Latency, Liquidity and Price Discovery

Ryan Riordan[a,1], Andreas Storkenmaier[a,1]

[a] *School of Economics and Business Engineering, Karlsruhe Institute of Technology, Germany*

Abstract

On April 23rd, 2007, Deutsche Boerse made the most important upgrade to their trading system since 2002. With the 8.0 release of Xetra, system latency was reduced from 50 ms to 10 ms round trip. Trading costs decreased by between 1 and 4 basis points. The liquidity increase is due to dramatically lower adverse selection costs that are only partially translated into higher liquidity. This is interpreted as a decrease in the competition between liquidity suppliers and specifically between liquidity suppliers sensitive to latency and those that are not. The adverse selection component of trades falls dramatically post-upgrade and price discovery shifts from trade-correlated to quote-correlated. These findings demonstrate the importance of latency in modern electronic securities markets.

Key words: Liquidity, Latency, Execution Speed, Exchange Systems, Algorithmic Trading
JEL G14

Email addresses: `Riordan@iism.uni-karlsruhe.de` (Ryan Riordan), `Storkenmaier@iism.uni-karlsruhe.de` (Andreas Storkenmaier)

[1] The authors are Ph.D. students at the Graduate School Information and Market Engineering at the Karlsruhe Institute of Technology, Germany funded by the Deutsche Forschungsgemeinschaft (German Research Foundation)

With the advent of fully electronic trading (Jain, 2005), the IT-systems exchanges use to match and report orders are becoming increasingly important. Investors are increasingly using technology to translate their investment decisions into orders. This increased use of technology has also driven an unprecedented increase in the number of orders per trade, volume, and the number of transactions. These dual trends in securities trading have increased the strain on the systems used to match and report trades. This increase in both systems capacity and algorithmic trading (AT) can have a number of effects on liquidity and information. Directly, we study the effects of latency on liquidity in this paper.

As more investors are using technology to mange their orders, exchange systems can come under great stress. An unfortunate example of this is the recent London Stock Exchange (LSE) system failure[2] that caused trading to cease for seven hours. This system failure, caused by a flood of orders in response to the Freddie Mac and Fannie Mae bailout, makes painfully clear the importance of IT-systems to the operations of financial exchanges worldwide. In response to increases in algorithmic and quantitative trading, exchanges[3] have been upgrading their infrastructure to reduce system latency and thereby increase the number of orders handled in a period of time. Exchanges also tout the liquidity-improving effects of reduced system latency and increased order-handling capacity. We test the hypothesis that reducing system latency increases liquidity as well as the equally important question as to how this affects the processing of information in a market microstructure context.

Latency is critical in electronic trading. Trading strategies that rely on short-term relative price differentials, like index arbitrage or correlated pairs trading, require near-simultaneous execution and face execution risks increasing in latency. Also, if some traders receive pricing relevant information before others the former can exercise the free-trading option offered by slower investors. Latency can also affect the compensation liquidity suppliers require for the free trading option they supply. While the effect of latency on liquidity is unclear, what is clear is that decreasing latency changes the competitive factors in the demand and supply of liquidity

[2]See September 10, 2008 Wall Street Journal - Back Online, LSE Faces Skeptics

[3]Deutsche Boerse rolled out their Xetra 8 system on April 23rd, 2007, the London Stock Exchange in June of 2007, the New York Stock Exchange staggered during Q4 2006 and Q1 2007, and the Toronto Stock Exchange rolled out Quantum in December of 2007

and how quotes are updated to reflect public information.

This is the first known study to isolate the effect of latency on liquidity and information processing in an electronic limit order market. The natural experiment provided by the April 23rd, 2007 Xetra 8.0 trading system upgrade at Deutsche Boerse, is used to test the hypothesis that reducing latency impacts liquidity and information. The Xetra 8.0 upgrade is unique in that a number of system changes were made simultaneously with the sole purpose of reducing latency. The upgrade included no market model, mechanism, or other meaningful microstructure changes beside a latency-reducing system upgrade. Deutsche Boerse also did not change the recording and dissemination of tick data.

Latency is commonly defined as the amount of time it takes for a trader to receive feedback about a submitted order. Using the example of a marketable buy order, latency is the amount of time that elapses between submitting the order and receiving confirmation that the order executed at a given quantity and price. Latency in an electronic order-driven market, in contrast to market-maker markets, is determined entirely by the IT-systems and algorithms supporting the operations of an exchange. In market-maker markets or market-maker segments in hybrid setups, latency is a function of the IT-systems routing an order to the floor and the time it takes for a human to process the order.

We use a panel-estimation technique that adjusts for cross-sectional variations to test for changes in liquidity. The results show that a reduction in latency has a positive effect on liquidity. The results hold across market capitalization (Mcap) quartiles (Q1, Q2, Q3, and Q4) and trade sizes. Interestingly, the increase in liquidity is driven by a reduction in the adverse selection component of trades — which falls dramatically. We test using a VAR framework (Hasbrouck (1991a) and Hasbrouck (1991b)) find that the permanent price impact falls post upgrade. The permanent price impact per trade falls from roughly 3 basis points (bps) on average to 0.6 bps for large stocks. Quotes have also become more informative with an increase in the total amount of information attributable to quote changes increasing from roughly 40 % to 90 %. This reduction in adverse selection costs translates into a small liquidity increase.

Only recently have researchers begun to focus on the effects of AT on market outcomes. Until recently, little data was available on AT trades and orders. Two recent studies; Chaboud et al. (2009) and Hendershott and Riordan (2009), are the first to have access to detailed AT data. Chaboud et al. (2009) study the effects of AT in FX markets using three widely traded currency pairs. Two of their findings are interesting: they find that AT is not related to volatility and that the variance in FX returns is not related to AT order flow. Although no direct evidence is found on average AT versus human information in this study the findings indicate that the relationship between trade-correlated and trade-uncorrelated information changes dramatically. Hendershott and Riordan (2009) look directly at the average level of AT and human information (measured as the variance of the random-walk component of returns) and find that AT trades and quotes are more informative human trades and quotes.

Our results contrast somewhat with both earlier empirical studies of transaction costs (Demsetz, 1968) and more recent studies (Bacidore et al. (2003), Battalio et al. (2003), Bennett and Wei (2006), and Boehmer (2005)) which find a trade-off between speed and cost. Boehmer (2005) states that there is a trade-off between costs and speed that is robust over time and insensitive to the econometric specification. The primary difference between these studies and the current one is that herein the effect of a system-wide latency reduction is studied. The previous studies are in effect studying two similar but unrelated effects.

The question of execution speed involves both a microstructure and an infrastructure component. Using the example of market versus limit orders, a market order executes faster than does a limit order, but pays the spread. A limit order takes longer to execute but receives the spread. Although there is a time component, it is not the driving factor. A limit order takes longer to execute but does not necessarily cost less as the order execution time increases. In fact, there is no relationship between the waiting time of two equivalent limit orders and the execution costs. Bacidore et al. (2003) address exactly this issue in their analysis of guaranteed and non-guaranteed orders. A guaranteed order waits on average 251.1 seconds for execution and costs 17 cents. A non-guaranteed order costs only 7 cents but executes in under $17 seconds$. The same study also notes that this difference in time to execution is the time it takes for the order to interact with the floor trader. The electronic transit time was roughly 6.3 seconds in their sample. This example highlights two components of latency. The first is a microstructure one, such as waiting for the floor to interact with an order. The second is an infrastructure effect and the focus of this work. Our findings in no way contradict these previous studies. Rather, they represent the study of two similar but distinct questions.

The literature also shows that traders unambiguously prefer fast execution to slower execution, holding costs and other factors equal. Blume (2001) cites a survey from 2000 from Sanford and Bernstein where 58% of online investors state that:

> "...immediacy of execution is more important than a favorable price.[4]"

Huang (2002) finds that the timeliness of information reflected in quotes is an important issue and that ECNs' quotes are more informative because of the speed with which they reflect information. Huang equates the speed of the trading system with the informativeness of quotes. Clearly, there are clientele that prefer fast execution and are willing to pay extra. The trade-off between costs and speed is actually a tradeoff between execution risk and profit. If the profit is minimally affected but the execution risk lowered considerably these traders are more likely to prefer faster execution. They are also much more likely to trade even when the per trade profit is very small.

Copeland and Galai (1983) made an important contribution towards the understanding of trading. They formulated one of the key costs incurred by liquidity suppliers, something they called the 'free-trading option'. The free-trading option is the option a liquidity supplier, or market-maker, supplies to the market when they provide firm commitments to trade. Both quotes and limit orders are firm commitments to trade and are exposed to the free-trading option risk. For a liquidity supplier the free-trading option is a cost that increases in time. The longer one of their free-trading options is on a market the more valuable it is. In the following hypothesis development section we show that by reducing the duration of the free-trading option which is the same as lowering latency, the cost to a liquidity supplier of the free-trading option, is lowered. To our knowledge, this is the only theoretical framework that has an explicit mechanism for latency.

Theoretical studies on limit order markets (Cespa and Foucault (2008), Foucault et al. (2005), Parlour (1998)) make some assumptions about the timeliness and speed of trade. They assume that an increase in speed also causes an increase in the informational efficiency of prices as in Cespa and Foucault (2008). They also assume that faster trade means higher cost, although this

[4]Weekly Notes. Bernstein Research (New York: Sanford C.Bernstein & Co., Inc., May 12, 2000)

assumption is attributed to market microstructure (limit vs. market order) rather than a direct market latency effect.

A study that is similar in nature to ours is the Hendershott and Moulton (2007) study of the effect of the introduction of the New York Stock Exchange's (NYSE) Hybrid system. While similar in theory to this study it is different in that the NYSE not only increased execution speed but simultaneously made changes to the market structure. Pre-hybrid automatic execution in a limit order environment is only available to orders under a volume 1,099. Post-hybrid orders can be automatically executed up to 1,000,000 thereby circumventing the NYSE specialist. These results differ in that the previous studies find an increase in both quoted and effective spread, whereas we find only a decrease in the effective spread. They attribute the increase to an increase in the adverse selection costs caused by anonymous trading.

Another recent study titled 'The Price of Latency' (Easely et al., 2007) studies the effect of a latency reduction at the NYSE in 1983. The results are interesting in that they find an increase in the price of stocks that switch from higher latency trading to lower latency trading. They also find a reduction in transaction costs of 13 basis points after the switch from higher to lower latency. Our study is somewhat different in that the focus is on studying an exclusively technological change at Deutsche Boerse (DB) with no accompanying market model or microstructure changes. Another differentiating factor is the trading era. In 1983 investors traded almost exclusively manually, in 2007 roughly 39 % of trade was algorithmically generated on Xetra[5].

White and Frame (2004) study studies of financial innovation. They report far too few empirical studies of financial innovations and analyses of their impacts. By studying an exchange system upgrade in a period of electronic low-latency trading, the focus is on the study of two parallel innovations and the interaction of both.

Section 1 provides an overview over Deutsche Boerse's Xetra System. Section 2 develops our hypotheses. Data and methodology are explained in section 3. Section 4 provides the spread analysis and section 5 presents the results from our information analysis. Section **??** discusses and interprets our results and section 6 finally concludes.

[5]See the 2007 Deutsche Boerse Annual Report p. 81

1. Xetra

The Xetra ("Exchange Electronic Trading") system is the electronic cash market stock trading system operated by Deutsche Boerse (DB). The Xetra system was originally introduced in 1997 and was the first fully-electronic stock exchange in Germany. Currently Xetra handles 97[6] percent of German equity trading by volume. Trading begins at 9:00 am and ends with a closing call auction at 5:30 pm during the sample period. The prices on Xetra are used to calculate the DAX ("Deutscher Aktien Index").

Xetra is organized as a centralized limit order book. Incoming orders are compared to existing orders stored in the book. If the price of the incoming order crosses the price of an existing order they are matched. Xetra follows a price and time priority matching rule meaning that orders are matched first based on price and then on time. As an example: given two limit orders with the same direction (sell or buy) and price but different submission times, the order with the lower submission time (i.e. the oldest of the two orders) will be executed to its entire quantity before any quantity of the second order is executed.

Xetra is a completely electronic trading system accessible worldwide. Xetra members are based mostly in Germany but there are a large number of foreign members based in the UK, France, and other countries. Presently there are over 260 participants banks and financial institutions from over 19 countries and more than 2,600 authorized traders. DB admits participants wishing to trade on Xetra based on regulations set and monitored by German and European financial regulators. After being admitted, participants can only connect electronically to Xetra; floor trading is operated separately with no direct interaction between the two trading segments.

Xetra is implemented as an electronic limit order book with trading split into phases as follows:

- Opening call auction with a random ending that opens trading at 9:00 AM;
- A continuous trading period;
- A two-minute intra-day call auction at 1:00 PM with a random ending;
- A second continuous trading period;

[6]See Deutsche Borse Annual Report 2007

- A closing call auction from 5:30 PM to 5:35 PM with a random ending.

The following analyses focus on trade occurring during the two continuous trading periods. Liquidity is provided by public limit orders displayed in the order book of each stock. Orders execute automatically when an incoming market, or marketable limit order, crosses with an outstanding limit order. Order execution preference is determined using price-time priorities. Three types of orders are permitted — limit, market, and iceberg orders. Iceberg orders are orders that display only a portion of their total size. Iceberg orders sacrifice time priority on the non-displayed portion. Pre-trade transparency includes the 10 best bids and ask prices and quantities but not the ID of the submitting participant (as on the Paris Bourse (Venkataraman, 2001)). Trade price and size are disseminated immediately to all participants. The tick size during our sample period is 1 Euro cent.

1.1. Xetra as a Trading System

Xetra is not only an electronic exchange used to trade blue chip German stocks but is also the underlying trading system used by the Irish, Vienna, and Budapest stock exchange as well as the Eurex derivatives exchange. The Xetra system supports a number of trading *modi*, including continuous double auctions, call auctions, and bi-lateral trading. Trading on Xetra can also be performed via designated sponsors as described in Klar and van den Bongard (2008). In fact, trading can be supported by more than one designated sponsor, similar to the Euronext system described in detail in Menkveld and Wang (2008).

1.2. Xetra 8.0

On April 23rd 2007 DB introduced its new trading system, its release, and the effects thereof, is the subject of this study. Xetra 8.0 was the first major system upgrade since the Xetra 7.0 release on August 20th 2002. The release of version 8.0 is interesting in that it provides an ideal opportunity to study the effect of trading system latency, or speed, in isolation. DB introduced no market model changes, they did not introduce different execution mechanisms, nor new order types. They made a series of system upgrades with the sole purpose of reducing the latency of the Xetra trading system. According to Market and Data Analytics at Deutsche Boerse which operates Xetra they did not change their data recording and dissemination.

The new Xetra trading system is designed to reduce the trading system latency from a minimum of 50 milliseconds to a minimum of 10 milliseconds. The most important upgrade was to the trade-matching algorithm and system used to match incoming orders against one another. Previously, each incoming order was stored by the matching algorithm on the physical hard-drive before being matched and reported. Post upgrade each order is matched in 'virtual' memory saving the computationally expensive operation of storing each order on a physical drive before matching. Other important upgrades include an increase in the networking bandwidth to members, and an internal networking upgrade.

The most important feature of the upgrade was the reduction in the market latency. The time between order entry and confirmation was reduced from 50 milliseconds to 10 milliseconds for 'speed-sensitive' traders. The following is a list of the most important latency reducing upgrades made:

- split market data streams
- improved caching
- memory based order matching
- increased network capacity

Each improvement focuses on a specific latency problem. Splitting market data streams avoids data bottlenecks that arise when market participants are forced to receive every order book update for every stock traded on the Xetra system. The new system allows participants to select the market segments for which they wish to receive order book data. Caching improves the speed of the trade matching algorithm as does memory based order matching. DB also dramatically improved their network capacity to deal with the increase in data and communications network requirements. See figure 1 for a diagram depicting the changes made to the Xetra system.

[INSERT FIGURE 1 HERE]

1.2.1. Algorithmic Trading on Xetra

The Xetra 8.0 release was targeted directly at reducing latency with the goal of increasing algorithmic trade. Recent research (Hendershott et al., 2008) presents the effects on liquidity of increased algorithmic trading. Exchanges themselves promote latency reductions and the resulting increase in algorithmic trade as positive for liquidity, which is confirmed in Hendershott et al. (2008). In Q1 DB reported[7] that 45% of trades in their own stock was executed by algorithmic traders up from 37% for the full year 2007. Liquidity has undoubtedly increased on Xetra, and worldwide (Chordia et al., 2007), but a causal relationship between algorithmic trading, latency, and liquidity is currently an open question.

To encourage automated trading DB introduced the 'Automated Trading Program' (ATP) in December 2006. The ATP program reduces the explicit trading costs for participants' orders that meet two of three conditions. Price, time of order submission, and/or quantity decision must be made by an algorithm or computer program. The costs are adjusted based on total order flow for a given month, discounts begin at a minimum of $250 million euros of executed volume.

1.3. Frankfurter Wertpapier Boerse

The Frankfurter Wertpaper Boerse (FWB) also operated by DB is organized as a traditional floor exchange. Each stock traded on the floor is associated with a lead broker. Orders routed to the floor are routed directly into limit order books managed by the lead brokers. Floor trading begins at 9:00 and ends at 10:00 PM. Trading is organized in a single continuous trading phase which is always supported by a market-maker. As the relevance of floor trading has fallen, most trades are being executed on Xetra in the public order book, sometimes however large trades are executed on the FWB.

2. Hypothesis Development

The literature is unclear as to the effect of reducing latency. Theoretically the model that best fits the real world scenario is the Copeland and Galai (1983) model. They introduce a method to value firm commitments to trade at the posted bid and ask prices as options. Although not the main purpose of the model they introduce a method to determine a theoretical impact of latency on liquidity. In their options valuation framework they allow for an effect of the duration of quotes on the value of the option provided to the market. As with most options they increase in

[7] http://www.reuters.com/article/companyNews/idUSWEB425820080507

value with an increase in their duration. In this context lowering the minimum amount of time it takes to revise or delete an order is the same as reducing the duration of the free-trading option. This brief analysis is designed to present a *stylized* impact of latency reduction on liquidity, specifically the supply thereof.

We use the original settings but calculate the value for different latencies using a constant volatility. The idea is to get a feel for the sensitivity of the value of the free-trading option to a reduction in latency. The value of a European call option is calculated with an annualized volatility of 20% with latencies of 50 milliseconds and 10 milliseconds and an efficient price of 100 and quoted prices $1ct$ above and below. The value of the options is the cost a liquidity provider incurs on average when submitting a buy or sell order $1ct$ above or below the efficient price. The analysis shows that the option, or cost to the liquidity supplier, of an order that lives for 50 milliseconds is a little more than 1.5 basis points. An order that lives for 10 milliseconds implies a cost of less than 0.09 basis points for the free-trading option.

From this, we derive one of the central hypotheses of this paper:

$H1$**: A system-wide reduction in latency will have a positive effect on liquidity.**

Hypothesis 1 formalizes the relationship between latency and liquidity. The Copeland and Galai (1983) model provides a theoretical lower bound on the liquidity improvement. By halving latency the maximum liquidity increase is between 30 % and 50 % depending on the distance of the ask from the stock price. The central assumption in hypothesis 1 is that liquidity suppliers are in fact speed-sensitive. If liquidity suppliers do not have systems in place to exploit the new exchange systems they may in fact reduce liquidity to compensate for the increased risk of being picked-off. The mechanism by which liquidity increases is not specifically determined but is joint with the following hypothesis:

$H2$**: The informativeness of quoted prices increases with speed.**

As market speed increases, or latency decreases, the informativeness of prices will increase (cf. Foucault et al. (2003)). In this context prices include bid, ask, and mid prices. The above hypothesis is tested in a VAR framework.

3. Data and Methodology

We use data from the TAQTIC data service operated by SIRCA[8] on behalf of Reuters. Data on market capitalization is collected directly from the DB web site, company annual reports, and compared with other public data sources (Yahoo! Finance, Google Finance, and OnVista). The sample period covers the 40 trading day prior to and post April 23rd, 2007 - the event date. This leaves a sample period between 22 February 2007 and 19 June 2007. This period is selected because it allows an analysis of the short-term and long-term changes around the event.

3.1. Data Source

SIRCA provides trade, order, quote, and order book data for a wide number of stocks trading on exchanges worldwide. Specifically we retrieved trade, best bid, and best ask for the stocks in our original sample. Each trade and quote is time-stamped to the millisecond and accessible via Reuters Instrument Code (RIC). All prices are reported in Euros. We selected 110 stocks, as reported by TAQTIC, the data access tool provided by SIRCA, that made up the HDAX index as of 22nd February 2007.

The first and last five minutes of the trading day are removed to avoid biases associated with the information processing and inventory management processes at those times. The data spans trading between 9:05 am and 5:25 pm local time with some exceptions. Xetra features intra-day auctions and volatility interruptions which halt continuous trading. Our analysis is focused on continuous trading so all data recorded outside these hours and during volatility interruptions are deleted. The opening, closing, and intra-day and volatility interruptions are identified via Reuters' qualifying code attached to special types of trades and periods of time in a trading day. These qualifiers are used to filter the data.

To further validate our results, and compensate for potential time trends we match Xetra trades and quotes with Frankfurter Wertpapier Boerse (FWB) trades and quotes. Since the stocks traded on Xetra and FWB are both the same this should compensate for any time trends in the variables. The differenced ($Xetra_{QuotedSpread} - FWB_{QuotedSpread}$) data will be independent of these.

[8] Securities Industry Research Centre of Asia-Pacific, http://www.sirca.org.au/

[INSERT FIGURE 2 HERE]

See Figure 2 for a graph of the natural logarithm of volumes over the sample period. The figure shows that volume does not shift as result of the Xetra upgrade and confirms the robustness of both the Xetra and Xetra-FWB instruments to this effect.

3.2. Sample Selection

The sample contains the 110 stocks that make up the DB's HDAX. The HDAX is a combination of three main indexes: the DAX, TecDAX, and MDAX[9]. They are the most actively traded and highest quality publicly traded German companies and present a broad cross-section of industries. The DAX contains the 30 largest and highest quality German blue-chip stocks determined by market capitalization, free-float, transparency regulations, and industry. The MDAX is made up of the next largest 50 companies, followed by the 30 technology stocks in the TecDAX. We take the index composition as of 22 February 2007. This is well before the Xetra 8.0 upgrade. All members of the HDAX meet certain minimum admission requirements, they must publish quarterly reports, adhere to IFRS or US-GAAP accounting standards, publish a financial calendar, hold one analyst conference per year, and provide ad-hoc disclosure information in German and English.

[INSERT TABLE 1 HERE]

Stocks that do not meet certain criteria are removed. The removal avoids effects related to size, trading frequency, and price and not to continuous trading pre and post Xetra upgrade. We modify criteria from Hendershott and Moulton (2007) to prepare the sample. A stock must trade above 1 euro and below 500 euros during the entire sample period. A stock has to be traded continuously throughout the study period and trade at least twice a day. Stocks that split or delist during the observation period are removed. Stocks that are dropped from the HDAX during the sample period are also removed. The final sample consists of 101 stocks. Table 1 reports the average price, trades, daily turnover, and turnover per trade for all stocks in the final sample. Pre and post Xetra 8 variable values are reported in Table 2.

[9]See the Deutsche Boerse website for a full description of the indexes

3.3. *Liquidity and Information Measures*

3.3.1. *Calculation*

Tick-by-tick observations are aggregated to a daily frequency for the regression analyses to capture the intra-day dynamics of each variable but avoid some of the noise associated with a higher sampling frequency (trade-by-trade or quote-by-quote). Descriptive statistics as reported in Table 2 are calculated on tick-by-tick data.

We use the now common Lee and Ready (1991) algorithm with contemporaneous quotes as proposed by Bessembinder (2003) to sign trades. Bessembinder (2003) compares different heuristics to infer trade direction with proprietary data featuring the trade direction and finds that a comparison of the trade with the contemporaneous quote using Lee and Ready's heuristic provides the best results. Given the current information technology and the period in which the data was collected using contemporaneous quotes should add no additional bias to our results.

3.3.2. *Spread Decomposition*

Several proxies for liquidity are presented. We calculate round-trip (full) spreads rather than half-spreads. Quoted spreads are the simplest and most common measure of trading costs and can easily be calculated using trade and order data. All calculations presented below are spreads relative to stock price and are reported in basis points (bps). In order to avoid distorted results from ambiguity in the raw data, intra-day observations featuring a quoted spread larger than 10%, an effective spread larger than 10%, or a realized spread larger than 10% or smaller than -10% are removed. The quoted spread on Xetra is created through public limit orders submitted by various participants. Let $\text{Ask}_{i,t}$ be the ask price for a stock i at time t and $\text{Bid}_{i,t}$ the respective bid price. $\text{Mid}_{i,t}$ denotes the mid quote then the quoted spread is calculated as follows:

$$\text{Quoted Spread}_{i,t} = (\text{Ask}_{i,t} - \text{Bid}_{i,t})/\text{Mid}_{i,t}$$

The effective spread is the spread paid when an incoming market orders trades against a limit order. The effective spread also captures institutional features of a market such as hidden liquidity and market depth. Let $\text{Price}_{i,t}$ be the execution price then the effective spread is defined as

$$\text{Effective Spread}_{i,t} = 2 * D_{i,t} * ((\text{Price}_{i,t} - \text{Mid}_{i,t})/\text{Mid}_{i,t})$$

$D_{i,t}$ denotes the trade direction, -1 for a market sell and $+1$ for a market buy order. The realized spread measures liquidity supplier revenues independent of the adverse selection costs imposed on the uninformed by the informed (Bessembinder and Kaufman, 1997). The realized spread is calculated with the mid-quote five minutes after the trade ($x = 5$) as follows.

$$\text{Realized Spread}_{i,t} = 2 * D_{i,t} * ((\text{Price}_{i,t} - \text{Mid}_{i,t+x})/\text{Mid}_{i,t})$$

Price impact is an approximate measure of the adverse selection component of the effective spread. The price impact is the effective spread minus the realized spread and measures the information content of a trade. It approximates the permanent impact of a trade under the assumption that information impacts are permanent and realized at the 5-minute mark whereas other effects such as inventory and explicit trading costs are transitory. Following a trade, liquidity suppliers adjust their beliefs about the fundamental value of an asset depending on the information content of a trade (cf. Glosten and Milgrom (1985)). The simple price impact of a trade is calculated as follows:

$$\text{Price Impact}_{i,t} = 2 * D_{i,t} * ((\text{Mid}_{i,t+x} - \text{Mid}_{i,t})/\text{Mid}_{i,t})$$

The price impact provides an indication of the information content of a trade. We apply more robust informationprice discovery measures, not dependent on the spread decomposition, in the following. In order to avoid distorted results due to ambiguity in the raw data, intra-day observations featuring a simple price impact larger than 10% or smaller than -10% are removed from the data. Ambiguity can arise due to system problems and unrecorded microstructure features.

3.3.3. Price Discovery

To further study and confirm the hypothesis that less trade-correlated information is present post-upgrade we perform the analysis laid out in Hasbrouck (1991a) and Hasbrouck (1991b). The results of the VAR analysis are the average cumulative impulse response function (CIRF) over 10 trades and the aggregate values per stock and day.

The permanent price impact of a trade (Hasbrouck, 1991a) is commonly used in price discovery research. We use the standard settings which include a forecast horizon of 10 trades. We test forecast validity above 10 trades and found no support for effects at lower frequencies. Let x_{t-i} be the trade direction, r_{t-i} denotes the quote midpoint changes. The full model is as follows:

$$r_t = \gamma_{0,r} + \sum_{i=0}^{10} \alpha_{t-i} x_{t-i} + \sum_{t=1}^{10} \beta_{t-i} r_{t-i} + u^r$$
$$x_t = \gamma_{0,x} + \sum_{i=1}^{10} \delta_{t-i} x_{t-i} + \sum_{t=1}^{10} \eta_{t-i} r_{t-i} + u^x$$

The estimation is restarted for each trading day and stock in the sample. The above VAR is inverted to get the vector moving average representation (VMA).

$$\begin{pmatrix} r_t \\ x_t \end{pmatrix} = \begin{pmatrix} a(L)\ b(L) \\ d(L)\ e(L)\ ) \end{pmatrix} \begin{pmatrix} u^r \\ u^x \end{pmatrix},$$

Following Hasbrouck (1991b) the sum of $\sum_{t=0}^{10} a$, where L are polynomial lag operators, is used to attain the cumulative impulse response function (CIRF). The CIRF is the permanent price impact of a trade and is generally interpreted as the private information content of a trade. Trades may contain information at lower frequencies than measured. This measure however has been used in a number of other studies with the same interpretation (Barclay and Hendershott (2003), Madhavan (2000)).

Using the VMA representation from above information can be decomposed into trade-correlated and uncorrelated portions (Hasbrouck, 1991b). The variance decomposition is as follows:

$$\sigma_w^2 = \left(\sum_{i=0}^{10} a_i\right)^2 \sigma_{u^r}^2 + \left(\sum_{i=0}^{10} b_i\right)^2 \sigma_{u^x}^2$$

The information content of quotes is the second term and the trade correlated portion the first term. All lags are summed to get the total contribution to price discovery of both portions. These results are reported below in basis points for the CIRF and in percent for the information content of quotes. By analyzing both of these measures, we study how information is impounded into prices.

3.4. Descriptives

The time-series means of each variable are calculated per stock. In Table 2 the cross-sectional means of the variables are reported. Trade prices on Xetra range from 3.16 euros to 351.70 euros with a sample mean of 70.98. The average stock trades 1,492 times in a day which translates roughly into three times a minute. Table 1 shows an interesting phenomenon, generally in order-driven markets without designated market makers the effective spread should be greater than the quoted spread. These results show that the quoted spread is on average larger than the effective spread. This indicates that traders monitor the market and trade when spreads are lower than average. This also indicates that a great deal of order splitting occurs and that it is worth-while to do so.

[INSERT TABLE 2 HERE]

Table 1 reports the descriptive statistics for the sample. Panel A reports the descriptive statistics for Xetra and Panel B for FWB. The following analysis of liquidity focuses on the differenced[10] results to mitigate any time effects in the sample. The table shows clearly that quoted spreads are smaller on Xetra, but that effective spreads are somewhat greater on Xetra. This coupled with a smaller price impact leaves a large realized spread for FWB trades. These results confirm other empirical studies that find that repeated interactions, by humans, lead to lower executions costs due to an ability to avoid informed trades (Hendershott and Moulton, 2007). Other clear differences are the number of trades per day and the average turnover per trade, both considerably higher on Xetra.

[10] i.e $QuotedSpread_{reported} = QuotedSpread_{Xetra} - QuotedSpread_{FWB}$

Table 2 reports the mean and standard deviations of quoted, effective, and realized spread, the price impact, and summed daily impulse response function (trade innovation). Table 2 shows a decrease in measures of trading costs (quoted and effective spreads) an increase in liquidity supplier revenues, and a corresponding decrease in price impact or trade correlated information post-upgrade. Most interestingly the robust information results show a decline from 4.42 bps to 1.19 bps. The results are consistent across all MCap categories.

4. Liquidity Analysis

Each equation is estimated once for each setting; once on the Xetra values, once on FWB values, and once on the difference. The focus of the analysis is on the differenced results but the finding hold for both settings (Xetra and differenced). To test the hypothesis that reducing latency has an effect on liquidity, regressions of the following form are estimated:

$$LM_{i,t} = \alpha_i + \delta Xetra8_{i,t} + \beta VDAX_t + \epsilon_{i,t} \tag{1}$$

where Liquidity Measure (LM) is the quoted spread, effective spread, realized spread, and price impact on date t for stock i. α_i are fixed cross-sectional effects for each individual stock. $Xetra8_{i,t}$ is a dummy variable that takes the value 0 before April 23rd 2007 and 1 otherwise. A daily volatility measure $VDAX_t$ is included as in Hendershott and Moulton (2007) to control for market wide volatility changes and the effects thereof on market-wide liquidity. For the $VDAX$ the daily opening value of DB's 3-Month $VDAX - New$ is used for each date in the sample period. Poolability tests show that the data are not poolable. A fixed-effects model is used that accounts for cross-sectional differences in stocks. The panel regressions are estimated with robust standard errors for within-groups estimators (Arellano, 1987) which are essentially White's robust standard errors (White, 1980) adjusted for panel data.

Table 3 shows the results of the panel regression from the fixed-effects model for the full sample and by market capitalization quartiles. We also estimate the above fixed-effects model for the Frankfurt floor exchange (FWB) and for the difference between Xetra and FWB. All results are reported in table 3.

[INSERT TABLE 3 HERE]

4.1. *Quoted Spread*

Table 2 presents the results of the quoted spread estimation for the full sample, pre and post Xetra 8.0 upgrade and separated in market capitalization quartiles. The results show that the quoted spread for the full sample decreases from 12.45 to 12.06 basis points (bps). Note that in the spread results quoted spreads are time-weighted and effective spreads are transaction weighted. This leads to the counterintuitive result that quoted spreads are greater than effective spreads. This doesn't impact the economic interpretation of the differences.

The results of the panel regression estimates in table 3 show that the Xetra upgrade has a significant negative effect on quoted spreads, when comparing with the FWB the results on quoted spread are even greater. The results across market capitalization quartiles are varied and generally show an increase in the quoted spread. The results for the difference are more consistent and larger than for Xetra alone. Only MCap quartile 4 is not statistically significant. Since quoted spreads only measure the trading costs for the smallest of trade sizes, a more accurate measure of execution costs is studied in the section below.

4.2. *Effective Spread*

The effective spread is the actual spread paid by a liquidity demander in a limit market. Table 2 reports effective spreads for the full sample and individual market capitalizations both pre and post Xetra 8.0 upgrade. In contrast to previous studies and the theoretical literature the findings show that effective spreads decrease from an average of 8.15 bps to an average of 6.71 bps. The decrease represents a 17% decrease in effective execution costs. The decrease in effective spreads can be found across market capitalization quartiles. The greatest decrease in effective spread is for small stocks with a decrease from 25.35 to 19.04 bps.

Table 3 shows a significant decrease -1.91 in effective spread after the release of the new Xetra system. The results hold when compared with the FWB results (-1.41 bps). In no case is the decrease for the FWB significant and only in the smallest MCap are the differenced results not statistically significant at the 5 % level. These results show clearly that transaction costs decline post-upgrade. Effective spreads are the most accurate measures of execution costs but are made up of at least two components, the realized spread and the price impact as an approximation of information. The interplay of these two components are important in understanding the drivers of liquidity changes.

4.3. *Realized Spread*

The realized spread represents the part of the effective spread that a liquidity supplier keeps. Essentially the realized spread is a liquidity supplier's revenue and is important to measure in this context. A reduction in the effective spread could mean that a trade is less information driven and hence the adverse selection costs imposed on a liquidity supplier are reduced. It could also mean that liquidity suppliers require less compensation for the services they provide for instance due to decreased fixed transaction costs. Surprisingly Table 2 shows that the realized spread increases by roughly 3 basis points. The realized spread *increases* across all market capitalization quartiles.

The results also hold in the panel estimation with an increase in realized spread of 5.39 bps after controlling for stock level variables, and volatility. The increase is consistent across MCap quartiles and when using the differenced variables. The realized spread also increases consistently across MCap quartile. These results are surprising in that it implies that execution costs should rise after the introduction of Xetra 8, when in fact they decline. This result is however consistent with the commonly held hypothesis that increased execution speed increases execution costs. The mechanism for this in certain scenarios could be that liquidity suppliers require more compensation for supplying liquidity if the chances of being exploited are greater (free-trading option) which may be the case when execution speed increases and arbitrageurs are employing algorithmic trading technologies. Regardless of the explanation it seems that the driver of the reduced execution costs is the interdependency between liquidity supplier competition and the information content (price impact) of trades that changes after the introduction of Xetra 8.0.

4.4. *Price Impact*

Table 2 shows that the price impact of trades decreases significantly following Xetra 8. On average, the price impact per trade decreases by 4 bps. The break down of price impact into market capitalization quartiles also shows a decrease in the price impact per trade, and an increase in the magnitude of the decrease. The results remain the same in a panel regression; the price impact for the smallest stocks decreases by 11.39 basis points. The panel results show that price impact is greatly affected by the upgrade, with an overall decrease of 7.29 bps and surprisingly a 2.72 bps decrease in market capitalization Q1 stocks. The FWB and differenced results show clearly that the information content of trade on the FWB did not change and hence is not driven by a time trend. In fact the Xetra-FWB results show a significant decrease of 7.17 bps.

4.5. Liquidity and Trade Size

The liquidity increase could in fact be driven by an increase in smaller trades or in liquidity for certain trades sizes. The results of the liquidity estimation regression in are reported in Table 4

[INSERT TABLE 4 HERE]

The largest liquidity increase is in large small cap trades. A dramatic fall of 16.82 bps for trade greater than 100,000 euros is found. The results are consistent with the previous analysis and confirm that liquidity increases across trade and MCap categories. For MCap groups Q2 to Q4 the Xetra-FWB estimations are not reported due to data limitations (not enough transactions).

To ensure that the results are not in fact driven by a shift in trade size a panel regression is performed on the average trade size. We report per trade turnover in Euros in Table 5.

[INSERT TABLE 5 HERE]

We find no statistically significant changes in average per trade turnover. In Q2 a small increase is found. In all other quartiles there are no statistically significant changes to turnover per trade. There is no change in the volume ratio between Xetra and FWB. It can be safely stated that the results are robust to these factors.

5. Information Content of Trades and Quotes

Lacking sharp theoretical predictions the statistical null hypothesis for the trade innovation measure is simply that there will be no difference between trade correlated and uncorrelated information post-upgrade. To test for differences in the amount of trade-correlated and uncorrelated information pre and post-upgrade the same regression as above[11] is estimated for the price impact (for Xetra) measured by the CIRF and the variance decomposition. The results of the estimation are presented in table 6.

[INSERT TABLE 6 HERE]

Table 6 reports a strongly negative effect of the Xetra upgrade on the CIRF. The Xetra results are highly significant (-21.04) with a coefficient of -3.16 (Table 6). These results confirm the results of the price impact analysis. Unfortunately due to data restrictions the CIRF cannot be calculated for the FWB. The simple price impact remains unchanged for the FWB pre and post-upgrade. The effect of the Xetra upgrade on the CIRF increases in absolute terms across MCap quartiles and indicates that liquidity suppliers are able to avoid informed trades post-upgrade in relatively small stocks that generally have larger adverse selection costs. To reconcile these results with the price impact note that price impact is calculated at the 5-minute mark and the CIRF is comparable in magnitude half of the effective spread.

Figure 3 presents the CIRF pre and post-upgrade for the entire sample. We include ninety-five percent confidence intervals that clearly show trade related information post-upgrade (in Red) is strictly and always smaller than pre-upgrade.

[INSERT FIGURE 3 HERE]

The forecast horizon in trades is reported on the X-axis (0 - 10). At 10 events the trade impact levels off, confirming the lag length selected. On the Y-axis the CIRF is represented in basis points. Table 6 reports a dramatic increase in the amount of information being impounded into prices via quotes. Quote-based information increases by 47.89 % and is highly significant at the 1% level. In Figure 4 the percentage of quote-based information across the sample period with 95% confidence intervals is graphed, similar to the above figures.

[INSERT FIGURE 4 HERE]

[11] $LM_{i,t} = \alpha_i + \delta Xetra8_{i,t} + \beta VDAX_t + \epsilon_{i,t}$

Visual inspection shows clearly that the ratio changes dramatically on the event day. The statistical test in Table 6 simply confirms what visual inspection already reveals. The amount of information being impounded into quotes post-upgrade is much higher.

We include a test of the equality of variance to better understand the information impounding process. In Panel A of Table 8, the mean and standard deviation of the quote-based information pre and post-upgrade is reported. The standard deviation falls by roughly half. An equality of means test is estimated and the results are reported in Panel B.

[INSERT TABLE 8 HERE]

Using the modified levene test the null hypothesis of equal variances pre and post-upgrade is rejected. This could be driven by AT but we cannot confirm. Imagine a situation where humans are doing most of the public information (trades, quotes, orders) processing. The level of attention of a human is variable throughout a day and bounded. On days where a great deal of information reaches the market humans may be unable to impound it all into quotes, leading to large swings in the ratio of public to private information. On days where there is a lot of activity the proportion of information processed will be lower than on days with less information. Limited attention and variability are not attributes one would typically associate with an algorithm. AT are more likely to be stable and should provide less variable information processing.

5.1. Information and Trade Size?

Technical innovations have been shown to have an effect on the amount, and nature of trade in electronic stock markets (Campbell et al. (1997) and Stoll (2006)). To analyze the effect of trade size we sort the sample into three trade size categories smaller than $25,000$ euros between $25,000$ and $100,000$ and greater than $100,000$. Table 7 reports the results of a regression of the simple price impact per trade on the Xetra 8 variable. We can only robustly compute the simple price impact for different trade sizes not the VAR based price impact. The trend towards an increase in the absolute decrease of the price impact remains across MCaps and trade size. The information content of large trades decreases significantly and indicates that post upgrade latency sensitive traders seem to be able to manage their exposure to informed trades deep into the book and not only at the inside quote. Due to data restrictions it was impossible to calculate values for the difference between Xetra and the FWB for large trades.

[INSERT TABLE 7 HERE]

Although there is an increase in the trade size, reported in Table 5, the post upgrade effect is largely insignificant statistically. It is interesting that as quotes become more informative (quote-based price discovery increases) participants are willing to trade in larger blocks. Perhaps pre-upgrade there were more smaller limit orders that were being picked-off, which would explain both the increase in trade size and the greater CIRF pre upgrade. What is clear is that a change in trade size is not driving the results.

6. Conclusion

In this paper we study the effect of reducing latency on market-wide price discovery and market liquidity. Common measures of information in a market microstructure sense and measures of the actually transaction costs paid by market participants are used. The results show that quote-based information increases in concert with liquidity. The Xetra 8.0 upgrade is ideal to test the hypotheses that latency reductions impact information and liquidity. We show that latency correlates positively with liquidity. This is in contrast to the results found in Hendershott and Moulton (2007) and theorized much early by Demsetz (1968). Although the results of this paper differ from the above studies this in no way negates or questions the results therein.

In the preceding we showed how reducing system latency can contribute to market liquidity. The results suggest that by reducing latency liquidity suppliers are either impounding more information into prices post upgrade or more price relevant information is in the public domain post-upgrade. By impounding more information into prices they are less-likely to have their stale quotes picked-off by more informed traders. This translates into a reduction in ex-post liquidity i.e. effective spreads.

The increase in liquidity is due to a dramatic drop in the information content of trades and a somewhat less dramatic increase in realized spread. Reducing latency seems to be a win-win situation for regulators, market participants, and exchange operators. Most important to regulators, posted prices are more efficient post upgrade and better reflect public information. As adverse selection costs fall market participants are also more likely to initiate trade. Also of importance to market participants, they can execute their orders at a lower cost. Exchange operators are content because they attract higher volumes and the fees related to those. One could imagine as the level of competition between algorithmic liquidity suppliers reaches that of humans prior to the Xetra 8.0 upgrade, liquidity will continue to increase. A warning is however in order. Even simple systems changes may have unexpected effects on markets. The mechanism by which liquidity improved was little understood before the upgrade and likely not the expected one. The effect of reducing latency could have been to reduce liquidity as seen by NYSE's Hybrid upgrade.

Future work in the area should focus on the theoretical underpinning of these results. Why is it that the NYSE Hybrid upgrade caused a decrease in liquidity and a similar execution speed increase at the DB had the opposite effect? Also further studies into similar recent upgrades (TSX Quantum, LSE Tradelect, Euronext Universal Trading Platform) could shed light on some of these differences. Finally detailed algorithmic trading data sets could also help to alleviate some of the suspicions with regards to algorithmic trading.

References

Arellano, J., 1987. Computing robust standard errors for within-groups estimators. Oxford Bulletin of Economics and Statistics 49 (4), 431–434.

Bacidore, J., Ross, K., Sofianos, G., 2003. Quantifying market order execution quality at the new york stock exchange. Journal of Financial Markets 6 (3), 281–307.

Barclay, M., Hendershott, T., 2003. Price discovery and trading after hours. Review of Financial Studies 16 (4), 1041–1073.

Battalio, R., Hatch, B., Jennings, R., 2003. All else equal?: a multidimensional analysis of retail, market order execution quality. Journal of Financial Markets 6 (2), 143–162.

Bennett, P., Wei, L., 2006. Market structure, fragmentation, and market quality. Journal of Financial Markets 9 (1), 49–78.

Bessembinder, H., 2003. Issues in assessing trade execution costs. Journal of Financial Markets 6 (3), 233–257.

Bessembinder, H., Kaufman, H., 1997. A cross-exchange comparison of execution costs and information flow for nyse-listed stocks. Journal of Financial Economics 46 (3), 293–319.

Blume, M., 2001. The structure of the us equity markets. Brookings-Wharton Papers on Financial Services.

Boehmer, E., 2005. Dimensions of execution quality: Recent evidence for us equity markets. Journal of Financial Economics 78 (3), 553–582.

Brown, M., Forsythe, A., 1974. Robust tests for the equality of variances. Journal of the American Statistical Association 69 (346), 364–367.

Campbell, J., Lo, A., MacKinlay, C., 1997. The Econometrics of Financial Markets. Princeton University Press, Princeton, NJ.

Cespa, G., Foucault, T., 2008. Insiders-outsiders, transparency, and the value of the ticker, Working Paper.

Chaboud, A., Chiquoine, B., Hjalmarsson, E., Vega, C., 2009. Rise of the machines: Algorithmic trading in the foreign exchange market, Working Paper.

Chordia, T., Roll, R., Subrahmanyam, A., December 2007. Why has trading volume increased?, Working Paper.

Copeland, T., Galai, D., 1983. Information effects on the bid-ask spread. Journal of Finance 38 (5), 1457–1469.

Demsetz, H., 1968. The cost of transacting. Quarterly Journal of Economics 82 (1), 33–53.

Easely, D., Hendershott, T., Ramadorai, T., 2007. The price of latency, Working Paper.

Foucault, T., Kadan, O., Kandel, E., 2005. Limit order book as a market for liquidity. Review of Financial Studies 18 (4), 1171–1217.

Foucault, T., Roëll, A., Sandas, P., 2003. Market making with costly monitoring: An analysis of the soes controversy. Review of Financial Studies 16 (2), 345–384.

Glosten, L., Milgrom, P., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics 14, 71–100.

Hasbrouck, J., 1991a. Measuring the information content of stock trades. Journal of Finance 46 (1), 179–207.

Hasbrouck, J., 1991b. The summary informativeness of stock trades: an econometric analysis. Review of Financial Studies 4 (3), 571–595.

Hendershott, T., Jones, C., Menkveld, A., 2008. Does algorithmic trading improve liquidity?, Working Paper.

Hendershott, T., Moulton, P., 2007. The shrinking new york stock exchange floor and the hybrid market, Working Paper, Haas School of Business and Fordham Graduate School of Business.

Hendershott, T., Riordan, R., 2009. Algorithmic trading and information, Working Paper presented at First FBF/IDEI-R Conference on Investment Banking and Financial Markets, Toulouse.

Huang, R., 2002. The quality of ecn and nasdaq market maker quotes. Journal of Finance 57 (3), 1285–1319.

Jain, P., 2005. Financial market design and the equity premium: Electronic versus floor trading. Journal of Finance 60 (6), 2955–2985.

Klar, J., van den Bongard, I., 2008. Determinants of the bid-ask spread and the role of designated sponsors: Evidence for xetra, Working Paper.

Lee, C., Ready, M., 1991. Inferring trade direction from intraday data. Journal of Finance 46 (2), 733–746.

Madhavan, A., 2000. Market microstructure: A survey. Journal of Financial Markets 3 (3), 205–258.

Menkveld, A., Wang, T., 2008. How do designated market makers create value for small-caps?, Working Paper.

Parlour, C., 1998. Price dynamics in limit order markets. Review of Financial Studies 11 (4), 789–816.

Stoll, H., 2006. Electronic trading in stock markets. Journal of Economic Perspectives 20 (1).

Venkataraman, K., 2001. Automated versus floor trading: An analysis of execution costs on the paris and new york exchanges. Journal of Finance 56 (4), 1445–1485.

White, H., 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. Econometrica 48 (4), 817–838.

White, L., Frame, W., 2004. Empirical studies of financial innovation: Mostly talk and not much action? Journal of Economic Literature 42 (1), 116–44.

Table 1: **Descriptive Statistics for Xetra and the Frankfurt floor (FWB)**: The sample consists of stocks listed in Deutsche Boerse's HDAX segment. The observation period comprises of 40 trading days before and after the introduction of Xetra 8 on 23 April 2007. Table 1 Panel A reports descriptive statistics for Xetra and Panel B reports statistics for the Frankfurt floor. Average measures are calculated on tick data. Daily turnover per instrument and daily trade count per instrument are calculated on a daily per instrument basis. Market capitalization is calculated as the product of shares outstanding and the average price. All spread measures are reported as relative measures in basis points. All monetary measures are reported in Euros.

Panel A: Descriptive Statistics Xetra

	Mean	Std. Dev.	Min.	Max.
Shares (1000)	234,690	490,060	6,350	436,130
Market Cap (MEUR)	10,584	17,277	356	80,236
Price (per Trade)	70.98	47.34	3.16	351.70
Quoted Spread	12.36	15.00	> 0.00	820.94
Effective Spread	7.50	10.14	> 0.00	581.24
Realized Spread	1.92	444.67	3992.34	997.35
Price Impact	5.58	44.49	-993.50	999.64
Turnover (1000 EUR)	71,816	153,330	48	2,718,500
Trade Count	1,492	1,868	9	20,467
Per Trade Turnover	48,150	94,650	3	22,055,446

Panel B: Descriptive Statistics FWB

	Mean	Std. Dev.	Min.	Max.
Shares (1000)	234,690	490,060	6,350	436,130
Market Cap (MEUR)	10,584	17,277	356	80,236
Price (per Trade)	58.13	43.72	3.17	351.00
Quoted Spread	15.62	18.79	> 0.00	982.57
Effective Spread	6.81	12.24	> 0.00	841.68
Realized Spread	4.82	48.09	-985.88	1000.00
Price Impact	1.99	47.53	-924.86	994.46
Turnover (1000 EUR)	1,059	2,335	0.1	52,875
Trade Count	90	145	1	2,020
Per Trade Turnover	11,829	45,192	4	10,934,900

Table 2: **Liquidity and Information Pre and Post Xetra 8 for Xetra**: The sample consists of stocks listed in Deutsche Boerse's HDAX segment. The observation period comprises of 40 trading days before and after the introduction of Xetra 8 on 23 April 2007. Table 2 reports descriptive statistics for quoted spreads, effective spreads, realized spreads, price impacts, and permanent impacts of trade innovations. Quoted spread, effective spread, realized spread, and price impact are calculated on tick data. Permanent price impact statistics are calculated using daily data. Results are reported for the entire sample and individually by a stock's market capitalization. Stocks are divided into four groups depending on their market capitalization. Market capitalization is calculated as the product of shares outstanding and the average price over the observation period for a single stock. All results are reported in basis points.

Liquidity and Information Pre and Post Xetra 8 Introduction										
	Quoted Spread		Effective Spread		Realized Spread		Price Impact		Permanent Price Impact	
	mean	Std.Dev.	mean	Std. Dev.	mean	Std. Dev.	mean	Std. Dev.	mean (daily)	Std. Dev. (daily)
Sample										
full	12.21	15.01	7.44	10.19	1.91	44.85	5.53	44.67	2.81	2.65
pre	12.45	15.98	8.15	11.40	0.49	48.17	7.65	48.06	4.42	2.73
post	12.06	14.39	6.71	8.71	3.39	41.06	3.32	40.71	1.19	1.19
MCAP Q1										
full	6.54	6.87	4.47	4.42	1.32	35.16	3.14	35.07	1.48	1.09
pre	6.23	6.93	4.77	4.68	0.58	36.21	4.20	36.13	2.29	0.94
post	6.71	6.83	4.15	4.11	2.09	34.03	2.06	33.91	0.66	0.38
MCAP Q2										
full	14.02	13.02	9.22	9.71	1.71	51.23	7.51	50.90	2.56	2.04
pre	13.40	13.66	9.77	10.60	-0.61	54.14	10.39	53.75	4.04	1.85
post	14.45	12.55	8.65	8.66	4.13	47.89	4.52	47.56	1.08	0.69
MCAP Q3										
full	22.70	19.22	14.27	14.83	2.80	63.08	11.47	62.52	3.06	2.31
pre	21.68	19.16	15.61	16.28	-0.40	70.12	16.01	69.46	4.86	1.89
post	23.43	19.23	12.82	12.92	6.25	54.27	6.57	53.62	1.27	0.82
MCAP Q4										
full	33.56	25.78	22.51	21.99	7.95	70.15	14.56	68.98	4.17	3.73
pre	33.69	27.17	25.35	24.39	5.29	79.29	20.05	77.97	6.59	3.52
post	33.45	24.59	19.04	18.03	11.19	56.84	7.85	55.33	1.78	1.94

Table 3: **Results Panel Regressions**: The sample consists of stocks listed in Deutsche Boerse's HDAX segment. The observation period comprises of 40 trading days before and after the introduction of Xetra 8 on 23 April 2007. Table 3 Panel A reports panel regression results for quoted spread, effective spread, and realized spread. Panel B reports regression results for price impact, permanent impact of trade innovation, and the variance decomposition in percent. The panel regression is performed on daily measures individually for each stock using the following: $LM_{i,t} = \alpha_i + \delta Xetra8_{i,t} + \beta VDAX_t + \epsilon_{i,t}$. Results are reported for the entire sample and individually by a stock's market capitalization. Stocks are divided into four groups depending on their market capitalization. Market capitalization is calculated as the product of shares outstanding and the average price. All results are reported in basis points except for variance decomposition which is in percent. *Xetra* denotes results for the Xetra, FWB denotes results for the Frankfurt floor, and Diff is Xetra-FWB. *** indicates significance at the 1% level, ** indicates significance at the 5% level, and * indicates significance at the 10% level. Robust t-statistics are reported in parentheses.

Regressions on Spread Measures									
	Quoted Spread			Effective Spread			Realized Spread		
	Xetra	FWB	Diff	Xetra	FWB	Diff	Xetra	FWB	Diff
Sample									
Xetra8	1.12***	-1.16	2.28***	-1.91***	-0.50	-1.41***	5.39***	-0.37	5.76***
t-Value	*(3.12)*	*(-1.46)*	*(3.67)*	*(-7.43)*	*(-1.06)*	*(-3.37)*	*(16.54)*	*(-0.75)*	*(10.70)*
MCAP Q1									
Xetra8	1.06***	-0.04	1.11***	-0.53***	-0.04	-0.49***	2.20***	-0.38	2.56***
t-Value	*(5.07)*	*(-0.25)*	*(8.20)*	*(-3.75)*	*(-0.39)*	*(-2.84)*	*(6.50)*	*(-1.14)*	*(4.84)*
MCAP Q2									
Xetra8	0.43	-1.34	1.77***	-1.38***	-0.32	-1.04***	5.28***	0.50	4.77***
t-Value	*(0.46)*	*(-1.27)*	*(6.42)*	*(-2.78)*	*(-0.68)*	*(-2.77)*	*(10.24)*	*(0.89)*	*(6.95)*
MCAP Q3									
Xetra8	1.83***	-0.72	2.55***	-1.73***	-0.62	-1.12***	6.85***	-0.79	7.69***
t-Value	*(3.48)*	*(-1.18)*	*(5.86)*	*(-4.15)*	*(-1.3)*	*(-2.59)*	*(15.43)*	*(-1.05)*	*(9.44)*
MCAP Q4									
Xetra8	1.16	-2.56	3.72	-4.05***	-1.05	-3.01*	7.34***	-0.82	8.17***
t-Value	*(1.26)*	*(-0.88)*	*(1.54)*	*(-6.75)*	*(-0.61)*	*(-1.81)*	*(10.56)*	*(-0.53)*	*(5.28)*

Table 4: **Results by Trade Size**: The sample consists of stocks listed in Deutsche Boerse's HDAX segment. The observation period comprises of 40 trading days before and after the introduction of Xetra 8 on 23 April 2007. Table 4 Panel A reports panel regression results for effective spreads by trade size. Panel B reports results for realized spreads by trade size. The regression is estimated on daily measures individually for each stock using the following: $LM_{i,t} = \alpha_i + \delta Xetra8_{i,t} + \beta VDAX_t + \epsilon_{i,t}$. Results are reported for the entire sample and individually by a stock's market capitalization. Different trade sizes are below 25,000 EUR, between 25,000 EUR and 100,000 EUR, and above 100,000 EUR. Stocks are divided into four groups depending on their market capitalization. Market capitalization is calculated as the product of shares outstanding and the average price. *Xetra* denotes results for Xetra, FWB are results for the Frankfurt floor, and Diff is $Xetra - FWB$. All results are reported in basis points. *** indicates significance at the 1% level, ** indicates significance at the 5% level, and * indicates significance at the 10% level. Robust t-statistics are reported in parentheses.

Panel A: Regressions on Effective Spreads by Trade Size

	< 25 kEUR			≥ 25 kEUR and ≤ 100 kEUR			> 100 kEUR		
	Xetra	FWB	Diff	Xetra	FWB	Diff	Xetra	FWB	Diff
Sample									
Xetra8	-1.58***	-0.54	-1.04**	-3.84***	-0.79	-2.75***	-7.19***	.	.
t-Value	*(-6.35)*	*(-1.12)*	*(-2.47)*	*(-3.97)*	*(-0.89)*	*(-5.01)*	*(-4.78)*		
MCAP Q1									
Xetra8	-0.33**	-0.05	-0.28*	-0.60***	0.02	-0.57**	-1.48***	-0.45	-0.86***
t-Value	*(-2.30)*	*(-0.46)*	*(-1.65)*	*(-4.85)*	*(0.04)*	*(-2.30)*	*(-3.95)*	*(-1.01)*	*(-2.78)*
MCAP Q2									
Xetra8	-1.08**	-0.40	-0.63	-2.21***	-0.56	-1.38**	-5.17***	.	.
t-Value	*(-2.21)*	*(-0.85)*	*(-1.64)*	*(-3.22)*	*(-0.70)*	*(-2.29)*	*(-5.06)*		
MCAP Q3									
Xetra8	-1.36***	-0.71	-0.67	-3.38***	-1.20	-2.78***	-10.86***	.	.
t-Value	*(-3.42)*	*(-1.42)*	*(-1.47)*	*(-5.86)*	*(-0.97)*	*(-2.70)*	*(-4.47)*		
MCAP Q4									
Xetra8	-3.61***	-1.00	-2.61	-9.48***	-2.03	-8.42***	-16.82***	.	.
t-Value	*(-6.05)*	*(-0.58)*	*(-1.58)*	*(-4.12)*	*(-1.05)*	*(-3.54)*	*(-4.63)*		

Panel B: Regressions on Realized Spreads by Trade Size

	< 25 kEUR			≥ 25 kEUR and ≤ 100 kEUR			> 100 kEUR		
	Xetra	FWB	Diff	Xetra	FWB	Diff	Xetra	FWB	Diff
Sample									
Xetra8	4.89***	-0.48	5.38***	8.51***	0.39	6.89***	9.66***	.	.
t-Value	*(15.90)*	*(-0.96)*	*(9.87)*	*(11.49)*	*(0.30)*	*(5.77)*	*(7.93)*		
MCAP Q1									
Xetra8	2.05***	-0.31	2.33***	2.51***	-1.413**	3.62***	3.69***	2.00	0.87
t-Value	*(7.01)*	*(-1.01)*	*(5.16)*	*(4.48)*	*(-2.28)*	*(5.08)*	*(5.61)*	*(1.05)*	*(0.49)*
MCAP Q2									
Xetra8	4.74***	0.22	4.51***	6.97***	1.16	5.46***	11.32***	.	.
t-Value	*(9.55)*	*(0.55)*	*(7.37)*	*(6.96)*	*(0.60)*	*(2.69)*	*(8.51)*		
MCAP Q3									
Xetra8	6.21***	-0.76	7.00***	10.60***	0.20	10.85***	11.69***	.	.
t-Value	*(14.29)*	*(-0.96)*	*(8.01)*	*(11.63)*	*(0.08)*	*(4.93)*	*(5.73)*		
MCAP Q4									
Xetra8	6.69***	-1.09	7.78***	14.38***	2.62	9.77**	15.44***	.	.
t-Value	*(9.64)*	*(-0.71)*	*(4.98)*	*(6.77)*	*(0.71)*	*(2.21)*	*(3.39)*		

Table 5: **Results Per Trade Turnover Xetra**: The sample consists of stocks listed in Deutsche Boerse's HDAX segment. The observation period comprises of 40 trading days before and after the introduction of Xetra 8 on 23 April 2007. Table 5 Panel A reports descripitive statistics for per trade turnover based on tick data. The mean and standard deviation are reported and calculated on each individual trade. Panel B reports panel regression results for daily per trade turnover. The panel regression is performed on daily measures individually for each stock using the following regression formula: $LM_{i,t} = \alpha_i + \delta Xetra8_{i,t} + \beta VDAX_t + \epsilon_{i,t}$. Results are reported for the entire sample and individually by a stock's market capitalization. Stocks are divided into four groups depending on their market capitalization. Market capitalization is calculated as the product of shares outstanding and the average price over the observation period for a single stock. All results are reported in Euros. The tables include results for Xetra only. All measures for the panel regressions are calculated as relative measures. *** indicates significance at the 1% level, ** indicates significance at the 5% level, and * indicates significance at the 10% level. Robust t-statistics are reported in parentheses.

Panel A: Average Per Trade Turnover - Descriptives

	Sample		MCAP Q1		MCAP Q2		MCAP Q3		MCAP Q4	
	mean	Std. Dev.	mean	Std. Dev.	mean	Std. Dev.	mean	Std. Dev.	mean	Std. Dev.
full	49,205	96,676	63,568	113,681	30,830	57,146	17,746	26,456	11,953	19,700
pre	48,021	95,423	62,540	113,425	29,890	51,362	17,167	25,078	11,817	20,122
post	50,441	97,951	64,622	113,932	31,809	62,575	18,370	27,853	12,120	19,170

Panel B: Average Per Trade Turnover - Panel Regressions

	Sample	MCAP Q1	MCAP Q2	MCAP Q3	MCAP Q4
Xetra8	742.29	1065.37	1857.478***	85.64	-52.24
t-Value	*(1.38)*	*(0.56)*	*(2.78)*	*(0.22)*	*(-0.21)*

Table 6: **Results Panel Regressions**: The sample consists of stocks listed in Deutsche Boerse's HDAX segment. The observation period comprises of 40 trading days before and after the introduction of Xetra 8 on 23 April 2007. Table 3 Panel A reports panel regression results for quoted spread, effective spread, and realized spread. Panel B reports regression results for price impact, permanent impact of trade innovation, and quote-based price discovery in percent. The panel regression is performed on daily measures individually for each stock using the following: $LM_{i,t} = \alpha_i + \delta Xetra8_{i,t} + \beta VDAX_t + \epsilon_{i,t}$. Results are reported for the entire sample and individually by stock market capitalization. Market capitalization is calculated as the product of shares outstanding and the average price over the observation period for a single stock. All results are reported in basis points expect for quote impounded information which is measured in percent. Xetra denotes results for the Xetra System, FWB denotes results for the Frankfurt floor, and Diff denotes Xetra-FWB. All measures for the panel regressions are calculated as relative measures. *** indicates significance at the 1% level, ** indicates significance at the 5% level, and * indicates significance at the 10% level. Robust t-statistics are reported in parentheses below panel regressions' coefficients.

Adverse Selection and Quote Fraction of Price Discovery					
	Price Impact			Trade Innovation	Var. Decomp.
	Xetra	FWB	Diff	Xetra	Xetra
Sample					
Xetra8	-7.29***	-0.13	-7.17***	-3.16***	47.89***
t-Value	*(-17.32)*	*(-0.42)*	*(-14.44)*	*(-21.04)*	*(72.04)*
MCAP Q1					
Xetra8	-2.72***	0.34	-3.04***	-1.60***	47.74***
t-Value	*(-6.96)*	*(1.01)*	*(-4.91)*	*(-16.05)*	*(34.80)*
MCAP Q2					
Xetra8	-6.66***	-0.82	-5.81***	-2.89***	50.87***
t-Value	*(-9.80)*	*(-1.13)*	*(-7.09)*	*(-13.39)*	*(41.08)*
MCAP Q3					
Xetra8	-8.58***	0.17	-8.81***	-3.54***	48.80***
t-Value	*(-18.93)*	*(0.29)*	*(-13.85)*	*(-22.33)*	*(35.11)*
MCAP Q4					
Xetra8	-11.39***	-0.22	-11.18***	-4.70***	44.17***
t-Value	*(-17.22)*	*(-0.36)*	*(-12.43)*	*(-16.48)*	*(48.68)*

Table 7: **Results Panel Regressions by Trade Size**: The sample consists of stocks listed in Deutsche Boerse's HDAX segment. The observation period comprises of 40 trading days before and after the introduction of Xetra 8 on 23 April 2007. Filters applied to HDAX stocks reduce the sample to 101 continuously trades stocks. Table 7 reports panel regression results for price impacts by trade size. The panel regression is performed on daily measures individually for each stock using the follwing regression formula: $LM_{i,t} = \alpha_i + \delta Xetra8_{i,t} + \beta VDAX_t + \epsilon_{i,t}$. The regression tests the null-hypothesis of no influence of the Xetra 8 update on liquidity and information measures. Results are reported for the entire sample and individually by a stock's market capitalization. Different trade sizes are below 25,000 EUR, between 25,000 EUR and 100,000 EUR, and above 100,000 EUR. Stocks are divided into four groups depending on their market capitalization. Market capitalization is calculated as the product of shares outstanding and the average price over the observation period for a single stock. Xetra denotes results for the Xetra System, FWB denotes results for the Frankfurt floor, and Diff denotes Xetra-FWB. All results are reported in basis points. All measures for the panel regressions are calculated as relative measures. *** indicates significance at the 1% level, ** indicates significance at the 5% level, and * indicates significance at the 10% level. Robust t-statistics are reported in parantheses below panel regressions' coefficients.

Regressions on Price Impacts by Trade Size									
	< 25 kEUR			≥ 25 kEUR and ≤ 100 kEUR			> 100 kEUR		
	Xetra	FWB	Diff	Xetra	FWB	Diff	Xetra	FWB	Diff
Sample									
Xetra8	-6.48***	-0.05	-6.42***	-12.35***	-1.18	-9.64***	-16.85***	.	.
t-Value	*(-16.38)*	*(-0.18)*	*(-12.84)*	*(-11.45)*	*(-0.96)*	*(-8.38)*	*(-8.80)*		
MCAP Q1									
Xetra8	-2.38***	0.26	-2.61***	-3.11***	1.44*	-4.20***	-5.16***	-2.44	-1.73
t-Value	*(-6.85)*	*(0.87)*	*(-4.86)*	*(-4.95)*	*(1.72)*	*(-5.74)*	*(-5.80)*	*(-1.34)*	*(-0.94)*
MCAP Q2									
Xetra8	-5.82***	-0.63	-5.15***	-9.18***	-1.72	-6.84***	-16.48***	.	.
t-Value	*(-8.85)*	*(-1.00)*	*(-6.40)*	*(-8.40)*	*(-0.88)*	*(-3.31)*	*(-9.27)*		
MCAP Q3									
Xetra8	-7.58***	0.06	-7.67***	-13.98***	-1.40	-13.63***	-22.55***	.	.
t-Value	*(-16.79)*	*(0.10)*	*(-11.19)*	*(-13.13)*	*(-0.62)*	*(-7.25)*	*(-8.90)*		
MCAP Q4									
Xetra8	-10.30***	0.08	-10.40***	-23.86***	-4.65	-18.19***	-32.25***	.	.
t-Value	*(-15.58)*	*(0.13)*	*(-10.59)*	*(-9.81)*	*(-1.56)*	*(-5.13)*	*(-6.48)*		

Table 8: **Equality of Variances of Quote Based Information**: The sample consists of stocks listed in Deutsche Boerse's HDAX segment. The observation period comprises of 40 trading days before and after the introduction of Xetra 8 on 23 April 2007. Table 8 Panel A reports descriptive statistics for mean and standard deviation of the variance decomposition before and after the introduction of Xetra 8. Panel B reports results of the Brown and Forsythe (1974) robust test for the equality of variances for the quote based information fraction before and after the introduction of Xetra 8. The tests and descriptives are based on daily per instrument data. The test is a modified Levene test and more robust against non-normality than the normal Levene test.

Panel A: Quote Based Information Fraction - Descriptives

	Sample		MCAP Q1		MCAP Q2		MCAP Q3		MCAP Q4	
	mean	Std. Dev.	mean	Std. Dev.	mean	Std. Dev.	mean	Std. Dev.	mean	Std. Dev.
pre	42%	15%	42%	12%	39%	14%	43%	15%	45%	19%
post	90%	8%	90%	6%	90%	8%	92%	7%	90%	10%

Panel B: Quote Based Information - Equality of Variances Before and After Event

	Sample	MCAP Q1	MCAP Q2	MCAP Q3	MCAP Q4
F-value	1,272.91	291.23	271.22	421.83	361.64
p-value	<.0001	<.0001	<.0001	<.0001	<.0001

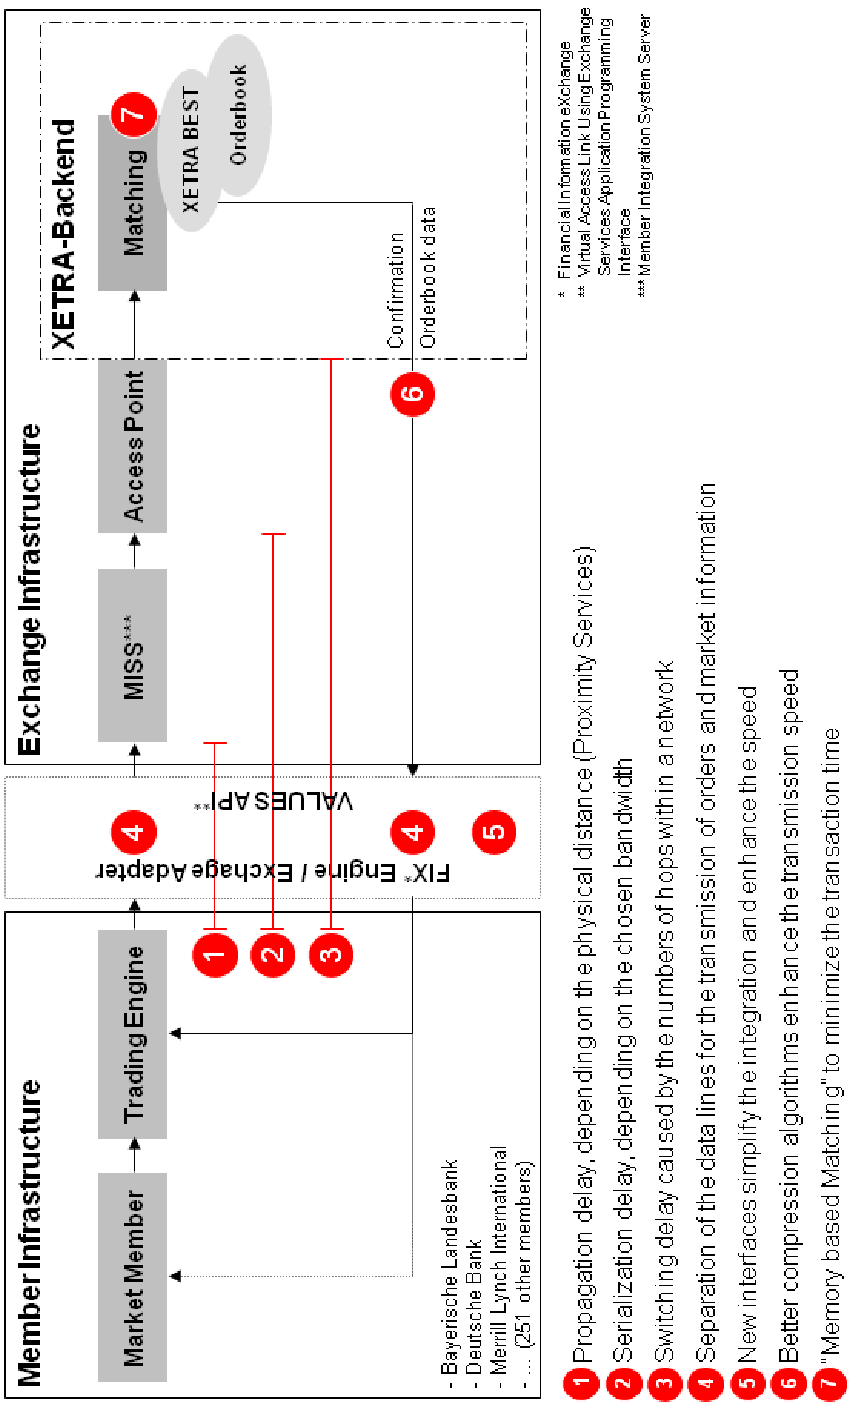

Figure 1: Xetra Latency Reducing Upgrades: This figure illustrates the latency reducing upgrades to Xetra.

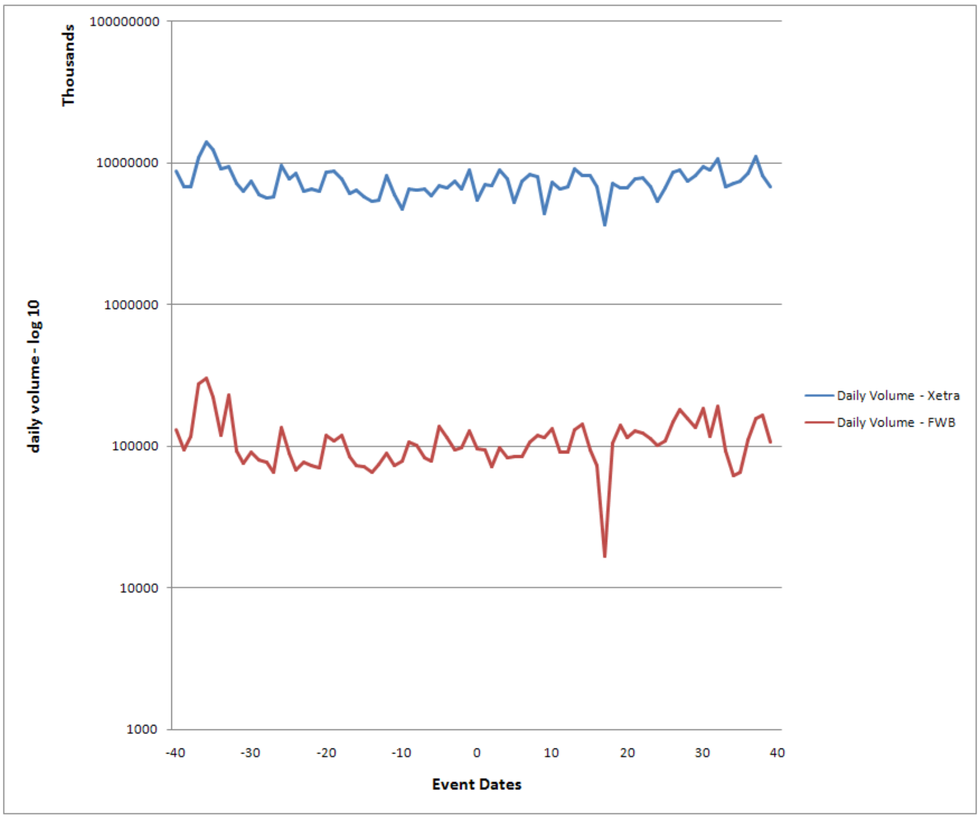

Figure 2: Daily Volume: Xetra and FWB: This figure graphs the daily log euro volume from the Frankfurt floor and Xetra markets. The event dates are on the x-axis and the daily log euro volume on the y-axis.

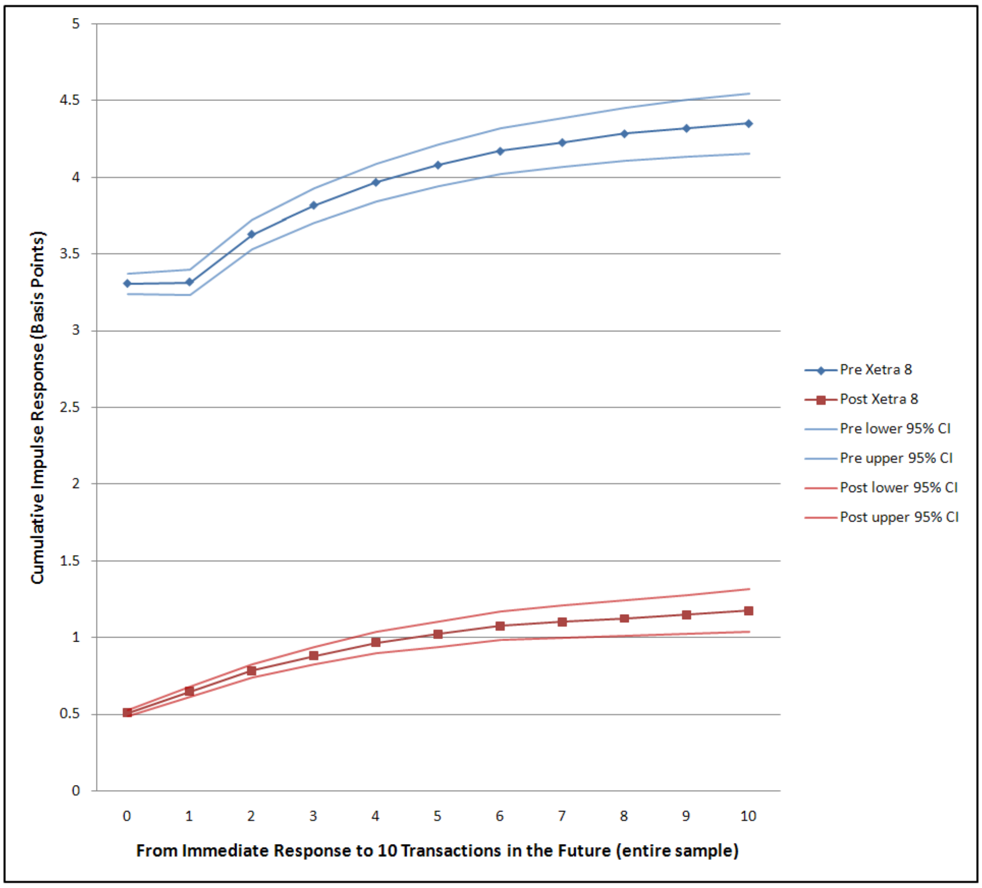

Figure 3: Cumulative Impulse Response (Basis Points) - Entire Sample: In this figure the cumulative impulse response of a trade is on the y-axis. The x-axis is the forecast horizon in trades. The blue lines are pre-event with 95% confidence intervals. The red values are post-events with 95% confidence intervals.

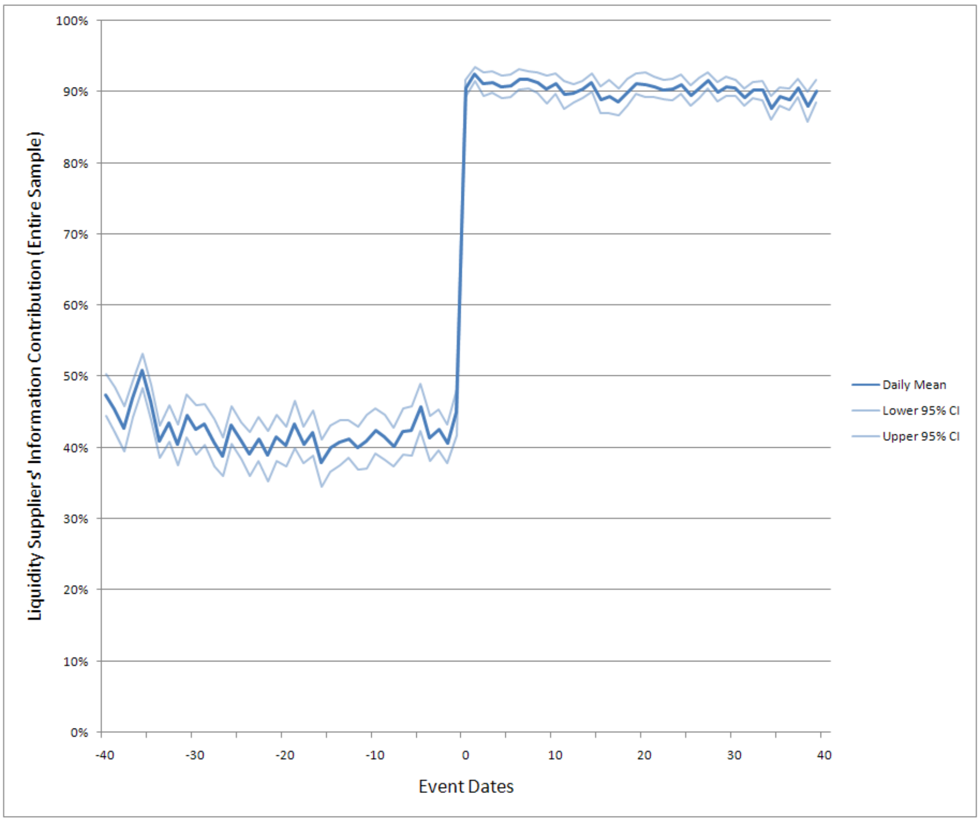

Figure 4: Quoted Based Information (Percent) - Entire Sample: The figure graphs the non-trade/quote-based correlated information with 95% confidence intervals. The events dates are on the x-axis and the quote-based contribution in percent on the y-axis.