

Loriana Pelizzon | Satchit Sagade | Katia Vozian

Resiliency: Cross-Venue Dynamics with Hawkes Processes

SAFE Working Paper No. 291

Leibniz Institute for Financial Research SAFE
Sustainable Architecture for Finance in Europe

info@safe-frankfurt.de | www.safe-frankfurt.de

Electronic copy available at: <https://ssrn.com/abstract=3711976>

Resiliency: cross-venue dynamics with Hawkes processes

Authors*: Loriana Pelizzon Satchit Sagade Katia Vozian

WORK IN PROGRESS – COMMENTS WELCOME.

First version: May 15th, 2019

This version: September 16th, 2020

Abstract: Market fragmentation and technological advances increasing the speed of trading altered the functioning and stability of global equity limit order markets. Taking market resiliency as an indicator of market quality, we investigate how resilient are trading venues in a high-frequency environment with cross-venue fragmented order flow. Employing a Hawkes process methodology on high-frequency data for FTSE 100 stocks on LSE, a traditional exchange, and on Chi-X, an alternative venue, we find that when liquidity becomes scarce Chi-X is a less resilient venue than LSE with variations existing across stocks and time. In comparison with LSE, Chi-X has more, longer, and severer liquidity shocks. Whereas the vast majority of liquidity droughts on both venues disappear within less than one minute, the recovery is not lasting, as liquidity shocks spiral over the time dimension. Over half of the shocks on both venues are caused by spiralling. Liquidity shocks tend to spiral more on Chi-X than on LSE for large stocks suggesting that the liquidity supply on Chi-X is thinner than on LSE. Finally, a significant amount of liquidity shocks spill over cross-venue providing supporting evidence for the competition for order flow between LSE and Chi-X.

JEL Classification: G10, G14

Keywords: liquidity, resiliency, fragmentation, competition, high-frequency data, Hawkes processes

* We are particularly grateful to Michael Schneider, Monika Gehde-Trapp, Peter Sarlin, Massimiliano Caporin, Lars Nordèn (discussant), Iivo Vehviläinen (discussant) for discussion and insightful comments. We are also thankful to participants at the Regulation and Operation of Modern Financial Markets 2019 conference in Reykjavik, Iceland, the Trans-Atlantic Platform Digging into Data Round 4 conference in Alexandria, U.S.A., and seminar participants at Deutsche Börse, Autorité des Marchés Financiers, Norwegian University of Science and Technology, Nordic Finance Network, and Helsinki Graduate School of Economics. Naturally, we are the sole responsible for errors and omissions. We gratefully acknowledge financial support from the "Digging into High Frequency Data" project funded by the Trans-Atlantic Platform for the Social Sciences and Humanities. Loriana Pelizzon and Satchit Sagade acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project number 329107530. The authors also thank EUROFIDAI for providing the BEDOFIH database.

Authors' affiliation: Loriana Pelizzon: Ca' Foscari University of Venice, Dorsoduro 3246, 30123, Venezia, Italy, Goethe University and Research Center SAFE, Theodor-W.-Adorno-Platz 3, 60323, Frankfurt am Main, Germany, pelizzon@safe.uni-frankfurt.de; Satchit Sagade: Goethe University and Research Center SAFE, Theodor-W.-Adorno-Platz 3, 60323, Frankfurt am Main, Germany, sagade@safe.uni-frankfurt.de. Katia Vozian: Hanken School of Economics and Helsinki GSE, Arkadiankatu 22, 00100 Helsinki, kvoziaan@hanken.fi.

1 Introduction

The contemporary equity trading landscape is best described by a competitive fragmented market with trading occurring on multiple trading venues¹. Market fragmentation is the result of investors with different preferences looking for the trading venue that can cater best to their preferences. For instance, uninformed passive investors, such as index funds, prefer block orders, while informed active investors, such as mutual funds, prefer order splitting strategies with low price impact. Regulatory initiatives (e.g. MiFID in the European Union) and technological advances increasing the speed of trading fuelled further the transition towards the current competitive equity trading landscape. The modern setup has several implications. Namely, how a trading venue is organized affects investors' trading costs as well as their portfolio choices. Consequently, trading costs and portfolio choices drive asset prices, which affect cost of capital, carrying welfare implications (Stiglitz, 1989) as well as market quality implications, including market resiliency. Market resiliency refers to the phenomenon of return to normal following a deviation. Knowing the properties of market resiliency in modern competitive fragmented markets becomes even more important for market participants and regulators. This paper studies market resiliency across competing venues of a modern fragmented equity market when liquidity is scarce.

We examine how resilient are modern competing trading venues to severe intraday liquidity shocks. Our analysis is motivated by two questions that are not yet sufficiently explored in the empirical literature. First, how fast does the liquidity recover on different trading venues following an intraday liquidity shock? Second, is the recovery only momentary considering that intraday liquidity shocks may repeat, spiralling within a venue and spilling over across venues? To this purpose, we build a unique dataset sourced from exchange feeds of two major competing trading venues: London Stock Exchange (LSE), a traditional venue, and Chi-X Europe, an alternative venue. Using this cross-venue dataset with millisecond resolution, we measure the threshold exceedance duration, i.e. TED, defined as the time for liquidity to recover following a shock. We employ a multivariate Hawkes point process methodology to uncover whether the recovery is only temporary, as shocks might propagate across time, forming spirals, or across venues, forming spillovers.

Our findings show that in 99% of cases liquidity gets back to normal within approximately one minute following a severe liquidity shock. This holds true irrespective of the type of venue yet varies by stock size. Zooming in reveals that the recovery is, however, only momentary, as more than half of the shocks propagate through time on the same venue and around 12% of shocks propagate cross-venue. In comparison with LSE, more frequent, longer, and more severe liquidity shocks were observed at Chi-X. In addition, we noted short instances of extreme illiquidity on Chi-X, with no ask or no bid quotes. Our findings suggest that on average Chi-X has lower market resiliency than LSE for FTSE 100 stocks, yet we do find variations of market resiliency across stocks and across time. The microfoundations of

¹ The term "trading venue" is used to designate individual exchanges and other venues on which a security of interest is traded. We use the term "market" to designate the aggregate trading activity over all trading venues where a security of interest is traded.

these findings might be related to different market models operated by these venues, trading mechanisms, and competition dynamics.

Our study contributes to the empirical literature on market resiliency, and ultra-fast market dynamics, as we aim to fill the gap of the empirical literature on cross-venue market resiliency in the high-frequency environment. With respect to the existing literature on market resiliency, our analysis has several distinctive features. The unique characteristics of our dataset support the proposed analysis thanks to (i) the time increment of one millisecond enabling a study of resiliency at a high-frequency resolution not done before in the literature, (ii) the timestamp accuracy coming from exchange feeds, which is critical for a reliable measurement of resiliency and cross-venue analysis. Moreover, our sample covers two major competing trading venues with different market microstructure for the same set of stocks over the same time span, with LSE as the primary listing exchange and Chi-X Europe as the alternative venue, accounting together for approximately 80%¹ of the equity market of FTSE 100 stocks during the observation period. This setup enables us to reliably compare the dynamics of venues having different market microstructure. Finally, in comparison with the existing literature, our methodology captures both the speed of recovery and how lasting it is. We do this by identifying potential propagation of shocks within a venue as well as across venues.

The rest of the paper is organized as follows. Section 2 presents the related literature. Section 3 presents the institutional structure of the two trading venues that we study, LSE and Chi-X Europe, and describes the data employed for the study. Section 4 describes the Hawkes methodology and its application for studying market resiliency. Section 5 presents and discusses the results. Section 6 documents the significance of parameter estimates and the robustness of our results. Section 7 concludes.

2 Related literature

Market resiliency subscribes to the literature on system resiliency. System resiliency refers to the system's "ability to anticipate, to absorb, to adapt, and to recover from a potential disruptive event" (Francis & Bekera, 2014). In the seminal literature on market microstructure, the term market resiliency has been first used by Kyle (1985) as a transactional property of markets that characterizes liquidity: "'Market liquidity' is a slippery and elusive concept, in part because it encompasses a number of transactional properties of markets. These include 'tightness' (the cost of turning around a position over a short period of time), 'depth' (the size of an order flow innovation required to change prices a given amount), and 'resiliency' (the speed with which prices recover from a random, uninformative shock)." Whilst the first two properties, i.e. tightness and depth are extensively studied in the literature, resiliency received less attention.

¹ This figure is computed based on trading volume data in shares from Reuters Eikon covering LSE, Chi-X, BATS, and Turquoise for 94 stocks of the FTSE 100 for the first quarter of 2015.

The concept of market resiliency evolved in the theoretical and empirical literature since its introduction in 1985. Foucault, Kadan, & Kandel (2005) measure market resiliency by “the probability that, after a liquidity shock, the spread reverts to its former level before the next transaction”. Obizhaeva & Wang (2013) understand market resiliency as the refresh rate of the limit order book, i.e. bid-ask updates, following a trade and the speed of convergence of the price, i.e. mid-quote, to a new steady state level following a trade. This contribution is important since it shows market resiliency as a reorganization capacity of the market, rather than as a recovery to a unique static steady state level. The empirical literature on market resiliency starts with the study of Degryse, De Jong, Van Ravenswaaij, & Wuyts (2005). The authors take a non-parametric event-study approach and measure market resiliency as the number of best limit updates, i.e. changes in the best bid or best ask required for the spread and depth to return to average following a shock and use data of 20 stocks on Paris Bourse. The study finds that following an aggressive order it takes around 20 best limit updates for the spread and depth to return to their initial level prior shock. Large (2007) uses a parametric approach where arriving orders and cancellations are viewed as a mutually-exciting 10-variate Hawkes point process with application on data for one stock, Barclays on LSE, with one-second snapshot frequency. Using this method, the author estimates two metrics that are used to measure market resiliency: (i) the probability of order replenishment after a liquidity demand shock (aggressive order) and (ii) the time it would take for the limit order book, i.e. LOB to replenish. The author finds that the LOB replenishes resiliently, i.e. reaching the level prior shock, less than 40 per cent of the time. In those cases where resilient replenishment does occur, the LOB replenishes quite fast, displaying a half-life of under 20 seconds. Gomber, Schweickert, & Theissen (2015) define market resiliency as the time for the liquidity metric XLM to return to the XLM level observed prior a shock induced by large transactions or Bloomberg ticker news items and empirically measure it using data for 21 stocks on XETRA with one-minute snapshot frequency. The study concludes that shocks pertaining to large transactions affect liquidity, yet resiliency is high such that liquidity returns quickly to the market. Kempf, Mayston, Gehde-Trapp, & Yadav (2015) understand market resiliency as the speed of return to the long-run average of spread and depth following a liquidity shock. Using an extended Ornstein-Uhlenbeck methodological approach and data covering 120 stocks of FTSE 100 on LSE with five-minute snapshot frequency, the authors find that the measures of market resiliency stay relatively stable over time, including in the period of the Global Financial Crisis when liquidity sunk. However, the marginal difference between consumption resilience and replenishment resilience was much higher during the Global Financial Crisis than in normal times. Among further findings, the authors document that algorithmic trading is associated with higher resiliency for the majority of stocks, yet less so for smaller stocks. The measurement approach of speed of return proposed by Kempf et al. (2015) has been employed in further empirical studies by Colliard & Hoffmann (2017) and Fález-Viñas (2019). Finally, Danielsson, Panayi, Peters, & Zigrand (2018) measure market resiliency as the Threshold Exceedance Duration, i.e. TED, that describes the time needed for a stock to recover following a liquidity shock of selected severity and employ an autoregressive survival model for detecting and forecasting Liquidity Resilience Profiles, i.e. LRP. The appeal of this resiliency measure is its modelability for different definitions of liquidity shocks while

preserving the validity of the drivers of TED. The study uses data of 40 stocks of CAC40 traded on Chi-X and recorded with one-second snapshot frequency. The authors find their measure of market resiliency to be time-varying and in particular, TED varies with the current and past state of the LOB and of the market. This methodology has been further employed in the studies on liquidity commonality and liquidity resilience by Panayi, Peters, & Kosmidis (2015), and on liquidity resilience and market making by Panayi, Peters, Danielsson, & Zigrand (2018).

The above review of the literature on the measurement of market resiliency shows that there is no standard way of measuring market resiliency. In particular, the cross-study comparability of the findings is limited due to (i) different definitions, (ii) different venues and assets underlying the studies, and (iii) the use of data with different frequency ranging from five minutes down to one-second snapshots. We consider these aspects in the design of our cross-venue analysis as follows. Whilst we are agnostic to the different types of measures discussed in the literature, we chose to apply a Hawkes point process approach with a counting process for liquidity shocks defined in line with the Threshold Exceedance Duration (TED). As shown by Panayi, Peters, & Kosmidis (2015), TED is compatible with ultra-high-frequency data. Hawkes processes are particularly well suited for capturing the dynamics of high-frequency data. We apply this methodology consistently on LSE and on Chi-X data for the same set of assets, i.e. FTSE 100 stocks over the same time period, i.e. Q1 2015 on data with one-millisecond snapshot frequency. Our approach has the advantage of capturing not only how quickly the market returns to the pre-shock level, but also whether that return was indeed meaningful or just fleeting.

Since our research interest lies in the cross-venue dynamics of a fragmented market, we further note the study of Féllez-Viñas (2019). She finds that market fragmentation is generally beneficial for market resiliency both under normal market circumstances as well as under stress conditions induced by a large trade or large information asymmetry. The author reconstructs the consolidated order book at one-minute snapshot frequency for 30 Spanish stocks of the IBEX 35 traded on BME, Bats, Chi-X, and Turquoise throughout May 2013 – February 2015 and studies market resiliency before the fragmentation of order flow and after. The study measures market resiliency as the speed of recovery of the relative quoted spread or of the depth as first proposed by Kempf et al. (2015). Relative to the work of Féllez-Viñas (2019), our study focuses on the cross-venue dynamics of market resiliency in the subsecond environment, rather than on the impact of market fragmentation on resiliency.

Finally, our work contributes to the literature on ultrafast market dynamics attempting to describe subsecond phenomena. In the presence of machines taking trading decisions, computational operating times used on trading venues go below normal human reaction times. As shown by Wolfe, Seppelt, Mehler, Reimer, & Rosenholtz (2020), human reaction times are in the range of 600 milliseconds in circumstances of a road hazard where decision is prompted. Johnson et al. (2013) find in price changes ultrafast extreme events (UEEs) with a duration below one second, i.e. beyond human response time, and show that these events have different characteristics than UEEs with a duration of above one second, i.e. within the human response time. Using a crowd model and testing it empirically on data, Johnson et al. (2013) show that UEEs are more frequent when the probability of having n agents

simultaneously using the same strategy is high and vice versa. We further note the studies of Cartledge, Szostek, De Luca, & Cliff (2012), Golub, Keane, & Poon (2012) and Menkveld (2018). Relative to this literature, our work focuses on market resiliency using one-millisecond snapshot frequency. This resolution allows us capturing durations of severe intraday liquidity shocks that last shorter than typical human response times and are likely driven by machine ecology as described by Johnson et al. (2013) and Beddington, J., Bond, P., Cliff, D., Houstoun, K., Linton, O., Goodhart, C., & Zigrand (2012).

3 Institutional structure and data

3.1 Market structure

The London Stock Exchange operates as an order-driven market with a limit order book. On the LSE, continuous trading starts at 8:00 a.m. and ends at 4:30 p.m. GMT. This phase is preceded by the opening auction, which starts at 07:50 a.m. and lasts for ten minutes, plus a random end time of up to 30 seconds. Continuous trading is followed by a closing auction, which starts at 4:30 p.m. and lasts for five minutes, plus a random end time of up to 30 seconds. Furthermore, the venue holds Exchange Delivery Settlement Price (EDSP) intraday auctions for the FTSE 100 Futures and Options Contracts on the third Friday of each month and Automatic Execution Suspension Period (AESP) auction calls that act as circuit breakers based on price monitoring. While EDSP auctions last about five minutes, AESP auctions generally last for five minutes, plus a potential extension of up to seven minutes. The market model of the London Stock Exchange for FTSE 100 stocks relies on the provision of liquidity by market makers. Market makers typically achieve this by submitting limit orders on both sides of the market, i.e. bid and ask, delivering guaranteed 2-way prices. Furthermore, the venue has designated specialists that are in charge of supporting market quality. LSE uses a standard symmetric pricing scheme where the liquidity maker and the liquidity taker pay the same level of fee per transaction. LSE has a downstairs market and upstairs market running in parallel. On the downstairs market, orders are executed through the stock exchange electronic trading system (SETS). On the upstairs market, orders are executed through broker-dealers, away from SETS limit order book and the resulting trades are reported to the SETS within three minutes of execution. This arrangement enables a delayed disclosure of large institutional trades.

Chi-X Europe is a London-based Recognised Investment Exchange (RIE). It was initially set up by a group of banks in 2007 as a Multilateral Trading Facility (MTF) aiming to compete with LSE and other incumbent venues following the change in regulation opening exchanges to competition. Chi-X Europe received a full exchange status in May 2013. In 2017, the venue has been acquired by CBOE. The trading session on Chi-X Europe starts at 8:00 a.m. and ends at 4:30 p.m. GMT. The venue does not have an opening, closing, EDSP, or AESP auction. Instead of circuit breakers, Chi-X Europe uses price collars pegged to the prices revealed on the incumbent exchange, which is the LSE for the FTSE 100 stocks of our interest. There are no market makers and designated specialists alike the ones on LSE. Chi-X Europe uses an asymmetric maker-taker fee structure where the liquidity maker and liquidity

taker pay different levels of fees. The maker receives a rebate whilst the taker gets charged. This structure aims to incentivize liquidity provision.

London Stock Exchange and Chi-X Europe are the two major competing trading venues for FTSE 100 stocks with different market microstructure. During the period of our study, there are no market microstructure, fee, or trading system changes on these venues. Both exchanges operate an electronic limit order book and have visible and hidden orders. Our analysis focuses on the visible order book. LSE and Chi-X Europe compete, among other things, on: (i) trading fee structures, e.g. different maker-taker pricing on LSE and Chi-X, and (ii) reduction in latency¹. Such differentiations explain the type of market participants that a venue attracts. As documented by Menkveld (2013), the participation rates of High-Frequency Traders (HFT) on an incumbent venue differ from those on alternative venues. Using data from January 2007 through June 2008, a period when Chi-X entered the market as an alternative venue, the author finds an HFT participation rate of 8.1% on NYSE Euronext versus 64.4% on Chi-X. Menkveld (2014) suggests that “new markets [i.e. alternative venues such as Chi-X] serve HFTs who seek low fees and high speed”. As such, whereas HFTs are likely present on both venues, one may conjecture that Chi-X as an alternative venue has more HFT activity, whereas LSE as an incumbent venue would likely attract more natural liquidity traders, e.g. institutional investors such as pension funds. In 2011, LSE introduced the Millennium Technology, enabling ultra-low latency connectivity and allowing the exchange to compete with Chi-X on speed. The literature on competition for order flow between LSE and Chi-X provides further evidence. Ibikunle (2018) shows that the share of price discovery in relation to informed trading is predominantly higher on Chi-X than on LSE throughout the trading day. The study argues that Chi-X’s price leadership allows attracting order flow away from LSE. The author also documents a higher presence of HFT activity on Chi-X than on LSE. The study uses data for 47 stocks of the FTSE 100 over 106 days of the second half of 2014, which is close to the period of our sample and makes the findings particularly relevant for our study. An earlier study by Riordan, Storkenmaier, & Wagener (2012) also argues that Chi-X has more informed trading in comparison with LSE. The study shows that Chi-X and LSE lead in trade and quote based price discovery with Chi-X having a larger fraction of information impounded into prices (44.6%) compared to LSE (34.6%). The authors use for their study data covering the quoting and trading activity of FTSE 100 stocks on LSE, Chi-X, BATS, and Turquoise over April/May 2010.

3.2 Data

¹ The roundtrip connectivity latency from connectivity providers within Interxion London for LSE and Chi-X was 10 microseconds for LSE and 283 microseconds for Chi-X in 2015 (Delaney, Fenick, Tyc, Saade, & Marsh, 2015). The roundtrip connectivity latency refers to the time from the source to the destination plus the time from the destination back to the source, excluding the processing time at destination. Interxion London is a data centre financial hub that offers ultra-low latency connectivity through proximity hosting. Exchange hosting, in comparison with proximity hosting, offers even lower roundtrip times. Moreover, differences in processing times, as well as differences in trade publishing times in the market data stream, contribute further to differences in speed from one venue to another.

We reconstruct the limit order book for each stock of our sample using visible messages from exchange feeds. This data is made available by Base Européenne des Données Financières à Hautes Fréquences (BEDOFIH). Our sample covers stocks composing the FTSE 100 index as per its composition on December 31, 2014, and covers two trading venues: London Stock Exchange and Chi-X Europe. The data resolution is one millisecond and the time period of the dataset is January 2, 2015, to March 31, 2015, i.e. 63 trading days.

We remove stocks that were subject to corporate splitting or that were delisted during the sample period. In total, 99 stocks are retained in the sample. We remove opening, closing, EDSP (297 instances observed for 99 stocks), and AESP auction phases (14 instances observed for 14 stocks) that occur on LSE and retain for our study only phases of continuous trading on LSE. We also remove these time segments from Chi-X Europe data for a consistent cross-venue comparison.

Our dataset has several unique features important for our analysis. The data has a time increment of one millisecond, allowing analysis at high frequency resolution. In addition, exchange feeds have a high timestamp accuracy, which is crucial for the validity of such analysis. Thanks to the granularity of the data, each event of the limit order book can be clearly seen. The period of our dataset, i.e. first quarter of 2015, is clean of significant market microstructural changes or high volatility events that could drive results. Finally, the sample covers two trading venues with the same set of stocks, yet with different market microstructure, creating an accurate setup for a cross-venue study for market resiliency. The list of stocks composing our sample and the subsamples of large, mid, and small capitalisation stocks used for enabling comparison with the existing literature are omitted here for brevity but available in the Appendix.

We describe the trading activity on LSE and Chi-X in terms of traded volume, quote updates, and spread levels. During the time period of our interest, the total traded volume on LSE and Chi-X accounted together for approximately 80% of the total traded volume on LSE, Chi-X, BATS, and Turquoise taken together. In average for one stock over the 63 trading days of our sample, the traded volume in number of shares was 4.98 million on LSE and 1.49 million on Chi-X, indicating a ratio of 3 to 1. Figure 1 shows the daily traded volume in millions of shares for the time period of our sample. For all market capitalisation subsamples and consistently for each day of our sample, the traded volume on LSE was substantially higher than the traded volume on Chi-X. The traded volume of stocks of firms within the highest market capitalisation bin account on a daily basis for roughly half of the total traded volume for all stocks of our sample.

Insert Figure 1 here

We describe the liquidity of the two venues in terms of the relative quoted spread. Figure 2 plots the 50th and 95th percentile of the hourly relative quoted spread distribution, averaged across all stocks of the sample. In average, consistently throughout all the days of our sample, the spread values are higher on Chi-X than on LSE. Whereas the values of the average daily median spread are relatively similar cross-venue, the values of the 95th percentile are more apart, indicating that Chi-X has more

extreme values of the spread than LSE. Consistently throughout all the hours of continuous trading, the average values of the 50th and 95th percentiles of the hourly spread distribution on Chi-X are higher than the respective values on LSE. We document intraday patterns for the bid-ask spread in line with the literature: highest in the morning (Biais, Hillion, & Spatt, 1995) and then decreasing throughout the day with a slight peak around noon (Tannous, Wang, & Wilson, 2013).

Insert Figure 2 here

We find on Chi-X instances of infinite spread where there is no best ask or best bid quote. These instances do not occur only in the hour of market opening following the LSE opening auction, but also throughout the trading day. Furthermore, this phenomenon is observed across several stocks and is therefore not particular to a specific stock. The absence of quote may be momentary or last for several milliseconds in a row. Table 2 shows the time duration of infinite spreads throughout the hours of continuous trading. For instance, within the trading hour from 9 a.m. to 9:59 a.m. there were in aggregate 44 minutes with infinite spread occurring in six stocks of the FTSE 100 stocks, which are actually a set of the most liquid instruments. We do not find such instances on LSE. We conjecture that this fact may be due to differences in market microstructure between Chi-X and LSE, particularly the market-making model charged with liquidity provision.

Insert Table 2 here

For all stocks of the sample there are instances of infinite spread in the time segment from 8 a.m. to 8:59 a.m. and the time segment 4 p.m. to 4:30 p.m. Furthermore, the average value of 95th percentile of the spread on Chi-X in the time segment 8 a.m. to 8:59 a.m. is far above the one on LSE. We conjecture that this effect is associated with the opening and closing auctions on LSE occurring in these hours. To exclude the eventuality that the results of our analysis of market resiliency are driven by such phenomena, we discard these time segments from our sample.

Higher trading volumes and narrower spreads at LSE do not necessarily mean price leadership and a larger fraction of informed trading on LSE over Chi-X. As argued by Ibikunle (2018) and Riordan, Storckenmaier, & Wagener (2012), the quoting activity on Chi-X is intense and the informativeness of quotes may attract order flow on Chi-X away from LSE. Therefore, as a proxy of quoting intensity, we report on Figure 3 the number of times that the relative quoted spread was updated each day. The order of magnitude of the number of updates on LSE and on Chi-X are very close to each other indicating a similar level of quoting intensity on both venues. Interestingly, we observe that the quoting intensity is not consistently higher on one venue in comparison with the other. On certain days of the sample, Chi-X has a higher quoting intensity than on LSE and the relation inverts on other days. This fact is indicative of market competition microfoundations where venues compete to attract order flow as documented by Ibikunle (2018) for LSE and Chi-X. Consistently through all days of the sample, stocks in the high market capitalisation bin show the highest level of spread updates and account for more than half of all spread updates observed in the sample.

Insert Figure 3 here

4 Methodology

To study empirically our research question on cross-venue dynamics of resiliency, we assume that the data follows a Hawkes point process and identify structural parameters from the data using maximum likelihood estimation. Our review of the empirical literature shows that there is no standard measure of market resiliency. Whilst we are agnostic to the different types of measures, we chose to apply a Hawkes point process given its suitability for high-frequency data. In this setup, liquidity shocks on LSE and liquidity shocks on Chi-X are two types of events described by a counting process with corresponding arrival intensity of such events. Modelling the arrival process of liquidity shocks with Hawkes processes permits the identification of both self-excitation, i.e. spirals¹ and cross-excitation, i.e. spillovers. Our methodology extends the measurement approach proposed by Schneider, Lillo, & Pelizzon (2018) with application to liquidity on the bond market and cross-asset spillovers. We make methodological adaptations to account for our interest in capturing market resiliency on the equity market instead of liquidity on the bond market. Furthermore, we make amendments to the methodology in order to study cross-venue spillovers instead of cross-asset spillovers.

First introduced by Hawkes (1971), Hawkes processes are point processes with a counting process of events and an associated vector conditional intensity process describing the arrival intensity of events considering the history of events. This type of processes has found application in numerous studies in finance as outlined in the survey by Bacry, Mastromatteo, & Muzy (2015) as shown particularly suitable for high-frequency data. Following the notation of Bowsher (2007) and Bacry et al. (2015), we write the D-variate counting process of events as $N(t) = \{N^i(t)\}_{i=1}^D$ with each element $N^i(t)$ counting the number of type i events that occurred up to and including time t . The complete observation of $N(t)$ up to and including time t for all types of events corresponds to the information set \mathcal{F}_t^N . Thus, the internal history – also known as natural filtration – of the D-variate point process $N(t)$ is $\{\mathcal{F}_t^N\}_{t \geq 0}$. The vector conditional intensity process associated with $N(t)$ is $\lambda(t) = \{\lambda^i(t)\}_{i=1}^D$ where $\lambda^i(t)$ is the \mathcal{F}_t^N -conditionally expected number of type i events per unit of time as the time interval tends to zero. The intensity vector for type i events can be written as:

$$\lambda^i(t) = \mu^i + \int_0^t \phi^{ii}(t-s) dN^i(s) + \sum_{j \neq i} \int_0^t \phi^{ij}(t-s) dN^j(s) \quad (1)$$

where μ^i is the exogenous baseline intensity and $\Phi(t) = \{\phi^{ij}(t)\}_{i \leq 1, j \leq D}$ is the response function corresponding to a $D \times D$ matrix-valued kernel, which for all elements $\phi^{ij}(t)$ is non-negative, causal, and L^1 -integrable. The kernel $\phi^{ij}(t)$ accounts for how process component j influences the occurrence

¹ We note that the definition of spiralling at ultra fast times that we employ here differs from the liquidity spirals described by Brunnermeier & Pedersen (2009) that arise due to the mutually reinforcing mechanism between an asset's market liquidity and the related traders' funding liquidity. Funding liquidity constraints are unlikely to change at the ultra short time durations that we study.

intensity of the process component i such that when $j = i$ this describes the self-excitation effect and when $j \neq i$ this describes the cross-excitation effect.

For the process to be asymptotically stationary, the spectral radius of $D \times D$ matrix $\Phi(t)$ made of the L^1 -norms $\|\phi^{ij}\| = \int_0^\infty \phi^{ij}(t)dt$ has to be strictly smaller than 1. In the presence of such a process, we may take the expectation value of the equation (1) to get

$$\bar{\lambda} = \mu + \gamma \bar{\lambda} \quad (2)$$

where $[\gamma]^{ij} = \int_0^\infty \phi^{ij}(t)dt$ is the norm of the kernel matrix $\Phi(t)$, $\bar{\lambda}$ the unconditional expectation of the arrival intensity of events, μ the exogenous baseline intensity. We write the equation (2) for a single element i and divide by $\bar{\lambda}^i$ to obtain the decomposable form

$$1 = \frac{\mu^i}{\bar{\lambda}^i} + \gamma^{ii} + \sum_{j \neq i} \gamma^{ij} \frac{\bar{\lambda}^j}{\bar{\lambda}^i} \quad (3)$$

where μ^i is the exogenous baseline intensity for element i , $\bar{\lambda}^i$ the unconditional expectation of the event arrival intensity for element i , $\bar{\lambda}^j$ the unconditional expectation of the event arrival intensity for element j , γ the norm of the kernel matrix Φ considering elements i and j . Each element on the right hand side of the equation (3) can be interpreted as follows: $\frac{\mu^i}{\bar{\lambda}^i}$ is the fraction of events due to baseline intensity for venue i , γ^{ii} is the fraction of events due to self-excitation intensity, i.e. spiralling for venue i , $\sum_{j \neq i} \gamma^{ij} \frac{\bar{\lambda}^j}{\bar{\lambda}^i}$ is the fraction of events due to cross-excitation intensity, i.e. spillovers from venue j to venue i . A higher γ^{ii} indicates that shocks on venue i for an individual stock are more likely to spread in time. This shows that the venue is less resilient over time.

A. Bivariate Hawkes process with exponential response function

We are interested in two types of events: liquidity shocks on LSE and liquidity shocks on Chi-X occurring in a stock. Narrowing down the multivariate Hawkes process to the case of the bivariate Hawkes process, i.e. $D = 2$, we re-write equation (1) for a bivariate Hawkes process with self-excitation and cross-excitation terms:

$$\begin{cases} \lambda^1(t) = \mu^1 + \int_0^t \phi^{11}(t-s)dN^1(s) + \int_0^t \phi^{12}(t-s)dN^2(s) \\ \lambda^2(t) = \mu^2 + \int_0^t \phi^{22}(t-s)dN^2(s) + \int_0^t \phi^{21}(t-s)dN^1(s) \end{cases} \quad (4)$$

For a bivariate Hawkes process, the kernel matrix $\Phi(t)$ takes the form:

$$\Phi(t) = \begin{pmatrix} \phi^{11}(t) & \phi^{12}(t) \\ \phi^{21}(t) & \phi^{22}(t) \end{pmatrix} \quad (5)$$

The Hawkes kernel is parametrized as an exponential kernel¹ so that the kernel components in equation (4) take the exponential form:

$$\phi^{ij}(t-s) = \sum_{k=1}^{p^{ij}} \alpha_k^{ij} e^{-\beta_k^{ij}(t-s)} \quad (6)$$

where the parameter α^{ij} can be interpreted as the amplitude of reaction in response to the occurrence of an event and the inverse of the parameter β^{ij} as the timescale of decay until the occurrence of the next event, such that the larger the β^{ij} the quicker the decay. The term P^{ij} gives the form of the exponential kernel, which we set as a double exponential kernel for the self-exciting components ($P^{ii} = 2$) and a single exponential kernel for the cross-excitation components ($P^{ij} = 1, i \neq j$) in line with the specification of Schneider et al. (2018). The double exponential kernel allows accounting for a fast and slow speed of propagation of events through the time dimension.

Considering the exponential parameterisation and the form of the kernel, i.e. double-exponential for self-excitation and single-exponential for cross-excitation, we can re-write (4) for data $\{t_i\}$ with $i = 1, 2, \dots, n$ and $\{t_j\}$ with $j = 1, 2, \dots, m$ corresponding to events of type 1 and type 2 respectively

$$\begin{cases} \lambda^1(t) = \mu^1 + \sum_{t_i < t} (\alpha_1^{11} e^{-\beta_1^{11}(t-t_i)} + \alpha_2^{11} e^{-\beta_2^{11}(t-t_i)}) + \sum_{t_j < t} \alpha_1^{12} e^{-\beta_1^{12}(t-t_j)} \\ \lambda^2(t) = \mu^2 + \sum_{t_j < t} (\alpha_1^{22} e^{-\beta_1^{22}(t-t_j)} + \alpha_2^{22} e^{-\beta_2^{22}(t-t_j)}) + \sum_{t_i < t} \alpha_1^{21} e^{-\beta_1^{21}(t-t_i)} \end{cases} \quad (7)$$

The stationarity condition for a bivariate Hawkes process requires that the matrix with entries equalling the L^1 -norms $\|\phi^{ij}\| = \int_0^\infty \phi^{ij}(t) dt$ of the 2×2 kernel matrix has a spectral radius greater or equal 1. With an exponential parameterization, the norm is given by

$$\gamma^{ij} = \sum_{k=1}^{p^{ij}} \frac{\alpha_k^{ij}}{\beta_k^{ij}} \quad (8)$$

and the matrix with entries equalling the L^1 -norms takes then the form

$$\begin{bmatrix} \frac{\alpha_1^{11}}{\beta_1^{11}} + \frac{\alpha_2^{11}}{\beta_2^{11}} & \frac{\alpha_1^{12}}{\beta_1^{12}} \\ \frac{\alpha_1^{21}}{\beta_1^{21}} & \frac{\alpha_1^{22}}{\beta_1^{22}} + \frac{\alpha_2^{22}}{\beta_2^{22}} \end{bmatrix} \quad (9)$$

Considering that the spectral radius of a square matrix $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ is the largest absolute value of its eigenvalues and takes the form $\frac{1}{2}(a+d + \sqrt{a^2 - 2ad + 4bc + d^2})$, the stationarity condition for our bivariate Hawkes process with exponential kernel becomes

¹ As shown by Bacry et al. (2015), most studies assume that the shape of the kernel components ϕ^{ij} has an exponential decay, some studies assume a power-law decay, and others take a non-parametric approach. We assume an exponential decay in line with the study of Schneider et al. (2018), Large (2007), and Muni Toke (2011a). Alike Muni Toke (2011a), we acknowledge the interest in exploring power-law shaped kernels and leave this to future research.

$$\left[\frac{\frac{\alpha_1^{11} + \alpha_2^{11} + \alpha_1^{22} + \alpha_2^{22}}{\beta_1^{11} + \beta_2^{11} + \beta_1^{22} + \beta_2^{22}}}{2} + \sqrt{\left(\frac{\alpha_1^{11} + \alpha_2^{11} + \alpha_1^{22} + \alpha_2^{22}}{\beta_1^{11} + \beta_2^{11} + \beta_1^{22} + \beta_2^{22}} \right)^2 - \left(\frac{\alpha_1^{11}}{\beta_1^{11}} + \frac{\alpha_2^{11}}{\beta_2^{11}} \right) * \left(\frac{\alpha_1^{22}}{\beta_1^{22}} + \frac{\alpha_2^{22}}{\beta_2^{22}} \right) + \frac{\alpha_1^{12}}{\beta_1^{12}} * \frac{\alpha_1^{21}}{\beta_1^{21}}} \right] < 1 \quad (10)$$

Finally, considering the above specifications, equation (3) becomes

$$\begin{cases} 1 = \frac{\mu^1}{\bar{\lambda}^1} + \gamma^{11} + \gamma^{12} \frac{\bar{\lambda}^2}{\bar{\lambda}^1} \\ 1 = \frac{\mu^2}{\bar{\lambda}^2} + \gamma^{22} + \gamma^{21} \frac{\bar{\lambda}^1}{\bar{\lambda}^2} \end{cases} \quad (11)$$

where $\bar{\lambda}^1 = \frac{\mu^1 * \gamma^{22} - \mu^1 - \mu^2 * \gamma^{12}}{\gamma^{11} + \gamma^{22} + \gamma^{12} * \gamma^{21} - \gamma^{11} * \gamma^{22} - 1}$, $\bar{\lambda}^2 = \frac{\mu^2 * \gamma^{11} - \mu^2 - \mu^1 * \gamma^{21}}{\gamma^{11} + \gamma^{22} + \gamma^{12} * \gamma^{21} - \gamma^{11} * \gamma^{22} - 1}$, $\gamma^{11} = \frac{\alpha_1^{11}}{\beta_1^{11}} + \frac{\alpha_2^{11}}{\beta_2^{11}}$, $\gamma^{22} = \frac{\alpha_1^{22}}{\beta_1^{22}} + \frac{\alpha_2^{22}}{\beta_2^{22}}$, $\gamma^{12} = \frac{\alpha_1^{12}}{\beta_1^{12}}$, $\gamma^{21} = \frac{\alpha_1^{21}}{\beta_1^{21}}$. The Hawkes parameters to be estimated are two baselines parameters (μ^1, μ^2), eight self-excitation parameters ($\alpha_1^{11}, \alpha_1^{22}, \alpha_2^{11}, \alpha_2^{22}, \beta_1^{11}, \beta_1^{22}, \beta_2^{11}, \beta_2^{22}$), and four cross-excitation parameters ($\alpha_1^{12}, \alpha_1^{21}, \beta_1^{12}, \beta_1^{21}$).

B. Maximum likelihood estimation

Following Ozaki (1979), Ogata (1978), and Ogata (1981), the log-likelihood for an intensity function $\lambda(t|\theta)$ of a bivariate Hawkes point process model where events are occurring in the interval $[0, T]$ is

$$\ell(\theta) = \ell^1(\theta^1) + \ell^2(\theta^2) \quad (12)$$

where $\theta = (\theta^1, \theta^2)$, and $\theta^i = (\mu^i, \alpha_1^{ii}, \beta_1^{ii}, \alpha_2^{ii}, \beta_2^{ii}, \alpha_1^{ij}, \beta_1^{ij})$ is a parameter vector of the intensity for type i events which varies in the parameter space Θ^i where $i \neq j$. Evaluating $\ell(\theta)$ gives

$$\begin{aligned} \ell(\theta) = & - \int_0^T \lambda^1(t|\theta^1) dt + \int_0^T \log \lambda^1(t|\theta^1) dN^1(t) - \int_0^T \lambda^2(t|\theta^2) dt \\ & + \int_0^T \log \lambda^2(t|\theta^2) dN^2(t) \end{aligned} \quad (13)$$

Plugging in the expression (7) into (13) and evaluating, we can write the log-likelihood function in extended form

$$\begin{aligned}
\ell(\theta) = & - \left[\mu^1 T + \frac{\alpha_1^{11}}{\beta_1^{11}} \sum_{i=1}^n (1 - e^{-\beta_1^{11}(T-t_i)}) + \frac{\alpha_2^{11}}{\beta_2^{11}} \sum_{i=1}^n (1 - e^{-\beta_2^{11}(T-t_i)}) \right. \\
& \left. + \frac{\alpha_1^{12}}{\beta_1^{12}} \sum_{j=1}^m (1 - e^{-\beta_1^{12}(T-t_j)}) \right] \\
& + \sum_{i=2}^n \log(\mu^1 + \alpha_1^{11} R_1^{11}(i) + \alpha_2^{11} R_2^{11}(i) + \alpha_1^{12} R_1^{12}(i)) \\
& - \left[\mu^2 T + \frac{\alpha_1^{22}}{\beta_1^{22}} \sum_{j=1}^m (1 - e^{-\beta_1^{22}(T-t_j)}) + \frac{\alpha_2^{22}}{\beta_2^{22}} \sum_{j=1}^m (1 - e^{-\beta_2^{22}(T-t_j)}) \right. \\
& \left. + \frac{\alpha_1^{21}}{\beta_1^{21}} \sum_{i=1}^n (1 - e^{-\beta_1^{21}(T-t_i)}) \right] \\
& + \sum_{j=2}^m \log(\mu^2 + \alpha_1^{22} R_1^{22}(j) + \alpha_2^{22} R_2^{22}(j) + \alpha_1^{21} R_1^{21}(j))
\end{aligned} \tag{14}$$

where the recursive terms are defined as follows: $R_1^{11}(i) = e^{-\beta_1^{11}(t_i-t_{i-1})}(1 + R_1^{11}(i-1))$, $R_2^{11}(i) = e^{-\beta_2^{11}(t_i-t_{i-1})}(1 + R_2^{11}(i-1))$, $R_1^{22}(j) = e^{-\beta_1^{22}(t_j-t_{j-1})}(1 + R_1^{22}(j-1))$, $R_2^{22}(j) = e^{-\beta_2^{22}(t_j-t_{j-1})}(1 + R_2^{22}(j-1))$, $R_1^{12}(i) = e^{-\beta_1^{12}(t_i-t_{i-1})}R_1^{12}(i-1) + \sum_{\{j':t_{i-1} \leq t_{j'} < t_i\}} e^{-\beta_1^{12}(t_i-t_{j'})}$, and $R_1^{21}(i) = e^{-\beta_1^{21}(t_j-t_{j-1})}R_1^{21}(j-1) + \sum_{\{i':t_{j-1} \leq t_{i'} < t_j\}} e^{-\beta_1^{21}(t_j-t_{i'})}$.

As the log-likelihood function corresponding to the exponential kernel can be computed recursively, it is suitable for maximum likelihood estimation (MLE).

C. Application

For the application of the Hawkes process to our case, we assume a proxy of liquidity $x^i(t)$, which is updated frequently yet at irregular intervals. In our application, $x^i(t)$ is the relative quoted bid-ask spread¹ of an individual stock on venue i at time t . The counting process $N^i(t)$ of events relies on the identification of liquidity shocks following the method of threshold exceedances. Threshold exceedances have been employed in the work on liquidity resilience by Panayi et al. (2015), Panayi, et al. (2018), and Danielsson et al. (2018), who propose the measure of Threshold Exceedance Duration (TED) and Liquidity Resilience Profile. TED is defined as “the length of time between the point at which the liquidity metric of choice² deviates from a threshold liquidity level, θ , (in the direction of less liquidity), and the point at which it returns to at least that level again” (Danielsson et al. 2018). Thus, liquidity shocks are defined as thresholds exceedances which are events occurring when the liquidity metric,

¹ Our specification here differs from that of Schneider et al. (2018), who use as a proxy for liquidity a variable derived from three metrics of liquidity through Principal Component Analysis: bid-ask spread, total quoted volume and inverse depth.

² This methodology can accommodate further liquidity metrics, e.g. Xetra Liquidity Measure. While recognizing that the depth captures further aspects of the LOB updates than the spread, we keep in mind that generally 95% of transactions hit the top of the limit order book. Therefore, we chose to focus on the top of the book and thus on the best bid-ask spread. A comparison of the performance of TED on both spread and XLM is shown in Panayi, Peters, & Kosmidis (2015).

$x^i(t)$, exceeds the threshold, θ^i , in the direction of less liquidity¹. In our application, θ^i corresponds to the 95th percentile² of the hourly empirical distribution of the relative quoted bid-ask spread of the respective day for a selected stock and a selected venue i , e.g. 9 a.m. 15-01-2019 for Lloyds on Chi-X. This leads to an hour-stock-venue specific threshold. Figure 4 shows a hypothetical set-up whereby the relative quoted spread fluctuates over time, exceeding the 95th percentile at time 2, and again at time 7. These two events are recorded as liquidity shocks. The spread returns underneath the 95th percentile at time 4, resulting in a TED of 2 milliseconds. The inter-event duration is 5 milliseconds. In the context of market resiliency, a high number of shocks may be indicative of a spiralling effect. Spiralling occurs when the return to “normal” is only momentary and another deviation from “normal” occurs soon after, indicating that the recovery is not meaningful.

Insert Figure 4 here

We motivate our choice of the threshold based on the literature as follows. We set the liquidity threshold level as the 95th percentile of the empirical distribution in line with the specification of Schneider, Lillo, & Pelizzon (2018). This threshold level serves our interest in capturing severe intraday liquidity shocks in contrast with any departure from the median or mean. Moreover, Panayi et al. (2015), Panayi, et al. (2018), and Danielsson et al. (2018) show how threshold exceedances work for a large range of liquidity thresholds, e.g. 50th percentile, 90th percentile, etc., and that the main drivers of TED remain valid whatever the threshold. We set the liquidity threshold as venue specific to acknowledge for individual venue-specific dynamics such as the fact that Chi-X often has higher threshold levels than LSE (as shown in Figure 2). In addition, we conduct a robustness analysis using a common cross-venue threshold. Our choice of an hourly time window is grounded in the strong intraday patterns observed across trading venues (see Figure 2, Biais et al. (1995), Tannous et al. (2013)). Furthermore, relying on hourly thresholds allows us to minimize any potential impact of intraday changes in tick sizes on TED³.

The above measurement specifications are in line with the characteristics of our dataset, e.g. high frequency data resolution, diurnal patterns and the focus of our study on severe liquidity events. We estimate the Hawkes parameters through MLE on a one-day-window⁴ separately for each of the 63 trading days and for each of the 99 stocks over the two venues of interest: LSE as venue i , and Chi-X as venue j . For each stock-day estimation, we check whether the Hawkes process is stationary using the stationarity condition derived in equation (10). We discard those estimations for which the process is non-stationary. We also discard estimations where the number of events on both LSE and Chi-X taken together is below 20 and the estimations where the diagonal of the hessian is negative, indicating

¹ Our specification of an event differs from that of Schneider et al. (2018), who define a shock as the instance “when there is a large and abrupt increase in illiquidity such that the speed of increase in illiquidity is over a threshold θ^i ”.

² The relative threshold approach has been also used in the empirical literature by Kempf et al. (2015) where the authors use the mean as a threshold.

³ We observe changes in tick size for 17 stocks out of 99 stocks.

⁴ We follow the specifications of Schneider et al. (2018) for the choice of the time window.

that the fit did not converge in a (local) minimum. Discarded cases amount to 3.13% of all estimations. Finally, using equation (11) we compute for each stock and day the fraction of events attributable to the baseline intensity, the fraction of events due to self-excitation on venue i , and the fraction of events on venue i attributed to cross-excitation from the paired venue j .

5 Results

First, we present and discuss the outcomes describing the liquidity shocks on LSE and Chi-X along with the severity and duration of these events. We continue by showing and discussing the results of the estimation of parameters in order to uncover whether shocks spiral within a venue and spill over across venues with implications for resiliency. All results are reported for both venues and for three market capitalisation subsamples.

A. *Liquidity shocks, severity, and duration*

Table 3 shows the summary statistics of the hourly number of shocks per stock-day on LSE and Chi-X, including market capitalisation subsamples. In average, more shocks occur per day on Chi-X (259 shocks) than on LSE (204 shocks). On both LSE and Chi-X, the number of shocks is the highest for the high market capitalisation bin (mean of 374 on LSE and 460 on Chi-X) and lowest for the low market capitalisation bin (mean of 111 on LSE and 152 on Chi-X). More than half of the shocks occurring in FTSE 100 stocks are happening in the high market capitalisation bin. This observation may be explained by the fact that generally, stocks with high market capitalisation are among the most actively traded and attract more HFTs, which means more quoting activity in large-cap stocks than in small-cap stocks. Figure 5 plots the daily number of shocks over the time of the sample for each market capitalisation subsample. We document that shocks occur generally more frequently on Chi-X than on LSE, independent of market capitalisation. The average number of shocks is higher on Chi-X than on LSE for nearly each day of the sample and for each of the market capitalisation subsamples. It is visible that the dynamic of the number of shocks follows the dynamic of the quoting activity shown in Figure 3. The distribution of the daily number of shocks is right-skewed with a long right tail. The maximum daily number of shocks observed on LSE is 1597 shocks (recorded for BP on February 4, 2015) and on Chi-X 1871 shocks (recorded for BP on February 2, 2015).

Insert Table 3 and Figure 5 here

Importantly, as shown in Figure 6, the number of shocks is higher on Chi-X than on LSE consistently throughout all the hours of the day. The number of shocks on both venues follows a similar diurnal pattern: higher at the opening, lower during lunch hours, higher after 2 p.m. GMT. Thus, the occurrence of shocks follows the intraday pattern of market activity. We document that the number of shocks is the highest in the afternoon after 2 p.m. GMT on both LSE and Chi-X. This clustering may be linked to the US market opening at 9:30 ET for NYSE and NASDAQ, corresponding to 2:30 p.m. GMT during most

of the period of our sample (except for the period March 8 to March 28, 2015 corresponding to 1:30 p.m. GMT due to different US and European daylight savings change date).

Insert Figure 6 here

Table 4 shows the summary statistics for the TED in milliseconds. A higher TED value is indicative of low market resiliency. The median duration of threshold exceedances on LSE is 116 milliseconds and on Chi-X 200 milliseconds confirming an intensive activity in the subsecond environment and below human reaction times. This difference in TED levels between LSE and Chi-X is not the result of differences in latencies since proximity and host connectivity latency for LSE, as well as for Chi-X, are known to be of the order of microseconds. 99% of instances of threshold exceedance occurring on LSE and Chi-X within the FTSE 100 stocks last on average less than 61 seconds on LSE and less than 52 seconds on Chi-X. Thus, measuring effectively market resiliency requires high-resolution data with a frequency more granular than 1 minute. The distribution of duration is highly right-skewed on both venues: on LSE (Chi-X) the mean TED is 3,920 milliseconds (3,387 milliseconds), and the median TED is 116 milliseconds (200 milliseconds). The maximum TED observed on LSE and Chi-X is approximately 57 minutes. To inspect the frequency distribution closer, we plot the histograms of the log-transformed TED distributions and inter-event distributions for LSE and Chi-X in Figure 7. Indeed, the fraction of shocks reverting within 100 milliseconds is larger on LSE than on Chi-X. This result provides evidence that following a shock LSE often reverts quicker than Chi-X. Considering the distribution of inter-event durations for LSE and Chi-X, we note that a small share of these ultra-low duration shocks are only momentary and liquidity dries up quickly again. Finally, the fractions of liquidity dry-ups with a duration above 100 milliseconds is higher on Chi-X than on LSE.

Insert Table 4 and Figure 7 here

Considering differences across stocks, TED increases with decreasing market capitalisation. The quickest to revert following an exceedance are stocks in the high market capitalisation bin. The longest lasting exceedances are observed in the low market capitalisation bin. This is true for both LSE and Chi-X. By comparison, the time to revert following an exceedance for stocks in the low market capitalisation bin (median TED of 146 milliseconds on LSE and 367 milliseconds on Chi-X) is more than double the time to revert of stocks in the high market capitalisation bin (median of 114 milliseconds on LSE and 168 milliseconds on Chi-X). This relation holds also for the other quantiles of the TED distribution.

Furthermore, we define a proxy metric to capture the severity of an exceedance. Our proxy metric is defined as the difference in relative terms between the maximum spread observed during the duration of the exceedance of the threshold, $\max(S)$, and the threshold itself, θ , as follows

$$severity = \frac{\max(S) - \theta}{\theta} \quad (15)$$

In the context of market resiliency, severer exceedances may affect the time it takes for the market to reorganize itself. Table 5 shows the summary statistics of the severity of threshold exceedances for each venue and each market capitalisation bin. We document that threshold exceedances occurring in the FTSE 100 stocks are severer on Chi-X than on LSE: the median severity on Chi-X is 0.111 and on LSE 0.004. The difference in the severity of exceedances between the two venues is even stronger, considering that the threshold values on Chi-X are already higher than the threshold values on LSE. This difference in severity across venues remains true for all subsamples of stocks based on market capitalisation.

Insert Table 5 here

For a broad comparison with the study of Degryse et al. (2005) that employs the number of best limit updates for the spread and depth to return, Table 6 shows the number of spread updates recorded between the time of threshold exceedance and the time when the spread returns below the threshold. Degryse et al. (2005) find that following an aggressive order it takes around 20 best limit updates for the spread to return to their initial level prior shock. Using a larger definition of shocks than aggressive orders, we find that on both venues the spread returns very quickly below threshold: for over 75% of observations of the timeseries, it takes one update for the spread to return below the threshold on LSE as well as on Chi-X. This remains true across all market capitalisation bins. At the right tail of the distribution, we observe an order of magnitude of 4-7 updates for the spread to return below the threshold on LSE and Chi-X. We conjecture that the significantly lower number of updates that we find in comparison with Degryse et al. (2005) is due primarily to the large technology leap between 1998, the year of the sample used by Degryse et al. (2005), and 2015, the year of our sample, reflecting a substantial intensification in market activity.

Insert Table 6 here

For a comparison with the study of Kempf et al. (2015) that employs the speed of return to normal, we discuss the methodological and empirical differences in Appendix B and show that this metric is not fitting the ultra-high-frequency data that we employ.

Zooming in for one individual stock, Lloyds, Figure 8 plots the occurrence and duration of threshold exceedance over time. We observe that short-lasting shocks are often followed by another shock. Furthermore, we note that shocks tend to cluster, in particular around 3 pm. The threshold exceedances occurring around this time tend to be severer than other instances occurring throughout the rest of the day.

Insert Figure 8 here

B. Estimation results

Before proceeding with the estimation of Hawkes parameters, we examine whether it is more reasonable to model the arrival intensity through a Hawkes process than through a Poisson process. A Poisson process assumes a constant arrival intensity of events. We evaluate the fittingness of the observed data to a Poisson process through a quantile by quantile plot, where the quantiles of the observed data are plotted against the quantiles of the exponential distribution. If the events corresponding to our data were to arrive with a constant intensity, then most of the points in a quantile-quantile plot would fit a straight line going through the origin. Figure 9 shows the quantile-quantile plot of inter-event durations against exponential quantiles for a selected stock and a subset of days of our sample. We see in Panel A that the data doesn't fit a straight line and we observe a significant number of data points of short inter-event times, which is indicative of clustering of events and thus justifies the use of Hawkes processes. After estimating the arrival intensity through a Hawkes process, we rescale the quantiles of inter-event durations by the estimated intensity and replot these against exponential quantiles. As shown in Panel B, we obtain a good fit for the vast majority of observations. The few observations with a less good fit are those with extremely high (above the 95th percentile, e.g. 840 seconds on January 2, 2015) and extremely low inter-event durations (below the 10th percentile, e.g. 35 milliseconds on January 2, 2015).

Insert Figure 9 here

The average daily fractions of events shown in Figure 10 correspond to the estimates for the baseline-, self-, and cross-excitation. The reported values are averages built across all stocks of the sample for each venue, LSE and Chi-X. Throughout the time period of our sample, on average, approximately one-third of shocks arrive randomly and are neither due to self-excitation nor to cross-excitation, irrespective of venue and of the time. More than half of the threshold exceedances occurring on either LSE or Chi-X are due to self-excitation. This result indicates a substantial spiralling effect. The mean fraction of events due to spiralling is higher on Chi-X (55%) than on LSE (53%). The difference in means is significant at 0.1% level. Finally, the mean fraction of events due to spillover from the competing venue is higher on LSE (13%) than on Chi-X (11%). The difference in means is significant at 0.1% level.

Insert Figure 10 here

In order to examine the cross-stock variation of the estimates, we present in Figure 11, Figure 12, and Figure 14 the estimates for each type of excitation, i.e. baseline excitation, self-excitation, and cross-excitation, across market capitalisation bins. The fraction of events due to baseline-excitation are lowest in the high market capitalisation bucket and highest in the low market capitalisation bucket, indicating that baseline intensity varies across stocks. We document also variations across time for the differences in baseline excitation between the two venues. The fraction of events attributable to baseline intensity is at times higher on LSE than Chi-X and at other times higher on Chi-X than on LSE. In absolute number of events, Chi-X has in average across all stocks for each day a higher number of events attributable to baseline excitation than LSE. The different dynamics observed in fractions versus absolute number of events are largely due to the fact that Chi-X overall has a higher number of shocks than LSE.

Insert Figure 11 here

Next, we look at the spiralling effect, which in our study is informative about the resiliency of liquidity. The fraction of events due to self-excitation is the highest for the high market capitalisation bin and lowest for the low market capitalisation bin, indicating that spiralling levels vary across stocks and across days. For stocks with highest market capitalisation self-excitation is higher on Chi-X (mean 61%) than on LSE (mean 57%) consistently across all days of the sample. This confirms that in average for these stocks more spiralling occurs on Chi-X than on LSE. This remains true also when considering the absolute number of events due to spiralling in the high market capitalisation subsample of stocks. For the medium and low market capitalisation bin, the superiority of the spiralling effect of one venue over the other changes across the days of the sample. In terms of absolute number of events, also for these stocks, there are more events due to spiralling on Chi-X than on LSE.

Insert Figure 12 here

Since a substantial fraction of events is due to self-excitation, we analyse the timescale of decay of self-excitation. This metric captures the time between an event occurring in a stock on venue i and a next event occurring in the same stock on the same venue as the result of self-excitation. The timescale of decay for self-excitation is estimated through the parameter $1/\beta^{ii}$. Figure 13 shows the timescale of decay for a fast propagation through time. The time for an event to spiral is of the order of milliseconds: the median time for a shock to be followed by another shock in the same stock and on the same venue is 25 milliseconds on LSE and 99 milliseconds on Chi-X¹. This is in line with our earlier observation that a larger fraction of shocks on LSE, in comparison with Chi-X, has ultra-low TEDs and inter-event durations (below 100 milliseconds). The timescale of decay for self-excitation varies across the days of the sample as well as across stocks.

Insert Figure 13 here

Such a large share of spiralling and a short timescale of decay may be due to several possible economic mechanisms. A first possible explanation is a positive autocorrelation in order flow due to algorithmic order splitting by investors (e.g. execution algorithms of meta-orders such as Time-weighted average price (TWAP) and Volume-weighted average price (VWAP)). A second possible rational for spiralling are market overreactions due to competing traders' response to a signal. A third explanation for spiralling is the delayed reaction to information in order flows due to speed differentials across traders. Finally, the hot potato effect associated with multiple traders simultaneously offloading excess positions may also trigger such effects. Relative to the existing literature on market resiliency, the spiralling effect is in line with the finding of Degryse et al. (2005) who observe that the probability of a particular order type to be followed by an order of the same type is higher than the probability of being followed by a different order type. This observation confirms the diagonal effect documented earlier by Biais et al.

¹ From our discussion with industry participants, the indicative half-life for an informed market move to be picked up by another market participant and exploited is 10 milliseconds for fast venues and 100 milliseconds for slow venues.

(1995) that may be the result of order splitting strategies or collective behaviour of investors imitating each other's reaction to a shock event. The stronger spiralling effect observed on Chi-X versus LSE, in particular for large stocks, might be indicative of different liquidity provision dynamics on the two venues that are the result of different market models. On LSE, market makers ensure the liquidity provision by submitting limit orders on both sides of the market, i.e. bid and ask. Market makers incur no trading fees. The liquidity maker and the liquidity taker pay the same level of fee per transaction. On Chi-X, there are no market makers and the liquidity provision is incentivized by the pricing scheme. The liquidity maker receives a rebate whilst the liquidity taker is charged. As such, a liquidity maker on Chi-X has an incentive to supply liquidity in frequent small servings, rather than in one single serving, in order to cumulate rebates.

The fraction of events due to cross-excitation are relatively similar across market capitalisation bins. Thus, the strong cross-bin differences observed for baseline excitation and self-excitation are not discernible for cross-excitation. For nearly all the days of the sample, the average daily fraction of events occurring on Chi-X and spilling over to LSE is larger than the fraction of events occurring on LSE and spilling over to Chi-X. The absolute number of events due to spillover is in average similar for LSE and Chi-X. When zooming into the sample on individual stocks, we do find variations in the direction of spillover across days: the spillover effect is at times stronger from Chi-X to LSE and on other days is stronger from LSE to Chi-X.

Insert Figure 14 here

Having found an economically significant daily fraction of spillovers occurring between LSE and Chi-X, we examine the timescale of decay of an event occurring on venue i until the occurrence of a next event on venue j through the parameter $1/\beta_{ij}$ ($i \neq j$). Figure 15 shows that the timescale of decay for spillovers is of the order of milliseconds. We further observe that the speed of spillovers varies across days of the sample and across stocks. Finally, for most stocks and days, illiquidity spillovers transmit quicker from Chi-X to LSE than from LSE to Chi-X. The median time for an event to spill over from Chi-X to LSE is 23 milliseconds and from LSE to Chi-X – 36 milliseconds.

Insert Figure 15 here

The direction of the spillover effect is likely related to the venues' price leadership and the relative contribution to price discovery of one venue over another. This mechanism is described in the studies by Ibikunle (2018) and Riordan et al. (2012). Ibikunle (2018) finds that Chi-X's price leadership allows it to attract order flow away from LSE. The author further provides empirical evidence that the share of informed trading is higher on Chi-X than on LSE throughout the trading day. Finally, the study also shows that Chi-X displays more HFT activity than LSE. These mechanisms are consistent with the stronger spillover effect that we document from Chi-X to LSE. As the time period of the sample analysed in the study by Ibikunle (2018) – second half of 2014 – is close to the time period of our sample – first quarter of 2015, we conjecture that the mechanism outlined by this study provides valid supporting evidence for the stronger spillover effect that we document from Chi-X to LSE. The observed spillover

effect and the ultra-fast time to spill over reflect the activity of market participants operating on both venues and having technological capabilities (like colocation) to react rapidly on venue j after observing an event (e.g. large trade) on venue i .

6 Significance and Robustness

A. Significance of parameter estimates

We test for every estimation whether the spiralling effect and the spillover effect are statistically significant. For this purpose, we compute from the inverse of the negative hessian the confidence intervals for the α^{ij} parameters and verify that these parameters are significantly different from zero under the normality assumption. The central limit theorem for maximum likelihood estimators of stationary point processes supports the normality assumption. Figure 16 shows the daily fraction of estimations with significant α^{ij} parameters at the 1% and 0.1% level. The fraction of estimations with significant spiralling effect and spillover effect are high and vary throughout the time period of our sample. The spiralling effect is significant at the 0.1% level for minimum 83% and maximum 97% of stocks. The spillover effect is significant at the 0.1% level for minimum 87% and maximum 99% of stocks. This result supports the evidence of illiquidity spiralling and illiquidity spillover effects.

Insert Figure 16 here

We further test whether LSE and Chi-X estimates are significantly different. All Chi-X spiralling estimates are significantly different from the corresponding LSE estimates at the 0.1% level.

B. Alternative specifications for a threshold common to LSE and Chi-X

In addition, as an alternative specification, we define a cross-venue threshold common to both venues as the minimum between the hour-stock-venue threshold for LSE and the hour-stock-venue threshold for Chi-X. This alternative definition allows us in particular to check the validity of the spillover effect. As expected, when changing the definition of the shock the total number of shocks increases and particularly so for Chi-X. The increase in events is observable in all market capitalisation bin and particularly so for high market capitalisation shocks for which the quoting activity is most intense. When running the maximum likelihood estimation, in comparison with the baseline, there are more estimations where the stationarity condition does not hold or the hessian is negative, indicating that the fit did not converge in a (local) minimum. In total, we discard 7.2% of estimations (in comparison with 3.1% under baseline). Figure 17 shows the fractions of events due to base-, self-, and cross-excitation per venue when using a cross-venue common threshold. Overall, the results that we obtain using this alternative specification remain qualitatively similar to the results obtained using hour-stock-venue thresholds. The fractions of events due to base-excitation and self-excitation do not change when employing a cross-

venue common threshold. However, we do observe an increase in the fractions of events on LSE that are spillovers from Chi-X: the mean of daily fractions of events increases to 18% when using a cross-venue common threshold in comparison with 13% when using an hour-stock-venue specific threshold. The mean of daily fractions of events on Chi-X that are spillovers from LSE (11%) does not change when using a common threshold. We conclude that also under an alternative specification a spillover from Chi-X to LSE occurs more often than from LSE to Chi-X.

Insert Figure 17 here

Next, we examine the spiralling effect by market capitalisation subsamples. Figure 18 plots the spillover-related fraction of events for each market capitalisation bin when using a common cross-venue threshold. We observe a substantial increase in the fractions of events spilling over from Chi-X to LSE for all subsamples. Whereas we do observe variations across individual stocks, we do not document substantial differences in the spillover effect across market capitalisation subsamples. In terms of absolute number of events, relative to the baseline, there are more days in high market capitalisation bin when there are more spillovers from Chi-X to LSE than from LSE to Chi-X.

Insert Figure 18 here

This robustness check provides supporting evidence for the validity of the spiralling effect and in particular of the spillover effect.

7 Conclusion

Global equity limit order markets have changed fundamentally due to fragmentation of markets and technological advances increasing the speed of trading. This phenomenon had a dramatic impact on the functioning and stability of financial markets. In this environment, knowing the properties of market resiliency becomes even more important for market participants and regulators. This paper studies how resilient are modern financial venues of a fragmented equity market when liquidity is scarce. We study the incidence, duration in the sense of TED, and severity of liquidity shocks occurring in FTSE 100 stocks on LSE, a traditional exchange, and Chi-X, an incumbent venue, with millisecond accuracy. We further employ multivariate Hawkes processes to identify the meaningfulness of recoveries from severe liquidity shocks by detecting spiralling effects within one venue as well as cross-venue spillover effects.

First, we find that liquidity shocks happen more often on Chi-X than on LSE consistently throughout continuous trading time. They are also more severe on Chi-X than on LSE. We observe instances of extreme illiquidity on Chi-X where no bid or no ask quote exists. There are no such instances on LSE thanks to market makers guaranteeing 2-way price delivery, whereas no market making mechanism exists on Chi-X. More than half of the shocks recorded for all stocks occur in large stocks on both Chi-X and LSE, however, large-cap stocks are also the quickest to revert to normal in comparison with mid- and small-cap stocks of the FTSE 100. The median duration of threshold exceedances on LSE is 116

milliseconds and 200 milliseconds on Chi-X, which is below known human reaction times and confirms an intensive activity in the subsecond environment on both venues. Among our most important findings is that 99% of shocks on both Chi-X and LSE last close to one minute. This indicates a substantial amount of quoting within the one-minute environment and suggests that to effectively measure market resiliency we need high-resolution data with a frequency below one minute.

Second, using multivariate Hawkes processes, we find that, on average, more than half of liquidity shocks occurring on Chi-X and LSE are due to spiralling effects. This result holds for large-, mid-, and small-cap stocks of the FTSE 100. The spiralling effect is stronger on Chi-X than on LSE and in particular for large stocks. This may be indicative of different liquidity provision dynamics on the two venues, suggesting that the liquidity supply on Chi-X, a venue without market making, is thinner than on LSE, a venue with market making. Finally, cross-excitation capturing cross-venue spillover effects is economically significant. We find that on average across stocks and days of our sample 13% of daily shocks on LSE are spillovers from Chi-X, whereas 11% of daily shocks on Chi-X are spillovers from LSE. The time to decay of the spiralling effect, as well as the time of decay of the spillover effect, are both below 100 milliseconds. The spiralling and spillover estimates are both statistically significant and robust to an alternative shock definition. We conjecture that the direction of the spillover effect is related to the price leadership of one venue over another when competing to attract order flow as posed by Ibikunle (2018). The observed spillover effect and the ultra-fast time to spill over may reflect the activity of market participants operating on both venues and having technological capabilities (like colocation) to react rapidly on one venue after observing an event (e.g. large trade) on the competing venue.

Overall, on average across the FTSE 100 stocks and in comparison with LSE, Chi-X has more liquidity shocks, longer lasting shocks, severer shocks, a larger spiralling effect of liquidity shocks and a significant spillover to LSE. The implications of different market resiliency on different venues may be costly when liquidity is scarce. This may be a source vulnerability, penalizing participants with cross-venue access engaged in cross-market arbitrage strategies (Menkveld & Yueshen, 2018) as well as participants with access to only the less resilient venue. While our results suggest that Chi-X is a less resilient venue than LSE for the FTSE 100 stocks, it is yet an unresolved question what specific aspects of market microstructure drive the difference in market resiliency between competing fragmented venues. We present several possible economic mechanisms that may explain the spiralling and spillover effects that we find and we hope that future research will shed further light on this aspect. A further interesting question for future research is integrating the severity of shocks in the Hawkes process. This question is important since incidence and duration alone may not offer a complete picture of market resiliency.

References

- Bacry, E., Mastromatteo, I., & Muzy, J.-F. (2015). Hawkes processes in finance. *Market Microstructure and Liquidity*, 1(1550005).
- Beddington, J., Bond, P., Cliff, D., Houstoun, K., Linton, O., Goodhart, C., & Zigrand, J. P. (2012). *The Future of Computer Trading in Financial Markets An International Perspective*. Retrieved from <http://www.bis.gov.uk/foresight>
- Biais, B., Hillion, P., & Spatt, C. (1995). An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse. *The Journal of Finance*. <https://doi.org/10.1111/j.1540-6261.1995.tb05192.x>
- Boneva, L., Linton, O., & Vogt, M. (2016). The Effect of Fragmentation in Trading on Market Quality in the UK Equity Market. *Journal of Applied Econometrics*. <https://doi.org/10.1002/jae.2438>
- Bowsher, C. G. (2007). Modelling security market events in continuous time: Intensity based, multivariate point process models. *Journal of Econometrics*, 141(2), 876–912. <https://doi.org/10.1016/j.jeconom.2006.11.007>
- Brunnermeier, M. K., & Pedersen, L. H. (2009). Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6). <https://doi.org/10.1093/rfs/hhn098>
- Cartlidge, J., Szostek, C., De Luca, M., & Cliff, D. (2012). Too fast too furious: Faster financial-market trading agents can give less efficient markets. In *ICAART 2012 - Proceedings of the 4th International Conference on Agents and Artificial Intelligence* (Vol. 2, pp. 126–135). <https://doi.org/10.5220/0003720301260135>
- Colliard, J. E., & Hoffmann, P. (2017). Financial Transaction Taxes, Market Composition, and Liquidity. *Journal of Finance*. <https://doi.org/10.1111/jofi.12510>
- Danielsson, J., Panayi, E., Peters, G., & Zigrand, J.-P. (2018). *Market Resilience*. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3169755>
- Degryse, H., De Jong, F., & Kervel, V. Van. (2015). The impact of dark trading and visible fragmentation on market quality. *Review of Finance*. <https://doi.org/10.1093/rof/rfu027>
- Degryse, H., De Jong, F., Van Ravenswaaij, M., & Wuyts, G. (2005). Aggressive orders and the resiliency of a limit order market. *Review of Finance*, 9(2), 201–242. <https://doi.org/10.1007/s10679-005-7590-6>
- Delaney, A., Fenick, W., Tyc, S., Saade, V., & Marsh, I. (2015). *TRADE EUROPE NOW! The Why and How of Electronic Execution in the EU and Beyond*. Retrieved from https://www.mckay-brothers.com/wp-content/uploads/2015/10/Interxion_TradingEurope_Paper.pdf
- Félez-Viñas, E. (2019). Effects of Market Fragmentation on Resiliency. *SSRN Electronic Journal*.

<https://doi.org/10.2139/ssrn.3315219>

- Foucault, T., Kadan, O., & Kandel, E. (2005). Limit order book as a market for liquidity. *Review of Financial Studies*, 18(4), 1171–1217. <https://doi.org/10.1093/rfs/hhi029>
- Foucault, T., & Menkveld, A. J. (2008). Competition for order flow and smart order routing systems. *Journal of Finance*. <https://doi.org/10.1111/j.1540-6261.2008.01312.x>
- Francis, R., & Bekera, B. (2014). A metric and frameworks for resilience analysis of engineered and infrastructure systems. *Reliability Engineering and System Safety*, 121, 90–103. <https://doi.org/10.1016/j.ress.2013.07.004>
- Golub, A., Keane, J., & Poon, S.-H. (2012). High Frequency Trading and Mini Flash Crashes. *SSRN Electronic Journal*, 1–22. <https://doi.org/10.2139/ssrn.2182097>
- Gomber, P., Schweickert, U., & Theissen, E. (2015). Liquidity dynamics in an electronic open limit order book: An event study approach. *European Financial Management*, 21(1), 52–78. <https://doi.org/10.1111/j.1468-036X.2013.12006.x>
- Gresse, C. (2017). Effects of lit and dark market fragmentation on liquidity. *Journal of Financial Markets*. <https://doi.org/10.1016/j.finmar.2017.05.003>
- Haslag, P. H., & Ringgenberg, M. (2015). *The Causal Impact of Market Fragmentation on Liquidity*. SSRN. <https://doi.org/10.2139/ssrn.2591715>
- Hawkes, A. G. (1971). Spectra of some self-exciting and mutually exciting point processes. *Biometrika*. <https://doi.org/10.1093/biomet/58.1.83>
- Ibikunle, G. (2018). Trading places: Price leadership and the competition for order flow. *Journal of Empirical Finance*, 49(September), 178–200. <https://doi.org/10.1016/j.jempfin.2018.09.007>
- Johnson, N., Zhao, G., Hunsader, E., Qi, H., Johnson, N., Meng, J., & Tivnan, B. (2013). Abrupt rise of new machine ecology beyond human response time. *Nature Scientific Reports*. <https://doi.org/10.1038/srep02627>
- Kempf, A., Mayston, D., Gehde-Trapp, M., & Yadav, P. K. (2015). *Resiliency: A dynamic view of liquidity* (CFR Working Paper NO. 15-04).
- Kyle, A. S. (1985). Continuous Auctions and Insider Trading. *Econometrica*, 53(6), 1315–1335. <https://doi.org/10.2307/1913210>
- Large, J. (2007). Measuring the resiliency of an electronic limit order book. *Journal of Financial Markets*, 10(1), 1–25. <https://doi.org/10.1016/j.finmar.2006.09.001>
- Menkveld, A. J. (2013). High frequency trading and the new market makers. *Journal of Financial Markets*, 16(4), 712–740. <https://doi.org/10.1016/j.finmar.2013.06.006>
- Menkveld, A. J. (2014). High-frequency traders and market structure. *Financial Review*, 49(2), 333–

344. <https://doi.org/10.1111/fire.12038>


- Menkveld, A. J. (2018). High-frequency trading as viewed through an electron microscope. *Financial Analysts Journal*, 74(2), 24–31. <https://doi.org/10.2469/faj.v74.n2.1>
- Menkveld, A. J., & Yueshen, B. Z. (2018). The Flash Crash: A cautionary tale about highly fragmented markets. *Management Science*.
- Muni Toke, I. (2011). “Market Making” in an Order Book Model and Its Impact on the Spread. *Econophysics of Order-Driven Markets* (Vol. 9). Springer. https://doi.org/10.1007/978-88-470-1766-5_4
- O’Hara, M., & Ye, M. (2011). Is market fragmentation harming market quality? *Journal of Financial Economics*. <https://doi.org/10.1016/j.jfineco.2011.02.006>
- Obizhaeva, A. A., & Wang, J. (2013). Optimal trading strategy and supply/demand dynamics. *Journal of Financial Markets*, 16(1), 1–32. <https://doi.org/10.1016/j.finmar.2012.09.001>
- Ogata, Y. (1978). The asymptotic behaviour of maximum likelihood estimators for stationary point processes. *Annals of the Institute of Statistical Mathematics*. <https://doi.org/10.1007/BF02480216>
- Ogata, Y. (1981). On Lewis’ Simulation Method for Point Processes. *IEEE Transactions on Information Theory*, 27(1), 23–31. <https://doi.org/10.1109/TIT.1981.1056305>
- Ozaki, T. (1979). Maximum likelihood estimation of Hawkes’ self-exciting point processes. *Annals of the Institute of Statistical Mathematics*, 31(1), 145–155. <https://doi.org/10.1007/BF02480272>
- Panayi, E., Peters, G. W., Danielsson, J., & Zigrand, J.-P. (2018). Designating market maker behaviour in limit order book markets Designating market maker behaviour in Limit Order Book markets. *Econometrics and Statistics*, (5), 20–44. <https://doi.org/10.1016/j.ecosta.2016.10.008>
- Panayi, E., Peters, G. W., & Kosmidis, I. (2015). Liquidity commonality does not imply liquidity resilience commonality : a functional characterisation for ultra-high frequency cross-sectional LOB data. *Quantitative Finance*, 15(10), 1737–1758.
- Riordan, R., Storkenmaier, A., & Wagener, M. (2012). Do Multilateral Trading Facilities Contribute to Market Quality? *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.1852769>
- Schneider, M., Lillo, F., & Pelizzon, L. (2018). Modelling illiquidity spillovers with Hawkes processes: an application to the sovereign bond market. *Quantitative Finance*. <https://doi.org/10.1080/14697688.2017.1403155>
- Stiglitz, J. E. (1989). Financial markets and development. *Oxford Review of Economic Policy*. <https://doi.org/10.1093/oxrep/5.4.55>
- Tannous, G., Wang, J., & Wilson, C. (2013). The Intraday Pattern of Information Asymmetry, Spread,

and Depth : Evidence from the NYSE. *International Review of Finance*, 215–240.
<https://doi.org/10.1111/irfi.12005>


Wolfe, B., Seppelt, B., Mehler, B., Reimer, B., & Rosenholtz, R. (2020). Rapid holistic perception and evasion of road hazards. *Journal of Experimental Psychology. General*.
<https://doi.org/10.1037/xge0000665>

Figure 1. Daily trading volume for each market capitalisation subsample. The figures for the trading volume correspond to the number of shares traded from January 2 to March 31, 2015 for all stocks of the sample, except for 5 stocks for which we do not find any records. X-axis: day of the sample. Y-axis: volume in millions of shares.

Panel A. High market capitalisation bin



Panel B. Medium market capitalisation bin



Panel C. Low market capitalisation bin





Figure 2. Relative quoted spread cross-venue. The 50th and 95th percentiles are calculated from the empirical distribution of the relative quoted spread for each stock-day-hour and subsequently averaged across days and stocks for each venue.

Panel A. Daily spread.

X-axis: day of the sample.

Y-axis: relative quoted spread in basis points.



Panel B. Intraday spread.

X-axis: hour of the day.

Y-axis: relative quoted spread in basis points.






Figure 3. Number of spread updates for each market capitalisation subsample. The figures correspond to the count of changes in relative quoted spread throughout continuous trading from 8am to 16:30, excluding LSE auction times from both LSE and Chi-X data. X-axis: day of the sample. Y-axis: number of spread updates.

Panel A. High market capitalisation bin



Panel B. Medium market capitalisation bin



Panel C. Low market capitalisation bin




Figure 4. Hypothetical illustration of threshold exceedances as liquidity shocks and Threshold Exceedance Duration. TED is the length of time between the point at which the relative quoted spread exceeds the 95th percentile of the empirical distribution of the spread for the respective hour, day, and stock, which occurs at millisecond 2, and the point at which the relative quoted spread goes back below the 95th percentile again, which occurs at millisecond 4. Liquidity shocks occur at millisecond 2 and millisecond 7. Thus, in this illustrative example, the TED equals 2 milliseconds and the inter-event duration is 5 milliseconds.

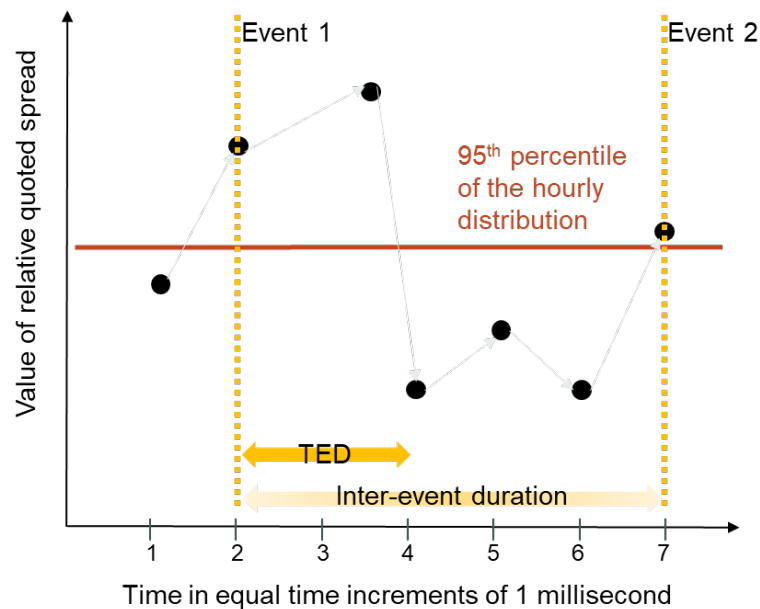


Figure 5. Average number of shocks per day of the sample for each market capitalisation subsample. X-axis: day of the sample. Y-axis: number of shocks.




Figure 6. Hourly number of threshold exceedances. Values correspond to averages build across all stocks of the sample per hour. X-axis: hour of the day. Y-axis: number of shocks.





Figure 7. Histogram of durations for LSE and Chi-X. The plotted data corresponds to binned log₁₀-transformed durations for LSE and Chi-X for all stocks and days of the sample. X-axis shows duration in milliseconds in log₁₀ base. Y-axis shows the frequency normalized by the total number of shocks. Note: the total number of shocks on Chi-X is larger than the total number of shocks on LSE.

Panel A. Threshold Exceedance Duration



Panel B. Inter-event duration





Figure 8. Threshold exceedance durations of Lloyds on LSE and Chi-X. The figure plots in both panels on X-axis the hour of the day and on Y-axis the day of the sample. Each line corresponds to a threshold exceedance duration and the colour code indicates the severity of the threshold exceedance.

Panel A. TED of Lloyds on LSE for day-hour threshold



Panel B. TED of Lloyds on Chi-X for day-hour threshold





Figure 9. Quantile-quantile plot of inter-event duration and exponential quantiles. The figure plots in both panels on X-axis the quantiles of the exponential distribution. In panel A on Y-axis the quantiles of the observed inter-event duration are shown whereas in panel B on Y-axis the quantiles of the inter-event duration rescaled by the Hawkes-estimated intensity are plotted. The observed inter-event durations are a subsample of data corresponding to the stock Lloyds for 5 trading days from January 2 to January 8, 2015 on LSE. Different colours indicate different days.

Panel A. Non-rescaled inter-event durations. X-axis: theoretical exponential quantiles in log10. Y-axis: sample quantiles in milliseconds.



Panel B. Rescaled inter-event durations. X-axis: theoretical exponential quantiles in log 10. Y-axis: intensity rescaled sample quantiles in log10.





Figure 10. Estimates of fractions of events due to baseline-, self-, and cross-excitation for each venue. The reported fractions are averages of daily estimates across all stocks. X-axis shows the fraction of events in decimals on a scale from 0 to 1. Y-axis shows the time period of the sample: January 2 to March 31, 2015.

Panel A. LSE



Panel B. Chi-X






Figure 11. Estimates of fraction of events and number of events due to base-excitation by market capitalisation bin. The reported figures are averages of estimates across all stocks of the respective bin. X-axis in the upper panels shows the fraction of events in decimals on a scale from 0 to 1 and in the lower panels the absolute number of events. Y-axis shows the time period of the sample: January 2 to March 31, 2015.

Panel A. High market capitalisation bin



Panel B. Medium market capitalisation bin




Panel C. Low market capitalisation bin




Figure 12. Estimates of fraction of events and number of events due to self-excitation by market capitalisation bin. The reported figures are averages of estimates across all stocks of the respective bin. X-axis in the upper panels shows the fraction of events in decimals on a scale from 0 to 1 and in the lower panels the absolute number of events. Y-axis shows the time period of the sample: January 2 to March 31, 2015.

Panel A. High market capitalisation bin



Panel B. Medium market capitalisation bin




Panel C. Low market capitalisation bin




Figure 13. Timescale of decay of self-excitation by market capitalisation bin. The reported figures correspond to the daily medians of estimated timescale of decay across all stocks of the respective bin. We interpret the figures as the time for a shock to spiral through the time dimension on venue i . The lower the value, the faster the shock propagates through time. X-axis shows the timescale of decay expressed in milliseconds. Y-axis shows the time period of the sample: January 2 to March 31, 2015.

Panel A. High market capitalisation bin



Panel B. Medium market capitalisation bin



Panel C. Low market capitalisation bin






Figure 14. Estimates of fraction of events and number of events due to cross-excitation by market capitalisation bin. The reported figures are averages of daily estimates across all stocks of the respective bin. X-axis in the upper panels shows the fraction of events in decimals on a scale from 0 to 1 and in the lower panels the absolute number of events. Y-axis shows the time period of the sample: January 2 to March 31, 2015.


Panel A. High market capitalisation bin






Panel B. Medium market capitalisation bin



Panel C. Low market capitalisation bin




bin
—●— chix
—●— lse



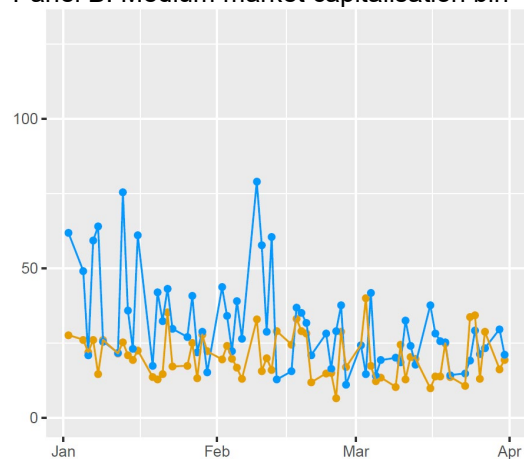
bin
—●— chix
—●— lse

Figure 15. Timescale of decay of cross-excitation by market capitalisation bin. The reported figures correspond to daily medians of estimated timescale of decay across all stocks of the respective bin. We interpret the figures as the time for a shock spilling over from venue i to venue j . The lower the value, the faster the shock propagates cross-venue. X-axis shows the timescale of decay expressed in milliseconds. Y-axis shows the time period of the sample: January 2 to March 31, 2015.

Panel A. High market capitalisation bin



Panel B. Medium market capitalisation bin



Panel C. Low market capitalisation bin





Figure 16. Significance of spiralling and spillover effects. The figure shows the daily fraction of estimations for which the self-excitation parameters describing the spiralling effect α^{ij} ($i = j$) and the cross-excitation parameters α^{ij} ($i \neq j$) describing the spillover effect are significantly different from zero at the 1% and 0.1% confidence level. X-axis shows the fractions of estimations with significant spiralling effect in Panel A and significant spillover effect in Panel B. Y-axis shows the time period of the sample: January 2 to March 31, 2015.

Panel A. Spiralling effect




Panel B. Spillover effect



Figure 17. Estimates of fractions of events due to baseline-, self-, and cross-excitation for each venue with alternative definition of a liquidity shock. The reported fractions are averages of daily estimates across all stocks. X-axis shows the fraction of events in decimals on a scale from 0 to 1. Y-axis shows the time period of the sample: January 2 to March 31, 2015.

Panel A. LSE



Panel B. Chi-X






Figure 18. Estimates of fraction of events and number of events due to cross-excitation by market capitalisation bin with alternative definition of a liquidity shock. The reported figures are averages of daily estimates across all stocks of the respective bin. X-axis in the upper panels shows the fraction of events in decimals on a scale from 0 to 1 and in the lower panels the absolute number of events. Y-axis shows the time period of the sample: January 2 to March 31, 2015.


Panel A. High market capitalisation bin



Panel B. Medium market capitalisation bin



Panel C. Low market capitalisation bin



bin
—●— chix
—●— lse






Figure 19. Measurement approach through TED, speed of return, speed of return from above, and speed of return from below.






Figure 20. TED versus data resolution. The figure plots the occurrence and duration of threshold exceedances of the quoted spread are computed on four datasets for a single stock (Aviva) for a subsample of one month (January 2015). The threshold is the monthly median. Panel A shows the results computed on the dataset with a time increment of 1 millisecond, panel B – for the dataset with a time increment of 1 second, panel C – for the dataset with a time increment of 1 minute, panel D – for the dataset with a time increment of 5 minutes. The TED is shown on each subplot with days of the sample on the Y-axis and hours of the day on the X-axis.


Panel A. 1 millisecond



Panel B. 1 second



Panel C. 1 minute



Panel D. 5 minutes




Table 1. Composition of terciles with stock names, codes, market capitalisation, and volume. The figures for market capitalisation correspond to the number of shares outstanding as per 31.12.2014 times the closing price on LSE on 31.12.2014. The figures for the volume correspond to the number of shares traded from January 2 to March 31, 2015 and are retrieved from Reuters Eikon, except for 5 stocks for which we do not find any records. We start with 103 stocks forming the FTSE 100 as per 31.12.2014 and eliminate Standard Life as it experiences a stock split midway through the sample period. We further remove Tullow Oil as it drops from the index on March 23, 2015. Finally, TUI had two listings and we include only the main listing on the LSE. The remaining 99 stocks are ranked based on the market capitalisation as of 31.12.2014 forming three terciles corresponding to high, mid, and low market capitalisation respectively.

Tercile	Name	LSE RIC	Market cap (mln £)	Volume LSE (mln)	Volume Chi-X (mln)
T1	Old Mutual	OML.L	170,257.80	519.24	105.52
T1	HSBC Holdings	HSBA.L	116,950.40	1,782.18	544.29
T1	Royal Dutch Shell 'A'	RDSa.L	108,075.90	436.65	115.76
T1	BP	BP.L	74,945.25	2,336.66	725.92
T1	GlaxoSmithKline	GSK.L	66,925.06	557.45	134.41
T1	British American Tobacco	BATS.L	65,241.64	181.70	61.77
T1	Vodafone Group	VOD.L	59,023.43	3,845.69	1,105.62
T1	AstraZeneca	AZN.L	57,542.46	168.27	67.90
T1	Royal Dutch Shell 'B'	RDSb.L	54,494.38	309.92	75.97
T1	SABMiller	SAB.L	54,239.39	-	-
T1	Lloyds Banking Group	LLOY.L	54,115.53	8,472.50	2,954.55
T1	Diageo	DGE.L	46,465.50	292.33	90.55
T1	Rio Tinto	RIO.L	42,424.35	319.61	94.67
T1	Barclays	BARC.L	40,173.08	2,517.74	742.92
T1	Glencore	GLEN.L	39,250.34	2,434.23	812.73
T1	Prudential Financial	PRU.L	38,311.25	258.95	83.88
T1	Reckitt Benckiser	RB.L	37,437.84	93.99	19.13
T1	National Grid	NG.L	34,500.07	556.56	136.98
T1	Unilever	ULVR.L	33,729.30	188.09	37.77
T1	BT Group	BT.L	32,685.87	1,169.61	358.72

Tercile	Name	LSE RIC	Market cap (mln £)	Volume LSE (mln)	Volume Chi-X (mln)
T1	BG Group	BG.L	29,531.76	-	-
T1	BHP Group	BLT.L	29,326.09	576.39	153.01
T1	Imperial Brands	IMT.L	27,143.08	-	-
T1	Shire	SHP.L	26,746.41	101.98	31.94
T1	Royal Bank of Scotland	RBS.L	25,107.09	728.82	209.13
T1	Associated British Foods	ABF.L	24,961.47	48.20	19.04
T1	Standard Chartered	STAN.L	23,809.90	652.17	204.08
T1	WPP	WPP.L	17,730.05	242.05	58.79
T1	Compass Group	CPG.L	17,288.88	240.59	52.97
T1	Anglo American	AAL.L	16,767.03	396.60	94.38
T1	Rolls-Royce Holdings	RR.L	16,377.89	438.01	96.48
T1	SSE	SSE.L	16,019.26	196.82	44.88
T1	Sky Ltd	SKYB.L	15,453.96	287.00	66.57
T1	Tesco	TSCO.L	15,352.46	1,798.13	551.97
T2	BAE Systems	BAES.L	14,882.14	478.50	117.31
T2	Legal & General Group	LGEN.L	14,771.18	720.27	198.63
T2	Aviva	AV.L	14,295.02	600.38	139.20
T2	ARM Holdings	ARM.L	13,972.85	-	-
T2	Centrica	CNA.L	13,862.76	1,245.52	356.71
T2	RELX	REL.L	12,492.34	254.23	54.03
T2	CRH	CRH.L	11,437.19	218.67	36.50
T2	Experian	EXPN.L	10,772.70	137.01	35.82
T2	Smith & Nephew	SN.L	10,620.29	189.06	31.76
T2	Next	NXT.L	10,425.97	25.84	7.74
T2	International Airlines Group	ICAG.L	9,910.96	539.49	204.27
T2	Pearson	PERSON.L	9,756.60	176.53	37.56
T2	Wolseley	WOS.L	9,585.41	45.52	17.72

Tercile	Name	LSE RIC	Market cap (mln £)	Volume LSE (mln)	Volume Chi-X (mln)
T2	Land Securities Group	LAND.L	9,146.38	133.71	25.63
T2	ITV	ITV.L	8,662.68	814.62	278.94
T2	Whitbread	WTB.L	8,656.42	28.07	9.05
T2	Kingfisher	KGF.L	8,018.21	499.96	173.58
T2	British Land Company	BLND.L	7,918.49	196.97	35.22
T2	Marks & Spencer Group	MKS.L	7,829.34	338.97	106.10
T2	London Stock Exchange	LSE.L	7,705.90	52.68	15.53
T2	InterContinental Hotels Group	IHG.L	7,679.26	43.29	14.38
T2	ANTOFAGASTA	ANTO.L	7,418.57	208.85	66.24
T2	TUI AG	TUIT.L	7,381.32	86.03	15.73
T2	Burberry Group	BRBY.L	7,273.07	78.46	26.43
T2	Capita Group	CPI.L	7,154.19	180.70	40.47
T2	Johnson Matthey	JMAT.L	6,961.05	34.49	11.41
T2	easyJet	EZJ.L	6,637.35	116.53	30.96
T2	United Utilities Group	UU.L	6,246.10	131.82	34.37
T2	Schroders	SDR.L	6,070.96	25.37	8.31
T2	Bunzl	BNZL.L	5,904.23	42.47	11.02
T2	Ashtead Group plc	AHT.L	5,798.52	160.95	41.62
T2	Aberdeen Asset Management	ADN.L	5,755.23	271.07	95.78
T3	GKN	GKN.L	5,650.99	329.82	99.10
T3	Fresnillo	FRES.L	5,644.60	104.57	33.25
T3	Carnival	CCL.L	5,372.49	59.67	21.50
T3	Dixons Carphone	DC.L	5,322.76	228.08	99.58
T3	Babcock	BAB.L	5,313.24	95.83	23.63
T3	Friends Life Group Ltd	FLG.L	5,159.95	-	-
T3	Sage Group	SGE.L	5,014.58	194.49	73.94
T3	Persimmon	PSN.L	4,835.94	82.40	23.04

Tercile	Name	LSE RIC	Market cap (mln £)	Volume LSE (mln)	Volume Chi-X (mln)
T3	Severn Trent	SVT.L	4,806.98	48.31	13.46
T3	Hargreaves Lansdown	HRGV.L	4,800.10	70.01	16.67
T3	Hammerson	HMSO.L	4,744.98	171.38	28.32
T3	Sainsbury (J)	SBRY.L	4,722.09	582.35	176.87
T3	Barrat Developments	BDEV.L	4,672.05	240.68	79.27
T3	Travis Perkins	TPK.L	4,618.41	43.24	11.89
T3	Taylor Wimpey	TW.L	4,483.27	1,038.57	262.38
T3	Coca-Cola HBC	CCH.L	4,474.51	44.74	9.77
T3	RSA Insurance Group	RSA.L	4,417.37	256.47	84.44
T3	INTU Properties	INTUP.L	4,398.24	175.59	53.18
T3	3I Group	III.L	4,377.00	101.06	47.84
T3	Direct Line Insurance Group	DLGD.L	4,369.50	284.99	58.05
T3	Smiths Group	SMIN.L	4,334.40	78.01	19.46
T3	G4S	GFS.L	4,311.88	233.06	82.57
T3	Wm Morrison Supermarkets	MRW.L	4,301.21	856.25	232.40
T3	Royal Mail	RMG.L	4,299.00	183.77	48.65
T3	Frasers Group	SPD.L	4,255.08	101.53	24.30
T3	St. James's Place	SJP.L	4,230.89	88.12	23.63
T3	Meggitt	MGGT.L	4,164.09	132.90	34.68
T3	Randgold Resources	RRS.L	4,058.59	39.34	15.47
T3	Weir Group	WEIR.L	3,949.68	75.77	29.80
T3	Mondi	MNDI.L	3,856.03	84.03	21.24
T3	Aggreko	AGGK.L	3,852.02	59.00	14.98
T3	Intertek Group	ITRK.L	3,766.18	42.57	15.51
T3	Admiral Group	ADML.L	3,684.28	54.55	16.46

Table 2. Instances with infinite spread on Chi-X for each trading hour. The column “Number of stocks” reports the count of stocks showing instances with no bid or no ask quote in the respective hour. The column “Aggregated duration in minutes” reports the cumulative duration of time in minutes with no bid or no ask quote summed over all stocks that are concerned.

Hour	Number of stocks	Aggregated duration in minutes
8	99	919
9	6	44
10	11	12
11	3	1
12	9	11
13	12	13
14	8	21
15	5	7
16	99	13

Table 3. Summary statistics of number of shocks per day for a stock. A shock is defined as the event when the relative quoted spread exceeds the 95th percentile of the hourly empirical distribution of the spread per stock-day. The reported values are averages build across days and stocks of the respective bin.

time series	mean	std	0.25	0.50	0.75	0.99
LSE	204	210	69	133	243	995
LSE - High market cap	374	265	166	299	539	1136
LSE - Medium market cap	119	94	58	95	155	467
LSE - Low market cap	111	78	50	90	152	361
Chi-X	259	253	95	164	322	1196
Chi-X - High market cap	460	320	199	391	647	1359
Chi-X - Medium market cap	156	123	73	122	204	596
Chi-X - Low market cap	152	95	83	129	199	461

Table 4. Summary statistics of threshold exceedance duration. The reported values are averages of TED build across the hourly distribution of TED for each day and stock of the respective bin. Values are in milliseconds.

time series	mean	std	0.25	0.50	0.75	0.99
LSE	3,920	14,041	3	116	2,087	60,362
LSE - High market cap	2,551	7,825	3	114	1,609	36,239
LSE - Medium market cap	5,865	20,112	2	102	3,309	84,435
LSE - Low market cap	6,663	20,884	3	146	3,951	93,665
Chi-X	3,387	12,004	10	200	1,911	51,706
Chi-X - High market cap	2,159	6,839	9	168	1,360	30,624
Chi-X - Medium market cap	5,014	17,171	7	225	2,887	72,216
Chi-X - Low market cap	5,584	16,869	19	367	3,511	76,657

Table 5. Summary statistics of severity of exceedances. Instances of infinite spread are excluded from the calculation of the figures for severity. The severity of a shock is defined as the difference in relative terms between the maximum spread observed during the duration of the exceedance of the threshold and the threshold itself. The reported values are averages build across individual exceedances across days and stocks of the respective bin. Occurrences with infinite spread are excluded.

time series	mean	std	0.25	0.50	0.75	0.99
LSE	0.197	0.275	0.001	0.004	0.333	0.999
LSE - High market cap	0.183	0.262	0.001	0.004	0.331	0.998
LSE - Medium market cap	0.238	0.330	0.001	0.003	0.494	0.999
LSE - Low market cap	0.201	0.252	0.001	0.010	0.335	0.997
Chi-X	0.216	0.548	0.001	0.111	0.334	0.999
Chi-X - High market cap	0.197	0.368	0.001	0.006	0.332	0.998
Chi-X - Medium market cap	0.247	0.422	0.001	0.163	0.495	0.999
Chi-X - Low market cap	0.245	0.966	0.001	0.198	0.400	1.001

Table 6. Summary statistics of number of times the relative quoted spread changed from the time of exceeding the threshold until the time when returning below the threshold. The reported values are averages build for each day and stock of the respective bin.

time series	mean	std	0.25	0.50	0.75	0.99
LSE	1.23	1.31	1	1	1	6
LSE - High market cap	1.25	1.42	1	1	1	6
LSE - Medium market cap	1.14	0.81	1	1	1	4
LSE - Low market cap	1.26	1.36	1	1	1	6
Chi-X	1.32	1.42	1	1	1	7
Chi-X - High market cap	1.33	1.52	1	1	1	7
Chi-X - Medium market cap	1.23	1.09	1	1	1	6
Chi-X - Low market cap	1.37	1.39	1	1	1	7

Appendix A. List of sample stocks and composition of subsamples

Insert Table 1 here

Appendix B. Additional analysis

A. TED and speed of return

For comparability with existing studies employing the speed of return as a measure of market resiliency, we show how the TED measure proposed by Danielsson et al. (2018) and the speed of return proposed by Kempf et al. (2015) perform on our dataset using the example of one stock.

In the following we provide the definition for the two measures of market resiliency: threshold exceedance duration, i.e. TED and speed of reversion. The figure below shows by the means of a hypothetical stock how the two metrics, TED and speed of return, relate to each other.

Insert Figure 19 here

Kempf et al. (2015) propose measuring resiliency through a measure of speed. This measure has as an underlying assumption that the spread is a stochastic variable S for which the arithmetic Ornstein-Uhlenbeck (OU) takes the continuous form

$$dS = k(\bar{S} - S)dt + \sigma dz \quad (16)$$

The corresponding discrete form of the OU process is obtained using the first order Euler's approximation and is given by

$$\Delta S = k(\bar{S} - S)\Delta t + \sigma \Delta z \quad (17)$$

The above can be rewritten as an autoregressive process of order one, i.e. AR(1) as follows

$$\Delta S_t = \alpha - kS_{t-1} + \varepsilon_t \quad (18)$$

where ΔS_t corresponds to the change in the stochastic variable within a time increment t , k is the speed of reversion to the long-run mean, α is the intercept and ε_t is the error term.

Kempf et al. (2015) propose the following extended specification with autoregressive elements for the measurement of resiliency

$$\Delta S_t = \alpha - kS_{t-1} + \sum_{\tau=1}^N \gamma_{t-\tau} \Delta S_{t-\tau} + \varepsilon_t \quad (19)$$

where S_t is the relative quoted spread with ΔS_t as the change within a time increment t , S_{t-1} as the relative quoted spread during the last period, $\Delta S_{t-\tau}$ as the autoregressive component describing the change within τ time increments t , i.e. lags with $\gamma_{t-\tau}$ the corresponding coefficient, k is the speed of reversion, α is the intercept, and ε_t is the error term.

In order to make this metric comparable to the TED metric proposed by Danielsson et al. (2018), we adopt the following specification:

$$\begin{aligned} \Delta S_t = & \alpha_{above} d_{S_{t-1} > \theta} + \alpha_{below} (1 - d_{S_{t-1} > \theta}) - k_{above} S_{t-1} d_{S_{t-1} > \theta} \\ & - k_{below} S_{t-1} (1 - d_{S_{t-1} > \theta}) + \sum_{\tau=1}^5 \gamma_{t-\tau} \Delta S_{t-\tau} + \varepsilon_t \end{aligned} \quad (20)$$

where $d_{S_{t-1} > \theta}$ is a dummy variable taking the value 1 if the relative quoted spread at time $t-1$ was above a threshold θ , which may be the median or the 95th percentile of the empirical distribution of the relative quoted spread, and value 0 otherwise, k_{below} is the speed of reversion from below, k_{above} is the speed of reversion from above, α_{below} and α_{above} are the corresponding intercepts, and the remaining terms are as introduced above. This specification is equivalent to Kempf et al. (2015) specification employed by the authors for studying consumption and replenishment resiliency.

Kempf et al. (2015) estimate the speed of return k by means of ordinary least squares (OLS). When attempting to run this estimation, we note that on millisecond level data, the timeseries of spread change ΔS_t show a share of non-zero observations below 1%. The spread change being the endogenous variable, the sparsity in this timeseries is not compatible with the mean-reverting Ornstein-Uhlenbeck process underlying this measure. Running an OLS estimation on timeseries showing this level of sparsity is not adequate and leads to spurious results.

In conclusion, we acknowledge the parametric measurement approach of speed of reversion proposed by Kempf et al. (2015), yet find after careful analysis that it is less compatible with the nature of our data. Thus, whilst speed of reversion may perform well on datasets with a time increment of five minutes or one minute, we conclude that TED is more robust and more suitable for datasets with high frequency time resolutions.

B. Market resiliency: duration at different time resolutions

For comparison with other studies on market resiliency using lower frequency data, we examine the effect of using lower frequency data on measuring TED. For this purpose, we generate four datasets of the relative quoted spread with the following time increment: 1 millisecond, 1 second, 1 minute and 5 minutes. The choice of these time increments is made for enabling comparison with existing studies on market resiliency, where datasets of 1 second, 1 minute, and 5 minutes have been used so far. We then calculate the TED for each dataset and plot for visual analysis.

Figure 20 shows the TED for each of the four datasets for a single stock, capturing along duration also the occurrence of exceedances. The subplots of the TED at 1-millisecond time increment and 1-second time increment are visually alike. The subplots of the TED at 1 millisecond and at 1 minute show a visible difference: when measured on data with 1-minute time increment, exceedances occur significantly less often and where they do occur, the duration of exceedances is altered. The comparison of the subplots of TED at 1-millisecond and TED at 5-minutes shows that the alteration of the number of exceedances and of the duration of exceedances becomes substantially more important as the time increment increases. As such, we document considerable activity at millisecond level and reason that an analysis of market resiliency on venues and stocks with a substantial presence of speed traders requires high data resolution.

Insert Figure 20 here

Recent Issues

No. 290	Nicola Fuchs-Schündeln, Dirk Krueger, Alexander Ludwig, Irina Popova	The Long-Term Distributional and Welfare Effects of Covid-19 School Closures
No. 289	Christian Schlag, Michael Semenischev, Julian Thimme	Predictability and the Cross-Section of Expected Returns: A Challenge for Asset Pricing Models
No. 288	Michele Costola, Michael Nofer, Oliver Hinz, Loriana Pelizzon	Machine Learning Sentiment Analysis, COVID-19 News and Stock Market Reactions
No. 287	Kevin Bauer, Nicolas Pfeuffer, Benjamin M. Abdel-Karim, Oliver Hinz, Michael Kosfeld	The Terminator of Social Welfare? The Economic Consequences of Algorithmic Discrimination
No. 286	Andreass Hackethal, Michael Kirchler, Christine Laudenbach, Michael Razen, Annika Weber	On the (Ir)Relevance of Monetary Incentives in Risk Preference Elicitation Experiments
No. 285	Elena Carletti, Tommaso Oliviero, Marco Pagano, Loriana Pelizzon, Marti G. Subrahmanyam	The COVID-19 Shock and Equity Shortfall: Firm-Level Evidence from Italy
No. 284	Monica Billio, Michele Costola, Iva Hristova, Carmelo Latino, Loriana Pelizzon	Inside the ESG Ratings: (Dis)agreement and Performance
No. 283	Jannis Bischof, Christian Laux, Christian Leuz	Accounting for Financial Stability: Bank Disclosure and Loss Recognition in the Financial Crisis
No. 282	Daniel Munevar, Grygoriy Pustovit	Back to the Future: A Sovereign Debt Standstill Mechanism IMF Article VIII, Section 2 (b)
No. 281	Kevin Bauer	How did we do? The Impact of Relative Performance Feedback on Intergroup Hostilities
No. 280	Konstantin Bräuer, Andreas Hackethal, Tobin Hanspal	Consuming Dividends
No. 279	Tobin Hanspal, Annika Weber, Johannes Wohlfart	Exposure to the COVID-19 Stock Market Crash and its Effect on Household Expectations