Available online at www.sciencedirect.com
ScienceDirect
Journal of Financial Markets 17 (2014) 65–93
Journal of FINANCIAL MARKETS
www.elsevier.com/locate/finmar

Leveling the trading field

David Easley[a], Terrence Hendershott[b], Tarun Ramadorai[c,*]

[a]Departments of Economics and Information Science, Cornell University, Ithaca, NY 14850, USA
[b]Haas School of Business, 545 Student Services Bldg. ♯1900, University of California, Berkeley, CA 94720, USA
[c]Said Business School, Park End Street, OX1 1HP Oxford, UK

Received 3 September 2012; received in revised form 7 June 2013; accepted 7 June 2013
Available online 9 July 2013

Abstract

We examine the impact on stock prices of a major upgrade to the New York Stock Exchange's trading environment. The upgrade improved information dissemination on the trading floor and reduced the latency in reporting trades and quotes. The portion of the upgrade that reduced latency for electronic orders had significant impacts on liquidity, turnover, and returns. A portfolio that is long stocks undergoing the upgrade in the first 20 days of the upgrade and short stocks receiving the upgrade later has a return of roughly 3% over the period. The abnormal return was a priced effect of the improved liquidity produced by the upgrade.

JEL classifications: G12; G14

Keywords: Market structure; Stock exchange; Trading volume; Transparency; Information; Latency; Technology upgrade

1. Introduction

Competition in financial markets is central to the fairness and efficiency of the markets (Macey and O'Hara, 1999). Since Congress's passage of the 1975 National Market System legislation the U.S. Securities and Exchange Commission has promoted competition and transparency in financial markets, often through encouraging technological innovation. However, the benefits of such innovations are difficult to measure. In this paper, we test the hypothesis that increasing competition among traders is accompanied by increased liquidity and price appreciations. Using

*Corresponding author. Tel.: +44 1865288800.
E-mail addresses: dae3@cornell.edu (D. Easley), hender@haas.berkeley.edu (T. Hendershott), tarun.ramadorai@sbs.ox.ac.uk (T. Ramadorai).

1386-4181/$ - see front matter
http://dx.doi.org/10.1016/j.finmar.2013.06.003

a major change to trading technology by the New York Stock Exchange (NYSE) in 1980 we find significant effects on both liquidity and asset prices arising from a technological innovation that increased the ability of traders off the floor of the NYSE to compete on a more equal footing with traders on the NYSE floor.[1]

In 1980, the NYSE significantly upgraded its trading platform—specifically, the technology used at the posts at which trading occurred. The specific technological changes increased the ability of traders off the floor of the exchange [traders who used the designated order turnaround (DOT) and intermarket trading systems (ITS)] to compete with on-floor traders (the specialists and floor brokers) by providing off-floor traders faster execution and more up-to-date information on recent trades and quotes.[2] This increase in transparency and reduction in transaction latency—the time that elapses between an investor making a trading decision and the execution and confirmation of the desired trade—allowed off-floor traders to condition their orders on more up-to-date information and reduced the free trading option that their limit orders provide.[3] If off-floor traders required compensation for this risk of adverse selection in the form of a higher return premium, then reducing these costs should result in abnormal returns at the time of the technological change.

The event that we examine is well suited to examining the effects of a change in competition among traders for several reasons. First, most stock exchange innovations simultaneously affect all stocks traded on the exchange. This means that inferring the effects of the innovation cannot be done in the usual manner, for example by examining returns relative to the market. In contrast, the NYSE upgrade was implemented in a staggered fashion, allowing for better identification of its effects on liquidity and returns. The upgrade was done room-by-room on the exchange floor so there appears to be no bias in selecting which stocks were included in the first or second phase. Second, the change in latency was expected to be large: from two minutes pre-upgrade to less than 20 seconds post-upgrade. Third, the staggered implementation allows us to view the effects of two distinct types of changes to the trading environment. The first change that was implemented was what NYSE's internal documents referred to as “Phase I” of the upgrades. This phase improved the dissemination of information on the trading floor and reduced the latency of reporting quotes and floor transactions, but it did not directly affect the order matching process. The second change that was implemented, “Phase II,” was a superset of Phase I, with the addition of technology that significantly reduced transaction reporting and processing times for off-floor traders who used the designed order turnaround (DOT) system or intermarket trading system (ITS).[4] Both phases reduced latency and thus increased transparency, but the first one affected only orders originating on the floor of the exchange while the second one affected orders originating off the floor of the exchange. Phase I also improved the display of information

[1]While technology has generally improved access to equity markets and fostered such competition, recent developments like high-frequency trading, market fragmentation, and dark pools have raised concerns about a level playing field for investors. Transparency and equal access continue to be problematic in most decentralized dealer and over-the-counter markets. For example, see Bessembinderr, Maxwell, Venkataraman (2006) and Edwards, Harris, and Piwowar (2007) for corporate bonds and Green, Hollifield, and Schurhoff (2006) and Harris and Piwowar (2006) for municipal bonds.

[2]During this upgrade, the NYSE replaced the specialists' trading posts that had been in place since the 1920s with cogwheel-shaped trading posts employing newer technology. Section 2 and the Appendix describe the upgrade process in detail.

[3]Ready (1999) and Stoll and Schenzler (2006) discuss how slower traders' orders provide a free trading option for those traders with lower latency.

[4]The Appendix provides descriptions of how the market operated before the upgrade and how the upgrade affected its operation.

available to floor traders around the trading posts. Therefore, Phase II provides a more straightforward test of the impact of a more level playing field for off-floor traders. During the 'Retrofit,' which commenced after all Phase II upgrades were completed, posts that were upgraded during Phase I received the additional upgrades.

Our empirical tests compare stocks traded at specialist posts that experienced Phase I between July 14, 1980 and August 11, 1980, with stocks traded at posts that experienced Phase II between November 10, 1980 and January 5, 1981. For ease of exposition, we call the first group of stocks First Phase stocks, and the second group of stocks Second Phase stocks.[5] Providing better information through reducing latency for off-floor traders decreases their disadvantage vis-à-vis on-floor traders. So we would expect off-floor traders to trade more following such improvements. While we cannot directly measure the trading behavior of this group, we can measure changes in overall stock turnover around the period of the upgrades. We find that the turnover of Second Phase stocks increased relative to that of First Phase stocks in the period in which Phase II occurs. The difference in turnover between the two groups of stocks disappears following the implementation of the Retrofit program. These findings receive further support when we examine transactions costs estimates from the Roll (1984) model estimated using intra-day transactions data and the Gibbs sampler methodology of Hasbrouck (2009). We find that transactions costs decline by approximately 13 bps for both Second Phase and First Phase stocks over the period from April 1980 to July 1981. We also find that there is a relative increase of approximately 4 bps in the differential between the average spread of matched Second Phase and First Phase stocks in the period around the implementation of Phase II.

The liquidity effects induced by the innovation should also induce an effect on returns. Our most significant finding is that a portfolio that is long Second Phase stocks and short First Phase stocks exhibits a substantial positive return in the 20-day window beginning with the introduction of Phase II. That is, when the off-floor investors' ability to compete in Second Phase stocks was improved, and the competition by off-floor investors in First Phase stocks was unchanged, Second Phase stocks appreciated in price relative to First Phase stocks. Over the 20 days following the commencement of the Phase II upgrade, the return on the long Second Phase–short First Phase portfolio is approximately 20 bps a day for a total return of between 3% and 4%. This excess return is robust to correction for Fama-French, momentum, and industry factors. It is also robust to risk adjustment by matching stocks across phases, where the matching is done on market capitalization and prior turnover. This suggests that the technological upgrade that increased competition had an economically significant impact on stock returns. We also show that the increases in returns for Second Phase stocks line up cross-sectionally with the transaction cost reductions experienced by these stocks.

Additionally, we do not find significant liquidity or return effects for Phase I. This could be because the timing and knowledge of on-floor trades is somewhat discretionary, so off-floor traders receiving this information with low latency is less important. It is also possible that the improvements in the display of information at the trading posts in Phase I disadvantaged off-floor traders, counterbalancing any positive effects from the faster reporting of floor trades and quotes. Finally, we examine the Retrofit program, which completed all phases of the upgrade program for all stocks. This program eliminated the return differential between First and Second Phase stocks for the matched sample portfolio over the six-month period following its commencement.

When interpreting our results it is worth noting that the NYSE did not design the perfect experiment to test for the effect on prices of making the trading environment more competitive

[5] We discuss the details of sample construction in Section 5.

and fair. Most importantly, we cannot directly measure competition, so we can only infer changes in competition from the nature of the upgrades. Second, the upgrade process was announced well in advance of its implementation. So we face the usual difficulty of determining whether anticipation of the effects of the innovation caused price changes prior to the actual innovation. The NYSE's unfortunate track record on carrying out announced innovations actually helps us here. As we show, there was little price anticipation of the first part of the upgrade process, and once this first upgrade had a demonstrable impact, the impact of subsequent parts of the upgrade process seem to have been anticipated. We discuss how we address both these issues more fully in Section 4.

We discuss the surprisingly large magnitude of the return effects we discover in Section 7. We first use a simple present discounted value calculation that computes the benefits from the reduction in the Roll spread [as calculated using the Gibbs sampler method of Hasbrouck (2009)] to show that the capitalized value of the reduction in spreads can explain approximately one-fifth of the abnormal return. The remainder of the return may be coming from a combination of factors that are difficult, or perhaps impossible, to quantify reliably. First, as noted in Ready (1999) and Stoll and Schenzler (2006), delays in trade execution give the specialist valuable discretion to execute trades at prices that are not favorable to customers. Stoll and Schenzler (2006) formalize this intuition, positing that this trading option is actually a look-back option, which allows the specialist to transact at the most favorable price over the window of discretion. The upgrade significantly reduced this trading delay and thus could have had an effect through a reduction in the look-back option.[6] Second, the upgrades may have increased investor confidence in the market as a fair game and this in turn may have increased participation in the market and thus provided a one-time boost in returns. Finally, the success of these innovations may have increased investors' expectations that the NYSE would continue to make changes to improve liquidity and level the playing field between different types of traders, which would also increase prices and thus provide another one-time boost in returns.

The organization of the paper is as follows: Section 2 describes the upgrade process in more detail. Section 3 describes literature related to our study. Section 4 presents our methodology for examining the impacts of the post upgrades. In Section 5, we describe the data. Section 6 presents the results of our estimation. Section 7 is a discussion. Section 8 concludes.

2. The upgrade process

To understand how these technological changes reduced latency and improved transparency, it is helpful to begin with an overview of how the market operated prior to the upgrades. Specialists had limit order books,[7] made quotes to buy or sell, executed trades, and reported the trade information. Specialist's quotes were displayed on flip cards and, according to our conversations with floor personnel, were not always up-to-date or firm. To find the best quote, a floor trader had to walk up to the specialist and ask for a quote. Information about trades was publicly displayed, but often there was a substantial lag between execution and display of the data. At any moment, only the specialist and traders at or around his post were in possession of up-to-date trade or quote information. Additional details of the NYSE trading environment and post upgrades are in Appendix A.

[6] We provide a possible quantification of the effect in Appendix B.

[7] The specialist kept the limit order book on paper until the subsequent introduction of DisplayBook over the 1983–1988 period. At his discretion, the specialist could share the limit order book information with other floor traders.

Phase I of the upgrades increased transparency by improving the dissemination of quotes and the reporting of transactions. The specific technological changes implemented during this Phase were as follows: First, new optical card readers and an expanded price display unit were installed at each post so that trade and quote information could be entered and displayed at the post, rather than having the information sent off the floor to be processed. This reporting innovation improved the volume and clarity of quote and trade information for both on- and off-floor traders. The improved display at the trading post only benefited on-floor traders. In this phase, the information that was put through the new system did not include designated order turnaround (DOT) system or intermarket trading system (ITS) orders. The new equipment was, according to the exchange's estimates, meant to reduce the publication time of (non-DOT and non-ITS) transactions and quote information from two minutes to 20 seconds.[8] This innovation reduced the lags in reporting times for on-floor trades, thereby increasing the transparency of the quote and transaction process on the floor.

Phase II was a superset of Phase I, with the addition of technology that significantly reduced transaction processing times for off-floor traders. In Phase II, the eight card readers installed in each new post began processing DOT and ITS cards. This phase of the upgrade process reduced the latency experienced by off-floor traders, and thus reduced the informational advantage that the specialist and on-floor traders had over off-floor traders. Finally, posts that were upgraded during Phase I received the Phase II upgrades during the Retrofit program.

While Phase I improved information dissemination for on-floor traders, both Phases I and II reduced latency and increased the transparency of the transaction process. The primary difference between the phases in terms of reducing latency is which group of trades was affected. In Phase I, only orders originating on the exchange floor were affected while in Phase II orders originating off the floor of the exchange were also affected. The market microstructure literature is mixed on whether increased transparency should increase or decrease returns.

3. Related literature

Our study is related to research on market structure. Garbade and Silber (1978) examine how the SEC-mandated introduction of the consolidated tape affected price differentials between cross-listed stocks on the NYSE and regional markets. This potentially increased the speed of NYSE participants' access to information on prices in regional markets, thus reducing latency. However, Garbade and Silber find no effect, and argue that this is due to the prior existence of private networks for information transmission.

Amihud, Mendelson, and Lauterbach (1997) analyze stocks on the Tel Aviv Stock Exchange that move (in batches selected on the basis of ‘perceived marketability’) from a once a day call auction to an opening call auction followed by iterated continuous trading. They show that these stocks experienced increases in liquidity, and generated an abnormal return of about 5.5% over a 35-day window. Muscarella and Piwowar (2001) also find that stocks on the Paris Bourse that moved from a call market to continuous trading earned an abnormal return of more than 5% over 40 days, and exhibited significant increases in volume and liquidity. The change from call to

[8]The delays reflected the time the floor reporter took in filling out the cards and feeding them into the card readers. The 20-second delay is similar to the 16-second median delay in 1990 found by Hasbrouck, Sofianos, and Sosebee (1993) for floor trades. Floor trades in 1990 were reported in the same manner as all trades in 1980–1981. The reduction in reporting time with the post-upgrade reflects the more limited space for personnel and equipment in the old post design.

continuous trading both significantly reduces latency and intertemporally fragments trading. The results of these authors suggest that the former effect dominates the latter.

Jain (2005) employs cross-country monthly data to examine the impacts on the equity premium when stock markets move from floor trading to electronic trading. Jain's results provide strong support for the hypothesis that the switch precedes declines in emerging country equity premiums, but the results are not significant for the developed countries in his sample.

Battalio, Ellul, and Jennings (2007) show that reputational effects between specialists and floor brokers mitigate adverse selection costs in trading. They show that when specialists are relocated, and floor brokers do not move, liquidity costs increase and increase more for stocks with greater adverse selection. These results show that conditional on the NYSE market structure, reputational trading on the floor improves liquidity. Our focus is different. Our results suggest that changing the NYSE market structure to reduce the advantage that on-floor traders have over off-floor traders improves liquidity and increases stock prices.

There are several differences between these studies and our paper. First, our evidence demonstrates that the trading environment matters in one of the largest and most studied stock markets in the world. Second, the NYSE applied its upgrade to all stocks based on their location on the trading floor, so there is no apparent sample selection issue in our study. Finally, the NYSE did not change its method of trading, but rather, it augmented its pre-existing facilities in ways that clearly changed latency. This narrower focus allows us to draw sharper conclusions about how changing competition affects liquidity and returns.

Our interpretation of the impact of the upgrade draws upon Ready (1999), Stoll and Schenzler (2006), and Cespa and Foucault (forthcoming). Ready (1999) and Stoll and Schenzler (2006) analyze the discretion of intermediaries to execute orders at a time of their choosing. This discretion provides these agents with a free option in the presence of execution latency. The second phase of the post-upgrade reduced the latency experienced by off-floor market participants; therefore, it should reduce the premium demanded by these investors to compensate them for bearing this execution risk. The phases differed in which traders experienced a reduction in latency: on-floor traders for Phase I and off-floor traders for Phase II. So the impact of these phases on returns could differ because of this difference.[9] Cespa and Foucault (forthcoming) analyze a related aspect of latency—the fact that insiders have access to post-trade information prior to outsiders. They show that reductions in this aspect of latency may also increase prices by reducing the informational advantage that insiders have over outsiders, thus reducing the risk that outsiders face in their trading with insiders. Easley, O'Hara, and Yang (2012) show that, when exchanges provide differential access to trade information, liquidity is reduced, volatility is increased, and prices are lowered. These theoretical analyses thus provide additional support for our empirical analysis, which shows that enhancing competition between off-floor and on-floor traders are accompanied by price improvements.

An important aspect of Phases I and II was that they increased transparency as they both resulted in faster publication of trade and quote information. This increased pre- and post-trade transparency and improved the flow of information. The impact of changes in transparency is a hotly debated topic in the literature.[10] Our results shed some light on this debate by showing that

[9]Off-floor traders could have hired floor brokers to attempt to obtain some of the on-floor traders' benefits. Our results suggest that either the marginal trader setting prices was not successfully doing this, possibly because agency problems in hiring and monitoring floor brokers prevented off-traders from sharing in the potential gains from lower latency on the floor.

[10]On the one hand, Bloomfield and O'Hara (1999) use an experimental financial market to show that increased post-trade transparency increases bid-ask spreads. Furthermore, the empirical findings of Madhavan, Porter, and Weaver

the increased transparency resulting from the reduced latency in the upgrade process resulted in significant price appreciation.

4. Methodology

4.1. Matched portfolios in calendar time

In order to obtain clean estimates of the impact of the post upgrades, in addition to the use of our full First Phase and Second Phase stock samples, we match First Phase stocks with Second Phase stocks, based on the market capitalization and turnover of these stocks. We do this to ensure that our estimates are uncontaminated by variation in characteristics across First and Second Phase stocks, as distinct from estimates of factor covariances that might drive expected returns, for which standard regression methods provide natural controls.

To implement our matching procedure, we form market capitalization deciles based on the size of all First and Second Phase stocks, and then further divide these into quintiles ranked by turnover. We then randomly order First Phase stocks. For the first stock in this order, we select the Second Phase stock located in the same market capitalization and turnover cell with the closest market capitalization. If no such stock exists, the First Phase stock is dropped. We repeat this process until all First Phase stocks are exhausted. We then drop any pair for which the market capitalization differential or turnover differential is greater than 50%. We repeat this process 1,000 times with independent randomizations of the order of First Phase stocks. Our matched sample is generated by the repetition with the greatest number of matches.

4.2. Evidence on liquidity

Our hypotheses about the price effects of increased competition are predicated on the assumption that latency reductions will improve liquidity and that liquidity is priced. We therefore first check whether the technological upgrades are accompanied by increases in turnover and reductions in transactions costs. We consider six time periods around the upgrades during which the trading environment for First Phase and Second Phase stocks differs: pre-upgrade, pre-Phase I, Phase I, Phase II, Retrofit, and post-upgrade. Differences in liquidity should persist as long as the trading environment differs for First Phase and Second Phase stocks. For example, if Phase II increases liquidity, then Second Phase stocks should be more liquid than First Phase stocks from when the Second Phase stocks receive the Phase II upgrade until the First Phase stocks receive the Phase II upgrade (during the Retrofit period).

We measure turnover for a stock i on date t as the ratio of share volume over the day divided by that day's reported shares outstanding for the firm, where data on volume and shares outstanding are from the CRSP database. We estimate transactions costs using the Fitch intra-day transactions data (described in Section 5). Our measure of transactions costs is the spread from the Roll (1984) model, estimated using the Gibbs sampling methodology introduced

(*footnote continued*)
(2005), who use data from the Toronto Stock Exchange, suggest that increased pre-trade transparency also increases bid-ask spreads. On the other hand, Flood, Huisman, Koedijk, Mahieu (1999) analyze an experimental multiple-dealer market, and find that increased pre-trade transparency leads to lower spreads and increased trading volume. This is supported by Boehmer, Saar, and Yu (2005), who find that the introduction of the NYSE Open Book generates increases in stock market liquidity.

by Hasbrouck (2009). The model we estimate is

$$\Delta p_\tau = c\Delta q_\tau + u_\tau, \tag{1}$$

where Δp_τ is the price change between successive transactions, Δq_τ is the change in trade direction between buys ($q_\tau = 1$) and sells ($q_\tau = -1$), and c is the estimated transaction cost (spread). We assume that $u_\tau \sim NIID(0, \sigma_u^2)$ and compute Δp_τ from transaction data in each period. The unknowns, c, σ_u^2, and the sequence of trade direction indicators are estimated using the Gibbs sampler. In our empirical specifications, we use 10,000 'sweeps' of the Gibbs sampler, of which the 2,000 initial sweeps are discarded, to remove the effect of the starting values that we employ. These starting values are 1% for c and 0.001 for σ_u^2, as recommended by Hasbrouck (2009). We estimate the model a total of 1,212 times (once for each of the 202 matches for each of the six time periods that we consider). The total number of transactions employed in our estimation across all 202 stocks over the six time periods is 2,009,243.[11]

Armed with our measures of transaction costs and turnover, our approach is to compute differences-in-differences of turnover and estimated transaction costs over various time periods before, during, and after the upgrades. Take turnover as an example: we first compute the average turnover of matched Second Phase stocks in the period when they receive the Phase II upgrade. We subtract from this estimate the average turnover of these stocks in the pre-upgrade period. This is the first level of differencing, in which we control for the 'normal' turnover of Second Phase stocks. We then compute this difference for the matched First Phase stocks, obtaining an estimate of the change in turnover of these stocks relative to pre-upgrade turnover in each period. Finally, we subtract off this difference in turnover for the First Phase stocks from the difference in turnover for the Second Phase stocks. This is our difference-in-difference estimator of the impacts of the upgrade phases on movements in turnover. We compute the p-values of these differences-in-differences using a heteroscedasticity and autocorrelation robust estimator.

4.3. Long-short portfolios in calendar time

There are at least two problems with using event-study methodology to test our hypotheses about returns. The first is that the events occur very close to one another in calendar time, with many of the stocks experiencing upgrades on the same day. This generates contemporaneous cross-sectional correlation in the data. As Kothari and Warner (2005) point out, "…if the test statistic in an event study is calculated ignoring cross-dependence in data, even a fairly small amount of cross-correlation in data will lead to serious misspecification of the test. In particular, the test will reject the null of no effect far more often than the size of the test…" One standard solution to this problem, implemented in papers such as Brav and Gompers (1997), Lyon, Barber and Tsai (1999), Fama (1998), and Mitchell and Stafford (2000), is to group stocks into a portfolio in calendar time, and to regress the returns on this portfolio on a set of factors. The "Jensen's alpha" of this portfolio is an estimate of the abnormal return.

The second problem with conducting inference using an event study is that the specific event that we investigate affects every stock in the market, over a reasonably short time period. This means that the benchmark that we use to evaluate abnormal performance is itself changing over time. This renders accurate measurement of abnormal returns difficult. A simple example

[11]The results for both turnover and the Roll spread are not sensitive to shortening the time periods to correspond to those for the return analysis. This is despite the fact that Gibbs estimates of the Roll spread are more difficult to estimate over the shorter time windows, particularly for the smaller stocks.

illustrates this point: assume that half the stocks in the market portfolio undergo a change that generates abnormal returns over a 60-day period. Abnormal returns for any remaining stocks, measured relative to the upward-moving market portfolio, will be biased downward. To insulate ourselves from this problem, we modify the calendar time procedure and measure the returns of stocks experiencing the two phases of the upgrade relative to one other. We implement this control by creating portfolios that are long Second Phase stocks and short First Phase stocks. One disadvantage of our long-short approach is that the relative returns on our portfolio cannot be uniquely attributed to returns on the long or short components of the portfolio. We discuss this when we examine returns. We also examine the long and short portfolio separately and we use our evidence on liquidity to help us sort out which component of the portfolio is responsible for the returns we observe.

It is conceivable that specialists, floor brokers, and investors anticipated the effects of these innovations, and so some of the effects may have taken place before the innovation process was actually implemented. It is also possible that it took market participants some time to understand how to best use the technology. Thus some of the effects of the innovations may not have appeared until after the innovations occurred. In order to control for both of these possibilities, we examine the returns to the long-short portfolios over various trading months (20-day windows) before, during, and after the upgrade process.[12]

We assess the impacts of the upgrades on returns, estimating the following equation:

$$(R_{SecondPhase,t} - R_{FirstPhase,t}) = \sum_{i=1}^{10} a_i D_i + \beta_{Rm}(R_{mt} - R_{ft}) + \beta_{SMB} SMB_t + \beta_{HML} HML_t + \beta_{UMD} UMD_t + u_t. \tag{2}$$

The left-hand side variable in both equations is created by subtracting the mean return across all First Phase stocks ($R_{FirstPhase,t}$) from the mean return across all Second Phase stocks ($R_{SecondPhase,t}$) each day. The ten dummy variables, $D_1, \ldots, D_{10}$, pick out the pre-and post-periods surrounding the rollouts of Phase I, Phase II, and the Retrofit program.[13] The returns on this long-short portfolio are risk-adjusted by regressing it on the Fama-French-Carhart four-factor model.[14] The coefficients $a_1, \ldots, a_{10}$ can therefore be viewed as the average abnormal returns to a trading strategy that goes long Second Phase stocks and short First Phase stocks over the periods selected by the dummy variables, $D_1, \ldots, D_{10}$.

In our regressions, the coefficients a_2, a_3 measure the difference between the returns of First and Second Phase stocks in the 20-day periods prior to and just following the implementation of Phase I for First Phase stocks. An effect in the period prior to the event (a_2) would be consistent with the market anticipating the impact of the change, while an effect in the period subsequent to the implementation of Phase I (a_3) corresponds to the unanticipated impact of this change. Estimated coefficients (a_2, a_3) greater than (less than) zero correspond to a negative (positive) return impact of improving the dissemination of quotes and the reporting of floor transactions.

Similarly, coefficients a_5, a_6 measure the difference between the returns of First and Second Phase stocks in the 20-day periods prior to and just following the implementation of Phase II for

[12]We discuss both anticipation and the use of 20-day windows in more detail in Section 6.2.

[13]We used only six time periods for liquidity, but prices are forward looking so the upgrades' impact on returns may be concentrated around the upgrade events, making narrower time windows appropriate for returns.

[14]Allowing for time variation in the loadings on the Fama-French-Carhart factors (that is, allowing for a different coefficient on each factor for each time period), does not qualitatively alter the results. The results are also robust to the inclusion of the five Fama-French industry portfolios as factors, suggesting that our results are not driven by differences in the industry group to which First Phase and Second Phase stocks belong.

Table 1
Summary statistics: full sample.

The column headings denote the summary statistics calculated over the time period prior to the upgrade (1/1/79–4/14/80), except for the statistics pertaining to market capitalization and price, which are calculated as on 4/14/80. Daily turnover is calculated as share volume divided by daily shares outstanding reported in CRSP. The turnover and return betas (from a market model, described in the text) are calculated only for those stocks with more than 250 return observations, over (1/1/79–4/14/80), and then averaged across the group of stocks denoted in the row headings. The row headings denote the group of stocks for which the statistics are calculated (First and Second Phase stocks, and 'Other' stocks, which are not in either of our Phases). The summary statistics are first calculated across all stocks in each of the phases, and then for the large, medium, and small stocks (sorted into terciles by market capitalization as on 4/14/80) separately.

Stocks	Number	Mean Mkt. cap ($MM)	Mean price ($)	Mean return (%)	Std. return	Mean beta return	Mean turnover (%)	Std. turnover (%)	Mean beta turnover
All:									
First Phase	159	486	20.451	0.062	2.258	1.020	0.187	0.220	1.058
Second Phase	149	835	24.318	0.064	2.250	1.060	0.197	0.235	1.031
Others	1118	554	21.518	0.058	2.227	0.985	0.173	0.199	0.985
Small:									
First Phase	50	40	10.796	0.033	2.662	1.011	0.180	0.238	1.009
Second Phase	44	42	12.517	0.047	2.979	1.190	0.233	0.342	1.242
Others	381	41	11.337	0.035	2.726	1.013	0.186	0.239	1.077
Medium:									
First Phase	58	192	19.050	0.078	2.317	1.061	0.219	0.263	1.143
Second Phase	42	200	23.232	0.082	2.233	1.068	0.208	0.228	1.203
Others	376	180	20.829	0.070	2.215	0.988	0.185	0.216	1.078
Large:									
First Phase	51	1257	31.510	0.071	1.795	0.983	0.158	0.154	1.008
Second Phase	63	1813	33.284	0.063	1.752	0.973	0.164	0.165	0.786
Others	361	1484	32.981	0.069	1.714	0.952	0.148	0.139	0.799

Second Phase stocks. Again, an effect in the period prior to the event (a_5) would be consistent with the market anticipating the impact of the change, while an effect in the period subsequent to the implementation of Phase II (a_6) corresponds to the unanticipated impact of this change. If the increased competition generated by Phase II had positive impacts on prices, either or both of these coefficients would be estimated to be positive and statistically significant.

Finally, coefficients a_8, a_9, and a_{10} correspond to the periods surrounding the initiation of the Retrofit program. As the Retrofit program essentially eliminated the difference between First Phase and Second Phase stocks (by providing the Phase II upgrades to First Phase stocks) we expect these coefficients (a_8, a_9, and a_{10}) to have the opposite sign to coefficients a_5 and a_6. The dates that we employ are based on the scheduled initiation of this program as per available NYSE documentation. Unfortunately, exact implementation dates for the Retrofit program for specific stocks are unavailable, making precise estimation of the Retrofit program's effects more difficult.

In our specifications, we use equal-weighted mean returns when computing the returns of the long Second Phase, short First Phase portfolio. We employ Newey and West (1987) standard errors throughout our regressions. These are robust to heteroscedasticity and autocorrelation. We also conduct the same analysis as before using the matched sample, that is, we re-estimate Eq. (2), substituting the difference between the equal-weighted returns of matched Second Phase–First Phase pairs as the left hand-side variables. We do this together with examining the returns of the long and short portfolio separately to assuage concerns about the robustness of our results. Given the differences in stock *characteristics* (which may not be adequately controlled for by the factor *covariances* that we estimate in our long-short return regressions) between the First Phase and Second Phase stocks in Table 1, it is important to show robustness to the use of this matched sample (e.g., Daniel and Titman, 1997).

4.4. Cross-sectional analysis

To examine whether any abnormal returns line up cross-sectionally with changes in transaction costs experienced by these stocks we estimate:

$$R_{SecondPhase,i} - R_{FirstPhase,i'} = \alpha + \beta_1(\Delta c_{SecondPhase,i} - \Delta c_{FirstPhase,i'}) + Controls + u_{ii'}. \tag{3}$$

The left-hand side variable in Eq. (3) is the difference in returns between each Second Phase stock i and its matched First Phase stock i', in the period following the implementation of Phase II. The right-hand side variable is the difference in the estimated Roll spread of each Second Phase stock i (measured as a change relative to the pre-upgrade period) and its match. The control variables are the differences for each pair of the return standard deviations, log (market capitalization), and average turnover. Our hypothesis is that the return movements, if any, on account of the off-floor transparency/latency improvements, should line up with the decreases in transaction costs experienced by each stock undergoing the transparency/latency improving upgrade. We therefore expect a negative coefficient, β_1.

5. Data

Our empirical tests compare First Phase stocks, which experienced Phase I between July 14, 1980 and August 11, 1980, with Second Phase stocks, which experienced Phase II between November 10, 1980 and December 8, 1980. These groups of stocks were chosen with two purposes in mind: our First Phase stocks experience Phase I well before our Second Phase stocks undergo Phase II. This separation in time allows the market time to adjust to the new trading

Table 2
Summary statistics: matched sample.

The column headings denote the summary statistics that are calculated over the time period prior to the upgrade (1/1/79–4/14/80), except for the statistics pertaining to market capitalization and price, which are calculated as on 4/14/80. Daily turnover is calculated as share volume divided by daily shares outstanding reported in CRSP. The turnover and return betas (from a market model, described in the text) are calculated only for those stocks with more than 250 return observations, over (1/1/79–4/14/80), and then averaged across the group of stocks denoted in the row headings. The row headings denote the group of stocks for which the statistics are calculated: First Phase stocks or Second Phase stocks. Each First Phase stock is matched with a Second Phase stock based on market capitalization and turnover. The summary statistics are first calculated across all matched stocks in each of the rooms, and then for the large, medium and small stocks (breakpoints computed using the market capitalization of Second Phase stocks as on 4/14/80, breakpoints then used to sort First Phase stocks) separately.

Stocks	Number	Mean Mkt. cap ($MM)	Mean price ($)	Mean return (%)	Std. return (%)	Mean beta return	Mean turnover (%)	Std. turnover (%)	Mean beta turnover
All:									
First Phase	101	453	19.861	0.056	2.297	1.059	0.183	0.215	1.145
Second Phase	101	426	22.665	0.067	2.232	1.019	0.183	0.229	1.058
Small:									
First Phase	33	41	10.741	0.042	2.831	1.178	0.195	0.252	1.353
Second Phase	33	43	13.750	0.052	2.796	1.099	0.207	0.320	1.225
Medium:									
First Phase	33	208	19.663	0.071	2.247	1.010	0.204	0.243	1.204
Second Phase	33	201	23.784	0.085	2.162	1.011	0.193	0.219	1.165
Large:									
First Phase	35	1073	28.646	0.055	1.840	0.998	0.152	0.155	0.908
Second Phase	35	1001	30.014	0.064	1.767	0.959	0.151	0.153	0.818

environment. Second, we choose our sample to ensure that all stocks within each group experience the respective upgrades within 20 trading days of each other.[15]

Using data from CRSP, Table 1 presents descriptive statistics for our stocks. There are 159 stocks in our First Phase group and 149 in the Second Phase group, leaving 1,118 "Other" stocks not in either of the two groups.[16] The average size of Second Phase stocks is U.S. $835 million, which is higher than that of First Phase stocks, at $486 million. The mean market capitalization of the other stocks lies in between these two numbers. However, in the two smallest terciles of stocks, the mean market capitalization is virtually identical for First Phase and Second Phase stocks. The mean turnover and turnover beta of First Phase stocks is similar to that of Second Phase stocks, and both turnover and return betas are close to one on average in both groups, across all size quintiles of stocks.[17]

Table 2 presents summary statistics from pairs of stocks matched on turnover and market capitalization across First Phase and Second Phase stocks. There are a total of 101 matched pairs of stocks (i.e., 202 stocks overall). The means of market capitalization and turnover across the two groups are very close, a result of the matching procedure.

To supplement the daily CRSP data, we use transaction data supplied by Francis Emory Fitch, Inc. The database is similar to the trade file in the NYSE's Trade and Quote (TAQ) database and consists of a time-ordered record of every common stock transaction on the NYSE for 1980 and 1981. Data on bid and ask quotes is unavailable for our sample period. Because the changes to the specialist post primarily affected continuous trading, the opening and closing transactions are excluded. Transaction data from Fitch is also used in Glosten and Harris (1988).

6. Results

6.1. Evidence on liquidity

Table 3 shows the evolution of turnover over the periods before, during, and after the upgrades. There are several points of interest here. First, mean turnover increases significantly for both First Phase and Second Phase stocks in the post-upgrade period relative to the pre-upgrade period. The increase in turnover is on the order of 0.49%, representing an increase of approximately 36% (33%) following the upgrades, for First Phase (Second Phase) stocks. Second, the difference-in-difference estimates of increases in turnover for the matched sample

[15]Our results are not sensitive to expanding either or both of the groups to include additional stocks experiencing the respective upgrades more than 20 days after the initial introduction of each phase. For example, expanding the set of Second Phase stocks to include those stocks experiencing upgrades on December 15, 1980 and January 5, 1981 leaves the results qualitatively unchanged.

[16]The other category is large for two reasons. First, we require stocks in the First and Second Phase portfolios to experience the upgrades close to the same time as the other stocks in same portfolio, but far from the time of the upgrades to stocks in the other portfolio. This requirement eliminates more than one-half of the stocks from our study; it places them in the other portfolio. Second, there was some reorganization of stocks' locations during the upgrade for which exact dates are not available (see Appendix A). See Battalio, Ellul, and Jennings (2007) for an analysis of how relocations impact "reputational" trading on the NYSE.

[17]It is not surprising that stocks in the different groups look fairly homogeneous as these groups are based on their location on the floor. Historically, the NYSE's limited floor space and the limited processing power of the human floor traders provided incentives for stocks with different characteristics to be uniformly located across the floor. This ensures "load balancing" of capacity so traders and locations on the floor are not overwhelmed when news arrives and/or trading activity increases. This incentive to assign stocks evenly across the floor is helpful in mitigating selection biases over the upgrade process.

Table 3
Turnover over time.

This table presents the average turnover (dollar trading volume as a percentage of stock market capitalization) for matched sample stocks in the time periods bracketed by the dates in the columns entitled 'Begin' and 'End.' The next two columns denote the set of stocks for which upgrade phases are implemented in each of the time periods, that is, whether Phase I and/or Phase II have been implemented for First and Second Phase stocks. The columns entitled 'Stock Sample' denote the set of stocks for which the measures are calculated. The column entitled "$\Delta a_i - \Delta a_1$" presents the difference-in-difference estimates for annualized average turnover in each time-period, and the column entitled p-value shows the statistical significance of this difference computed using a heteroscedasticity and autocorrelation-robust Newey-West covariance matrix. Thus, in the time period between 7/14/80 and 11/7/80, Phase I was implemented for First Phase stocks, and the annualized average turnover over this period was 0.196 for First Phase stocks, and 0.208 for Second Phase stocks. $\Delta a_i - \Delta a_1$ is computed by subtracting from this difference (0.208−0.196), the prevailing difference between annualized average turnover in the Pre-Phase I period (0.144−0.140). Differences significant at the 10% level are marked with asterisks.

			Implementation status			Stock sample			
Period	Begin	End	First Phase	Second Phase	Coeff.	First Phase	Second Phase	$\Delta a_i - \Delta a_1$	p-Value
Pre-upgrade	4/14/80	6/12/80			$a_{1,p}$	0.140	0.144		
Pre-Phase I	6/13/80	7/11/80			$a_{2,p}$	0.177	0.180	−0.001	0.981
Phase I	7/14/80	11/7/80	I		$a_{3,p}$	0.196	0.208	0.008	0.260
Phase II	11/10/80	1/16/81	I	I & II	$a_{4,p}$	0.196	0.229	0.029*	0.006
Retrofit	1/19/81	2/13/81	I & II	I & II	$a_{5,p}$	0.151	0.168	0.013	0.219
Post-upgrade	2/16/81	7/13/81	I & II	I & II	$a_{6,p}$	0.190	0.191	−0.003	0.674

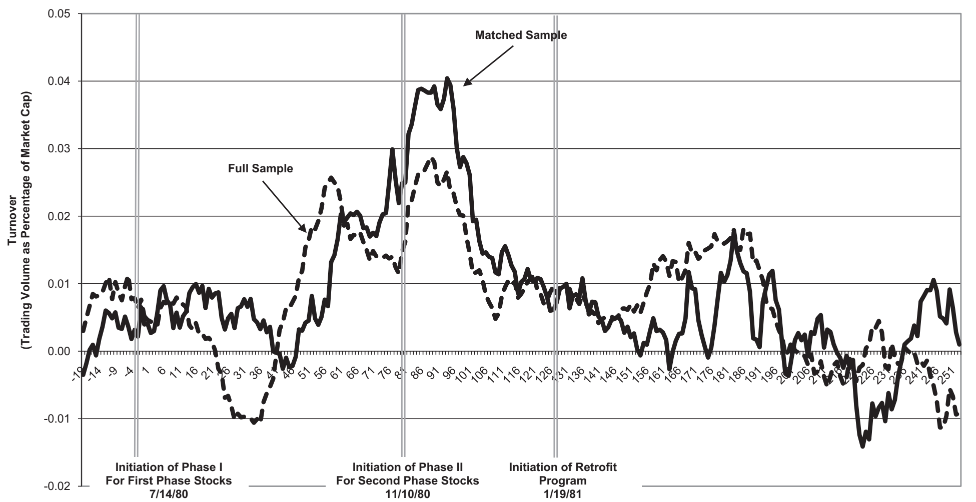

Fig. 1. Calendar time turnover of long-short portfolio. This figure plots the 20-day centered moving average of the difference in turnover (dollar trading volume as a percentage of stock market capitalization) between Second Phase and First Phase stocks for the full and matched samples. For the full sample, turnover values for First Phase and Second Phase stocks are adjusted by subtracting off their pre-period means, as reported in Table 1. The vertical bars show the implementation dates of various phases highlighted in the legend beneath each graph.

Table 4
Transactions costs over time.

This table presents transactions costs estimates for matched sample stocks from the Roll (1984) model of transactions prices (estimated using a Gibbs sampler as outlined in the text), for stocks in the time periods bracketed by the dates in the columns entitled 'Begin' and 'End.' The next two columns denote the set of stocks for which upgrade phases are implemented in each of the time periods, that is, whether Phase I and/or Phase II have been implemented for First and Second Phase stocks. The columns entitled 'Stock Sample' denote the set of stocks for which the measures are calculated. The column entitled "$\Delta a_i - \Delta a_1$" presents the difference-in-difference estimates for transactions costs in each time-period, and the column entitled p-value shows the statistical significance of this difference computed using a heteroscedasticity and autocorrelation-robust Newey-West covariance matrix. Thus, in the time period between 7/14/80 and 11/7/80, Phase I was implemented for First Phase stocks, and the average estimated spread was 0.447% for First Phase stocks, and 0.420% for Second Phase stocks. $\Delta a_i - \Delta a_1$ is computed by subtracting from this difference (0.420−0.447), the prevailing difference between the estimated spreads in the Pre-Phase I period (0.538−0.562). Differences significant at the 10% level are marked with asterisks.

			Implementation status		Stock sample				
Period	Begin	End	First Phase	Second Phase	Coeff.	First Phase	Second Phase	$\Delta a_i - \Delta a_1$	p-Value
Pre-Upgrade	4/14/80	6/12/80			$a_{1,p}$	0.562	0.538		
Pre-Phase I	6/13/80	7/11/80			$a_{2,p}$	0.505	0.470	−0.012	0.594
Phase I	7/14/80	11/7/80	I		$a_{3,p}$	0.447	0.420	−0.004	0.839
Phase II	11/10/80	1/16/81	I	I & II	$a_{4,p}$	0.492	0.427	−0.042*	0.046
Retrofit	1/19/81	2/13/81	I & II	I & II	$a_{5,p}$	0.486	0.450	−0.012	0.587
Post-Upgrade	2/16/81	7/13/81	I & II	I & II	$a_{6,p}$	0.431	0.405	−0.003	0.906

reveal that the period around the implementation of Phase II for Second Phase stocks exhibits a relative increase in turnover for Second Phase stocks. This difference is significant at the 1% level, and is on the order of 0.03%, which is approximately 20% of the pre-upgrade average turnover of the Second Phase stocks.[18] The transparency/latency improving upgrade appears to be accompanied by movements in trading volume of the predicted sign.

The plot of the difference in turnover between Second Phase and First Phase stocks is given in Fig. 1. Prior to the beginning of Phase I, there is little difference in turnover between First and Second Phase stocks. Leading up to the introduction of Phase II for Second Phase stocks, the turnover for Second Phase stocks increases relative to that for First Phase stocks for both the entire sample and the matched sample. This turnover differential persists until the Retrofit program begins.

Table 4 shows the results for the matched sample estimates of transaction costs from the Roll (1984) model estimated using the Gibbs sampler. The average estimated spread in the pre-upgrade period is on the order of 0.55%, which corresponds roughly to the average spread for all but the smallest stocks reported in Hasbrouck (2009). Similar to the turnover estimates, we see evidence of a large decline in transaction costs in the post-upgrade period relative to the pre-upgrade period. The estimated transaction cost moves to around 0.42% in the post-upgrade period, a decline of 0.13, which is 24% of the pre-upgrade spread. Furthermore, the difference-in-difference estimator reveals that the Phase II implementation generated a large difference between the estimated spread for Second Phase stocks and that for First Phase stocks, with

[18] Another natural test is to compare differences in turnover between adjacent periods rather than comparing to the pre-upgrade period. Relative to Phase I implementation, in Phase II implementation First Phase stocks exhibit no change in turnover and Second Phase stock's turnover increases by 0.021%, with a p-value of 0.047.

Second Phase stocks' transactions costs lower by 4.2 bps.[19] This estimator also reveals that Phase I produced no relative decline in spread for First Phase stocks.

Some caution is in order in interpreting our measures of changes in turnover and transaction costs. First, we have shown that between the pre-upgrade period and the post-upgrade period, turnover increases, and transaction costs decline. We believe that the upgrade was responsible for some of this change, but as we have no base to use to normalize this change, we cannot attribute all of the movement to the upgrades. If there were secular changes over this period, then they too may have contributed to our measures. In fact, during Phase I there was a cumulative value-weighted return to the market of about 20%. One would expect this to go hand-in-hand with an increase in turnover and perhaps a decline in transaction costs. Tables 3 and 4 show that there is an increase in turnover and a decline in transaction costs in the pre-Phase I and Phase I stages for both First Phase and Second Phase stocks. Most importantly, there is also an increase in estimated transaction costs for both First Phase and Second Phase stocks during Phase II. During Phase II, the overall market was quite volatile (there was a 9.8% decline in the NYSE Composite Index during nine trading days of this phase), and this may have contributed to the increase in estimated transaction costs.

We view the difference-in-difference method as providing a control for market-wide events. Although overall transaction costs increase during Phase II, they increase significantly less for Second Phase stocks. So we find significant relative results during the implementation of Phase II. These changes remain significant after including the secular changes affecting the market as a whole over this period. We believe that Second Phase stocks have higher relative turnover increases and lower relative transaction cost changes because of the beneficial effects of Phase II for these stocks. After examining whether the changes in liquidity for Second Phase stocks is associated with higher relative returns, we conduct cross-sectional analysis on the changes in liquidity and changes in returns to provide additional support for this interpretation.

6.2. Calendar time results

The time interval over which we measure abnormal returns following an upgrade should be approximately the same length as the interval over which we require that all stocks in a portfolio receive the upgrade. Table A.1 reveals that the first three introductions of Phase II occur within 20-trading day periods of one another. Fig. 2 plots the abnormal returns to portfolios of matched stocks (Table 2) receiving the Phase II upgrade on 11/10/1980, 11/17/1980, and 12/08/1980, lining these abnormal returns up around the implementation dates of the upgrades. Abnormal returns are created by subtracting off the returns of matched stocks identified using the procedure described in Section 5.

Fig. 2 reveals that the use of 20-day windows captures the essential features of the dynamics of returns over the period of the upgrades. The roughly 20-day gap between the upgrades for the stocks receiving them is mirrored in their abnormal return patterns. The figure is also consistent with a learning story, i.e., as soon as the first group of stocks receives the upgrade on 11/10/1980, we see simultaneous movements in the returns of stocks that are slated to receive the upgrades over the subsequent weeks. It seems as though the implementation of the upgrades for the first

[19]Relative to Phase I implementation as opposed to the pre-upgrade period, the difference-in-difference transactions costs estimate for Phase II implementation is 3.8 bps with a p-value of 0.076. Both difference-in-difference estimates are being driven by the larger increase in estimated transaction costs for First Phase stocks during Phase II. Overall transaction costs were rising during this period, but they increased less for Second Phase stocks than they did for First Phase stocks. We discuss this point in the cautionary notes outlined in the next paragraph.

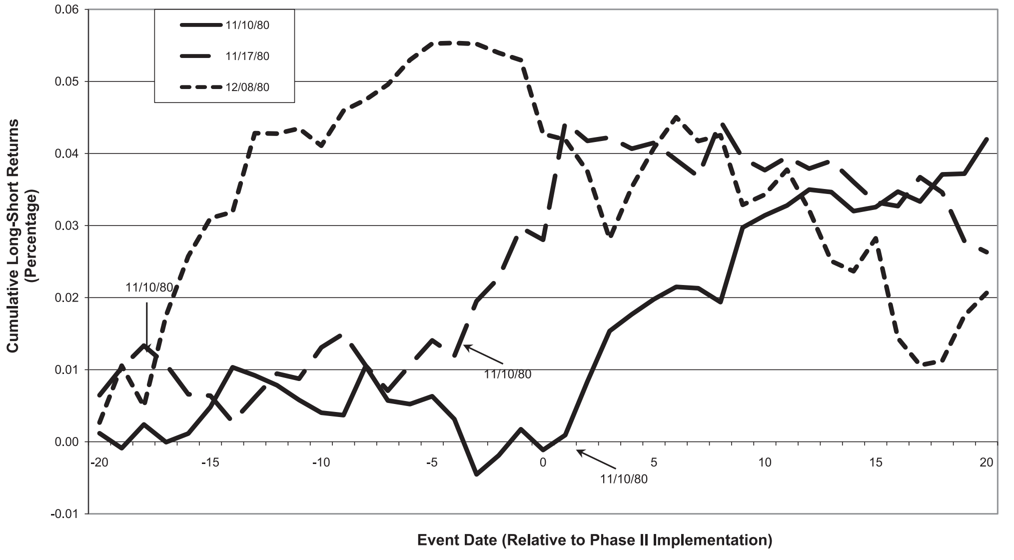

Fig. 2. Returns to long-short portfolio centered on phase implementation dates. This figure plots the cumulative difference in abnormal equal-weighted returns between Second Phase and First Phase stocks for the matched sample. The Phase II upgrade was implemented for different stocks in the Second Phase sample on November 10, 1980, November 17, 1980, and December 8, 1980 (see Appendix A). Three portfolios are formed of Second Phase stocks receiving the upgrade on the same date. The implementation date is defined as day 0. Returns are calculated net of the matched sample returns. The arrows indicate the event date for each line corresponding to the calendar date November 10, 1980, i.e., the date Phase II upgrades began.

group of stocks served the same function as an announcement, generating price movements for the stocks about to receive the upgrades in subsequent weeks.

Therefore, we construct the portfolios so that all stocks in any portfolio receive the upgrade within 20 trading days (roughly one calendar month) of each other and we measure the abnormal returns over 20-trading day periods. The choice of 20 days is intended to ensure that the events are far enough apart so that the two portfolios can serve as controls for one another, and so that the time period over which we measure returns is relatively concentrated around the upgrade event.

Table 5 presents estimates of Eq. (2) using both the full and matched samples. In this table, a_6 (the coefficient on the dummy variable for the introduction of Phase II) is positive and statistically significant, for all stocks in both the full and matched samples, and for the long sides of the portfolios (but not the short sides), as well as for the long-short portfolio. These results lend strong support to the hypothesis that the liquidity improvements generated by the transparency/latency improving upgrade, i.e., improving off-floor traders' ability to compete, were priced into stocks.

The coefficients translate into relatively large abnormal returns to a trading strategy that goes long Second Phase stocks and short First Phase stocks during the introduction of Phase II. For the 20 days following the rollout of Phase II, a_6 shows that the average Second Phase stock outperforms the average First Phase stock by 19.2 bps per day. This translates into a cumulative abnormal return of 3.84% over the 20-day period, after risk-adjustment using the four Fama-French factors.[20] Importantly,

[20]Consistent with the returns in Fig. 2, looking at horizons beyond the 20-day window after the Second Phase lowers the estimated effect of the Second Phase as a_7 is less than zero, although not statistically significant. Aggregating a_6 and a_7 gives a cumulative abnormal return for the second phase of $3.68\% = 20 \times 19.2 - 8 \times 0.02$ when adjusting for the fact that the period corresponding to a_7 is only 8 trading days. The coefficient for the 20 days prior to Phase II (a_5) is positive,

Table 5
Long-short, short, and long portfolio returns: four-factor model adjusted returns.

This table presents results from estimating three related regressions, namely the long second phase, short first phase portfolio regressed on the determinants presented below: $(R_{SecondPhase,t} - R_{FirstPhase,t}) = \sum_{i=1}^{10} a_i D_i + \beta_{Rm}(R_{mt} - R_{ft}) + \beta_{SMB} SMB_t + \beta_{HML} HML_t + \beta_{UMD} UMD_t + u_t$, as well as two additional regressions showing the results of separate regressions for the long and the short sides of the portfolio above. The D_i select the events and time periods denoted in the first column, and the remaining regressors are the four Fama-French factors. The first (second) group of three columns presenting results are for the full (matched) sample, and the subheadings show the portfolio for each set of reported coefficients, that is, Second Phase minus First Phase (Long-Short), First Phase (Short), Second Phase (Long). In each case, the returns are the equal-weighted average return across all stocks in each group, that is, First Phase stocks or Second Phase stocks on day t. Newey-West t-statistics are robust to heteroscedasticity and autocorrelation, and presented in italics beneath coefficients. Coefficients significant at the 10% level are marked with asterisks.

	Implementation status			Full sample			Matched sample		
Period	First Phase	Second Phase	Coeff.	Long-Short	Short	Long	Long-Short	Short	Long
Pre-Upgrade			a_1	0.004	0.043*	0.047*	0.012	0.039*	0.051*
1/1/79–6/12/80				*0.282*	*4.084*	*4.006*	*0.726*	*3.067*	*3.858*
Phase I (−19,0)			a_2	0.025	0.034	0.059	0.014	0.043	0.057
6/13/80–7/11/80				*0.421*	*0.699*	*1.823*	*0.155*	*0.792*	*1.448*
Phase I (+1,+20)	I		a_3	0.062	−0.014	0.048	0.111	−0.047	0.064
7/14/80–8/8/80				*0.713*	*−0.286*	*0.856*	*1.178*	*−0.952*	*0.887*
Between Phases	I		a_4	−0.012	0.037	0.025	−0.016	0.042	0.026
8/11/80–10/9/80				*−0.383*	*1.661*	*0.878*	*−0.212*	*1.336*	*0.655*
Phase II (−19,0)	I		a_5	0.038	0.070*	0.108*	0.100	−0.016	0.084*
10/10/80–11/7/80				*0.667*	*1.967*	*2.451*	*1.627*	*−0.297*	*2.415*
Phase II (+1,+20)	I	I & II	a_6	0.192*	−0.038	0.155*	0.152*	0.017	0.170*
11/10/80–12/7/80				*3.517*	*−0.695*	*3.375*	*2.615*	*0.287*	*3.370*
Post-Phase II	I	I & II	a_7	−0.020	0.071	0.050	−0.045	0.104	0.060
12/8/80–12/17/80				*−0.254*	*0.801*	*0.914*	*−0.561*	*1.250*	*1.336*
Retrofit (−19,0)	I	I & II	a_8	−0.034	0.065*	0.031	−0.040	0.056	0.016
12/18/80–1/16/81				*−0.647*	*2.237*	*0.615*	*−0.502*	*1.134*	*0.298*
Retrofit (+1,+20)	I & II	I & II	a_9	−0.051	0.121*	0.070	−0.079	0.140*	0.061
1/19/81–2/13/81				*−0.669*	*4.430*	*0.952*	*−1.417*	*5.300*	*0.827*
Post-Upgrade	I & II	I & II	a_{10}	0.003	0.063*	0.066*	−0.033	0.070*	0.037
2/16/81–7/13/81				*0.103*	*2.971*	*2.298*	*−1.368*	*2.579*	*1.209*
			$R_m - R_f$	0.048*	1.127*	1.175*	−0.294	1.182*	1.153*
				1.907	*60.919*	*56.634*	*−0.962*	*47.620*	*52.289*
			SMB	−.067*	0.644*	0.578*	−0.080	0.736*	0.656*
				−1.706	*22.563*	*23.097*	*−1.193*	*15.400*	*20.341*
			HML	−.109*	0.177*	0.068*	−0.114*	0.217*	0.103*
				−2.170	*5.201*	*1.747*	*−2.027*	*5.136*	*2.587*
			UMD	−0.046	.162*	−0.208*	−0.024	−0.181*	−0.205*
				−1.476	*−7.661*	*−8.266*	*−0.658*	*−6.542*	*−7.343*
			R^2	0.082	0.945	0.939	0.042	0.924	0.918
			N	639	639	639	639	639	639

these returns are virtually all driven by the long side of the portfolio rather than the short side, which

(*footnote continued*)
but not statistically significant. The sum of a_6 and a_7 and a_5 through a_7 are statistically different from zero at the 10% level, but not at the 5% level.

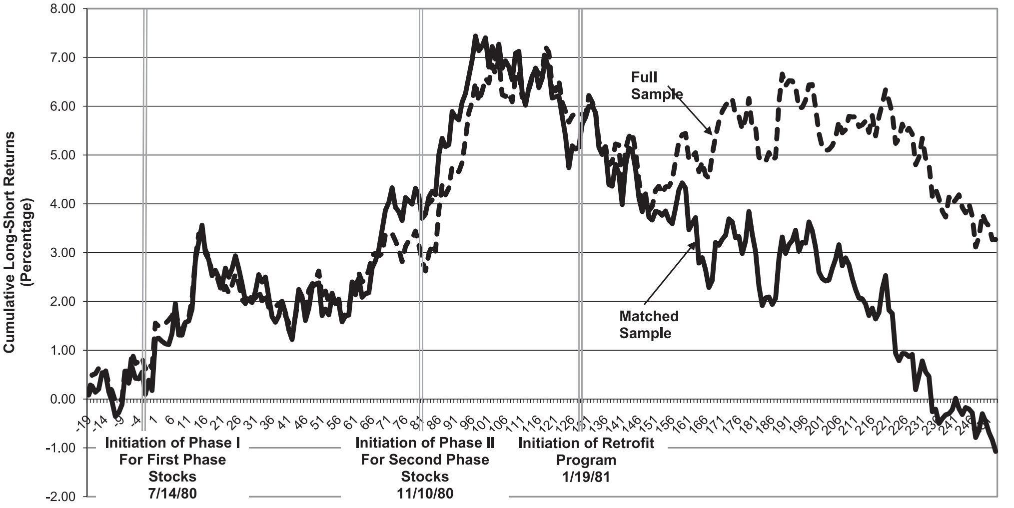

Fig. 3. Calendar time returns to long-short portfolio. This figure plots the cumulative difference in abnormal equal-weighted returns between Second Phase and First Phase stocks for the full and matched samples. The vertical bars show the implementation dates of various phases highlighted in the legend.

lends additional support to our preferred interpretation. The relatively large size of the return response could be attributed to the fact that the price, being a forward-looking variable, responds to what is likely a permanent reduction in future trading costs.

In contrast, as with the liquidity results, the return effects for First Phase stocks (not experienced by Second Phase stocks until much later) are weak. In Table 5 for the long-short portfolio, a_2 and a_3 are positive, but insignificant. In short, there is little significant evidence of abnormal returns on First Phase stocks relative to Second Phase stocks following the publication time-reducing upgrade in Phase I.

Finally, the return differential that we detect between the two groups does not persist outside of the period of the upgrades—a_{10} is statistically insignificant for all the specifications. This is reassuring, as it offers additional verification that the return differentials between groups that we are picking up are not a consequence of improper risk adjustment, or the different characteristics of stocks in the two groups.

6.3. Matching results

The final three columns of Table 5 present results from regressing the return differential between First Phase and Second Phase *matched* stocks on time period-specific intercepts. The results in the table confirm that while it is smaller in magnitude than for the full sample, a_6 is significantly positive for all stocks, with a cumulative effect of 3.04% over the 20 days following the introduction of Phase II.[21] This reduction in size could be attributed to better controlling for stock characteristics. Table 2 also reveals that the matching procedure is accomplished at a cost, as the size of the matched sample falls from 149 to 101 stocks, although this reduction in sample size does not affect our ability to reject the null of no effect. Finally, the matched sample analysis in Table 5 reveals that none of the a_2 and a_3 coefficients is statistically significant, which supports the conclusions from the full sample that the effects of Phase I are not significant.

[21] Including the four Fama-French factors does not impact the matched sample results.

Table 6
Cross-sectional evidence on returns and transactions costs.

For the matched sample stocks this table presents estimates of

$$R_{SecondPhase,i} - R_{FirstPhase,i'} = \alpha + \beta_1(\Delta c_{SecondPhase,i} - \Delta c_{FirstPhase,i'}) + Controls + u_{ii'}$$

The left-hand side variable is the difference in returns between Second Phase stock i and its matched First Phase stock i' in the period following the implementation of Phase II. The right-hand side variable is the difference between the Roll spread (measured as a change relative to the pre-upgrade period) of the Second Phase stock i and its match. The control variables are the differences for each pair of the return standard deviations, log (market capitalization), and average turnover. Newey-West t-statistics are robust to heteroscedasticity and autocorrelation, and presented in italics beneath coefficients. Coefficients significant at the 10% level are marked with asterisks.

	Matched Ret. difference Second Phase–First Phase	Matched Ret. difference Second Phase–First Phase	Matched Ret. difference Second Phase–First Phase
Intercept	0.113*	0.193*	0.132*
	1.641	*2.574*	*1.836*
Roll spread difference	−1.263*		−1.302*
	−4.202		*−3.781*
Return Std. difference		0.063	0.025
		0.589	*0.253*
Average turnover difference		0.014	0.011
		0.828	*0.656*
Log(Mcap) difference		0.501	0.732*
		1.311	*1.632*
R^2	0.151	0.029	0.161
N	101	101	101

The plot of the cumulative difference in abnormal returns between Second Phase and First Phase stocks given in Fig. 3 confirms the results in Table 5. Prior to the beginning of Phase I, there is no movement in the return differential, suggesting that market participants did not anticipate the effects of (or simply were not aware of) the Phase I upgrade. Following the initial implementation of Phase I, First Phase stocks underperform Second Phase stocks by roughly 2% over 20–60 trading days. However, this difference is not statistically significant. In the 20 days prior to the introduction of Phase II for Second Phase stocks, the return differential between Second Phase and First Phase stocks rises 1% (2%) for the entire (matched) sample. Next, after the introduction of Phase II, there is a substantial increase in the return differential of 3–4%. This replicates the result that a_6 is significant and positive in Table 5. Finally, over the period subsequent to the initiation of the Retrofit program, ending one year after the commencement of the upgrade process, the return differential for the matched sample reverts to zero. The full sample return differential, however, does not fully revert, which may be driven by the difference in characteristics between First Phase and Second Phase stocks, such as market capitalization, as is evident in Table 1.

6.4. *Cross-sectional results linking liquidity and returns*

Tables 3 and 4 show that the liquidity of Second Phase stocks improves with the upgrade and Table 5 shows that Second Phase stocks have higher returns. To further link the improvements in liquidity with higher returns, Table 6 presents results from the cross-sectional analysis of liquidity and returns in Eq. (3). The analysis reveals that a 1 bps reduction in the spread (adjusted by subtracting off the change in spread of a matched stock) is accompanied by a 1.3 bps per day (or $1.3 \times 20 = 26$ bps over the upgrade period) increase in the returns of a stock relative to its match. Thus there appears to be a substantial pricing of liquidity effect.

7. Discussion of the magnitude of liquidity and price effects

While our findings on competition's impact on liquidity and returns are consistent with the Amihud and Mendelson (1986) arguments that lower trading costs translate into higher stock prices, the magnitude of the return effects we find is large relative to the decline in the spread. We therefore investigate the reasons for these large return effects based on effects the upgrade may have had beyond reducing trading costs by narrowing the bid-ask spread.

Our first step is to undertake a simple present discounted value calculation that computes the benefits from the reduction in the Roll spread (as calculated using the Gibbs sampler) identified in the paper. Table 4 shows that the reduction in the Roll spread is 4.2 bps for the average stock. Assuming a 48% annualized rate of turnover for the average stock in our sample (from the summary statistics in our paper) and using the average risk-free rate between 1926 and 1980, which is 2.77%, the present discounted value of this reduction in trading costs (Amihud and Mendelson, 1986) is 73 bps.

The capitalized value of the spread reduction explains only about one-fifth of the relative return experienced by Second Phase stocks. The remainder of the return may be coming from a variety of factors. One possibility is that by reducing the delay in executing trades originating off the floor, Phase II reduced the discretion of specialists to execute trades at prices unfavorable to customers (Ready, 1999; and Stoll and Schenzler, 2006). Stoll and Schenzler (2006) formalize this intuition, positing that this trading option is actually a look-back option, which allows the specialist to transact at the most favorable price over the window of discretion. The post upgrades significantly reduced this trading delay and in Appendix B we provide a couple of possible quantifications of the benefits that accrue to investors as a consequence of this reduction. A second possibility is that this innovation may have increased investors trust in the fairness of the trading system (Guiso, Sapienza, and Zingales, 2008) or reduced the ambiguity that they perceive about the trading process (Easley and O'Hara, 2010). Either of these possibilities would have increased investor participation and thus generated abnormal returns around the time of the innovations. Finally, prior to this event, the NYSE had a long track record of announcing trading technology improvements that never materialized. The implementation of a large-scale improvement in trading technology such as the post upgrades is likely to have gone a long way towards convincing investors that future improvements are more likely. For example, in 1984 the NYSE upgraded the DOT system to SuperDOT, an expansion that would not have been possible without the post-upgrade. Forward-looking investors may have priced the capitalized value of future changes in liquidity into prices at the time of the post-upgrade.

8. Conclusion

Phase II of the 1980 upgrades to the NYSE's trading environment resulted in a reduction in latency for off-floor traders. This change improved off-floor traders' ability to compete with on-floor traders and generated positive returns for Second Phase stocks relative to First Phase stocks, which were reversed during the Retrofit when First Phase stocks received the Phase II upgrades.

The competition enhancing upgrades also generated relatively greater turnover and relatively lower transaction costs for Second Phase stocks. These relative effects on liquidity were also reversed during the Retrofit. It is not possible to directly infer effects on absolute returns, turnover or transaction costs from these results. But the consistency of these relative effects, the large impact on returns we found in the event study, and the large pre-to-post-upgrade increase in turnover and decline in transaction costs, strongly suggest that the upgrade process had a positive effect on liquidity and returns for the market as a whole.

Taken together, the results of this study indicate that the latency that traders experience is important for market participants and exchanges alike. The results also suggest that leveling the playing field between the public and intermediaries leads to higher liquidity and higher prices. Many markets continue to favor exchange members, for example, the New York Stock Exchange and the International Securities Exchange. Furthermore, markets' responses to new regulations such as Reg NMS in the U.S. and the European Union Markets in Financial Instruments Directive (MiFID) may lead to internalization of order flow and less competition. For example, local market participants trading on Euronext.Liffe argue that the rules should be changed to reduce competition from off-floor traders (Bradbery, 2007). After the introduction of its Hybrid electronic trading system and the consequent reduction in the role of the specialist and floor brokers, the NYSE proposed to increase the advantage of floor participants. Our results suggest that improving access for off-floor participants improves liquidity and prices, and that regulators should carefully consider the costs and benefits of markets that grant advantages to any group of market participants. The importance of leveling the playing field may be particularly acute for decentralized dealer markets, e.g., fixed income and derivatives, where research has shown that intermediaries exercise significant market power when dealing with small customers.

The significant price effects that we uncover in our study also support the arguments of Frame and White (2004) that technological innovations warrant further investigation in finance. Finally, the clean identification that arises from staggered introductions of new market systems should encourage regulators and exchanges to implement such changes as an aid to measuring their effects.

Appendix A. Details of the 1980–1981 NYSE systems and the specialist post upgrades

In the 1960s the securities industry's inability to keep up with paperwork created a crisis that caused a number of firms to go bankrupt and forced the NYSE to close on Wednesday afternoons to allow firms time to catch up. The NYSE launched an automation program that included upgrades to the clearing process, which eliminated the need for paper stock certificates, as well as improvements to trading that allowed electronic order submission [referred to as Designated Order Turnaround (DOT)]. A Market Data System (MDS) with a Common Message Switch (CMS) was also introduced to connect member firms and NYSE systems. CMS forwarded orders to DOT, which processed them (sequenced, attached addresses for routing, etc.) and passed them on to the specialist posts, and reporting of trade execution and quotes followed the same path in the opposite direction. Except for the addition of an electronic limit order book called DisplayBook in the 1980s, at the time of the post-upgrade NYSE trading generally followed the procedures described in Hasbrouck, Sofianos, and Sosebee (1993). Below we provide additional technological details on different aspects of the trading and information dissemination processes and the post upgrades.

1. *DOT*. DOT was the NYSE's first electronic order submission system. It was introduced in 1976 initially to execute 100 share market orders, and was later expanded to allow for limit

orders. Based on summary data available from the NYSE, DOT usage (up to 599 share market orders and up to 500 share limit orders at the time of the upgrade) grew rapidly in the late 1970s. The date closest to the post upgrades for which DOT information was available is the third quarter of 1978 when roughly 30,000 DOT orders were executed per day, and DOT orders participated in 40% of NYSE transactions. DOT trades were significantly smaller than floor broker trades, and DOT was subsequently expanded and renamed SuperDOT.

2. *Floor trading*. Prior to DOT, orders were typically sent to floor brokers as follows: First, a firm's trading desk would telephone orders to the firm's floor booth. Telephones on the trading floor would contact a floor broker to walk the order to the appropriate specialist post. The floor broker would either leave the order with the specialist, or remain in person to represent the order. All trades and quotes were reported by an NYSE employee called a floor reporter who stood by the specialist on the trading floor. The specialist or the floor broker representing the seller then called out the terms of the trade (ticker symbol, price, size, and the seller's badge number).
3. *Trade and quote reporting*. Prior to the post-upgrade, the floor reporter recorded the trade and quote information by filling in boxes on a card. The card was placed into the NYSE's pneumatic tube system and entered in the MDS off the floor. During busy times, queues of trades and quotes built up for reporting. Following the upgrade, the floor reporter entered the information on a "mark-sense" card, which was then fed into one of the several optical card readers at the post. The automation of the reporting system was particularly helpful during busy times.
4. *Linkages and competition among exchanges*. Legislation to promote a national market system required the NYSE to link its systems with the regional exchanges to produce the consolidated tape (CT) and consolidated quote (CQ) across all markets. In 1978, the Intermarket Trading System (ITS) began to link market centers together by enabling electronic order routing between markets. According to the ITS Plan, to avoid a "trade-through" complaint when a market (or marketable limit) order arrived in a market with an inferior quote the market maker receiving the order could do one of two things: he or she could send through ITS an irrevocable "commitment to trade" to the best-quote market to try to have the order executed there; or he or she could execute the order in the receiving market at the best ITS quote ("quote-matching"). If the order size was larger than the best-quote size, the market maker needed to satisfy the superior market before trading any of the order in his or her market at an inferior quote.

   A market maker receiving a commitment to trade executed it if the quote was still available or canceled it if the quote had changed. The commitment could also expire if not acted upon in one or in two minutes. Market makers were supposed to execute incoming commitments immediately or promptly cancel them. NYSE floor brokers could directly access ITS by filling "mark-sense cards" and feeding them into the readers by the specialist posts. According to the NYSE's "A Review of the Intermarket Trading System" in May 1981, the changes in ITS volume in 1980 "suggest that ITS grew rapidly as trading floor participants learned to integrate ITS into the usual trading process." ITS was involved in approximately 5% of NYSE transactions and a much larger share of regional exchange trading. The NYSE participated in about 90% of ITS trades. The NYSE received significantly more incoming ITS trades than ITS trades sent from the NYSE. The average response time for executed ITS commitment was about 40 seconds. According to NYSE factbooks, NYSE market share declined from just above 88% to just below 88% during the 1979–1981 period. The NYSE market share declined more quickly over the subsequent years reaching 83% in 1985.

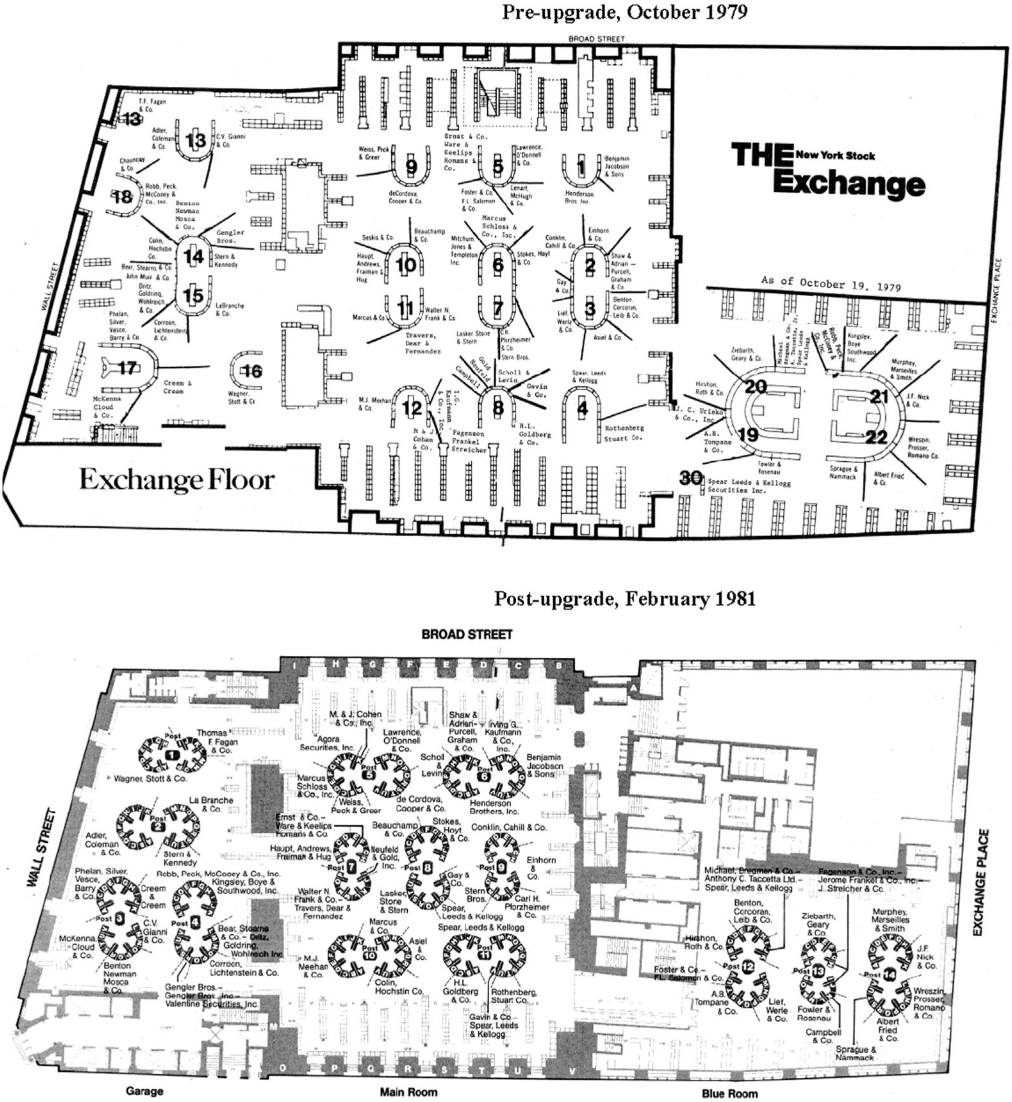

Fig. A.1. The NYSE trading floor layout, pre- and post-upgrade.

5. *Trading posts*. Prior to the 1980–1981 post-upgrade the post design from the 1920s made for pure human and paper-based trading was utilized. Over time, electronic equipment had been added, but the old posts were not designed for the new equipment, nor the work flow to utilize them optimally. For example, the ITS and DOT systems had been added separately over time. An important component of the upgrade was to link these systems together and better integrate them with other NYSE systems via the new optical card readers at the trading posts. The 1980–1981 post upgrades used cogwheel-shaped trading posts as opposed to the older U-shaped posts. Ten large posts and four smaller posts were installed, with large (small) posts holding 22 (16) specialist positions. Fig. A.1 shows the location of specialist posts on the NYSE trading floor on October 19, 1979 prior to the upgrades, and what the floor looked like

Table A.1
Post-upgrade schedule.

Remove old post	Install new post	Dates	Equipment available	Sample group
21 & 22	14	July 14, 1980	Phase I	First Phase
19 & 20	12	August 4, 1980	Phase I	First Phase
	13	August 11, 1980	Phase I	First Phase
15 & 17	3	August 25, 1980	Phase I	
16	4	September 8, 1980	Phase I	
14 & 18	2	September 15, 1980	Phase I	
13	1	October 6, 1980	Phase I	
10 & 11	7	October 13, 1980	Phase I	
2 & 3	9	October 27, 1980	Phase I	
5 & 9	5	November 10, 1980	Phase II	Second Phase
1	6	November 17, 1980	Phase II	Second Phase
8 & 12	10	December 8, 1980	Phase II	Second Phase
4	11	December 15, 1980	Phase III	
6 & 7	8	January 5, 1981	Phase IV	
		January 19, 1981	Begin Retrofit Program	
		First Quarter, 1981	Begin Turnaround Printer Pilot	

in February 1981 following the completion of the upgrades. Table A.1 and Fig. A.1 show how the upgrade began in the Blue Room before moving to the Garage and then through the Main Room. The new post design was of different size and capacity than the older posts, which led to some reorganization of the floor. The upgrades were also accompanied by changes in the location of specialist firms to consolidate their position on the floor. For example, Robb, Peck, McCooey, and Company, located on old posts 18 and 21 prior to the upgrade, moved to new post 4 following the upgrade.

The upgrades were conducted on weekends beginning with July 12–13, 1980 and continuing at one-to-three week intervals through the first weekend in 1981. The detailed schedule of the upgrades is presented in Table A.1. We used the supporting documents accompanying Fig. A.1, lists of stock locations on specific posts, to identify the dates at which specific stocks received the upgrades. For example, if a stock was located on post 16 according to the October 1979 list and on post 4 according to the February 1981 list, we classify it as having experienced Phase I of the post-upgrade on September 8, 1980. As some stocks switched locations, and given the consolidation of specialists across posts, we are not able to unambiguously identify the upgrade date associated with every stock on the exchange floor. These stocks are roughly half of the stocks receiving the Phase I and II upgrades and are included in the other category.

6. *Phases of the post-upgrade*. The upgrade was deployed in several phases. During each phase, functionality implemented in the previous phase was made operational, along with the additional functionality specific to the phase. For example, the Phase II implementation was actually the simultaneous implementation of Phases I and II. At the conclusion of the phases, a Retrofit program was rolled out and posts at which only earlier phases were implemented were brought up to date. For example, posts at which only Phase I was originally implemented received Phase II functionality during the Retrofit program.

   In Phase I of the upgrade, new card readers containing last trade and quote information were installed at each post. As described above these enabled dissemination of information more

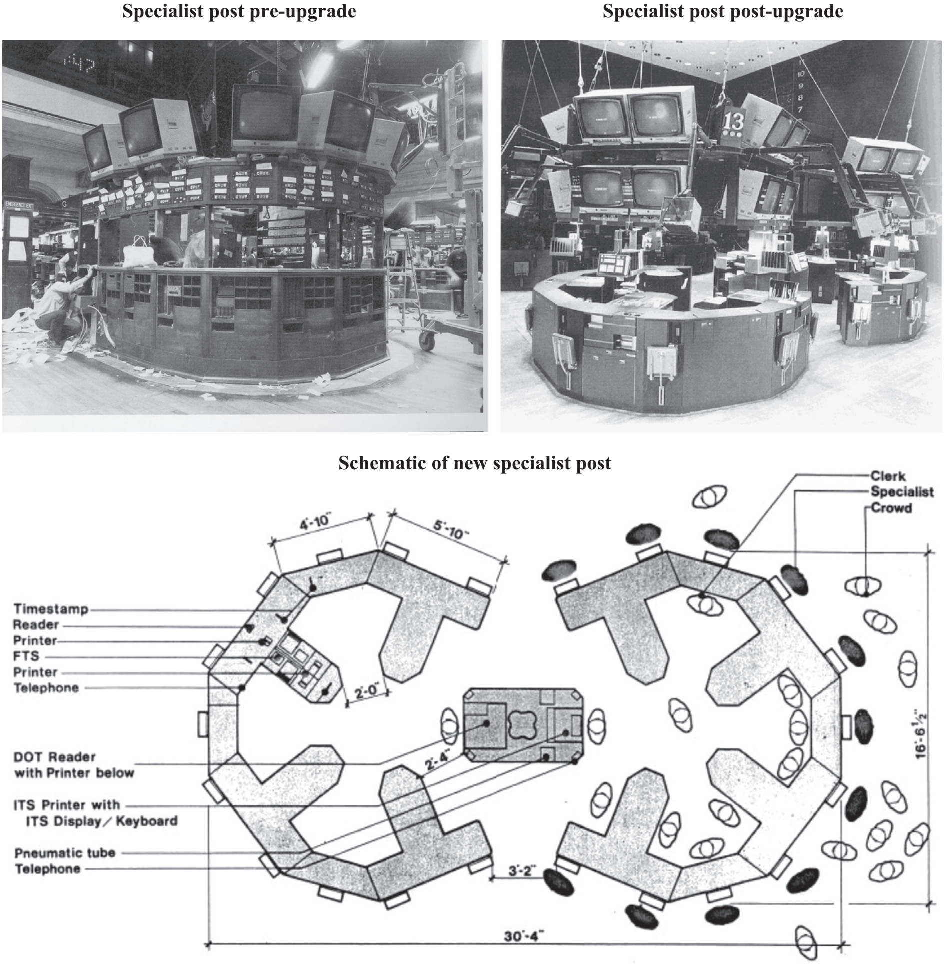

Fig. A.2. Specialist posts.

directly from the floor, but did not impact the matching of orders. Concurrently, an expanded price display unit (a CRT screen with the capacity to display up to 12 stocks) was installed in each post. In this phase, the information from floor trading was put through the new system, but DOT or ITS orders were affected. The better display technology improved the information available to floor traders. In Phase II, the card readers installed in each new post began processing DOT and ITS cards. The linking of DOT and ITS to the NYSE's other systems was expected to significantly reduce processing times for these trades. The change did not affect the information available to floor traders around the post. Phase III involved the installation of mini-printers and ITS universal mark-sense card readers, which were anticipated to cut ITS processing time for *outgoing* commitments by half, and to reduce transcription errors by card operators. Unfortunately, in the NYSE archives there is no available documentation about Phase IV, which was implemented at old posts 6 and 7 on January 5, 1981. Discussions with NYSE employees suggest that Phases III and IV did not significantly impact the trading or informational environment.

Finally, the Retrofit program brought on-line all the existing systems at all posts. Details of exact implementation dates for the Retrofit program are unavailable. Fig. A.2 contains a schematic of the new trading post, detailing the functionality installed during the various phases, and photographs of the new and old posts.

Appendix B. Look-back option calculations

To examine impacts of the post upgrades that may not show up directly in measures of the bid-ask spread, we follow the analysis of Ready (1999) and Stoll and Schenzler (2006), who show that delays in trade execution give the specialist valuable discretion to execute trades at prices that are not favorable to customers. Stoll and Schenzler (2006) formalize this intuition, positing that this trading option is actually a look-back option, which allows the specialist to transact at the most favorable price over the window of discretion. The post upgrades significantly reduced this trading delay and therefore our second step is to attempt to quantify the benefits that accrue to investors as a consequence of this reduction. Finally, we further discuss how the reduction in the specialist's trading options impacts liquidity risk (Acharya and Pedersen, 2005).

To quantify the costs to off-floor traders of delays in trade execution suggested by Ready (1999) and Stoll and Schenzler (2006), we calculate look-back call option values for a newly issued option at time $t=0$, assuming that the stock does not pay dividends using the formula of Goldman, Sosin, and Gatto (1979). Writing σ for the standard deviation of stock returns, T for the time of expiration, r for the risk-free rate of interest, S_0 for the current stock price, and $N(.)$ for the cdf of the normal distribution, the look-back call value C_0 is

$$C_0 = S_0 N(a_1) - S_0\left(\frac{\sigma^2}{2r}\right)N(-a_1) - S_0 e^{\{-rT\}[N(a_2)-(\sigma^2/2r)N(-a_3)]}. \tag{B.1}$$

where

$$a_1 = \frac{(r+(\sigma^2/2))T}{\sigma\sqrt{T}}; \quad a_2 = a_1 - \sigma\sqrt{T}; \quad a_3 = \frac{(-r+(\sigma^2/2))T}{\sigma\sqrt{T}}.$$

We price these look-back options for all the stocks in our matched sample for the period prior to the post upgrades, and the period following the post upgrades. In the first case, we price options with T set to two minutes, and in the second, we price options with T set to 20 seconds (0.33 minutes). The two minutes and 20 seconds were the NYSE's estimates of the average delays before the upgrade and what the average delay was expected to be after the upgrade (according to the NYSE's internal documents). For both calculations we use the pre-upgrade short-horizon return volatility for each stock calculated from the Fitch data used for the Roll spread estimation. Table B.1 Panel A provides the resulting look-back call option values (in basis points of the stock price).

The average difference between the two-minute look-back option value and the 20 second look-back option value is 3.75 bps of the stock price, which is statistically significant. Using the above 48% annualized rate of turnover and risk-free rate provides a capitalized value for the reduction in the look-back option of 65 bps on the price of the average stock.

Adding the reductions in look-back option values and the spreads together, we arrive at a 73 +65 = 138 bps gain for the average stock. However, this figure is likely to under-estimate the true figure for two reasons. First, trading and volatility are positively correlated [e.g., Bessembinder, Chan, and Sequin (1996) find a contemporaneous correlation between volume and volatility of about 0.40]. If trading occurs more when volatility is higher, then using the average volatility will

Table B.1
Lookback straddle option value calculations
The table presents estimates of look-back straddle option values calculated using the formula of Goldman, Sosin and Gatto (1979). The columns show the time period (in minutes) used for the option calculations, the differences between the two values, the difference multiplied by the annualized rate of turnover of the stocks in our sample (48%, taken from Table 1), and the capitalized present value of the difference, computed using the average risk-free rate between 1926 and 1980, of 0.0277. Panel A shows this difference between 2 minutes and 20 second options, and Panel B between 5 minutes and 20 second options.

Look-back option value	$T=2$	$T=0.33$	Difference	Turnover adjusted	Capitalized
Panel A					
Mean	5.599	1.850	3.749	1.800	64.964
Median	5.283	1.778	3.424	1.643	59.329
Std. Dev.	2.184	0.708	1.681	0.807	29.138
Panel B					
Look-back option value	$T=5$	$T=0.33$	Difference	Turnover adjusted	Capitalized
Mean	9.905	1.850	8.055	3.866	139.581
Median	9.377	1.778	7.512	3.606	130.178
Std. Dev.	3.485	0.708	3.028	1.453	52.464

underestimate the look-back option costs investors will face. Second, during fast markets, the pre-upgrade delays were significantly longer than two minutes. This is basically due to queuing effects on the floor (in queuing models, holding the processing time of a trade constant, an increase in the arrival rate of trades causes the waiting time to increase). These queuing delays in fast/volatile markets could cause the average order processing time to be significantly higher than two-minute time (i.e., 5 minutes or even higher). In such situations, the present discounted value of the look-back option value reduction is 1.4% for the average stock (see Table B.1 Panel B).

References

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

Amihud, Y., Mendelson, H., 1986. Asset pricing and the bid-ask spread. Journal of Financial Economics 17, 223–249.

Amihud, Y., Mendelson, H., Lauterbach, B., 1997. Market microstructure and securities values: evidence from the Tel Aviv exchange. Journal of Financial Economics 45, 365–390.

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

Bessembinder, H., Chan, K., Sequin, P., 1996. An empirical examination of information, difference of opinion, and trading activity. Journal of Financial Economics 40, 105–134.

Bessembinder, H., Maxwell, W., Venkataraman, K., 2006. Market transparency, liquidity externalities, and institutional trading costs in corporate bonds. Journal of Financial Economics 82, 251–288.

Bloomfield, R., O'Hara, M., 1999. Market transparency: who wins and who loses? Review of Financial Studies 12, 5–35.

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.

Bradbery, A., 2007. Traders fret algorithmic systems. Wall Street Journal , February 21.

Brav, A., Gompers, P., 1997. Myth or reality? The long-run underperformance of initial public offerings: evidence from venture and nonventure capital-backed companies. Journal of Finance 52, 1791–1821.

Cespa, G., Foucault, T., 2013. Sale of price information by exchanges: does it promote price discovery? Management Science (forthcoming).

Daniel, K., Titman, S., 1997. Evidence on the characteristics of cross-sectional variation in stock returns. Journal of Finance 52, 1–33.
Easley, D., O'Hara, M., 2010. Microstructure and ambiguity. Journal of Finance 65, 1817–1846.
Easley, D., O'Hara, M., Yang, L., 2012. Differential access to price information in financial markets. Working paper, Cornell University.
Edwards, A., Harris, L., Piwowar, M., 2007. Corporate bond market transaction costs and transparency. Journal of Finance 62, 1421–1451.
Fama, E., 1998. Market efficiency, long-term returns, and behavioral finance. Journal of Financial Economics 49, 283–306.
Flood, M.D., Huisman, R., Koedijk, K.G., Mahieu, R.J., 1999. Quote disclosure and price discovery in multiple-dealer financial markets. Review of Financial Studies 12, 37–59.
Frame, W.S., White, L.J., 2004. Empirical studies of financial innovation: mostly talk and not much action? Journal of Economic Literature 42, 116–144.
Garbade, K.D., Silber, W.L., 1978. Technology, communication and the performance of financial markets: 1840–1975. Journal of Finance 33, 819–832.
Glosten, L., Harris, L., 1988. Estimating the components of the bid-ask spread. Journal of Financial Economics 21, 123–142.
Goldman, M., Sosin, H., Gatto, M., 1979. Path dependent options: buy at the low, sell at the high. Journal of Finance 34, 1111–1127.
Green, R., Hollifield, B., Schuroff, N., 2006. Financial intermediation and the costs of trading in an opaque market. Review of Financial Studies 20, 275–314.
Guiso, L., Sapienza, P., Zingales, L., 2008. Trusting the stock market. Journal of Finance 63, 2557–2600.
Harris, L., Piwowar, M., 2006. Secondary trading costs in the municipal bond market. Journal of Finance 61, 1361–1397.
Hasbrouck, J., 2009. Trading costs and returns on U.S. equities: estimating effective costs from daily data. Journal of Finance 64, 1445–1477.
Hasbrouck, J., Sofianos, G., Sosebee, D., 1993. New York Stock Exchange systems and trading procedures. Working paper.
Jain, P.K., 2005. Financial market design and the equity premium: electronic versus floor trading. Journal of Finance 60, 2955–2985.
Kothari, S.P., Warner, J.B., 2005. The econometrics of event studies. In: Eckbo, B.Espen (Ed.), Handbook of Corporate Finance: Empirical Corporate Finance. Handbooks in Finance Series. Elsevier/North-Holland (Chapter 1).
Lyon, J., Barber, B.M., Tsai, C., 1999. Improved methods of tests of long-horizon abnormal stock returns. Journal of Finance 54, 165–201.
Macey, J., O'Hara, M., 1999. Regulating exchanges and alternative trading systems: a law and economics perspective. Journal of Legal Studies 28, 17–54.
Madhavan, A., Porter, D., Weaver, D., 2005. Should securities markets be transparent? Journal of Financial Markets 8, 265–287.
Mitchell, M., Stafford, E., 2000. Managerial decisions and long-term stock price performance. Journal of Business 73, 287–329.
Muscarella, C.J., Piwowar, M., 2001. Market microstructure and securities values: evidence from the Paris Bourse. Journal of Financial Markets 4, 209–230.
Newey, W.K., West, K.D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703–708.
Ready, M., 1999. The specialist's discretion: stopped orders and price improvement. Review of Financial Studies 12, 1075–1112.
Roll, R., 1984. A simple explicit measure of the effective bid-ask spread in an efficient market. Journal of Finance 39, 1127–1139.
Stoll, H.R., Schenzler, C., 2006. Trades outside the quotes: reporting delay, trading option or trade size? Journal of Financial Economics 79, 615–653.