Levelling the Trading Field *

David Easley
Cornell University

Terrence Hendershott
University of California, Berkeley

Tarun Ramadorai
University of Oxford and CEPR

This version: November 19, 2009

Abstract

We examine the impact on stock prices of a major upgrade to the New York Stock Exchange's trading environment. The upgrade increased competition among traders by reducing the execution latency experienced by traders off the floor of the exchange, relative to traders on the floor of the exchange. The upgrade was sequentially implemented across groups of stocks, and a portfolio that is long stocks undergoing the upgrade in the first 20 days of the upgrade process and short stocks receiving the upgrade later has a return of between three and four percent over the period, depending on the method of risk adjustment. This return differential is eliminated when all stocks are upgraded. The upgrade also had significant impacts on liquidity and turnover.

JEL Classifications: G12, G14
Keywords: market structure, stock exchange, trading volume, transparency, information, latency, technology upgrade

---

* We thank Thierry Foucault, Larry Glosten, Joel Hasbrouck, Charles Jones, Albert Menkveld, Maureen O'Hara, Vish Viswanathan, Bruce Weber, and seminar participants at the 3rd Annual Central Bank Workshop on the Microstructure of Financial Markets, the 2007 AMF-SEC Conference on the Structure and Regulation of Financial Markets, the 2006 ISB CAF Winter Finance Conference on Microstructure of International Financial Markets, the 2006 Workshop on Information Systems and Economics, the Copenhagen School of Business, the Georgia Institute of Technology, HEC-Paris, the NASDAQ Economic Advisory board, the Stockholm School of Economics and the University of California at Davis for comments. We are indebted to Archivist Janet Lindeh and the NYSE for providing us with the information used in this study. We thank Abhishek Saxena and especially Tian Liang for excellent research assistance. Hendershott gratefully acknowledges support from the National Science Foundation. Part of this research was conducted while Hendershott was the visiting economist at the New York Stock Exchange. Previous versions of this paper were titled "The Impact of Trading Technology: Evidence from the 1980 NYSE Post Upgrades" and "The Price of Latency."

David Easley: dae3@cornell.edu; Terrence Hendershott: hender@haas.berkeley.edu; Tarun Ramadorai: tarun.ramadorai@sbs.ox.ac.uk

1. Introduction

Competition in financial markets is central to efficiency and fairness (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 in financial markets, often through encouraging technological innovation. However, the benefits of such competition-promoting innovation on financial markets are difficult to measure. In this paper, we test the hypothesis that increasing competition among traders is accompanied by price appreciations and increases in measures of liquidity. Using a major change to trading technology by the NYSE in 1980 we find significant effects on both asset prices and liquidity 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 New York Stock Exchange (NYSE) significantly upgraded its trading platform – specifically, the technology used at the posts at which specialists processed transactions. 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 (ITS) systems) 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 could 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 abnormal returns cannot be done in the usual manner, i.e. 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 stock returns. Second, the change in latency and consequent change in transparency for off-floor traders was expected to be large: from two minutes pre-upgrade to less than 20 seconds post-upgrade. Third, the staggered implementation allows us

[1] While equity markets now foster such competition, most decentralized dealer and over the counter markets continue to lack transparency and equal access. For example, see Bessembinder, Maxwell, Venkataraman (2006) and Edwards, Harris, and Piwowar (2007) for corporate bonds and Green, Hollifield, and Schuroff (2006) and Harris and Piwowar (2006) for municipal bonds.

[2] During this upgrade, NYSE replaced the specialists' trading posts that had been in place since the 1920's 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.

to view the effects of two distinct types of changes to the trading environment. The first change that was implemented was what NYSE internal documents referred to as "Phase I" of the upgrades. This phase reduced the dispersion of information among on-floor participants by improving the dissemination of quotes and the reporting of floor transactions. As a consequence, this change increased the informational advantage that on-floor traders had over off-floor traders, but also potentially increased competition among traders on the floor. The second change that was implemented, "Phase II," was a superset of Phase I, with the addition of technology that significantly reduced transaction processing times 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.[4] Our most significant finding is that a portfolio which is long Second Phase stocks and short First Phase stocks exhibits a substantial positive return in the twenty 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 basis points a day for a total return of four percent. 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 find that Phase I (which reduced the dispersion of information among on-floor participants by improving the dissemination of quotes and the reporting of floor transactions) generated a reduction in returns of two percent for First Phase stocks relative to Second Phase stocks. However, this effect is not precisely estimated – at best, it is significant at the ten percent level. 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.

Providing better information and reducing latency for off-floor traders decreases their disadvantage vis-à-vis on-floor traders. Therefore, we would expect off-floor traders to trade more following such improvements. While we cannot directly measure the trading behavior of this group, we

[4] We discuss the details of sample construction in the Data section and the Appendix.

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. As with returns, there is no longer any difference in turnover between the two groups of stocks 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 transactions data and the Gibbs sampler methodology of Hasbrouck [2006]. We find that transactions costs decline by approximately 13 basis points 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 four basis points in the differential between the average spread of matched Second Phase and First Phase stocks in the period around the implementation of Phase II. Finally, we show that the increases in returns for Second Phase stocks line up cross-sectionally with the transactions costs reductions experienced by these stocks.

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 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. Finally, a note on the magnitude of the effects we discover: Ready (1999) and Stoll and Schenzler (2006) 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 we attempt to quantify the benefits that accrue to investors as a consequence of this reduction. Using 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 (1999)), as well as the reduction in the value of the look-back option, we are able to account for over half of the approximately four percent magnitude of the detected effect. We attribute the remainder to the decrease in investors' liquidity risk induced by the look-back option, and increased investors expectations that the NYSE would continue to make changes improving liquidity and leveling the playing field between different types of traders.

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. Section 5 describes the data employed. Section 6 presents the results of our estimation. Section 7 provides discussion. Section 8 concludes.

2. The Upgrade Process

To understand how these technological changes reduced latency and improved transparency, it is helpful to have a view of how the market operated prior to the upgrades. Figure A.2, Specialist Post Pre-Upgrade, and Figure A.3, The NYSE Trading Floor Pre-Upgrade, shows a typical specialist's post and the trading floor prior to the upgrade. Specialists had limit order books,[5] 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 and 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 long 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.

Phase I of the upgrades reduced the dispersion of information among on-floor participants by improving the dissemination of quotes and the reporting of floor transactions. The specific technological changes implemented during this Phase were as follows: First, new Chatsworth card readers, which were used to read punch cards containing last trade and quote information, were installed in each post. Concurrently, an expanded price display unit was installed at each post to improve the volume and clarity of quote and trade information available to floor members. In this phase, the information that was put through the new system did not include designated order turnaround system (DOT) or intermarket trading system (ITS) orders. The new equipment was, according to the exchange's estimates, meant to reduce the publication time of floor (non-DOT and non-ITS) transactions and quote information to floor members from two minutes to 20 seconds. This change has ambiguous affects on trader competition because while it could intensify competition among on-floor traders, it increased the advantages of on-floor traders over off-floor traders.

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. NYSE internal documents state that the exchange anticipated the implementation of Phase II to reduce the average reporting delays of ITS and DOT trades from 2 minutes to less than 20 seconds, reducing turnaround times and latency for these orders, which in turn

[5] 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.

would improve transparency for off-floor traders.[6] This phase of the upgrade process reduced the latency experienced by off-floor traders, and thus reduced the informational advantage that on-floor had over off-floor traders. Finally, posts that were upgraded during Phase I received the Phase II upgrades during a Retrofit program, which commenced after all Phase II upgrades were completed.

3. Related Literature

Our study is related to research on market structure changes by Garbade and Silber [1983], Amihud, Mendelson and Lauterbach [1997], Muscarella and Piwowar [2001], Jain [2005], Battalio, Ellul and Jennings [2007], Ready [1999], Stoll and Schenzler [2006], and Cespa and Foucault [2007].

Garbade and Silber [1983] 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 percent 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 five percent over 40 days, and exhibited significant increases in volume and liquidity. The change from call to 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

[6] In interviews, market participants and NYSE employees that were present at the time of the upgrades confirm that these stated time reductions were in fact realized.

structure to reduce the advantage that on-floor traders have over off-floor traders improves liquidity and increase 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 change latency. This narrower focus allows us to draw sharper conclusions about how changing competition affects liquidity and returns.

Our interpretation of the impact of Phase II draws upon Ready [1999], Stoll and Schenzler [2006], and Cespa and Foucault [2007]. 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.[7] Cespa and Foucault [2007] analyze a different aspect of latency, the fact that insiders (for us, those on the floor of the exchange) 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, and thus reducing the risk that outsiders face in their trading with insiders. Cespa and Foucault's [2007] theoretical analysis thus provides additional support for our empirical analysis, which shows that enhancing competition between off-floor and on-floor traders are accompanied by price improvements.

The relationship of Phase I to the literature is more complex. This phase of the upgrade resulted in faster publication of trade and quote information on the floor of the exchange. This increased pre- and post-trade transparency for on-floor traders, improving the flow of information on the floor, and, consequently, reducing the dispersion of information among on-floor participants. Transparency is a hotly debated topic in the literature.[8] Because the informational dispersion and transparency changes are

---

[7] Off-floor traders could have hired floor brokers to attempt to obtain some of 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.

[8] 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 [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 et. al. [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 [2006], who find that the introduction of the NYSE Open Book generates increases in stock market liquidity.

simultaneous during Phase I, the competitive effects of Phase I cannot be interpreted as unambiguously as the competition results for Phase II.

4. Methodology

4.1. Event Study

We begin with an event-study (Fama, Fisher, Jensen and Roll [1969]), the standard method to analyze the price impacts of events occurring over multiple dates, and involving multiple stocks. Generally, researchers identify two dates for each event: the announcement date, and the implementation date. However there are several issues with identifying the announcement date of the post upgrades. First, the NYSE's archives contain few references to official announcements, and when they exist, they are often wrong. Indeed, the NYSE announced the implementation date of the post upgrade incorrectly at least once: the first announcement that we could identify occurs on January 29, 1979, when the exchange announced the beginning of a pilot for the design of new specialist posts. On July 2, 1979 the results of the pilot (on Post 9) were announced. The July 2 announcement also contained an implementation schedule which ran from April-May 1980 to July-Oct, 1980. This differs by several months from the July 11, 1980 memo which gives the final installation schedule as beginning on the following Monday, July 14, 1980.

Second, the NYSE had a track record of announcing changes that were never implemented, or when implemented, failed to deliver the benefits that were advertised (see Keith and Grody [1988]). Examples include: the Centralized Clearing Systems (CCS), the Alternative Trading System (ATS), the Centralized Exchange Network Trading and Unified Reporting (CENTAUR), Locked-In Trade (LIT), and Floor Derived Clearance (FDC). These repeated instances of failed announcements of technology improvements by the NYSE made the credibility of their pronouncements questionable.

For these reasons, we line stocks up according to the day on which they resume trading after upgrade implementation. Note that our use of implementation rather than announcement dates would bias measured abnormal returns towards zero if price effects occurred at the announcement rather than implementation dates. Since the post upgrades were staggered in time, the implementation sequence indicated in the upgrade schedule (see Table A.1. in the Appendix) can be used to study the impacts of Phase I of the upgrades, which commences on 7/14/1980, separately from the off-floor competition enhancing Phase II of the upgrade, which commences on 11/10/1980. We are able to reliably identify the date on which Phase I was implemented for 327 stocks, and for 286 stocks for which Phases II, III and IV were implemented (Phases III and IV include Phases I, II, and minor additional improvements such as the installation of mini-printers for outgoing ITS commitments).

We measure abnormal returns using a market model, and sum these abnormal returns over windows of time surrounding the events, to form cumulative abnormal returns (CARs). These CARs are a measure of abnormal price increases. Tests of whether Phases I and II impact stock prices are, in this framework, tests of whether the CARs are statistically different from zero prior to and following the implementation of each phase.

4.2. Long-Short Portfolios in Calendar Time

There are at least two problems with using event-study methodology to test our hypotheses. 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], Barber, Lyon 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 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 that experience the same change during this 60 day window, 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, and we provide evidence on liquidity that helps 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 (twenty day windows) before, during, and after the upgrade process.[9]

We begin by assessing the impacts of the upgrades on returns, estimating the following equations:

$$(R_{SecondPhase,t} - R_{FirstPhase,t}) = \sum_{i=1}^{10} a_i D_i + u_t \ , \tag{1}$$

and

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

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 ($R_{SecondPhase,t}$) each day. The ten dummy variables $D_1,...,D_{10}$ pick out the pre-and post-periods surrounding the rollouts of Phase I, Phase II, and the Retrofit program. Equation (2) risk-adjusts the return on this long-short portfolio by regressing it on the Fama-French-Carhart four-factor model.[10] The coefficients $a_1,...,a_{10}$ from equation (2) 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, $D_1,...,D_{10}$.

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 greater than (less than) zero correspond to a negative (positive) return impact of reducing the dispersion of information among on-floor participants by 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 Second

---

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

[10] 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.

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 off-floor 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 also separately estimate the results for terciles of stocks sorted by pre-event-period market capitalization. Finally, we employ Newey and West (1987) standard errors throughout our regressions. These are robust to heteroskedasticity and autocorrelation.

4.3. Matched Portfolios in Calendar Time

In order to obtain clean estimates of the impact of the post upgrades, we attempt to control for the factors that drive returns on the right-hand side of equation (2). However, it is possible that our results are driven by differences in the characteristics of the stocks in the two groups that we do not adequately control for with our factor proxies. To insure ourselves against this possibility, we match First Phase stocks with Second Phase stocks, based on the market capitalization and turnover of these stocks. 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 percent. 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.

We then conduct the same analysis as before, that is, we re-estimate equation (1), substituting the difference between the equal-weighted returns of matched Second Phase-First Phase pairs as the left hand-side variables:

$$(R_{MatchSecondPhase,t} - R_{MatchFirstPhase,t}) = \sum_{i=1}^{10} \mathrm{a}_i D_i + u_t \qquad (3)$$

4.4. Evidence on Liquidity

Our hypotheses about the price effects of increase competition are predicated on the assumption that latency reductions will improve the ability of off-floor market participants to transact with one another. We therefore check whether the technological upgrades are accompanied by increases in turnover and reductions in transactions costs. We define the time periods around the upgrades differently for the liquidity analysis than for the return analysis, compressing the ten periods into six periods during which the trading environment for First Phase and Second Phase stocks differs (these are: pre-upgrade, pre-Phase I, Phase I, Phase II, Retrofit and post-upgrade). Prices are forward looking so the upgrades' impact may be concentrated around the upgrade events, making narrower time windows appropriate for returns. In contrast, 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.

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 transactions data (described in the Data section). Our measure of transactions costs is the spread from the Roll [1984] model, estimated using the Gibbs sampling methodology introduced by Hasbrouck [2006]. The model we estimate is:

$$\Delta p_\tau = c\Delta q_\tau + u_\tau , \qquad (4)$$

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 transactions cost (spread). We assume that $u_\tau \sim NIID(0, \sigma_u^2)$ and compute Δp_τ from transactions 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 [2006]. 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 transactions costs and turnover, our approach is to compute differences-in-differences of turnover and estimated transactions 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 heteroskedasticity and autocorrelation robust estimator.

We also conduct a cross-sectional test, in which we line up the price effects of the off-floor competition improvement with the transactions cost changes. We estimate:

$$R_{SecondPhase,i} - R_{FirstPhase,i'} = \alpha + \beta_1(\Delta c_{SecondPhase,i} - \Delta c_{FirstPhase,i'}) + Controls + u_{ii'} \qquad (5)$$

The left-hand side variable in equation (5) 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 transactions costs experienced by each stock undergoing the transparency/latency improving upgrade. We therefore expect a negative coefficient β_1. We now turn to describing the data we employ in our analysis.

[11] 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.

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 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.[12]

Using data from CRSP, Table I presents descriptive statistics for our stocks. The table shows that there are 159 stocks in our First Phase group, and 149 in our Second Phase group, leaving 1,118 'Other' identified stocks not in either of the two groups.[13] 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.[14]

[TABLE I AROUND HERE]

Table II 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.

[TABLE II AROUND HERE]

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) data base and consists of a

[12] 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.

[13] 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). See Battalio, Ellul and Jennings [2007] for an analysis of how relocations impact "reputational" trading on the NYSE.

[14] 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.

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 used in Glosten and Harris (1988).

The next section presents and discusses the results from estimating the specifications outlined in the previous section.

6. Results

6.1. Event Study Results

Figure 1 presents results from a market-model event study for Phase I and Phase II. Date zero in the figure represents the implementation date of each upgrade. Panel A (B) of the figure plots the cumulative abnormal returns for 30 days before and 60 days after the implementation of Phase I (Phase II). The graph suggests that there is positive anticipation of Phase I, with the CAR rising to 1.17 percent over the 20 day window prior to its implementation. However, in the 60 days following Phase I, there is no real pattern to the residuals, which fluctuate around zero.

[FIGURE 1 AROUND HERE]

Panel B offers an interesting contrast, with little evidence of systematic return movements prior to the implementation of the improved transparency/latency phase, but a very strong response in the 60 day window after the implementation – the CAR stands at 3.58 percent on the 60^{th} day post-implementation. These CARs are significant at the one percent level, using the jackknife estimator (Giacotto and Sfiridis [1996]) to compute standard errors.

While these results suggest that there may be positive impacts around both phases, we must address the issues of contemporaneous correlation and the shifting benchmark to draw clear inferences. We now turn to the results from the calendar time portfolio analysis.

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. in the Appendix reveals that the first three introductions of Phase II occur within 20 trading day periods of one another. Figure 2 plots the event time abnormal returns to portfolios of matched stocks (see Table II) receiving the Phase II upgrade on 11/10/1980, 11/17/1980 and 12/08/1980. Abnormal returns are created by subtracting off the returns of matched stocks identified using the procedure described in the Data section.

[FIGURE 2 AROUND HERE]

The figure 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 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 III presents estimates of equation (1). In this table, a_6 (the coefficient on the dummy variable for the introduction of Phase II) is statistically significant, for all stocks, as well as for all three terciles of stocks sorted by market capitalization. These 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 20.4 basis points per day. This translates into a cumulative abnormal return of 4.08 percent over the 20 day period. 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. These results provide additional support for the prediction suggested by our event study, that improving off-floor traders' ability to compete increases stock prices.

[TABLE III AROUND HERE]

Table IV estimates equation (2), which modifies equation (1) by incorporating the four Fama-French factors. The Table confirms the results in Table III – the coefficient a_6 remains statistically significant for all stocks and all three size terciles of stocks.

[TABLE IV AROUND HERE]

In contrast, the effect of the reductions in reporting time for trade and quote information on the floor, experienced by First Phase stocks (not experienced by Second Phase stocks until much later) are weak. In Tables III and IV, a_2 and a_3 are positive, though only significant for the medium size tercile of stocks in Table III. In short, although we are unable to find a great deal of evidence, the decrease in abnormal returns on First Phase stocks relative to Second Phase stocks following the publication-time-reducing upgrade provides some support for the hypothesis that Phase I adversely affected First Phase stocks.

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 rooms that we are picking up are not a consequence of improper risk-adjustment, or the different characteristics of stocks in the two rooms.

6.3. Matching Results

Table V presents results from equation (3) – the regression of the return differential between First Phase and Second Phase matched stocks on time period-specific intercepts. The results in the table confirm that a_6 is significantly positive for all stocks, with a magnitude close to those reported in Tables III and IV. However, there is some attenuation of statistical significance for the result for small and medium size stocks. These results could be attributed to insufficient power to detect a difference from zero in the return differential. Table II reveals that the matching procedure is accomplished at a cost, as the size of the matched sample in the small (medium) tercile falls from 94 to 66 (100 to 66) stocks. Finally, the matched sample analysis in Table V reveals that only one of the eight a_2 and a_3 coefficients is statistically significant, which supports the conclusions in Tables III and IV that the effects of Phase I are weak at best.

[TABLE V AROUND HERE]

The plot of the cumulative difference in abnormal returns between Second Phase and First Phase stocks given in Figure 3 confirms the results in Tables III and V. 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 two percent over 20 to 60 trading days. While this increase is not statistically significant, the impacts on returns are large in magnitude. 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 one (two) percent for the entire (matched) sample. Next, after the introduction of Phase II, there is a substantial increase in the return differential of three to four percent. This replicates the result that a_6 is significant and positive in Tables III and V. 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 – this may be driven by the difference in characteristics between First Phase and Second Phase stocks, such as market capitalization, which is evident in Table I.

[FIGURE 3 AROUND HERE]

6.4. Evidence on Liquidity

Table VI 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-differences estimates of increases in turnover for the matched sample 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. The transparency/latency improving upgrade appears to be accompanied by movements in trading volume of the predicted sign.

[TABLE VI AROUND HERE]

The plot of the difference in turnover between Second Phase and First Phase stocks given in Figure 4 is analogous to Figure 3, and shows that the pattern in returns is mimicked by the patterns in turnover. 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.

[FIGURE 4 AROUND HERE]

Table VII shows the results for the estimates of transactions 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 [2006]. Similar to the turnover and trade frequency estimates, we see evidence of a large decline in transactions 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 24%. Furthermore, the difference-in-differences 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 Second Phase stocks' transactions costs lower by 4.2 basis points.

[TABLE VII AROUND HERE]

Some caution is in order in interpreting our measures of changes in turnover and transactions 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 VI and VII 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. Interpretation of our differences-in-differences measures is more straightforward. We find significant relative results during the implementation of Phase II. These changes are immune to the secular changes affecting the market as a whole over this period. However, there still remains the question of which group of stocks was responsible for the results. We believe that Second Phase stocks have higher relative returns, higher relative turnover increases and lower relative transaction cost changes because of the beneficial effects of Phase II on competition for these stocks. The cross-sectional analysis that we now turn to provides additional support for this interpretation.

Table VIII presents results from our cross-sectional analysis. The analysis reveals that a one basis point reduction in the spread (adjusted by subtracting off the change in spread of a matched stock) is accompanied by a 1.3 basis point per day (or 1.3*20 = 26 basis points over the upgrade period) increase in the returns of a stock relative to its match.

[TABLE VIII AROUND HERE]

7. Discussion of the Magnitude of Liquidity and Price Effects

While our findings on competition's impact on return and liquidity 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 using a series of computations based on effects the upgrade 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. Next, 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 (as in Acharya and Pedersen (2005)).

Table VII shows that the reduction in the Roll spread is 4.2 basis points 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 basis points.

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:

$$C_0 = S_0 N(a_1) - S_0 (\frac{\sigma^2}{2r}) N(-a_1) - S_0 e^{\{-rT\}[N(a_2) - (\frac{\sigma^2}{2r}) N(-a_3)]} \qquad (6)$$

Where

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

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 IX Panel A shows that the resulting look-back call option values (in basis points of the stock price) have the following characteristics:

[TABLE IX AROUND HERE]

The average difference between the two-minute lookback option value and the twenty second look-back option value is 3.75 basis points 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 basis points 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 basis point 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 (for example, 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 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 would 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 IX Panel B.

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 were 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.

To summarize the impacts of the changes due to the event, first, it decreased the Roll spread by 4.2 basis points which translates into a 0.73% increase in price. Second, it decreased on-floor traders' look-back option by 3.75 or 8.06 basis points, depending on whether we account for an increase in expected delays during fast markets. This translates into a capitalized value of 0.65% to 1.4%. Accounting for the positive correlation between volume and volatility could push this number higher. Finally, it increased investors' expectations that the NYSE would continue to make changes improving liquidity and leveling the playing field between off-floor investors and on-floor traders. Taking all these factors together, it seems plausible that the post upgrade could result in a 4% increase in the price of the average stock.

8. Conclusion

Phase II of the 1980 upgrades to the NYSE's trading environment resulted in a reduction in latency and improvements in transparency for off-floor traders. These changes 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 transactions 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 and transparency 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 has 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 important 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
Details of the 1980-1981 NYSE Specialist Post Upgrades from NYSE Archive Documents

In the 1980-81 post upgrades, the exchange replaced trading posts that had been in place since the 1920's with a cogwheel shaped trading post that utilized new technology. Ten large posts and four smaller posts were installed, with large (small) posts holding 22 (16) specialist positions.

Figure 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 in February 1981 following the completion of the upgrades. The upgrades were accompanied by some changes in the location of specialist firms which were made in order 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 post 4 following the upgrade.

The upgrades were conducted on weekends beginning with 12-13 July 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 Figure 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.

The upgrade was deployed in several phases. At each phase, functionality implemented in the previous phase was made operational, along with additional functionality specific to the phase. For example, the Phase II implementation was actually an implementation of Phases I and II simultaneously. 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 Chatsworth card readers, which were used to read punch cards containing last trade and quote information, were installed in each post. Concurrently, an expanded price display unit (a CRT screen with the capacity to display six ITS stocks, or 12 non-ITS stocks, or combinations of the two) was installed in each post. In this phase, the information that was put through the new system did not include designated order turnaround system (DOT) or intermarket trading system (ITS) orders. Transactions involving these off-floor orders constituted a significant fraction of total daily volume.[15] In Phase II, the eight card readers installed in each new post began processing DOT and ITS cards. 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 environment.

Finally, the Retrofit program was rolled out to bring on-line all the existing systems installed, but not active at each new post during the previous installation phases. Details of exact implementation dates for the Retrofit program are unavailable. Figure 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. Figure A.3 provide photographs of the floor pre-upgrade (1955 and 1978) and post-upgrade (1981).

---

[15] Detailed records of DOT and ITS use during the sample period are not available. DOT 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 (599 share market and 500 share limit orders at the time of the upgrade) was growing rapidly in the late 1970s. In the third quarter of 1978 there were roughly 30,000 DOT orders per day and DOT orders participated in 40 percent of NYSE transactions. Extrapolating a linear trend from 1977-1978 would project the number of DOT orders to be about 50,000/day at the time of the upgrade. In November 1980, ITS transactions constituted around five percent of trading volume.

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	

Figure A.1.
The NYSE Trading Floor Layout, Pre- and Post-Upgrade

Pre-Upgrade, October 1979

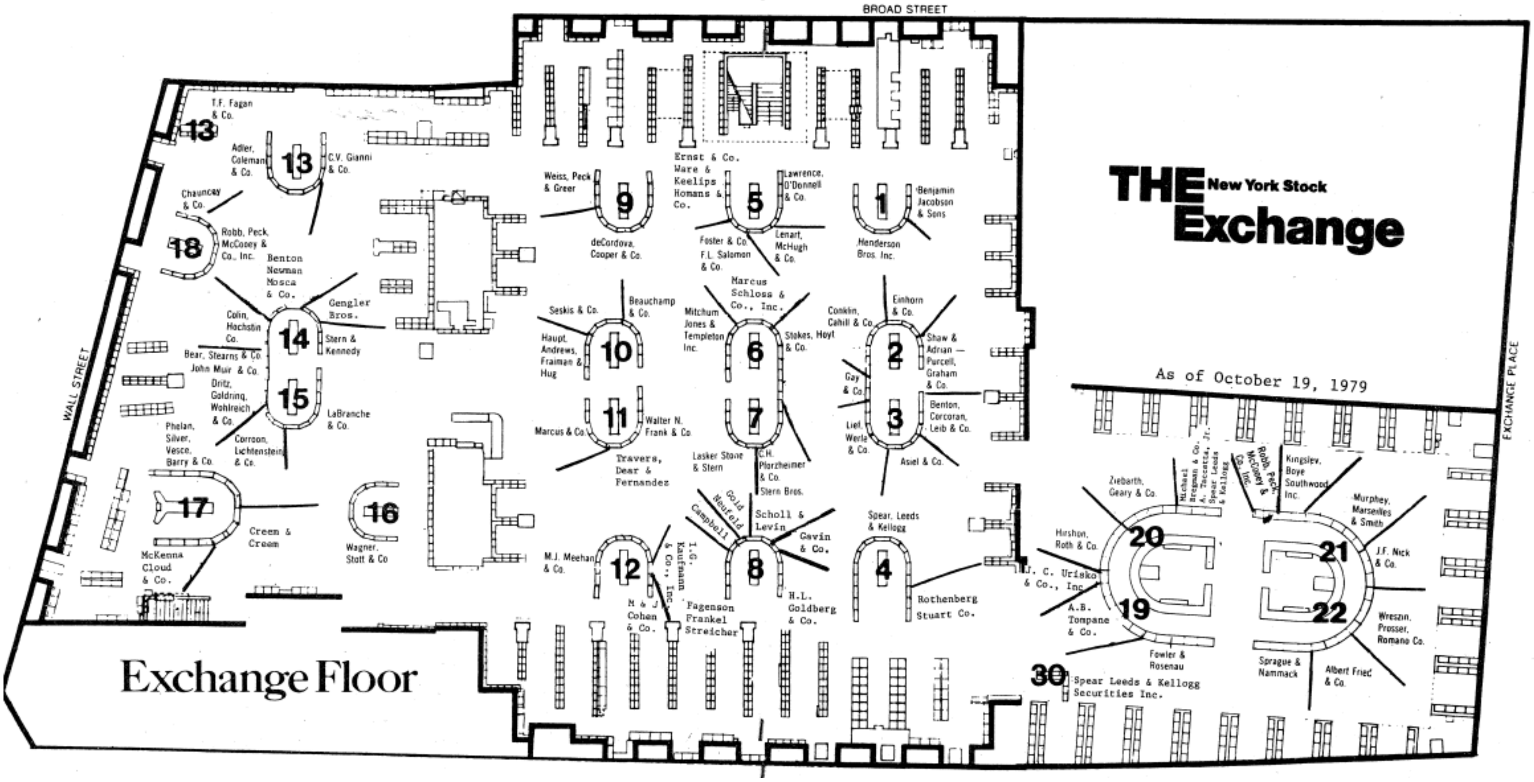

Post-Upgrade, February 1981

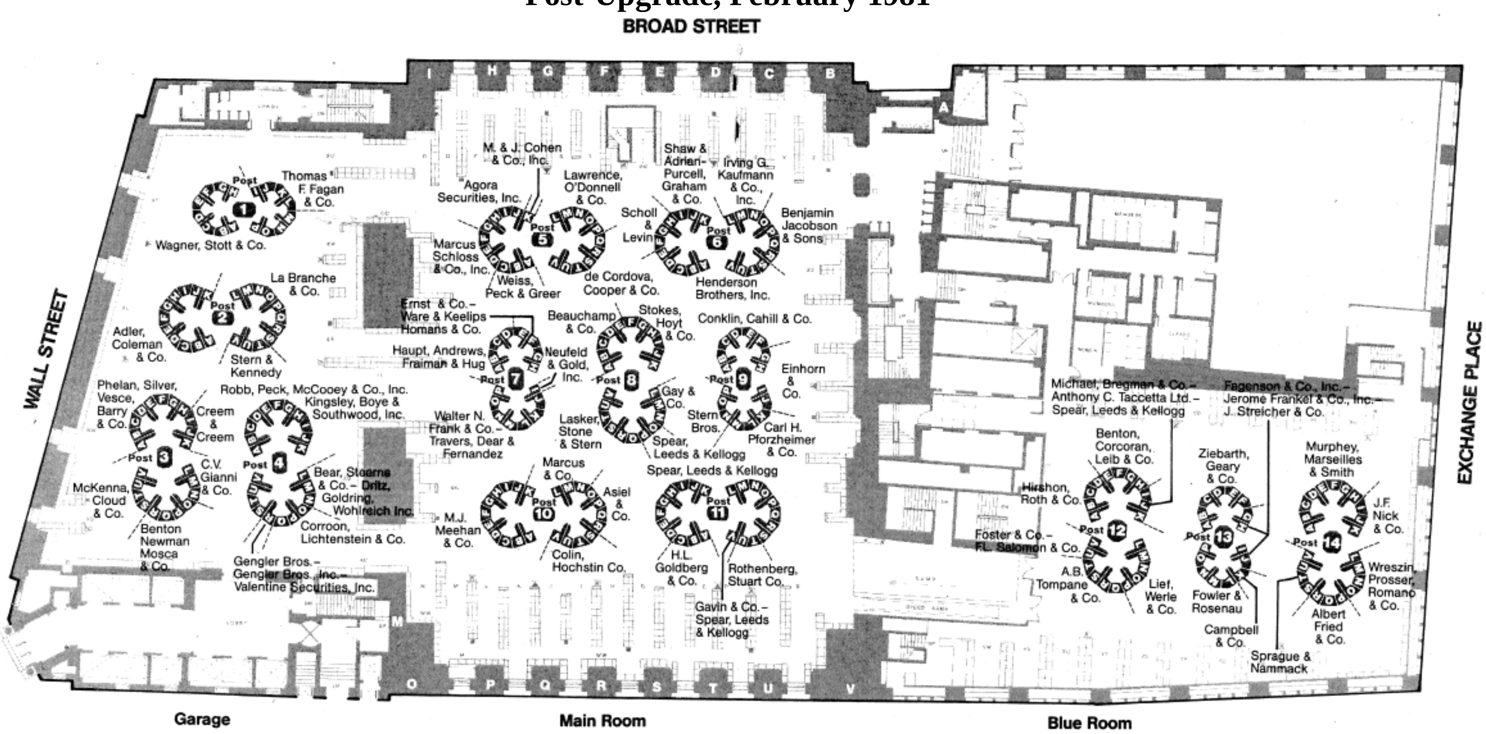

Figure A.2.
Specialist Posts

Specialist Post Pre-Upgrade

Specialist Post Post-Upgrade

Schematic of New Specialist Post

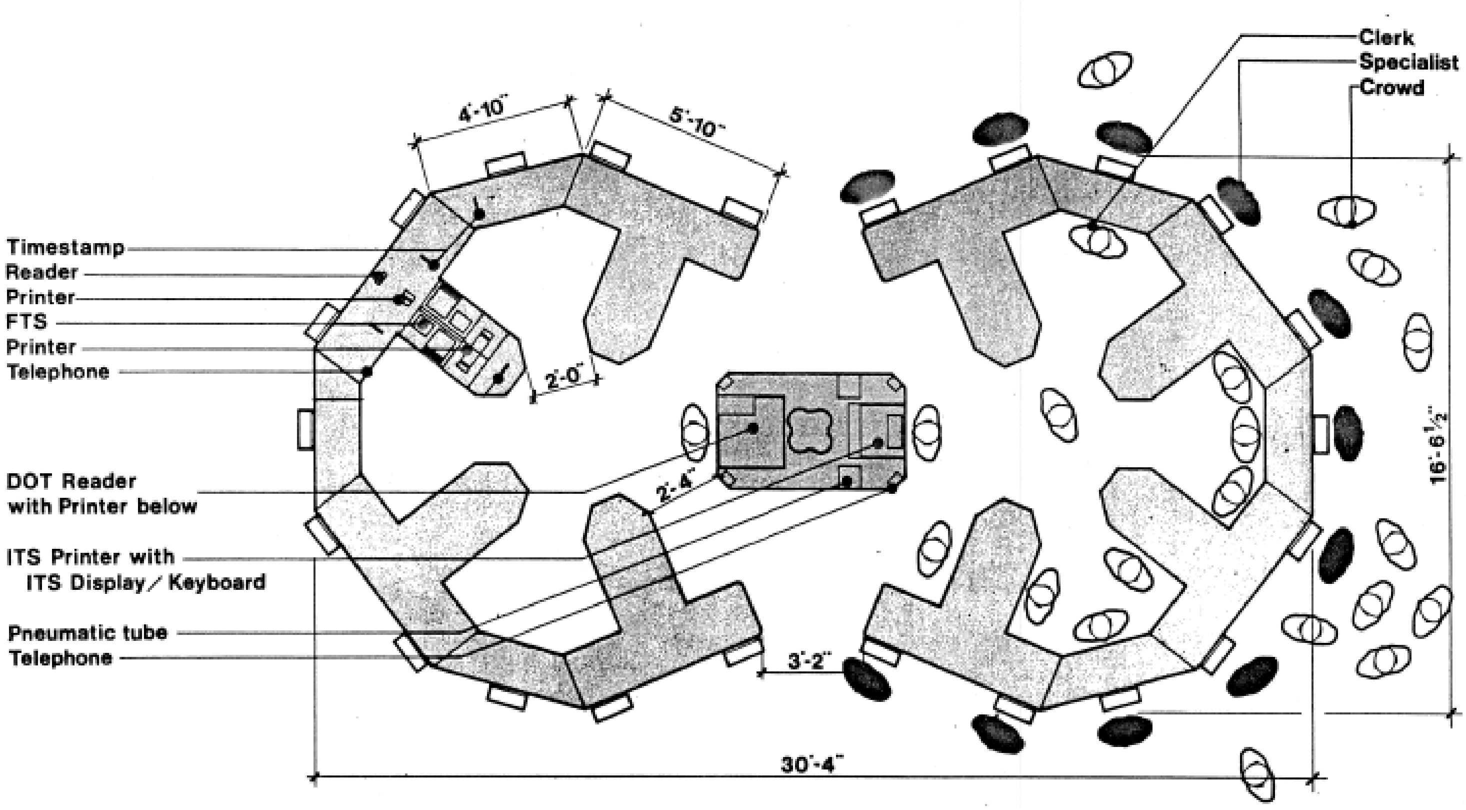

Figure A.3.
The NYSE Trading Floor, Pre- and Post-Upgrade

Pre-Upgrade

1955

1978

Post-Upgrade, 1981

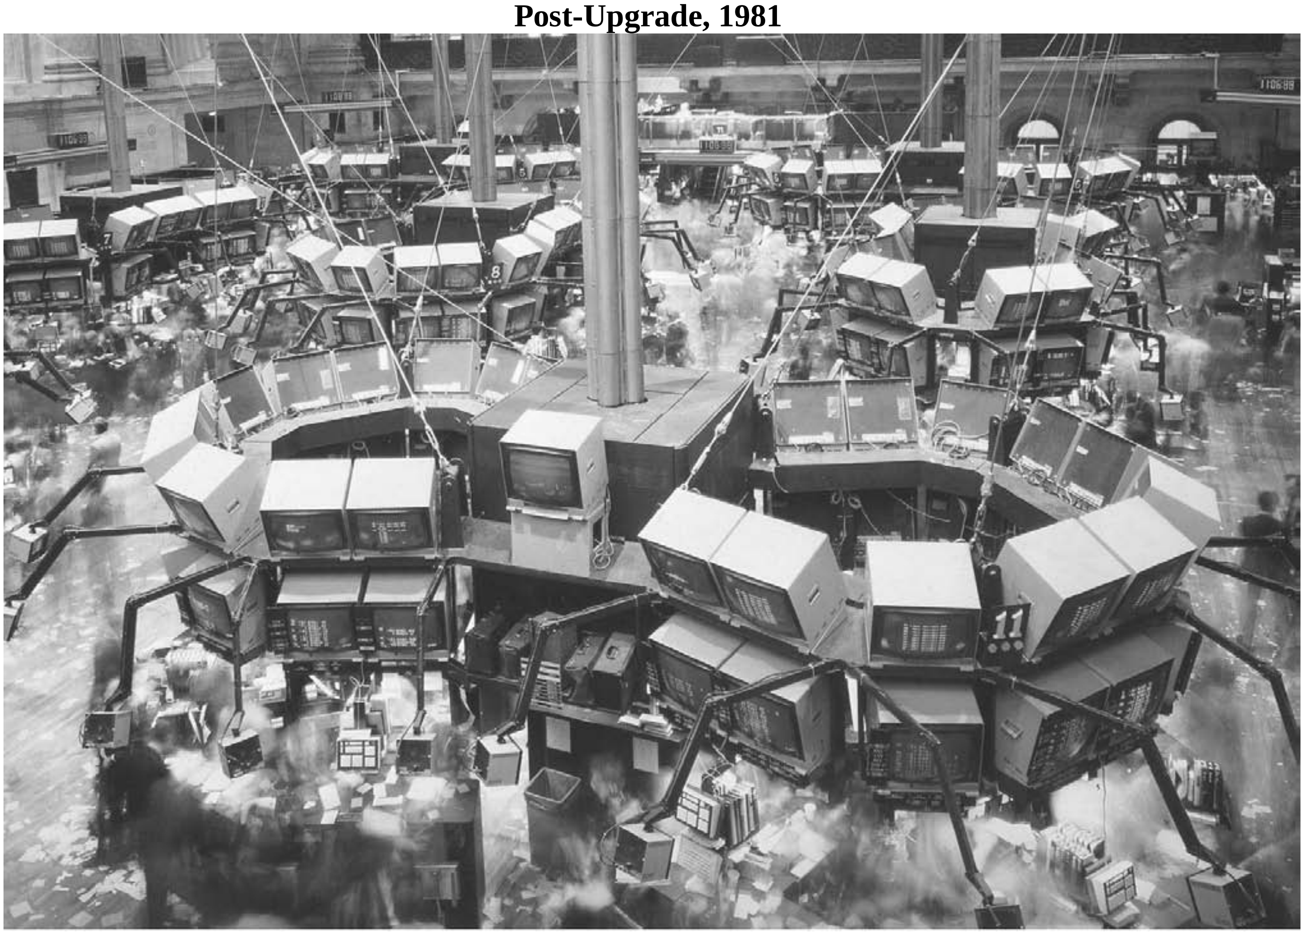

References

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

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

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

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

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

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

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

Boehmer, E., G. Saar and L. Yu, 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., and P. Gompers, 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., and Foucault, T., 2007, Insiders-Outsiders, transparency, and the value of the ticker, working paper.

Edwards, A., L. Harris and M. Piwowar, 2007, Corporate bond market transaction costs and transparency, Journal of Finance 62, 1421-51.

Harris, L., and M. Piwowar, 2006, Municipal bond liquidity, Journal of Finance 61, 1361-1397.

Fama, E., 1998, Market efficiency, long-term returns, and behavioral finance, Journal of Financial Economics 49, 283-306.

Fama, E., Fisher, L., Jensen, M., Roll, R., 1969, The adjustment of stock prices to new information. International Economic Review 10, 1-21.

Flood, M.D., R. Huisman, K. G. Koedijk, and R. J. Mahieu, 1999, Quote disclosure and price discovery in multiple-dealer financial markets, Review of Financial Studies 12, 37-59.

Frame, W.S., and L. J. White, 2004, Empirical studies of financial innovation: Mostly talk and not much action? Journal of Economic Literature 42, 116-144.

Garbade, K. D., and W. L. Silber, 1978, Technology, communication and the performance of financial markets: 1840-1975, Journal of Finance 33, 819-832.

Giacotto, C., and J. M. Sfiridis, 1996, Hypothesis testing in event studies: the case of variance changes, Journal of Economics and Business 48, 349-370.

Glosten, L., and L. Harris, 1988. Estimating the components of the bid-ask spread, Journal of Financial Economics 21, 123–142.

Glosten, L., R. Jaganathan, and D. Runkle, 1993, On the relation between the expected value and the volatility of the nominal excess return on stocks, Journal of Finance 48, 1779-1801.

Goldman, M., H. Sosin, and M. Gatto, 1979, Path dependent options: buy at the low, sell at the high, Journal of Finance 34, 1111-1127

Green, R., B. Hollifield, and N. Schuroff, 2006, Financial intermediation and the costs of trading in an opaque market, Review of Financial Studies 20, 275-314.

Hasbrouck, J., 2006, Trading Costs and returns on U.S. equities: Estimating effective costs from daily data, unpublished working paper.

Jain, P. K., 2005, Financial market design and the equity premium: Electronic versus floor trading, Journal of Finance 60, 2955-2985.

Keith, C., and A. Grody, 1988, Electronic automation at the New York Stock Exchange, in Bruce R. Guild and James Brian Quinn (eds.), Managing Innovation: Cases from the Services Industries (Series on Technology and Social Priorities, National Academies Press), 82-107.

Kothari, S. P., and J. B. Warner, 2005, The econometrics of event studies, in B. Espen Eckbo (ed.), Handbook of Corporate Finance: Empirical Corporate Finance (Handbooks in Finance Series, Elsevier/North-Holland), Chapter 1.

Lyon, J., B. M. Barber, and C. Tsai, 1999, Improved methods of tests of long-horizon abnormal stock returns, Journal of Finance 54, 165-201.

Macey, J. and M. O'Hara, 1999, Regulating exchanges and alternative trading systems: A law and economics perspective, Journal of Legal Studies 28, 17-54.

Madhavan, A., D. Porter, and D. Weaver, 2005, Should securities markets be transparent?, Journal of Financial Markets 8, 265-287.

Mitchell, M., and E. Stafford, 2000, Managerial decisions and long-term stock price performance, Journal of Business 73, 287-329.

Muscarella, C. J. and M. S. Piwowar, 2001, Market microstructure and securities values: Evidence from the Paris Bourse, Journal of Financial Markets 4, 209-230.

Newey, W. K., and K. D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, Econometrica 55, 703-708.

Pastor, L., and R. Stambaugh, 2003, Liquidity risk and expected stock returns, Journal of Political Economy, 111, 642-685.

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., and C. Schenzler, 2006, Trades outside the quotes: reporting delay, trading option or trade size? Journal of Financial Economics 79, 615-653.

Table I
Summary Statistics: Full 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 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**	**Mean Price**	**Mean Return**	**Std. Return**	**Mean Beta Return**	**Mean Turnover**	**Std. Turnover**	**Mean Beta Turnover**
		($MM)	**($)**	**(%)**			**(%)**	**(%)**	
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

Table II
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**	**Mean Price**	**Mean Return**	**Std. Return**	**Mean Beta Return**	**Mean Turnover**	**Std. Turnover**	**Mean Beta Turnover**
		($MM)	**($)**	**(%)**	**(%)**		**(%)**	**(%)**	
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

Table III
Long-Short Portfolio: Returns

This table presents results from estimating:

$$(R_{SecondPhase,t} - R_{FirstPhase,t}) = \sum_{i=1}^{10} a_i D_i + u_t$$

The D_i select the events and time periods denoted in the first column. The other columns denote the stock sample used in the analysis, that is, 'All' stocks, and terciles of stocks sorted by market capitalization, 'Small', 'Medium' and 'Large' stocks. $R_{Group,t}$ is 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 heteroskedasticity and autocorrelation, and presented in italics beneath coefficients. Coefficients significant at the 10% level are highlighted in **bold**.

			Implementation Status			Stock Sample			
Period	**Begin**	**End**	**First Phase**	**Second Phase**	**Coeff.**	**All**	**Small**	**Medium**	**Large**
Pre-Upgrade	1/1/79	6/12/80			a_1	0.002	0.011	0.010	-0.008
						0.150	*0.342*	*0.425*	*-0.458*
Phase I (-19,0)	6/13/80	7/11/80			a_2	0.032	0.086	-0.031	0.073
						0.583	*0.861*	*-0.396*	*1.332*
Phase I (+1,+20)	7/14/80	8/8/80	**I**		a_3	0.088	-0.020	**0.253**	0.039
						0.961	*-0.221*	*1.873*	*0.390*
	8/11/80	10/9/80	**I**		a_4	0.000	0.042	0.082	**-0.095**
						0.006	*0.665*	*1.268*	*-2.152*
Phase II (-19,0)	10/10/80	11/7/80	**I**		a_5	0.029	0.037	0.105	-0.031
						0.570	*0.444*	*0.890*	*-0.341*
Phase II (+1,+20)	11/10/80	12/7/80	**I**	**I & II**	a_6	**0.204**	**0.256**	**0.228**	**0.124**
						3.804	*2.658*	*2.100*	*2.198*
	12/8/80	12/17/80	**I**	**I & II**	a_7	-0.015	-0.087	0.151	-0.175
						-0.152	*-0.756*	*0.641*	*-1.583*
Retrofit (-19,0)	12/18/80	1/16/81	**I**	**I & II**	a_8	-0.046	0.036	-0.129	-0.021
						-0.785	*0.363*	*-1.424*	*-0.259*
Retrofit (+1,+20)	1/19/81	2/13/81	**I & II**	**I & II**	a_9	-0.093	**-0.243**	-0.074	0.028
						-1.196	*-1.834*	*-0.594*	*0.367*
Post-Upgrade	2/16/81	7/13/81	**I & II**	**I & II**	a_{10}	-0.007	0.061	-0.025	-0.041
						-0.212	*1.055*	*-0.479*	*-0.998*
					R^2	0.021	0.012	0.022	0.015
					N	639	639	639	639

Table IV
Long-Short Portfolio: Four-Factor Model Adjusted Returns

This table presents results from estimating:

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

The D_i select the events and time periods denoted in the first column. The other columns denote the stock sample used in the analysis, that is, 'All' stocks, and terciles of stocks sorted by market capitalization, 'Small', 'Medium' and 'Large' stocks. $R_{Group,t}$ is the equal-weighted average return across all stocks in each group, that is, First Phase stocks or Second Phase stocks on day t.

$R_{mt} - R_{ft}, SMB_t, HML_t, UMD_t$ are the four Fama-French factors. Newey-West t-statistics are robust to heteroskedasticity and autocorrelation, and presented in italics beneath coefficients. Coefficients significant at the 10% level are highlighted in **bold**.

			Implementation Status			Stock Sample			
Period	**Begin**	**End**	**First Phase**	**Second Phase**	**Coeff.**	**All**	**Small**	**Medium**	**Large**
Pre-Upgrade	1/1/79	6/12/80			$\mathbf{a_1}$	0.004	0.003	0.008	0.000
						0.282	*0.106*	*0.338*	*-0.004*
Phase I (-19,0)	6/13/80	7/11/80			$\mathbf{a_2}$	0.025	0.089	-0.059	0.051
						0.421	*0.815*	*-0.737*	*0.969*
Phase I (+1,+20)	7/14/80	8/8/80	**I**		$\mathbf{a_3}$	0.062	-0.021	0.188	-0.002
						0.713	*-0.207*	*1.418*	*-0.025*
	8/11/80	10/9/80	**I**		$\mathbf{a_4}$	-0.012	0.026	0.022	**-0.072**
						-0.383	*0.373*	*0.328*	*-1.833*
Phase II (-19,0)	10/10/80	11/7/80	**I**		$\mathbf{a_5}$	0.038	0.043	0.093	0.005
						0.667	*0.480*	*0.771*	*0.049*
Phase II (+1,+20)	11/10/80	12/7/80	**I**	**I & II**	$\mathbf{a_6}$	**0.192**	**0.221**	**0.166**	**0.181**
						3.517	*1.774*	*1.712*	*3.179*
	12/8/80	12/17/80	**I**	**I & II**	$\mathbf{a_7}$	-0.020	-0.093	0.199	-0.141
						-0.254	*-0.713*	*1.250*	*-1.188*
Retrofit (-19,0)	12/18/80	1/16/81	**I**	**I & II**	$\mathbf{a_8}$	-0.034	0.047	-0.075	-0.080
						-0.647	*0.476*	*-0.961*	*-0.961*
Retrofit (+1,+20)	1/19/81	2/13/81	**I & II**	**I & II**	$\mathbf{a_9}$	-0.051	-0.209	-0.022	0.043
						-0.669	*-1.622*	*-0.187*	*0.606*
Post-Upgrade	2/16/81	7/13/81	**I & II**	**I & II**	$\mathbf{a_{10}}$	0.003	0.060	-0.011	-0.039
						0.103	*1.044*	*-0.235*	*-0.892*
					$\mathbf{R_m\text{-}R_f}$	**0.048**	**0.110**	0.005	**0.074**
						1.907	*2.428*	*0.131*	*2.800*
					SMB	**-0.067**	**-0.039**	0.045	**-0.026**
						1.907	*2.428*	*0.131*	*2.800*
					HML	**-0.109**	**0.073**	-0.229	**-0.107**
						1.907	*2.428*	*0.131*	*2.800*
					UMD	-0.046	0.067	0.009	**-0.191**
						-1.476	*1.087*	*0.182*	*-5.106*
					$\mathbf{R^2}$	0.082	0.046	0.057	0.074
					N	639	639	639	639

Table V
Matched Long-Short Portfolio: Returns

This table presents results from estimating the following equation:

$$(R_{MatchSecondPhase,t} - R_{MatchFirstPhase,t}) = \sum_{i=1}^{10} a_i D_i + u_t \cdot$$

Each stock in Phase I is matched with a stock in Phase II. The dummy variables D_i select the different events and time periods denoted in the first column. The remaining columns denote the group of stocks within each of the phases which are used in the analysis, that is, 'All' stocks, and terciles of stocks sorted by market capitalization, 'Small', 'Medium' and 'Large' stocks. For the matched sample, each First Phase stock is matched with a Second Phase stock. $R_{MatchGroup,t}$ is the equal-weighted average return across each stocks in each group, that is, matched First Phase stocks or matched Second Phase stocks on day t. $R_{mt} - R_{ft}, SMB_t, HML_t, UMD_t$ are the four Fama-French factors. Newey-West t-statistics are robust to heteroskedasticity and autocorrelation, and presented in italics beneath coefficients. Coefficients significant at the 10% level are highlighted in **bold**.

			Implementation Status			Matched Stock Sample			
Period	**Begin**	**End**	**First Phase**	**Second Phase**	**Coeff.**	**All**	**Small**	**Medium**	**Large**
Pre-Upgrade	1/1/79	6/12/80			$\mathbf{a_1}$	0.007	-0.002	0.021	0.004
						0.437	*-0.058*	*0.723*	*0.225*
Phase I (-19,0)	6/13/80	7/11/80			$\mathbf{a_2}$	0.009	0.030	-0.024	0.018
						0.155	*0.223*	*-0.239*	*0.361*
Phase I (+1,+20)	7/14/80	8/8/80	**I**		$\mathbf{a_3}$	0.116	0.099	**0.193**	0.058
						1.178	*0.718*	*1.754*	*0.615*
	8/11/80	10/9/80	**I**		$\mathbf{a_4}$	-0.009	0.012	0.048	-0.084
						-0.212	*0.143*	*0.615*	*-1.515*
Phase II (-19,0)	10/10/80	11/7/80	**I**		$\mathbf{a_5}$	0.103	0.154	0.210	-0.047
						1.627	*1.489*	*1.373*	*-0.592*
Phase II (+1,+20)	11/10/80	12/7/80	**I**	**I & II**	$\mathbf{a_6}$	**0.162**	0.143	0.152	**0.191**
						2.896	*1.521*	*1.474*	*2.441*
	12/8/80	12/17/80	**I**	**I & II**	$\mathbf{a_7}$	-0.012	-0.097	0.220	-0.152
						-0.170	*-0.495*	*0.816*	*-1.504*
Retrofit (-19,0)	12/18/80	1/16/81	**I**	**I & II**	$\mathbf{a_8}$	-0.068	-0.025	-0.194	0.011
						-0.802	*-0.168*	*-1.506*	*0.108*
Retrofit (+1,+20)	1/19/81	2/13/81	**I & II**	**I & II**	$\mathbf{a_9}$	-0.105	**-0.233**	-0.022	-0.061
						-1.417	*-2.248*	*-0.210*	*-0.645*
Post-Upgrade	2/16/81	7/13/81	**I & II**	**I & II**	$\mathbf{a_{10}}$	-0.044	0.038	-0.081	**-0.086**
						-1.368	*0.630*	*-1.398*	*-2.044*
					$\mathbf{R^2}$	0.021	0.008	0.021	0.018
					N	639	639	639	639

Table VI
Turnover over Time

This table presents the average turnover (dollar trading volume as a percentage of stock market capitalization) 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 "Δa_i-Δa_1" presents the difference-in-differences 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 heteroskedasticity 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. Δa_i-Δa_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 highlighted in **bold**.

			Implementation Status			Stock Sample			
Period	**Begin**	**End**	**First Phase**	**Second Phase**	**Coeff.**	**First Phase**	**Second Phase**	**Δa_i-Δa_1**	**p-value**
Pre-Upgrade	4/14/80	6/12/80			$\mathbf{a_{1,p}}$	0.140	0.144		
Pre-Phase I	6/13/80	7/11/80			$\mathbf{a_{2,p}}$	0.177	0.180	-0.001	0.981
Phase I	7/14/80	11/7/80	**I**		$\mathbf{a_{3,p}}$	0.196	0.208	0.008	0.260
Phase II	11/10/80	1/16/81	**I**	**I & II**	$\mathbf{a_{4,p}}$	0.196	0.229	**0.029**	0.006
Retrofit	1/19/81	2/13/81	**I & II**	**I & II**	$\mathbf{a_{5,p}}$	0.151	0.168	0.013	0.219
Post-Upgrade	2/16/81	7/13/81	**I & II**	**I & II**	$\mathbf{a_{6,p}}$	0.190	0.191	-0.003	0.674

Table VII
Transactions Costs over Time

This table presents transactions costs estimates 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 "Δa_i-Δa_1" presents the difference-in-differences estimates for transactions costs in each time-period, and the column entitled p-value shows the statistical significance of this difference computed using a heteroskedasticity 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 percent for First Phase stocks, and 0.420 percent for Second Phase stocks. Δa_i-Δa_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 highlighted in **bold**.

			Implementation Status			**Stock Sample**			
Period	**Begin**	**End**	**First Phase**	**Second Phase**	**Coeff.**	**First Phase**	**Second Phase**	**Δa_i-Δa_1**	**p-value**
Pre-Upgrade	4/14/80	6/12/80			$\mathbf{a_{1,p}}$	0.562	0.538		
Pre-Phase I	6/13/80	7/11/80			$\mathbf{a_{2,p}}$	0.505	0.470	-0.012	0.594
Phase I	7/14/80	11/7/80	**I**		$\mathbf{a_{3,p}}$	0.447	0.420	-0.004	0.839
Phase II	11/10/80	1/16/81	**I**	**I & II**	$\mathbf{a_{4,p}}$	0.492	0.427	**-0.042**	0.046
Retrofit	1/19/81	2/13/81	**I & II**	**I & II**	$\mathbf{a_{5,p}}$	0.486	0.450	-0.012	0.587
Post-Upgrade	2/16/81	7/13/81	**I & II**	**I & II**	$\mathbf{a_{6,p}}$	0.431	0.405	-0.003	0.906

Table VIII
Cross-Sectional Evidence on Returns and Transactions Costs

Table IX 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 heteroskedasticity and autocorrelation, and presented in italics beneath coefficients. Coefficients significant at the 10% level are highlighted in **bold**.

Explanatory Variables	**Return Difference**	**Return Difference**	**Return Difference**
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

Table IX
Lookback Straddle Option Value Calculations

Table IX presents estimates of lookback 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 I), and the capitalized present value of the difference, computed using the average risk-free rate between 1926 and 1980, of 0.0277. Panel A of the table calculates this difference between 2 minute and 20 second options, and Panel B between 5 minute and 20 second options.

Panel A

Lookback Option Value	T=2	T=0.33	Difference	Turnover Adjusted	Capitalized
Mean	5.599	1.850	3.749	1.800	64.964
Median	5.283	1.778	3.424	1.643	59.329
Standard Deviation	2.184	0.708	1.681	0.807	29.138

Panel B

Lookback 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
StDev	3.485	0.708	3.028	1.453	52.464

Figure 1
Event Study

Panel A (Panel B) of this figure plots the cumulative equal-weighted abnormal return from a market model across all 327 stocks for which we can identify the implementation date of Phase I between July 14, 1980 and October 27, 1980 (286 stocks for which we can identify the implementation date of Phases II, III and IV between November 10, 1980 and January 5, 1981). See Table A.1. in the Appendix for the detailed post upgrade implementation schedule. The vertical bar shows the implementation date of the phase (stocks are lined up in event time) highlighted in the legend beneath each graph. The grey (black) line cumulates the return in the window prior to (following) the implementation of the upgrade.

Panel A: Phase I

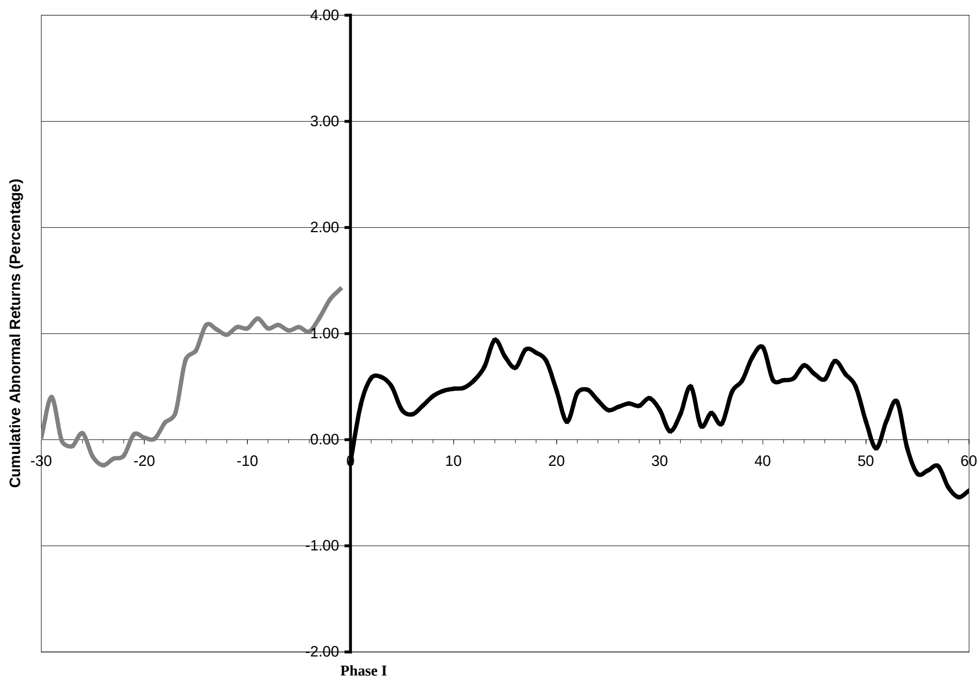

Panel B: Phase II

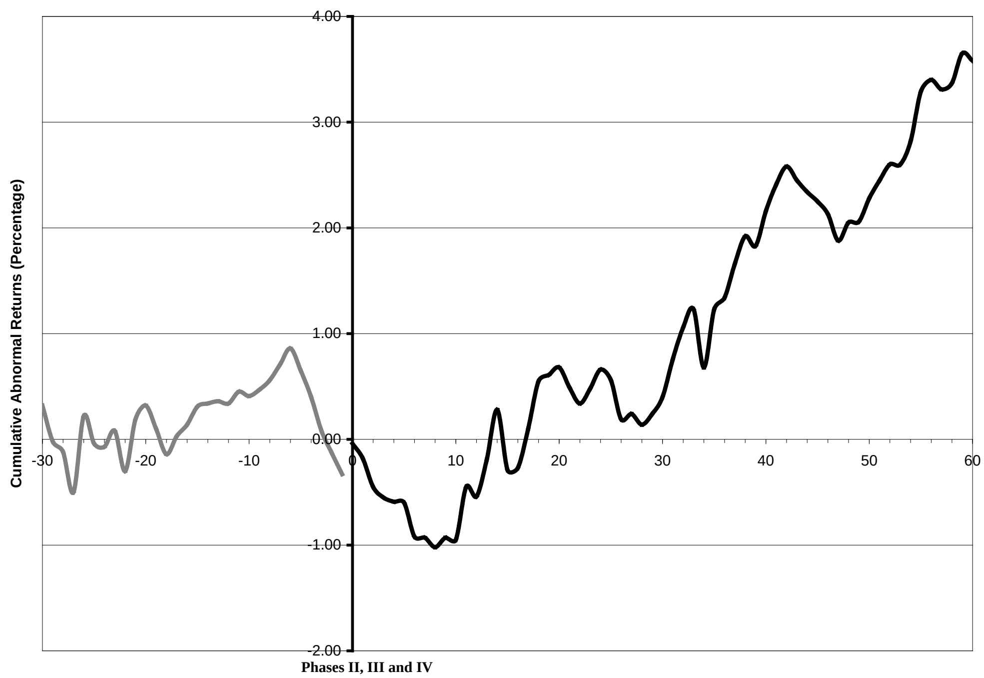

Figure 2
Event Time Returns to Long-Short Portfolio

This figure plots the cumulative difference in abnormal equal-weighted returns between Second Phase and First Phase stocks for the 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). 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 in the graph indicate the event date for each line corresponding to the calendar date November 10, 1980, i.e., the date at which Phase II upgrades began.

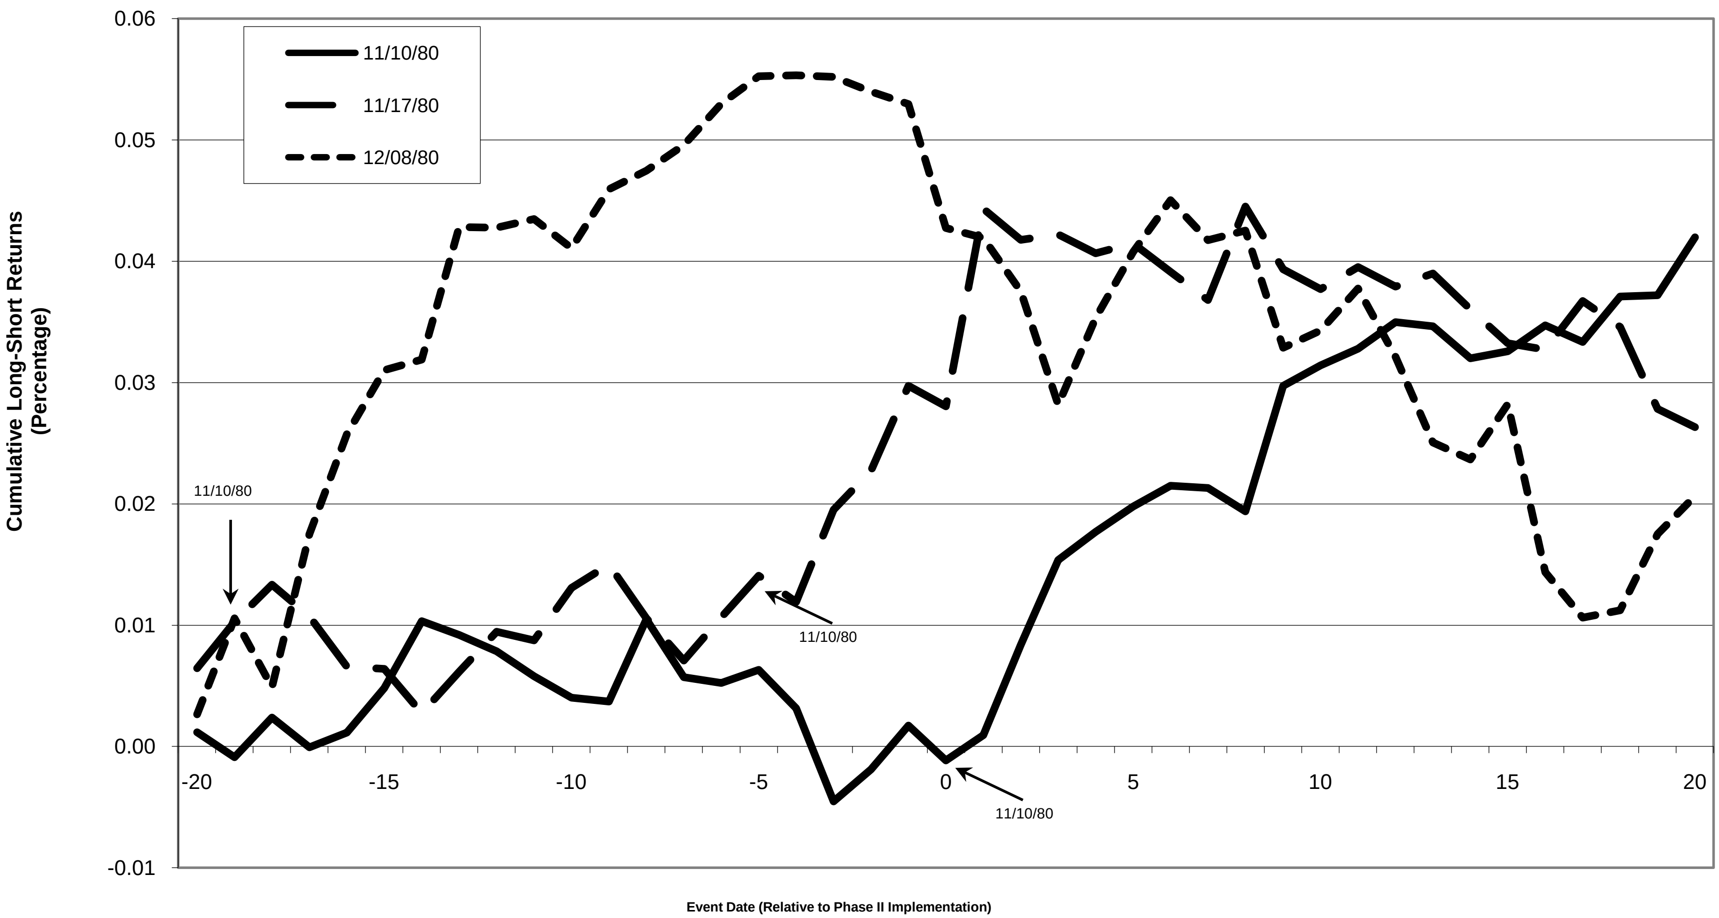

Figure 3
Calendar Time Returns to Long-Short Portfolio

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 beneath each graph.

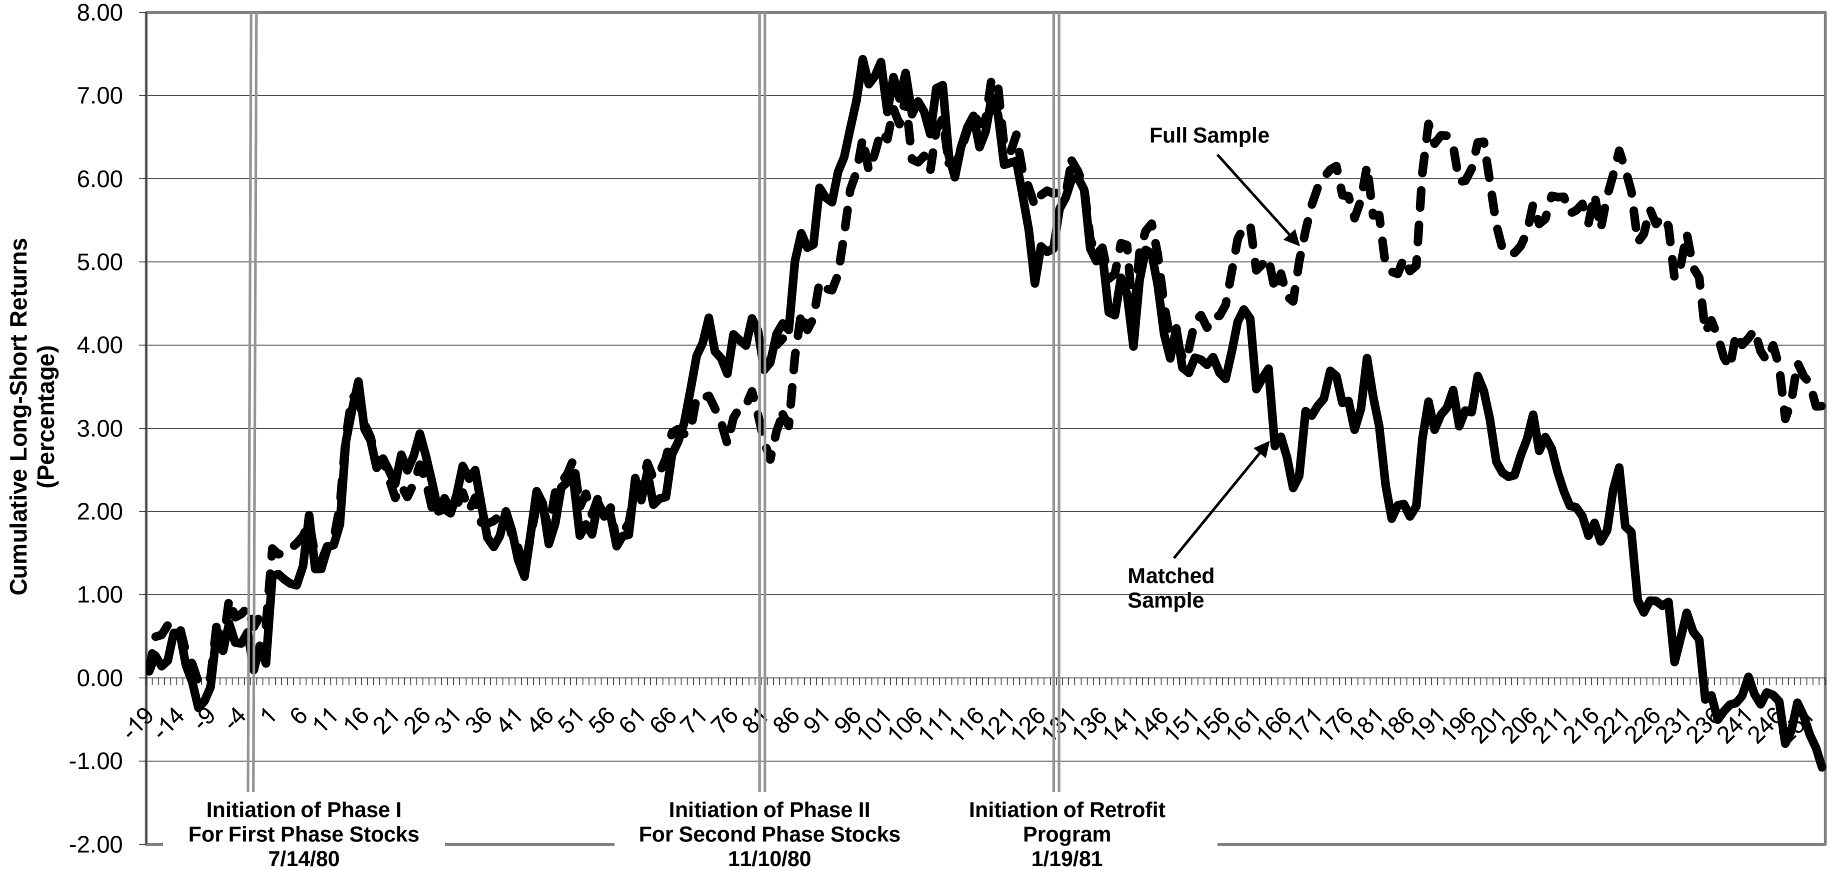

Figure 4
Calendar Time Turnover of Long-Short Portfolio

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 I. The vertical bars show the implementation dates of various phases highlighted in the legend beneath each graph.