

Measuring the resiliency of an electronic limit order book[☆]

Jeremy Large

All Souls College, University of Oxford, Oxford, OX1 4AL, UK

Available online 7 November 2006

Abstract

An electronic limit order book is *resilient* when it reverts to its normal shape promptly after large trades. This paper suggests a continuous-time impulse response function based on intensities, which formalizes resiliency in terms of a time-frame and probability of order book replenishment. This is then estimated for trading on an LSE order book, using an appropriate parametric model which views orders and cancellations as a mutually-exciting ten-variate Hawkes point process. Consistent with findings in the related literature, in over 60 per cent of cases, the order book does not replenish reliably after a large trade. However, if it does replenish, it does so with a fairly fast half life of around 20 s. Various other dynamics are quantified.

© 2006 Elsevier B.V. All rights reserved.

JEL classification: C32; C51; C52; G10

Keywords: Market microstructure; Limit order book; Resiliency; Point process; Conditional intensity

1. Introduction

Traders often prefer an electronic limit order book which is resilient, so that the supply of unmatched offers to trade (i.e., limit orders), when substantially depleted in a large

[☆] The paper benefited greatly from the comments of seminar participants at ECARES, Brussels (October 2005) and Nuffield College, University of Oxford (February 2004), among them David Hendry. I'm particularly grateful to Clive Bowsher, Bruce Lehmann, Bent Nielsen, Neil Shephard and an anonymous referee for their considerable support and valuable suggestions. I'm also grateful to the Bendheim Center for Finance for accommodating me at Princeton University during part of the writing of this paper. I gratefully acknowledge financial support from the UK Economic and Social Research Council, and the US-UK Fulbright Commission.

E-mail address: Jeremy.Large@Economics.Ox.Ac.Uk.

transaction, then replenishes rapidly at or near previously prevailing prices. This resilient replenishment is out of the exchange's hands, being realized through the order submissions of participating traders, which are momentary events of unpredictable timing.

To quantify resiliency, this paper models the random timing of these central events directly. It uses the categorization of event types in [Biais et al. \(1995\)](#) as a way to distinguish resiliency events and the shocks that precede them from other orders. A vector of intensities thus emerges, one for each category of event. By estimating 'mutually-exciting' interactions among these intensities, the paper quantifies resiliency in easily interpretable terms. Specifically, it assumes that if a limit order book lacks resiliency it replenishes too unreliably, or too late. These two dimensions, conditional probability and delay, are distinguished by estimating impulse response functions in continuous time which can be judged either too small or too slow. Four such functions are distinguished, describing respectively the response at the bid or ask to a large purchase or sale. Comparing the four functions adds a third dimension, price, to the resiliency study.

For Barclays equity on the London Stock Exchange, the estimated impulse response functions all have rather short half-lives: of under 20 s. However, they are fairly slight, only implying a resilient response to big trades in under 40 per cent of cases. Thus when it occurs, the resilient response to a shock is fast—but it occurs quite infrequently. This is consistent with the dynamic illiquidity reported in electronic foreign exchange markets by [Danielsson and Payne \(2002\)](#). It also corroborates the long memory reported in [Degryse et al. \(2005\)](#), whereby wide spreads sometimes narrow only in the distant (by microstructure standards) future. The relationship to the findings of [Coppejans et al. \(2004\)](#) is difficult to define, because they study the shape of the order book at a lower, five minute, frequency. However, it is consistent with the periodic liquidity crises they report following sharp price movements. The study also finds that after spreads widen, liquidity supply is equally likely to be restored on either side of the book. This accords with a theory of market order price impact, whether information-bearing or not.

The paper's fully parametric model is rich enough to capture and control for the various dynamics of trades, orders and order cancellations. This leads to an abundance of parameters. However, its closed-form log-likelihood function has an additive separability, permitting the econometrician to estimate and interpret only a subset containing the parameters pertaining to resiliency. Of course, when the model is used for other purposes, other parameter subsets can be estimated. Taking an interest in the drivers of large market orders, the paper does report some extra parameters, finding among other things evidence of copycat market order submission. Some additional results emerge as a byproduct of the resiliency estimation; for example, we learn how order cancellation is sometimes followed by very rapid replacement.

The model is most applicable on markets which are not yet highly resilient: on markets such the Chicago Board of Trade's (CBOT's) market for the 10 Year US Treasury Bond Future, spreads are very seldom wider than the minimum price increment (under 1 per cent of the time), rendering the measures of resiliency proposed here somewhat singular and less useful. The paper is closely related to [Bauwens and Hautsch \(2004\)](#), [Engle \(2000\)](#), [Engle and Russell \(1998\)](#), and [Engle and Lunde \(2003\)](#). [Hasbrouck \(1999\)](#) and [Bisière and Kamionka \(2000\)](#) also study the limit order book as a multivariate point process.

The paper proceeds as follows. Section 2 discusses the concept of resiliency and its econometric interpretation in continuous time. Section 3 proposes an approach to identifying those events on an order book that characterize resiliency. Section 4 presents

the econometric details of the model, while Section 5 brings it to the data. The results are discussed in Section 6, and Section 7 concludes. An Appendix provides a complete account of the results, and discusses data issues in more detail.

2. Order book resiliency

The taxonomy due to Kyle (1985) for the “slippery and elusive concept” of market liquidity remains current in the literature: he suggests its principal elements to be *tightness*—the narrowness of the bid-ask spread; *depth*—the rate at which the potential volume of a market order grows as its execution price deteriorates; and *resiliency*—“the speed with which prices recover from a random, uninformative shock.” The particular importance of resiliency in the case of an electronic limit order book, where the price-smoothing agency of a specialist or market maker is absent, has recently been noted by theorists and econometricians alike. In theoretical studies, Foucault et al. (2005) and Rosu (2004) propose conditions under which resiliency is high, while Degryse et al. (2005) and Coppejans et al. (2004) study the phenomenon empirically using non-parametric and VAR-based techniques.

As mentioned in the Introduction, the model here is different. Limit order books recover from large trades through a moderate number of instantaneous events that bring spreads down to normal levels, namely the submissions of fresh limit orders. The question of interest can thus be framed as how fast these events occur, or indeed how likely they are to occur at all. This motivates the paradigm of Engle (2000), which proposes treating high frequency market data as the realization of a point process in continuous time: that is, as a succession of distinct, momentary events that occur at random times. Assuming it is simple with intensity adapted to its natural filtration,¹ such a process can be fully specified by giving its (typically vector-valued) intensity’s functional dependence on that filtration. Define its intensity at time t conditional on its natural filtration up to but not including time $s \leq t$, $\mathcal{F}_s$, as

$$\lambda(t|\mathcal{F}_s) = \lim_{\delta t \downarrow 0} \frac{\Pr[N(t + \delta t) > N(t)|\mathcal{F}_s]}{\delta t}, \quad (1)$$

where $N(t)$ is the number of events to have occurred up to and including time t . Refer to $\lambda(t|\mathcal{F}_t)$ simply as $\lambda(t)$, the value of the process λ at time t . Also define $\mathcal{F}_{s+}$, the natural filtration up to and including time s .

Within this framework, resiliency is captured simply by the way large trades alter the future intensities of fresh limit order submissions (these events are defined precisely in Section 3). In particular, if intensities are for a time greatly increased following the liquidity demand shock embodied by a large trade, then resiliency is high. The rate at which they tend back to average levels gives a measure of how fast the resilient order book responds to the shock. More formally, if a shock of type m , denoted $\mathcal{E}_m$, happens at time s , then its impact on these intensities, say λ_r , at a later time t can be defined as

$$\lambda_r(t|\mathcal{F}_{s+}) - \lambda_r(t|\mathcal{F}_s), \quad (2)$$

¹A point process is simple if conditional on $\mathcal{F}_t$ as δt becomes smaller, the probability of observing two or more events in $(t, t + \delta t)$, divided by the probability of observing one event in $(t, t + \delta t)$, tends to zero.

where $\lambda_r(t|\mathcal{F}_{s+}) = \lambda_r(t|\mathcal{F}_s, \mathcal{E}_m)$ includes knowledge of $\mathcal{E}_m$. If resiliency is invariant over time, i.e., if $\mathcal{E}_m$ has the same effect whatever the value of s , then there exists a function G_{rm} such that

$$G_{rm}(t-s) = \lambda_r(t|\mathcal{F}_{s+}) - \lambda_r(t|\mathcal{F}_s). \quad (3)$$

Thus defined, G_{rm} serves as a continuous time analogue to the impulse response function used by Coppejans et al. (2004). In the parametrization proposed later it rises before dying away in the short or medium term. From the definition of intensity, it follows that (if the integral is finite)

$$\int_0^\infty G_{rm}(u) du \quad (4)$$

is the expected increase in the number of future resiliency events at time s compared to the counterfactual world where $\mathcal{E}_m$ had not occurred. This concept is complemented by a related notion of G_{rm} 's 'half-life,' which gives a convenient description of the rate at which the order book reverts to normal after a shock.

3. Identifying resiliency-related events

During continuous trading in January 2002 the shape of the limit order book on the London Stock Exchange SETS platform for Barclays Plc stock, whose price averaged £22.40, underwent 117.5 thousand changes, or 10.5/min. The times of these changes were recorded to the nearest second. Each such change was due to a single order submission or cancellation. Among these were all the shocks that widened spreads, as well as all the events through which the order book replenished resiliently.

As a precursor to the econometric implementation, criteria are needed identifying the subset of events which pertain to resiliency. Kyle (1985)'s definition of resiliency, cited in Section 2, concerns "random, uninformative" liquidity demand shocks, ruling out shocks which are informative or expected. Via the estimated variation in their intensities, the current model finds liquidity demand shocks at times more expected, at time less expected. However, no shock is fully unexpected. Similarly, large trades can involve more or less private information, and this can be measured in their long-term market impact. However, long-term impact is not observed by the traders who in the near term replenish the limit order book resiliently. Nor is private information identified in this reduced-form econometric approach. Given these issues, the paper studies resiliency after *any* liquidity demand shock, not just those that are "random" and "uninformative".

Define an *aggressive* order to be one that produces exceptional changes in the economic outcomes available to traders. By this criterion, a liquidity demand shock is an aggressive market order. Equally, an aggressive limit order, where it follows an aggressive market order, is characteristic of resilient replenishment. Biais et al. (1995) propose a price-based technique to identify aggressive market and limit orders. Their categorization is applied to the current data and recorded in Table 1.² They think of aggressiveness as a relative term. Here, however, we will judge an order to be aggressive if and only if it changes the prevailing best bid or best ask: limit orders by improving on the current best quoted price; market orders by completely filling the best-priced opposing limit orders. In the current

²On the buy side and on the sell side we combine two categories in Biais et al. (1995).

Table 1

Table showing the categorization and characteristics of events. MO refers to market orders, LO to limit orders. The average volume statistics are in numbers of shares

Type number	Order of submission or cancellation	Buy or sell?	Triggers immediate execution?	Moves prices ?	Named	Number in sample	Average volume
1	S	B	✓ (MO)	✓	Market buy that moves the ask	3,673	4,830
2	S	S	✓ (MO)	✓	Market sale that moves the bid	4,351	4,790
3	S	B	× (LO)	✓	Bid between the quotes	5,292	3,390
4	S	S	× (LO)	✓	Ask between the quotes	5,281	3,940
5	S	B	✓ (MO)	x	Market buy that doesn't move the ask	5,885	2,660
6	S	S	✓ (MO)	x	Market sale that doesn't move the bid	5,114	2,550
7	S	B	× (LO)	x	Bid at or below best bid	25,121	2,840
8	S	S	× (LO)	x	Ask at or above best ask	22,433	3,110
9	C	B	—	—	Cancelled bids	20,921	2,740
10	C	S	—	—	Cancelled asks	19,424	2,820

data, such orders (the top four entries in Table 1) number only 18.6 thousand, 16 per cent of all events.

An alternative means of identifying aggressive orders is via their offered or traded volume of shares. A high volume is indicative of aggressiveness. A quantile cutoff could be defined above which large orders are viewed as aggressive. However, as the choice of such a cutoff (or related notion) is not immediate, here I favor the [Biais et al. \(1995\)](#) methodology to identify aggressiveness, nevertheless including volume as a conditioning variable of secondary importance.

Aggressive and unaggressive activity together constitute ten types of event, producing a 10-variate point process. They are numbered as shown in Table 1. Thus, events of type 1 are liquidity demand shocks due to buying, and events of type 2 are those due to selling. Similarly, events of type 3 are resiliency events involving bids, and events of type 4 are resiliency events through asks. Four impulse response functions will therefore be of interest: from types 1 and 2 to types 3 and 4—i.e., $G_{3,1}$, $G_{4,1}$, $G_{3,2}$ and $G_{4,2}$.

4. Econometric model

[Hasbrouck \(1999\)](#) also categorizes the events on a limit order book to give a multivariate point process, but does not parameterize its intensity. In a univariate framework, [Engle \(2000\)](#) parameterizes the intensity λ in (1) indirectly using an Autoregressive Conditional Duration (ACD) model, introduced in [Engle and Russell \(1998\)](#). The ACD model is extended to multiple event types in [Engle and Lunde \(2003\)](#), and used to generate an




Fig. 1. Illustration of the self-exciting nature of the point process specification in (5). In this example crosses indicate the times of past events.

explanatory variable in the study of [Engle and Lange \(2001\)](#). The current paper departs from [Engle \(2000\)](#) by, as in the Latent Factor Intensity (LFI) model of [Bauwens and Hautsch \(2004\)](#), giving a direct specification of the intensity rather than one via the durations between events. This provides an analytic representation of the expression in (3). The model is due to [Bowsher \(2005\)](#), which draws from [Hawkes \(1971\)](#). The data is viewed as a 10-variate point process, driven by a 10-variate counting process, N , over positive times, $\mathbb{R}^+$ (so that N_t is a 10-vector whose entries record the number of elapsed events of each type before time t), with an associated intensity vector, λ , which is adapted to the process' natural filtration.³ As shown in [Bowsher \(2005\)](#), the model is fully specified by the equation for λ ,

$$\lambda(t; \theta) = \mu(t; \theta) + \int_{(0,t)} W(t-u; \theta) dN(u), \quad (5)$$

where $\mu > 0$ is a deterministic 10-vector-valued function on $\mathbb{R}^+$, $W > 0$ is a 10×10 matrix-valued function on $\mathbb{R}^+$, and the integration is Stieltjes. The vector θ contains the parameters. In later discussions of $\lambda(t)$, $\mu(t)$ and $W(t)$, θ is left implicit. The function W contains the autoregressive exciting effects, whereby events in the past increase current intensities. [Fig. 1](#) illustrates, this for the simpler univariate case. Note that W appears horizontally inverted in this figure. The multivariate case is more complex, but somewhat clarified by inferring directly from (5) the following characterization of λ_r , for any $r \leq 10$:

$$\lambda_r(t) = \mu_r(t) + \sum_{m=1}^{10} \left(\int_{(0,t)} W_{rm}(t-u) dN_m(u) \right), \quad (6)$$

which contains only univariate Stieltjes integration. So in the 10-variate case, on each intensity, λ_r , 10 of these autoregressive exciting effects (i.e., $W_{rm} : m \leq 10$) operate additively and independently, each capturing the past occurrences of a different type of event, m . These are added to a non-dynamic underlying part, μ_r . The diagonal

³The econometric model works identically for any dimension, M , not just $M = 10$.

entries of W will be called self-exciting effects, while off-diagonal entries are cross-exciting effects. Note that if $\mu \equiv 0$ then the exciting effects in W would never be realized, and so $\lambda \equiv 0$. For non-trivial results we therefore require that μ is somewhere positive.⁴

As will be explained in detail later, W alone determines the impulse response functions discussed in the Introduction. It has 100 entries, as each of the ten types of event potentially responds to every other type, whether aggressive or unaggressive.

In the current application, however, resiliency will be captured by only four of these 100 entries, where aggressive limit orders (buy and sell) respond to aggressive market orders (buy and sell). I follow [Bowsher \(2005\)](#) in incorporating intraday seasonality into μ via a piecewise linear spline. This controls for time of day effects in a simultaneous estimation.

This modeling approach to the limit order book is closest to that of [Bisière and Kamionka \(2000\)](#), which also estimates a point process model using, as this paper, the order segmentation of [Biais et al. \(1995\)](#). However, their specification is via the durations between events, whose distributions depend only on the type of event that began them and the concurrent state of the limit order book. By contrast, the very last event to have occurred has no special status in the current model. The advantages of directly specifying the intensity process also accrue to the LFI model of [Bauwens and Hautsch \(2004\)](#). At the cost of some computational complexity, the LFI generalizes the current framework to the case where λ is adapted to a wider filtration than that of the observed point process. Other related papers include [Dufour and Engle \(2000\)](#), [Russell \(1999\)](#), [Heinen and Rengifo \(2003\)](#) and [Hollifield et al. \(2004\)](#). Finally, note that a re-parametrization here in terms of durations is not straightforward: at any time t the conditional duration until the next event substantially depends on the type of that event, as well as on the timings and types of all prior events.

4.1. Resiliency in the econometric model

Section 2 showed how a point process model can quantify resiliency by estimating rises in the intensities of resiliency events subsequent to a shock. Resiliency depended on the function $G_{rm}(t-s)$, which was defined as

$$\lambda_r(t|\mathcal{F}_s, \mathcal{E}_m) - \lambda_r(t|\mathcal{F}_s),$$

where a shock of type m , $\mathcal{E}_m$, occurs at s , and λ_r is the intensity of resiliency events. As discussed in Section 3, the liquidity demand shocks denoted $\mathcal{E}_m$ are identically events of types 1 and 2 (so m takes two values: 1 and 2), while resiliency events are types 3 and 4 (so r takes values 3 and 4). It is mathematically convenient to embed these four functions, $G_{3,1}(t)$, $G_{4,1}(t)$, $G_{3,2}(t)$ and $G_{4,2}(t)$, within the complete matrix of response functions, which has 100 entries like the matrix $W(t)$.

⁴Avoiding the notation of Stieltjes integration, (6) may be written as

$$\lambda_r(t) = \mu_r(t) + \sum_{m=1}^{10} \left(\sum_{i=1}^{\infty} W_{rm}(t-t_i) \mathbb{I}_{z_i=m} \mathbb{I}_{t_i < t} \right), \quad (7)$$

where t_i denotes the time, and z_i the type (from 1 to 10) of the i th event in the pooled point process. $\mathbb{I}$ denotes the indicator function.

Definition 1. Let G be a 10×10 matrix of functions defined on $\mathbb{R}^+$. Denoting a representative column thereof by G_m , for any $t > s$,

$$G_m(t-s) \equiv \lambda(t|\mathcal{F}_s, \mathcal{E}_m) - \lambda(t|\mathcal{F}_s), \quad (8)$$

should this be invariant to $s \geq 0$, where $\mathcal{E}_m$ is an event of type m at time s .

Interpretation in terms of causality. Suppose that a shock of type m occurs at time s . Then G 's m th column, $G_m(t-s)$, contains the expected increase at time $t > s$ in the 10 intensities compared to a counterfactual world in which the shock has not occurred. This *counterfactual comparison*, if estimated to be significantly positive, is suggestive of a causal connection from the shock to future events. However, it does not rule out the possibility of an unobserved common cause, nor that the shock be caused by subsequent events, via traders' expectations. I will therefore use in this context the more moderate language that the shock *precipitates* future effects on the limit order book.

Proposition 4.1. In the current model G exists and satisfies the integral equation

$$G(t) = W(t) + \int_{(0,t)} W(t-u)G(u)du. \quad (9)$$

Furthermore,

$$\int_0^\infty G(t)dt = \left\{ Id - \int_0^\infty W(t)dt \right\}^{-1} \int_0^\infty W(t)dt. \quad (10)$$

Proof. See Appendix B. $\square$

Discussion. G is therefore determined by W . Eq. (9) shows it is the sum of two parts. The first part is W itself, and captures the *direct* exciting effects of a shock on future intensities. The second part, a convolution of G and W , captures *indirect*, chain reaction or domino effects, whereby the shock precipitates intermediate events, which themselves augment future intensities. This distinction is illustrated in Fig. 2. As noted in Section 2, integrating G over $\mathbb{R}^+$ gives a matrix whose $\{r, m\}$ element is the expected increase in future events of type r given an event of type m . The second part of Proposition 4.1 shows how this expectation (for any m and r) can be estimated given an estimate of W .




Fig. 2. Illustrates the response in the intensity of events of type r after a shock of type m at time s . *Direct* and *indirect* (i.e., chain reaction or domino) effects are distinguished.

Using the event definitions of Section 3, the impulse response functions of interest are $G_{3,1}$, $G_{4,1}$, $G_{3,2}$ and $G_{4,2}$. To estimate any one of these requires W to be estimated in its entirety. However, concentrating only on their direct component substantially reduces the number of pertinent parameters, since then only $W_{3,1}$, $W_{4,1}$, $W_{3,2}$ and $W_{4,2}$ need be estimated. If indeed only a small part of W is of interest, it will be desirable to reduce the computational task by disregarding its irrelevant parts. The following corollary of [Bowsher \(2005\)](#) provides conditions for the parametrization under which this is feasible.

Corollary 4.2. *Suppose that the data is observed over the interval $[0, T]$. Let λ_r be the univariate intensity of the r 'th event type, $r = 1 \dots 10$. Suppose that θ can be partitioned into 10 sets, denoted θ_r , such that for all r , λ_r depends on θ_r but on no other parameters. Suppose that $\theta_1, \theta_2, \dots, \theta_{10}$ are variation-free. Then MLEs for θ_r can be found by maximizing*

$$\int_0^T -\lambda_r(u) du + \sum_{i: t_{i,r} < T} \log \lambda_r(t_{i,r}), \quad (11)$$

where $\{t_{i,r}\}$ is the sequence of the times of the events of type r .

Proof. See [Bowsher \(2005, Theorem 3.1\)](#). This follows from the additive separability of the log-likelihood function. $\square$

I therefore provide a parametrization, θ , which has an appropriate partition so that this Corollary applies. Then, to estimate $W_{3,1}$, $W_{4,1}$, $W_{3,2}$ and $W_{4,2}$, (11) need only be maximized for event types 3 and 4. Thus, 20 of the 100 entries in the matrix W are estimated. While only four of these are of interest for resiliency, the other 16 provide information about order book dynamics which are somewhat incidental to its resiliency, but are reported for Barclays in the empirical implementation. As mentioned in the Introduction, the intensities of Barclays' liquidity demand shocks, types 1 and 2 events, are also studied in the same way.

4.2. Parametrization

To complete the model, parametrizations of μ and W are now given which satisfy the conditions of Corollary 4.2. For each r , μ_r is taken as a daily seasonal piecewise linear spline with knots at 8am (when the trading day begins), 8:30am, 9am, 12pm, 3pm and 4:30pm (the end of the trading day).⁵ Set $W(0) \equiv 0$. An exponential decay is fitted to W 's off-diagonal entries: for $u > 0$

$$W_{rm}(u) = \alpha_{rm} \beta_{rm} e^{-\beta_{rm} u}, \quad (12)$$

whenever $r \neq m$, so that W_{rm} has area α_{rm} and a half life of $\ln(2)/\beta_{rm}$. This describes cross-exciting effects, where an event precipitates an event of a different type.

The sum of two exponentials is used for diagonal entries where due to self-exciting effects an event precipitates more of its own type of event, so

$$W_{rr}(u) = \alpha_{rr} \beta_{rr} e^{-\beta_{rr} u} + \gamma_r \delta_r e^{-\delta_r u}. \quad (13)$$

Impose the constraints that α_{rm} , β_{rm} , γ_r and δ_r are strictly positive for all m and r . Finally, β_{rr} and δ_r are identified by requiring that $\beta_{rr} > \delta_r$. This model choice reflects an important

⁵ $\mu(t)$ is set to zero after the end of the trading day.

feature found in the order book data: that self-exciting effects tend to have a quick impact, but also a longer lasting autoregressive intensity component. The longer lasting component is captured by the expression $\gamma_r \delta_r e^{-\delta_r u}$.

It was mentioned in Section 3 that an alternative way to identify aggressive orders, not pursued here, is to look for those with exceptionally high traded or offered share volumes. To assess the importance of share volume here, I additionally allow α_{rm} to depend on the share volume of the preceding event, say v_i :

$$\alpha_{rm} \equiv \eta_{rm} e^{\phi_{rm} v_i}, \quad (14)$$

where η_{rm} and ϕ_{rm} are further parameters. An identical formulation is adopted for γ_r using further parameters $\tilde{\eta}_r$ and $\tilde{\phi}_r$: $\gamma_r \equiv \tilde{\eta}_r e^{\tilde{\phi}_r v_i}$. When estimated for the current data, ϕ_{rm} and $\tilde{\phi}_r$ are sometimes positive and sometimes negative, but are mainly close to zero. In presenting the results v_i is normally set to 5,000.

Finally, note that the parameter set, θ , thus defined, has a partition so that Corollary 4.2 applies, with

$$\theta_r = \{\eta_{rm}; \phi_{rm}; \beta_{rm}; \tilde{\eta}_r; \tilde{\phi}_r; \delta_r; \text{knots of } \mu_r : m = 1, \dots, 10\}, \quad (15)$$

for in the model the intensity λ_r depends on these parameters alone.

4.3. Inference

Ogata (1978) provides theoretical results on the asymptotic distribution theory of the MLEs, $\hat{\theta}$, for the parameters, θ , of Corollary 4.2, provided that the process is univariate and stationary. He shows that as the observed period grows without limit, the MLEs are consistent and asymptotically normal. This provides a basis for inference since the MLEs' covariance matrix can then be estimated with the usual numerical calculations of second derivatives.

However, the current model is multivariate, as well as non-stationary, since it depends on the time of day via the spline μ . This paper follows Bowsher (2005) in assuming that the MLEs are similarly consistent in this more general problem, though to my knowledge this remains unproven in the literature. Bowsher (2005) addresses the problem via simulation, with favorable results.

4.4. Mis-specification testing

If the rule determining the intensity λ of a univariate point process is known, then a procedure to stretch or compress the time scale can be performed, which transforms the point process into a Poisson process of parameter 1.⁶ Mis-specification tests can concentrate on the goodness of this candidate Poisson process as derived from the data and estimation.

The time-varying multiplicative factor to be applied to stretch or compress the time scale is simply the intensity. This is intuitive, for it means that if at a given moment events are arriving with intensity λ then they would arrive with intensity 1 if time was passing at rate λ . Unfortunately, the true λ is not available: however it may be replaced by the estimated

⁶Theorem 4.1 (p.14) of Bowsher (2005) shows this, provided that with probability one any realization of the process produces events of all types indefinitely (see its Eq. (17)).

intensity as implied by the MLEs, $\hat{\theta}$. For given event type, r , define the time-deformed sequence of durations by

$$\left\{ \int_{t_{(i-1),r}}^{t_{i,r}} \lambda_r(t; \hat{\theta}) dt : i = 1, 2, \dots \right\}, \quad (16)$$

where $\{t_{i,r}\}$ is the sequence of the times of the events of type r . The mis-specification test adopted here is to check that (16) lacks autocorrelation at lags of 15 and 25 using Box–Ljung,⁷ and lacks excess dispersion using the statistic proposed by Engle and Russell (1998), $\sqrt{N}((\hat{\sigma}^2 - 1)/\sqrt{8})$, where $\hat{\sigma}^2$ is the sample variance of the sequence in (16) and N is its length. Further assessment of fit is provided by QQ plots against an exponential distribution.

5. Empirical implementation

In January 2002 there were 22 days of trading on the LSE SETS electronic limit order book for Barclays. On each day, the list of all events was viewed as a 10-variate point process, using the categorization proposed in Section 3. Viewing each trading day as an independent realization of the proposed data-generating process, the conditional log-likelihood function in Corollary 4.2 was numerically maximized. This was performed in Ox (see Doornik, 2001), applying the MaxBFGS algorithm with numerical derivatives in turn to each of the event types 1, 2, 3 and 4 (each maximization took 4–6 h). The estimated model passes nine of twelve specification tests performed on it at the 10 per cent confidence level, and a further one at the 1 per cent confidence level. QQ plots of the time-deformed durations are presented in Fig. 7. Some reassurance that global, not local, maxima were found in the numerical maximization is given by the striking similarities between the separately and independently estimated dynamics of bids and asks (types 3 and 4), as well as between those of market purchases and sales (types 1 and 2). A sample of the estimated intensity path over time is presented in Fig. 3.

A problem arises regarding the granularity of the data, which as previously noted is rounded to the nearest second: 40 per cent of events therefore have the same time-stamp as some other event. The treatment of these multiple events, which is discussed in detail in the Appendix, rules out exciting effects between events with the same timing. This is most costly on the description of exciting effects that attenuate fast. However, it emerges from the estimation that the exciting effects pertinent to resiliency have half-lives greater than these, making them relatively immune to the rounding.

6. Results for the Barclays limit order book

The results are framed in terms of *direct* effects as captured by W and exclude the indirect or knock-on effects that are additionally included in G , as in (9). Full results are extensive, and are reported in Tables 2–5. Figs. 4 and 5 provide more accessible bubble chart representations of the main results for the reader, including those about resiliency. These two figures each contain two charts. All four charts are interpretable in the same way. Each relates to the causes of a single event type $r = 1, 2, 3$ or 4. They are estimated

⁷I also check for autocorrelation in the squared durations at lags of up to 25.




Fig. 3. The estimated intensity of aggressive asks on the morning of 2 January 2002 (events per hour).

separately and independently. They show discs, each of which represents a single exciting effect, i.e., an element of the matrix W , say W_{rm} . Its position on the vertical axis gives its magnitude, α_{rm} , while its position on the horizontal axis (a log scale) gives its half life, $\ln 2/\beta_{rm}$. The magnitude measures the average number of events of type r precipitated by each event of type m . The diameter of the disc shows the number of observed events of type m , the type that precipitates the effect. The discs on the left-hand charts are labeled with the type of the precipitating event, m .⁸ The shadings of the discs on adjacent charts are related: passing from one chart across to the other, the same shading is used for pairs of exciting effects which are buy-sell mirrors of each other. Given this, for simplicity the right-hand charts are not labeled: m is not given for each disc, but can be inferred by comparison with the chart to the left. All effects are those of orders for 5,000 shares. The information in Figs. 4 and 5 is discussed next. Subsequent statements concerning the dynamics of the limit order book are made in ordinary language, but they have a precise probabilistic interpretation concerning W .

6.1. Barclays' quantified resiliency

Barclays' resiliency is captured in Fig. 4, whose charts both contain a pair of isolated discs in their top left quadrant. These four discs represent the four impulse response functions describing resiliency: through the exciting effects of aggressive market orders (buy and sell) on aggressive limit orders (buy and sell). Their lateral position measures the half-life of resilient reaction, while their height measures their magnitudes. The parameter values they record are: $\alpha_{4,1} = 0.17$; $\ln 2/\beta_{4,1} = 15$ s; $\alpha_{3,1} = 0.15$; $\ln 2/\beta_{3,1} = 19$ s; $\alpha_{4,2} = 0.19$; $\ln 2/\beta_{4,2} = 16$ s; $\alpha_{3,2} = 0.17$; $\ln 2/\beta_{3,2} = 12$ s.

⁸In Fig. 5 some unimportant discs are not labeled.

Table 2

Estimated parameters determining the intensity of bids between the quotes

This table reports results for the arrival of events of type $r = 3$. The model is defined by

$$\lambda_r(t) = \mu_r(t) + \sum_{m=1}^{10} \left(\int_{(0,t)} W_{rm}(t-u) dN_m(u) \right),$$

where $\mu_r(t)$ is a piecewise linear intraday spline and for $m \neq r$

$$W_{rm}(u) = \alpha_{rm} \beta_{rm} e^{-\beta_{rm} u} \text{ and } W_{rr}(u) = \alpha_{rr} \beta_{rr} e^{-\beta_{rr} u} + \gamma_r \delta_r e^{-\delta_r u},$$

subject to the constraint that $\beta_{rr} > \delta_r$. Furthermore

$$\alpha_{rm} \equiv \eta_{rm} e^{\phi_{rm} v_i} \text{ and } \gamma_r \equiv \tilde{\eta}_r e^{\tilde{\phi}_r v_i},$$

where v_i is the share volume of the shock. All parameters except ϕ_{rm} and $\tilde{\phi}_r$ are > 0 . Throughout, upper and lower 95% confidence intervals are given, using the basis for inference set out in Section 4.3. The table consists of four stacked arrays. Setting $v_i = 5,000$, the top array reports α_{3m} and $\ln(2)/\beta_{3m}$ for all m . In its lowest line, ‘autoregressive intensity effect,’ it reports γ_3 and $\ln(2)/\delta_3$. The second array reports $1,000\phi_{3m}$ for all m , and, in its lowest line, $1,000\tilde{\phi}_3$. The third array reports specification test results as detailed in Section 4.4. The last array reports the knots of μ_3 .

Type of the precipitating event	Number of events	Effect of event of 5,000 shares (number of events)			Half-life of effect (s)		
		Lower	Estimate	Upper	Lower	Estimate	Upper
Cancelled bids	20,921	0.02	0.02	0.03	1.2	1.3	1.4
Bids at or below best bid	25,121	0.01	0.01	0.02	23.47	32.08	43.83
Bids between the quotes	5,292	0.04	0.06	0.09	43.5	69.2	110.0
Market buys that don't move ask	5,885	0.01	0.02	0.03	10.45	17.81	30.33
Market buys that do move ask	3,673	0.14	0.15	0.17	16.7	19.1	21.8
Cancelled asks	19,424	0.03	0.03	0.04	58.0	84.1	121.6
Asks at or above best ask	22,433	n/a	n/a	n/a	n/a	n/a	n/a
Asks between the quotes	5,281	0.02	0.03	0.06	29.8	53.2	94.8
Market sales that don't move the bid	5,114	n/a	n/a	n/a	n/a	n/a	n/a
Market sales that do move the bid	4,351	0.16	0.17	0.19	10.8	12.4	14.1
Autoregressive intensity effect	5,292	0.18	0.22	0.26	548	774	1094
Type of the precipitating event	Elasticity of exiting magnitude to a 1,000 increase in share volume						
				Lower	Estimate	Upper	
Cancelled bids				−0.23	−0.18	−0.13	
Bids at or below best bid				0.04	0.04	0.05	
Bids between the quotes				0.02	0.03	0.05	
Market buys that don't move ask				0.05	0.06	0.07	
Market buys that do move ask				−0.03	−0.01	0.01	
Cancelled asks				0.02	0.03	0.05	
Asks at or above best ask				n/a	n/a	n/a	
Asks between the quotes				0.06	0.08	0.11	
Market sales that don't move the bid				n/a	n/a	n/a	
Market sales that do move the bid				0.01	0.02	0.02	
Autoregressive intensity effect				−0.02	0.01	0.03	
Specification testing							

Table 2 (continued)

Specification testing			
Time-adjusted durations			
Mean (should be 1)			0.98
Var. (should be 1)			1.02
Robust Ljung–Box (15 lags) of squares <i>p</i> -value			0.24
Robust Ljung–Box (25 lags) of squares <i>p</i> -value			0.56
Test of excess dispersion <i>p</i> -value			0.28
Likelihood			13,917
Intraday spline			
Time	Intensity at node (events/h)		
	Lower	Estimate	Upper
08:00	22.1	25.1	28.6
08:30	7.1	9.3	12.0
09:00	3.3	4.5	6.3
12:00	0.1	0.5	2.6
15:00	2.5	3.5	5.1
16:30	1.3	3.1	7.0

Table 3

Estimated parameters determining the intensity of asks between the quotes

This table reports results for the arrival of events of type $r = 4$. The model is defined by

$$\lambda_r(t) = \mu_r(t) + \sum_{m=1}^{10} \left(\int_{(0,t)} W_{rm}(t-u) dN_m(u) \right),$$

where $\mu_r(t)$ is a piecewise linear intraday spline and for $m \neq r$

$$W_{rm}(u) = \alpha_{rm} \beta_{rm} e^{-\beta_{rm} u} \text{ and } W_{rr}(u) = \alpha_{rr} \beta_{rr} e^{-\beta_{rr} u} + \gamma_r \delta_r e^{-\delta_r u},$$

subject to the constraint that $\beta_{rr} > \delta_r$. Furthermore

$$\alpha_{rm} \equiv \eta_{rm} e^{\phi_{rm} v_i} \text{ and } \gamma_r \equiv \tilde{\eta}_r e^{\tilde{\phi}_r v_i},$$

where v_i is the share volume of the shock. All parameters except ϕ_{rm} and $\tilde{\phi}_r$ are > 0 . Throughout, upper and lower 95% confidence intervals are given, using the basis for inference set out in Section 4.3. The table consists of four stacked arrays. Setting $v_i = 5,000$, the top array reports α_{4m} and $\ln(2)/\beta_{4m}$ for all m . In its lowest line, ‘autoregressive intensity effect,’ it reports γ_4 and $\ln(2)/\delta_4$. The second array reports $1,000\phi_{4m}$ for all m , and, in its lowest line, $1,000\tilde{\phi}_4$. The third array reports specification test results as detailed in Section 4.4. The last array reports the knots of μ_4 .

Type of the precipitating event	Number of events	Effect of event of 5,000 shares (number of events)			Half-life of effect (s)		
		Lower	Estimate	Upper	Lower	Estimate	Upper
Cancelled bids	20,921	0.03	0.04	0.05	64.1	89.7	125.2
Bids at or below best bid	25,121	n/a	n/a	n/a	n/a	n/a	n/a
Bids between the quotes	5,292	0.02	0.04	0.08	115.5	275.7	638.2
Market buys that don’t move ask	5,885	n/a	n/a	n/a	n/a	n/a	n/a
	3,673	0.15	0.17	0.18	12.4	14.5	16.9

Table 3 (continued)

Type of the precipitating event	Number of events	Effect of event of 5,000 shares (number of events)			Half-life of effect (s)		
		Lower	Estimate	Upper	Lower	Estimate	Upper
Market buys that do move ask							
Cancelled asks	19,424	0.02	0.03	0.03	1.4	1.5	1.6
Asks at or above best ask	22,433	0.02	0.02	0.03	33.0	46.4	65.1
Asks between the quotes	5,281	0.05	0.07	0.10	38.6	58.0	87.0
Market sales that don't	5,114	0.01	0.01	0.02	5.7	9.8	16.9
move the bid							
Market sales that do move the bid	4,351	0.17	0.19	0.20	14.4	16.0	17.8
Autoregressive intensity effect	5,281	0.12	0.16	0.21	665	975	1431
Type of the precipitating event	Elasticity of exciting magnitude to a 1,000 increase in share volume						
				Lower	Estimate	Upper	
Cancelled bids				0.01	0.03	0.05	
Bids at or below best bid				n/a	n/a	n/a	
Bids between the quotes				0.01	0.03	0.05	
Market buys that don't move ask				n/a	n/a	n/a	
Market buys that do move ask				0.01	0.02	0.03	
Cancelled asks				−0.17	−0.12	−0.06	
Asks at or above best ask				0.04	0.05	0.06	
Asks between the quotes				0.06	0.07	0.09	
Market sales that don't move the bid				−0.01	0.01	0.03	
Market sales that do move the bid				−0.01	0.00	0.00	
Autoregressive intensity effect				0.00	0.04	0.07	
Specification testing							
Time-adjusted durations							
Mean (should be 1)							0.98
Var. (should be 1)							1.01
Robust Ljung–Box (15 lags) of squares <i>p</i> -value							0.10
Robust Ljung–Box (25 lags) of squares <i>p</i> -value							0.02
Test of excess dispersion <i>p</i> -value							0.44
Likelihood							13,999
Intraday spline							
Time	Intensity at node (events/h)						
				Lower	Estimate	Upper	
08:00				23.5	26.6	30.1	
08:30				5.6	7.6	10.4	
09:00				1.6	2.6	4.2	
12:00				0.0	0.0	0.0	
15:00				1.3	2.2	3.7	
16:30				1.0	2.4	5.8	

Table 4
Estimated parameters determining the intensity of aggressive market purchases
This table reports results for the arrival of events of type $r = 1$. The model is defined by

$$\lambda_r(t) = \mu_r(t) + \sum_{m=1}^{10} \left(\int_{(0,t)} W_{rm}(t-u) dN_m(u) \right),$$

where $\mu_r(t)$ is a piecewise linear intraday spline and for $m \neq r$

$$W_{rm}(u) = \alpha_{rm} \beta_{rm} e^{-\beta_{rm} u} \text{ and } W_{rr}(u) = \alpha_{rr} \beta_{rr} e^{-\beta_{rr} u} + \gamma_r \delta_r e^{-\delta_r u},$$

subject to the constraint that $\beta_{rr} > \delta_r$. Furthermore

$$\alpha_{rm} \equiv \eta_{rm} e^{\phi_{rm} v_i} \text{ and } \gamma_r \equiv \tilde{\eta}_r e^{\tilde{\phi}_r v_i},$$

where v_i is the share volume of the shock. All parameters except ϕ_{rm} and $\tilde{\phi}_r$ are > 0 . Throughout, upper and lower 95% confidence intervals are given, using the basis for inference set out in Section 4.3. The table consists of four stacked arrays. Setting $v_i = 5,000$, the top array reports α_{1m} and $\ln(2)/\beta_{1m}$ for all m . In its lowest line, ‘autoregressive intensity effect,’ it reports γ_1 and $\ln(2)/\delta_1$. The second array reports $1,000\phi_{1m}$ for all m , and, in its lowest line, $1,000\tilde{\phi}_1$. The third array reports specification test results as detailed in Section 4.4. The last array reports the knots of μ_1 .

Type of the precipitating event	Number of events	Effect of event of 5,000 shares (number of events)			Half-life of effect (s)		
		Lower	Estimate	Upper	Lower	Estimate	Upper
Cancelled bids	20,921	0.01	0.01	0.01	2.8	3.5	4.4
Bids at or below best bid	25,121	0.00	0.00	0.01	6.31	10.13	16.26
Bids between the quotes	5,292	0.06	0.07	0.09	14.7	18.6	23.6
Market buys that don't move ask	5,885	0.13	0.13	0.14	7.38	8.11	8.92
Market buys that do move ask	3,673	0.08	0.09	0.10	12.3	15.0	18.2
Cancelled asks	19,424	0.00	0.00	0.00	4.5	10.1	22.5
Asks at or above best ask	22,433	0.00	0.00	0.00	27.6	49.1	87.4
Asks between the quotes	5,281	0.00	0.00	0.00	6.7	8.9	11.8
Market sales that don't move the bid	5,114	n/a	n/a	n/a	n/a	n/a	n/a
Market sales that do move the bid	4,351	n/a	n/a	n/a	n/a	n/a	n/a
Autoregressive intensity effect	3,673	0.25	0.31	0.37	951	1509	2396

Type of the precipitating event	Elasticity of exciting magnitude to a 1,000 increase in share volume		
	Lower	Estimate	Upper
Cancelled bids	0.03	0.04	0.05
Bids at or below best bid	0.03	0.04	0.06
Bids between the quotes	0.03	0.04	0.04
Market buys that don't move ask	0.03	0.03	0.04
Market buys that do move ask	0.02	0.03	0.03
Cancelled asks	−4.90	−4.05	−3.20
Asks at or above best ask	−5.73	−4.08	−2.42
Asks between the quotes	−1.63	−1.36	−1.08
Market sales that don't move the bid	n/a	n/a	n/a
Market sales that do move the bid	n/a	n/a	n/a
Autoregressive intensity effect	−0.01	0.02	0.04
Specification testing			

Table 4 (continued)

Specification testing			
Time-adjusted durations			
Mean (should be 1)			0.98
Var. (should be 1)			0.98
Robust Ljung–Box (15 lags) of squares p -value			0.68
Robust Ljung–Box (25 lags) of squares p -value			0.38
Test of excess dispersion p -value			0.63
Likelihood			8,450
Intraday spline			
Time	Intensity at node (events/h)		
	Lower	Estimate	Upper
08:00	1.3	2.3	4.0
08:30	6.7	8.3	10.3
09:00	1.4	2.2	3.6
12:00	0.2	0.8	3.8
15:00	2.4	3.5	5.0
16:30	4.3	6.4	9.5

Table 5

Estimated parameters determining the intensity of aggressive market sales

This table reports results for the arrival of events of type $r = 2$. The model is defined by

$$\lambda_r(t) = \mu_r(t) + \sum_{m=1}^{10} \left(\int_{(0,t)} W_{rm}(t-u) dN_m(u) \right),$$

where $\mu_r(t)$ is a piecewise linear intraday spline and for $m \neq r$

$$W_{rm}(u) = \alpha_{rm} \beta_{rm} e^{-\beta_{rm} u} \text{ and } W_{rr}(u) = \alpha_{rr} \beta_{rr} e^{-\beta_{rr} u} + \gamma_r \delta_r e^{-\delta_r u},$$

subject to the constraint that $\beta_{rr} > \delta_r$. Furthermore

$$\alpha_{rm} \equiv \eta_{rm} e^{\phi_{rm} v_i} \text{ and } \gamma_r \equiv \tilde{\eta}_r e^{\tilde{\phi}_r v_i},$$

where v_i is the share volume of the shock. All parameters except ϕ_{rm} and $\tilde{\phi}_r$ are > 0 . Throughout, upper and lower 95% confidence intervals are given, using the basis for inference set out in Section 4.3. The table consists of four stacked arrays. Setting $v_i = 5,000$, the top array reports α_{2m} and $\ln(2)/\beta_{2m}$ for all m . In its lowest line, ‘autoregressive intensity effect,’ it reports γ_2 and $\ln(2)/\delta_2$. The second array reports $1,000\phi_{2m}$ for all m , and, in its lowest line, $1,000\tilde{\phi}_2$. The third array reports specification test results as detailed in Section 4.4. The last array reports the knots of μ_2 .

Type of the precipitating event	Number of events	Effect of event of 5,000 shares (number of events)			Half-life of effect (s)		
		Lower	Estimate	Upper	Lower	Estimate	Upper
Cancelled bids	20,921	0.01	0.01	0.01	5.6	7.2	9.1
Bids at or below best bid	25,121	n/a	n/a	n/a	n/a	n/a	n/a
Bids between the quotes	5,292	0.00	0.00	0.00	5.1	6.4	7.9
Market buys that don’t move ask	5,885	0.07	0.10	0.15	799.61	1141.26	1592.55

Table 5 (continued)

Type of the precipitating event	Number of events	Effect of event of 5,000 shares (number of events)			Half-life of effect (s)		
		Lower	Estimate	Upper	Lower	Estimate	Upper
Market buys that do move ask	3,673	0.02	0.02	0.03	2.0	2.5	3.1
Cancelled asks	19,424	n/a	n/a	n/a	n/a	n/a	n/a
Asks at or above best ask	22,433	0.01	0.01	0.01	7.2	9.7	13.2
Asks between the quotes	5,281	0.05	0.06	0.07	11.6	14.7	18.7
Market sales that don't move the bid	5,114	0.15	0.16	0.17	8.0	8.8	9.8
Market sales that do move the bid	4,351	0.09	0.10	0.12	10.8	12.4	14.3
Autoregressive intensity effect	4,351	0.21	0.25	0.29	724	957	1264

Type of the precipitating event	Elasticity of exciting magnitude to a 1,000 increase in share volume		
	Lower	Estimate	Upper
Cancelled bids	0.03	0.04	0.05
Bids at or below best bid	−0.66	−0.51	−0.35
Bids between the quotes	−1.93	−1.61	−1.28
Market buys that don't move ask	−0.09	0.01	0.11
Market buys that do move ask	−0.02	0.00	0.03
Cancelled asks	−0.28	−0.18	−0.08
Asks at or above best ask	0.00	0.02	0.04
Asks between the quotes	0.07	0.08	0.09
Market sales that don't move the bid	0.00	0.01	0.01
Market sales that do move the bid	0.02	0.02	0.03
Autoregressive intensity effect	−0.05	−0.02	0.01

Specification testing			
Time-adjusted durations			
Mean (should be 1)			0.97
Var. (should be 1)			0.95
Robust Ljung–Box (15 lags) of squares <i>p</i> -value			0.00
Robust Ljung–Box (25 lags) of squares <i>p</i> -value			0.00
Test of excess dispersion <i>p</i> -value			0.87
Likelihood			10,765

Intraday spline			
Time	Intensity at node (events/h)		
	Lower	Estimate	Upper
08:00	4.5	6.1	8.2
08:30	8.1	9.8	12.0
09:00	2.3	3.3	4.6
12:00	#N/A	0.0	#N/A
15:00	2.1	3.0	4.2
16:30	3.1	4.8	7.4




Fig. 4. A bubble chart representation of aggressive limit order dynamics.




Fig. 5. A bubble chart representation of aggressive market order dynamics.

A likelihood ratio test rejected at 5 per cent the hypothesis that all four magnitudes were 0.2. Rather, they are all estimated to be lower. Therefore, viewed individually both the bid and the ask have less than a 20 per cent chance of replenishing reliably after a large market order. Jointly, they resiliently replenish the limit order book less than 40 per cent of the time. This is consistent with the dynamic illiquidity reported in Danielsson and Payne (2002) and the periodic liquidity crises following sharp price movements predicted by Coppejans et al. (2004), as well as the long memory reported in Degryse et al. (2005), whereby wide spreads sometimes narrow only in the medium or long term.

The estimated half-lives are all below 20 s. Moreover, a likelihood ratio test rejected at 5 per cent the hypothesis that all four half-lives were 20 s. Concluding that half-lives are all below 20 s, it follows that *if* the order book does replenish, it has more than a 50 (75) per

cent chance of doing so within 20 (40) s. This is quite fast: for example, it is too fast to be captured by 5 min sampling, and is substantially faster than typical half-lives considered in the calibration exercise in [Obizhaeva and Wang \(2005\)](#).

In a likelihood ratio test, it was not possible to reject at 10 per cent the hypothesis that all four effects had identical magnitudes of 0.17. We may therefore conclude that there is little evidence that the bid reacts any more or less than the ask, following either a large market purchase or a large market sale. Therefore, after an aggressive market sale, only half the time does the market's resilient response (where there is one) involve *replacing* bids thereby executed. The rest of the time, it is rather new price-improving asks that narrow spreads, which moreover may typically arrive faster (see [Fig. 4](#)). The same holds for a market buy. Thus, market orders often move the quotes: they have price impact.

Discussion. In the data we see large market orders that cause spreads to become suddenly unusually wide. Then limit orders can, for a given level of execution risk and delay, achieve unusually favorable prices. This attracts limit orders to the book, which narrow spreads: an effect present in the theoretical work of [Rosu \(2004\)](#) and [Foucault et al. \(2005\)](#). Limit orders at a given price tick are served in time priority. Hence they race to be early to the market at times of wide spreads. However there may be a winner's curse in this race if the initial market order resulted from asymmetric information. The results of this section suggest that the race to replace liquidity supply is brief, lasting only tens of seconds after the liquidity demand shock. It takes place equally fast and often on both sides of the book, though a priori we might have expected more hesitation to winner's curse effects on the side where the market order trades.

6.2. Other dynamics of aggressive limit orders

After what sequence of events is liquidity supply likely to increase on a limit order book? When are spreads likely to narrow? These are questions about the dynamics of aggressive bidding and asking. These dynamics exhibit significant buy-sell symmetry. As depicted in [Fig. 6](#), the background intradaily splines are of very similar shapes and sizes. Furthermore, all the exciting effects on aggressive bids are reflected in similar exciting effects on asks, of similar magnitudes and half-lives, as depicted in [Fig. 4](#). Full results are reported in [Tables 2 and 3](#).

6.2.1. Aggressive limit orders following cancellations

A cancellation has a 2–3 per cent chance of precipitating a price-improving limit bid, and a similar chance of precipitating a price-improving limit ask. Where this concerns the replacement of a cancelled order with one on the same side of the book, the half life of this effect is a little as 1–2 s. This supports a conjecture that traders sometimes cancel a stale limit order and replace it rapidly with one at a more pertinent price. This effect has a negative elasticity with respect to share volume. Thus, if cancelled, an order for 1,000 shares is roughly twice as likely to be replaced as one for 5,000 shares. Where a cancellation precipitates a limit order on the *other* side of the book, this effect is relatively slow, with a half life of around 80–90 s.

6.2.2. Aggressive limit orders after aggressive limit orders

Aggressive, price-improving limit orders have a 6–7 per cent chance of precipitating further similar price improvements, with a half life of approximately one minute. However, orders which




Fig. 6. The estimated intraday intensity splines (events/h).




Fig. 7. QQ plots of time-adjusted durations, i.e., residuals, against an exponential distribution of parameter 1.

enhance liquidity supply without moving prices are much less likely to do so. This is consistent with a conjecture that a price-improving limit bid (ask) carries more information than one that does not move prices, and therefore is more likely to be followed by price rises (falls).

In addition, with a much slower half life of 13–17 min, aggressive, price-improving limit orders have a 16–22 percent chance of precipitating further similar price improvements. This second self-exciting effect, specified in (13), suggests positive medium-term autocorrelation in limit order submission intensities.

6.3. The dynamics of aggressive market orders

When are spreads likely to widen? Can this be predicted? These questions concern the dynamics of aggressive market orders. Like those of aggressive limit orders, their observed dynamics show great buy-sell symmetry. As indicated in Fig. 6, the underlying splines are very similar to each other, although fundamentally different from those reported for limit orders. They fall to a minimum around noon, before rising again in the afternoon. Despite some isolated asymmetries, dynamic effects on each side of the market mirror each other as noted in Section 6.2 for aggressive limit orders. All the effects now discussed are reported in Tables 4 and 5.

6.3.1. Aggressive market orders following limit orders

A price-improving limit bid precipitates aggressive market purchases 7 per cent of the time. This effect has a half life of 19 s. If it is large, for example comprising more than 5,000 shares, it has a negligible chance of triggering an aggressive market sale causing its execution. However, if it is sized at 1,000 shares, it has a 5 per cent chance of being hit. A natural interpretation of this is that 5 per cent of the time, sellers are watching and waiting for a small improvement in the best bid before submitting a market order. However, if this improvement involves a large volume of shares, such as 5,000, then in a way reminiscent of Easley and O'Hara (1992), it signals future rises in price, which deter the waiting sellers from acting. These effects are closely mirrored on the other side of the market. They are reported in Fig. 5 (see also Tables 4 and 5).

6.3.2. The effect on market orders of market orders themselves

The largest exciting effect on market orders is a self-exciting effect at a timescale of 15–30 min (the half life of this effect is estimated at 25 min for aggressive market purchases, and 16 min for aggressive market sales). Each aggressive market order has a 25–30 per cent chance of precipitating repeats over this time frame. This therefore describes medium-term positive autocorrelation in market order frequencies.

Further, more moderate short-run dynamics are captured by the model. These are cointegration effects, where market buys (sales) precipitate renewed market buying (selling). The effect on aggressive market purchases of unaggressive market purchases is large and rapid, with a magnitude of 13 per cent and a half life of 8 s. Thus, if a market order clears some liquidity at the best ask, but leaves some limit orders unfilled, it is often followed by another order that fills the remaining asks at the best ask. By contrast, the short-run effect (as opposed to the aforementioned medium-term effects) on aggressive market purchases of prior aggressive purchasing is slower and smaller. Traders are less likely to copy a market order if they must trade at still less advantageous prices. The volume of the market order is a relatively insignificant factor for these effects (except indirectly, in that it determines whether a market sale fills all, or just some, of the limit bids at the best bid). All this is recorded in Fig. 5 (see also Tables 4 and 5).

7. Conclusion

To quantify resiliency, a dynamic model of limit order book activity is proposed, and is partially estimated for Barclays shares on the LSE's electronic limit order book. The model takes the form of a multivariate point process in continuous time with an adapted intensity. It is reasonably well specified and throws up a number of striking dynamic effects. Aggressive market orders are submitted above all in the wake of similar but unaggressive market orders. They are also sometimes submitted to hit small, price-improving limit orders on the other side of the book. Aggressive limit orders are typically placed to replenish liquidity following aggressive market orders, but can also be submitted to update a stale limit order. There is some evidence of positive medium-term autocorrelation in event intensities. Finally, the resiliency of the limit order book is quantified in three respects: its magnitude, its delay and its trade direction. Resilient replenishment follows a shock less than 40 per cent of the time, and when it does follow, it is equally likely to be at the bid as at the ask. It is quite fast, with a half-life of under 20 s.

Appendix A. Events falling within the same second

Forty per cent of the events observed on the SETS limit order book for Barclays Plc over the studied period were marked with the same time stamp as another event. This was purely due to the rounding of time stamps to the nearest second. However, the statistical framework of this paper is one of simple point processes where contemporaneous events never occur.

Three approaches exist to resolve these problems. The choice among them will have greatest impact on the estimation of exciting effects that attenuate on timescales near to one second, such as the tendency, documented in Section 6.2.1, of traders to resubmit cancelled orders in the next 1–2 s. It is likely that the estimated magnitude and rate of decay of these effects would differ from the estimates herein reported if the data were more precisely timed.

On the first approach, the data is thinned so that only one event at each time stamp is retained. Given the prevalence of this problem, this approach seems inappropriate here. Second, a uniform random variable with support equal to $[-0.5, 0.5]$ (in s) can be added to the time stamp of each event. However, this approach sometimes imposes an incorrect ordering on the order book events. In this setting this is unhelpful, since the order of events is very important. These two approaches ensure that the data is a simple point process.

The paper adopts a third approach: the data is not altered prior to maximizing the likelihood. It therefore contains contemporaneously timed events. Since in the parametrization $W(0) = 0$, these events have no effect one on the other. The diagnostic testing procedure of Section 4.4 now yields many durations of length zero, as well as a corresponding shortage of durations near zero (at levels < 0.05). The following procedure was therefore applied: any term of the sequence in (16) which was less than 0.05 was discarded, and 0.05 was subtracted from the remaining terms. Recall that if E is an exponentially distributed random variable of parameter 1, then $(E - 0.05)$, truncated to $\mathbb{R}^+$, is also.⁹ The resulting sequence should therefore also be i.i.d. exponential. This

⁹Redraw if $E < 0.05$.

procedure cuts out mis-shapings in the empirical distribution near zero. As it only disregards a few terms, it retains much of any original autocorrelation.

Appendix B. Proof of Proposition 4.1

The proposition may be derived thus:

$$G_m(t-s) := \lambda(t|\mathcal{F}_s, \mathcal{E}_m) - \lambda(t|\mathcal{F}_s) \quad (17)$$

$$= E \left\{ \mu(t-s) + \int_{[s,t)} W(t-u) dN_u | \mathcal{F}_s, \mathcal{E}_m \right\} \quad (18)$$

$$- \left\{ \mu(t-s) + \int_{[s,t)} W(t-u) dN_u | \mathcal{F}_s \right\} \quad (19)$$

$$= W_m(t-s) + \int_s^t W(t-u) \{ \lambda(u|\mathcal{F}_s, \mathcal{E}_m) - \lambda(u|\mathcal{F}_s) \} du \quad (20)$$

$$= W_m(t-s) + \int_s^t W(t-u) G_m(u-s) du, \quad (21)$$

where $W_m(t-s)$ is the m th column of W and appears in (20) since with probability 1

$$dN_s | \mathcal{F}_s, \mathcal{E}_m \equiv (0, 0, \dots, 1, \dots, 0)', \quad (22)$$

where 1 appears in the m th entry. From (21) it follows that G is a transformation of W . Stacking (21) over m , and integrating it over t , setting s to zero, and changing the order of the double integration,

$$\int_0^\infty G(t) dt = \int_0^\infty W(t) dt + \int_0^\infty \left\{ \int_u^\infty W(t-u) dt \right\} G(u) du. \quad (23)$$

Rearranging,

$$\int_0^\infty G(t) dt = \left\{ Id - \int_0^\infty W(t) dt \right\}^{-1} \int_0^\infty W(t) dt. \quad (24)$$

References

- Bauwens, L., Hautsch, N., 2004. Dynamic latent factor models for intensity processes. Working paper, CORE.
- Biais, B., Hillion, P., Spatt, C., 1995. An empirical analysis of the limit order book and the order flow in the Paris Bourse. *Journal of Finance* 50, 1655–1689.
- Bisière, C., Kamionka, T., 2000. Timing of orders, orders aggressiveness and the order book at the Paris Bourse. *Annales d'Economie et de Statistique* 60, 43–72.
- Bowsher, C., 2005. Modelling security market events in continuous time: intensity based, multivariate point process models. Nuffield College Economics Discussion Paper No. 2005-W26.
- Coppejans, M., Domowitz, I., Madhavan, A., 2004. Resiliency in an automated auction. Working paper, Duke University.
- Danielsson, J., Payne, R., 2002. Liquidity determination in an order driven market. Working paper, London School of Economics.
- Degryse, H., de Jong, F., van Ravenswaaij, M., Wuyts, G., 2005. Aggressive orders and the resiliency of a limit order market. *Review of Finance* 9, 201–242.
- Doornik, J., 2001. An Object-Oriented Matrix Programming language. Timberlake Consultants Ltd., London.
- Dufour, A., Engle, R.F., 2000. Time and the price impact of a trade. *Journal of Finance* 55, 2467–2498.

- Easley, D., O'Hara, M., 1992. Time and the process of security price adjustment. *Journal of Finance* 47, 577–605.
- Engle, R.F., 2000. The econometrics of ultra-high-frequency data. *Econometrica* 68, 1–22.
- Engle, R.F., Lange, J., 2001. Predicting VNET; a model of the dynamics of market depth. *Journal of Financial Markets* 4, 113–142.
- Engle, R.F., Lunde, A., 2003. Trades and quotes: a bivariate point process. *Journal of Financial Econometrics* 1, 159–188.
- Engle, R.F., Russell, J.R., 1998. Autoregressive Conditional Duration: a new model for irregularly spaced transaction data. *Econometrica* 66, 1127–1162.
- Foucault, T., Kadan, O., Kandel, E., 2005. Limit order book as a market for liquidity. *Review of Financial Studies* 18, 1171–1217.
- Hasbrouck, J., 1999. Trading fast and slow: security market events in real time. Working paper, Stern School of Business, New York University.
- Hawkes, A.G., 1971. Spectra of some self-exciting and mutually exciting point processes. *Biometrika* 58, 83–90.
- Heinen, A., Rengifo, E., 2003. Comovements in trading activity: a multivariate autoregressive model of time series count data using copulas. Working paper, Catholic University of Louvain.
- Hollifield, B., Miller, R.A., Sandås, P., 2004. Empirical analysis of limit order markets. *Review of Economic Studies* 71, 1027–1063.
- Kyle, A., 1985. Continuous auctions and insider trading. *Econometrica* 53, 1315–1336.
- Obizhaeva, A., Wang, J., 2005. Optimal trading strategy and supply/demand dynamics. Working paper, MIT.
- Ogata, Y., 1978. The asymptotic behaviour of maximum likelihood estimators for stationary point processes. *Annals of the Institute of Statistical Mathematics* 30, 243–261.
- Rosu, I., 2004. A dynamic model of the limit order book. Ph.D. Thesis, MIT.
- Russell, J.R., 1999. Econometric modeling of multivariate irregularly-spaced high-frequency data. Mimeo, University of Chicago, Graduate School of Business.