute value of the autocorrelation of five-minute midquote returns, the five-minute/30-minute variance ratio of midquote returns, or the pricing error for stock i less its NASDAQ match on day t; and the explanatory variables are as described above.

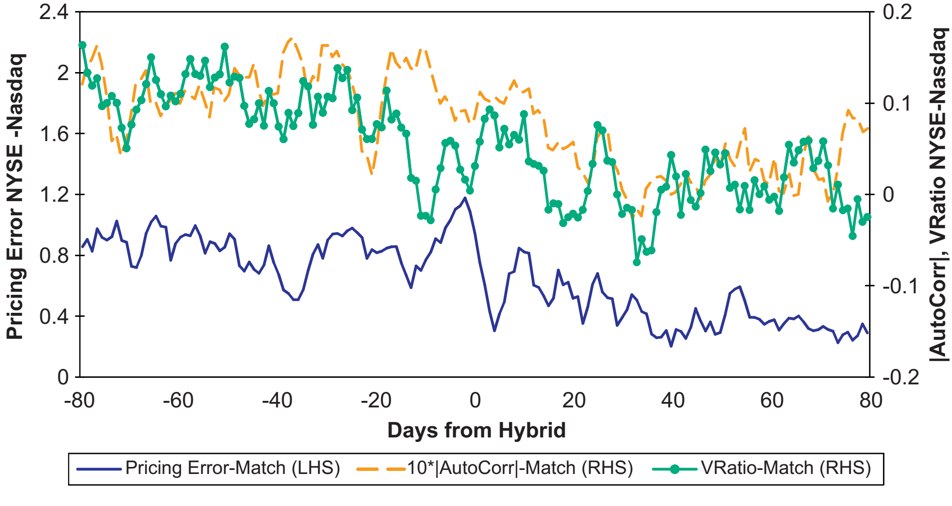

Fig. 9. NYSE–NASDAQ price efficiency measures. This chart graphs the backward-looking five-day moving average of the pricing error (in basis points), absolute value of the five-minute quote return autocorrelation (scaled by a factor of 10), and five-minute/30-minute variance ratio (in basis points) for the NYSE stocks minus their NASDAQ matches over the four-month window surrounding the Hybrid activation date for each NYSE stock. Efficiency measures are calculated for all 400 NYSE stocks and their NASDAQ matches from TAQ data.

Table 8
Panel regressions of efficiency measures on Hybrid.

Analysis periods are the 40-day window (first three columns) and 8-month window (last three columns) surrounding each stock's Hybrid activation date. Absolute value of five-minute quote return autocorrelation (|*Corr*|), five-minute/30-minute variance ratio (*VRatio*), in basis points, and pricing error (*PrError*), in basis points, 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: the daily difference in price, turnover, and stock volatility. Coefficients for volatility, control variables, constant, and stock fixed effects are not reported. Full sample is 400 stocks; Quartile 1 comprises the largest 100 stocks. *t*-Statistics, reported in parentheses below coefficient estimates, are robust to time series and cross-sectional correlation.

Dependent variable	40-day window			8-month window		
	|*Corr*| - *Match*	*VRatio* - *Match*	*PrError* - *Match*	|*Corr*| - *Match*	*VRatio* - *Match*	*PrError* - *Match*
Full sample						
Hybrid	−0.003	0.02	−0.34	−0.006	−0.05	−0.44
	(−0.9)	(0.5)	(−3.0)	(−3.9)	(−2.5)	(−5.0)
Observations	16,000	16,000	16,000	64,000	64,000	64,000
Adj. R^2	0.03	0.05	0.24	0.01	0.03	0.21
Quartile 1						
Hybrid	−0.002	0.03	−0.05	−0.006	−0.04	−0.03
	(−0.5)	(0.5)	(−0.3)	(−2.3)	(−0.9)	(−0.3)
Quartile 2						
Hybrid	−0.001	0.08	−0.18	−0.002	−0.05	−0.29
	(−0.1)	(1.3)	(−1.0)	(−0.9)	(−1.5)	(−2.1)
Quartile 3						
Hybrid	−0.003	−0.01	−0.53	−0.005	−0.05	−0.83
	(−0.5)	(−0.2)	(−2.6)	(−1.9)	(−1.5)	(−4.5)
Quartile 4						
Hybrid	−0.005	−0.03	−0.61	−0.012	−0.08	−0.65
	(−1.0)	(−0.4)	(−2.0)	(−4.0)	(−2.0)	(−2.8)

Table 8 presents the results from estimating Eq. (8) over both the 40-day window (first three columns) and the eight months surrounding the Hybrid introduction (last three columns). Over the 40-day window, the pricing error shows a statistically significant decline in the full sample of 0.34 basis points relative to NASDAQ, suggesting that price efficiency improves with the introduction of Hybrid. This represents a 7% decrease in the pricing error relative to the pre-Hybrid average pricing error of 4.79 basis points.[24] The reductions in pricing error are concentrated in the smaller stocks. While the return autocorrelation and variance ratio show no statistically significant changes over the 40-day window around Hybrid, over the longer horizon all three measures suggest that efficiency increases with the introduction of Hybrid. These findings are consistent with Goettler,

[24] Table 7 shows that the volatility of the efficient price increases with Hybrid's introduction. Therefore, if the pricing error is calculated relative to the efficient price volatility, the decrease in the pricing error due to Hybrid's introduction is even larger.

Parlour, and Rajan's (2009) prediction that price efficiency improves when private-value traders choose to acquire more information on the common value.

7. Hybrid's impact on the NYSE floor

Up to this point we have used the publicly available TAQ and Dash-5 data to measure the cost of immediacy. Fig. 3 shows that floor trading declined with the introduction of Hybrid. To separate the effect of Hybrid on trades involving floor participants versus those not involving the floor, we use the NYSE's CAUD data. The CAUD database matches buyers and sellers for each NYSE trade, providing information about whether the parties on each side of a trade are electronic orders (system orders) or floor participants (specialists and floor brokers). 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 800 shares of ABC may execute with a floor broker buy order for 200 shares 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[25]:

Pure Floor = Specialist and Floor Brokers, or Floor Brokers only;
Pure System = System participants only; or
Floor and System Interaction = Specialist and System participants, 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).[26] 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 the same 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 9 shows that the most significant switch in the 40-day window surrounding Hybrid activation is a roughly 8% change towards pure system trading away from mixed-initiator floor-system interaction trades. This likely stems from Hybrid's faster execution making system trading more attractive and precluding floor participants from joining what would otherwise have been pure system trades all along.[27] The fraction of trading that is pure floor is nearly unchanged in the 40-day window surrounding Hybrid's introduction, remaining under 2% in all quartiles. Floor-initiated

[25]Market participant types are determined from the SOURCE, ACCT, and TYPE codes in the CAUD file. Percentage (CAP) executions are included as floor broker executions. Incoming Intermarket Trading System (ITS) executions are included as system participant executions.

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

[27]Consistent with floor participants' inability to join trades at the point of sale, Boni and Rosen (2006) and Sofianos and Abrokwah (2007) document that price improvement virtually disappears with the introduction of Hybrid.

Table 9
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 five columns, within an 8-month window in the last five columns. Cross-sectional means are presented for the full sample of 400 stocks and by quartile, with Quartile 1 comprising the largest 100 stocks. *Pure System* trades involve only system participants; *Pure Floor* trades involve only floor brokers and/or the specialist; *Floor and System Interaction* trades involve some combination of floor and system participants. Floor and System Interaction Trades are further categorized by which type of participant inititated 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					8-month window				
			Floor and System Interaction					*Floor and System Interaction*		
	Pure System (%)	*Pure Floor (%)*	*Floor-Initiated (%)*	*System-Initiated (%)*	*Mixed Initiator (%)*	*Pure System (%)*	*Pure Floor (%)*	*Floor-Initiated (%)*	*System-Initiated (%)*	*Mixed-Initiator (%)*
Full sample										
Pre-Hybrid	70.5	1.4	4.7	8.4	15.0	68.1	1.7	4.9	8.9	16.4
Post-Hybrid	78.2	1.4	3.5	10.5	6.4	80.6	1.1	3.0	9.9	5.4
Change	7.7**	0.0	−1.3**	2.1**	−8.6**	12.5**	−0.5**	−1.9**	1.0**	−11.0**
Quartile 1										
Pre-Hybrid	67.0	1.4	5.1	7.0	19.5	63.9	1.7	5.3	7.4	21.5
Post-Hybrid	75.0	1.6	4.1	11.3	8.0	77.8	1.3	3.6	10.6	6.6
Change	8.0**	0.2	−1.0**	4.3**	−11.5**	13.9**	−0.4**	−1.7**	3.2**	−14.9**
Quartile 2										
Pre-Hybrid	70.7	1.3	4.8	8.1	15.0	67.9	1.7	5.0	8.8	16.6
Post-Hybrid	78.6	1.4	3.3	10.4	6.3	80.8	1.1	2.9	9.7	5.5
Change	7.9**	0.0	−1.6**	2.3**	−8.7**	12.9**	−0.6**	−2.1**	1.0*	−11.1**
Quartile 3										
Pre-Hybrid	72.3	1.5	4.5	8.8	12.9	69.9	1.7	4.6	9.5	14.3
Post-Hybrid	79.8	1.3	3.2	9.9	5.8	81.8	1.1	2.7	9.6	4.9
Change	7.4**	−0.1	−1.3**	1.1**	−7.1**	11.8**	−0.6**	−1.9**	0.1	−9.5**
Quartile 4										
Pre-Hybrid	71.9	1.3	4.6	9.7	12.6	70.6	1.5	4.8	9.9	13.2
Post-Hybrid	79.4	1.3	3.4	10.4	5.6	82.0	1.0	2.8	9.7	4.5
Change	7.5**	−0.1	−1.2**	0.8	−7.0**	11.4**	−0.5**	−2.0**	−0.2	−8.7**

interaction trades decrease about 1% to 2% while system-initiated interaction trades increase a similar amount. The last five columns reveal similar patterns over the eight-month period surrounding the introduction of Hybrid.

We run regressions of the following form to study how Hybrid affects trades with different combinations of participants:

$$SprdDecomp_{i,k,t} = \alpha_i + \sum_{k=1}^{5} \gamma_k Hybrid_{i,t} \times Type_k + \sum_{k=2}^{5} \beta_k Type_k + \delta\ Volatility_t + \sum_{m=1}^{3} \lambda_m ControlVariable_{i,t,m} + \varepsilon_{i,t}, \tag{9}$$

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[28]; 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; $Type_k$ is a dummy variable indicating the trade type (pure system, pure floor, floor-initiated interaction, system-initiated interaction, or mixed-initiator interaction); $Volatility_t$ is the opening value of CBOE's VIX index on day t; and $ControlVariable_{i,t,k}$ represents three stock-level control variables: the log of the absolute daily price difference, the daily turnover difference, and the daily stock volatility difference for each NYSE and NASDAQ matched stock pair. 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 10 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 eight months 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 4. 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 eight-month analysis as well as in the period immediately surrounding Hybrid introduction.

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 (Benveniste, Marcus, and Wilhelm, 1992; Chan and Weinstein, 1993; Battalio, Ellul, and Jennings, 2007) attenuating adverse selection. After Hybrid implementation, 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 by selectively choosing which system orders to execute against, a task made more difficult in a market with more worked orders (Back and Baruch, 2007). Losing this ability could explain the 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

[28] We do not include the impulse response measure in this analysis because there are too few trades in many of the who-trades-with-whom categories to allow robust daily estimation.

Table 10
Panel regressions of spreads on Hybrid by who trades with whom type.

Analysis periods are the 40-day window (first three columns) and 8-month window (last three columns) surrounding each stock's Hybrid activation date. Effective spread (*ESpread*), five-minute price impact (*PImpact*), and five-minute realized spread (*RSpread*) 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: the daily difference in price, turnover, and stock volatility. Coefficients for control variables and stock fixed effects are not reported. All dependent variables are in basis points. Sample is 400 stocks. *t*-Statistics, reported in parentheses below coefficient estimates, are robust to time series and cross-sectional correlation.

Dependent variable	40-day window			8-month window		
	ESpread - Match	*PImpact - Match*	*RSpread - Match*	*ESpread - Match*	*PImpact - Match*	*RSpread - Match*
Hybrid × *Pure System*	0.49	0.08	0.41	0.58	0.66	−0.10
	(5.1)	(0.9)	(4.3)	(4.3)	(6.8)	(−0.9)
Hybrid × *Pure Floor*	1.33	1.90	−0.61	1.66	2.18	−0.53
	(7.9)	(3.8)	(−1.2)	(8.7)	(7.8)	(−1.9)
Hybrid × *Floor-Initiated Interaction*	1.57	−0.72	2.25	2.00	−0.03	1.98
	(12.5)	(−2.5)	(7.2)	(12.3)	(−0.1)	(8.1)
Hybrid × *System-Initiated Interaction*	0.93	2.22	−1.31	1.10	2.81	−1.74
	(6.4)	(8.9)	(−5.6)	(7.2)	(15.7)	(−9.0)
Hybrid × *Mixed-Initiator Interaction*	0.71	1.13	−0.45	0.93	1.40	−0.51
	(5.5)	(4.9)	(−1.9)	(6.0)	(8.6)	(−2.9)
Pure Floor	−0.13	−0.88	0.72	−0.12	−0.93	0.75
	(−1.3)	(−2.5)	(2.1)	(−1.5)	(−5.9)	(4.6)
Floor-Initiated Interaction	−0.61	−0.87	0.26	−0.59	−0.86	0.27
	(−9.1)	(−4.0)	(1.1)	(−8.9)	(−5.7)	(1.6)
System-Initiated Interaction	−0.12	−2.07	1.94	−0.07	−2.21	2.13
	(−1.8)	(−11.0)	(11.4)	(−1.3)	(−18.1)	(18.8)
Mixed-Initiator Interaction	0.93	0.84	0.09	0.97	1.09	−0.12
	(12.8)	(6.8)	(0.6)	(14.9)	(11.9)	(−1.0)
Observations	74,546	74,546	74,546	296,820	296,820	296,820
Adj. R^2	0.42	0.04	0.05	0.36	0.03	0.04

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 eight-month period surrounding Hybrid's introduction (last three columns of Table 10) 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 spread 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 eight-month period, Hybrid introduction leads to an increase in the price impact of pure system trades of 0.66 basis points, with a *t*-statistic of 6.8. 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 faster speed leading to greater adverse selection in the limit order book.

8. The impact of changes in off-floor order submission strategies

Hybrid's increase in transparency due to faster execution and the decrease in floor traders' advantages may impact the order submission strategies of off-floor traders. With more up-to-date information about transactions and quotes, off-floor traders may choose to monitor the market more closely and increase the rate of their order submissions, even if the amount of trading remains constant. Off-floor traders may also become more patient because with better information and greater anonymity they better optimize their limit order placement. In addition, the decline in the floor's last-mover advantage should decrease the competition faced by limit order submitters, which would make them more patient. Foucault, Kadan, and Kandel (2005) and Rosu (2009) show that an increase in order arrival rate should lead to a decrease in the cost of immediacy. These papers also show that an increase in traders' patience should cause an increase in cost of immediacy.

We use data from the NYSE's electronic order submission database, SOD, to test for changes in order submission strategies. We measure the arrival rate of all orders, the fraction of shares placed by marketable and nonmarketable orders, and the fraction of orders that are cancellations of outstanding orders. Because we do not have similar order-level data for our matched NASDAQ stocks and because order arrival rates have steadily increased over time, we introduce a time trend variable in these analyses. We use the following regression specification to examine how order submission strategies are affected by the Hybrid introduction:

$$DepVar_{i,t} = \alpha_i + \beta_0\, trend_t + \beta_1\, Hybrid_{i,t} + \beta_2\, Hybrid_{i,t} \times trend_t + \gamma\, Volatility_t + \sum_{k=1}^{3} \delta_k ControlVariable_{i,t,k} + \varepsilon_{i,t}, \tag{10}$$

where $DepVar_{i,t}$ is the order arrival rate (in thousands of orders per day), the ratio of shares in non-marketable limit to the total number of shares in all orders placed, or the fraction of shares cancelled to total shares placed; $Hybrid_{i,t}$ is an indicator variable taking the value of one if the stock is in Hybrid mode on day t, otherwise zero; $trend_t$ is number of days that day t is from the beginning of the sample period; $Volatility_t$ is the opening value of CBOE's VIX index on day t; and $ControlVariable_{i,t,k}$ represents three stock-level control variables: the log of the daily price, the daily turnover, and the daily stock volatility for each NYSE stock.

Table 11 provides the order submission strategy regression results with the first three columns presenting results for the 40-day window surrounding the Hybrid introduction and the last three columns extending the analysis to the eight months surrounding the Hybrid introduction. There is limited support for an increase in order arrival rates: Only the largest quartile has a positive and significant coefficient on the Hybrid dummy variable and only in the eight-month sample period. In addition, the variable that interacts Hybrid's introduction with the time trend variable is negative and significant in that same regression, suggesting that the Hybrid's increase in order arrival rates declines over time.

Table 11
Panel regressions of orders and cancellations on Hybrid.

Analysis periods are the 40-day window (first three columns) and 8-month window (last three columns) surrounding each stock's Hybrid activation date. The Order Arrival Rate (in thousands of orders per day) and ratios of non-marketable limit order shares to total shares placed (*NonMkt%*) and cancelled shares to total shares placed (*Cancel%*) 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, a time trend, the time trend interacted with the Hybrid indicator (*Hybrid* × *Time trend*), and the following control variables for each stock: the log daily price, daily turnover, and daily volatility. Coefficients for volatility, time trend, control variables, constant, and stock fixed effects are not reported. Full sample is 400 stocks; Quartile 1 comprises the largest 100 stocks. *t*-Statistics, reported in parentheses below coefficient estimates, are robust to time series and cross-sectional correlation.

Dependent variable	40-day window			8-month window		
	Order Arrival Rate	*NonMkt%*	*Cancel%*	*Order Arrival Rate*	*NonMkt%*	*Cancel%*
Full sample						
Hybrid	5.53	−0.03	−0.05	10.08	−0.04	0.00
	(0.9)	(−0.3)	(−1.2)	(2.3)	(−0.9)	(−0.2)
Hybrid × *Time trend*	−0.022	0.000	0.001	−0.058	0.001	0.000
	(−0.5)	(0.7)	(1.7)	(−1.8)	(1.8)	(1.2)
Observations	16,000	16,000	16,000	64,000	64,000	64,000
Adj. R^2	0.89	0.32	0.32	0.85	0.30	0.30
Quartile 1						
Hybrid	19.10	−0.22	−0.10	30.89	−0.19	−0.07
	(1.5)	(−1.9)	(−1.4)	(2.8)	(−3.1)	(−1.8)
Hybrid × *Time trend*	−0.107	0.002	0.001	−0.201	0.001	0.000
	(−1.1)	(1.8)	(1.2)	(−2.6)	(3.2)	(1.8)
Quartile 2						
Hybrid	−2.51	0.28	0.13	−0.94	0.05	0.04
	(−0.4)	(1.9)	(2.0)	(−0.2)	(0.7)	(1.2)
Hybrid × *Time trend*	0.03	0.00	0.00	0.01	0.00	0.00
	(0.8)	(−1.6)	(−1.6)	(0.5)	(−0.2)	(−0.7)
Quartile 3						
Hybrid	−5.41	−0.06	−0.05	−0.31	0.02	0.02
	(−1.0)	(−0.4)	(−0.6)	(−0.1)	(0.3)	(0.7)
Hybrid × *Time trend*	0.045	0.001	0.001	0.009	0.000	0.000
	(1.2)	(0.7)	(1.1)	(0.4)	(0.4)	(0.3)
Quartile 4						
Hybrid	−2.79	0.10	−0.03	−0.02	−0.03	0.02
	(−0.7)	(0.8)	(−0.5)	(0.0)	(−0.4)	(0.5)
Hybrid × *Time trend*	0.031	0.000	0.001	0.017	0.001	0.000
	(1.1)	(−0.3)	(1.2)	(1.4)	(1.5)	(0.8)

These findings combined with the theoretical prediction that higher arrival rates should decrease spreads provide no support for changes in order arrival rates explaining our results.

If traders become more patient with Hybrid's introduction, their increased patience would increase the spread due to higher costs of immediacy, as in Foucault, Kadan, and Kandel (2005) and Rosu (2009). Table 5 finds such an increase in the realized spread, although only for the 40-day sample period. However, the coefficients on the Hybrid variables in Table 11 do not provide evidence that off-floor trader's order submission choices reflect increased patience. One possible explanation is that the fraction of non-marketable orders is a poor proxy for traders' patience. Overall, our analysis of Hybrid's effect on order submission strategies does not provide evidence that such changes are responsible for our results.

9. Conclusion

The NYSE's introduction of its Hybrid Market increases automation and speeds up electronic trading by an order of magnitude: The execution time for market orders drops from 10 seconds to less than one second. We show that the Hybrid Market raises the cost of immediacy (the effective spread) by about 10% relative to its pre-Hybrid level and that this increase is attributable to higher adverse selection. The increase in adverse selection is accompanied by information being incorporated into prices more efficiently.

Price efficiency is a public good that can inform corporate investment and financing decisions. A faster market can also enhance welfare by reducing risk-averse traders' uncertainty about the probability and price at which execution may occur. Furthermore, faster trading can facilitate more complex trading strategies.

The cost of immediacy is a zero-sum transfer from liquidity demanders to liquidity suppliers. Calculating the aggregate welfare effects of changes in the cost of immediacy in markets with heterogeneous market participants requires a structural model with numerous assumptions about traders' utilities and strategies. If all traders follow the same strategies before and after Hybrid's introduction, the adverse-selection-driven increase in the cost of immediacy implies transfers from uninformed to informed traders. Greater losses by uninformed traders can hinder risk-sharing. However, if Hybrid lowers the cost of information acquisition and more extreme private-value traders choose to become informed, then risk-sharing could be enhanced. The potential positive and negative welfare impacts of changes in the cost of immediacy preclude sharp determinations of the overall welfare impact of increasing execution speed.

The SEC's goal with Reg NMS is to enhance competition between markets. 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 laudable goals of Reg NMS are 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). Increasing the speed of execution and the efficiency of prices likely gives investors greater confidence of fair treatment, but the move to faster electronic trading raises the cost of immediacy via adverse selection. Thus it may be challenging to meet both of the SEC's goals.

Appendix. Hasbrouck decompositions

Hasbrouck (1991a, b) 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 are 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)].

We construct a VAR with two equations to separate price moves into trade-related and trade-unrelated components. The first equation describes the trade-by-trade evolution of the quote midpoint, while the second equation describes the persistence of order flow. Define $q_{j,t}$ to be the buy-sell indicator for trade t in stock j (+1 for buys, −1 for sells), and define $r_{j,t}$ 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 trading, and the first term captures price changes that are unrelated to trading (sometimes referred to as public information). As discussed in Hasbrouck (1991a, b), this method is robust to price discreteness, lagged adjustment to information, and lagged adjustment to trades. The VAR is estimated for each stock each day. The random-walk component is assumed to be the efficient price, making the square root of the variance of the random-walk component represent the volatility of the efficient price. The volatility of the efficient price is in transaction time, so it is then converted to a five-minute volatility. The impulse response is the permanent impact of a trade innovation:

$$\sum_{i=0}^{\infty} b_i.$$

Using a similar VAR, Hasbrouck (1993) decomposes price changes into their random walk (permanent price change) and transitory (pricing error) changes and calculates a lower bound on the pricing error. Because the pricing error has zero mean, its volatility is used to measure the magnitude of the pricing error. By using midquote prices/return in our VAR, we remove the effects of the increase in spreads and focus on the efficiency of quotes.

Estimating the VAR on a daily basis occasionally results in large outliers (e.g., impulse responses that differ from the mean by more than 10 standard deviations). Therefore, we winsorize at the first and 99th percentiles all variables estimated from the VAR: the impulse response, the volatility of the efficient price, and the pricing error.

References

Acharya, V., Pedersen, L.H., 2005. Asset pricing with liquidity risk. Journal of Financial Economics 77, 375–410.

Alizadeh, S., Brandt, M., Diebold, F., 2002. Range-based estimation of stochastic volatility models. Journal of Finance 57, 1047–1091.

Back, K., Baruch, S., 2007. Working orders in limit-order markets and floor exchanges. Journal of Finance 62, 1589–1621.

Barclay, M., Hendershott, T., 2003. Price discovery and trading after hours. Review of Financial Studies 16, 1041–1073.

Barclay, M., Hendershott, T., McCormick, D.T., 2003. Competition among trading venues: information and trading on electronic communication networks. Journal of Finance 58, 2637–2665.

Baruch, S., 2005. Who benefits from an open limit-order book? Journal of Business 78, 1267–1306.

Battalio, R., Ellul, A., Jennings, R., 2007. Reputation effects in trading on the New York Stock Exchange. Journal of Finance 62, 1243–1271.

Battalio, R., Hatch, B., Jennings, R., 2003. All else equal? A multi-dimensional analysis of retail market order execution quality. Journal of Financial Markets 6, 143–162.

Benveniste, L., Marcus, A., Wilhelm, W., 1992. What's so special about the specialist? Journal of Financial Economics 32, 61–86.

Bessembinder, H., 2003. Issues in assessing trade execution costs. Journal of Financial Markets 6, 233–257.

Boehmer, E., 2005. Dimensions of execution quality: recent evidence for U.S. equity markets. Journal of Financial Economics 78, 463–704.

Boehmer, E., Jennings, R., Wei, L., 2007. Public disclosure and private decisions: equity market execution quality and order routing. Review of Financial Studies 20, 315–358.

Boehmer, E., Kelley, E., 2009. Institutional investors and the informational efficiency of prices. Review of Financial Studies 22, 3563–3594.

Boehmer, E., Saar, G., Yu, L., 2005. Lifting the veil: an analysis of pre-trade transparency at the NYSE. Journal of Finance 60, 783–815.

Boni, L., Rosen, M., 2006. Hello Hybrid, goodbye price improvement! UNX Strategy Paper.

Bunge, J., 2009. NYSE adjusts charges in bid to draw traders. Wall Street Journal, February 3, C5.

Chakrabarty, B., Moulton, P.C., 2011. The role of market design in alleviating attention constraints. Cornell University Working Paper.

Chan, Y., Weinstein, M., 1993. Bid-ask spread and market structure. Financial Analysts Journal, 57–62.

Cardella, L., Hao, J., Kalcheva, I., 2010. The floor trader vs. automation: a survey of theory and empirical evidence. University of Arizona Working Paper.

Chordia, T., Roll, R., Subrahmanyam, A., 2008. Liquidity and market efficiency. Journal of Financial Economics 87, 249–268.

Davies, R., Kim, S.S., 2009. Using matched samples to test for differences in trade execution costs. Journal of Financial Markets 12, 173–202.

Foucault, T., Kadan, O., Kandel, E., 2005. Limit order book as a market for liquidity. Review of Financial Studies 18, 1171–1217.

Foucault, T., Roell, A., Sandas, P., 2003. Market making with costly monitoring: an analysis of the SOES controversy. Review of Financial Studies 16, 345–384.

Glosten, L., 1989. Insider trading, liquidity, and the role of the monopolist specialist. Journal of Business 62, 211–235.

Goettler, R., Parlour, C., Rajan, U., 2009. Informed traders and limit order markets. Journal of Financial Economics 93, 67–87.

Handa, P., Schwartz, R., Tiwari, A., 2004. The economic value of a trading floor: evidence from the American Stock Exchange. Journal of Business 77, 331–355.

Harris, L., Hasbrouck, J., 1996. Market vs. limit orders: the SuperDOT evidence on order submission strategy. Journal of Financial and Quantitative Analysis 31, 213–231.

Hasbrouck, J., 1991a. Measuring the information content of stock trades. Journal of Finance 46, 179–207.

Hasbrouck, J., 1991b. The summary informativeness of stock trades: an econometric analysis. Review of Financial Studies 4, 571–595.

Hasbrouck, J., 1993. Assessing the quality of a security market: a new approach to transaction-cost measurement. Review of Financial Studies 6, 191–212.
Hasbrouck, J., Sofianos, G., 1993. The trades of market-makers: an analysis of NYSE specialists. Journal of Finance 48, 1565–1594.
Jain, P., 2005a. Financial market design and the equity premium: electronic versus floor trading. Journal of Finance 60, 2955–2985.
Jain, P., 2005b. Institutional design and liquidity at stock exchanges around the world. University of Memphis Working Paper.
Jones, C., Kaul, G., Lipson, M., 1994. Information, trading, and volatility. Journal of Financial Economics 36, 127–154.
Kandel, E., Tkatch, I., 2008. Demand for the immediacy of execution: time is money. Hebrew University Working Paper.
Keynes, J.M., 1936. The General Theory of Employment, Interest, and Money, Harvest Book edition Harcourt Brace, New York.
Leach, C., Madhavan, A., 1993. Price experimentation and security market structure. Review of Financial Studies 6, 375–404.
Lee, C., Ready, M., 1991. Inferring trade direction from intraday data. Journal of Finance 46, 733–747.
Madhavan, A., Cheng, M., 1997. In search of liquidity: block trades in the upstairs and downstairs markets. Review of Financial Studies 10, 175–203.
Madhavan, A., Smidt, S., 1993. An analysis of daily changes in specialist inventories and quotations. Journal of Finance 48, 1595–1628.
Moulton, P., 2006. Who trades with whom? Cornell University Working Paper.
NYSE Group, 2006a. Hybrid Market training program. Booklet available online at: 〈https://www.nyse.com/pdfs/hm_booklet.pdf〉.
NYSE Group, 2006b. NYSE Hybrid FAQ. Document available online at: 〈http://exchanges.nyse.com/archives/2006/09/faqs_updated.php〉.
NYSE Group, 2007. NYSE completes Hybrid Market phase III activation. Press release available online at: 〈http://www.nyse.com/pdfs/HybridPhaseIV1.24.07.pdf〉.
Pastor, L., Stambaugh, R.F., 2003. Liquidity risk and expected stock returns. Journal of Political Economy 111, 642–685.
Rosu, I., 2009. A dynamic model of the limit order book. Review of Financial Studies 22, 4601–4641.
Securities and Exchange Commission, 2005. Regulation NMS. Available online at: 〈http://www.sec.gov/rules/final/34-51808.pdf〉.
Sofianos, G., Abrokwah, K., 2007. Executing quality at the new, fast NYSE. October 2, 2007, Goldman Sachs Equity Execution Strategies Report.
Sofianos, G., Werner, I., 2000. The trades of NYSE floor brokers. Journal of Financial Markets 3, 139–176.
The Economist, 2007a. Marketplaces on the move, September 13.
The Economist, 2007b. Ahead of the tape, June 21.
The Economist, 2008. The battle of the bourses, May 29.
Thompson, S., 2011. Simple formulas for standard errors that cluster by both firm and time. Journal of Financial Economics 99, 1–10.
Tobin, J., 1969. A general equilibrium approach to monetary theory. Journal of Money, Credit and Banking 1, 15–29.
Venkataraman, K., 2001. Automated versus floor trading: an analysis of execution costs on the Paris and New York Exchanges. Journal of Finance 56, 1445–1885.
Werner, I., 2003. NYSE order flow, spreads, and information. Journal of Financial Markets 6, 309–335.