

Critical reflexivity in financial markets: a Hawkes process analysis

Stephen J. Hardiman,* Nicolas Bercot, and Jean-Philippe Bouchaud
Capital Fund Management,
23 rue de l'Université, 75007 Paris, France

(Dated: February 7, 2013)

We model the arrival of mid-price changes in the E-Mini S&P futures contract as a self-exciting Hawkes process. Using several estimation methods, we find that the Hawkes kernel is power-law with a decay exponent close to -1.15 at short times, less than $\approx 10^3$ seconds, and crosses over to a second power-law regime with a larger decay exponent ≈ -1.45 for longer times scales in the range $[10^3, 10^6]$ seconds. More importantly, we find that the Hawkes kernel integrates to unity independently of the analysed period, from 1998 to 2011. This suggests that markets are and have always been close to criticality, challenging a recent study which indicates self-reflexivity (endogeneity) has increased in recent years as a result of increased automation of trading. However, we note that the scale over which market events are correlated has decreased steadily over time with the emergence of higher frequency trading.

I. INTRODUCTION

The “excess volatility puzzle” is the name coined to describe Shiller’s [1] and LeRoy and Porter’s [2] observation that *stock prices move too much to be justified by changes in subsequent dividends*, in contradiction with the Efficient Market Hypothesis. Although hotly debated ever since among theorists (see [3] for a short review), many empirical studies have confirmed that a large fraction of the volatility of asset prices cannot be explained by changes in their fundamental value. For example, most large price jumps seem to be endogenously generated rather than related to the arrival of news in the market place [4–6]. Although excess volatility, crises and crashes seem to be as old as financial markets themselves [7, 8], the debate has recently been stoked by the rise of “high frequency trading” which is blamed for having made markets more volatile and prone to large magnitude events such as flash-crashes without the need for exogenous triggers. Another closely related debate concerns the financial transaction tax (FTT) as a tool to curb the excessive volatility of markets. (Most academic studies however suggest that a FTT would in fact *increase* volatility [9].)

Keynes [10], Soros [11] and others (see e.g. [12, 13]) have argued that a substantial part of the market activity is of endogenous, “self-reflexive” origin. However, real progress has been hobbled by the lack of statistical tools that would enable one to define and measure precisely the degree of self-reflexivity in financial markets. Recently, Filimonov and Sornette (FS) [14] have applied the so-called self-exciting Hawkes process formalism to market events, with the aim of quantifying precisely endogeneity and exogeneity in market activity. Hawkes processes were originally introduced to model the clustered occurrences of earthquakes. Intuitively, earthquakes re-

lax accumulated stresses between tectonic plates – but by doing so other regions are brought closer to their yield stress, leading to an inter-twined pattern of foreshocks and aftershocks. Hawkes processes describe the arrival of these seismic events as a Poisson point process but with a history dependent intensity, that increases when a seismic event just occurred. The same picture lends itself naturally to the arrival of events in the market place. Not surprisingly, the Hawkes formalism has generated a flurry of activity as a means of modelling volatility and market dynamics [15–18].

The results of Filimonov and Sornette are twofold. First, their Hawkes process analysis over the period 1998–2010 on the E-mini S&P futures contract suggests that the degree of self-reflexivity has increased steadily in the last decade, an effect they attribute to the increased deployment of high-frequency and algorithmic trading. Second, when calibrated over much shorter time intervals (10 minutes), the Hawkes process analysis is found to be able to detect precursors of the flash-crash that happened on May 6th, 2010. Devising such early warning tools is evidently beneficial (in particular for market regulators). The work of FS is, in this respect, in a similar vein to that of Easley, O’Hara and Lopez de Prado who developed their “VPIN” indicator [19].

The aim of the present paper is to revisit the first claim of FS that concerns the secular increase of self-reflexivity in financial markets. (Although for the moment we reserve judgement on the second claim, which is of course of high interest as well). Our results are markedly different from those of FS. We find that the degree of self-reflexivity has been extremely stable over time and in fact equal to the critical value beyond which the calibrated Hawkes process becomes ill-defined. We also find, in agreement with the recent work of Bacry et al. [17], that the “influence kernel” of the Hawkes process (see below for a precise definition) decays as a power-law in time, with two different regimes and a crossover time of around five minutes. This is the main explanation for the discrepancy between our own analysis and that of FS,

*Electronic address: stephen.hardiman@cfm.fr

who chose an exponentially decaying influence kernel.

We consider our result very plausible as they fit well with results previously reported in the literature. For example, as we show in the next section, the power-law decay of the influence kernel is related to the long-range memory of the market activity, which is one of the best documented stylized facts of financial markets. Furthermore, that markets operate in the vicinity of a critical point has been repeatedly emphasized in the econophysics literature, see e.g. [6, 13, 20–25]. Within a Hawkes process description, criticality is related to both long-range memory and the prevalence of endogeneity over exogeneity in a precise manner. Finally, many studies conclude that the statistics of financial markets have been remarkably stable over time and assets (see e.g. [13, 26] and refs. therein) – a result that would have been hard to reconcile with the strong increase of self-reflexivity reported by FS.

II. HAWKES PROCESSES: DEFINITION AND SOME USEFUL RESULTS

A. The general case

A one-dimensional Hawkes process is a non-homogeneous Poisson process $N(t)$ with an intensity $\lambda(t) = E \left[\frac{dN(t)}{dt} \right]$ described by a constant base intensity μ and a self-exciting term that is a weighted sum over previous events:

$$\lambda(t) = \mu + \int_{-\infty}^t \phi(t-s) dN(s) \quad (1)$$

The “influence kernel” $\phi(\tau) \geq 0$ describes the effect on the instantaneous event rate $\lambda(t)$ of a past event that took place at time $t - \tau$. Assuming stationarity, and taking the expectation of both sides of Eq. (1) allows one to compute the *average* intensity Λ , which is the solution of:

$$\Lambda = \mu + \Lambda \int_0^{\infty} \phi(\tau) d\tau. \quad (2)$$

Defining $n \equiv \int_0^{\infty} \phi(\tau) d\tau$, the above equation only has a meaningful solution provided $n < 1$, in which case $\Lambda = \mu/(1-n) \geq \mu$. When $n = 0$, the Hawkes process is a homogeneous Poisson process with intensity μ . For $n > 1$, the intensity of the process defined by Eq. (1) explodes in finite time. The critical case $n = 1$ is special and will be discussed below. Hawkes processes can be mapped into the well known branching processes [27]: exogenous “mother” events occur with intensity μ , and give rise to endogenous “child” events, which themselves give birth to more child events. The quantity n is the branching ratio of the process, i.e. the expected number of additional events that any single event triggers. As is well known, branching processes become critical when

n reaches unity, and give rise to infinite families with strictly positive probability when $n > 1$.

The functional form of $\phi(\tau)$ further defines the type of Hawkes process. A popular form, chosen by FS in their analysis of market data, is the exponential kernel:

$$\phi(\tau) = \alpha e^{-\beta\tau} \longrightarrow n = \frac{\alpha}{\beta}. \quad (3)$$

Note that β^{-1} defines the memory time of the Hawkes process, beyond which past events have a negligible influence on the instantaneous event rate. However, other forms are possible too. For example, a power-law kernel with a short-time cutoff τ_0 :

$$\phi(\tau) = \Theta(\tau - \tau_0) \varphi_0 \frac{\tau_0^\epsilon}{\tau^{1+\epsilon}}, \quad \epsilon > 0, \quad (4)$$

where $\Theta(\cdot)$ is the Heaviside function. In this case, the branching ratio is $n = \varphi_0/\epsilon$.

A quantity that can easily be measured empirically is the auto-covariance of the process, defined as:

$$\nu(\tau) = \frac{1}{h^2} E \left[\left(\int_0^h dN(s) - \Lambda h \right) \left(\int_\tau^{\tau+h} dN(s) - \Lambda h \right) \right], \quad (5)$$

where h is some small time interval.

Bacry et al. showed that in the small frequency limit, the Fourier transforms of ν and ϕ are related through the following equation [17]:

$$\hat{\nu}(\omega) = \Lambda \left| \frac{1}{1 - \hat{\phi}(\omega)} \right|^2. \quad (6)$$

Letting $\omega = 0$ above leads to:

$$\int_{-\infty}^{\infty} \nu(\tau) d\tau = \Lambda \left(\frac{1}{1-n} \right)^2 \longrightarrow n = 1 - \sqrt{\frac{\Lambda}{2 \int_0^{\infty} \nu(\tau) d\tau}}. \quad (7)$$

Now the autocorrelation function can either decay *faster* than τ^{-1} for large τ , corresponding to a short-memory process for which $\int_0^{\infty} \nu(\tau) d\tau < \infty$, or *slower* than τ^{-1} corresponding to a long-memory process for which now $\int_0^{\infty} \nu(\tau) d\tau$ is infinite. From the above equation, one therefore concludes that *long-memory of the autocorrelation necessarily implies criticality*, i.e. $n = 1$.

B. The critical case

Let us study this critical case in more detail. Assuming that the Hawkes kernel is a power-law as above, with an exponent $0 < \epsilon < 1$, one finds that the small $\omega\tau_0$ expansion of $\hat{\phi}(\omega)$ reads:

$$\hat{\phi}(\omega) = n - A(\omega\tau_0)^\epsilon \quad (8)$$

where A is a complex number, related to φ_0 and ϵ . (Note that when $\epsilon > 1$, the first term of the expansion of

$\hat{\phi}$ becomes regular, i.e. $\propto \omega$.) Similarly, for an auto-covariance $\nu(\tau)$ which decays asymptotically as a power-law $\propto \tau^{-\alpha}$ for some $0 < \alpha < 1$, one finds:

$$\hat{\nu}(\omega) = B\omega^{\alpha-1} \quad (9)$$

with some real constant B . Note that $\alpha < 1$ corresponds to long-range memory, and $\hat{\nu}(\omega \rightarrow 0) \rightarrow \infty$, leading to $n = 1$ as alluded to above. Inserting these expressions into Eq. (6), we obtain, for $\omega \rightarrow 0$,

$$B\omega^{\alpha-1} \approx |A|^{-2}\omega^{-2\epsilon}, \quad (10)$$

leading to:

$$\alpha = 1 - 2\epsilon, \quad B = |A|^{-2} \quad (11)$$

The above relation implies that ϵ must in fact lie in the restricted interval $]0, \frac{1}{2}[$ for α to be in the range $]0, 1[$, as assumed. But since $\Lambda = \mu/(1-n)$, the limiting case $n = 1$ only makes sense if simultaneously $\mu = 0$, i.e. the total number of exogenous events grows sub-linearly with the sample length. Interestingly, the above construction has been made rigorous in [28] in their study of ‘‘Hawkes processes without ancestors’’.

Another way to think about this limit, which may be more appropriate for financial applications, is to assume that $n = 1 - \delta$, $\delta \ll 1$. The autocorrelation function then decays as a power-law $\propto \tau^{-\alpha}$ up to a (long) crossover time $\tau^* \sim \delta^{-1/2\epsilon}$, beyond which it decays much faster, in order to make $\nu(\tau)$ integrable.

What happens when $\epsilon > \frac{1}{2}$ but $\varphi_0 = \epsilon$ such that $n = 1$ (see Eq. (4))? Naively, the above analysis predicts $\alpha < 0$, i.e. an autocorrelation that increases without bound as $\tau \rightarrow \infty$. This is of course impossible, and signals that in this case, the critical Hawkes process is non stationary, to wit, the average instantaneous rate increases with time. This means that an initial time is needed to make the process well defined.

The reason we have focused so much attention on this critical case is that, as we will show below, the empirical data clearly suggests that the relevant Hawkes process is indeed critical and stationary, i.e. precisely the Hawkes processes without ancestors considered by Brémaud & Massoulié [28].

III. POWER-LAW VERSUS EXPONENTIAL HAWKES KERNEL

As mentioned above, FS estimated the parameters of a Hawkes processes on the midpoint changes of the E-mini S&P Futures contract in the period 1998-2010 assuming an exponential kernel. However, this choice can only be justified if the dynamics of the market can be characterized by a single time scale, β^{-1} . This is at variance with many empirical studies that clearly elicit a scale-free, long memory of market activity. This strongly suggest the use of a power-law influence kernel. As a matter

of fact, Bacry et al. [17] have developed a non-parametric estimation technique that yields a power-law kernel over several decades for incoming market orders on both the Bund and the Dax futures contracts. They find that the exponent ϵ (defined by Eq. (4)) is small.

We will revisit this non-parametric determination of the kernel in Sect. IV, but for now we wish to repeat the analysis of [14] with the hypothesis of a simple power-law kernel, as given by Eq. (4). To allow for changes to the high frequency end of the Hawkes kernel over time resulting from improved temporal resolution of market events and reduced trading latency, we consider a power-law kernel with smooth exponential cut-off for short lags. Furthermore, to expedite the calculation of the likelihood of the Hawkes process, we approximate the power-law as a sum of exponential functions with power-law weights ¹.

$$\phi(\tau|\mu, n, \epsilon, \tau_0) = \frac{n}{Z} \left(\sum_{i=0}^{M-1} \left(\frac{1}{\xi_i} \right)^{1+\epsilon} e^{-\frac{\tau}{\xi_i}} - S e^{-\frac{\tau}{\xi-1}} \right) \quad (12)$$

The scales of the exponentials are given by:

$$\xi_i = \tau_0 m^i \quad (13)$$

and S and Z are constants chosen such that $\phi(0) = 0$ and $\int_0^\infty \phi(\tau) d\tau = n$. For the empirical calibration, we have fixed $M = 15$ and $m = 5$. For values of ϵ close to zero as we expect, the resulting function describes closely a power-law form with tail exponent $\tau^{-(1+\epsilon)}$ but the negative exponential term provides a smooth cut-off at short lags.

Given observations of empirical market events (mid-point changes) at times t_1, t_2, \dots, t_n in an interval $[0, T]$ we can fit the model by maximising the log-likelihood [29, 30] over the set of parameters $\theta = (\mu, n, \epsilon, b)$.

$$\log L(t_1, t_2, \dots, t_n | \theta) = - \int_0^T \lambda(t|\theta) dt + \int_0^T \log \lambda(t|\theta) dN(t) \quad (14)$$

The data we use consists of a sequence of times in the period 1998-2011 when the mid-point of the bid and ask price of the E-mini S&P contract was observed to change. Changes in this mid-point may have been the result of executed market orders or limit orders that have been placed or cancelled. Our data is time-stamped to the millisecond, but since sometimes multiple events occur during the same millisecond - the limit of our temporal resolution - we employ a randomisation procedure. We give each event a real number time-stamp distributed randomly and uniformly within the millisecond that it was reported. A similar operation is performed in [14], but within the second resolution that they have.

¹ By approximating with a sum of exponentials we can take advantage of a recurrence relation which reduces the time complexity of the log-likelihood calculation in Eq. (14) from $\mathcal{O}(N^2)$ to $\mathcal{O}(N)$

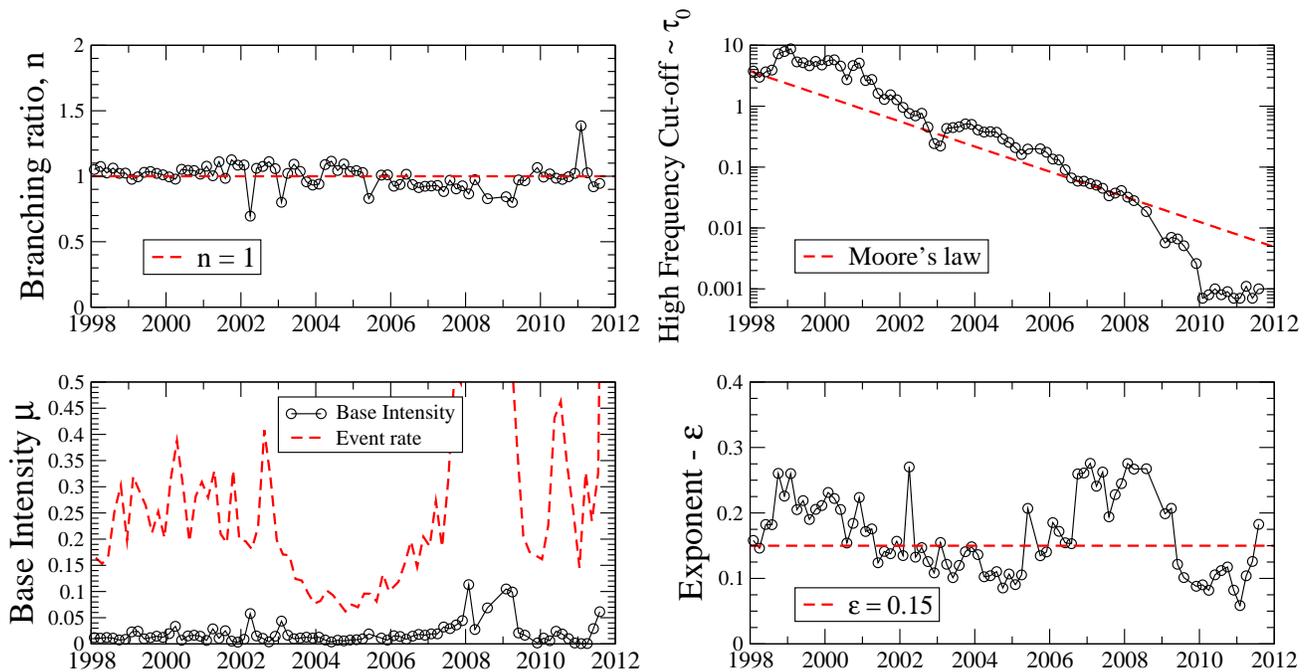


FIG. 1: Results for the MLE fitting of the Hawkes process with power-law kernel to mid-point changes in the E-mini S&P Futures contract. The fit is performed over two-month periods between Jan. 1998 and Dec. 2011. Top-Left: The branching ratio n appears to change little and fluctuates around the critical value $n = 1$ for the 14-year period studied. Top-Right: The high frequency cut-off τ_0 , on the other hand, decreases exponentially over time as a result of the increasing frequency at which events are correlated, this trend is not too far out of line with Moore's law, the oft-cited observation that the processing speed of computers doubles every 18 months. Bottom-Left: The base intensity μ (in events per second) that has remained roughly constant over time, apart from spikes during crisis periods. Bottom-Right: The exponent ϵ entering the power-law kernel $\phi(\tau)$, that has also remained roughly constant around 0.15 – see below.

In fitting the model we only consider market events occurring during Regular Trading Hours (09:30 to 16:15). Furthermore, we take account of the ‘U-shape’ intra-day activity profile. Without appropriate detrending, the fitting procedure would interpret the slow variation in event rate during the day as arising from the shape of the Hawkes kernel. To this end, we consider a detrended Hawkes process model

$$\lambda(t|\theta) = \frac{1}{w(t)} \left[\mu + \int_0^t w(s)\phi(t-s|\theta)dN(s) \right] \quad (15)$$

where $w(t)$ is a periodic function defined over the trading day that dictates the weight to be appropriated to events occurring at different hours. The aim is to give less weight to events occurring during the morning and late afternoon when high activity is to be expected. The weights are chosen to be the reciprocal of the empirical average daily event rate at that moment of the day $w(t) = \frac{1}{R(t)}$ normalised such that $E[R(t)] = E[w(t)] = 1$.

The model parameters as a function of time are calculated between 1998 and 2011 on non overlapping two-month periods using an estimate of the function $R(t)$ made for each year. All regular trading hours in each

two-month period are concatenated to form a single continuous point-process for which the log-likelihood of Eq. (14) and Eq. (15) is maximised with respect to the four parameters $\theta = \{\mu, n, \epsilon, \tau_0\}$. The results are shown in Fig. 1.

Our key observation is that the branching ratio appears to change little over the course of the fourteen years studied but fluctuates about the critical value $n_c = 1$ corresponding to the onset of stationarity. This means, that within the Hawkes framework, the majority of events observed in the market can be attributed to the long-range influence kernel, both now and in 1998. However, we see that the short lag cut-off of the kernel has decreased exponentially over time. This can be attributed to the steady emergence of higher frequency trading despite little change in the event rate. It may surprise the reader that the event rate has not increased significantly in the 14-year period (the event rate, for example in 2009 is comparable with that in 1999) despite the growth of the electronic futures market, but understand that the events studied are mid-point changes, not trades or order-book events and as such it is a better proxy for volatility than transaction quantity or volume.

Note that for the stationary critical Hawkes process with $n = 1$ described in Sect. II B, the base intensity μ must be zero. The small “excess” base intensity observed in our results can be explained as arising from the contribution of the part of the power-law kernel that extends beyond the two-month window on which the fit is performed. The integral of this remainder is approximately $\epsilon \left(\frac{\tau_0}{T}\right)^\epsilon$ where $T \approx 2$ months. Indeed, since each two month period is fitted independently, without prior event history, there must be some non-zero base intensity to explain the observation of any event at all!

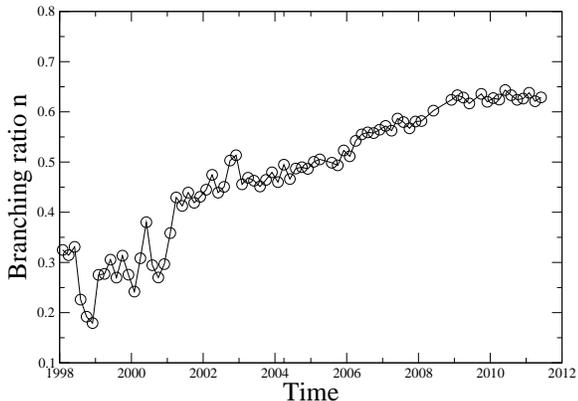


FIG. 2: The analysis of [14] on a critical power-law Hawkes processes. We generate synthetic power-law Hawkes processes with parameters $\{\mu, \epsilon, n\} = \{0.02, 0.15, 1.0\}$ but with a cut-off τ_0 drawn from Fig. 1 and attempt to fit these generated processes to a model with an exponential kernel. Furthermore we also impose the intra-second randomisation procedure employed in [14]. Despite the fact that the generated datasets represent critical Hawkes processes with $n = 1.0$, the analysis (incorrectly) reveals a sub-critical branching ratio which has been steadily increasing since 1998, as reported in [14], Fig. 4-D.

Why is our result $n \approx 1$ over the whole period so different from the one presented in FS [14]? To reconcile our result with theirs, we propose that a sub-critical branching ratio $n < 1$ which increases over time can arise from the inadequate choice of an exponential short-memory model to fit a critical, power-law Hawkes process with a time dependent high-frequency cut-off τ_0 . To this end, we have performed MLE fits of the exponential kernel model to synthetic data that we generated using the power-law kernel hypothesis. We fixed $\{\mu, \epsilon, n\} = \{0.02, 0.15, 1.0\}$ (values which are approximately representative of the full fourteen years studied) but allowed the cut-off parameter τ_0 to vary according to the MLE results of Fig. 1. For each two-month time period, we generated an ensemble of 100 power-law Hawkes processes from which, following FS’s procedure, we drew a 30-minute window. After randomising the time-stamps within the second, as in [14], the window was fit to the exponential kernel model

in the same way as FS. The resulting branching ratio $\frac{\alpha}{\beta}$ was then averaged over all 100 windows in each ensemble. The results of this analysis are shown in Fig. 2, and resemble much the data presented in FS (see their Fig. 4, panel D). Despite the fact that the dataset is generated from a critical Hawkes process with $n = 1.0$, the analysis with the exponential kernel model on 30-minute windows reveals a sub-critical branching ratio which increases steadily over time.

IV. NON-PARAMETRIC ESTIMATION OF THE HAWKES KERNEL

To complement and back up the results presented in Section III, we also estimate the shape of the Hawkes kernel for several times in the 14-year period using non-parametric means. In [17], the authors present a methodology for estimating the kernel of multivariate Hawkes process based on Eq. (6) and the determination of the empirical auto-covariance of the event rate, Eq. (5). We have repeated their analysis in the monivariate case for our data set, choosing $h = 0.001s$ for 2009 (the maximum temporal resolution of our data), $h = 0.01s$ for 2005 and $h = 0.1s$ for 1998. Following [17] the auto-covariance is also discretely sampled at a rate $\Delta = h$ and the Fourier transform is performed in practice with DFT. The resulting kernel estimate is subsequently averaged with logarithmic bins to reduce noise for presentation.

To minimise the effects of intra-day seasonality, event rates are again detrended by dividing by the daily event rate as averaged for 5-minute bins over each year studied. Furthermore the empirical auto-covariance is estimated on 45-minute windows, during which we can expect the event rate to be relatively stationary anyway. The auto-covariance estimate used to produce the kernel in the range 1 ms – 1000s in Fig. 3 is based on an average over all such 45-minute windows in a full year (1998, 2006 or 2009).

The non-parametric procedure confirms the MLE results and also captures the change in the high-frequency cut-off between the two years. Note that the strong oscillations in the 1998 kernel estimate for $\tau < 1s$ are not random noise, but the result of a spurious correlation in how the data (although time-stamped to the millisecond) is collected or aggregated. However, the MLE estimation of the kernel shape seems robust to these data anomalies.

As an additional check of the goodness of fit of the power-law Hawkes process, we consider the residual process. We re-express the point process in terms of a transformed time $t \rightarrow t^*$ which is the integral of the conditional intensity.

$$t^* = \int_{-\infty}^t \lambda(s) ds \quad (16)$$

where $\lambda(\cdot)$ is calculated using the model on events prior to time t . If the Hawkes process with our chosen kernel,

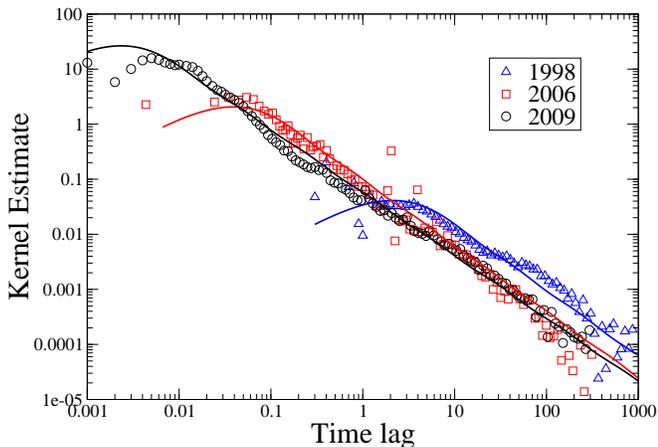


FIG. 3: Non-parametric kernel estimation of the Hawkes kernel in 1998, 2006 and 2009. The power-law form with decay exponent $\epsilon \approx 0.15$ is reproduced in all three years. However the high frequency cut-off of the kernels are considerably different. The kernel estimates match well with the MLE fits (solid lines) using the average parameter estimates for each year as taken from Fig. 1

$\phi(\epsilon, n, \tau_0)$, is indeed a good description of the data, then the transformed process $N(t^*)$ should be Poisson with unit intensity. We assess this fit by analysis of the arrival times between events in the transformed time t^* , which if unit variance Poisson should be described by the law $P(\Delta t^*) = \exp(-\Delta t^*)$.

The results of this analysis for periods in 1998 and 2009 are given in Fig. 4. The PDFs presented are produced from the inter-arrival times of the residual process for all points in the given two month periods. The kernel parameters used in the analysis are extracted from the MLE results of Fig. 1. The Kolmogorov-Smirnov distance (i.e. the maximum deviation between the empirical CDF and the Poisson expectation) is 0.011 in 1998 and 0.031 in 2009.

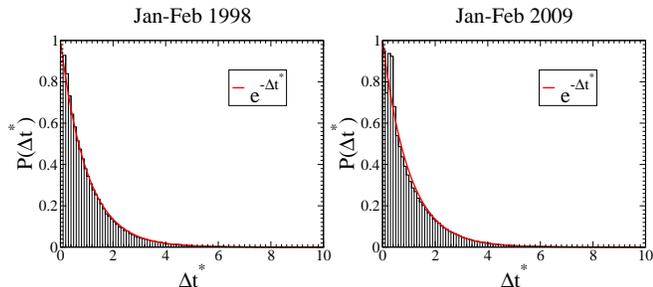


FIG. 4: The empirical distribution for the inter-arrival times of events in the transformed time of Eq. (16). The parameters for the power-law kernel used are extracted from the MLE results of Fig. 1

V. THE SHAPE OF THE HAWKES KERNEL AT LARGE LAGS

The results in the previous section demonstrate that a Hawkes process described by a power-law kernel with an exponent $\epsilon \approx 0.15$ is indeed a good fit to the data in the millisecond to 5-minute region. We will now show that there is a marked regime change in the nature of the correlations for larger time horizons. To see this, we use Detrended Fluctuation Analysis (DFA) [31], a commonly used tool for detecting long-range correlations in the presence of trends and non-stationarity. More precisely, we compute:

$$F(L) = \sqrt{\frac{1}{T} \sum_{t=1}^T (N(t) - f_L(t))^2} \quad (17)$$

where $N(t)$ is the empirical event counting process formed by concatenating all regular trading hour periods in 1998-2011. $N(t)$ is discretised into T 0.1 second bins and $f_L(t)$ is a piecewise function formed by the linear regression to the points $N(t)$ in non-overlapping windows of length L which cover the period studied. If $N(t)$ behaved as a Brownian motion with uncorrelated fluctuations, as for a simple Poisson process, we would expect $F(L) \sim L^H$ with a Hurst exponent $H = \frac{1}{2}$. In the presence of long-ranged correlations induced by a power-law Hawkes kernel, one finds:

$$H = \frac{1}{2} + \epsilon, \quad (0 < \epsilon < \frac{1}{2}) \quad (18)$$

Note that this result is not intuitive, since the larger the value of ϵ , the faster the decay of the Hawkes kernel, but the stronger the deviation from the Poisson result $H = \frac{1}{2}$. This paradoxical behaviour was in fact already apparent in the relation $\alpha = 1 - 2\epsilon$ derived in Sect. II B above.

Fig. 5 clearly shows two regions. For small time windows ($L < 1000s$), the Hurst exponent is found to be $H \approx 0.63$, consistent with the observation of $\epsilon \approx 0.15$ in the previous section. For longer time windows ($L > 1000s$), the diffusion of the integrated event rate becomes more super-diffusive and leads to a stronger Hurst exponent of $H \approx 0.95$. Such a change to a more super-diffusive regime has been observed before for measures of market volatility [32, 33]. However these previous studies have reported the cross-over at the scale of one day, whereas in our dataset for event rate it is happening much faster (10 - 20 minutes). We have repeated the same analysis for absolute value of five-minute returns and confirmed that the cross-over time is indeed longer in that case, on the scale of a day. We attribute the difference to the fact that the distribution of five-minute returns is well known to have fat-tails, much fatter than the distribution of the number of price changes over five minutes. The extra noise brought about by these tails reduce short-time correlations, and therefore bring the value of H closer to $1/2$.

The critical power-law Hawkes kernel implied by the stronger autocorrelation $H \approx 0.95$ corresponds to a larger decay exponent of $\epsilon \approx 0.45$. To visualise the shape

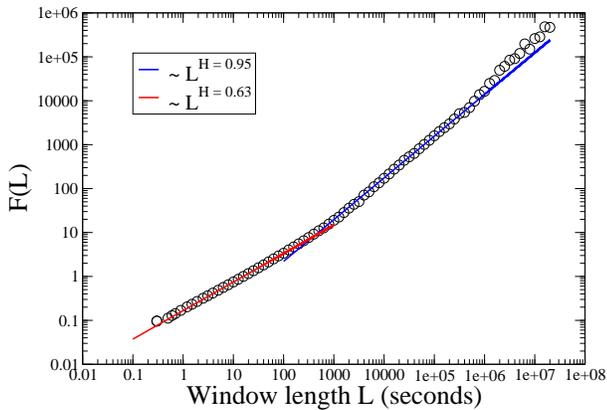


FIG. 5: Detrended Fluctuation Analysis on the counting process for market events in the period 1998-2011. We identify two regimes, one at high frequencies corresponding to a Hurst exponent of $H \approx 0.63$ and one at much lower frequencies corresponding to $H \approx 0.95$.

of the Hawkes kernel at these large time scales, we estimate the auto-covariance of event rate but now using 5-minute binned data over the full 14-year period. To do this, we again treat the event rate dataset as one large continuous time series by concatenating all regular trading periods. The resulting kernel estimate is shown in Fig. 6. Two results can be read from this figure: a) again, the Hawkes kernel indeed integrates to unity, confirming the critical nature of the process implied by the long-memory of event rates (as discussed in Sect. II); b) the exponent $\epsilon = 0.45$ extracted from the DFA analysis correctly describes the time dependence of the Hawkes kernel in the region ~ 1000 seconds $\rightarrow 10^6$ seconds (40 days). The value $\epsilon = 0.45$ leads to $\alpha \approx 0.1$, which is indeed similar to the value of the exponent for the decay of volatility correlations found in the literature (see e.g. [34, 35]).

VI. CONCLUSION

In this study, we have confirmed that Hawkes processes provide a useful framework to describe the high-frequency activity in financial markets. However, contrary to Filimonov & Sornette [14], but in agreement with Bacry et al. [17], we have found that the Hawkes kernel (that describes how past events influence future activity) decays as a power-law, rather than exponentially. The branching ratio we extract from the calibration of our model to the mid-point changes in the E-mini S&P over 14 years is close to critical and has

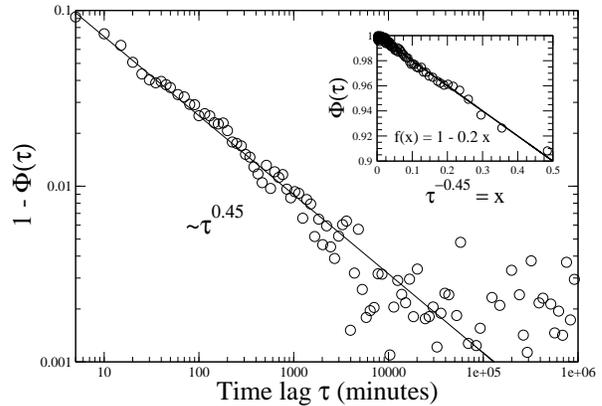


FIG. 6: The result of non-parametric kernel estimation on the auto-covariance of event rate as calculated using 5-minute bins for the full 14-year period. We plot $1 - \Phi(\tau) = 1 - \int_0^\tau \phi(s)ds$ in a log-log representation in the main figure and $\Phi(\tau)$ as a function of $\tau^{-0.45}$ in the inset, clearly demonstrating that $\phi(s)$ integrates to unity and has a power-law tail. This tail exponent $\epsilon = 0.45$ is consistent with the DFA analysis in Fig. 5.

remained constant over time, contrary to the claim of Filimonov & Sornette that self-reflexivity has substantially increased since 1998. Nonetheless, we uncover that the short lag cut-off of the power-law kernel which we fit has significantly decreased over time, with the emergence of higher frequency trading and higher temporal resolution in the market. We have argued that this explains the discrepancy between our results and those of Filimonov & Sornette. Using both a non-parametric method and Detrended Fluctuation Analysis, we have found that the Hawkes kernel must in fact be described by *two* power-laws: one decaying as $\approx \tau^{-1.15}$ for short time scales ($\tau < 1000$ seconds), crossing over to a faster decay $\approx \tau^{-1.45}$ for longer time scales. This long time regime is consistent with the slow decay of volatility correlations reported in the literature.

The picture which emerges from our analysis is one in which bursts of diverging trading activity are just as inevitable now as they were in 1998 (and probably as far back as financial markets existed), but that the time-scale over which they may occur has substantially shortened in the last decade. We have shown that the Hawkes kernel must indeed be critical whenever the event arrival is a long-memory process. Based on this observation and on the results of Bacry et al. [17] on the Bund and the DAX futures, we conjecture that our critical market scenario is common to most actively traded markets. Interestingly, financial markets appear to be well described by a critical Hawkes process “without ancestors”, a theoretical scenario put forth by Brémaud & Massoulié [28]. In non-technical terms, this means that the rate of endoge-

nous events in financial markets is overwhelming compared to that of exogenous (news related) events. A detailed mechanism explaining why markets are precisely poised at criticality is however still lacking.

Acknowledgments

We thank Marc Potters, Julianus Kockelkoren, Emeric Henry, Yves Lempérière, Cyril Deremble and Iacopo

Mastromatteo for many helpful discussions as well as Emmanuel Bacry and Jean-François Muzy.

-
- [1] R. J. Shiller, “Do stock prices move too much to be justified by subsequent changes in dividends?,” *American Economic Review* **71** no. 3, (June, 1981) 421–436.
- [2] S. F. Leroy and R. D. Porter, “The present value relation: Tests based on implied variance bounds,” *Econometrica* **49** no. 3, (May, 1981) 555–574.
- [3] S. F. Leroy, “Excess volatility,” *UCSB working paper* (2005). Available at <http://www.econ.ucsb.edu/~sleroy/downloads/excess.pdf>.
- [4] D. M. Cutler, J. M. Poterba, and L. H. Summers, “What moves stock prices,” *Journal of Portfolio Management* **15** no. 3, (Spring, 1989) 4–12.
- [5] R. C. Fair, “Events that shook the market,” *The Journal of Business* **75** (2002) 713–732.
- [6] A. Joulin, A. Lefevre, D. Grunberg, and J.-P. Bouchaud, “Stock price jumps: news and volume play a minor role,” *Wilmott Magazine* (September–October, 2008) 1–7.
- [7] C. MacKay, *Memoirs of Extraordinary Delusions and the Madness of Crowd*. reprinted by L. C. Page, Boston (1932), 1852.
- [8] C. M. Reinhart and K. S. Rogoff, *This time is different*. Princeton University Press, 2009.
- [9] A. Pomeranets, “Financial transaction taxes: International experiences, issues and feasibility,” *Bank of Canada Review* (Autumn 2012) 3. (and refs. therein).
- [10] J. M. Keynes, *The general Theory of Employment, Interest and Money*, ch. 12. McMillan, London, 1936.
- [11] G. Soros, *The New Paradigm for Financial Markets: The Credit Crisis of 2008 and What it Means*. PublicAffairs, 2008.
- [12] D. Sornette, “Endogenous versus exogenous origins of crises,” in *Extreme Events in Nature and Society*, S. Albeverio, V. Jentsch, and H. Kantz, eds. Springer, Heidelberg, 2005.
- [13] J.-P. Bouchaud, “The endogenous dynamics of markets: Price impact, feedback loops and instabilities,” in *Lessons from the 2008 Crisis*, A. Berd, ed. Risk Books, Incisive Media, London, 2011.
- [14] V. Filimonov and D. Sornette, “Quantifying reflexivity in financial markets: Toward a prediction of flash crashes,” *Physical Review E* **85** (2012) 056108.
- [15] P. Embrechts, T. Liniger, and L. Lin, “Multivariate hawkes processes: an application to financial data,” *Journal of Applied Probability* **48A** (August, 2011) 367–378.
- [16] E. Bacry, S. Delattre, M. Hoffmann, and J.-F. Muzy, “Modeling microstructure noise with mutually exciting point processes,” *Quantitative Finance* **13** (2012) 65–77.
- [17] E. Bacry, K. Dayri, and J.-F. Muzy, “Non-parametric kernel estimation for symmetric hawkes processes. application to high frequency financial data,” *European Physics Journal B* **85** no. 5, (2012) 157.
- [18] G. Bormetti, L. M. Calcagnile, M. Treccani, F. Corsi, S. Marmi, and F. Lillo, “Modelling systemic cojumps with hawkes factor models,” [arXiv:1301.6141](https://arxiv.org/abs/1301.6141) [q-fin.ST].
- [19] D. Easley, M. L. de Prado, and M. O’Hara, “The microstructure of the ‘Flash Crash’: Flow toxicity, liquidity crashes and the probability of informed trading,” *The Journal of Portfolio Management* **37** no. 2, (2011) 118–128.
- [20] D. Challet, M. Marsili, and Y.-C. Zhang, *Minority Game*. Oxford University Press, 2005.
- [21] J.-P. Bouchaud, Y. Gefen, M. Potters, and M. Wyart, “Fluctuations and response in financial markets: The subtle nature of ‘random’ price changes,” *Quantitative Finance* **4** (April, 2012) 176–190.
- [22] F. Caccioli, M. Marsili, and P. Vivo, “Eroding market stability by proliferation of financial instruments,” *European Physics Journal B* **71** (October, 2009) 467–479.
- [23] B. Tóth, Y. Lempérière, C. Deremble, J. de Lataillade, J. Kockelkoren, and J.-P. Bouchaud, “Anomalous price impact and the critical nature of liquidity in financial markets,” *Physical Review X* **1** (2011) 021006.
- [24] F. Patzelt and K. Pawelzik, “Unstable price dynamics as a result of information absorption in speculative markets,” [arXiv:1211.6695](https://arxiv.org/abs/1211.6695) [q-fin.TR].
- [25] D. R. Parisi, D. Sornette, and D. Helbing, “Universality class of balanced flows with bottlenecks: granular flows, pedestrian fluxes and financial price dynamics,” [arXiv:1205.2915](https://arxiv.org/abs/1205.2915) [q-fin.ST].
- [26] G. Zumbach and C. Finger, “A historical perspective on market risks using the DJIA index over one century,” *Wilmott Journal* **2** (August, 2010) 193–206.
- [27] T. E. Harris, *The Theory of Branching Processes*. Dover Publications, 2002.
- [28] P. Brémaud and L. Massoulié, “Hawkes branching point processes without ancestors,” *Journal of Applied Probability* **38** no. 1, (2001) 122–135.
- [29] I. Rubin, “Regular point processes and their detection,” *IEEE Transactions on Information Theory* **IT-18** (September, 1972) 547–557.
- [30] T. Ozaki, “Maximum likelihood estimation of hawkes

- self-exciting point process,” *Annals of the Institute of Statistical Mathematics* **31** (1979) 145–155. Part B.
- [31] C.-K. Peng, S. V. Buldyrev, S. Havlin, M. Simons, H. E. Stanley, and A. L. Goldberger, “Mosaic organization of DNA nucleotides,” *Physical Review E* **49** no. 2, (1994) 1685–1689.
- [32] Y. Liu, P. Gopikrishnan, P. Cizeau, M. Meyer, C.-K. Peng, and H. E. Stanley, “The statistical properties of the volatility of price fluctuations,” *Physical Review E* **60** (1999) 1390–1400.
- [33] F. Wang, K. Yamasaki, S. Havlin, and H. E. Stanley, “Scaling and memory of intraday volatility return intervals in stock markets,” *Physical Review E* **73** (2006) 026117.
- [34] J.-F. Muzy, J. Delour, and E. Bacry, “Modelling fluctuations of financial time series: from cascade process to stochastic volatility,” *European Physics Journal B* **17** (October, 2000) 537–548.
- [35] R. Chicheportiche and J.-P. Bouchaud, “The fine structure of volatility feedback,” [arXiv:1206.2153](https://arxiv.org/abs/1206.2153) [q-fin.ST]. (and refs. therein).