



THE LONDON SCHOOL  
OF ECONOMICS AND  
POLITICAL SCIENCE ■



# Market Resilience

Jon Danielsson  
Efsthathios Panayi  
Gareth W. Peters  
Jean-Pierre Zigrand

SRC Discussion Paper No 78

April 2018



Systemic Risk Centre

Discussion Paper Series

**Abstract**

We propose a method to capture the notion of resilience, the dynamic aspect of liquidity in the limit order book, through the Threshold Exceedance Duration (TED) metric that we introduce. This measures the duration of liquidity 'droughts.' We illustrate the explanatory power of a survival regression framework for the duration of 'droughts' in terms of observable state variables reflecting the shape and evolution of the limit order book using Chi-X data. Finally, we introduce a method to summarise exceedance duration information across different thresholds, called Liquidity Resilience Profile, which enables the comparison and the ranking of liquidity resilience.

Keywords: Liquidity Measures, Resilience, Limit Order Book, Liquidity Provision, Optimal Trade Execution

This paper is published as part of the Systemic Risk Centre's Discussion Paper Series. The support of the Economic and Social Research Council (ESRC) in funding the SRC is gratefully acknowledged [grant number ES/K002309/1].

Jon Danielsson, Department of Finance and Systemic Risk Centre, London School of Economics and Political Science

Efstathios Panayi, University College London and Systemic Risk Centre, London School of Economics and Political Science

Gareth W. Peters, Herriot-Watt University and Systemic Risk Centre, London School of Economics and Political Science

Jean-Pierre Zigrand, Department of Finance and Systemic Risk Centre, London School of Economics and Political Science

Published by  
Systemic Risk Centre  
The London School of Economics and Political Science  
Houghton Street  
London WC2A 2AE

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior permission in writing of the publisher nor be issued to the public or circulated in any form other than that in which it is published.

Requests for permission to reproduce any article or part of the Working Paper should be sent to the editor at the above address.

© Jon Danielsson, Efstathios Panayi, Gareth W. Peters and Jean-Pierre Zigrand, submitted 2018

# Market Resilience\*

Jon Danielsson  
London School of Economics  
and Systemic Risk Centre, LSE  
j.danielsson@lse.ac.uk

Efstathios Panayi  
University College London  
and Systemic Risk Centre, LSE  
efstathios.panayi.10@ucl.ac.uk

Gareth W. Peters  
Heriot-Watt University  
and Systemic Risk Centre, LSE  
g.peters@hw.ac.uk

Jean-Pierre Zigrand  
London School of Economics  
and Systemic Risk Centre, LSE  
j.p.zigrand@lse.ac.uk

April 27, 2018

---

\*We thank Thierry Foucault, Oliver Linton and Lorian Pelizzon. We also thank the participants at the *Empirical Microstructure Workshop* organised by the Cambridge - INET Institute, at the *Computerised Trading at Low & High Frequency* workshop at the Newton Institute of the University of Cambridge, at the *7th International Conference in Computational and Financial Econometrics* in London, at the *Forecasting Financial Markets Conference* in Marseille, at *Computational Intelligence for Financial Engineering & Economics* in London, *Recent Advances in Algorithmic and High Frequency Trading* at

# Market Resilience

## Abstract

We propose a method to capture the notion of resilience, the dynamic aspect of liquidity in the limit order book, through the Threshold Exceedance Duration (TED) metric that we introduce. This measures the duration of liquidity ‘droughts.’ We illustrate the explanatory power of a survival regression framework for the duration of ‘droughts’ in terms of observable state variables reflecting the shape and evolution of the limit order book using Chi-X data. Finally, we introduce a method to summarise exceedance duration information across different thresholds, called Liquidity Resilience Profile, which enables the comparison and the ranking of liquidity resilience.

**Keywords:** Liquidity Measures, Resilience, Limit Order Book, Liquidity Provision, Optimal Trade Execution

---

UCL and at seminars at the LSE, SAFE/Goethe Universität Frankfurt, Tilburg, Banque de France and City University London for helpful comments. The support of the Economic and Social Research Council (ESRC) in funding the Systemic Risk Centre is gratefully acknowledged [grant number ES/K002309/1].

# 1 Introduction

Market resilience has received relatively little empirical attention compared to other aspects of modern computer-based and high-frequency trading environments. This is problematic, since resilience is of key interest to market participants, both as an overall characteristic of a well functioning market and as a dynamic attribute that can change dramatically over time, not least in times of market stress. In modern computer-based trading environments, such as the equities markets studied in this paper, static liquidity notions on their own are not informative any longer given the fast flickering of the order book and the relatively minor capital and inventories held by market makers. The speed of order book replenishment has replaced capital to a large extent, and the ease of trading is captured by the dynamic notion of liquidity resilience. Liquidity in such markets without dedicated market maker capital is ephemeral and dependent on algorithms to submit new limit orders, and is therefore vulnerable to a temporary malfunctioning of algorithms.

Addressing the deficiency of empirical work on dynamic liquidity provides the key motivation for this paper. This paper is very narrowly focused on resilience. It does not attempt to explain the dynamics of the entire limit order book (LOB), nor does it explain the presence or sudden disappearance of liquidity. It only looks at resiliency in the sense of explaining how quickly liquidity gets replenished following an illiquidity event, and it empirically relates this duration to the current and past state of the LOB as well as to wider market variables. This means the methods introduced herein are simple and of direct practical use, and the approach is agnostic as to the specific choice of a liquidity measure. One could implement any number of the popular candidates. Applying our methodology to Chi-X LOB data, we find that *liquidity resilience is time-varying and the state of the LOB is very informative about the level of resilience in market liquidity, as is the recent past state of the LOB and some higher latency market variables.*

Our main quantity of interest is what we term the *threshold exceedance duration* (TED), that is, for how long liquidity disappears after a shock of a certain magnitude. We explain and forecast the level of the TED through a number of covariates summarizing the current and past state of the LOB, and this is achieved through the development of a survival time model, similar in structure to that used by Lo, MacKinlay and Zhang (2002). We also relate the TED across a range of liquidity thresholds to the state of the LOB

through the notion of *liquidity resilience profiles*, or LRP in short. A LRP is a curve that tells us the best estimate of the expected TED for exceedances of illiquidity beyond ever worse thresholds. The LRP captures a stock's resilience in great granularity and can be used also to compare the resilience of multiple stocks, for instance stock *A* compared to stock *B* may exhibit liquidity that replenishes more quickly for mild exceedances but not for large exceedance etc.

In other words, we explain and forecast the duration of the aftermath of liquidity shocks given the current and past states of the book, as summarised in a small number of intuitive variates. For instance, a liquidity shock lasts much longer when the shock is due to the pulling of limit-orders (LOs) compared to the case where aggressive marketable orders walk through the book and cause a liquidity disappearance, or when the LOB is made up of a larger number of different LOs, when passive traders recently lost to aggressive traders, when ask and bid volumes are larger or older, when spreads are larger, when tick size is larger, when the time since the last exceedance is smaller, when the previous day has been quieter, when the index has more activity, when the day is a major US macro announcement day etc.

While we could just as easily have forecasted the incidence of threshold exceedances, in our view, the LOB is resilient when such exceedance events are short in duration, rather than when such exceedance events are rare. This is because their frequency pertains to market depth and the amount of new information hitting the market more than to the resilience of the book. The further advantage of modelling durations only is that the regression model is extremely flexible, simple to implement and requires very few assumptions.

## 1.1 Modelling approach

Our specific empirical approach is based on a survival time regression models utilising a number of exogenous covariates, derived from the history of the LOB. The specific empirical model is a distributed lag, auto-regressive survival model, enabling us to capture the dynamic conditional expected duration of liquidity shocks.

The number of potential covariates being extremely large, and in the absence of theoretical work that can be tested, we based our intuition on the limited number of academic papers, on common sense and on discussions with practitioners involved in electronic trading and market making.

On the face of it, it would appear that forecasting TEDs in a limit order

book ought to be very difficult since liquidity events in LOBs are complex outcomes of myriads of decisions about aggressively taking liquidity with market orders, providing LOs, cancelling LOs etc., done with very low latency by robots whose strategies interact with each other. For instance, these algorithms use machine learning tools, sniff each other out and interact with buy-side VWAPs. Some algorithms are known to be trend-following at certain scales and others are known to be contrarian at various other scales. If this is not bad enough, the algos may follow non-linear decision rules and a master robot switches various algos on and off depending on its own rules and depending on the High Frequency Trader's (HFT) position and risk limits. None of these elements are observable. And yet we find good panel fits for the expected TEDs based on a few variates capturing the past and current states of the limit order book as well as a number of market and economy wide variables, presumably because the decisions taken by the market participants and their algorithms in turn depend themselves on such variables.

This paper is not set up to model the incidence of TEDs throughout the day (even though the frequency of TEDs in the near past is a covariate in the model), but, rather, models the duration of such exceedances, once they have occurred. We do this for three reasons. First, periods of higher incidence are fairly predictable - the start of the trading day, when the market is slowly feeling its way towards what the true values can be at that time, and periods around important economic announcements. Secondly, and more importantly, the model is intended to inform decision making in the millisecond environment, where potentially hundreds of such exceedances could occur on a daily basis. In this environment the intensity is of secondary importance compared to the expected duration and our model can help liquidity motivated market participants (or, more likely, the algorithms operated by such participants) estimate when they could next expect sufficient liquidity to return to the LOB. Finally, the duration model is extremely simple and fast to implement and requires very few assumptions about the data-generating process.

## 1.2 Data and liquidity measures

The data consists of four months of Chi-X visible order book data, spanning the first four months of 2012, a total of 82 trading days. Chi-X, prior to its merger with BATS, was a pan-European multilateral trading facility (MTF). Chi-X data has a number of advantages. Over the time of analysis, it did

not have the sometimes complex circuit-breakers that could have disturbed the analysis (see Brugler and Linton (2014)). For this period, it also commanded a large market share of the European equity trading volume, in many cases comparable to the volumes traded on the primary markets. While our data set contains trading instructions for all major European markets, in the interest of tractability we focus our attention on the stocks in the CAC40 index. For each asset we were able to recreate the limit order book for we have every limit order submission, execution and cancellation, along with millisecond timestamps and individual order IDs. For our empirical analysis, we consider the first 5 levels of the LOB of price quotes on the bid and ask.

A large number of liquidity or illiquidity measures have been proposed both in the academic literature and by practitioners. While the choice of liquidity measures is often hotly debated, in this work we are agnostic as to the choice of liquidity measure since our interest is in the modelling of the lifecycle of liquidity shocks, however defined. As a practical matter in the empirical application below, we picked two frequently used measures, the spread  $\mathcal{M}_t := S_t = P_t^{a,1} - P_t^{b,1}$  as the difference between the best ask and the best bid, and the Xetra Liquidity Measure (XLM). The latter captures the market impact costs of a round-trip for a certain order size (opening and closing a position in one point of time), whereas the spread does not capture this depth dimension of liquidity.<sup>1</sup>

### 1.3 Results

We find clear diurnal patterns in the frequency of these liquidity droughts over our sample period. An exceedance of a given liquidity threshold is more likely to occur close to the start of the day, and the exceedance frequency generally decreases throughout the day. Conversations with market participants have helped us interpret this phenomenon with regard to market activity. Market makers may be more uncertain about the price of assets after the opening auction, and are generally unwilling to offer competitive spreads, but also change their quotes frequently. This eventuates in a larger number of liquidity threshold exceedances, compared to later in the day, when prices and trends have been established. We find a second cluster of exceedances

$${}^1\mathcal{M}_t := \frac{XLM_t(R)}{R} = \frac{\sum_{i=1}^k TV_t^{a,i}(P_t^{a,i} - P_t^m) + (R - \sum_{i=1}^k TV_t^{a,i})(P_t^{a,k+1} - P_t^m)}{R} + \frac{\sum_{i=1}^k TV_t^{b,i}(P_t^m - P_t^{b,i}) + (R - \sum_{i=1}^k TV_t^{b,i})(P_t^m - P_t^{b,k+1})}{R} \text{ where } k = \max(n : \sum_{i=1}^n TV_t^{a,i}(P_t^{a,i} - P_t^m) < R) \text{ and } TV_t^{b,i} = \mathbf{1}_n \cdot \mathbf{V}_t^{b,i} \text{ is the total volume at level } i \text{ and } P_t^m \text{ is the mid-price.}$$



at 13:30 UK time, 8:30 EST coinciding with the release of economics reports in the United States, eg weekly jobless claims, retail sales, core PPI, housing starts, non-farm payrolls.

When it comes to explaining the variation in the TED, with typical  $R^2$  around 20%, we find that the explanatory power of the model is good for both the durations of exceedances above median liquidity levels, as well as for higher threshold levels. The explanatory power is slightly higher when considering the spread as the liquidity measure of choice, rather than the XLM, but both produce very satisfactory results.

Most of the explanatory variables are very highly significant at all thresholds.

If one believes that market liquidity, including the bid-ask (BA) spread, is set in equilibrium to a meaningful extent by competing passive liquidity providers, resilience as defined here measures the speed by which an exceedance beyond a given  $c$  attracts new competitive LOs that re-establish the profit-maximising spreads at this level  $c$ . Liquidity resilience is therefore a function of  $c$ . If the level  $c$  is high and compensates more than commensurately for the level of informational asymmetry and other risk factors, the TED will be short, and the regressions inform us on some of the factors determining how short. If the chosen level  $c$  does not adequately compensate given the state of the LOB, then the TED will be long and may require either a period with exceptionally low informational asymmetries or some non-informational happenstance that leads algorithms to quote LOs at such low  $c$  (including hedging needs, liquidity related trades, algorithms following trading signals not related to market making, algorithmic malfunctioning, spillovers from other trading venues, or non-market makers using LOs instead of marketable orders so as not to pay the spread etc.).

Our baseline explanations in this summary therefore consider informational factors for illiquidity exceedance durations, to which a disparate set of non-informational and market-microstructure explanations will be added. Information tends to come in clusters, and we find that lagged durations are significant in explaining current durations and that the TED process is persistent and clustering, suffering from bursts of informational clusters. We find that liquidity supplying algorithms are not purely reactive, they try to tease out informational events before they fully hit the LOB and adjust spreads accordingly. Consistent with the idea that market makers pull LOs in anticipation of informational events, we show that TEDs are much longer if the exceedances are due to pulled LOs than to aggressive marketable orders.

We also find that exceedances that get resolved in a way that the mid-price at the end is the same as the mid-price prior to the exceedance, which we deem “non-informational,” get resolved more quickly than exceedances during which the mid-price moves. Similarly, TEDs tend to be longer following episodes where LO suppliers in the aggregate (we don’t have individual trader IDs) lose more money over a 10 seconds interval preceding the exceedance.

Informational factors are furthermore strengthened through market-wide variates that directly speak to informational richness. For instance, moments with more overall CAC40 market activity (including LO submissions, cancellations and amendments) are followed by longer TEDs for individual stocks. Furthermore, days with US macro announcements lead to significantly longer TEDs on the CAC40 market. Also, larger exceedances - reflecting bigger pieces of news - take longer to heal.

Consistent with the informational narrative of news coming in bursts followed by periods of relative unformativeness, we find that if the previous day has seen many exceedances or has been an announcement day, the following day brings about smaller expected TEDs.

The remainder of this paper is organised as follows: In Section 2, we discuss the notions of LOB liquidity and we provide an overview of resilience related work. In Section 3, we provide a formal definition of our resilience measure and detail the statistical framework we adopt. In Section 5, we discuss features of the dataset used in this paper and in Section 6, we present a detailed set of results and assess the model fit. Section 8 concludes.

## 2 Related work

### 2.1 Liquidity resilience

In the context of a financial market, we can consider resilience to be a quality of the market that allows various qualities to recover after a shock. The seminal paper of Kyle (1985) refers to resilience as ‘the speed with which prices recover from a random, uninformative shock’. This is similar to the interpretation of Obizhaeva and Wang (2012), who suggest that in a resilient market there is a swift convergence of the price of an asset to a new steady state, after a market order. Garbade and Garbade (1982) describe a resilient market as one in which ‘new orders pour in promptly in response to a temporary order imbalance’, whereas Harris (2002) suggests that in such a market, ‘un-

informed traders cannot change prices substantially'. These interpretations of resilient markets differ somewhat, in that the two former definitions are related to price evolution, while the latter two concern order replenishment.

We now focus on the resilience of liquidity, which itself does not have a universally accepted definition, but is only loosely understood as being related to the return to some former level of either the price (Kyle, 1985), the volume (Garbade and Garbade, 1982) or liquidity metric of interest (Foucault, Kadan and Kandel, 2005). The theoretical models of Kyle (1985) and Glosten and Milgrom (1985) suggest that liquidity fluctuations may arise as a result of information asymmetries. Dong, Kempf and Yadav (2007) note that resilience has received much less attention compared to other aspects of liquidity, citing the extensive research in depth and tightness. There have still, however, been a number of attempts to define and model resilience, including by Foucault et al. (2005). They analyze in a dynamic equilibrium model the determinants of liquidity resilience, which they define as the number of orders required for the spread to recover to a competitive level. Waiting costs are the main determinant of resilience, based on the intuition of Demsetz (1968). They identify three liquidity resilience regimes for the LOB, which they relate to the proportion of patient and impatient traders (traders that predominantly submit limit or market orders, respectively). They find that resilience increases with the proportion of patient traders, while resilience is reduced by a reduction in the tick size.

The resilience model of Large (2007) uses a parametric model which views orders and cancellations as a mutually-exciting ten-variate Hawkes point process, which formalises resilience in terms of a time-frame and probability of order book replenishment. Kempf, Mayston and Yadav (2007) use a mean reversion model of liquidity, where resilience is proportional to the intensity of the mean reversion. They find that resilience is not significantly correlated with either the spread or the depth of the LOB, highlighting the need for separate treatment. Finally, similar to this paper, Gomber, Schweickert and Theissen (2011) employ the XLM measure to quantify the resilience of liquidity, which they define as the difference in the XLM after large transactions. They find that the liquidity measure generally returns to close to its pre-trade level within 2-3 minutes of a large transaction. Notions of resilience have been incorporated in the dynamic LOB model of Roşu (2009).

While these previous attempts to capture liquidity resilience have exhibited the importance of the concept, they have generally tied the definitions to particular measures of liquidity, like the inside spread (Foucault et al., 2005),

or the volume at the top of the LOB (Large, 2007). Our first contribution is a more general definition of resilience as the duration of deviations from a particular liquidity threshold level. Different liquidity measures and threshold levels of liquidity are appropriate for different applications (for example, an algorithmic execution setting would have different ‘trigger’ level to a setting where liquidity is monitored for systemic risk purposes). Our definition can accommodate any liquidity measure and threshold level of interest.

In addition, the aforementioned papers have generally tried to model the concept through theoretical processes. Several authors have studied resilience empirically, such as Degryse, De Jong, Van Ravenswaaij and Wuyts (2005); Large (2007). Degryse et al. (2005) present an event study on the Paris Bourse, estimating the impact of aggressive orders through a non-parametric approach. Coppejans, Domowitz and Madhavan (2004) use a vector autoregressive model to estimate whether liquidity shocks dissipate swiftly. Large (2007) studies the intensity of order replenishment after liquidity has been consumed on the LSE using a multivariate Hawkes process. Compared to these papers, our use of a regression model with covariates coming from the structure of the LOB allows us to assess the contribution of each covariate to the explanatory power of the model. It also gives us the ability to model different scenarios that had not previously occurred in the dataset, by modifying the covariate values. We also assess the predictive power of our model, and explain how such a forecasting model could then be readily incorporated into an execution model.

Our work complements the literature in determining optimal order sizes of e.g. Bertsimas and Lo (1998) and Obizhaeva and Wang (2012), as it provides an estimated time between subsequent tranches, given a certain state of the LOB. We show that resilience is neither infinite, as assumed in the former, nor constant throughout the trading day, as in the latter, but rather relates to the level of structural variables of the LOB. While Alfonsi, Fruth and Schied (2010) has considered the exponential decay of the impact of market orders on LOB volume and the spread, we show that resilience is dependent on the state of the LOB, and that incorporating LOB variables can improve the explanatory power of the model.

Our notion of liquidity resilience therefore extends the standard notion of resilience in three key dimensions, first by making it time-varying and showing that a substantial part of its variation can be explained by the structure of the LOB; second by explicitly relating it to certain liquidity levels of interest, and thirdly, by modelling and forecasting the time required

for a market to recover following a liquidity shock. We show that the state of the LOB is an important determinant of the level of liquidity resilience, while previous work by Kyle (1985) and Foucault et al. (2005) only considered exogenous factors to be relevant (informational asymmetries and waiting costs, respectively).

This approach enables us to provide a comparison of resilience, for example, in what could be considered to be ‘normal’ and ‘stressed’ scenarios, enabling us to classify assets according to whether they show similar behaviour in terms of these exceedances in particular market conditions, and finally, study the behaviour of resilience across time.

We also contribute to the empirical understanding of resilience, as our methodology enables us to identify the market conditions that are most likely to result in an environment of slow liquidity replenishment. Developing such an understanding is key to implementing measures to reduce the duration and/or severity of liquidity droughts, as these droughts have been shown to contribute to the amplification of small shocks into full-blown financial crises Brunnermeier (2008). At the same time, by forecasting the return of liquidity to the markets following an event, our methods are also of great interest to practitioners involved in optimal order execution.

## 2.2 Survival analysis

Survival analysis models the time until a particular event (or events) will occur, such as the failure of some component or the death of an individual. It has a unique advantage in that it can incorporate censored observations, which occur when one cannot obtain the true value of an observation because the terminating event has not occurred inside the observation window. It has been used frequently in medical studies, for example, in modelling the time to a single outcome, such as the lifetime after a kidney transplant (Lambert, Collett, Kimber and Johnson, 2004), or the time to one of a range of outcomes, such the time to response to treatment or the time to recurrence in cancer patients (Bradburn, Clark, Love and Altman, 2003).

In the context of the LOB, survival regression frameworks have mainly been employed to explain the variation in the lifetimes of limit orders. Al-Suhaibani and Kryzanowski (2000) used such a framework for the Saudi Stock Market, under an assumption of a Weibull specification for the time to execution of limit orders, where cancelled or expired orders were considered to be censored. A similar formulation was used by Cho and Nelling (2000)

for the New York Stock Exchange. Lo et al. (2002) used a more general formulation with a Generalised Gamma Distribution, similar to that used in this paper, particularly regarding the transformation of the structure of the LOB to a set of summary statistics.

### 3 A model of liquidity resilience

Within the extensive empirical market microstructure literature on liquidity, the study of market resilience has received relatively little attention. In historical LOB trading environments where trading was done by human beings, resilience might have been less of a concern than in the current computer-based and high-frequency trading environments, because algorithms might react differently to anomalies than humans and because algorithms can be optimised to trade quickly and keep risk parameters and capital usage to tight bounds. Since algorithms and resilience are more important than before, it makes studying resilience more important and that provides, the main objective of our work.

We extend existing work that mostly assumes that liquidity is a static concept (see e.g. Bertsimas and Lo, 1998; Obizhaeva and Wang, 2012; Alfonsi et al., 2010), arguing that resilience varies throughout the trading day in partially predictable way, based on the state of the LOB. In this, our work resembles the time-dependent resilience used in the execution model of Alfonsi, Fruth and Schied (2008), but we further relate resilience to the level of structural variables of the LOB.

We start by introducing a new notion of liquidity resilience that we term *threshold exceedance duration* (TED), which is based on the duration of liquidity shocks exceeding a threshold — threshold exceedance. In this, our main interest is on the duration of threshold exceedances rather than the incidence of threshold exceedance. The reason is that in our view, the LOB is resilient when such exceedance events are short in duration, rather than when such exceedance events are rare, because their frequency pertains more to market depth and the amount of new information hitting the market, rather than to the resilience of the book and the quality of liquidity provision. This is in contrast with intensity models, designed to answer different questions, such as the ones proposed by Large (2007) and Toke (2011), where they require a description of the behavior of the entire stock price process.

The literature has identified a number of different notions of market liq-

uidity, all with their particular pros and cons (see for instance the textbook by Amihud, Mendelson and Pedersen (2013) or Foucault, Pagano and Röell (2013)), and our TED framework allows the inclusion of most arbitrary liquidity notions.

Each liquidity threshold exceedance event has a lifetime, the TED, which we estimate and forecast with a statistical survival model. The model incorporates a regression component, which serves to capture the state of the LOB at the point of exceedance. One of the interesting applications of the model is the ability to forecast extreme TEDs, which can be thought of as extended periods of relative illiquidity. While the survival regression framework can be used to estimate expected TEDs over a particular threshold, a quantile regression framework can additionally produce upper and lower probabilistic bounds on the TED. These can be thought of as a confidence interval around the expected TED, and will be analyzed in a follow-on paper.

The TED model is not specific to any particular threshold level of liquidity. Indeed, some end-users may be concerned with very high-frequency liquidity behaviour and hence very low thresholds, others are more interested in the recovery of liquidity after larger and less frequent events, for example fund managers placing large orders, while a regulator might only be interested in extreme disruptions to liquidity supply, and hence exceedances over a very large threshold. For this reason, we estimate the model across a range of thresholds, terming the joint relationship between thresholds and the TED as *liquidity resilience profiles*, or LRPs.

One benefit of summarizing resilience behaviour in the form of LRPs is that one can then directly compare liquidity resilience across assets. In this way, one can potentially group assets according to their similarity in the speed of return to high levels of liquidity after a shock, which may be an indication of the presence of high-frequency market makers in these assets.

### 3.1 Threshold exceedance duration — TED

Our starting point is the TED, which is formally defined by:

**Definition 1.** *The threshold exceedance duration (TED) is the length of time between the point at which a liquidity measure,  $\mathcal{M}_t$ , deviates from a threshold liquidity level,  $c$ , (in the direction of less liquidity), and the point at which it returns to at least that level again. The random starting time and duration (in ms) of the TED of the  $i$ th upcrossing in a given trading day are denoted*

by  $T_i$  and  $\tau_i$  respectively, where  $i$  refers to the  $i^{\text{th}}$  exceedance. If we want to make the relationship between the TED and the threshold explicit, we write  $\tau_i(c_j)$ .

In a LOB where liquidity is resilient at particular threshold level, we then expect TED to be low with the market quickly returning to good levels of liquidity after a shock. The unconditional average of the TED therefore represents the typical level of liquidity resilience for an asset while the TED at any given time captures the particular level of resilience given overall market conditions.

**Figure 1.** Threshold exceedance duration (TED)

The figure shows how a particular liquidity measure,  $\mathcal{M}_t$ , evolves over time for two exceedance thresholds, with  $c_{\text{high}}$  more extreme than  $c_{\text{low}}$ , and the duration time,  $\tau_2(c_{\text{high}})$  and  $\tau_1(c_{\text{low}})$ .




Figure 1 shows a hypothetical market whereby  $\mathcal{M}_t$  fluctuates over time, exceeding a low threshold,  $c_{\text{low}} = 3$  at time 1, and a higher threshold,  $c_{\text{high}} = 6$  at time 8. In these cases, the first TED for the lower threshold is  $\tau_1 = 3$  and the second is  $\tau_1 = 4$ , where the exceedances happen at times  $T_1 = 1$  and  $T_2 = 8$ , respectively.

More formally, an exceedance (or birth) occurs at time  $T_i$ , when the liquidity measure  $\mathcal{M}_t$  deviates from a threshold,  $c_j$ , in the direction of lower liquidity for the  $i^{\text{th}}$  time. The exceedance ends (or dies) at time  $T_i + \tau_i$ , when liquidity metric returns below the threshold level. Mathematically, we can



define both the TED and the exceedance time as:

$$\tau_i(c_j) = \inf \{ \tau : \mathcal{M}_{T_i+\tau} \leq c_j, \tau > 0 \} \quad (1)$$

$$T_i(c_j) := \inf \{ t : \mathcal{M}_t > c_j, t > T_{i-1}(c_j) + \tau_{i-1}(c_j), t > T_0 \} \quad (2)$$

## 3.2 Liquidity thresholds

Our empirical study is based on the time a liquidity measure,  $\mathcal{M}_t$ , spends exceeding a particular liquidity threshold. We denote the threshold by  $c_j$  where  $j$  refers to the threshold level. We can define these thresholds in different, but interchangeable ways. In particular, we can either define them based on the absolute level of a liquidity level (for example, 3 ticks in the case of the spread) or in a relative way, where the threshold is obtained from the frequency distribution of the values the liquidity measure can take (for example, the median daily spread). There is no theoretical reason to pick either approach and since these, of course, uniquely map to each other, it is best to let empirical considerations determine the choice.

We opted for the relative approach, primarily because that enables us to capture the entire set of possible thresholds efficiently whilst absolute thresholds might span the actual liquidity of the asset under question inefficiently in terms of the number of unique values one needs to consider. Furthermore, the choice of relative liquidity levels facilitates the comparison of liquidity resilience profiles across assets, by reducing the range of liquidity values for different assets to a common set of thresholds.

In practice, we start by empirical observations of the chosen liquidity measure,  $\mathcal{M}_t$ , on a particular day and use percentiles of its empirical distribution as liquidity thresholds. Therefore, for example,  $c_{.50} := c_{.50}(\text{day } d)$  and  $c_{.90} := c_{.90}(\text{day } d)$  denote the constant 50<sup>th</sup> and 90<sup>th</sup> percentiles respectively of the exceedance durations on the given day  $d$ . Recomputing the quantiles each day allows us to interpret the thresholds in a coherent way across the sample. For instance, when computing the expected TED at a moment in time beyond  $c_{.90}$  we get an estimate of an expected duration that we know is rather serious. If we kept the quantile constant in absolute terms, then the conditional expected TED at the fixed  $c_{.90}$  at that moment may be meaningless and misleading as it would not express a large but not catastrophic exceedance on the day if markets find themselves in a much more eventful or very quiet day.

### 3.3 Survival model

A natural way to estimate the conditional expectation of the TED in (1) is by a survival regression where the probability of TED exceeding a certain duration is forecast by a survival model that uses a set of covariates as exogenous explanatory variables. Some of the covariates are derived from the state of the LOB. This section has as a sole purpose to motivate the linear survival regressions we run of the log duration of exceedances on explanatory variates and the convenient interpretation of the regression coefficients.

Survival regressions have a number of advantages. They explain the duration of exceedances in a tractable fashion since exceedances are modelled directly, given an exceedance, as opposed to the outcome of a larger model. They are designed to answer this precise question, and do not also attempt to model the entire data generating process (e.g. general intensity based specifications), including the frequency of exceedances. This makes the approach useful also in practice. Say an agency broker tries to slice and dice an order and optimises the feeding of the orders, then a simple regression-based approach allows the broker to better time its orders.

Several authors have previously adapted survival models for modelling the lifetimes of limit orders, for example Lo et al. (2002), and we adopt some of the techniques, particularly regarding the transformation of the structure of the LOB to a set of summary statistics.

In a nutshell, survival regression modelling aims to explain the variation in a dependent observed positive random variable (or in our case multiple random variables) as a function of explanatory covariates. We seek to model the durations  $\tau_i(c_j)$  and their survival function is

$$S(\tau; \boldsymbol{\beta}) := 1 - F(\tau; \boldsymbol{\beta}) = \Pr(\tau_i(c_j) \geq \tau)$$

Here  $\boldsymbol{\beta}$  is a vector of coefficients in the regression model, which parameterise the survival distribution. Two most commonly used classes of survival models are proportional hazards models and the accelerated failure time (AFT) models. AFT models can be viewed as those representable with a location coefficient that is log-linear in the chosen variates and whose survival function between populations satisfy the relationship 3 below. We use the latter approach in this paper, as AFT models have a number of relevant advantages:

- The log-linear formulation of such models (since variates act linearly on

the log of the failure time) emphasizes that the roles of the regression parameters and dispersion parameters are clearly separated.

- The regression parameters in an AFT model are also robust towards neglected covariates.

The AFT model has the distinctive feature that the model covariates have a multiplicative influence on the survival time. A unit change in, say, covariate  $x^{(k)}$  to  $x^{(k)} + 1$  produces generic survival times  $\ln \tau_1^{\text{TED}}$  and  $\ln \tau_2^{\text{TED}}$  respectively that will satisfy the following relationship between the two survival functions  $S_1(\tau)$  and  $S_2(\tau)$ :

$$S_1(\tau) := \mathbb{P}r(\tau_1 \geq \tau) = S_2(e^{\beta_k \tau}) := \mathbb{P}r(\tau_2 \geq e^{\beta_k \tau}) \quad (3)$$

The conditional means will for instance be related by  $\mu_2 = e^{\beta_k} \mu_1$  and the quantiles at any level  $p \in [0, 1]$  will be related by  $\frac{q_2 - q_1}{q_1} = e^{\beta_k} - 1 \approx \beta_k$ . These characteristics will allow us to interpret the regression coefficients with ease. The AFT regression framework thus relates the exceedance times at each threshold on a combination of common covariates, through a rescaling of time.

Within this modelling approach we can choose the very flexible three parameter distributional family that is the Generalised Gamma distribution (hereafter g.g.d.). We assume that the TED random variables are conditionally independent, given the LOB covariates:

$$\tau_i(c_j) \stackrel{i.i.d.}{\sim} F(\tau; k, a, b) = \frac{\gamma\left(k, \left(\frac{\tau}{a}\right)^b\right)}{\Gamma(k)}$$

with the incomplete gamma function defined as:  $\gamma(x, y) = \int_0^y t^{x-1} e^{-t} dt$ . The g.g.d. family includes the exponential model ( $b = k = 1$ ), the Weibull distribution (with  $k = 1$ ), the Gamma distribution (with  $b = 1$ ) and the Lognormal model as a limiting case (as  $k \rightarrow \infty$ ). The resulting density for the generalised gamma distribution is analytic and given by

$$f_{\tau}(\tau; k, a, b) = \frac{b}{\Gamma(k)} \frac{\tau^{bk-1}}{a^{bk}} \exp\left(-\left(\frac{\tau}{a}\right)^b\right)$$

with parameter ranges  $k > 0$ ,  $a > 0$  and  $b > 0$  and a support of  $\tau \in (0, \infty)$ . One can write the survival function explicitly in closed form.

Now relate this statistical model assumption to a set of explanatory variables (covariates) from lagged values of the LOB. Under the AFT framework, the regression structure involves constant and nonstochastic terms  $k$  and  $b$  as well as the following loglinear form for the time-varying location coefficient  $a(\mathbf{x}_t)$ :

$$a(\mathbf{x}_t) = \exp \left( \beta'_0 + \sum_{s=1}^p x_t^{(s)} \beta_s \right). \quad (4)$$

Each of the covariates is a transform from the LOB (current or past) for which the illiquidity measure is observed or from market conditions more broadly (current or present).<sup>2</sup> Under this AFT model with this location regression structure, we observe that the conditional mean of the survival times is also related directly to this linear structure where for the  $i$ -th exceedance of the threshold, we have (see Lo et. al. (2002)):

$$\mathbb{E}[\tau_i(c_j)|\mathbf{x}_{T_i}] = a(\mathbf{x}_{T_i}) \left( \frac{1}{k} \right)^{\frac{1}{b}} \frac{\Gamma(k + \frac{1}{b})}{\Gamma(k)}$$

Next we relate this statistical model assumption to a set of explanatory variables (covariates) from state of the LOB. To achieve this, it is beneficial to work on the log scale with  $\ln(\tau)$ , i.e. with the log-generalized gamma distribution (hereafter l.g.g.d.): this parameterisation improves identifiability and estimation of parameters. A discussion on this point is provided in significant detail in Lawless (1980). It follows that the Log Weibull is extreme value distributed, log lognormal is normal, etc.

One great advantage of the current approach is that for elliptical distributions the conditional expectations function for log-durations is in fact linear

$$\mathbb{E}[\ln \tau_i(c_j)|\mathbf{x}_{T_i}] = \beta_0 + \sum_{s=1}^p x_{T_i}^{(s)} \beta_s$$

and not merely a linear approximation to the true CEF that we are seeking. While we have run the regressions with a variety of distributions for the error terms,<sup>3</sup> for the CEF little is gained and some simplicity lost by using

---

<sup>2</sup>We note that we also considered models with interactions between the covariates, but interaction terms were not found to be significant in the majority of our models.

<sup>3</sup>We did estimate a version of the general AFT model allowing for such innovations. However, the estimated gamma coefficients did suggest that the innovation distribution

distributions other than the normal, and the results in this paper, pertaining mainly to conditional expectations and not to tail risk properties and the like, are all derived in the Gaussian case.

With an elliptical distribution such as the normal, conditional expectation functions of log-duration are affine and the estimated AFT model boils down to affine regressions in log-durations (see Lawless (1980))<sup>4</sup>:

$$\ln(\tau_i^{(\ell)}) = \beta_0 + \sum_{s=1}^N x_{T_i^{(\ell)}}^{(\ell,s)} \beta_s + \varepsilon_{T_i^{(\ell)}}^{(\ell)}, \quad \varepsilon_{T_i^{(\ell)}}^{(\ell)} \sim N(0, \sigma^2) \quad (5)$$

where we denoted securities by  $\ell$  and the  $N$  covariates (including fixed effects dummies) by  $(x_{T_i^{(\ell)}}^{(\ell,1)}, \dots, x_{T_i^{(\ell)}}^{(\ell,N)})$ .

We assume that for any threshold  $c_j$  that we choose, the liquidity measure will eventually return to this level after an exceedance at time  $T_i$ . This assumption is necessary, in order to ensure that the density  $f$  for the TED observations normalizes to unity on its support. We impose an upper bound  $\tau_i \leq T_D - T_i$ , where  $T_D$  denotes the end of the observation window. If a return of the liquidity measure  $\mathcal{M}_t$  to the threshold  $c$  has not occurred by time  $T_D$ , we consider the observation censored. Censored observations are accounted for separately in the model estimation. Each  $\tau_i$  therefore has a potential maximum observation time  $T_D - T_i$ , i.e. the time remaining until the end of the observation period  $T_D$  (1 minute prior to the end of the trading day, after which the spread will not be calculated). These maximum observation times are thus known at the time of birth.

## 4 Covariates

The survival model is estimated with a number of conditionally exogenous covariates that are obtained from four sources. First we use variates that capture the salient features of the *current state* of the LOB. Second we use

---

was at least approximately normal and certainly symmetric. By estimating the log-duration model with normal innovations, the model residuals were by and large symmetrically distributed, 99% empirical quantiles the skewness in  $[-0.46, 0.74]$ , with the 99% empirical quantiles of the residual kurtosis in  $[1.8, 3.7]$ . For this reason, we opted to use normal innovations in the estimation.

<sup>4</sup>In much of the applied work, AFT models are *defined* by the log-linearity of duration in the variates plus the noise term, even if the noise term is not elliptical.

variates that summarise the state of the book in the *recent past*, allowing for autoregressive components. Third, we look at the *event-specific variables* obtaining during the exceedance, and lastly, we use variates that are not LOB related and that represent the *wider market conditions*, such as overall market activity, macro announcements etc.

The particular model requires us to specify the distribution of the random shock, the conditional mean structure of the TED and the variance equation. We assume the TED random variables are conditionally independent, given the LOB covariates.

In the following, a ‘level’ of the LOB is defined as one in which there is at least 1 resting limit order. Thus the first 5 levels of the bid are the 5 levels closest to the quote mid-point, where there is available volume for trading.

Covariates pertaining to the shape of the LOB Model (sampled at  $T_i$  just after exceedance):

**Ask** is the total number of different ask orders (not the amount of stock bid or offered) in the first 5 levels of the LOB at time  $t$ , obtained according to  $\sum_{i=1}^5 |V_t^{a,i}|$  (where  $|\cdot|$  is the number of orders at a particular level)

**Bid** is the total number of bids in the first 5 levels of the LOB at time  $t$  is obtained according to  $\sum_{i=1}^5 |V_t^{b,i}|$

**askVolume** is the total ask volume (in 1000s of shares) in the first 5 levels of the LOB at time  $t$ , obtained according to  $\sum_{i=1}^5 TV_t^{a,i}$

**bidVolume** is the total bid volume (in 1000s of shares) in the first 5 levels of the LOB at time  $t$ , obtained according to  $\sum_{i=1}^5 TV_t^{b,i}$

**bidModified** is the number of bids in the first 5 levels of the LOB that had received price or size revisions (and were thus cancelled and resubmitted with the same order ID)

**askModified** is the number of asks in the first 5 levels of the LOB that had received price or size revisions

**bidAge** is the average age (in m) of bids in the first 5 levels at time  $t$

**askAge** is the average age (in m) of asks in the first 5 levels at time  $t$

**spreads** is the instantaneous value of the spread at the point at which the  $i$ -th exceedance occurs, which is given by  $P_t^{a,1} - P_t^{b,1}$

**Ticksiz**e is either 0.1, 0.5 or 1 cent and measures the tick size regime in which a stock finds itself at the moment of exceedance

The following are variates capturing the local dynamics of the LOB:

**lask**, **lbid** For the previously defined covariates, we also include exponentially weighted lagged versions. For example, in the case of the  $x_t^{(s)}$  covariate, the respective lagged covariate value is then given by:

$$z_t^{(s)} = \sum_{n=1}^d w^n x_{t-n\Delta}^{(s)} \quad (6)$$

where for a time  $t$ , we consider  $w = 0.75$  is the weighting factor,  $d = 5$  is the number of lagged values we consider and  $\Delta = 1s$  is the interval between the lagged values. These covariates are hereafter denoted with the ‘l’ prefix, e.g. **lask** =  $\sum_{n=1}^d w^n \text{ask}_{T-n\Delta}$ .

**prevexceed** is the number of previous TED observations in the interval  $[t - \delta, t]$ , with  $\delta = 1s$ .

**timelast** is the time since the last exceedance (in m).

**log.duration**, ten layers of past  $\ln(\text{TED})$  in ms,  $\ell = 1, \dots, 10$ .

The event specific variates are:

**MObuy** is a dummy variable indicating if the exceedance occurred as a result of a marketable order to buy.

**MOsell** is a dummy variable indicating if the exceedance occurred as a result of a marketable order to sell.

**cancelbuy** is a dummy variable indicating if the exceedance occurred as a result of a cancelled buy limit order. There is no **cancelsell** variable since the sum of **mobuy**, **mosell**, **cancelbuy** and **cancelsell** would be colinear with the constant.

**samemidsread** is 1 if

$$\mathcal{M}_{T_i+\tau_i} = \mathcal{M}_{T_i-}$$

and

$$P_{T_i+\tau_i}^{mid} = P_{T_i-}^{mid}$$

and 0 otherwise, i.e. it's 1 if the LOB returns to the same mid and spread (and thus the spread could not have tightened due to the unexecuted portion of an order).

One use of this is in differentiating between exceedances of the following sort, that one may perhaps label “informational”<sup>5</sup>

$$\begin{array}{cc}
 * & * \\
 * & * \\
 * & * \\
 T_i & T_i + \tau_i
 \end{array}$$

and between uninformed imbalance exceedances

$$\begin{array}{ccc}
 * & & \\
 * & & * \\
 * & * & * \\
 T_i & T_i + \tau_i &
 \end{array}$$

This variate is not adapted to the information filtration of the agents. Section 6.3.1 will provide a more detailed explanation.

**profLO** captures the few last (undiscounted) profits of the passive traders when trading against market orders. The intuition is that if passive traders have over the last minutes lost out to aggressive traders, they may be less willing to offer tight liquidity quickly. We use the methodology of Bessembinder (2003). Define  $\Delta_1 > \Delta_2 \geq 0$ , and  $I(t) := \{s \in [t - \Delta_1, t - \Delta_2] : D_s \neq 0\}$  and  $D_s = \begin{cases} 0 & \text{if no LO gets hit at } s \\ 1 & \text{if the LO supplier sold} \\ -1 & \text{if LO supplier bought} \end{cases}$

Then **profLO** :=  $\sum_{i \in I(t)} D_i \times \{(\text{average price achieved for that MO } i) - M_t\} \times (\# \text{ shares executed by that MO } i)$ .

If the metric **profLO**  $\ll 0$  that means that on average a series of market orders came in and the sign of  $D_i$  on average was the opposite of

---

<sup>5</sup>Here a buy MO comes in and lifts the ask, or the LO ask is cancelled and replaced by a higher one, and subsequently the bid is cancelled and placed back higher up to close the spread. The mid is moving up.



$\{(\text{average price achieved for that MO } i) - M_t\}$ . So the buy orders have been on average executed at ask prices below the final mid at the time of the exceedance and sell orders at bid prices above the final mid. The mid  $M_t$  at the exceedance time may be argued to reflect the information at that stage better and proxies for the underlying true value at time  $t - \Delta_1$ . In other words, the limit order submitters on average sold at a loss to market order traders over the last  $\Delta_1 = 10s$ .

Other market related (and not related to the current LOB) variates are:

**lag.Return** is the return for that asset the previous day.

**indexact** is the activity in the entire CAC40 index in the previous second (in 1000s of limit order submissions, amendments, cancellations, executions), comprising the total number of events (limit orders, executions, cancellations). This also capture the sense of “volatility” (which if defined through execution prices would not be insightful at such high frequency).

**lag.low.volume** = 1 if volume across all French stocks on all markets (not only on Chi-X, gathered via Bloomberg) was low, ( $\leq 15\%$ ) else 1

**lag.high.volume** = 1 if this volume was high, ( $\geq 85\%$ ) else 1

**Announcement** = 1 for a stock if there is a US macro announcement today at 13:30 London time (8:30 ET), such as weekly jobless claims, retail sales, core PPI, housing starts and non-farm payrolls.

**lag.Announcement** = 1 if there was such an announcement yesterday

**num.lag.TED** is the number of TEDs yesterday at yesterday’s  $c$  corresponding to the same quantile

## 5 Data

The data is provided by Chi-X, which (prior to its merger with BATS) was a pan-European multilateral trading facility (MTF) and commanded a large market share of the trading volumes for the assets traded on it. Indicatively, for the week starting the 2nd of January 2012, it had 25.5% of the CAC40

trading volume, 26.9% for the stocks on the DAX index and 33.4% of the FTSE100<sup>6</sup> Chi-X did not operate with circuit breakers during the trading sample, which makes Chi-X data particularly clean for the purposes of our study given the number of complex circuit-breaker mechanisms operated by other exchanges (see e.g. Brugler and Linton, 2014).

The Chi-X trading platform enables market participants to post limit orders, with a specified price and size, or several types of pegged orders, which automatically adjust their prices according to market conditions. It also allows the posting of iceberg orders, only a portion of which is displayed. For every order, there is an option to specify a ‘time in force’ after which if an order is still resting in the LOB, it will be cancelled.

The exchange has both a visible and a hidden order book and orders are routed to each book according to the type and size of the order: Limit orders, pegged orders and part of each iceberg order is displayed in the visible book; orders meeting MiFID large in scale requirements are routed to the visible book, but remain non-displayed; orders in the hidden order book are executed at the mid-price.

We use an 82 day trading sample (January 2nd to April 27, 2012) of all limit order submissions, executions and cancellations in the visible order book, focusing our attention CAC40 stocks. Both limit order submissions and executions in our dataset may be the result of pegged, limit or iceberg orders, however, the data only indicates the resulting submission of the limit order. In addition, a cancellation may be automatic (as a result of a time in force option), or as a result of a manual cancellation request, but this is not indicated in the data. We do not attempt to infer this information here, and in any case we have sufficient information to rebuild the LOB without it.

The TED observations considered are those occurring between 08:01 and 16:29 London time daily, to avoid market opening and closing effects. We also note that while the continuous trading hours on Chi-X are not necessarily the same as those in the national exchanges where the assets trade, for these French stocks the opening hours coincide. Hence, we do not have any additional considerations that would result from the sudden submission or withdrawal of liquidity from the primary exchange.

We also investigated whether liquidity resilience is affected by various economic factors such as economic announcements and recent market activity.

---

<sup>6</sup><http://www.liquidmetrix.com/LiquidMetrix/Battlemap>.

## 5.1 Diurnal patterns

It is well documented that there are strong diurnal patterns in limit order trading environments, for example on volume and even volatility. The presence or absence of diurnal patterns for the TED is relevant to how one constructs the empirical model.

As an example, consider Cr dit Agricole in Figure 2. All stocks we considered exhibited such a pattern.

The figure shows that there are clear diurnal patterns in the frequency of these liquidity droughts over our sample period. A liquidity drought is more likely to occur close to the start of the day, and we note a second concentration of exceedances around 13:30 London time, or 08:30 ET. While the concentration of liquidity droughts in the morning may be due to well documented market open effects (see, e.g. Biais, Hillion and Spatt (1995)), this does not explain the very distinct mid-day clustering.

Figure 3 shows the aggregate proportion of time that a threshold was exceeded in each half hour interval of the trading day, for a range of thresholds. In other words, it shows the relative illiquidity ratio over the day for three threshold levels, corresponding to exceedances over the median liquidity level ( $c_{.50}$ ), as well as the 90<sup>th</sup> and 99<sup>th</sup> thresholds ( $c_{.90}$  and  $c_{.99}$ , respectively). For the most extreme threshold level,  $c_{.99}$ , only the 11 most liquid stocks provided enough observations to calculate the aggregate time.

The more extreme the threshold level, the more extreme the diurnal pattern.

## 6 Regression Results and Discussion


While the TED approach is agnostic as to the choice of liquidity measure, for presentational purposes we need to focus on only a selection, and we opted to focus our attention for most parts on the spread, with some results for the XLM (Xetra Liquidity Measure). It would be straightforward to consider other liquidity measures.

### 6.1 Model estimation choice

We estimate the duration model as a panel. We have performed the regressions both with fixed effects and without. The fixed effects we have considered

**Figure 2.** Crédit Agricole  $c_{.90}$  TEDs

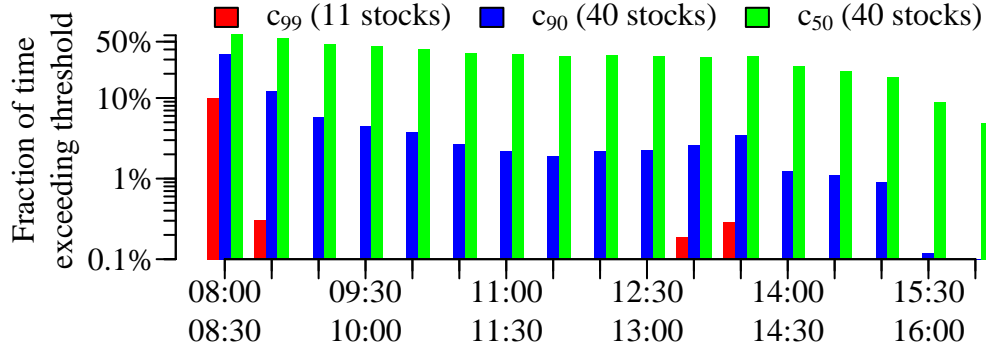
The duration of time during all the trading days in our dataset that the inside spread  $\mathcal{M}$  of Crédit Agricole (stock symbol ACAp) is above the 9th decile threshold. Time is on the x-axis starting at 8:01 am in the morning and ending at 16:29. Each day corresponds to one row on the y-axis. Rows are coloured red if there is a US economic announcement at 13:30 London time (08:30 ET), and blue otherwise.



are a stock specific dummies, day dummies as well as time-bucket-of-the-day dummies. Neither the estimated coefficients of the explanatory variables nor adjusted  $R^2$ s are much affected when adding or omitting the fixed effects. Stock fixed effects are mostly positive and statistically significant at the 0.1% level. Many daily dummies - though by no means all - are significant and

**Figure 3.** Intra-day threshold exceedances

The fraction of time during every half hour interval of the trading day that the spread exceeds thresholds  $c_{.50}$ ,  $c_{.90}$  and  $c_{.99}$ .



the signs of their coefficients are varying with the days. In the table below we exhibit the final fixed effects that capture the diurnal patterns where we split the day into 8 buckets. As expected, the values are positive and largest for the buckets at the start of the day, with a slight upwards hump for the bucket that captures the US announcements.

The quantile-based thresholds  $c$  are re-estimated every day for each stock.

Adjusted  $R^2$  are inverse U-shaped in  $c$  when  $c$  moves from  $c_{.10}$  to  $c_{.99}$ , with  $R^2$ s moving from .170 to .205 for intermediate  $c$  and down to 0.187 for  $c_{.99}$ .

Notice that the interpretations for small quantiles and for large quantiles are different. For instance, the exceedance duration beyond  $c_{.99}$  signifies the duration of illiquidity being very severe before reverting below this level  $c_{.99}$ . We expect that duration to be small since limit order suppliers are expected to replenish the LOB. The duration of an exceedance beyond  $c_{.10}$  is expected to be large since it measures the duration during which market illiquidity, having exceptionally been nearly zero, went slightly worse before returning to the already exceptionally high level of liquidity. The latter duration may well be useful to a trader who wishes to time his trades to happen during exceptionally liquid periods for instance, but TEDs beyond large quantiles such as  $c_{.90}$  and  $c_{.99}$  are probably more interesting to most since they measure the duration of a really bad episode before returning to a better (but still poor) level.

The regression results below are from a panel regression with all fixed

effects.

## 6.2 Regression results

We show the panel regression results in Table 1 below, not showing the firm fixed-effects and the day fixed-effects. The color coding is: 0.1% significant, 1% significant, 5% significant, not significant.

**Table 1.** Regression coefficients with  $\ln(TED)$  as a dependent variable

	c=0.1	c=0.2	c=0.3	c=0.4	c=0.5	c=.6	c=.7	c=.8	c=.9	c=.99
adj-R2	0.170	0.169	0.188	0.199	0.205	0.205	0.204	0.197	0.190	0.187
(Intercept)	3.18	2.49	1.77	1.57	1.27	0.93	0.63	0.39	-0.051	-0.22
‘ask-scale-0.01‘	4.65	3.95	2.92	2.17	1.56	1.07	0.55	0.004	-0.68	-0.90
‘askAge-scale-0.001‘	-2.82	-0.52	0.63	1.17	1.77	2.90	3.37	3.35	4.55	4.31
‘askModified-scale-0.01‘	4.64	5.53	6.11	6.43	6.56	6.75	7.07	7.57	7.76	7.90
‘askVolume-scale-0.001‘	0.87	1.71	1.86	4.15	5.11	4.95	4.75	4.85	5.09	2.16
‘bid-scale-0.01‘	4.64	3.97	2.95	2.32	1.82	1.24	0.77	0.39	-0.37	-1.02
‘bidAge-scale-0.001‘	-3.68	-2.13	-0.64	0.064	0.82	1.32	1.56	2.68	3.53	3.25
‘bidModified-scale-0.01‘	4.43	5.45	6.00	6.82	7.18	7.41	7.70	8.62	8.74	7.96
‘bidVolume-scale-0.001‘	1.04	1.63	1.57	2.32	2.52	2.73	2.56	2.09	2.31	1.50
‘cancelbuy-scale-0.01‘	4.83	4.71	2.96	1.82	1.17	1.30	-0.40	-1.62	-2.49	-1.68
‘indexact-scale-0.01‘	14	14	12	11	9.05	7.75	6.80	5.80	4.06	0.57
‘lask-scale-0.01‘	-5.46	-4.76	-3.61	-3.00	-2.58	-2.01	-1.45	-0.72	0.027	-0.90
‘laskAge-scale-0.001‘	1.46	0.35	1.03	1.11	0.86	-0.53	-1.12	-0.079	0.95	1.19
‘laskModified-scale-0.01‘	-7.17	-7.70	-7.23	-7.86	-7.72	-7.88	-8.08	-8.07	-8.00	-9.44
‘laskVolume-scale-1e-04‘	5.01	3.49	15	6.38	-6.35	-15	-26	-33	-26	-9.41
‘lbid-scale-0.01‘	-5.37	-4.64	-3.54	-3.07	-2.73	-2.16	-1.63	-1.20	-0.43	-1.28
‘lbidAge-scale-0.001‘	2.38	1.54	0.96	1.32	1.12	1.44	1.34	1.83	2.94	1.15
‘lbidModified-scale-0.01‘	-6.84	-7.33	-7.05	-7.66	-7.78	-7.89	-8.09	-8.23	-8.08	-8.77
‘lbidVolume-scale-0.001‘	-0.003	-0.26	0.45	1.09	0.30	-0.26	-0.94	-1.56	-1.98	1.75
‘lspreads-scale-0.001‘	-0.18	-3.02	-2.25	1.07	2.29	3.83	3.44	2.58	1.80	-0.52
mobuy	-1.60	-1.59	-1.60	-1.61	-1.61	-1.59	-1.55	-1.45	-1.14	-0.15
mosell	-1.61	-1.60	-1.60	-1.59	-1.58	-1.56	-1.54	-1.44	-1.14	-0.13
‘prevexceed-scale-0.001‘	-14	-12	-9.15	-8.01	-7.29	-6.75	-6.59	-5.83	-4.77	-4.36
‘profLO-scale-1e-08‘	71	15	0.65	3.91	-2.81	-3.49	-3.38	-3.91	-2.96	-3.35
‘samemidspread-scale-0.1‘	-13	-8.99	-6.55	-4.92	-3.55	-2.43	-1.34	-0.14	0.88	-0.20
‘spreads-scale-0.001‘	32	10	4.64	3.59	1.99	1.02	0.54	0.48	0.53	0.64
‘ticksize-scale-0.1‘	6.06	7.10	7.82	7.55	8.35	9.17	9.94	8.78	9.11	4.60
‘timelast-scale-0.1‘	-8.31	-6.96	-5.51	-4.12	-2.90	-1.87	-1.09	-0.71	-0.32	0.006
log.duration.1	0.19	0.20	0.21	0.21	0.21	0.22	0.22	0.22	0.22	0.22
log.duration.2	0.089	0.090	0.09	0.09	0.10	0.10	0.10	0.10	0.10	0.10
log.duration.3	0.054	0.056	0.059	0.060	0.061	0.061	0.062	0.062	0.061	0.058
log.duration.4	0.038	0.039	0.043	0.044	0.044	0.045	0.045	0.046	0.045	0.042
log.duration.5	0.028	0.030	0.033	0.035	0.035	0.035	0.036	0.037	0.036	0.032
log.duration.6	0.024	0.026	0.029	0.031	0.031	0.033	0.034	0.035	0.034	0.026
log.duration.7	0.022	0.023	0.026	0.027	0.027	0.028	0.027	0.028	0.029	0.026
log.duration.8	0.020	0.022	0.025	0.026	0.026	0.027	0.027	0.028	0.026	0.020
log.duration.9	0.021	0.022	0.025	0.025	0.026	0.026	0.027	0.028	0.026	0.022
log.duration.10	0.023	0.024	0.028	0.028	0.028	0.029	0.029	0.029	0.029	0.025
‘num.lag.TED-scale-1e-05‘	-3.98	-2.70	-2.49	-2.94	-2.35	-3.14	-2.69	-4.12	-6.89	-18
‘lag. Announcement-scale-0.01‘	-12	-4.83	-14	-22	0.27	0.60	-5.12	-6.96	-6.61	-27
‘Announcement-scale-0.001‘	-63	-140	191	143	24	121	144	182	148	292
lag.Return	-0.032	0.38	0.49	1.04	0.75	0.74	0.41	0.25	0.46	0.49
‘lag.low.volume-scale-0.01‘	0.13	-2.44	-0.43	1.46	-1.12	0.23	3.16	3.63	2.37	1.05
‘lag.high.volume-scale-0.01‘	4.09	3.38	1.92	2.10	1.25	-0.51	0.39	-1.85	-1.92	-2.55
season.480.490	0.67	0.83	0.88	0.89	0.85	0.85	0.86	0.85	0.82	0.40
season.490.540	0.52	0.56	0.58	0.57	0.55	0.53	0.52	0.49	0.43	0.073
season.540.600	0.40	0.44	0.47	0.47	0.45	0.43	0.41	0.36	0.25	0.059
season.600.660	0.42	0.44	0.46	0.46	0.44	0.42	0.39	0.34	0.23	-0.005
season.660.720	0.46	0.46	0.46	0.43	0.40	0.38	0.34	0.29	0.20	-0.17
season.720.780	0.48	0.49	0.47	0.45	0.42	0.40	0.37	0.31	0.24	0.14
season.780.840	0.43	0.41	0.37	0.35	0.32	0.31	0.29	0.27	0.23	0.37
season.840.900	0.18	0.17	0.16	0.14	0.12	0.11	0.11	0.10	0.12	0.35

### 6.3 Importance of covariates

The empirical results on the covariates, their magnitude and signs provide a rich source of information for the analysis of liquidity resilience.

Informational factors, including anticipatory actions taken by LO suppliers, form the baseline explanations for illiquidity exceedance durations, to which a disparate set of non-informational explanations will be added. The TED process is persistent and clustering, suffering from bursts of informational clusters.

We provide our intuition for each finding. These explanations must by necessity be tentative, given that we do not observe the IDs of the traders, nor do we observe the contemporaneous LOBs on all other trading venues. While we are not always able to discriminate between competing explanations, we still find strong results using very little information, which by and large has the distinct advantage of allowing algorithmic traders to implement methods such as ours with publicly available information in near real time and better manage their trading or liquidity provision.

The discussion below applies to the spread as a liquidity measure, but we obtain largely similar results from the XLM measure. Below, we simply refer to each covariate by name. The variates that are not only statistically but also economically more significant have two stars: \*\*.

#### 6.3.1 Details on LOB explanatory variables

We now interpret the regression coefficients one by one.

**\*\* MO buy (-), MO sell (-).**

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
MOBuy	-1.6	-1.59	-1.6	-1.61	-1.61	-1.59	-1.55	-1.45	-1.14	-0.15
MOSell	-1.61	-1.6	-1.6	-1.59	-1.58	-1.56	-1.54	-1.44	-1.14	-0.13

The coefficients of the market-order (which included marketable limit-orders) dummies are uniformly around  $-1.6$  for both variables and across all  $c$  up to around  $.8$  beyond which the absolute magnitudes become smaller. For instance the value is  $-1.14$  for  $c_{.90}$  which translates to  $\frac{E[\tau|\cdot, \text{MObuy}=1]}{E[\tau|\cdot, \text{MObuy}=0]} = \exp(-1.14) = 0.32$ .

If the exceedance has been caused by a market(able) order blowing through the book, as opposed to being caused by pulled limit orders, then the expected exceedance duration is significantly smaller.



It is as if much of the time the market decided that a market order either reflects a retail trade, and thus does not carry private information that matters materially to the LO suppliers (so that new LOs replace the executed ones), or that it reflects an aggressive release of private information or uncertainty, and therefore the information innovation is fully digested in one movement. Relatedly there is the complementary possibility that the pulling of limit orders (and the absence of new LOs placed across the touch) occurs in anticipation of a private information event, explaining the longer expected duration. Indeed, a liquidity provider would not replace a limit order by an identical one immediately after pulling it, and other liquidity providers may infer an informational event and not step in immediately, or may in fact not immediately find out about the liquidity event related to another participant's pulling of LOs (whereas if a marketable order takes out LOs, the liquidity provider whose LOs were hit would find out very quickly).

**\*\* samemidspread (-,0)** suggests that uninformative exceedances (where the mid does not move) lead to shorter exceedance durations for most thresholds but the highest. For the highest exceedances, the average TED does not differ much if the mid moves or does not, the mid movements being dwarfed by the movements in the spread.

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
e-0.1	-13	-8.99	-6.55	-4.92	-3.55	-2.43	-1.34	-0.14	0.88	-0.2

For instance the value is -2.43 for  $c_{.60}$  which translates to  $\frac{E[\tau||\cdot, \text{samemidspread}=1]}{E[\tau||\cdot, \text{samemidspread}=0]}$   
 $= \exp(-2.43 * .1) = 0.78$ .

Beyond the informational rationale, this variate also captures a mechanical effect in that it further informs on the interpretation of the variates **MObuy** and **MOsell**, since a further possibility for the fast replenishment of the book following a marketable order has to do with the way Chi-X partly executes marketable limit orders. If the size of a marketable limit order exceeds the available size of sitting limit orders at that price, a portion of the order gets filled whereas the remaining bit is added as a limit order into the book just thereafter. If this happens, the data shows a widening of the spread (when the incoming marketable limit order gets filled by the entire level at that price point) followed by a reduction of the spread to the minimum tick size and a movement of the midprice (when the remaining order gets written in the book as a new limit order at the best bid or ask). The variable **samemidspread** is equal to zero if this is going on. **Samemidspread** therefore “corrects” the strong negative coefficients in **MObuy** and **MOsell** for this

mechanical splitting effect. Since the closing of the liquidity gap would be extremely quick in the case of a split marketable limit order and also comes with `samemidpread= 0`, it would not be a very prevalent effect in the data since outages resolve much faster if the book returns to the old midspread. Movements of the midspread following a liquidity outage are therefore more likely to be informational than mechanical.

**\*\* Ask (+,-) and Bid (+,-)** are both positive for small  $c$  and decrease with larger threshold levels  $c$ , turning mildly negative around  $c = .9$ .

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
ask e-0.01	4.65	3.95	2.92	2.17	1.56	1.07	0.55	0.004	-0.68	-0.9
bid e-0.01	4.64	3.97	2.95	2.32	1.82	1.24	0.77	0.39	-0.37	-1.02

The larger the number of different bids and asks (irrespective of volume) on the book at the first 5 levels at the moment of the exceedance, the longer the liquidity blow-out for  $c$  up to around .9 and the shorter the liquidity blowout for extreme exceedances.

For instance, for  $c_{.9}$ ,  $\frac{E[\tau|.,Ask=x+10]}{E[\tau|.,Ask=x]} = \exp(-0.68 * .1) = 0.934$ , so when ten more asks with different IDs are added to the first five levels of the LOB, the expected duration of a severe exceedance ( $c_{.9}$ ) is reduced by 6.6%. For exceedances above an already decent level, say  $c_{.6}$  we get  $\frac{E[\tau|.,Ask=x+10]}{E[\tau|.,Ask=x]} = \exp(1.07 * .1) = 1.1$ , so ten more asks means the TED is increased by 10%.

The general pattern for **Ask** and **Bid** has an interesting change of sign at a level just north of  $c_{0.8}$ , as if **Ask** and **Bid** were like a measure of mean-reversion of the level of measured liquidity towards some competition induced level of  $c_{0.8}$  that corresponds to a normal arrival of information. In a sense, an exceedance of a low  $c$ , say  $c_{0.4}$ , from below can be seen to occur without new information and *after* a period of exceptionally little news that pushed the spread so low in the recent past in the first place, while an exceedance of  $c_{.9}$  occurs because of a current or anticipated informational event.

If an exceedance is beyond, say  $c_{0.9}$ , having larger **Ask** and **Bid** in the LOB at that moment reflects more competition and faster healing of the LOB towards that competitive level as that information gets digested by more liquidity suppliers ready to compete and close the spread. If the exceedance is beyond, say  $c_{0.4}$  coming from a level below that, larger **Ask** and **Bid** makes a return below that low  $c_{0.4}$  slower because the increased number of bids or asks at the first 5 levels were what caused the return to an informationally more normal market, with spreads returning towards the mean competitive level of  $c_{0.8}$  reflecting normal levels of news again, rather than the lower and less competitive level of news corresponding to a spread of  $c_{0.4}$ .

**profLO (0-)** being insignificant for  $c$  below  $c_{.5}$  and uniform and negative for  $c_{.5}$  and above illustrates that liquidity suppliers tend to replenish the book after an exceedance faster if they have profited from filling market orders in the last few seconds, and similarly they replenish more slowly if they suffered losses to the aggressive traders.

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
e-08	71	15	0.65	3.91	-2.81	-3.49	-3.38	-3.91	-2.96	-3.35

This magnitude of the effect is very small, though. For instance, for  $c_{.9}$ ,  $\frac{E[\tau|.,\text{profLO}=x-1M]}{E[\tau|.,\text{profLO}=x]} = \exp(2.96 * .01) = 1.03$ , so if in the last 10 seconds market makers lost 1M euros more, the duration of exceedances is expected to increase by a mere 3%. This may be due to the fact that market makers do not react to past losses, or, more likely, that market makers measure profits differently and/or that those different measures are already captured by a variety of regression variates.

**AskVolume (+)** and **Bidvolume (+)** are both positive.

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
askV 0.001	0.87	1.71	1.86	4.15	5.11	4.95	4.75	4.85	5.09	2.16
bidV0.001	1.04	1.63	1.57	2.32	2.52	2.73	2.567	2.09	2.31	1.5

If the LOB has lots of volume sitting at the first five levels and yet we have an exceedance, this would mean the exceedance is serious and will therefore need more time to resolve itself.

For instance, for  $c_{.9}$ ,  $\frac{E[\tau|.,\text{profLO}=x+10000]}{E[\tau|.,\text{profLO}=x]} = \exp(5.09 * .001 * 10) = 1.03$ , so when total **AskVolume** in the first five levels increases by  $10 \times 1000 = 10000$  shares, expected TED rises by 3%.

**AskAge (-,0,+)** and **BidAge (-,0,+)** are positive from about  $c_{.4}$  and  $c_{.6}$  onwards.

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
askAge 0.001	-2.82	-0.52	0.63	1.17	1.77	2.9	3.37	3.35	4.55	4.31
bidAge 0.001	-3.68	-2.13	-0.64	0.064	0.82	1.32	1.56	2.68	3.53	3.25

For the most relevant range  $c \geq c_{.6}$ , the older the limit orders on the book, the longer the duration, perhaps because the quotes have become stale and once hit will not be replaced quickly. Perhaps the age reflects an otherwise quiet market, so if a sudden burst occurs, it either catches limit order suppliers unawares, or they don't view a slow market (otherwise the limit orders would have been hot or replaced earlier) as profitable. Notice, however, that this is beyond and above the degree of activity or quietness which we control for through **indexact**, the activity (incl. LO submissions, amendments, cancellations, executions) in the entire CAC40 market in the previous second

(see below for more analysis of `indexact`). In terms of magnitude, for  $c_9$ ,  $\frac{E[\tau|\cdot, \text{AskAge}=x+10m]}{E[\tau|\cdot, \text{AskAge}=x]} = \exp(4.5 * .001 * 10) = 1.046$ , so that an average age at the first five levels that is older by 10m leads to an expected duration that is longer by 4.6%. It follows that the TED is quite sensitive to the age of the LOB.

**Spreads (+,+,0,+)** are positive and significant for both small and very large  $c$ , and they are very close to zero as well as not significantly different from zero for  $c = 0.7, 0.8$ .

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
spr. 0.001	32	10	4.64	3.59	1.99	1.02	0.54	0.48	0.53	0.64

A positive value reflects the intuition that the wider the spread just after an event, the longer we would expect the spread exceedance to last, on average, possibly because the larger spread reflects a larger piece of news.

Overall spreads do not play as much a role as one might have guessed, perhaps because for the bid-ask spread as a metric, all it takes is one newly placed quote close to the touch to end the illiquidity event (with the spread as measure), there is no need to add quotes first far from the touch and then move towards the touch. In that sense if a liquidity event occurs and the limit order suppliers decides that it either was not an informational event or that the information is now out, then competition will make sure that the new limit orders will be placed tightly immediately without any tâtonnement.

Also, other variables already capture informativeness.

**\*\* askModified (+)** and **bidModified (+)** are very close one to the other and they are positive, increasing with  $c$ .

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
askM 0.01	4.64	5.53	6.11	6.43	6.56	6.75	7.07	7.57	7.76	7.9
bidM 0.01	4.43	5.45	6	6.82	7.18	7.41	7.7	8.62	8.74	7.96

For instance, with  $c_9$ ,  $\frac{E[\tau|\cdot, \text{askModified}=x+1]}{E[\tau|\cdot, \text{askModified}=x]} = \exp(7.76 * .01 * 1) = 1.08$ : one more modified ask order present in the first 5 layers leads to an 8% increase in expected TED, which is a very strong effect.

The increased frequency of such changes can be a reflection of uncertainty surrounding an anticipated arrival of information. In that sense, these orders can be interpreted as fleeting liquidity not representing a firm willingness to provide liquidity. If the event then does occur, it is more likely to be informational the more modifications are present in the LOB - or interpreted as such - and the time it takes to narrow the spread again is longer.

**\*\* Ticksize (+)** is pretty uniform with a slight inverse U-shape, indicating that if a stock suffers an exceedance while in a larger tick size regime,

it will take longer to return to that level, except for very large exceedances.

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
ts 0.1	6.06	7.1	7.82	7.55	8.35	9.17	9.94	8.78	9.11	4.6

For instance, with  $c_{.9}$ ,  $\frac{E[\tau|., \text{Ticksize}=x+.1]}{E[\tau|., \text{Ticksize}=x]} = \exp(9.11 * .1 * .1) = 1.095$ : a tick size increase of .1cent leads to a 10% increase in expected TED, which is a very strong effect.

This is commonsensical as larger tick sizes reduce the updating frequency of the LOB since the larger tick sizes effectively correspond to larger costs of jumping the queue.

There are also possibly mechanical reasons, since a larger tick size may make returning to the old spread impossible. It can also possibly be viewed as a cross-sectional finding saying that higher priced equities are less liquid (since the number of potential investors is lower), though this argument needs to be relativised because our regressions have firm fixed effects.

This is in spirit contrary to the theoretical results of Foucault et al. (2005). However, our measure and their measure of resilience are not quite identical: they measure market resiliency by the probability that, after a liquidity shock, the spread reverts to its former level before the next transaction. Also, our analysis is not exactly a ceteris paribus result since a stock goes through different liquidity regimes when its stock price changes, which is itself an informative event. Finally, the quantile thresholds  $c$  in our analysis are adjusted daily to that day's liquidity behaviour (though we have day dummies and firm fixed effects). The two sets of results may in theory be compatible since our definition does not consider transactions as the defining characteristic but the spread itself, so one might imagine market order decisions that come in more slowly after an exceedance when tick sizes are large.

### 6.3.2 Dynamic serial structure

Of the covariates, several capture directly the serial dynamic component of the data generating process and they are all highly significant (capturing effects over and above the fact that some days or intervals within a day can exhibit more or less resilience for other reasons, as captured by day-dummies and intra-day bucket dummies). Akin to volatilities that exhibit GARCH clustering effects (Engle (1982), Bollerslev (1986)) or durations that also exhibit clustering (e.g. Hawkes self-exciting processes used in Engle and Russell (1998) where they find strong duration between trade clustering, and also high serial dependence), TEDs exhibit clustering (large TEDs are likely

to be followed by large TEDs). The signs in parentheses indicate the sign of the regression coefficient.

**Log.duration (+)** are the lagged log-durations, captures the autoregressive structure of the model, and always comes with a positive coefficient. TEDs are persistent. Coefficients decrease with the lag (we use 10 lags) from .2 to .03, so that the effect of an exceedance duration on subsequent durations is positive and temporary, getting slightly washed out over time.

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
-1	0.19	0.2	0.21	0.21	0.21	0.22	0.22	0.22	0.22	0.22
-2	0.089	0.09	0.09	0.09	0.1	0.1	0.1	0.1	0.1	0.1
-3	0.054	0.056	0.059	0.06	0.061	0.061	0.062	0.062	0.061	0.058
-4	0.038	0.0397	0.043	0.044	0.044	0.045	0.045	0.046	0.045	0.042
-5	0.028	0.03	0.033	0.035	0.035	0.035	0.036	0.037	0.036	0.032
-6	0.024	0.026	0.029	0.031	0.031	0.033	0.034	0.035	0.034	0.026
-7	0.022	0.023	0.026	0.027	0.027	0.028	0.027	0.028	0.029	0.026
-8	0.02	0.022	0.025	0.026	0.026	0.027	0.027	0.028	0.026	0.02
-9	0.021	0.022	0.025	0.025	0.026	0.026	0.027	0.028	0.026	0.022
-10	0.023	0.024	0.028	0.028	0.028	0.029	0.029	0.029	0.029	0.025

The positive sign indicates that the expected TED over a particular threshold will be larger, when the duration of similar exceedances in the near past has been longer. For a given lag, the magnitude of the coefficient is a nearly flat function of the threshold  $c$ .

**\*\* Prevexceed (-)**. At the time of an exceedance, a higher *number* of exceedances in the past second is associated with a *shorter* exceedance duration under the model. This effect is weaker the larger  $c$ .

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
0.001	-14	-12	-9.15	-8.01	-7.29	-6.75	-6.59	-5.83	-4.77	-4.36

E.g. for  $c_{.9}$ ,  $\frac{E[\tau|.,\text{Prevexceed}=x+10]}{E[\tau|.,\text{Prevexceed}=x]} = \exp(-4.77 * 0.001 * 10) = 0.95$ : 10 more exceedances in the last second reduces the expected TED by 5%.

There may be a purely mechanical effect at work here in that if there have been many exceedances above a given  $c$  in the last second, they must have been of short duration.

It can also capture the fact that private information is now out and priced in, so the liquidity related aftershocks are quickly absorbed.

Since the log durations of the previous 10 exceedances are regressors in their own right (confirming the interpretation that short exceedances are more likely following short exceedances), the **prevexceed** variate captures the complementary flavour of exceedances that happen in fast-moving mar-

kets. One further possible explanation for this phenomenon may be that limit order suppliers scan the markets and allocate capital and CPU power to those securities that have suffered exceedances in the near past (and that have therefore exhibit larger bid-ask spreads and possibly juicier profit opportunities to the limit order suppliers).<sup>7</sup>

**\*\* Timelast (-)**. A larger `timelast`, i.e. the longer it has been since the last exceedance, the smaller the expected current exceedance duration, with the effect shrinking with larger  $c$ .

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
0.1	-8.31	-6.96	-5.51	-4.12	-2.9	-1.87	-1.09	-0.717	-0.32	0.006

E.g.  $c_{.9}$ ,  $\frac{E[\tau|., \text{Timelast}=x+1m]}{E[\tau|., \text{Timelast}=x]} = \exp(-0.32*0.1) = 0.968$ : at this exceedance moment, if no exceedance had occurred not for  $xm$  but for  $x + 1m$ , the expected TED would be reduced by 3.2%.

This complements lagged TEDs in addressing the dynamic serial dependence of exceedances and also has a clustering flavour. Since longer past exceedances would lead to longer future exceedances by serial dependence, absence of past exceedances for a longer while may reflect so resilient a replenishment that no exceedances occurred in the first place.

### 6.3.3 Algorithm choice and market factors

HFTs are large players in the Chi-X market we are studying here and we would expect them to have a strong impact on liquidity resilience. As a part of our research, and since we do not have HFT flags, we consulted a number of HFTs so we could determine the factors which affect their approach to trading. What became clear is that both the number and types of trading algorithms depend on market conditions. For example, one trader told us he selected his algorithms at the start of his day based on market conditions at the time. If this is indeed the case, it would provide an explanation for the changing significance and magnitude of coefficients in the model over time.

While we are inclined to agree with this explanation, it does not directly

---

<sup>7</sup>This interpretation can be seen as a corroboration of the limited information processing capacity of algorithms, CPUs and bandwidth, and the fact that fast traders need to make a choice as to which of the different competing tasks should be given priority. This interpretation is very much in the spirit of \*\*\* Simon (1957) that although boundedly rational agents experience limits in formulating and solving complex problems and in processing (receiving, storing, retrieving, transmitting) information, they otherwise remain ‘intendedly rational’ (also \*\*\* Williamson, 1981) and allocate priority to stocks that have suffered exceedances in the near past.

lead to a testable hypothesis, due to the absence of information regarding the components of traders' strategies. However, we can test the hypothesis indirectly by relating the magnitude and choice of coefficients to prevailing market factors.

**Num.Lag.TED (-).** The larger the *number of exceedances during the previous day*, the lower the expected duration today.

1e-05	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
	-3.98	-2.7	-2.49	-2.94	-2.35	-3.14	-2.69	-4.12	-6.89	-18

For  $c_{.9}$ ,  $\frac{E[\tau|.,Num.Lag.TED=x+1000]}{E[\tau|.,Num.Lag.TED=x]} = \exp(-6.89*0.00001*1000) = 0.933$ : had there been 1000 more exceedances beyond 90th percentile yesterday (computed using yesterday's histogram), the expected TED beyond today's 90th percentile would be reduced by 6.7%.

This effect is stronger for larger  $c$ . This may capture clustering of news in the sense that if yesterday was a busy informational day that generated lots of exceedances, it is more likely that innovations will have settled and been digested overnight, bringing back a stronger willingness to make tight spreads the next day. The limited attention assumption may also have some power here in the sense that limit order suppliers may review the previous day's trading and reallocated priorities to those stocks that have exhibited more action.

**\*\* indexact (+)** is positive and diminishing with  $c$ , with nearly zero for  $c = .99$ .

0.01	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
	14	14	12	11	9.05	7.75	6.8	5.8	4.06	0.57

For example, with  $c_{.9}$ ,  $\frac{E[\tau|.,Indexact=x+1000]}{E[\tau|.,Indexact=x]} = \exp(4.06 * 0.01) = 1.041$ : one thousand more limit orders submissions, cancellations, amendments, executions in the last second lead to an increase of 4% in the expected TED.

When the aggregate stock market is in a more active phase, individual LOBs get replenished more slowly, possibly reflecting more risk to market makers for a variety of reasons, including informational ones. Given that placed LOs are akin to free options that the liquidity suppliers make available to the aggressive traders, the BA spread is akin to a delayed-payment option premium paid to the liquidity suppliers. In faster markets, option values naturally increase, so that a return to a lower spread takes longer.

**lag.Return (+)** being positive indicates that the better the previous daily return is the larger the exceedance duration today.

0.01	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
	14	14	12	11	9.05	7.75	6.8	5.8	4.06	0.57



For instance, for  $c_9$ ,  $\frac{E[\tau|\cdot, \text{lag.Return}=x+.03]}{E[\tau|\cdot, \text{lag.Return}=x]} = \exp(-0.46 * .03) = 1.01$ . If the previous day's return is 3% higher, the expected duration is 1% longer.

Perhaps the market makers believe that good news get reflected in prices quickly and adverse selection disappears while bad news trickle more slowly.

**lag.high.volume (+,-)** suggests that low exceedances resolve faster when the previous day was a high volume day while larger exceedances take longer to resolve in that case (even though the thresholds  $c$  are computed daily).

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
0.01	4.09	3.38	1.92	2.1	1.25	-0.51	0.39	-1.85	-1.92	-2.55

E.g. with  $c_9$ ,  $\frac{E[\tau|\cdot, \text{Indexact}=1]}{E[\tau|\cdot, \text{Indexact}=0]} = \exp(-1.92 * 0.01) = 0.98$ , suggesting that if the previous day was a high volume day ( $\geq 85\%$ ), today's expected TED is reduced by 2%.

**Announcement (-,+)** is positive from  $c_3$  onwards and would correspond to the common-sensical intuition that days with US macro or monetary announcements tend to extend the exceedance durations (beyond already larger thresholds since thresholds are re-estimated every day).

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
0.001	-63	-140	191	143	24	121	144	182	148	292

For instance, for  $c_9$ ,  $\frac{E[\tau|\cdot, \text{Announcement}=1]}{E[\tau|\cdot, \text{Announcement}=0]} = \exp(148 * .001) = 1.16$ , saying that announcement days lead to expected TEDs that are 16% longer than on non-announcement days.

**lag.Announcement (-,0,-)** says that except for the median exceedances where no effect can be felt, a day after US macro announcements is a day where exceedance durations are shorter (beyond already lower thresholds), perhaps because most of the uncertainty related to the announcements has been removed.

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
0.01	-12	-4.83	-14	-22	0.27	0.6	-5.12	-6.96	-6.61	-27

For instance, for  $c_9$ ,  $\frac{E[\tau|\cdot, \text{lag.Announcement}=1]}{E[\tau|\cdot, \text{lag.Announcement}=0]} = \exp(-6.61 * .01) = 0.94$ , saying that if yesterday was an announcement day, then the expected TEDs today are 6% shorter (again beyond today's thresholds) than after a non-announcement day.

**num.lag.TED (-)** shows that a larger number of TEDs beyond a given threshold yesterday reduces exceedance durations beyond that same threshold today.

	c0.1	c0.2	c0.3	c0.4	c0.5	c0.6	c0.7	c0.8	c0.9	c0.99
e-05	-3.98	-2.7	-2.49	-2.94	-2.35	-3.14	-2.69	-4.12	-6.89	-18

For instance, for  $c_9$ ,  $\frac{E[\tau|\cdot, \text{num.lag.TED}=x+1000]}{E[\tau|\cdot, \text{num.lag.TED}=x]} = \exp(-6.89 * 1e - 5 * 1000) = 0.93$ , saying that if the previous day had 1000 more TEDs beyond yesterday's 90% quantile (irrespective of duration) then the expected TED today is 7% shorter. However, if only 100 more exceedances yesterday are recorded, the shortening is only about .007, or 0.7%.

## 7 Liquidity resilience profiles, LRPs

### 7.1 Definition

The level of the exceedance threshold,  $c_j$ , depends on the underlying applications, of which we have described a subset. In order to understand the relationship between the TED and a range of thresholds we now define the notion of the liquidity resilience profile:

**Definition 2.** *The liquidity resilience profile, LRP, is a curve of the expected log of the TED observations as a function of the liquidity threshold,  $c$ . More precisely, for any threshold  $c$  it tells us the expected log duration of an exceedance of  $c$ , given that the limit order book is in the state that typically obtains if that threshold  $c$  is exceeded.*

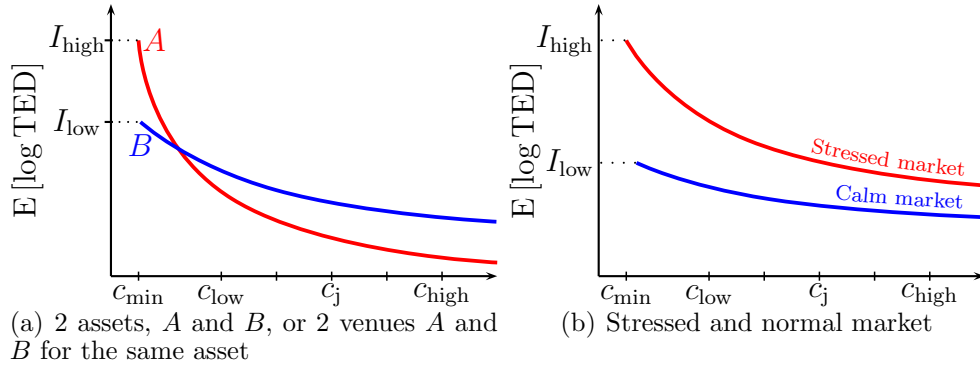
We would expect the LRP curves to be different for different assets and LOB regimes. One would generally expect a monotonically decreasing function of the observed TEDs for increasing thresholds  $c_j$  because once we have an exceedance over a given threshold, we must have had an exceedance over a lower threshold. However, this does not necessarily imply that the same will hold for the expected TEDs also, as they rely on the estimated model at each threshold.

The LRPs also enable a brokerage firm to compare estimates of the total time for execution of a block trade across thresholds and across assets, when execution is dependent on a threshold level of liquidity. Alternatively, it enables firms to identify suitable threshold levels of liquidity, such that the block will be filled in a given time horizon.

We illustrate the LRPs in Figure 4. In order to compare resiliences across assets and regimes, when plotting the LRPs we view them not as mappings from  $c$  to expected log duration, but from  $j \in [0, 1]$  to expected log duration through the composite mapping  $j \mapsto c_j \mapsto$  expected log duration. Panel (a) shows the LRP for two different assets,  $A$  and  $B$ , or for two different

**Figure 4.** Liquidity resilience profiles

The figure shows a hypothetical relationship between the exceedance threshold level and the TED, for levels,  $I$  and two slopes.



trading venues  $A$  and  $B$  for the same asset. For very low thresholds, the liquidity shocks for  $A$  last much longer than for  $B$ , whilst the opposite holds true for very large liquidity shocks. Therefore, the intercept,  $I$  for  $A$  is much higher, but  $A$  also has a much more negative slope. This means that for asset  $A$ , liquidity is generally replenished quickly after an adverse liquidity shock, bringing the LOB back to normal liquidity levels, but very high liquidity levels take much longer to achieve.

Panel (b) shows the LRP for the same asset but in different market conditions, the first in a calm market and the second in a stressed market. In stressed market conditions, the LOB is less resilient, particularly at the lowest thresholds. There is thus less liquidity replenishment activity at all thresholds, and this can occur through a reluctance of market makers to replenish activity in uncertain stressed conditions.

### 7.1.1 Applications

The LRPs are informative about the level of liquidity replenishment in the LOB and can identify assets for which we would expect a swift return to a high liquidity level after a shock. To the extent that liquidity replenishment in the millisecond environment is predominantly the domain of high-frequency LO suppliers, LRPs can also indicate the presence or absence of such traders in particular assets. This is in line with the theoretical predictions of Foucault et al. (2005), who suggested that large spreads would be

more common in markets dominated by impatient traders (those submitting aggressive market orders, rather than passive limit orders).

One use of the LRP is in aiding the comparison of liquidity resilience across assets. In particular, given an assumption about the LOB regime, LRPs for different assets enable an ordering of the assets by their liquidity resilience. This is helpful because it indicates the ex-ante relative risk of trading each asset in that LOB regime. The risk here refers to the duration of the disturbance caused by trading.

We can also compare profiles for the same asset, but for different LOB regimes. This enables a regulator to perform scenario analysis, in order to understand the effects of stressed market conditions on market composition, for example. Poor liquidity replenishment (and thus high values for the expected TED) would be associated with the absence of high-frequency market makers.

In the case where one is interested in the duration of exceedances over particular levels of a liquidity measure, the LRP may be presented as a function of the absolute liquidity threshold level, rather than the relative levels used here.

## 7.2 LRP Results

By estimating the model across different thresholds, we can construct the liquidity resilience profiles (LRPs).

Figure 5 shows the LRP for each asset, averaged across all the days. It is as expected declining with extreme thresholds,


Figure 6 shows the average LRP for Crédit Agricole for the first 10 days and last 10 days of the sample, as well as across all the days. In the first 10 days, the volatility of stock returns was 1.4%, whilst in the last 10 days it was 3.1%, indicating that the Figure shows both stressed and non-stressed time periods.

### 7.2.1 Liquidity resilience asset ordering

Figures 7 and 8 show example LRPs of certain CAC40 stocks, broken down by market capitalisation and by share price. Comparing the intercept terms gives us an ordering of the assets by the relative likelihood of the asset returning to an extremely high (and exceptional) level of liquidity in a short period of time. The same ordering can be performed for any threshold level of

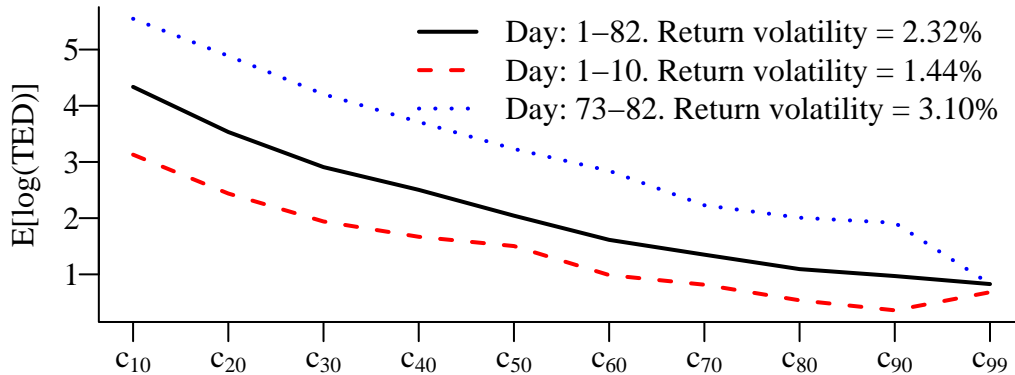
**Figure 5.** LRPs for all assets and the average

The figure shows the average LRP for each asset and the average across all assets. Some LRPs end in  $c_{80}$  or  $c_{90}$  because the stock has a low trading volume and there are not enough exceedances at extreme thresholds.



**Figure 6.** LRP for Crédit Agricole and different market conditions


The figure shows the average LRP across all days and for the first and last 10 days, as well as the stock return volatility for those days.




liquidity, and we facilitate comparison of the profiles of assets, by presenting thresholds as deciles of the liquidity measure, rather than absolute values.

We can compare the shapes of the LRPs visually or define orderings.

**Definition 3** (First and Second Order Resilience Dominance). *First Order Resilience Dominance (FRD) between  $LRP_A$  and  $LRP_B$  (say, between asset A and asset B across all periods and regimes, or between the same asset in a*




(a) Some stocks with low market cap




(b) Some stocks with high market cap

**Figure 7.** Liquidity Resilience Profile for some of the stocks with low (left) and high(right) market capitalisations in the CAC40



(a) Stocks with low share prices



(b) Stocks with high share prices

**Figure 8.** Liquidity Resilience Profile for some of the stocks with low (left) and high(right) share prices in the CAC40

*low vol environment A and in a high vol environment B, or between a given asset on one trading venue A vs on another one B) is be defined as*

$$AFRDB \Leftrightarrow LRP_A(c_j) \leq LRP_B(c_j) \quad \forall j \geq \underline{j} > 0$$

*and the weaker Second Order Resilience Dominance (SRD) – capturing not absolute superiority but better resilience after large exceedances – is defined*

as

$$\text{ASRDB} \Leftrightarrow \int_q^{+\infty} \text{LRP}_A(c_j) dj \leq \int_q^{+\infty} \text{LRP}_B(c_j) dj \quad \forall j \geq \underline{j}$$

for some small  $\underline{j} > 0$  chosen to make sure LRPs and integrals are defined.

Next, we can cluster assets according to whether they show similar behaviour in terms of TED durations across different thresholds in certain market conditions. Particularly in stressed market conditions (for example, when volumes are low and spreads are high), a market maker may focus her efforts on assets whose LRPs do not differ greatly from those in normal conditions, as those assets will be more likely to respond quickly to further liquidity shocks.


### 7.2.2 Liquidity resilience in different regimes

In this section we amend the definition of LRP slightly by choosing a different conditioning set. Whereas in the general definition,  $\text{LRP}(c)$  is computed given the typical variates that obtain if indeed there is an exceedance of  $c$ , we can further condition on additional variate constellations.

For instance, in Figure 9, we observe directly how a change in the LOB regime (through the LOB covariates) has an impact on the Liquidity Resilience Profile by further conditioning on the LOB being the one of a normal, a typical extremely quiet and a typical extremely volatile day, respectively. In particular, we see that the curve is shifted upwards in the case of the stressed scenario, corresponding to a higher expected log TED. This indicates that, at least in the case of Crédit Agricole, the relationship between the durations of exceedances over different thresholds does not change when we move to a stressed scenario.

However, this relationship between the liquidity profiles at different regimes does not have to hold in general, and we provide as an example the liquidity profiles for BNP Paribas in Figure 9. In this case, one has to consider carefully the liquidity resilience behaviour of the asset, particularly when operating in a regime characterised by poor liquidity replenishment.

The ranking of assets by their LRP can also be performed for different assumptions regarding the LOB regime. In this case, one can observe whether the rankings of the assets are consistent between regimes. One would prefer to operate in assets that have a high level of liquidity resilience for the prevailing regime. This would have implications for brokerages, for example, who have a large volume of orders to fill for different assets, and have some flexibility



**Figure 9.** Liquidity resilience profile for Crédit Agricole (left) and BNP Paribas (right) in both the normal LOB regime (in which covariates take the values of a median day) and extreme scenarios (in which covariates take the quantile values (10% and 90% for the ‘Extreme 1’ and ‘Extreme 2’ cases, respectively) of such an extreme day).


about the timing of their orders. They can then select particular assets to focus on, depending on the ranking of the assets in that regime.

As far as market quality is concerned, regulators would be interested in comparing LOB resilience across regulatory regimes (say different tick sizes, order-to-trade ratios etc.), and a measure such as Second order Resilience Dominance would allow them to choose amongst alternative regimes in a coherent fashion.

### 7.2.3 Results for the XLM liquidity measure

For comparison with previous results, we also provide in Figure 10 the LRPs obtained for a normal LOB regime for the subset of assets of the CAC40 that we consider here. For most assets, the shape of the LRP in the XLM case seems broadly similar to its counterpart in the spread case, given the same assumption about the LOB regime.





**Figure 10.** Liquidity Resilience Profile for a subset of assets in the CAC40, using the XLM as the liquidity measure and assuming a normal LOB regime

## 8 Conclusion

We shed light in this paper on the narrow but important question of what drives resilience in modern computer-based trading environments. A market is resilient if an evaporation of liquidity quickly resolves again, that is if the exceedance duration beyond an illiquidity threshold is short. Our resilience methods work with any liquidity measure, and for concreteness in this paper we focus on the frequently-used measures that are the bid-ask spread and the round-trip cost, the results being very similar for both.

The reason that resilience of liquidity crucially matters in today’s markets is because static liquidity notions on their own are not informative any longer given the fast flickering of the order book and the relatively minor capital and inventories held by market makers. The speed of order book replenishment has replaced capital to a large extent, and the ease of trading is captured by the dynamic notion of liquidity resilience.

We use survival analysis which allows us to quickly and conveniently characterise the effects of the shape of the current and past limit-order book and of the overall market conditions on the duration of an illiquidity event. The estimated parameters provide directly interpretable guidance as to how long the exceedance is expected to be, allowing a market participant to readjust its trades accordingly so as to minimise trading costs, market impact and

slippage.

Applying our methods to millisecond data from Chi-X, we show that the effect of the shape of the LOB on exceedance durations is consistent with informational considerations. At the moment of the exceedance, exceedances are predicted to be longer if the LOB, its recent history as well as overall market variates, are in constellations that are consistent with those one would expect if informational events were expected by the market. Market participants can adapt our methods to better navigate liquidity and minimize trading costs, market impact and the timing of trades.

These findings can then be used to construct what we call *Liquidity Resilience Profiles* (LRP) that condense the resilience behaviour of an asset on a given venue (and can be made conditional on a given overall market environment) for any given exceedance size. These profiles, generally decreasing in the size of the exceedance, can be used to order the resilience between different assets on the same venue, or between different venues for the same asset, or between different macro environments for the same stock on the same venue. Such information allows traders to better choose venues to minimize trading costs, or it allows market supervisors to gauge the quality of markets in various stressed conditions.

## References

- AL-SUHAIBANI, M. AND L. KRYZANOWSKI, “An exploratory analysis of the order book, and order flow and execution on the Saudi stock market,” *Journal of Banking & Finance* 24 (2000), 1323–1357. 2.2
- ALFONSI, A., A. FRUTH AND A. SCHIED, “Constrained portfolio liquidation in a limit order book model,” *Banach Center Publ* 83 (2008), 9–25. 3
- , “Optimal execution strategies in limit order books with general shape functions,” *Quantitative Finance* 10 (2010), 143–157. 2.1, 3
- AMIHUD, Y., H. MENDELSON AND L. H. PEDERSEN, *Market Liquidity: Asset Pricing, Risk, and Crises* (CUP, 2013). 3
- BERTSIMAS, D. AND A. W. LO, “Optimal control of execution costs,” *Journal of Financial Markets* 1 (1998), 1–50. 2.1, 3
- BESSEMBINDER, H., “Issues in assessing trade execution costs,” *Journal of Financial Markets* 6 (2003), 233–257. 4
- BIAIS, B., P. HILLION AND C. SPATT, “An empirical analysis of the limit order book and the order flow in the Paris Bourse,” *Journal of Finance* 50 (1995), 1655–1689. 5.1
- BOLLERSLEV, T., “Generalised autoregressive conditional heteroskedasticity,” *Journal of Econometrics* 51 (1986), 307–327. 6.3.2
- BRADBURN, M., T. CLARK, S. LOVE AND D. ALTMAN, “Survival analysis part II: multivariate data analysis—an introduction to concepts and methods,” *British journal of cancer* 89 (2003), 431. 2.2
- BRUGLER, J. AND O. LINTON, “Circuit Breakers on the London Stock Exchange: Do they improve subsequent market quality?,” Working Paper, 2014. 5
- BRUNNERMEIER, M. K., “Deciphering the liquidity and credit crunch 2007–08,” Technical Report, National Bureau of Economic Research, 2008. 2.1
- CHO, J.-W. AND E. NELLING, “The probability of limit-order execution,” *Financial Analysts Journal* (2000), 28–33. 2.2

- COPPEJANS, M., I. DOMOWITZ AND A. MADHAVAN, “Resiliency in an automated auction,” Technical Report, Working Paper, Duke University, 2004. 2.1
- DEGRYSE, H., F. DE JONG, M. VAN RAVENSWAAIJ AND G. WUYTS, “Aggressive orders and the resiliency of a limit order market,” *Review of Finance* 9 (2005), 201–242. 2.1
- DEMSETZ, H., “The Cost of Transacting,” *The Quarterly Journal of Economics* 82 (1968), 33–53. 2.1
- DONG, J., A. KEMPF AND P. YADAV, “Resiliency, the neglected dimension of market liquidity: Empirical evidence from the New York Stock Exchange,” *Available at SSRN 967262* (2007). 2.1
- ENGLE, R. F., “Autoregressive conditional heteroskedasticity with estimates of the variance of the United Kingdom Inflation,” *Econometrica* 50 (1982), 987–1007. 6.3.2
- ENGLE, R. F. AND J. R. RUSSELL, “Autoregressive conditional duration: a new model for irregularly spaced transaction data,” *Econometrica: Journal of the Econometric Society* (1998), 1127–1162. 6.3.2
- FOUCAULT, T., O. KADAN AND E. KANDEL, “Limit order book as a market for liquidity,” *Review of Financial Studies* 18 (2005), 1171–1217. 2.1, 6.3.1, 7.1.1
- FOUCAULT, T., M. PAGANO AND A. RÖELL, *Market Liquidity: Theory, Evidence, and Policy* (Oxford University Press, 2013). 3
- GARBADE, K. D. AND K. GARBADE, *Securities markets* (McGraw-Hill New York, 1982). 2.1
- GLOSTEN, L. R. AND P. R. MILGROM, “Bid, ask and transaction prices in a specialist market with heterogeneously informed traders,” *Journal of financial economics* 14 (1985), 71–100. 2.1
- GOMBER, P., U. SCHWEICKERT AND E. THEISSEN, “Liquidity dynamics in an electronic open limit order book: An event study approach,” Technical Report, CFR working paper, 2011. 2.1

- HARRIS, L., *Trading and exchanges: Market microstructure for practitioners* (Oxford University Press, USA, 2002). 2.1
- KEMPF, A., D. MAYSTON AND P. YADAV, “Resiliency in limit order book markets: A dynamic view of liquidity,” in *EFA 2007 Ljubljana Meetings Paper* (2007). 2.1
- KYLE, A. S., “Continuous auctions and insider trading,” *Econometrica: Journal of the Econometric Society* (1985), 1315–1335. 2.1
- LAMBERT, P., D. COLLETT, A. KIMBER AND R. JOHNSON, “Parametric accelerated failure time models with random effects and an application to kidney transplant survival,” *Statistics in medicine* 23 (2004), 3177–3192. 2.2
- LARGE, J., “Measuring the resiliency of an electronic limit order book,” *Journal of Financial Markets* 10 (2007), 1–25. 2.1, 3
- LAWLESS, J. F., “Inference in the generalized gamma and log gamma distributions,” *Technometrics* 22 (1980), 409–419. 3.3
- LO, A. W., A. C. MACKINLAY AND J. ZHANG, “Econometric models of limit-order executions,” *Journal of Financial Economics* 65 (2002), 31–71. 1, 2.2, 3.3
- OBIZHAEVA, A. AND J. WANG, “Optimal trading strategy and supply/demand dynamics,” *Journal of Financial Markets* (2012). 2.1, 3
- ROŞU, I., “A dynamic model of the limit order book,” *Review of Financial Studies* 22 (2009), 4601–4641. 2.1
- TOKE, I. M., ““Market Making” in an Order Book Model and Its Impact on the Spread,” in *Econophysics of Order-Driven Markets* (Springer, 2011), 49–64. 3



THE LONDON SCHOOL  
OF ECONOMICS AND  
POLITICAL SCIENCE ■



The London School of Economics  
and Political Science  
Houghton Street  
London WC2A 2AE  
United Kingdom

Tel: +44 (0)20 7405 7686  
[systemicrisk.ac.uk](http://systemicrisk.ac.uk)  
[src@lse.ac.uk](mailto:src@lse.ac.uk)