Summer 10-6-2011

FROM MINUTES TO SECONDS AND BEYOND: MEASURING ORDER-BOOK RESILIENCY IN FRAGMENTED ELECTRONIC SECURITIES MARKETS

Michael Chlistalla

Follow this and additional works at: http://aisel.aisnet.org/ecis2011

FROM MINUTES TO SECONDS AND BEYOND: MEASURING ORDER-BOOK RESILIENCY IN FRAGMENTED ELECTRONIC SECURITIES MARKETS

Michael Chlistalla, Goethe-University Frankfurt, Grüneburgplatz 1, 60323 Frankfurt am Main, Germany, chlistalla@wiwi.uni-frankfurt.de

Abstract

This study sheds light on the resiliency of competing electronic order-driven markets by investigating the time path of liquidity after a large endogenous shock by means of an intra-day event study. We confirm earlier results that large trades, which qualify as shocks for the purpose of this analysis, are timed as they occur when liquidity in the market is extraordinarily high. Moreover, we find that adverse liquidity effects peak within the first-minute interval after a large trade. In contrast to previous researchers' results, we discover that liquidity recovers not within minutes, but within a few seconds so that the adverse effect already diminishes significantly towards the end of the first post-event minute. Three minutes later, basically no residue from the liquidity shock remains observable. We attribute this finding to the fact that temporary order-book imbalances are immediately detected by algorithmic trading engines and high-frequency traders whose order management (IT-) systems are continuously eagle-eying the stock markets for such opportunities. Interestingly, our results differ depending on the market under scrutiny.

Keywords: Electronic Trading Systems, Resiliency, Market Quality.

1.   Introduction

European securities markets are increasingly competing for investors' order-flow. Competition between execution venues is important as it drives down trading costs; the lower the trading costs, the more attractive is the market for investors. However, a competitive market requires by definition that there be more than one supplier in the market, and as such, competitive markets are always fragmented (Lee, 2002). In securities markets, fragmentation arises when all orders do not interact with each other via a single execution mechanism but across a multitude of execution venues.

This leads to two interesting questions: Firstly, what does "attractive" in this context refer to, i.e. when is a market attractive for traders? And secondly, what are the implications of fragmentation for the attractiveness of an execution venue to traders? Before answering these questions, we shall have a brief look at the technological developments that continue to shape the securities industry: beginning in the late 1990s, the electronification of execution venues across Europe enabled market participants (banks, brokers as well as their institutional and retail clients) to remotely access electronic order-books (Gsell and Gomber, 2009). Since the first decade of this century, electronification goes much further: algorithmic trading, defined in general as the use of computer algorithms to manage the trading process (Hendershott et al., 2011), is increasingly used to divide large trades into several smaller trades in order to minimize market impact.[1] Hendershott et al. (2011) show that algorithmic trading may have a positive impact on market liquidity.

Market liquidity is widely regarded to be the most important criterion for market quality (Clemons and Weber, 1992), giving investors the opportunity to buy and sell stocks immediately and without adversely affecting the price (Gomber et al., 2004). Coming back to the question of attractiveness, we can state that the more liquid a market, the higher its attractiveness for traders and investors. Aside from the explicit costs (such as exchange fees, taxes etc.), overall execution costs are determined implicitly by a market's liquidity (Davydoff et al., 2002); hence, the higher the liquidity in a certain stock, the lower the implicit trading costs for investors demanding immediacy (Irvine et al., 2000). In terms of quantifying liquidity, pertinent literature unanimously suggests that liquidity cannot be captured by a single measure (Bernstein, 1987; Dong et al., 2007) and thus identifies four main dimensions of market quality (immediacy, width, depth, and resiliency), which have been researched extensively both theoretically and empirically (Garbade, 1982; Kyle, 1985; Harris, 1990). Yet, to the best knowledge of the authors, no academic studies have empirically assessed and compared the recovery of liquidity (resiliency) across electronic trading system-based execution venues – the most common organization of equities trading in Europe nowadays – in the context of newly arising competition between these execution venues.

Referring to the second question above – the implications of fragmentation for the attractiveness of a market – a key issue is how (i.e. based on which information) investors choose among alternative markets. Although being an important decision criterion for investors facing this choice, resiliency seems to have been neglected in previous empirical academic research. Following the argumentation by Hitt and Brynjolfsson (1996) that "a firm with a unique access to IT may be in a position to earn higher profits from that access", it might well be conceivable that a firm with such a unique access is one that employs IT in order to specifically exploit resiliency differentials between execution venues.

In this paper, we study the time path of a specific resiliency measure (the Exchange Liquidity Measure, XLM) after a shock that adversely affects liquidity on a stock's home market. We then compare how this time path evolves to similar order-book situations on alternative trading venues. The objective is to detect whether different types of execution venues are equally resilient in their capacity to restore liquidity after such an adverse liquidity shock. In doing that, our work addresses and contributes to the academic literature on the impact of market fragmentation and competition between electronic markets

---

[1] A subset of algorithmic traders, known as high-frequency traders, whose trading strategies are based on very short holding periods, accelerates the trading process even more.

as well as on overall liquidity in order-driven markets. In the field of IS literature, our study contributes to the research on the value of IT as it exemplifies how a quicker replenishment of liquidity attributable to the growing use of algorithmic trading increases market efficiency. Thereby, the advent of high-frequency databases facilitates analyzing the role of time in the trading process. The remainder of this paper is organized as follows: The next section presents previous literature related to our research question, both from the realms of IS- as well as from market microstructure literature. The subsequent section describes our dataset and methodology. The fourth section reports our findings while the last section concludes.

2.   Related Literature

As mentioned above, our research contributes both to the field of IS, in particular on the general value of IT, and to the field of market microstructure in securities trading.

In terms of the general value of IT, reviews of current academic literature (e.g. Kohli and Grover, 2008) suggest that researchers have disengaged from the question *whether* IT does create value, as numerous studies have found that there *is* a relationship between IT and some aspect of firm value, be it financial, intermediate (e.g., process-related) or affective (e.g., perception-related). Melville et al. (2004) evaluate accumulated knowledge of IT business value research – on the one hand, from studies that emphasize focal firm dynamics and on the other hand from studies that include factors in the competitive environment. They propose that the greater the degree of competition in an industry, the greater the extent to which firms achieve efficiency gains via IT. With newly arising competition among execution venues in Europe and an ongoing arms race of IT, it is a highly relevant issue to look at how market participants can exploit efficiency gains by means of IT given the increased competitive pressure. In the sense of Hitt and Brynjolfsson (1996), the resiliency of a market may be interpreted as an indicator for the unique capability of market participants and their IT to achieve just that across a multitude of fragmented execution venues: resiliency as defined in the following measures the amount of time it takes until a temporary liquidity shock in the market is corrected; how quickly this correction occurs is likely to strongly depend on the IT deployment.

In terms of market microstructure in securities trading, there are extensive amounts of research in both theoretical and empirical literature focusing on market quality in general and on liquidity in specific. As stated earlier, liquidity is regarded as the most important criterion for market quality. However, only resiliency takes account of the time dimension of liquidity. Resiliency may very generally be defined as the rate at which pricing errors caused by temporary order-flow shocks are corrected in the market (Dong et al., 2007). Yet, different perceptions of resiliency exist: Garbade (1982) looks at resiliency from an *order replenishment* perspective. Black (1971), Kyle (1985), Holthausen et al. (1987), and Harris (2003) analyze resiliency from the *price recovery* perspective. Gomber et al. (2004) propose a *liquidity perspective*, where a market is said to be resilient when liquidity quickly reverts to normal levels after an adverse shock. This study also takes on the liquidity perspective.

As resiliency provides a key insight into the nature of the market, there are a number of relevant studies which investigate the role and importance of resiliency in the case of electronic limit order-books, with the better part of research in this context consisting of empirical studies and only a minor part constituting theoretical work: Degryse et al. (2005) provide an excellent literature overview which covers both theoretical and empirical research in terms of the interaction of resiliency with aggressive orders, tick size, firm size and order flow. This literature review may be complemented on the theoretical side by the work of Foucault et al. (2005), which is an academic view on limit order-books and liquidity implications investigating how traders' impatience affects order-book strategies, bid-ask spreads and market resiliency, and by Large (2007), who proposes an intensity model for order arrivals and uses that model to study order-book resiliency for a single stock traded on the London Stock Exchange. On the empirical side, we shall mention Dong et al. (2007), who investigate the main features of resiliency (trading activity, tick size, information asymmetry, and volatility) and its effect on stock returns. They are able to confirm that these are all significant determinants of resiliency, which in turn is only weakly related

with two other price and quantity dimensions of liquidity (spread and depth) and thus provides significant new information on market quality.

Lastly, it is equally important to position this research piece within the academic literature on latency, algorithmic (AT) and high-frequency trading (HFT). Particularly in recent years, latency[2] has grown to be of paramount importance for traders pursuing strategies that rely on short-term relative price differentials; yet, latency reductions were sought for all along over the last decades. In this regard, Easley et al. (2009) study the impact of a 1980 major upgrade to the NYSE's trading environment which reduced the execution latency experienced by traders off the floor relative to traders on the floor of the exchange. The authors conclude that leveling the playing field improved liquidity and that the technological upgrade which increased competition had an economically significant impact on stock returns. Nowadays, as more data becomes available on AT/HFT trades and orders, research in this field begins to emerge: Riordan and Storkenmaier (2009), for instance, study the effects of latency on liquidity and on price discovery in an electronic limit order market. Analyzing an important system upgrade at Deutsche Börse in 2007, they are able to confirm that a reduction of system latency leads to a decrease of the effective spread and thus contributes to market liquidity. Hasbrouck and Saar (2011) investigate the impact of "low-latency traders" (in contrast to HFT) on traditional market quality measures and find that their activity improves short-term volatility, spreads, and displayed depth in the limit order-book. One of the first studies to analyze whether AT improves market quality is Hendershott et al. (2011), who find that AT improves liquidity and enhances the informativeness of quotes. Many other studies[3] have since tried to shed light on the influence of AT and HFT on market quality in general and liquidity in specific.

Our study is distinguished from previous empirical analyses in two ways. Firstly, in terms of methodological details and available data quality: while Dong et al. (2007), for instance, draw on minute-by-minute data and Gomber et al. (2004) use one-minute order-book snapshots, data points in our study are time-stamped to the millisecond and the granularity of snapshots used is on a second-by-second basis, thus accounting for the recent increases in speed on electronic trading venues. Secondly and more importantly, we position ourselves very differently in terms of the research question asked: while most (not all) previous studies focus on quote-driven markets in the U.S., none of them have, to the best of our knowledge, ever addressed the influence of the newly arising competition between order-driven securities markets in Europe on resiliency (i.e. the time dimension of liquidity) as an indicator for market quality.

3.   Dataset and Methodology

In this section our research approach will be presented. To begin with, the basic characteristics of the Euronext Paris, Chi-X and Turquoise market models will be examined, followed by the description of our dataset, i.e. sample selection procedure and data source. Eventually, we will elaborate on the methodology applied.

---

[2] There is no commonly agreed definition of latency, neither in academic research nor among practitioners: Riordan and Storkenmaier (2009), for instance, define latency in an electronic order-driven market as the time that elapses between submitting an order and receiving confirmation that the order was executed. Hasbrouck and Saar (2011) define latency as the sum of three components: the time it takes for information to reach the trader, the time it takes for the trader's algorithms to analyze the information, and the time it takes for the generated action to reach the exchange and get implemented. Exchange operators typically define latency much more narrowly, usually as the processing delay from the moment an order enters the operator's IT system until an acknowledgement leaves the operator's IT system.

[3] E.g., Jarnecic and Snape (2010), Jovanovic and Menkveld (2010), Brogaard (2010), Cvitanic and Kirilenko (2010), Menkveld (2010).

3.1   Market model characteristics

Euronext is a pan-European stock exchange based in Paris with subsidiaries in several European markets. Euronext Paris (ENP) operates a centralized hybrid market (i.e. quote- and order-driven) using an electronic trading system, where securities that are sufficiently liquid or securities with a designated liquidity provider are traded continuously following price-time priority[4]. The stocks we study are constituents of the blue-chip index CAC-40, for which ENP is the primary ("home"). All orders are anonymous in the order-book. Trading of French stocks traditionally concentrated on Euronext; there were no regional exchanges in France and even though some blue-chip stocks had been cross-tradable on European exchanges for some ten years, equity trading had mainly focused on the home market. Today, with multi-market competition across Europe, ENP's market share in French CAC-40 stocks, for which ENP is the home market, has fallen to around 67% on average.

UK-based **Chi-X Europe** is one of the new market entrants and has gained a considerable market share in European blue-chip stocks, which amounted to around 18% in French CAC-40 equities on average in September 2010. Chi-X launched its fully electronic trading system in March 2007 and currently serves 15 European markets. Its specific characteristics make Chi-X relevant for academic research: it was the first trading venue in Europe to adopt a maker-taker fee structure (0.30bps are charged for aggressive executions, while passive executions receive a 20bps rebate); it is currently the alternative trading venue in Europe with the largest market share; and it is a stand-alone provider of trading services in contrast to other venues that are part of a structure integrating trading, clearing and settlement services into a single group. Trades on Chi-X are matched in price-time priority by a fully electronic proprietary matching engine. All orders are anonymous on the order-book. Stocks are traded in their official local currency, i.e. EUR for the CAC-40 stocks used in this analysis.

Turquoise is another new player in the European equities trading arena. As an authorized investment firm regulated by the UK's FSA, Turquoise provides equities trading covering 18 European markets. Originally established by a consortium of European investment banks, Turquoise has been majority-owned by the London Stock Exchange Group since February 2010. Its market share amounted to around 4% in French CAC-40 equities on average in September 2010. Trades on Turquoise are matched according to price-transparency-time-priority by a fully electronic matching engine. This means that visible limit orders are matched before dark orders. All orders are anonymous on the order-book. Stocks are traded in their official local currency, i.e. EUR for the CAC-40 stocks used in this analysis. Applicable tick sizes are those defined by the primary markets at the relevant time. Aggressive orders are charged 0.28 bps, while passive orders are rebated between 0.20 and 0.24 bps depending on a member's level of the trading activity (Turquoise 2010).

As presented above, all three market structures exhibit similar market design characteristics for our sample of CAC-40 stocks in a way that these stocks are traded continuously in an electronic order-book and trading is organized in an order-driven manner. All three venues feature visible as well as non-displayed order types whereby the latter imposes a limitation to our dataset that will be addressed in the next subsection.

3.2   Dataset

3.2.1   Sample selection

Our sample period ranges from November 1st, 2009 to April 30th, 2010. This period was chosen as it represents a relatively stable trading environment without any severe exogenous shocks (e.g. any event similar to the Lehman bankruptcy). The sample of instruments in this analysis comprises those 15

---

[4] Euronext's *Internal Matching Service* is an optionally available service that uses a trading algorithm privileging a price-member-time priority, which allows orders at the best price originating from the same member firm subscribing to this service to be matched in the central order book.

stocks from the French blue-chips index CAC-40 that are also constituents of the Eurostoxx50 index. They represent some of Europe's most liquid and thus most heavily traded stocks. All stocks are tradable on Euronext Paris, Chi-X and Turquoise. We checked whether any of them had dropped out or been replaced in the CAC-40 index during the sample period. This was not the case. Likewise, none of the stocks implemented or cancelled an ADR program during the sample period. In order to avoid any distorting effects, we additionally ensured that no other new execution venue commenced its trading activities in French stocks during the sample period, which was also not the case.

3.2.2   Data source

For the purpose of our study, we use data retrieved from the Thomson Reuters Tick History (TRTH) database. Those data include the ten highest bid and ten lowest ask limits on each side of the electronic order-book. Every change in the order-book generates a new database entry and is time-stamped to the millisecond.

The first and last five minutes of each trading day are removed to avoid biases associated with the inventory management and information processing at those times. Opening, closing and potential intraday auctions are identified via Reuters' qualifying code indicating auctions. We use these qualifiers to filter the data to exclude auction periods. For our purposes we use limit order-book snapshots sampled either every minute or every second (see later section for specification).

Although all three execution venues feature non-displayed order types in their market models, publicly available order-book data lack this hidden liquidity. We can thus measure the change in displayed liquidity following an endogenous shock, but not the change in overall (hidden and displayed) liquidity. The change in displayed liquidity could therefore underestimate or overestimate overall liquidity changes when e.g. order flow has shifted from displayed to hidden.

3.3   Methodology

In order to investigate and to compare the resiliency of the three markets, we analyze both the immediate effect of an endogenous shock that adversely affects liquidity and the time path of liquidity recovery. Two steps are therefore required to perform the analysis: firstly, the determination of the relevant events (large trades) and secondly, the measurement of the time path of liquidity.

3.3.1   Determination of endogenous shocks

An endogenous shock is represented by a large trade that hits the market; the shock is endogenous[5] for it originates from within the market. Large trades typically consume liquidity from a central limit order-book and thus adversely affect the amount of liquidity available. The size of the shock is directly related to the trade size (Gomber et al., 2004). However, whether a particular trade order is "large" or "small" is a relative notion, depending on the number of traders present in the market ready to fill that order (Massimb and Phelps, 1994).

In order to determine a dataset of endogenous shocks, the easiest way would be to select the 100 largest trades per stock and per market during the sample period as the events relevant for the purpose of this analysis. However, a cross-market comparison reveals that on average, the 100 largest trades are 3.5 times larger on Euronext Paris (the stocks' home market) than on Chi-X and even 5.5 times larger on ENP than on Turquoise. So, in order to guarantee the comparability and validity of the results, a modified approach is chosen: firstly, an upper boundary (*UB*) of the range of large trades qualifying as endogenous shocks is determined as $UB_{All}^i = \min(LT_{ENP}^i, LT_{CHI}^i, LT_{TQ}^i)$, where LT_j^i represents the largest trade per stock *i* and per market *j*. From this upper boundary *UB* downwards, the 100 largest trades per stock and per market during the sample period are determined, or more precisely the exact instant in which they occur. As the event window which shall be analyzed extends from 15 minutes prior to a large trade

---

[5] In contrast to exogenous shocks, which originate from outside the market, such as e.g. corporate events, news, etc.

until 15 minutes after the trade, we exclude trades that occur within the first or last 15 minutes of each market's continuous trading session as well as any transactions that occur within 15 minutes before or after any intraday call auction. By making use of the TRTH qualifying codes we ensure that only order-book trades are selected (and e.g., no reported OTC-trade is included in the dataset).

Table 1 shows the upper boundary (UB) for all stocks that are part of this study. The first column is the Reuters Identifier Code (RIC) for the 15 stocks in our sample. The second column presents the upper boundary of the interval (UB_{All}^i) per stock.

RIC	UB_{All}^i	RIC (cont'd)	UB_{All}^i (cont'd)	RIC (cont'd)	UB_{All}^i (cont'd)
AIRP.PA	368,287	FTE.PA	343,332	SGOB.PA	207,660
AXAF.PA	526,171	OREP.PA	385,110	SOGN.PA	528,985
CAGR.PA	168,006	SASY.PA	400,822	TOTF.PA	432,175
CARR.PA	254,147	SCHN.PA	529,932	UNBP.PA	192,680
DANO.PA	485,633	SGEF.PA	437,682	VIV.PA	238,329

Table 1.   Upper boundary of large trades (in EUR)

3.3.2   Liquidity measurement

In order to illustrate possible changes in order-book liquidity we apply three variables, namely (i) the quoted bid-ask spread, and (ii) the quoted volume at the top of the book, i.e. the number of shares at the top of the book for both sides multiplied by the associated quote and (iii) the Exchange Liquidity Measure (XLM) as developed by Gomber and Schweickert (2002). We use the third measure to capture the depth of the order-book, i.e. the order-book liquidity beyond the best bid and ask. The XLM measures the execution costs of a (synthetic) round-trip transaction and incorporates the information from all orders within the order-book to calculate the weighted average price at which an order of given size (EUR-denominated in our case) could be executed immediately at time t. The measure is given in basis points (bps) for a given euro transaction volume. The interpretation of the measure is straightforward: the lower it is, the more liquid the market for a certain stock. The weighted average price is denoted by $P_{B,t}(V)$ and $P_{S,t}(V)$, respectively, where the index (B, S) indicates the type of the transaction (buyer-initiated or seller-initiated) and V denotes the order size. Let MQ_t denote the quote midpoint at time t. Execution costs for a buy and a sell order in bps are calculated by

$$XLM_{B,t}(V) = \frac{P_{B,t}(V) - MQ_t}{MQ_t} \cdot 10{,}000 \quad and \quad XLM_{S,t}(V) = \frac{MQ_t - P_{S,t}(V)}{MQ_t} \cdot 10{,}000$$

For the execution costs of a round-trip transaction at time t both measures are added. A similar measure has been suggested by Irvine et al. (2000) who considered spreads alone not to be sufficient measures for market liquidity. For our analyses, we assume round-trip transactions of V=100 kEUR.

3.3.3   Intraday event analysis

In order to analyze the impact block trades have on the liquidity and on the resiliency of the market, we employ an intraday event study approach. In our general model, the observation immediately prior to the large trade is defined as t_0. The impact of the large trade is measured as the difference in liquidity (operationalized by either the quoted spread, the quoted volume, or by the XLM) from t_0 to t_1. We define a pre-event window ranging from 15 minutes prior to the large trade until t_0 [-15:0] and a post-event window ranging from t_0 to 15 minutes after the large trade [0:15]. For the analysis, a complete series of XLM(V) observations per large trade is required. In a slightly modified specification ("adapted model"), we will additionally have an in-depth look at the two intervals [-1:0] and [0:1], i.e. the one minute immediately prior to and after the event, respectively.

4. Results

The following Table 2 presents the results of the general model: the size of the liquidity shocks and the corresponding t-statistics are shown in the column labeled [0:1]. The unit of measurement is bps for the XLM and EUR cents for the quoted spread. Let's focus on Euronext (ENP) first: here, we can see that the values are positive and statistically significant. This is absolutely intuitive as large trades, by definition, adversely affect liquidity and thus lead to an increasing XLM measure. In the next interval, [1:2], the values turn negative which indicates a recovery of liquidity. This is again exactly what we expect to see, because in resilient markets, liquidity quickly reverts to its normal level. If we now turn to the columns labeled [0:3] and [0:15], i.e. look at the order-book's liquidity three minutes (15 minutes respectively) after the event, we would anticipate a lower value than in the [0:1] interval as we expect liquidity to recover progressively the farther away from the event. Contrary to that expectation, we can observe that there is nearly no change in liquidity between [0:1], [0:3] and [0:15]. This holds true for both ENP and Chi-X, only Turquoise (TQ) exhibits a different behavior: here, the recovery rate (defined as the proportion of the adverse liquidity effect compensated until minute t) is about 30% after three minutes and about 35% after 15 minutes. Apparently, refilling the order-book and thus restoring a sufficient depth of the book takes longer in a less heavily traded market which highlights that the Turquoise market share is well below the one of Chi-X).

Euronext Paris		[-15:0]	[0:1]	[1:2]	[2:3]	[3:4]	[4:5]	[0:3]	[0:15]	[-15:15]
	$\triangle$quoted spread	-0.57***	0.43***	-0.13***	0.02	0.11***	-0.05	0.32***	0.33***	-0.24***
		-8.04	7.32	-2.86	0.47	2.13	-0.96	5.42	5.35	-3.98
	$\triangle$XLM (100)	-1.486***	1.040***	-0.176**	0.146**	0.159**	-0.113	1.010***	0.997***	0.496***
		-11.66	9.62	-2.43	1.98	2.18	-1.52	8.97	8.65	-5.00
Chi-X		[-15:0]	[0:1]	[1:2]	[2:3]	[3:4]	[4:5]	[0:3]	[0:15]	[-15:15]
	$\triangle$quoted spread	-0.36***	0.23***	0.05	-0.04	0.04	0.02	0.24***	0.19***	-0.16***
		-6.21	4.18	1.26	-0.90	0.83	0.45	4.34	3.22	-2.84
	$\triangle$XLM (100)	-2.772 ***	1.922***	-0.009	0.003	0.016	-0.059	1.931***	1.734***	-1.120***
		-15.30	13.89	-0.08	0.03	0.17	-0.63	13.85	11.98	-6.62
Turquoise		[-15:0]	[0:1]	[1:2]	[2:3]	[3:4]	[4:5]	[0:3]	[0:15]	[-15:15]
	$\triangle$quoted spread	-0.17	0.52***	-0.12	0.31***	-0.21**	0.07	0.85***	0.49***	-0.18
		-1.55	5.09	-1.16	0.01	-2.06	0.67	5.84	3.83	-1.39
	$\triangle$XLM (100)	-6.288***	5.019***	2.788**	-1.357	-0.454	0.457	3.483**	3.217***	-1.388
		-4.17	3.27	1.96	-0.96	-0.31	0.31	2.32	2.43	-0.75

Table 2.   Event study results ("General model" specification). The upper entry in each cell shows the change (delta) in the liquidity measure indicated in the first row over the period given in the first line. The lower entry in each cell represents the t-statistic. Asterisks represent significance at the 1%-, 5%-, and 10%-level, respectively.

A possible explanation for these first results is that the liquidity shock must have had its peak already *within* the first-minute interval [0:1] and the (remaining) effect we observe *at the end* of this interval is already compensated by the recovery that has taken place between t_0 and t_1. This seems to confirm what Gomber et al. (2004) suspected in a similar study that focused on the resiliency of the Xetra order-book. There, the authors discuss the probability of understating the size of a liquidity shock (as measured by the difference in XLM from t_0 to t_1) due to the chosen observation frequency. Having selected a frequency of one minute, they identify the risk of understating the impact of a certain transaction when new limit orders are submitted during the delay between t_0 and the subsequent observation. The authors admit that such an understatement is quite likely given the constant flow of small limit orders[6]. Hence,

---

[6] The authors argue that these limit orders are unlikely to be submitted in direct response to the liquidity shock because traders need some time to observe the market and to react to the liquidity change. While this argumentation may indeed have been valid at the time of the analysis, we would consider this explanation highly unlikely nowadays. In recent years, speed in order execution has become paramount and a plethora of algorithmic trading engines and high-frequency traders are conti-

it seems rational to submit the interval [0:1] to a closer scrutiny. This shall first be done by reference to some graphical representations. Figure 1 provides a graphical representation of the results of the intra-day event study, exemplified for the market ENP and the XLM (100,000) measure. The horizontal axis measures the time, relative to the time of the large transaction t_0, in minutes. Liquidity tends to increase in the minutes prior to the shock, with a very pronounced rise clearly visible in the minute immediately prior to the shock. The large trade, occurring at t_0, has an immediate adverse effect on liquidity, yet this effect lasts only until briefly after the event.
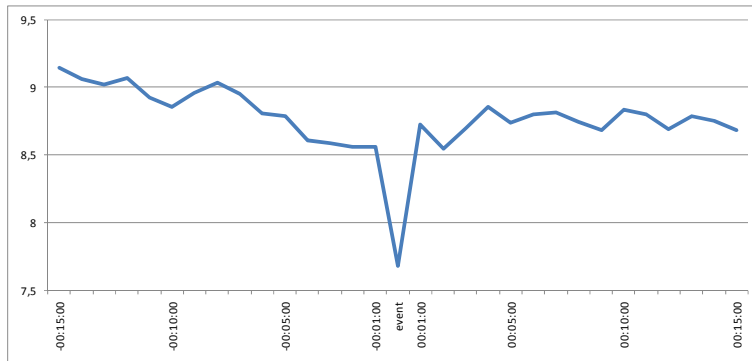

Figure 1.   *Time path of liquidity (XLM) on ENP with one minute observation frequency over the entire event window*

Figure 2 is a graphical representation of the same analysis. The difference here is that the interval [0:1] is depicted with an observation frequency of one second (orange), while the remaining data points remain depicted in one-minute frequency (blue)[7]. It is obvious that in fact the adverse effect on liquidity is much more pronounced than was visible before – and that at the same time much of it is being absorbed *within* the first minute following the event and would "go lost" if analyses were only conducted on a one-minute-interval basis.
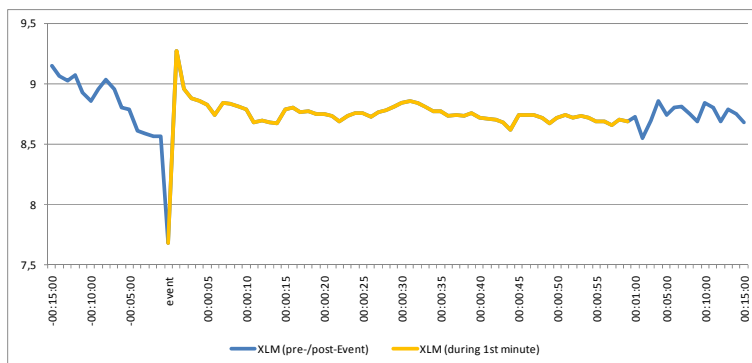

Figure 2.   *Time path of liquidity (XLM) on ENP with one second observation frequency during interval [0:1]*

The following Figure 3 shows the same detailed graph of the time path of liquidity on Chi-X, while Figure 4 represents TQ. In the case of Chi-X, the average values of the liquidity measure XLM over the entire observation period are nearly in the same order of magnitude (~11.51 bps) as on ENP (~8.76

---

nuously monitoring every stock market to literally immediately detect any potential opportunities from temporary imbalances or inefficiencies. It is therefore unlikely today that in an interval of up to nearly one minute, only little changes occur in the order book because "traders need time to observe and to decide what to do".

[7] The horizontal axis again measures the time. The 15-minute interval prior to the event is depicted with an observation frequency of 1 min.; the one-minute interval subsequent to the liquidity shock is depicted with an observation frequency of 1 sec.; and the 14 minutes following the first post-event minute is again depicted with an observation frequency of 1 min.

bps), while there is a significant difference in size concerning TQ (~34.88 bps)[8]. This means that the costs a round-trip transaction of 0.1m EUR size in any of the sample stocks entails on average on TQ are three times higher than on Chi-X and almost four times higher than on ENP.
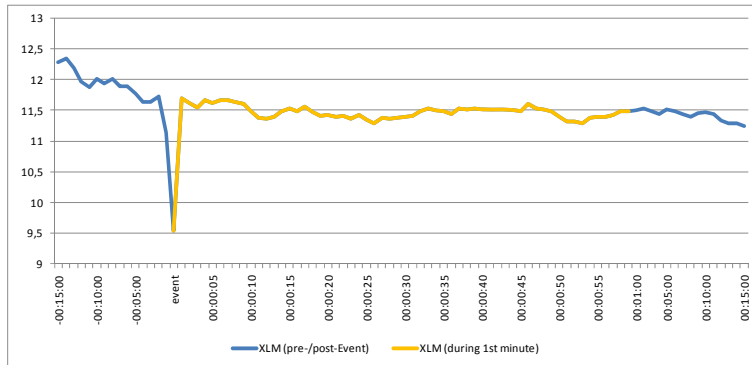

Figure 3.   *Time path of liquidity (XLM) on Chi-X with one second observation frequency during interval [0:1]*

In terms of comparing the resiliency of these three markets, what is far more interesting than the mere magnitude of the liquidity measure is the fact that the liquidity charts themselves differ considerably. What we observe on ENP (Figure 2) was to be expected against the background of previous research in this area: first, large trades are timed to situations where liquidity is extraordinarily high; second, the shock negatively affects liquidity; third, in the aftermath of the shock liquidity recovers to its normal level. The important difference in our results is that this recovery occurs far quicker than in earlier studies. Figure 3 (showing Chi-X) differs distinctly from this "standard" sequence: while again a liquidity increase precedes the shock (which confirms timing of trades), the market seems to be capable of exploiting very precisely the improved liquidity without any overshooting. An explanation for this behavior is not immediately evident; it might to some extent be related to the fact that Chi-X is an institutional investors only market and constitutes another starting point for further research.
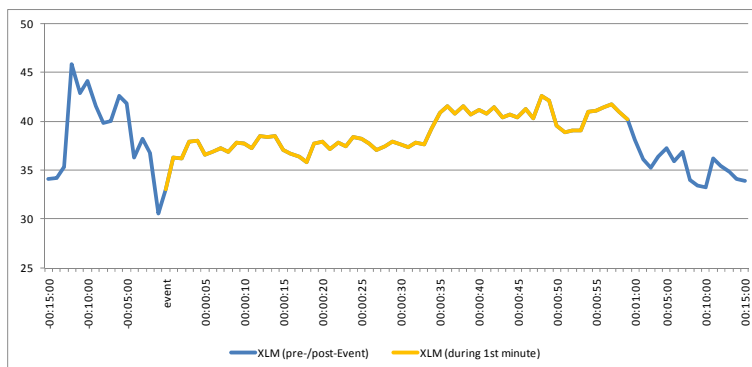

Figure 4.   *Time path of liquidity (XLM) on TQ with one second observation frequency during interval [0:1]*

The Turquoise chart, by contrast, depicted in Figure 4 above, exhibits a completely different, seemingly erratic, appearance. Here, the adverse liquidity effect does not peak immediately after the shock, but rather seems to increase *within* the first post-event minute. Hitting its highest value at 00:00:49 past the shock, the liquidity measure afterwards declines (i.e. recovers) again. There are two possible explanations for this behavior: first, compared with the two other markets in this analysis, TQ exhibits a considerably lower market share, which entails generally less order-flow to replenish the book. Second, the lack of an explicit immediate adverse effect could possibly be traced back to the fact that in relation to

---

[8] These values correspond to relative liquidity costs for a virtual 100,000 EUR round-trip transaction of 87.60 EUR on ENP, 115.10 EUR on Chi-X and 348.80 EUR on TQ, respectively. The measure hence not only allows for the comparison of securities according to their trading costs, but also for an international comparison of the empirical order book liquidity.

all trades on that venue, those trades identified as "large in scale" for the purpose of this analysis are not sufficiently differentiated from all other trades. This constitutes once more a potential source for a more detailed future analysis.

In light of these insights and returning to statistical analysis, we shall now switch to the "adapted model", which means that we will drill into the one minute immediately after the event. The following Table 3 presents the results of the adapted model. The grayed-out columns (labeled [0:1s] and [0:5s] respectively) show the size of the liquidity shocks and the corresponding t-statistics *immediately* (one second) and five seconds after the event. Compared with column [0:1], which repeats the results from the general model, we can state that for both ENP and Chi-X the adverse movement is much more pronounced in the one-second interval than in the one-minute interval. The recovery rate of liquidity (in terms of the XLM) on ENP amounts to 33% within five seconds after the shock and to 35% within sixty seconds after the shock. After that (e.g., until three or fifteen minutes past the shock), recovery is negligible. On Chi-X, the recovery rate of liquidity is 7% within five seconds after the shock and 12% within sixty seconds after the shock. Further recovery until three minutes after the shock is negligible, while after fifteen minutes, a total recovery rate of 27% can be observed. In relation to TQ, the results differ: here, the adverse liquidity effect does not peak at [0:1s], but rather amidst the first-minute interval (at t=+49 sec.). Still, like with ENP and Chi-X, liquidity again recovers significantly within the first minute after the event. What is also interesting is that XLM and quoted volumes strongly behave inversely proportional (not depicted). While hardly surprising, as much of what runs into the calculation of the XLM measure is the top of the order-book, this is yet another sign that transactions are precisely timed by market participants to occur when liquidity in the market is extraordinarily high.

Euronext Paris	[-15:0]	[0:1s]	[0:5s]	[0:1]	[0:3]	[0:15]
Δ quoted spread	-0.57***	0.38***	0.39***	0.43***	0.32***	0.33***
	-8.04	6.47	6.69	7.32	5.42	5.35
Δ XLM (100)	-1.486***	1.587***	1.059***	1.040***	1.010***	0.997***
	-11.66	15.62	9.33	9.62	8.97	8.65
Chi-X						
	[-15:0]	[0:1s]	[0:5s]	[0:1]	[0:3]	[0:15]
Δ quoted spread	-0.36***	0.31***	0.31***	0.23***	0.24***	0.19***
	-6.21	5.27	5.66	4.18	4.34	3.22
Δ XLM (100)	-2.772 ***	2.180***	2.031***	1.922***	1.931***	1.734***
	-15.30	16.69	16.44	13.89	13.85	11.98
Turquoise						
	[-15:0]	[0:1s]	[0:5s]	[0:1]	[0:3]	[0:15]
Δ quoted spread	-0.17	0.33***	0.40***	0.52***	0.85***	0.49***
	-1.55	4.91	4.21	5.09	5.84	3.83
Δ XLM (100)	-6.288***	3.734***	4.023**	5.019***	3.483**	3.217***
	-4.17	4.35	2.79	3.27	2.32	2.43

Table 3.   *Event study results ("Adapted model" specification). For legend refer to Table 2.*

5.   CONCLUSION / FURTHER RESEARCH

In this analysis, we investigate how different electronic order-driven markets behave in terms of their resiliency after a large endogenous liquidity shock. We can confirm earlier findings that large trades, which trigger the liquidity shock, are timed by investors as they occur when liquidity in the market is exceptionally high. These large trades have an immediate adverse effect on liquidity, yet this effect lasts only until briefly after the event. Most of what happens in consequence of a shock (adverse liquidity movement and subsequent recovery) takes place within the first-minute interval after the trade occurred; any visible effect already diminishes significantly towards the end of the first post-event minute, whereas this seems to depend on the market a stock is traded on. Three minutes later, basically no effect from the liquidity shock remains observable. We attribute this finding to the fact that temporary order-book imbalances are immediately detected by algorithmic trading engines and high-frequency traders who are continuously eagle-eying the stock markets for such opportunities.

This study may be potential starting points for further research: considering the increasing usage of smart order routing technologies at the sell side's trading desks to identify best-suited execution venues, it is likely that when a large trade occurs in one market, another large trade also occurs (at more or less the same instant) in another market. A possible extension could therefore be to examine whether this is the case and by what factors this timing and slicing of orders is influenced. As this study represents a short-term analysis, another interesting perspective could be to analyze the recovery of liquidity over a long-term period in order to see how resilient one market is over time and across various market phases (heavy trading, different volatility levels, etc.).

References

Bernstein, P. (1987). Liquidity, Stock Markets, and Market Makers. Financial Management 16, 54-62.

Black, F. (1971). Toward a Fully Automated Stock Exchange. Financial Analysts Journal 27, 28-44.

Clemons, E. and Weber, B. (1992). Turmoil, Transparency, and Tea: Valuing the Impact of IT on London's Stock Exchange. In: R. Kauffman, R. Banker, and M. Mahmood, eds., The Strategic and Economic Value of Information Technology Investments. Idea Group Publishing, 202-235.

Coppejans, M., Domowitz, I. and Madhavan, A. (2003). Resiliency in an automated auction. Unpublished working paper. Barclays Global Investors, ITG Inc.

Davydoff, D., Gajewski, J.F., Gresse, C. and Grillet-Aubert, L. (2002). Trading Costs Analysis: A Comparison of Euronext Paris and the London Stock Exchange. OEE Research Paper.

Degryse, H., de Jong, F., van Ravenswaaij, M., and Wuyts, G. (2005). Aggressive Orders and the Resiliency of a Limit Order Market. Review of Finance 9 (2), 201-242.

Dong, J., Kempf, A. and Yadav, P.K. (2007). Resiliency, the Neglected Dimension of Market Liquidity: Empirical Evidence from the New York Stock Exchange. European Finance Association.

Easley, D., Hendershott, T. and Ramadorai, T. (2009). Levelling the Trading Field. Working Paper.

Foucault, T., Kadan, O. and Kandel, E. (2005). Limit Order-book as a Market for Liquidity. Review of Financial Studies 18 (4), 1171–217.

Garbade, K.D., (1982). Pricing Corporate Securities as Contingent Claims, MIT Press.

Gomber, P. and Schweickert, U. (2002). The Market Impact - Liquidity Measure in Electronic Securities Trading. Working Paper.

Gomber, P., Schweickert, U. and Theissen, E. (2004). Zooming in on Liquidity. EFA 2004, Maastricht.

Harris, L. (1990). Liquidity, trading rules and electronic trading systems. NYU Salomon Center Series in Finance and Economics, Monograph 1990-1994.

Hasbrouck, J. and Saar, G. (2011). Low-Latency Trading. Johnson School Research Paper 35-2010.

Hendershott, T., Jones, C.M. and Menkveld, A.J. (2011). Does Algorithmic Trading Improve Liquidity? The Journal of Finance 66 (1), 1-33.

Hitt, L. and Brynjolfsson, E. (1996). Productivity, profit and consumer welfare: three different measures of information technology value. MIS Quarterly 20 (2), 121-142.

Holthausen, R.W., Leftwich, R.W. and Mayers, D. (1987). The Effect of Large Block Transactions on Security Prices. Journal of Financial Economics 19, 237-367.

Irvine, P.J., Benston, G.J. and Kandel, E. (2000). Liquidity Beyond the Inside Spread: Measuring and Using Information in the Limit Order-book. WP, Emory University and Hebrew University.

Kohli, R. and Grover, V. (2008). Business value of IT: An essay on expanding research directions to keep with the times, Journal of the Association for Information Systems, 9 (1), 23-39.

Kyle, A. (1985). Continuous auctions and insider trading. Econometrica, 53, 1315-1335.

Large, J. (2007). Measuring the resiliency of an electronic limit order-book. J. of Fin. Markets 10, 1-25.

Lee, R. (2002). Capital Markets that Benefit Investors – A survey of the Evidence on Fragmentation, Internalization and Market Transparency. Oxford Finance Group.

Massimb, M.N. and Phelps, B.D. (1994). Electronic Trading, Market Structure and Liquidity. Financial Analysts Journal, 50 (1), 39-50.

Melville, N., Kraemer, K. and Gurbaxani, V. (2004). Information Technology and Organizational Performance: An Integrative Model of IT Business Value. MIS Quarterly, 28 (2), 283-322.