

PAPER: Interdisciplinary statistical mechanics

Endogenous liquidity crises

Antoine Fosset^{1,2}, Jean-Philippe Bouchaud^{2,3}
and Michael Benzaquen^{1,2,3,4}

¹ Ladhyx, UMR CNRS 7646, Ecole Polytechnique, 91128 Palaiseau Cedex, France

² Chair of Econophysics & Complex Systems, Ecole Polytechnique, 91128 Palaiseau Cedex, France

³ Capital Fund Management, 23-25, Rue de l'Université 75007 Paris, France
E-mail: michael.benzaquen@polytechnique.edu

Received 3 December 2019

Accepted for publication 19 February 2020

Published 2 June 2020



Online at stacks.iop.org/JSTAT/2020/063401
<https://doi.org/10.1088/1742-5468/ab7c64>

Abstract. Empirical data reveals that the liquidity flow into the order book (limit orders, cancellations and market orders) is influenced by past price changes. In particular, we show that liquidity tends to decrease with the amplitude of past volatility and price trends. Such a feedback mechanism in turn increases the volatility, possibly leading to a liquidity crisis. Accounting for such effects within a stylized order book model, we demonstrate numerically that there exists a second order phase transition between a stable regime for weak feedback to an unstable regime for strong feedback, in which liquidity crises arise with probability one. We characterize the critical exponents, which appear to belong to a new universality class. We then propose a simpler model for spread dynamics that maps onto a linear Hawkes process which also exhibits liquidity crises. If relevant for the real markets, such a phase transition scenario requires the system to sit below, but very close to the instability threshold (self-organised criticality), or else that the feedback intensity is itself time dependent and occasionally visits the unstable region. An alternative scenario is provided by a class of non-linear Hawkes process that show occasional 'activated' liquidity crises, without having to be poised at the edge of instability.

Keywords: market microstructure, market impact, models of financial markets, quantitative finance, agent-based models

⁴ Author to whom any correspondence should be addressed.

Contents

1. Introduction	2
2. Destabilizing feedback effects: empirical analysis	4
2.1. Average event rates	4
2.2. A generalized Q-Hawkes model	5
2.3. Calibration strategy	6
2.4. Results	8
3. An agent-based model for liquidity crises	10
3.1. The Santa Fe model with feedback	10
3.2. Numerical simulations	11
3.3. Phase transition and finite size scaling	11
3.4. Discussion	14
4. A state-dependent Hawkes model for spread dynamics	15
4.1. A simple model	15
4.2. The stable regime	16
4.3. Linear spread growth	17
4.4. The explosive regime	18
4.5. A stabilizing mechanism	18
5. Non-linear Hawkes models and metastability	19
5.1. A model with quadratic feedback	19
5.2. A continuous time description	20
6. Conclusion	22
Acknowledgments	23
Appendix A. Empirical data	23
Appendix B. Finite size scaling method	24
Appendix C. More on the linear spread model	25
Appendix D. A model with price feedback on the spread	26
References	28

1. Introduction

Why are financial markets so prone to liquidity crises and crashes? It is now well established that market volatility is too high to be explained by fluctuations of fundamental value. In particular, a large fraction of large price jumps (say, $4\text{-}\sigma$ events at the 1 min

time scale [1], or major daily moves [2, 3]) cannot be explained by significant news. These jumps seem to be rather the result of endogenous feedback loops that lead to liquidity seizures. The memory of most spectacular ones is still vivid, such as the infamous S & P500 flash crash of May 6, 2010 [4], or the Treasury bond flash crash of October 15, 2014.

These recent events have triggered a large amount of controversy, in particular in the general press, pointing fingers at electronic markets and high frequency traders. However, financial markets have always been unstable. For example on May 28, 1962, the US stock market suffered a flash crash of severity similar to the that of May 6, 2010 [5]. This happened with good old market makers and, obviously, no HFT. Upon closer scrutiny one finds that the frequency of large price moves is remarkably stable over time, once rescaled by volatility, see e.g. [6].⁵ It is found to decay as a power-law of the amplitude of the price move—the so-called ‘inverse cubic-law’ [9].

A plausible general scenario is that of destabilising feedback loops resulting in liquidity breakdown. Consider for example the classic Glosten–Milgrom model [10] relating liquidity to adverse selection. When liquidity providers believe that the quantity of information revealed by trades exceeds some threshold, there is no longer any value of the bid–ask spread that allows them to break even—liquidity vanishes (see e.g. [6], chapter 16). Whether real or perceived, the risk of adverse selection is detrimental to liquidity. This creates a clear amplification channel that can lead to liquidity crises⁶. To illustrate this point, imagine that the price has recently experienced a burst of volatility. This creates anxiety for liquidity providers, who fear that some information about the future price, unbeknownst to them, is the underlying reason for the recent price changes. The consequence is an increased reluctance to provide liquidity: such liquidity providers become more likely to cancel their existing limit orders and less likely to refill the limit order book with new limit orders. Less liquidity is likely to amplify the future price moves, thereby creating an unstable feedback loop which might result in a runaway trajectory.

The present paper attempts to capture such feedback effects both empirically and through stylized models for the dynamics of order books. In section 2, we empirically show that event rates in the limit order book are indeed affected by past volatility. Using tick-by-tick order book data from the EURO STOXX contract, we calibrate a generalisation of the self-exciting Hawkes processes [12, 13]—nowadays commonly used in finance but initially introduced to reproduce seismic activity. In particular, we show that market orders and cancellations tend to increase when recent price changes are large, in turn diminishing the available liquidity, much as argued above. We then turn to the modelling part. In section 3, to study the aggregate outcomes of such feedback in a minimal setting, we consider an extended version of the *Santa Fe order book model* [14–16].⁷

⁵Note however that the frequency of co-jumps has escalated in the past decades bearing witness of a significant increase of the level of synchronization of large price movements across assets, see [7, 8].

⁶For an alternative/complementary view on liquidity crises see also [11], where we have shown that liquidity dry outs may also be understood as the result of lag effects on latent liquidity revealing.

⁷The *Santa Fe model* stands among the first *zero intelligence* order book models reproducing some statistical properties, such as the mean bid–ask spread and mean volume profiles near the best quotes. Note however that the model is too simple to account for volatility levels, volume profile queues far from the best, or to solve the diffusivity puzzle [6].

The original model consists in a collection of N queues that evolve with constant additive limit order and market order arrival probability rates, and a constant cancellation rate per existing limit order. We introduce, in a minimal fashion, the effect of interest to us by letting past prices changes feed back into the event rates. Our numerical results strongly suggest the existence of a genuine phase transition from a stable regime to an unstable regime in which liquidity crises arise, as feedback intensity is increased. We perform a finite size scaling and determine the corresponding critical exponents. In section 4 we present a simpler model, more prone to analytical treatment, setting aside the dynamics of the order book, and restricting our attention to the dynamics of the spread. We argue that phase transition scenario can be rather generic, but require markets to sit very close to the instability threshold. Another possibility is that the feedback parameter itself is time dependent and occasionally visits the unstable phase. In section 5, we explore an alternative scenario (activation), in which occasional liquidity crises arise without having to be poised at the edge of instability or having a time dependent feedback parameter. In section 6, we conclude.

2. Destabilizing feedback effects: empirical analysis

In this section, we provide an empirical analysis of feedback effects within order book dynamics. Consider an electronic market with three event types only: limit order deposition (LO), limit order cancellation (C) and market orders (MO).

It is already well documented that these events strongly interact with one another. A very useful framework to describe these interactions is provided by Hawkes self-exciting point processes [17], which have already been applied to order book events in [12, 13, 18–20]. Here we want to extend these studies to account not only for activity feedback but for *price feedback* as well, in the spirit of the quadratic Hawkes (Q-Hawkes) model of Blanc *et al* [21].⁸

2.1. Average event rates

Figure 1 displays the average order size, frequency of events and order rate (= order size \times frequency) as function of the rescaled volume at best for each event type, on the EURO STOXX contract between 2016/09/12 and 2017/04/28. The volume at best has been rescaled by the average limit order size in the same time bin, in order to eliminate intra-day seasonalities. In terms of time scales, we find that for EURO STOXX the average time between two events is $\tau_e = 0.03\text{s}$, whereas the average time between two price changes is $\tau_p = 7\text{s}$. In addition to the expected bid–ask symmetry, figure 1(c) reveals that the total rate of cancellations and market orders are roughly proportional to the size of the queue, whereas limit order posting does not show any appreciable dependence on the volume at best. This observation motivates the specification of the Q-Hawkes model that we calibrate below.

⁸See also [22] for a recent analysis of the complex interplay between intraday volatility spikes and negative stock market jumps.




Figure 1. Average order size, frequency of events and order rate (= order size \times frequency) as function of volume at best rescaled by the average limit order size, on the EURO STOXX contract between 2016/09/12 and 2017/04/28.

2.2. A generalized Q-Hawkes model

For the sake of simplicity, we focus on events (LO, C, MO) at the best quotes only, bid (b) and ask (a) (we do not distinguish between placing limit orders at the current best or inside the spread). We therefore introduce the following six-dimensional process that counts all such events:

$$N_t = \left(N_t^{C,b}, N_t^{LO,b}, N_t^{MO,b}, N_t^{MO,a}, N_t^{LO,a}, N_t^{C,a} \right).$$

We further assume that the time dependent intensities λ_t of these six processes follow the following Q-Hawkes dynamics:

$$\lambda_t = Q_t \left(\alpha_0 + \int_0^t \phi(t-s) dN_s + \int_0^t L(t-s) dP_s + \int_0^t \int_0^t K(t-s, t-u) dP_s dP_u \right)_+, \quad (1)$$

where dP_s is the price change at time s in tick units, $Q_t = \text{diag} (Q_t^b, 1, Q_t^b, Q_t^a, 1, Q_t^a)$ with $Q_t^{b/a}$ the volume at the best bid/ask in units of average limit order size. Equation (1) assumes that cancellations and market orders are multiplicative while limit order event types are additive, as mentioned above. Note that all kernels L, K are six-dimensional vectors and ϕ a six-dimensional matrix.

The first term on the RHS of equation (1) accounts for a stationary exogenous intensity α_0 and the second is the classical Hawkes kernel accounting for event interactions⁹. The third and fourth terms were introduced in [21] and are new in the context of limit order book modelling. The third term is a linear feedback term from past price changes, modelling the fact that up or down price moves directly impact the rate of cancellations, market orders and limit orders. The fourth is a quadratic feedback term on the rate of order book events, which does not depend on the sign of past price changes. In [21], it was proposed to write the kernel as $K(t-s, t-u) = K_d \psi(t-s) \delta(s-u) + K_1 Z(t-s) Z(t-u)$, with:

⁹Whereas the Hawkes contribution is not the focus of the present paper, including it is essential to obtain a reasonable explanatory power (see below).

- a diagonal (in time) contribution $\psi(t-s)$ with weights \mathbf{K}_d , which represents the feedback of past volatility on current activity, since it can be written in terms of:

$$\Sigma^2(t) = \int_0^t \psi(t-s)(dP_s)^2,$$

- a rank-one contribution, with weights \mathbf{K}_1 , which amounts to coupling the square of past trends, as measured by:

$$R(t) = \int_0^t Z(t-s)dP_s.$$

This is the so-called Zumbach effect: past trends, independently of their sign, lead to an increase in future activity.

In the following we will choose $\psi(s) = Z(s) = e^{-\beta s}$.¹⁰ One of the main empirical findings of the present study is that these two effects (volatility feedback and Zumbach effect) are indeed present and large, and capture the destabilizing feedback loop

trends & volatility \rightarrow lower liquidity \rightarrow more trends & volatility

as surmised in the introduction.

2.3. Calibration strategy

The Hawkes contribution $\alpha_H = \int \phi(s)ds$ has been studied in several papers in the past (see e.g. [23]) and is now rather well understood. We first calibrate a Hawkes process without the price feedback term, i.e. setting \mathbf{L} and \mathbf{K} to zero in equation (1). We use the non-parametric technique introduced in [12, 24], expecting bid/ask symmetry. This means that the coefficients $\mathbf{Q}\tilde{\alpha}_0$ only depend on the type of events (and not their ‘sign’), and that the matrix $\mathbf{Q}\alpha_H$ has a block-symmetry: the couplings of $b \rightarrow b$ are equal to those of $a \rightarrow a$, and that $b \rightarrow a$ is equivalent to $a \rightarrow b$. Our results are qualitatively similar to those reported in the literature [12, 13, 23, 24]. The matrix structure of the norm of the Hawkes feedback kernel is shown in figure 2 for the EURO STOXX contract.

Reintroducing the quadratic coupling term \mathbf{K} leads to a much more complicated structure for the nonparametric calibration problem (see [21]), in particular in the present multidimensional setting. We have not yet been able to implement satisfactorily such a scheme, so we devised a simplified protocol to get some partial information on the structure of the price feedback terms. The idea is to capture the effect of local trends on the liquidity of the order book. Hence we define the net flux of orders at the bid $x = b$ or at the ask $x = a$ as:

$$dJ_t^x := dN_t^{\text{LO},x} - dN_t^{\text{MO},x} - dN_t^{\text{C},x}.$$

From this we define the total flux and the signed flux as:

$$dJ_t^{b+a} = dJ_t^a + dJ_t^b, \quad dJ_t^{b-a} = dJ_t^b - dJ_t^a.$$

¹⁰We assume for the sake of simplicity that the memory of trends is the same as that of volatility.




Figure 2. Fitting parameters resulting from the Hawkes non-parametric calibration on the EURO STOXX futures contract between 2016/09/12 and 2017/04/28. $Q\tilde{\alpha}_0$ stands for the exogenous intensity from the non-parametric Hawkes fit.

We also introduce the forward realized flux and the forward Hawkes flux on time scale β'^{-1} :

$$F_{\beta'}^x(t) = \int_t^{+\infty} e^{-\beta'(s-t)} dI_s^x, H_{\beta'}^x(t) = \int_t^{+\infty} e^{-\beta'(s-t)} \lambda_s^{H,x} ds,$$

where $x = (b + a, b - a)$ and $\lambda_s^{H,x}$ is the expected future activity, as predicted by the Hawkes contribution¹¹. In the absence of other feedback mechanisms, one would expect any conditional expectation of $F_{\beta'}^x(t)$ should simply be $H_{\beta'}^x(t)$.

This is what we test now, by considering two conditioning variables suggested by the Q-Hawkes formalism, namely past trends and past realised volatility, as measured by the following exponential moving averages:

$$R_{\beta}(t) := \underbrace{\int_0^t e^{-\beta(t-s)} dP_s}_{\text{past trend}}, \quad \Sigma_{\beta}^2(t) := \underbrace{\int_0^t e^{-2\beta(t-s)} (dP_s)^2}_{\text{past volatility}}.$$

By symmetry, we expect that the conditional expectations of $F_{\beta'}^{b-a}(t)$ and $F_{\beta'}^{b+a}(t)$ write:

$$\mathbb{E}_c[\beta' F_{\beta'}^{b+a}(t) | R, \Sigma, H] = C_0 + 2\beta C_1 R_{\beta}^2(t) + 2\beta C_2 \Sigma_{\beta}^2(t) + \beta' H_{\beta'}^{b+a}(t) \tag{2}$$

$$\mathbb{E}_c[\beta' F_{\beta'}^{b-a}(t) | R, \Sigma, H] = \sqrt{\beta} C_3 R_{\beta} + \beta' H_{\beta'}^{b-a}(t), \tag{3}$$

i.e. the asymmetric part of the liquidity flow depends on the sign of the past trend, whereas the symmetric part of the flow depends both on the past volatility and on the

¹¹ More explicitly, $\lambda_t^{H,x} := \lambda_t^{H,LO,x} - \lambda_t^{H,MO,x} - \lambda_t^{H,C,x}$, where λ^H is the Hawkes intensity process calibrated above. In order to speed up the computation of $F_{\beta'}^x$, we approximate the non-parametric Hawkes kernels by sums of exponentials.

J. Stat. Mech. (2020) 063401




Figure 3. Correlation between the trend and the total liquidity flux (left) and signed liquidity flux (right) in the plane β, β' for the EURO STOXX futures contract between 2016/09/12 and 2017/04/28. Note that the color scale is not the same in the left and in the right graph: the directional effect is weaker than the impact on the total (unsigned) liquidity.

past trend squared (i.e. the Zumbach effect). C_0, C_1, C_2 and C_3 are numerical constants. The normalisation factor β comes from the fact that $R_{\beta} \sim \beta^{-1/2}$ and $\Sigma_{\beta}^2 \sim \beta^{-1}$. Note that the regression coefficients in front of the calibrated Hawkes contribution are fixed to unity, as they should be for consistency.

2.4. Results

We determine β and β' by looking at the maximum absolute correlations of $F_{\beta'}^{b+a}$ with R_{β}^2 , see figure 3 and appendix A. We find $\beta = 0.001$ and $\beta' = 0.02$, corresponding to a negative correlation ≈ -0.3 , indicating that *trends indeed reduce liquidity*. Note that the correlations between $F_{\beta'}^{b-a}$ with R_{β} are one order of magnitude smaller, and in fact change sign depending on the time scales: the short time response to an up trend is adding liquidity at the ask, but the long time response is in fact removal of liquidity at the ask. This could reflect the behaviour of different actors in the market (high frequency traders/market makers vs longer term traders).

Fixing $\beta = 10^{-3}$ (i.e. trends measured over 1000 s, similar to the time scale found in [21]) and $\beta' = 0.02$ (market response over the next 50 s), we find the regression coefficients C_i given in table 1 for several futures contracts, again using the period between 12/09/2016 and 28/04/2017. The quality of the regressions in the case of the EURO STOXX is illustrated in figure 4 (similar plots are obtained for the BUND, BOBL and SCHATZ, not shown). We see that both the trend (Zumbach) effect, parameterised by C_1 and the volatility effect, parameterised by C_2 , are both important to reproduce

Table 1. Values of the coefficients C_0, C_1, C_2 for the symmetric part of the liquidity flow and C_3 for the antisymmetric part, as defined in equations (2) and (3). We fixed $\beta = 10^{-3}$ and $\beta' = 2 \times 10^{-2}$.

	$C_0(10^{-2})$	C_1	C_2	C_3
EUROSTOXX	78	-8.9	-6.7	-0.03
BUND	72	-1.7	-2.8	0.16
BOBL	13	-4.0	-0.29	0.29
SCHATZ	0.42	-2.5	0.001	0.50

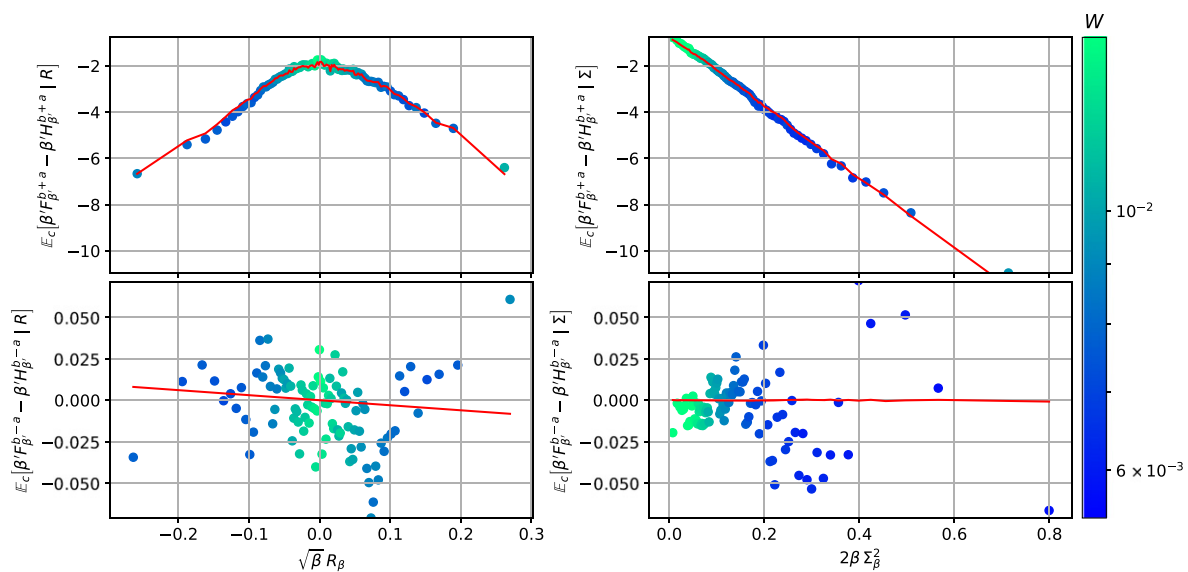


Figure 4. Regressions of the incoming flux with the trend and the volatility on the EURO STOXX futures contract between 2016/09/12 and 2017/04/28. The red curves correspond to the in-sample prediction of the linear regression. Each empirical point has a weight computed from the fraction of time spent in the corresponding state.

the future liquidity flow. The directional effect, measured by C_3 , is much weaker, as indeed suggested by figure 3, so we will neglect it in the following.

The conclusions of this calibration exercise are that:

- large recent price trend and volatility indeed tend to increase the rate of market orders and cancellations and lead to a decrease in liquidity. This is the main take-away message of this section;
- the quadratic feedback terms \mathbf{K} in equation (1) is the dominant effect; in the following section we will thus neglect the linear term and set $\mathbf{L} = 0$.

The question is now whether such a quadratic feedback is enough to induce self-generated liquidity crises, i.e. an unstable feedback loop that wipes out all the volume in the order book and leads to crash. We explore this question in the next sections by analyzing different models, starting from the zero-intelligence order book model

proposed by the Farmer and collaborators [14–16], and ending by much simpler models amenable to some analytical predictions.

3. An agent-based model for liquidity crises

The so-called *Santa Fe model* [14–16] stands among the first purely stochastic order book models, where zero-intelligence agents place their orders at random (see also [18, 25]). It was shown that this model is able to reproduce some empirical properties of order books, such as the mean bid–ask spread and mean volume profiles near the best quotes. However the model fails to account for the empirical relation between spread and volatility (see [26, 27], and [6], chapter 8); in fact prices are found to be strongly mean reverting, partly because of the absence of long-range correlations in the flow of market orders in the model—see the detailed discussion of this point in [6, 28, 29].

In spite of these shortcomings, the Santa Fe model is an interesting starting point for modelling order book dynamics. It consists in a collection of queues that evolve with constant additive limit and market order arrival Poisson rates, and a constant cancellation rate per existing limit order. Note that while real data is not fully consistent with additive depositions and multiplicative cancellations (see figure 1), this simplifying hypothesis allows for easier analytical treatment, and leads to a well defined steady state order book where queues are neither empty nor of infinite size.

Here, we present an extension of the Santa Fe model where the feedback of past price changes on event rates is taken into account. As suggested by the empirical results of the previous section, we only retain, for simplicity, the quadratic feedback term on cancellations, neglecting all others. We also keep the initial Santa Fe specification of an additive (rather than multiplicative) rate for market orders. Numerical simulations suggest that this brings no qualitative changes to our main conclusions, which are as follows.

- (a) There exists a critical value of the feedback parameter α_K such that for $\alpha_K < \alpha^*$, an infinite size order book *never empties*, while for $\alpha_K > \alpha^*$ such infinite size order book *empties with probability 1*.
- (b) The transition appears to be of second order nature, which means that as the transition point is approached some scaling behaviour is observed. For example, the average time $\bar{\tau}$ needed for the liquidity crisis to appear in an infinite order book diverges as $(\alpha_K - \alpha^*)^{-\zeta}$ with $\zeta \approx 3$ when $\alpha_K \downarrow \alpha^*$. For a book of finite size N , this time is always finite, but diverges as N^η with $\eta \approx 3$ when $\alpha_K = \alpha^*$.

3.1. The Santa Fe model with feedback

Consider a grid of prices with unit tick size, with all orders of unit size¹². This grid is divided into three parts: the bid side $\mathcal{B}_t = \{p \leq b_t\}$, the ask side $\mathcal{A}_t = \{p \geq a_t\}$ and the spread $\mathcal{S}_t = \{b_t < p < a_t\}$ where b_t and a_t respectively denote the best bid and the best

¹²One could introduce a distribution of order size at the expense of extra complexity. We expect that if this distribution is broad enough, the character of the phase transition could change.

ask. Market orders can only fall at the best bid and best ask; they do so with total rate 2μ , with probability $1/2$ to fall on the bid and $1/2$ to fall on the ask.

Bid limit orders fall uniformly with rate λ per tick size in $\mathcal{B}_t^+ = \{p \leq \min(b_t + 1, a_t - 1)\}$ and ask limit orders uniformly with the same rate λ in $\mathcal{A}_t^+ = \{p \geq \max(a_t - 1, b_t + 1)\}$. Orders cannot be placed inside the spread at a distance higher than one tick of the best prices.

Cancellations occur with a rate ν_t per outstanding limit order, which means that the probability that a given queue loses one order is proportional to the size of the queue. We assume that ν_t is given by a Q-Hawkes process of the type we considered in the previous section, where we retain only the Zumbach term, i.e.

$$\nu_t = \nu_0 + \alpha_K \left(\int_0^t \sqrt{2\beta} e^{-\beta(t-s)} dP_s \right)^2. \quad (4)$$

The case $\alpha_K = 0$ recovers the Santa Fe specification. Note that the dynamics of the different price levels are independent from one another, but described by the same parameters λ for the deposition rate and ν_t for the cancellation rate. In other words, the speed-up of cancellations when the price is trending affects *all* price levels.

3.2. Numerical simulations

To simulate the model, we take a price grid of size N ticks, and as initial condition, the equilibrium order book provided by the Santa Fe model with $\alpha_K = 0$. Then, to make the system evolve one can notice that, conditioned to the past, the system follows a multidimensional, non-homogeneous Poisson process, which is well known and easy to implement. Furthermore, for computing the integral $\int_0^t e^{-\beta(t-s)} dP_s = \sum_{T_n \leq t} e^{-\beta(t-T_n)} \Delta P_{T_n}$ efficiently, we use the usual recursive formula to speed up the algorithm, see [30].

Figure 5 displays typical results in the stable phase. Note that at some point the spread opens and triggers a cascade of cancellations that empties the order book. At some point in time denoted τ_c a liquidity crisis arises, that is here defined as the first time one side of the order book is completely empty.

3.3. Phase transition and finite size scaling

Exploring the parameter space (α_K, β) reveals that for $\alpha_K \gtrsim \alpha_m(\beta)$ liquidity crises arise with high probability. Figure 6 displays the crisis probability, defined as $\mathbb{P}[\tau_c \leq T]$, as function of α_K and β for $T\nu_0 = 200$ and $N = 280$. As expected, large feedback intensities α_K lead to unstable markets. The crossover value $\alpha_m(\beta)$ decreases as β increases, i.e. when the time scale over which trends are considered as dangerous by liquidity providers gets shorter. As expected, longer integration timescales β^{-1} lead to more stable order book, or in other terms, longer memory is a stabilising factor.

Although suggestive, figure 6 cannot be used to conclude on the existence of a true phase transition in the model, between a phase where liquidity crises never happen from a phase where liquidity crises always happen, provided one waits long enough. Mathematically, the question is about the behaviour of $\mathbb{P}[\tau_c \leq T]$ in the double limit

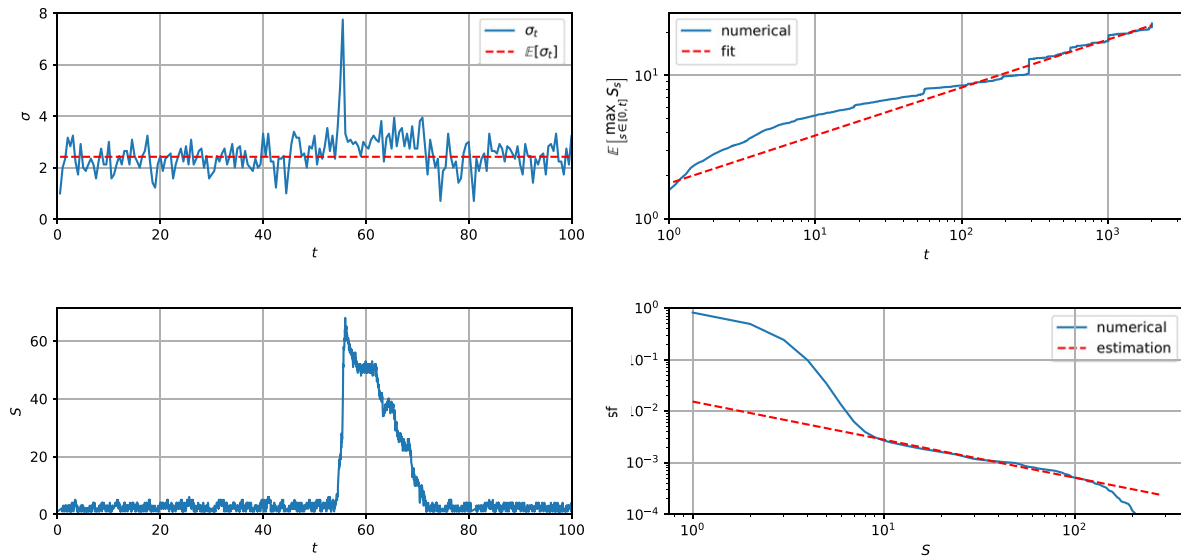


Figure 5. Properties of trajectories close to, but below the instability transition ($\lambda = 10$, $\nu_0 = 0.5$ and $\mu = 20$, and $\alpha_K = 0.2$, $\beta = 1$, $T = 2000$, $N = 1000$). (a) One trajectory of the volatility exhibiting a cluster of high volatility. (b) Average maximum of the spread S as a function of time and its fit by a power-law $t^{1/\eta}$ with $\eta = 3$ for $t \geq 100$. (c) The spread trajectory corresponding to (a). (d) The spread survival function (sf), also called complementary cumulative probability distribution, decays as a power law $S^{-\kappa}$, with a cut-off that diverges as one approaches the transition α^* . The dotted line corresponds to $\kappa = 0.74$.




Figure 6. Stability map: crisis probability $\mathbb{P}[\tau_c \leq T]$ for $T = 200$, $N = 280$, $\lambda = 10$, $\nu_0 = 1$ and $\mu = 20$. The blue region correspond to a stable order book, whereas the red region corresponds to liquidity crises. The crossover line $\alpha_m(\beta)$ is the white sliver between the two.

$N \rightarrow \infty$ and $T \rightarrow \infty$. Clearly, for finite N , there is always a nonzero probability (perhaps very small) that the order book completely empties if one waits long enough, even when $\alpha_K = 0$. Hence:

$$\lim_{N \rightarrow \infty} \lim_{T \rightarrow \infty} \mathbb{P}_N[\tau_c \leq T, \alpha_K] = 1, \forall \alpha_K.$$

If on the other hand the limit $N \rightarrow \infty$ is taken first, one may be in a situation where, for a fixed value of β

$$\lim_{T \rightarrow \infty} \lim_{N \rightarrow \infty} \mathbb{P}_N[\tau_c \leq T, \alpha_K] = \begin{cases} 1, & \text{when } \alpha_K > \alpha^*, \\ 0, & \text{when } \alpha_K < \alpha^*, \end{cases} \quad (5)$$

where α^* depends on the parameters of the model, in particular β .

Since numerical simulations can only be done for finite N and T , a common strategy is to use finite size scaling to extrapolate to infinite sizes and waiting times. If a genuine, continuous phase transition occurs at some $\alpha_K = \alpha^*$, one expects the following behaviour to hold for large enough N and T :

$$\mathbb{P}_N[\tau_c \leq T, \alpha_K] = F(T(\alpha_K - \alpha_m(T, N))^\zeta); \quad \alpha_m(T, N) = \alpha^* - \frac{1}{T^{1/\zeta}} g\left(\frac{N^\eta}{T}\right), \quad (6)$$

with $F(u)$ a monotonic regular function going from 0 for $u \rightarrow -\infty$ to 1 for $u \rightarrow +\infty$, and $g(v)$ another function that goes to a constant g_∞ when $v \rightarrow \infty$ and to $+\infty$ as $v \rightarrow 0$. This scaling form has the following interpretation:

- when $1 \ll T \ll N^\eta$, $\alpha_m \approx \alpha^*$. As α_K increases, $\mathbb{P}_N[\tau_c \leq T, \alpha_K]$ evolves from 0 (no crises) to 1 (crises) in a region of width $T^{-1/\zeta}$ around α^* ;
- when $T \gg N^\eta$, α_m becomes negative, meaning that $\mathbb{P}_N[\tau_c \leq T, \alpha_K]$ is close to 1 for any α_K if one waits long enough.

The comparison between T and N^η has the following interpretation: for $T \ll N^\eta$, the system cannot ‘feel’ the boundaries of the order book because the spread has never grown so large: $S(T) \ll N$. For $T \gg N^\eta$ on the other hand, it is highly probable that the spread S has been as large as size of the order book N , meaning that a liquidity crisis has taken place. This suggests a direct way to measure η , from the dynamics of the spread that behaves as a power-law of time (see figure 5), with an exponent which should equal $1/\eta$ for consistency. This gives $\eta \approx 3$, which is compatible with the finite size scaling analysis reported in figure 7 (left inset).

A convenient method to pin down the values of α^* and the exponent ζ is to study the variance of the first crisis time, defined as

$$\chi(\alpha_K, T, N) = \mathbb{V}[\min(\tau_c, T)]$$

for a fixed value of β and different values of α_K , T and N . This quantity is expected to peak close to the phase transition, since for small α_K , τ_c is nearly always larger than T and $\chi \rightarrow 0$, whereas for large α_K , τ_c is small and χ is also small. The finite size scaling assumption for this quantity amounts to:

$$\chi(\alpha_K, T, N) = T^\gamma G(T(\alpha_K - \alpha_m(T, N))^\zeta), \quad (7)$$

where α_m is given by equation (6) with $\alpha^* \approx 0.06$ and $G(u)$ is a humped function that goes to zero for $u \rightarrow \pm\infty$. The details and justification of this procedure to find the

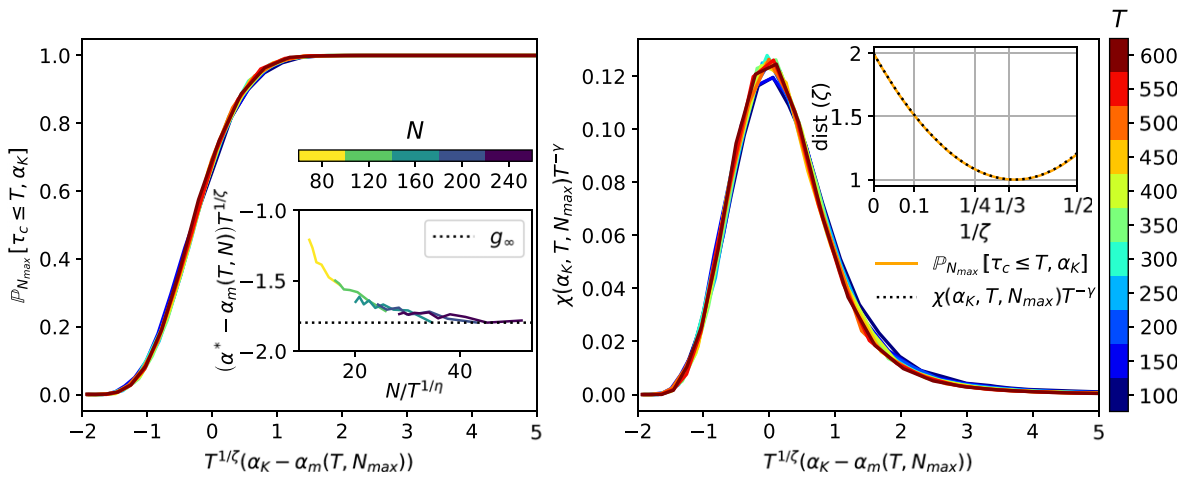


Figure 7. Left: rescaling of $\mathbb{P}_N[\tau_c \leq T, \alpha_K]$ according to equation (6), for $\beta = 0.5$, $\lambda = 10$, $\nu_0 = 1$, $\mu = 20$, $N = N_{\max} = 240$, and with $\zeta = 3$ and $\eta = 3$. Inset: numerical determination of $g(v)$, also well described by our scaling hypothesis with $\eta = 3$, giving $\alpha^* \approx 0.063$ and $g_\infty = \lim_{+\infty} g \approx -1.8$. Right: finite size scaling of χ , with $\alpha_m(T, N_{\max}) := \operatorname{argmax}_\alpha \chi(\alpha_K, T, N_{\max})$. Inset: average pairwise distance between the different curves as a function of $1/\zeta$, showing a minimum for $\zeta \approx 3$ for both quantities \mathbb{P} and χ .

different exponents is described in appendix B. We find $\gamma \approx 2$ and $\zeta \approx \eta \approx 3$ ¹³. Figure 7 shows how all the different curves collapse on top of each other when these parameters are fixed. We also show the quality of this rescaling as a function of ζ in the inset of figure 7 (right), clearly favouring the value $\zeta = 3$.

We note that to the best of our knowledge, the numerical value of the exponents ζ, η do not seem to relate to an identified phase transition. It would be very interesting to explore further the nature of this transition and (if possible) compute analytically the value of these exponents.

3.4. Discussion

Although not perfect, we consider the rescaling sufficiently convincing to support our interpretation that the observed liquidity transition is a second order phase transition. This interpretation is further supported by the fact that a similar finite size scaling with the same value of the exponents ζ, η (but different values of α^*) holds for different values of the time scale β and rates λ, ν_0 and μ , and is also robust against changes in the specification of the model. This universality is a landmark of second order phase transitions.

Although our numerical evidence for such a phase transition is satisfactory, we have not found a way to bolster our results by a rigorous mathematical analysis. Indeed, even if highly stylized, the Santa Fe model with feedback is in fact quite complex. Hence, the

¹³Note that the value $\gamma = 2$ is not unreasonable since for $\alpha_K = \alpha^*$ one expects that τ_c is larger than T with some probability $p \in]0, 1[$, leading to $\mathbb{V}[\min(\tau_c, T)] \propto T^2$.

existence of this phase transition, and its second-order nature, can only be considered as conjectures at this stage. In order to make some progress, we have studied even simpler models, where the existence of a phase transition can be ascertained mathematically. This is what we discuss in the following sections.

4. A state-dependent Hawkes model for spread dynamics

In this section we introduce and discuss a family of simple models for which the liquidity crisis transition observed within the Santa Fe model can be analyzed in more details. This however comes at the price of setting aside the dynamics of the order book, and restrict our attention to the dynamics of the spread.

4.1. A simple model

The simplest class of models consist in retaining the feedback of the spread dynamics on itself, forgetting about the price dynamics which is the main driver of the instability in the context of the Santa Fe model. Hence we also move away from the quadratic Hawkes model we calibrated on real data in section 2. The destabilizing mechanism we imagine is that spread opening events are likely to lead to more spread opening events. A simplified model reinstating the price feedback mechanism is presented in appendix D.

We consider an order book filled with limit orders of size unity that can be cancelled or executed only at the best. As soon as a price slot is filled, no further limit order can be placed. Limit orders can thus only be placed in front of the best (inside the spread) provided the spread is open, i.e. $S_t := a_t - b_t \geq 2$. Note that with such simplifying hypotheses, there is no gap in the order book (apart from the spread itself) and market orders play the exact same role as cancellations. The model can thus be entirely characterized by the spread dynamics. We assume that the event intensities read:

$$\lambda_t^+ = \lambda_0^+ + \alpha \int_0^t \beta e^{-\beta(t-s)} dS_s^+ \quad (8a)$$

$$\lambda_t^- = \mathbb{1}_{\{S_t \geq 2\}} \lambda_0^- \quad (8b)$$

where λ^+ is the intensity of events that increase the spread, i.e. orders that are cancelled or executed by a market order and λ^- is the rate of limit orders reducing S_t by falling inside the spread. Only spread opening events contribute to the feedback on λ^+ , i.e. $dS_t^+ := \max(dS_t, 0)$. This highly stylized model has the advantage of being analytically tractable, while giving valuable insights on the possible phase transitions that can take place in order book models with feedback.

For $\alpha < 1$, using linear Hawkes theory, one can show (see e.g. [24]) that there exists a martingale process M_t such that:

$$S_t = S_0 + \int_0^t \left[(1 - \alpha e^{-(1-\alpha)\beta s}) \frac{\lambda_0^+}{1 - \alpha} - \mathbb{1}_{\{S_s \geq 2\}} \lambda_0^- \right] ds + M_t. \quad (9)$$

Introducing the parameter $\alpha_c = 1 - \lambda_0^+/\lambda_0^-$ one can distinguish between the different regimes¹⁴:

- $0 \leq \alpha < \alpha_c$ —the system is Hawkes-stable and the spread has a stationary distribution.
- $\alpha_c < \alpha < \alpha^* = 1$ —the system is Hawkes-stable but the spread increases on average linearly with t .
- $\alpha \geq \alpha^* = 1$ —the system is Hawkes-unstable, or ‘explosive’.

The terminology Hawkes-(un)stable refers to the stability transition of a linear Hawkes process, that is, the transition between a regime where the intensities reach a stationary state from a regime where the number of events grows exponentially with time.

4.2. The stable regime

Let us first discuss the stable regime $\alpha < \alpha_c$. In the stationary state, we can prove that the probability for the spread to be open is given by:

$$\mathbb{P}[S \geq 2] = \frac{1 - \alpha_c}{1 - \alpha},$$

which goes to 1 as $\alpha \uparrow \alpha_c$. This result reproduces very well our numerical data. Although we have not been able to prove the result mathematically, numerical simulations also suggest that the full distribution of the spread is exactly geometric in this model:

$$\mathbb{P}[S = n] = \frac{1 - \alpha_c}{1 - \alpha} (1 - r) r^{n-2}; \quad n \geq 2,$$

where r depends of α and β , see figure 8 (left). This result should in principle follow from the following equation that describes the evolution of the two-dimensional density function $\rho_t(S_t, X_t = \int_0^t \beta e^{-\beta(t-s)} dS_s^+)$:

$$\begin{aligned} \partial_t \rho_t &= [\lambda_0^+ + \alpha(x - \beta)] \rho_t(S - 1, x - \beta) \mathbb{1}_{\{S \geq 2, x \geq \beta\}} - [\lambda_0^+ + \alpha x] \rho_t(S, x) \\ &+ \lambda_0^- \rho_t(S + 1, x) - \lambda_0^- \mathbb{1}_{\{S \geq 2\}} \rho_t(S, x) + \beta \partial_x (x \rho_t(S, x)), \end{aligned} \quad (10)$$

see appendix C. Setting the left-hand side to zero gives the stationary joint distribution of S and X . However, we have not been able to make much analytical progress, except in the limit $\beta \rightarrow 0$ where a geometric distribution for S indeed follows with $r = (1 - \alpha_c)/(1 - \alpha)$. Unfortunately, this approximation does not hold for $\beta \sim 1$ but works well for small β . Note that in the presence of a price feedback mechanism, the spread distribution acquires a power-law tail as we observed within the extended Santa Fe model (see appendix D).

¹⁴These results are general to any Hawkes kernel ϕ provided $\|\phi\| = \alpha < 1$.

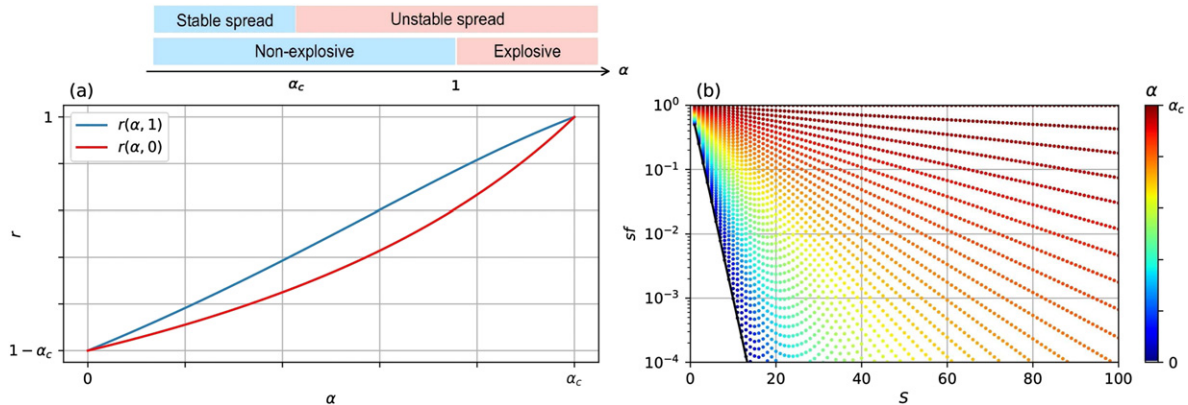


Figure 8. Properties of the spread in the linear model for $\alpha < \alpha_c$, using a set of parameters $\lambda_0^+ = 1$, $\lambda_0^- = 0.5$ and $\beta = 1$. (a) Plot of $\alpha \mapsto r(\alpha, \beta)$ with $\beta = 1$ and the theoretical results of the limit $\beta \rightarrow 0$. (b) Log survival function (sf) of the spread for different values of α , suggesting an exact geometrical distribution for all α . The black curve corresponds to the theoretical equilibrium distribution when $\alpha = 0$.

4.3. Linear spread growth

In the interesting regime $\alpha_c < \alpha < \alpha^* = 1$ phase, one finds that the spread grows on average linearly in time, with a drift V that vanishes when $\alpha \downarrow \alpha_c$:

$$\lim_{t \rightarrow \infty} \frac{1}{t} \mathbb{E}[S_t] = V; \quad V := \lambda_0^+ \frac{\alpha - \alpha_c}{(1 - \alpha)(1 - \alpha_c)}.$$

On top of this average drift, the spread has diffusive fluctuations with some diffusion constant D defined as:

$$D(\alpha) := \lim_{t \rightarrow \infty} \frac{1}{t} [\mathbb{E}[S_t^2] - \mathbb{E}[S_t]^2] = \lambda_0^- + \frac{\lambda_0^+}{(1 - \alpha)^3}.$$

One can thus compute the probability that the spread exceeds some threshold N before time T , corresponding to an empty book in the Santa Fe model. Using standard first passage time results for the one dimensional Brownian motion [31], one has, for large N and T (and keeping the same notation as in section 3):

$$\mathbb{P}_N[\tau_c \leq T, \alpha] = \int_0^T du \frac{N}{\sqrt{2\pi D(\alpha)u^3}} e^{-\frac{(N+Vu)^2}{2D(\alpha)u}}. \tag{11}$$

While the spread will eventually exceed N for large enough time, it is easy to see that:

$$\lim_{T \rightarrow \infty} \lim_{N \rightarrow \infty} \mathbb{P}_N[\tau_c \leq T, \alpha] = 0,$$

for all $\alpha < \alpha^* = 1$. In other words, the second order transition observed in the Santa Fe model with feedback is absent in the present setting. While a linear increase of the spread is interesting, it can hardly be called a liquidity ‘crisis’. Similarly the susceptibility χ

J. Stat. Mech. (2020) 063401

can be easily computed using equation (11) and one finds:

$$\chi(\alpha, T, N) = \mathbb{V}[\min(\tau_c, T)] = \int_0^T du \frac{N(T-u)^2 e^{-\frac{(N+Vu)^2}{2D(\alpha)u}}}{\sqrt{2\pi D(\alpha)u^3}} - \left(\int_0^T du \frac{N(T-u) e^{-\frac{(N+Vu)^2}{2D(\alpha)u}}}{\sqrt{2\pi D(\alpha)u^3}} \right)^2. \tag{12}$$

This gives us the same result as above:

$$\lim_{T \rightarrow \infty} \lim_{N \rightarrow \infty} \chi(\alpha, T, N) = 0,$$

i.e. no liquidity ‘crisis’. Nevertheless, it is interesting to notice that χ can be written exactly in the scaling form that we used to analyze the Santa Fe model. Indeed, the result of equation (12) can be transformed into:

$$\chi(\alpha, T, N) = T^\gamma \mathcal{G}(NT^{-1/\eta}, T^{1/\zeta}(\alpha - \alpha_c)),$$

with:

$$\mathcal{G}(x, y) = x \int_0^1 \frac{du(1-u)^2}{\sqrt{2\pi D(\alpha_c)u^3}} \exp\left(-\frac{[x + uy\Lambda]^2}{2D(\alpha_c)u}\right) - \left(x \int_0^1 \frac{du(1-u)}{\sqrt{2\pi D(\alpha_c)u^3}} \exp\left(-\frac{[x + uy\Lambda]^2}{2D(\alpha_c)u}\right) \right)^2,$$

$\Lambda := \lambda_0^+(1 - \alpha_c)^{-2}$ and $\gamma = \eta = \zeta = 2$. These exponents should be compared with the values found numerically for the generalized Santa Fe model: $\gamma \approx 2$, $\eta \approx \zeta \approx 3$.

4.4. The explosive regime

When $\alpha > \alpha^* = 1$, the model becomes Hawkes unstable, which means in the present context that the spread increases exponentially with time. Although formally the spread never diverges in finite time, in practice there is a ‘liquidity crisis’ as soon as $T(\alpha - \alpha^*) \propto \log N$, i.e. when the spread reaches the boundary of the order book. This would look numerically akin to a second order phase transition with exponents $\zeta = 1$ and $\eta = 0$, quite far from the results reported for the Santa-Fe model.

4.5. A stabilizing mechanism

One could expect some stabilizing mechanisms to arise when the spread becomes too large. A way to include the latter in our simple setting by substituting equation (8b) with:

$$\lambda_t^- = \lambda_0^-(S_t - 1), \tag{13}$$

meaning that there is an increased probability to introduce limit orders inside the spread when it is large. The model remains analytically tractable; the bottom line is that the Hawkes stable regime $\alpha_c < \alpha < 1$ disappears: our specification is indeed able

to stabilise the spread in the whole region $\alpha < 1$. The Hawkes unstable regime $\alpha > 1$ of course subsists and is associated to liquidity crises in an otherwise stable market ($\alpha < 1$).

5. Non-linear Hawkes models and metastability

We have studied in the previous section a simple spread dynamics model that maps onto a linear Hawkes process. In these models, the spread becomes unstable and grows linearly in time before the Hawkes process (i.e. the activity of the process) becomes itself explosive. One can stabilize the spread dynamics, as in the last subsection above, such that sudden liquidity crises in this model are associated to the Hawkes explosive transition.

For this picture to be correct, however, real financial markets must sit below, but very close to the Hawkes instability threshold α^* , or else one must argue that α itself is time dependent, and occasionally visits the explosive region $\alpha > \alpha^*$ before decreasing back below α^* , allowing the market to re-stabilise. The same remark in fact applies to the generalised Santa Fe model studied in section 3: if liquidity crises are indeed related to the existence of a second order phase transition, one must argue that financial markets are for some reason close to the critical point—a phenomenon called ‘self-organized criticality’ [32]—or that the parameters fluctuate over time and occasionally push the system in the unstable phase.

Although many models in mathematical finance are tweaked such that their parameters become time dependent, we feel that this common procedure might in fact hide the inadequacy of such models. In this section, we want to explore an alternative scenario. We introduce a class of non-linear Hawkes process that show occasional liquidity crises without either being poised at the edge of instability ($\alpha \uparrow \alpha^*$) or having a time dependent feedback parameter α .

5.1. A model with quadratic feedback

Let us consider again the simplified framework of section 4.1 and generalize the feedback on spread opening events as:

$$\lambda_t^+ = \lambda_0^+ + \alpha X_t + \epsilon X_t^2; \quad X_t := \int_0^t \beta e^{-\beta(t-s)} dS_s^+. \quad (14)$$

When $\epsilon = 0$, this Hawkes process is non-explosive provided $\alpha < \alpha^* = 1$. But as soon as $\epsilon > 0$, the process has a non zero probability to explode, even when $\alpha < \alpha^*$. However, interestingly, these ‘liquidity crises’ only happen with a rate that is *exponentially small* in $1/\epsilon$, and therefore interrupt very long periods of apparent market stability—a phenomenon called ‘metastability’ in the physics literature. This is confirmed by direct numerical simulations of the model equation (14) in figure D1. In the following section, we give an analytical description of this phenomenon.

5.2. A continuous time description

In the ‘slow’ limit $\beta \rightarrow 0$ one can write an approximate SDE for X_t . Start from the exact expression $dX_t = -\beta X_t dt + \beta dS_t^+$. When β is small, λ_t^+ is slowly varying and one can approximate dS_t^+ by $\lambda_t^+ dt + \sqrt{\lambda_t^+} dB_t$, where B_t is a Brownian motion (for more rigorous statements, see [33]). Hence:

$$dX_t = \beta [\lambda_0^+ - (1 - \alpha)X_t + \epsilon X_t^2] dt + \beta \sqrt{\lambda_t^+(X_t)} dB_t. \tag{15}$$

Let us write the deterministic part of this equation as minus the derivative of some ‘potential’ $\mathcal{V}(X)$, to wit:

$$\mathcal{V}(X) = \frac{\beta(1 - \alpha)}{2} \left(X - \frac{\lambda_0^+}{1 - \alpha} \right)^2 - \frac{\beta\epsilon}{3} X^3. \tag{16}$$

Such a potential is drawn for $\alpha < 1, \epsilon = 0$ and $\alpha < 1, \epsilon > 0$ in figure 9. One sees clearly that for $\epsilon = 0$ the equilibrium $X_{\text{eq}} = \lambda_0^+ / (1 - \alpha)$ (that corresponds to the average intensity of the Hawkes process) is stable. But as soon as $\epsilon > 0$ the potential reaches a maximum for some value X^* beyond which it plunges towards $-\infty$. In the limit $\epsilon \rightarrow 0$, one finds that X^* is given by:

$$X^* \approx \frac{1 - \alpha}{\epsilon},$$

corresponding to:

$$\mathcal{V}(X^*) \approx \frac{\beta(1 - \alpha)^3}{6\epsilon^2},$$

which diverges when $\epsilon \rightarrow 0$. This picture allows one to describe the dynamics of the model for $\epsilon > 0$ in intuitive terms: for a very long time, X_t will oscillate around its equilibrium value X_{eq} until some rare fluctuation of the Brownian noise dB_t is able to bring X_t close to the top of the high barrier X^* . In such rare circumstances, X_t escapes the stable valley and runs all the way to $+\infty$ in finite time, corresponding to a ‘liquidity crisis’.

The theory of high barrier crossing under the influence of noise is very well understood. In the present case, the final formula for the average first escape time τ_c (corresponding to the ‘emptying of the book’ as in section 3) is given by [6, 34]:

$$\mathbb{E}[\tau_c] \approx 2\pi \left(\frac{D(X_{\text{eq}})/D(X^*)}{|\mathcal{V}''(X^*)\mathcal{V}''(X_{\text{eq}})|} \right)^{1/2} \times \exp \left(\int_{X_{\text{eq}}}^{X^*} dx \frac{\mathcal{V}'(x)}{D(x)} \right), \tag{17}$$

with $D(X) := \frac{\beta^2}{2} (\lambda_0^+ + \alpha X + \epsilon X^2)$. The expansion to second order in ϵ gives:

$$\log \mathbb{E}[\tau_c] \underset{\epsilon \rightarrow 0}{\approx} \begin{cases} -\frac{2}{\beta} \left[\frac{1 - \alpha + \log \alpha}{\epsilon} + \frac{\lambda_0^+}{\alpha^2} \log \frac{1}{\epsilon} \right] - \frac{1}{2} \log \frac{1}{\epsilon} & \text{if } \alpha > 0 \\ \frac{1}{\beta\epsilon} \left(\log \frac{1}{\epsilon\lambda_0^+} - 2 \right) & \text{if } \alpha = 0. \end{cases} \tag{18}$$




Figure 9. Plots of the potential as given by equation (16) with $\lambda_0^+ = 1$ and $\alpha = 0$.

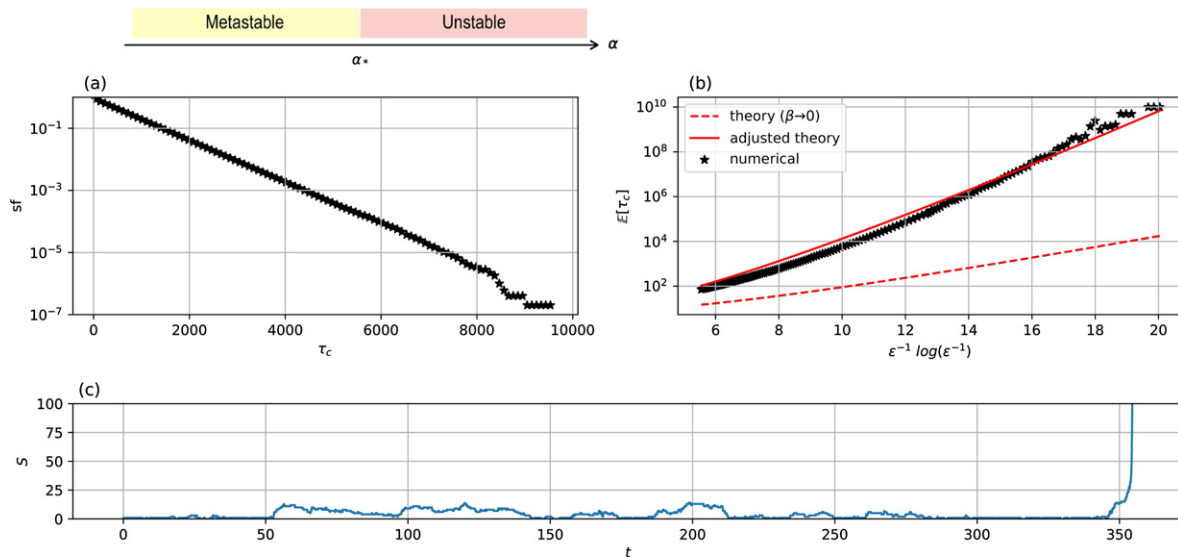


Figure 10. Properties of the time of metastability when $\alpha = 0$ and $\epsilon > 0$. (a) Survival function (sf) of the time before explosion for $\epsilon = 0.2$, which is found to be exponential. (b) Evolution of the average metastability time with ϵ . The dotted red curve is the continuous time prediction given by equation (18). The plain red curve is obtained by multiplying the term in the exponential by an empirical factor 2.5. (c) Typical metastable trajectory. The set of parameters is the same that in (a): $\lambda_0^+ = 1$, $\lambda_0^- = 0.5$, $\beta = 1$ and $\epsilon = 0.2$.

Hence, as announced, the time before a crisis is exponentially large in ϵ^{-1} , with logarithmic corrections for $\alpha = 0$. Another prediction of this approach is that in the limit $\epsilon \rightarrow 0$, the time-to-crisis becomes a Poisson variable with mean $\mathbb{E}[\tau_c]$, as indeed found numerically (see figure 10(a)).

Our analytical result compares well with our numerical results in terms of the overall dependence on ϵ , but the numerical prefactor inside the exponential is off by a factor ~ 2.5 . This can be traced to the fact that our numerical simulations are in a regime

where $\beta/\lambda_0^+ = O(1)$, whereas the theoretical analysis is done in a regime where $\beta/\lambda_0^+ \rightarrow 0$ (see [35], and [6], section 5.4, where a similar phenomenon is present).

6. Conclusion

Let us summarise what we have achieved in this study. Using tick-by-tick order book data on futures contracts, we were able to show that event rates are strongly affected by past price moves. In particular, large price trend and/or volatility tends to increase the rate of market orders and cancellations, which subsequently leads to a decrease in liquidity. This, in turn, contributes to increasing volatility, which may lead to a destabilising feedback loop and a liquidity dry-out.

Let us stress that the work presented in this paper is relevant for both effects, explaining ‘regular’ excess volatility and understanding extreme endogenous price jumps such as flash crashes. Building on such empirical evidence, we introduced an extension of the stylized Santa Fe model which accounts for the feedback of past price changes on event rates. Numerical simulations of our model revealed the existence of a second order phase transition, and more precisely a critical value of the feedback parameter below which an infinite size order book never empties, and above which it empties with probability one. We performed a finite size scaling analysis in order to determine the critical exponents, which to the best of our understanding do not appear to be in any of the known universality classes for 1D phase transitions.

In order to bolster our results with analytical arguments, we then considered an even simpler model, where the existence of a phase transition can be verified mathematically. Setting aside the dynamics of the order book, and focusing our attention to the dynamics of the spread, we presented a model which can be mapped onto a linear Hawkes process in which spread opening events are likely to lead to more spread opening events. We exhibited three dynamic spread regimes as function of feedback intensity: stable, linearly increasing, and exploding spread. We argued that the second regime could be stabilised and that in such a case, and within some parameter range, a phase transition from stable to unstable spread exists, much like in the extended Santa Fe model presented before, but with major quantitative differences.

We then pointed out that for this picture to be relevant, real financial markets would have to sit below, but very close to the critical point, consistent with the idea of self-organised criticality (SOC), a concept first introduced in [32] and developed by many in the context of game theory [36] and financial markets [37–39]. Another option would be for the feedback parameter to be itself time dependent and occasionally visit the unstable phase.

Finally, we presented an alternative scenario which needs no *a priori* proximity to the instability threshold, nor a time dependent feedback parameter. The model is a non-linear Hawkes process for which liquidity crises are ‘activated’ events within a metastable phase. A continuous time description allowed us to derive the typical crisis frequency as function of the model’s parameters, and show that this time can be much longer than the microscopic time of the model. It is quite likely that both effects (proximity of a transition and metastability) do actually coexist in financial markets.

Several improvements/extensions of this work would be of interest. In particular, achieving proper calibration of the Q-Hawkes model presented in equation (1) in section 2, perhaps through a non-parametric procedure, would help buttress our empirical results. A deeper mathematical analysis aimed at deriving the critical exponents of the extended Santa Fe model presented in section 3 would be highly valuable to ascertain the new universality class we exhibited numerically. For the sake of completeness, it would naturally also be of interest to couple our second order phase transition scenario to a mechanism that draws the systems towards the critical point, building on ideas inspired e.g. by the minority game, see [36]. Research should probably also be devoted to thinking about which empirical test could help discriminating between the second order phase transition and activation scenario.

An important point on which we decided not to insist too much in the present study is the effect of memory timescales, apart from the empirical section in which we thoroughly analyzed the intensity of the response functions as function of forward and backward memory timescales. Indeed, as shown for example in [11], lag effects can be extremely important destabilising factors that must be taken into account. Including these lag effects within the present framework is certainly a relevant extension worth investigating.

Acknowledgments

We thank Jonathan Donier, Iacopo Mastromatteo, José Moran, Mehdi Tomas, Stephen Hardiman and Mathieu Rosenbaum for fruitful discussions. This research was conducted within the *Econophysics & Complex Systems Research Chair*, under the aegis of the Fondation du Risque, the Fondation de l'École polytechnique, the École polytechnique and Capital Fund Management.

Appendix A. Empirical data

We used tick-by-tick (or event-by-event) data for four futures contracts (EUROSTOXX, BUND, BOBL & SCHATZ) over around 160 trading days provided by CFM. We have chosen these assets because of their high liquidity and because they are all large tick (the spread is equal to its minimal value of 1 tick more than 99% of the time). Before doing any specific inference on the data, we preprocess it in the following way.

- We load data from 9 am to 4 pm.
- Separate events displaying the same timestamps are shuffled within the millisecond without breaking causality.
- We restrict to the best queues only.
- We use the mid-price changes in tick units.

The number of events after cleaning is of the order of one million per day for the most liquid asset, and 50 000 for the least liquid. First, we perform a non-parametric

Hawkes calibration that gives the parameters $Q\tilde{\alpha}_0$ and $Q\phi$, as defined in equation (1). Then, we turn to the contribution of the trend and the volatility. To do so, we compute $\beta' F_{\beta'}^{b+a}$, $\beta' H_{\beta'}^{b+a}$, $\beta' F_{\beta'}^{b-a}$, $\beta' H_{\beta'}^{b-a}$, R_β and Σ_β^2 , as defined in section 2.3. For practical reasons, we approximate the kernels with a sum of three exponentials, in the spirit of [40], which allows for a fast algorithm thanks to the recursivity of the exponentially weighted moving averages. We associate a weight to each of these quantities that is the fraction of inter-event time, and bin the data in $100 \times 100 \times 100$ -sized windows for $F_{\beta'}^{b+a}$ with $H_{\beta'}^{b+a}$, R_β , Σ_β^2 and $F_{\beta'}^{b-a}$ with $H_{\beta'}^{b-a}$, R_β , Σ_β^2 . We aggregate the weights to get a weight for each bin, and perform the regressions given in equations (2) and (3) using a very standard generalised least square method [41]. We take the values of β and β' that maximise the absolute correlation $\text{Cor}(F_{\beta'}^{b+a}, R_\beta^2)$.

Appendix B. Finite size scaling method

Here, we discuss the method used to do the finite size scaling in section 3.3. First, let us recall the framework. The susceptibility writes:

$$\chi(\alpha_K, T, N) = T^\gamma G(T(\alpha_K - \alpha_m(T, N))^\zeta) = T^\gamma \mathcal{G}(NT^{-1/\eta}, T^{1/\zeta}(\alpha_K - \alpha^*)), \tag{B1}$$

where the function \mathcal{G} satisfies:

- $\lim_{|y| \rightarrow +\infty} \mathcal{G}(x, y) = 0$
- $\forall x, y \mapsto \mathcal{G}(x, y)$ has a unique maximum, denoted $y^*(x)$.

First, we determine γ . We introduce $\alpha_m(T, N) = \text{argmax}_{\alpha_K} \chi(\alpha_K, T, N)$ and we assume that $\lim_{T, N \rightarrow \infty} \alpha_m(T, N) = \alpha^*$. The idea is to look at $\max_{\alpha_K} \chi(\alpha_K, T, N) = \chi(\alpha_m(T, N), T, N)$:

$$\chi(\alpha_m(T, N), T, N) = T^\gamma \max_y \mathcal{G}(NT^{-1/\eta}, y) \xrightarrow{N \rightarrow +\infty} T^\gamma \lim_{x \rightarrow +\infty} \max_y \mathcal{G}(x, y). \tag{B2}$$

If $N_{\max}^\eta \gg T$, then $\chi(\alpha_m(T, N_{\max}), T, N_{\max}) \approx T^\gamma \lim_{x \rightarrow +\infty} \max_y \mathcal{G}(x, y)$ on our range of T . Note that the validity of such a hypothesis depends on the value of η , which we shall determine and self-consistently validate below. Then we compute the value of $\gamma \approx 2$ from a linear regression of $\log \chi(\alpha_m(T, N_{\max}), T, N_{\max})$ vs $\log T$.

Then we determine ζ . If T, N_{\max} are large enough that $N_{\max}^\eta \gg T$, then $T^{1/\zeta}(\alpha_K - \alpha_m(T, N_{\max})) \approx T^{1/\zeta}(\alpha_K - \alpha^*)$ and $\chi(\alpha_K, T, N_{\max}) \approx T^\gamma \lim_{x \rightarrow +\infty} \mathcal{G}(x, T^{1/\zeta}(\alpha_K - \alpha^*))$. So we plot $\chi(\alpha_K, T, N_{\max})$ as a function of $T^{1/\zeta}(\alpha_K - \alpha_m(T, N_{\max}))$ for different values of T and α and we tune the exponent ζ to make all the curves collapse together, see figure 7. We can do this experiment numerically by minimising the distance between the curves as a function of ζ . Adding the fact that we expect regular rational values we deduce the most likely exponent, $\zeta = 3$, see right inset of figure 7.

Finally we compute α^* and η . By definition of $\alpha_m(T, N)$, one has $T^{1/\zeta}(\alpha_m(T, N) - \alpha^*) = y^*(NT^{-1/\eta})$. Thus, if one plots $T^{1/\zeta}(\alpha_m(T, N) - \alpha^*)$ as a function of $NT^{-1/\eta}$ for different values of T and N , one should find a set of parameters η and α^* such that all the curves collapse together. This leads to $\alpha^* \approx 6.3 \times 10^{-2}$ and $\eta \approx 3$, which is compatible with the direct result on the spread dynamics shown in figure 5, where one observes that $S(t) \sim t^{1/3}$. But since the finite size-finite time crossover should occur when $S(T) \sim N$, one finds that $T^{1/3} \sim N$, again leading to $\eta \approx 3$.

Appendix C. More on the linear spread model

Here, we focus on the linear case ($\epsilon = 0$), see equation (D3). In order to remain very general we rewrite the equation as $\lambda_t^+ = \lambda_0^+ + (\phi * dS^+)_t = \lambda_0^+ + \int_0^t \phi(t-s) dS_s^+$. Point process theory teaches us that there exists two independent martingales M_t^- and M_t^+ such that $S_t^\pm = \lambda_t^\pm dt + dM_t^\pm$. One can write:

$$\lambda_t^+ = \lambda_0^+ + (\phi * \lambda^+)_t + (\phi * dM^+)_t. \tag{C1}$$

Assuming that $\|\phi\| < 1$ one can define the resolvent $\phi_R = \sum_{n \geq 1} \phi^{*n}$, with $\phi^{*(n+1)} = \phi * \phi^{*n}$. Note that $(\delta + \phi_R) * \phi = \phi_R$ with δ the Dirac function. This enable to invert the above equation and obtain:

$$\lambda_t^+ = \left(1 + \int_0^t \phi_R(s) ds\right) \lambda_0^+ + (\phi_R * dM^+)_t. \tag{C2}$$

Combining the previous equations and introducing the martingale M_t with $dM_t = dM_t^+ - dM_t^- + (\phi_R * dM^+)_t dt$, one gets:

$$S_t = S_0 + \int_0^t \left[\left(1 + \int_0^s \phi_R(u) du\right) \lambda_0^+ - \mathbb{1}_{\{S_s \geq 2\}} \lambda_0^- \right] ds + M_t. \tag{C3}$$

Equation (9) is the particular case with $\phi(t) = \alpha\beta e^{-\beta t}$. Choosing such a kernel one can derive the Fokker-Planck equation for the joint distribution of the variables $(S_t, X_t = \beta \int_0^t e^{-\beta(t-s)} dS_s^+)$ given in equation (10). While we did not manage to solve this equation, we can compute the Laplace transform of the variable X_t at equilibrium:

$$\mathbb{E} [e^{-uX}] = \int_{\mathbb{R}^+} \rho_X^{st}(x) e^{-ux} dx = \exp \left(\int_0^u \frac{\lambda_0^+ (1 - e^{-\beta v})}{\alpha (1 - e^{-\beta v}) - \beta v} dv \right), \tag{C4}$$

from which we can get the cumulants. In particular, one has: $\mathbb{E}[X] = \lambda_0^+ / (1 - \alpha)$ and $\mathbb{V}[X] = \beta \lambda_0^+ / (2(1 - \alpha)^2)$. Interestingly, we can get the full stationary solution ρ^{st} in two simple limit:

- $\alpha = 0$: $\rho^{st}(S, x) = (1 - r)r^S \rho_X^{st}(x)$ with $r = \lambda_0^+ / \lambda_0^-$.
- $\beta \rightarrow 0$: $\rho^{st}(S, x) = \delta \left(x - \frac{\lambda_0^+}{1 - \alpha} \right) (1 - r)r^S$ with $r = \frac{1 - \alpha_c}{1 - \alpha}$.

Note that the spread is geometrically distributed in both cases.

Appendix D. A model with price feedback on the spread

The model presented in section 4 displays a very simple destabilizing mechanism in the midprice reference frame, in which spread opening events trigger more spread opening events. Here, we re-introduce price dynamics to illustrate the effects of ‘volatility’ in the spread opening mechanism, bringing the model one step closer to the empirical study presented in section 2 and the Santa Fe model of section 3. In order to do so, we write the intensity of cancellations/market orders as:

$$\lambda_t^+ = \lambda_0^+ + \alpha \left(\int_0^t \sqrt{2\beta} e^{-\beta(t-s)} dP_s \right)^2 \quad (\text{D1a})$$

$$\lambda_t^- = \lambda_0^- \mathbb{1}_{\{S_t \geq 2\}}. \quad (\text{D1b})$$

Each event takes place with equal probability at the bid b_t or the ask a_t . The dynamics of the bid and ask are thus such that

$$\mathbb{E}[da_t | da_t > 0, \mathcal{F}_t] = -\mathbb{E}[db_t | db_t < 0, \mathcal{F}_t] = \lambda_t^+ dt/2$$

and

$$\mathbb{E}[da_t | da_t < 0, \mathcal{F}_t] = -\mathbb{E}[db_t | db_t > 0, \mathcal{F}_t] = -\lambda_t^- dt/2.$$

Using the formalism introduced in [C1] we are able to solve the dynamics of the system and show that there exists a martingale M_t such that:

$$S_t = S_0 + \int_0^t \left[(2 - \alpha e^{-(2-\alpha)\beta s}) \frac{\lambda_0^+}{2-\alpha} - \frac{\mathbb{1}_{\{S_s \geq 2\}}}{2-\alpha} \lambda_0^- \right. \\ \left. \times \left(2 - 2\alpha - \alpha e^{-(2-\alpha)\beta(t-s)} \right) \right] ds + M_t. \quad (\text{D2})$$

Calling again $\alpha_c = 1 - \lambda_0^+/\lambda_0^-$, one obtains the same regimes as in section 4.1 only replacing $\alpha^* = 1$ by $\alpha^* = 2$:

- $0 \leq \alpha < \alpha_c$ —the system is non-explosive and the spread has a stationary distribution.
- $\alpha_c < \alpha < 2$ —the system is non-explosive but the spread increases on average linearly with t .
- $2 \leq \alpha$ —the system is explosive.

Note that this transition is similar to the Z-Hawkes transition that was presented in [C1]. In the $\alpha_c < \alpha < 2$ phase, the spread grows again linearly with time:

$$\mathbb{E}[S_t] \sim Vt, \quad V := 2\lambda_0^+ \frac{\alpha - \alpha_c}{(1 - \alpha_c)(2 - \alpha)}.$$

In the $\alpha < \alpha_c$ phase, one finds again $\mathbb{P}[S \geq 2] = (1 - \alpha_c)/(1 - \alpha)$. Interestingly however, simulating numerically equation (D1a) we observe that the spread distribution is asymptotically fat tailed instead of geometric (see figure D1a). Such a power law tail is

also observed in our extended Santa Fe model close to the critical point. We note that the mid-price $P_t = (a_t + b_t)/2$ behaves like a diffusion in the two phases, with an average diffusivity D_P :

$$D_P := \lim_{t \rightarrow \infty} \frac{1}{t} \mathbb{E} \left[\sum_{s < t} (\Delta P_s)^2 \right] = \frac{\lambda_0^- \mathbb{P}[S \geq 2] + \lambda_0^+}{2 - \alpha},$$

One can also show that joint the probability density function $\rho(t, S_t, X_t = \int_0^t \sqrt{2\beta} e^{-\beta(t-s)} dP_s)$ now solves:

$$\begin{aligned} \partial_t \rho = & \left[\lambda_0^+ + \alpha(x - \hat{\beta})^2 \right] \rho(t, S - 1, x - \hat{\beta}) + \left[\lambda_0^+ + \alpha(x + \hat{\beta})^2 \right] \rho(t, S - 1, x + \hat{\beta}) \\ & + \lambda_0^- \left[\rho(t, S + 1, x - \hat{\beta}) + \rho(t, S + 1, x + \hat{\beta}) \right] - 2 \left[\lambda_0^+ + \alpha x^2 \right. \\ & \left. + \lambda_0^- \right] \rho(t, S, x) + \beta \partial_x(x\rho), \end{aligned} \tag{D3}$$

where $\hat{\beta} = \sqrt{\beta/2}$ and with the boundary condition $\forall S < 1, \rho(t, S, x) = 0$. The proof of such results uses the same techniques as in the previous appendix but is slightly more complex. First of all, there exists four martingales $M_t^{\pm, b}, M_t^{\pm, a}$ and $M_t^{\pm, a}$ such that:

$$db_t^\pm = \lambda_t^\pm dt/2 + dM_t^{\pm, b}, da_t^\pm = \lambda_t^\pm dt/2 + dM_t^{\pm, a} \tag{D4}$$

Note that we have:

$$\begin{aligned} dS_t &= da_t^+ + db_t^+ - da_t^- + db_t^- \\ dP_t &= (da_t^+ - da_t^- + db_t^- - db_t^+) / 2 \\ d[P]_t &= (dP_t)^2 = (da_t^+ + da_t^- + db_t^- + db_t^+) / 4. \end{aligned} \tag{D5}$$

We then use equation (D1(a)) in a more general framework:

$$\lambda_t^+ = \lambda_0^+ + \int_0^t \int_0^t K(t-s, t-u) dP_s dP_u, \tag{D6}$$

where K is symmetric. Calling $\alpha = \text{Tr } K = \int_0^{+\infty} K(t, t) dt$, one can rewrite:

$$\begin{aligned} \lambda^+ = \lambda_0^+ + \int_0^t K(t-s, t-s) d[P]_s + M_t^P &= \lambda_0^+ + (\phi * (\lambda^+ + \lambda^-))_t \\ &+ \frac{1}{2} (\phi * (M^+ + M^-))_t + M_t^P, \end{aligned} \tag{D7}$$

where $\phi(t) = K(t, t)/2$, $M_t^\pm = M_t^{\pm, a} + M_t^{\pm, b}$ and $M_t^P = \int_0^t \left(\int_0^{s-} K(t-s, t-u) dP_u \right) dP_s$, that is a martingale. Introducing the resolvent $\phi_R = \sum_{n \geq 1} \phi^{*n}$ and the martingale:

$$dM_t = [(\phi_R * (M^+ + M^-))_t / 2 + ((\delta + \phi_R) * M^P)_t] dt,$$

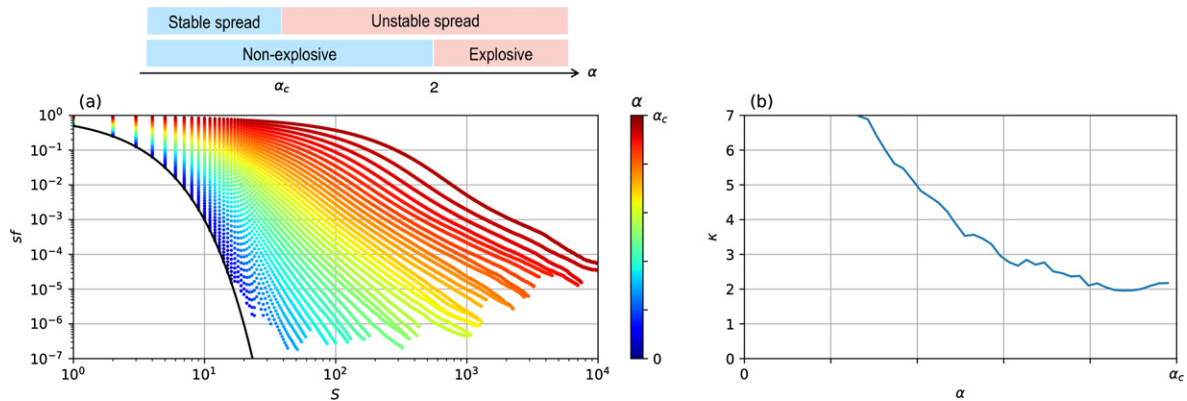


Figure D1. Properties of the spread for $\alpha < \alpha_c$ and $\beta = 1$, $\lambda_0^+ = 1$, $\lambda_0^- = 0.5$. (a) Survival function (sf) of the spread. The black curve corresponds to the theoretical equilibrium distribution when $\alpha = 0$. For large α , the survival function decays asymptotically as a power-law $S^{-\kappa}$. (b) Tail exponent κ of the survival function, as a function of α . κ appears to saturate around 2 when $\alpha \rightarrow \alpha_c$.

we solve the equation:

$$\lambda_t^+ = \left(1 + \int_0^t \phi_R(s) ds\right) \lambda_0^+ + (\phi_R * \lambda^-)_t + dM_t, \tag{D8}$$

and deduce the dynamics of the spread:

$$S_t = S_0 + \int_0^t \left[\lambda_0^+ \left(1 + \int_0^s \phi_R(u) du\right) - \lambda_0^- \mathbb{1}_{\{S_s \geq 2\}} \times \left(1 - \int_0^{t-s} \phi_R(u) du\right) \right] ds + M_t. \tag{D9}$$

This gives the condition of stability: if $\alpha < \alpha_c$ then $\mathbb{P}[S \geq 2] = (1 - \alpha_c)/(1 - \alpha)$. Then we get the diffusivity of the price:

$$\lim_{t \rightarrow +\infty} \frac{1}{t} \mathbb{E}[[P]_t] = \lim_{t \rightarrow +\infty} \frac{1}{2} (\mathbb{E}[\lambda_t^+ + \lambda_t^-]) = \frac{\lambda_0^+ + \lambda_0^- \mathbb{P}[S \geq 2]}{2(1 - \|\phi\|)}$$

One can check that $\|\phi\| = \alpha/2$.

References

- [1] Armand J, Lefevre A, Grunberg D and Bouchaud J-P 2008 Stock price jumps: news and volume play a minor role *Wilmott Magazine* (London: Wiley)
- [2] Cutler D M, Poterba J M and Summers L H 1989 What moves stock prices? *J. Portfolio Manag.* **15** 4–12
- [3] Fair R 2002 Events that shook the market *J. Bus.* **75** 713–32
- [4] Kirilenko A, Kyle A S, Samadi M and Tuzun T 2017 The flash crash: high-frequency trading in an electronic market *J. Finance* **72** 967–98
- [5] Zweig J 2010 Back to the future: lessons from the forgotten ‘flash crash’ of 1962 *The Wall Street Journal*
- [6] Bouchaud J-P, Bonart J, Donier J and Gould M 2018 *Trades, Quotes and Prices: Financial Markets under the Microscope* (Cambridge: Cambridge University Press)

- [7] Maria Calcagnile L, Bormetti G, Treccani M, Marmi S and Lillo F 2018 Collective synchronization and high frequency systemic instabilities in financial markets *Quant. Finance* **18** 237–47
- [8] Bormetti G, Maria Calcagnile L, Treccani M, Corsi F, Marmi S and Lillo F 2015 Modelling systemic price cojumps with Hawkes factor models *Quant. Finance* **15** 1137–56
- [9] Gopikrishnan P, Meyer M, Amaral LA N and Stanley H E 1998 Inverse cubic law for the distribution of stock price variations *Eur. Phys. J. B* **3** 139–40
- [10] Glosten L R and Milgrom P R 1985 Bid, ask and transaction prices in a specialist market with heterogeneously informed traders *J. Financ. Econ.* **14** 71–100
- [11] Dall’Amico L, Antoine F, Bouchaud J-P and Benzaquen M 2019 How does latent liquidity get revealed in the limit order book? *J. Stat. Mech.* **013404**
- [12] Achab M, Bacry E, Muzy J-F and Rambaldi M 2018 Analysis of order book flows using a non-parametric estimation of the branching ratio matrix *Quant. Finance* **18** 199–212
- [13] Rambaldi M, Bacry E and Lillo F 2017 The role of volume in order book dynamics: a multivariate Hawkes process analysis *Quant. Finance* **17** 999–1020
- [14] Daniels M G, Farmer J D, Gillemot L, Iori G and Smith E 2003 Quantitative model of price diffusion and market friction based on trading as a mechanistic random process *Phys. Rev. Lett.* **90** 108102
- [15] Smith E, Farmer J D, Gillemot L and Krishnamurthy S *et al* 2003 Statistical theory of the continuous double auction *Quant. Finance* **3** 481–514
- [16] Farmer J D, Patelli P and Zovko I I 2005 The predictive power of zero intelligence in financial markets *Proc. Natl Acad. Sci.* **102** 2254–9
- [17] Hawkes A G 1971 Spectra of some self-exciting and mutually exciting point processes *Biometrika* **58** 83–90
- [18] Huang W, Lehalle C-A and Rosenbaum M 2015 Simulating and analyzing order book data: the queue-reactive model *J. Am. Stat. Assoc.* **110** 107–22
- [19] Wu P, Rambaldi M, Muzy J-F and Bacry E 2019 Queue-reactive Hawkes models for the order flow (arXiv:1901.08938)
- [20] Morariu-Patrichi M and Mikko S P 2018 State-dependent Hawkes processes and their application to limit order book modelling (arXiv:1809.08060)
- [21] Blanc P, Donier J and Bouchaud J-P 2017 Quadratic Hawkes processes for financial prices *Quant. Finance* **17** 171–88
- [22] Bates D S 2019 How crashes develop: intradaily volatility and crash evolution *J. Finance* **74** 193–238
- [23] Bacry E, Mastromatteo I and Muzy J-F 2015 Hawkes processes in finance *Market Microstructure and Liquidity* **1** 1550005
- [24] Bacry E, Dayri K and Muzy J-F 2012 Non-parametric kernel estimation for symmetric Hawkes processes. application to high frequency financial data *Eur. Phys. J. B* **85** 157
- [25] Cont R and De Larrard A 2012 Order book dynamics in liquid markets: limit theorems and diffusion approximations (arXiv:1202.6412)
- [26] Madhavan A, Richardson M and Roomans M 1997 Why do security prices change? A transaction-level analysis of nyse stocks *Rev. Financ. Stud.* **10** 1035–64
- [27] Wyart M, Bouchaud J-P, Kockelkoren J, Potters M and Vettorazzo M 2008 Relation between bid–ask spread, impact and volatility in order-driven markets *Quant. Finance* **8** 41–57
- [28] Tóth B, Lemperiere Y, Deremble C, De Lataillade J, Kockelkoren J and Bouchaud J-P 2011 Anomalous price impact and the critical nature of liquidity in financial markets *Phys. Rev. X* **1** 021006
- [29] Mastromatteo I, Toth B and Bouchaud J-P 2014 Agent-based models for latent liquidity and concave price impact *Phys. Rev. E* **89** 042805
- [30] Toke I M 2011 An introduction to Hawkes processes with applications to finance *Lectures Notes from Ecole Centrale Paris, BNP Paribas Chair of Quantitative Finance* **193**
- [31] Redner S 2001 *A Guide to First-Passage Processes* (Cambridge: Cambridge University Press)
- [32] Bak P, Tang C and Wiesenfeld K 1987 Self-organized criticality: an explanation of $1/f$ noise *Phys. Rev. Lett.* **59** 381
- [33] Jaisson T and Rosenbaum M *et al* 2015 Limit theorems for nearly unstable Hawkes processes *Ann. Appl. Probab.* **25** 600–31
- [34] Peter H, Talkner P and Borkovec M 1990 Reaction-rate theory: fifty years after kramers *Rev. Mod. Phys.* **62** 251
- [35] Godreche C, Bouchaud J P and Mézard M 1995 Entropy barriers and slow relaxation in some random walk models *J. Phys. A: Math. Gen.* **28** L603
- [36] Challet D and Zhang Y-C 1998 On the minority game: analytical and numerical studies *Physica A* **256** 514–32

- [37] Giardina I and Bouchaud J-P 2003 Bubbles, crashes and intermittency in agent based market models *Eur. Phys. J. B* **31** 421–37
- [38] Alfi V, Cristelli M, Pietronero L and Zaccaria A 2009 Minimal agent based model for financial markets I *Eur. Phys. J. B* **67** 385–97
- [39] Emanuele Biondo A, Pluchino A and Rapisarda A 2015 Modeling financial markets by self-organized criticality *Phys. Rev. E* **92** 042814
- [40] Bochud T and Challet D 2007 Optimal approximations of power laws with exponentials: application to volatility models with long memory *Quant. Finance* **7** 585–9
- [41] Mardia K V, Kent J T and Bibby J M 1979 Multivariate analysis *Probability and Mathematical Statistics* (New York: Academic)