



City Research Online

City St George's, University of London

Citation: Bormetti, G., Calcagnile, L. M., Treccani, M., Corsi, F., Marmi, S. & Lillo, F (2015). Modelling systemic price jumps with Hawkes factor models. Quantitative Finance, 15(7), pp. 1137-1156. doi: 10.1080/14697688.2014.996586

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/19455/>

Link to published version: <https://doi.org/10.1080/14697688.2014.996586>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Modelling systemic cojumps with Hawkes factor models

Giacomo Bormetti^{a,b,*}, Lucio Maria Calcagnile^{a,b,*},
Michele Treccani^{c,b,*}, Fulvio Corsi^a, Stefano Marmi^{a,b}, Fabrizio Lillo^{a,d,e,b}

January 25, 2013

^a *Scuola Normale Superiore, Piazza dei Cavalieri 7, Pisa, 56126, Italy*

^b *QUANTLab¹, via Pietrasantina 123, Pisa, 56122, Italy*

^c *LIST S.p.A., via Pietrasantina 123, Pisa, 56122, Italy*

^d *Dipartimento di Fisica e Chimica, Università degli Studi di Palermo, Viale delle Scienze Ed. 18, Palermo, 90128, Italy*

^e *Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA*

Abstract

Instabilities in the price dynamics of a large number of financial assets are a clear sign of systemic events. By investigating a set of 20 high cap stocks traded at the Italian Stock Exchange, we find that there is a large number of multiple cojumps, i.e. minutes in which a sizable number of stocks displays a discontinuity of the price process. We show that the dynamics of these jumps is not described neither by a multivariate Poisson nor by a multivariate Hawkes model, which are unable to capture simultaneously the time clustering of jumps and the high synchronization of jumps across assets. We introduce a one factor model approach where both the factor and the idiosyncratic jump components are described by a Hawkes process. We introduce a robust calibration scheme which is able to distinguish systemic and idiosyncratic jumps and we show that the model reproduces very well the empirical behaviour of the jumps of the Italian stocks.

Keywords: cojumps, Hawkes processes, systemic shocks, high frequency data

1 Introduction

Modelling the dynamics of security prices is of paramount importance for risk control, derivative pricing, trading, and to understand the behaviour of markets. The most popular class of stochastic models for security prices is certainly the class of jump-diffusion models, where the price is described by a continuous diffusive part plus a jump component, modelling discontinuities in asset prices. In the majority of models the jump component is described by a compound Poisson process. In a multi-asset setting the classical jump diffusion models assumes that the jump component of different assets are independent Poisson processes and the only dependence is possibly in the size of the jumps. An incomplete list of recent studies on jumps following this approach includes theoretical work on nonparametric jump identification [1, 2, 3, 4, 5, 6, 7], as well as empirical analysis [8, 9, 10, 11, 12, 13], and applications to asset

*These authors contributed equally to this work.

¹www.quantlab.it

pricing [14, 15, 16, 17, 18]. General review on jumps are [19] and [20]. However, there has been less research on cojumps (i.e. simultaneous jumps in two or more stock prices) and most are of empirical nature such as [21, 22, 23, 24]; exceptions are the papers of [25, 26] and [27] on cojump estimation and modelling.

Poisson jump processes have the great advantage of being analytically tractable. However one might argue that the jump component of the price process could be described instead by a process where jumps are clustered in time. The same objection emerges when one considers many assets. Markets are nowadays more and more interconnected and it is a priori reasonable to expect that some sort of synchronization between the jumping times of different asset is present.

This synchronization effect had its most spectacular appearance during the May 6, 2010 Flash Crash. According to the SEC-CFTC, the crash started from a rapid price decline in the E-Mini S&P 500 market [28, 29]. However in a very short time the price drop propagated towards ETF, stock indices and their components, and derivatives. For example, the Dow Jones Industrial Average plunged about nine percent, only to recover those losses within minutes. The contagion effect can be extremely rapid in liquid markets [30, 31] and leads to a strongly synchronized discontinuous movement of the price of many assets. This type of systemic events can not be described by a model where prices jump following independent processes.

In this paper we show that indeed the dynamics of jumps of a portfolio of stocks deviates significantly from a collection of independent Poisson processes. The deviation that we observe is twofold. On one side, by considering individual assets, we find evidence of time clustering of jumps, clearly inconsistent with a Poisson process. This means that the intensity of the point process describing jumps depends on the past history of jumps, and a recent jump increases the probability that another jump occurs.

The second deviation from the Poisson model is probably more important, especially in a systemic context. We find a strong evidence of a high level of synchronization between the jumping times of a portfolio of stocks. In other words, we find a large number of instances where several stocks (up to 20, the size of our set) jump at the same time (by using a resolution of a minute). This evidence is absolutely incompatible with the hypothesis of independence of the jump processes across assets.

Taking together these empirical deviations from the independent Poisson model, there is a need of a suitable modelling of the multi-asset jump process and this is the main methodological contribution of this paper. In order to model the time clustering of jumps for individual assets we propose the use of a class of self-exciting point processes, termed Hawkes processes. These processes were introduced more than forty years ago [32], and have been widely employed to model earthquake data [33, 34, 35]. In the last years, Hawkes processes have experienced an increasing popularity in mathematical finance and econometrics. One of the first applications to financial time series is due to [36], and a wide literature review in this context is collected in [37]. Among more recent developments not covered by the latter reference we mention [38] where these processes are applied to the order flow in a continuous double auction market, [39] for the modelling of trades-through orders in a limit order book, [40] where Hawkes processes are used to introduce a new stochastic model for the variation of tick-by-tick asset price both in one and two dimension able to reproduce the strong microscopic mean reversion and the Epps effect, and [41], which introduces a measure of the market activity providing a direct access to the level of its endogeneity and as a potential predictor of market micro instabilities.

In this paper we use Hawkes processes for modelling the dynamics of jumps of individual assets and we show that they describe well the time clustering of jumps. However the direct extension of the application of Hawkes processes to describe the dynamics of jumps in a multi-asset framework is highly problematic and inconsistent with data. In fact, from a methodological point of view, even by using a simple two parameter kernel (for example, ex-

ponential) of the process, the number of parameters to estimate a Hawkes process on N assets is $N(2N + 1) \sim O(N^2)$, which is clearly too high. Moreover, even when we consider $N = 2$ stocks, we empirically find that a bivariate Hawkes model is unable to describe the empirical data, especially to replicate the high number of *synchronous* jumps that we observe. This is due to the fact that the kernel structure of Hawkes is more suited to model lagged jumps than synchronous jumps.

For this reason, the main methodological contribution of this paper is the introduction of Hawkes factor models to describe systemic cojumps. We postulate the presence of an unobservable point process describing a market factor. When this factor jumps, each asset jumps with a given probability, which is different for each stock. In general, an asset can jump also by following an idiosyncratic point process. In order to capture also the time clustering of jumps, we model the point processes as Hawkes processes. We show how to estimate this model and discriminate between systemic and idiosyncratic jumps. We show that the model is able to reproduce both the longitudinal and the cross sectional properties of the multi-asset jump process.

Recent approaches sharing some aspects with the current paper are discussed in [27, 42] and [24]. The former design a model of asset returns able to capture periods of crisis characterized by contagion and consider it to solve the problem of optimal investment-consumption for a log-utility investor. The jump-diffusion component of the dynamics is described in terms of a class of multi-dimensional Hawkes models, and the authors discuss an estimation methodology based on the Generalized Method of Moments. Not surprisingly, when they estimate the model on real data, they face the problem of the curse of dimensionality and limit most of the study to the two asset case. Our one factor approach solves the calibration issue in a natural way and therefore represents a viable alternative to their model. By using data sampled at a frequency of eleven minutes the authors of [24] find empirical support to the hypothesis that stocks tend to be involved in systematic cojumps², rejecting an assumption of independence in jump arrival times among stocks. We find strong evidence of this result also at the frequency of one minute and furthermore we extensively investigate the properties of self and mutual excitation possessed by the jumps series of different stocks.

The paper is organized as follows. In Section 2 we present our dataset and in Section 3 we summarize our jump detection method (detailed in Appendices). The identification of jumps is a delicate topic and therefore we are careful in using a robust identification method in order to minimize the number of false positives. In Section 4 we provide empirical evidence of the large number of systemic cojumps and we describe some simple statistical properties. In Section 5 we present the statistics we use to test our models in their ability of reproducing the multiple jumps of a single stock in a given time window and the cross jumps, i.e. the occurrences of jumps in two different stocks in a given time window. Section 6 discusses how Hawkes processes fit the jump process in a univariate and in a multivariate setting. In Section 7 we present the Hawkes factor model approach, showing how to estimate it and the results obtained in the investigated dataset. Finally, in Section 8 we draw some conclusions.

2 Data description and data handling

The analyses reported in this work are performed on tick-by-tick transaction data for the period from 5th March 2012 to 9th July 2012 that have been made available to us by LIST S.p.A.³.

²From a terminological viewpoint, authors of [24] define a systematic cojump as an instance when one stock and a market index jump at the same time. In this paper we use the term systemic cojump to indicate a sizeable number of stocks jumping simultaneously. In Section 7 we show how to identify systemic cojumps in a self consistent way.

³www.list-group.com

Company	ISIN code	trades per day ($\times 10^3$)
Assicurazioni Generali	IT0000062072	7.5
Mediobanca	IT0000062957	3.6
Banca Popolare di Milano	IT0000064482	4.9
Saipem	IT0000068525	5.4
Intesa Sanpaolo	IT0000072618	18.1
Mediaset	IT0001063210	3.1
Banca Monte dei Paschi di Siena	IT0001334587	9.2
Fiat	IT0001976403	9.6
Enel	IT0003128367	10.7
Eni	IT0003132476	10.4
UBI Banca	IT0003487029	3.0
Telecom Italia	IT0003497168	6.3
Finmeccanica	IT0003856405	4.7
Prysmian	IT0004176001	2.8
Banco Popolare	IT0004231566	4.2
Pirelli & C.	IT0004623051	4.6
Fiat Industrial	IT0004644743	5.1
UniCredit	IT0004781412	27.4
Tenaris S.A.	LU0156801721	3.3
STMicroelectronics N.V.	NL0000226223	4.3

Table 1: List of the twenty investigated stocks, with the ISIN code and the average number of transactions per day.

In this work we investigate 20 among the most liquid stocks of the FTSE MIB index of the Italian stock exchange Borsa Italiana. We sample prices at a frequency of one minute, taking the last executed price, to obtain 1-minute logarithmic returns. The total number of days in the period is 88, with 505 intraday returns each day.

In Table 1 we list the twenty stocks, along with the ISIN code and the average number of trades per day. From the last column we see that the assets are characterized by a certain level of heterogeneity, since the trade activity varies between 2.8 and 27.4 thousand transactions per day. A low level of activity implies in general a higher probability for the absence of transactions inside a given sampling interval, and this effect has important implications for the methodology that we use for the detection of jumps.

The process from tick-by-tick data to jump identification can be described as follows. First, anomalous values in tick-by-tick price data are detected and removed. The algorithm for the outliers detection that we use is due to Brownlees and Gallo [43] and it is explained in Section A.1 of the Appendix. We find no outliers at all for 13 stocks and few units for the other 7, for the vast majority concentrated on the very first minute of the day. Removing outliers can in principle introduce a distortion in the analysis. However, the number of anomalous prices that we identify is extremely low, and the probability of an actual distortion of returns, computed at the minute level, is negligible. Cleaned prices are then sampled at a frequency of one minute and logarithmic returns are calculated carrying forward to the next sampling instant the last price observed within the sampling window.

In the identification of jumps, an important care should be devoted to the way in which intervals without trades are treated. In fact, jump detection methods typically compare returns with local volatility. In periods of low liquidity (or missing observations) one risks to underestimate volatility and to identify relatively small price returns as jumps. There are two different

situations which lead to having no price observations in a given time interval.

The first one occurs in correspondence of volatility auction phases. According to the rules of the market and the accompanying instructions [44, 45], whenever the price exits a reference interval, reaching too high or too low levels, the continuous double auction phase of the exchange is suspended and a volatility auction starts. In such a phase we consider the related returns as *not available*. We do not have direct information on volatility auctions, but we infer their presence from market data. Since the investigated stocks are among the most liquid ones, we are reasonably sure to be able to distinguish volatility auction phases from continuous double auction phases of low liquidity. For more details, see the discussion in Section A.2 of the Appendix.

The second mechanism for missing price observations is when the stock is available for trades, but still there are no transactions in that particular minute. There are several ways to treat these cases. To fix the notation, let i and $i + 1$ label two consecutive sampling instants between which no trade has been made. The most common way to deal with the missing observation of the price between times i and $i + 1$ is to bring forward the last recorded price. This means setting the price $p_{i+1} = p_i$ and the log-return $r_{i+1} = 0$. Sampling intervals with no transactions in them are therefore given a zero return. This is our first method of dealing with Missing Observations, which we call MO1. An alternative method of treating gaps in available data made of missing observations is simply to consider the corresponding returns as not available data ($r_{i+1} = \text{NA}$), thus avoiding to give them a numerical value. The next available return will be at the first sampling time after the next available observation. It can be defined as the usual difference of the sampled logarithmic prices after and before the gap (method MO2) or as this quantity rescaled by the square root of the time between the two sampling times (MO3). For instance, for the price series $p_0, p_1, -, -, p_4, p_5$ we construct the following three returns series:

$$\text{MO1} \quad \log \frac{p_1}{p_0}, 0, 0, \quad \log \frac{p_4}{p_1}, \log \frac{p_5}{p_4};$$

$$\text{MO2} \quad \log \frac{p_1}{p_0}, \text{NA}, \text{NA}, \quad \log \frac{p_4}{p_1}, \log \frac{p_5}{p_4};$$

$$\text{MO3} \quad \log \frac{p_1}{p_0}, \text{NA}, \text{NA}, \frac{1}{\sqrt{3}} \log \frac{p_4}{p_1}, \log \frac{p_5}{p_4}.$$

The reason for the rescaling of the over-gap return in method MO3 is that we do not want to identify as a jump a price change that is compatible with a diffusive random movement of the efficient price.

Once the returns are defined according to one of the three possible methodologies discussed above, they are checked in order to eliminate possible large values due to stock splits or merges. We further check for the presence of stock splits in our data through the detection of returns greater than 0.2 in absolute value. This would detect, for example, a 3-for-2 split or a 4-for-5 merge. In our data we do not find any such return.

The last step to perform is the removal of daily seasonalities, by filtering out the intraday pattern from raw returns. As it is well known, intraday returns show significant seasonal behaviour, as the dynamics of markets is greatly variable during the day. Opening and closing periods generally show a higher volatility than the rest of the day, since traders are more active during these phases. If these daily seasonalities are not properly filtered out, spurious jump detection may happen in correspondence of these periods. We describe the details of how we deal with this issue in Section A.3.

3 Jump identification

After performing the data cleaning procedures explained in Section 2 and detailed in Appendix A, we proceed with the detection of jumps. Following a standard approach, we estimate

the local volatility σ and then test whether the ratio between absolute returns and local volatility is above a given threshold, that is, if

$$\frac{|r|}{\sigma} > \theta.$$

In this work we take $\theta = 4$ as in [46] and return series are obtained from historical prices as reported in Section 2. In order to estimate the local volatility, we can follow several strategies available in literature. We base our approach on the realized absolute variation and the realized bipower variation, whose asymptotic theory is treated for example in [47]. Our estimators of local volatility based on these quantities are defined by the exponentially weighted moving averages

$$\hat{\sigma}_{\text{abs},t} = \mu_1^{-1} \alpha \sum_{i>0} (1-\alpha)^{i-1} |r_{t-i}| \quad (1)$$

and

$$\hat{\sigma}_{\text{bv},t}^2 = \mu_1^{-2} \alpha \sum_{i>0} (1-\alpha)^{i-1} |r_{t-i}| |r_{t-i-1}|, \quad (2)$$

where $\mu_1 = \sqrt{\frac{2}{\pi}} \simeq 0.797885$ and the parameter of the exponential averages is $\alpha = \frac{2}{M+1}$, with $M = 60$, which corresponds to a characteristic time of 30.5 minutes. In order to avoid biased estimates due to the presence of jumps in past returns, we actually use modified versions of the estimators (1) and (2) that use only returns where no jump is detected. In fact the presence of a jump in the estimation window leads to an overestimation of the volatility. We refer the reader to the Appendix A.4 for the details on the implementation of our estimators and the asymptotic properties of the realized absolute variation and the realized bipower variation.

Since we construct three series of returns (which differ only in correspondence and proximity of missing observations), for each of which we obtain the volatility estimates $\hat{\sigma}_{\text{abs}}$ and $\hat{\sigma}_{\text{bv}}$, we have six volatility-normalised series $\{\frac{r_t}{\sigma_t}\}$, each of which with its own set of detected jumps. In order not to be sensitive to the choice of how to treat missing observations nor to the choice of volatility estimation, we take as final series of jumps the intersection of the six sets of jumps. This should also minimize spurious jump detection, as suggested in [48].

3.1 Performance of the jump detection methods on simulated data

In the next section we will discuss the results that we have found about identified jumps. Preliminarily, we conduct a simulation study to assess the effectiveness of our entire procedure from data cleaning to jump detection. The usual approach is to assume the observed log-price $p(t)$ as coming from the sum

$$p(t) = X(t) + \epsilon(t) \quad (3)$$

of an efficient log-price $X(t)$, following a simple jump-diffusion model

$$dX(t) = \mu(t) dt + \sigma(t) dW(t) + \kappa(t) dJ(t), \quad (4)$$

where the three terms on the right hand side are the drift, diffusion and jump components, and a microstructure noise $\epsilon_t \sim \text{i.i.d. } \mathcal{N}(0, \eta^2)$. Following Fan and Wang [49], for the dynamics of the volatility $\sigma(t)$ we take the Geometric Ornstein-Uhlenbeck model

$$d \log \sigma^2(t) = -(0.6802 + 0.1 \log \sigma^2(t)) dt + 0.25 dW'(t), \quad (5)$$

with correlation between the Wiener processes $W(t)$ and $W'(t)$ equal to -0.62. Simulating model of Equations (3) and (4) we fix the drift term μ equal to 0, and the microstructure noise standard deviation η to $10^{-5} dt$. As far as the process J_t is concerned, we set its rate equal to

three jumps per day, and the size κ as a multiple of the spot volatility bootstrapped from the sample of the empirical data. Using the Euler discretization scheme we obtain simulated high-frequency log-prices from the model at a frequency of one minute, which corresponds to the sampling frequency of our data. The time convention that we adopt fixes the one day horizon equal to one, so we have $dt = 1/1440$. We simulate data for a period 50 times longer than that of the sample data, that is, $88 \times 50 = 4400$ days. We preliminarily draw a Monte Carlo simulation of the process (5) and then we rescale it by adding the intraday volatility pattern. The seasonal volatility is used in (4) to generate one realization of the process X_t . Finally, for each day we discard some observations in order to reproduce the observed intertrade times, thus introducing artificial missing observations to investigate the role of methods MO1, MO2, MO3 of Section 2 in the jump detection tests. This operation is performed through a random sampling from the empirical distribution of the intertrade times two examples of which are shown in Figure 11.

Table 2 shows the performance of our jump detection procedure for simulations calibrated on two stocks: Intesa Sanpaolo, the second most liquid stock among our data, and Banco Popolare, which has a much lower liquidity. The performance is presented in terms of size and power of the test, for the jumps series coming from methods MO1, MO2, MO3, with the two different estimates of volatility $\hat{\sigma}_{\text{abs}}$ and $\hat{\sigma}_{\text{bv}}$. In the last column and row we report the intersections between the different methods. Simulation calibrated on Intesa Sanpaolo shows small differences among the three methods of dealing with the missing observations in sampled data. This is expected, since there are relatively few points on which the corresponding return series differ. Also taking the intersection of jumps detected in methods MO1, MO2, MO3 has performances comparable to those of each of the three methods.

A benchmark simulation calibrated on the stock Banco Popolare shows instead a very different picture, namely it highlights the significance of the way of filling the missing observations for a stock of low liquidity. If such missing observations are filled with zeroes (method MO1), volatility estimates from past returns are lower than for methods MO2 and MO3. This translates into detecting much more jumps, both true and spurious. In absolute terms, average counts of right and false positives are respectively 189 and 143 for method MO1, 146 and 45 for method MO2, 147 and 10 for method MO3, 134 and 7 for the intersection of all methods. Thus, results of this simulation clearly show that taking the intersection of jumps detected in methods MO1, MO2, MO3 drastically reduces the number of false positives, although this is accompanied also with a significant reduction of right positives. Taking a conservative point of view, we are ready to miss the detection of some true jumps as long as this procedure allows minimizing false positive hits. This simulation also demonstrates in a very clear way how the method MO3 is effective in detecting much fewer false positives than method MO2, still detecting the same amount of true jumps.

4 Basic statistics of jumps of individual stocks and of systemic jumps

In this section we present some basic statistical properties of detected jumps. In particular we will consider the restricted set of jumps identified simultaneously by all methods, in order to minimize the number of false positives. We then consider some simple statistical characterization of cojumps, i.e. instances in which the price of at least two stocks jumps in the same minute.

	MO1		MO2		MO3		$\bigcap_{\text{MO1,MO2,MO3}}$	
	Size	Power	Size	Power	Size	Power	Size	Power
σ_{abs}	0.036%	63.1%	0.033%	62.5%	0.029%	62.7%	0.028%	62.1%
σ_{bv}	0.039%	63.4%	0.037%	62.7%	0.033%	62.7%	0.032%	62.3%
$\bigcap_{\text{abs,bv}}$	0.031%	61.6%	0.029%	61.0%	0.026%	61.0%	0.025%	60.6%
	MO1		MO2		MO3		$\bigcap_{\text{MO1,MO2,MO3}}$	
	Size	Power	Size	Power	Size	Power	Size	Power
σ_{abs}	0.369%	72.9%	0.110%	56.4%	0.027%	58.0%	0.019%	52.3%
σ_{bv}	0.390%	74.7%	0.135%	59.7%	0.031%	58.3%	0.025%	55.1%
$\bigcap_{\text{abs,bv}}$	0.324%	71.7%	0.102%	55.2%	0.023%	55.8%	0.016%	50.9%

Table 2: Size and power of jump tests, determined on a simulation calibrated on real data of Intesa Sanpaolo (top) and Banco Popolare (bottom). The size is the ratio between the number of false positives and the number of minutes when the jump is not present. The power is one minus the ratio between the number of right positives and the number of real jumps.

4.1 Jumps of individual stocks

In the previous sections we have introduced several methods to identify jumps in price time series. By using Monte Carlo simulations of a realistic model of price dynamics, we have shown also that different methods display different ability to identify jumps correctly, both in terms of false positives and in terms of false negatives. The use of the intersection among the different methods improves significantly the identification procedure, at least on simulated data.

Table 3 shows the number of jumps for the stock Monte dei Paschi di Siena, detected with the six different methods. For this specific stock, the number of jumps ranges from 200 to 281, showing a significant dependence on the used method. The table shows also the size of the intersection of the sets of detected jumps. When one considers jumps detected by all methods the number of jumps falls to 178. This is the restricted set of events that we will consider. We notice that, in this case, by estimating the volatility as $\hat{\sigma}_{\text{abs}}$ we find more jumps than when using $\hat{\sigma}_{\text{bv}}$, regardless of what the method of treating missing observation is. However this is not always the case. For the stock Intesa Sanpaolo it is the other way round: more jumps are detected when we estimate local volatility by $\hat{\sigma}_{\text{bv}}$ than when we use $\hat{\sigma}_{\text{abs}}$, in methods MO1, MO2 and MO3. It is not clear to us whether there are features of the return series (such as quantity and position of the missing observations) which are systematically responsible for the first or the second scenario to happen.

	MO1	MO2	MO3	$\bigcap_{\text{MO1,MO2,MO3}}$
$\hat{\sigma}_{\text{abs}}$	281	228	217	205
$\hat{\sigma}_{\text{bv}}$	260	208	200	190
$\bigcap_{\text{abs,bv}}$	239	196	186	178

Table 3: Summary table of the number of detected jumps by using different methods and of the size of their intersection. The investigated stock is Monte dei Paschi di Siena.

We extend this analysis to the whole set of 20 stocks. In Table 4 we report simple statistics about the jumps identified on the 20 stocks by using the intersection of the six methods. The number of jumps detected in the 88 days of our sample varies between 59 and 188 across the

twenty stocks (that is, between 0.67 and 2.14 jumps per day per stock), with an overall average value of 108, corresponding to 1.23 jumps per day per stock⁴.

Five stocks show a statistically significant difference in the number of positive and negative jumps (at 10% significance level), when one assumes a null model in which stocks have the same probability of jumping up and down. In such a model, the number of jumps in either direction has a binomial distribution. Interestingly, in all five cases the number of negative jumps is greater than that of positive jumps. Statistically significant asymmetry in the jumping direction is also found in the overall jump count (at 1% significance level). The twenty stocks have also variable proportions of “single jumps”, that is, jumps that do not occur simultaneously to jumps of other stocks, with two having even a greater number of cojumps than single jumps.

ISIN	jumps	jumps up	jumps down	single jumps	cojumps
IT0000062072	103	48 (47%)	55 (53%)	53 (51%)	50 (49%)
IT0000062957	63	29 (46%)	34 (54%)	38 (60%)	25 (40%)
IT0000064482	121	60 (50%)	61 (50%)	97 (80%)	24 (20%)
IT0000068525	93	46 (49%)	47 (51%)	56 (60%)	37 (40%)
IT0000072618	127	67 (53%)	60 (47%)	55 (43%)	72 (57%)
IT0001063210	59	28 (47%)	31 (53%)	44 (75%)	15 (25%)
IT0001334587	178	73 (41%)	105 (59%)	150 (84%)	28 (16%)
IT0001976403	123	61 (50%)	62 (50%)	76 (62%)	47 (38%)
IT0003128367	188	81 (43%)	107 (57%)	107 (57%)	81 (43%)
IT0003132476	155	66 (43%)	89 (57%)	95 (61%)	60 (39%)
IT0003487029	70	28 (40%)	42 (60%)	41 (59%)	29 (41%)
IT0003497168	129	74 (57%)	55 (43%)	79 (61%)	50 (39%)
IT0003856405	95	50 (53%)	45 (47%)	74 (78%)	21 (22%)
IT0004176001	74	41 (55%)	33 (45%)	46 (62%)	28 (38%)
IT0004231566	103	50 (49%)	53 (51%)	72 (70%)	31 (30%)
IT0004623051	115	47 (41%)	68 (59%)	85 (74%)	30 (26%)
IT0004644743	100	51 (51%)	49 (49%)	65 (65%)	35 (35%)
IT0004781412	118	49 (42%)	69 (58%)	57 (48%)	61 (52%)
LU0156801721	59	27 (46%)	32 (54%)	32 (54%)	27 (46%)
NL0000226223	86	39 (45%)	47 (55%)	51 (59%)	35 (41%)
total	2159	1015	1144		
average	108.0	50.8 (47%)	57.2 (53%)		

Table 4: Number of detected jumps for the twenty stocks, with direction and cojumping information. Bold values are inconsistent (at 10% significance level for single stocks, at 1% significance level for the total counts) with a null assumption of equal probability of jumping up and jumping down. Single jumps occur at a time when no other stock jumps, while cojumps occur simultaneously with jumps in other stocks.

We investigate the distribution of jumps during the trading day. The left panel of Figure 1 shows the histogram of the time of the day when a jump occurs. The figure is obtained by including all the jumps of the 20 stocks, but we count only once a minute where multiple stocks jump simultaneously. We observe that there is no clear periodicity in the number of jumps, indicating that our careful intraday pattern removal is quite effective. Moreover, apart a spike

⁴It is worth noticing that in a recent study [46] authors found more than seven jumps per stock per day when investigating a relatively large set of US stocks. Even if the threshold used is the same as ours, i.e. $\theta = 4$, they use only one method, they do not detail the data cleaning procedure, and they investigate a much larger universe of stocks. These might be reasons for the difference in the typical number of detected jumps.

observed at the beginning of the day, there is no evidence of minutes of the day when it is more likely that one stock jumps.

4.2 Systemic jumps




Figure 1: Left panel: intraday distribution of the time of the day when a jump occurs. The x axis is the number of minutes from the beginning of the trading day. Notice that if more than one stock jump in a given minute, we count the minute only once in the histogram. This is done in order to avoid to count N times an event when N stocks jump simultaneously. Right panel: same figure as the left panel, but considering only cojumps of more than 5 stocks.

The main topic of this paper is systemic cojumps, i.e. minutes when the price of a (possibly large) number of stocks displays a jump. For this reason it is important to investigate how frequently cojumps occur and to compare the observed statistics with those expected under some null hypotheses.

We start with some simple visualization of the occurrences of multiple cojumps. Figure 2 shows the dynamics of the number of cojumps, indicating also the number of stocks that display a jump in a given minute. We notice that there are several occurrences where more than 8 stocks jump simultaneously (big circles). For example, we observe one case each when 10, 12, 13, 15, 16, 17, and 20 stocks jump simultaneously. Also the number of cases where more than 3 stocks jump (medium-sized circles) is quite high. For example, we observe 22 cases with 4 stocks, 15 cases with 5, 2 cases with 6, 4 cases with 7, 5 cases with 8, 3 with 9, and 2 with 11. Finally, there is a significant background of cases where two (136) or three (44) stocks jump simultaneously.

By considering all the 240 events in which multiple stocks jump, we find only 7 cases in which not all the jumping stocks follow the same direction. There are 6 cases in which 2 stocks jump in opposite direction and 1 case in which two stocks jump down and one stock jumps up. The fact that when several stocks simultaneously jump they all move in the same direction suggests that a single common factor explains the jumping probability. In Section 7 we will develop this idea more formally by introducing models that capture this important feature of real data.

It is worth noticing that the Figure 2 also indicates that there are no specific times of the day (for example corresponding to pre-announced news or opening of other markets) where the systemic cojumps are more frequent. This qualitative statement can be made more precise by drawing the histogram of the time of the day when a systemic cojump with more than 5 stocks jumping simultaneously (see right panel of Figure 1). Therefore systemic jumps do not occur at preferential moments of the trading day.




Figure 2: Time series of the cojumps observed in the set of 20 investigated stocks. The horizontal axis is the time of the day and the vertical axis is the day. The size of the circle codes the number of stocks simultaneously jumping in a given minute.

In order to compare the pattern observed in Figure 2 for the systemic cojumps with a null hypothesis, we perform a bootstrap analysis. Specifically, we construct a bootstrap replica independently for each stock. Therefore our replicas are consistent with a model of 20 independent but not identical Poisson processes for the jumps. In the next section we will use this model as the simplest benchmark model. In the left panel of Figure 3 we show the analogous of Figure 2 for one of the bootstrap replicas of the real data. As it can be seen, in this replica there are no cases where more than three stocks jump simultaneously, a result that is clearly inconsistent with the real data. To be more quantitative, in the right panel of Figure 3 we show the histogram of the number of stocks jumping in a minute (in which at least one stock jumps) for real data (solid line). We compare it with the curve obtained by taking the 0.01% confidence interval from bootstrap analysis. It can be seen that already the observed number of cojumps of two stocks is incompatible with the one observed in bootstrap test at the 0.01% confidence. Moreover in 10,000 bootstrap replicas we never observe a cojump with more than 4 stocks (observed only once) while in real data we have several cases with the price of many stocks simultaneously jumping.

In conclusion, our preliminary analysis shows that (i) jumps are relatively frequent, even when one considers relatively strict detection criteria, (ii) there is a large number of cojumps, i.e. minutes when a sizable number of stocks (up to 20!) simultaneously jump, (iii) these cojumps show no clear timing inside the day, and, more important, (iv) they are absolutely not compatible with a null model of independent but not identical Poisson processes. In Section 7 we will introduce models able to describe the jump dynamics of a set of stocks.




Figure 3: Left panel. Time series of the cojumps observed in a bootstrap replica of the real data preserving the jump intensity of each stock. The horizontal axis is the time of the day and the vertical axis is the day. The size of the circle codes the number of stocks simultaneously jumping in a given minute. Right panel. Histogram of the number of stocks simultaneously cojumping in a minute (solid line). For comparison the dashed line shows the 0.01% confidence interval for the counts under the null hypothesis of independent but not identical Poisson processes.

5 A multi-scale statistical test based on multiple jumps and cross jumps detection

The empirical observation of a large number of cojumps requires a rigorous statistical test to compare the observed behaviour of jumps with the prediction of multivariate point processes. The main problem in working with jumps is that their number is relatively small, and therefore the statistics must be suited to work on small samples. Moreover, we want to use a multi-scale statistical test, i.e. a test that is able to identify deviations from models at different time scales. A correlated point process can be undistinguishable from a Poisson process if the observation time is much larger than the correlation time of the process. Therefore we use a test that considers simultaneously different time scales. Finally, even if the main topic of this paper is cojumps of different stocks, we are also interested in investigating deviation from a Poisson behaviour at the individual stock level. Therefore we shall introduce two related statistics, one for consecutive jumps of the same asset and one for jumps of different stocks occurring at close times.

Since the distributional and time correlation analysis is unfeasible with small samples of jumps we will consider the following statistics that measures the frequency of multiple jumps occurring in the same time windows.

Specifically, when considering individual stocks, we define a *multiple jump* (MJ) as an event which occurs when *at least two jumps of the same stock price* happen inside a time window of a fixed length w . Let us call s_i the number of jumps inside the i -th window. An estimator of the MJ probability in a window of length w over a sampling period of length N is given by

$$\hat{p}_w^{\text{MJ}} = \frac{\sum_{i=1}^{\lfloor \frac{N}{w} \rfloor} \mathbf{1}_{s_i \geq 2}}{\lfloor \frac{N}{w} \rfloor}, \quad (6)$$

where $\mathbf{1}_A$ is the indicator function of the event A , and the symbol $\lfloor N/w \rfloor$ corresponds to the integer part of the ratio N/w .

The same idea can be extended to capture also the cross sectional clustering of jumps, in particular the evidence of a large number of simultaneous jumps of different stocks. The second notion which we will work with is therefore that of *cross jumps* (CJ) between two stocks defined

when *both stocks jump at least once* inside a given time window of fixed length w . With the same notation as before, the estimator of the CJ probability between stock l and k is given by

$$\hat{p}_w^{\text{CJ}} = \frac{\sum_{i=1}^{\lfloor \frac{N}{w} \rfloor} \mathbf{1}_{s_i^l \geq 1} \mathbf{1}_{s_i^k \geq 1}}{\lfloor \frac{N}{w} \rfloor}. \quad (7)$$

Our statistical procedure consists in estimating these quantities in real data and compare them, at each time scale, with the 99% and 95% confidence bands of the tested model obtained analytically (when possible) or via Monte Carlo simulations. It should be noted that, since the number of windows that we are considering is greater than one, we are indeed performing a multiple hypothesis test and we have to correct the significance level accordingly. Among the possible approaches discussed in literature, we decide to adopt the most conservative one, namely the Bonferroni correction. This correction amounts to divide the significance level of the single hypothesis by the total number of tested hypotheses in order to achieve a global significance at least of the pre-fixed level. Therefore when the empirical points of $\hat{p}_w^{\text{MJ}}$ and $\hat{p}_w^{\text{CJ}}$ fall inside the confidence bands of a given model we can not reject the null hypothesis with the given confidence level.

5.1 A benchmark case: the Poisson model

In order to show how our testing procedure works, we consider here an important benchmark case, in which the jumps process of each stock is described by an independent Poisson process. Under this model, the mean and variance of both previous estimators can be computed. For the MJ estimator we have

$$\mathbb{E}[\hat{p}_w^{\text{MJ}}] = p_{w,\lambda}, \quad \text{Var}[\hat{p}_w^{\text{MJ}}] = \frac{p_{w,\lambda} - p_{w,\lambda}^2}{\lfloor \frac{N}{w} \rfloor},$$

where $p_{w,\lambda} = \mathbb{P}(\{s \geq 2\}) = 1 - \text{Pois}_{\lambda w}(s=0) - \text{Pois}_{\lambda w}(s=1)$, and λ is the intensity of the Poisson process.

Analogously for the CJ estimator, we obtain

$$\mathbb{E}[\hat{p}_w^{\text{CJ}}] = q_{w,\lambda_l} q_{w,\lambda_k}, \quad \text{Var}[\hat{p}_w^{\text{CJ}}] = \frac{q_{w,\lambda_l} q_{w,\lambda_k} - q_{w,\lambda_l}^2 q_{w,\lambda_k}^2}{\lfloor \frac{N}{w} \rfloor},$$

where $q_{w,\lambda_i} = \mathbb{P}(\{s^i \geq 1\}) = 1 - \text{Pois}_{\lambda_i w}(s=0)$.

Since both quantities in Equation (6) and (7) correspond to the sum of a large number of indicator functions, the Central Limit Theorem implies that their distribution is well approximated by a Normal law whose p -values are readily available, allowing the analytical computation of the confidence bands. As usual, the value of the intensity of the Poisson process given by the maximum likelihood estimator is $\hat{\lambda} = \#\text{jumps}/N$.

In Figure 4 we show the result of our test on the Italian asset Assicurazioni Generali, for which $\lambda = 2.4 \times 10^{-3} \text{ min}^{-1}$. The filled circles correspond to the empirical values of the estimator, the solid line to the theoretical mean, while the dashed and dotted lines establish the boundaries of the 99% and 95% confidence bands, respectively, adjusted with the Bonferroni correction. The figure shows that the Poisson model is rejected for both levels. In our data sample this is the typical situation, which occurs for 18 stocks out of 20 (the two exceptions are represented by the assets Mediobanca and Finmeccanica). We therefore conclude that there is a strong evidence of time clustering of jumps and violation of the univariate Poisson model.




Figure 4: MJ probability test under Poisson null for the Italian asset Assicurazioni Generali.

5.2 Block bootstrap

At this stage, the natural step to proceed would be to assume a more sophisticated model for the jump process, and to test it as an alternative hypothesis. The Poisson hypothesis can be relaxed in several respects: among the possible alternatives we can weaken the assumption of identically and independently distributed waiting times, for example modelling them either as Markov processes or as realizations of a time inhomogeneous Poisson process. In the next Section we will consider the class of point processes known as self-exciting processes. However, before doing so we approach the problem in a model independent way and try to grasp the typical time scale at which the clustering of jumps takes place.

Our approach is simple and based on a block bootstrap technique. We fix the length of the block L_B and we draw $N_B = 10^4$ synthetic copies sampling with repetition from the original sequence of zeros and ones which corresponds to the point process of the jumps. On each copy we evaluate the estimator, we sort all the values and compute the bootstrap estimate of the confidence bands. The comparison of the empirical estimate with the bootstrap bands allows us to ultimately accept or reject the model-independent null hypothesis. For $L_B = 1$, our procedure destroys any existing autocorrelation, and the confidences we find correspond to the theoretical Poisson levels. In this case, for 18 stocks out of 20, the null is rejected. When we increase the value of L_B , the clustering structure with typical scale equal to or smaller than L_B is preserved. Therefore, we expect that for values of L_B below a critical level, dependent on the specific asset, we will reject the null, but when the length of the block is larger than the critical level, the estimate will be compatible with the null. Figure 5 confirms the expected behaviour for Generali⁵. We see that by increasing the value from $L_B = 1$ to $L_B = 7$ the level of the bootstrap confidences increases and the empirical points become fully compatible with the hypothesis. For the specific asset under consideration we have evidence of a time scale of few minutes, but the same result still holds for all the assets.

6 Modelling jumps with Hawkes processes

The evidences that we have described above play in favour of a null model which should be able to describe the tendency of jumps to cluster over a time scale of few minutes. A natural

⁵We report the analysis only for low values of the time windows lengths.




Figure 5: MJ probability test under block bootstrap null for the Italian asset Assicurazioni Generali.

extension of the homogeneous Poisson process, where the intensity is itself stochastic and tends to increase when a new jump arrives, is the class of point processes known as *Hawkes processes*.

6.1 Univariate case

In this section we provide the main results needed for the remaining of the paper, but for a complete mathematical treatment of Hawkes and more general point processes we refer the reader to the comprehensive textbook [50].

A univariate point process N_t ⁶ is called a Hawkes process if it is a *linear self-exciting process*, defined by the intensity

$$I(t) = \lambda(t) + \int_{-\infty}^t \nu(t-u) dN_u = \lambda(t) + \sum_{t_i < t} \nu(t-t_i), \quad (8)$$

where λ is a deterministic function called the base intensity, ν is a positive decreasing weight function, and t_i are the jumping times. The most common parametrization of ν is given by $\nu(t) = \sum_{j=1}^P \alpha_j e^{-\beta_j t}$, for $t > 0$, where $\alpha_j \geq 0$ are scale parameters, $\beta_j > 0$ control the strength of decay, and the positive integer P is the order of the process. A particular advantage of the linear Hawkes process of order $P = 1$ is that the log-likelihood function can be computed as

$$\mathcal{L}(t_1, \dots, t_n) = (1 - \lambda)t_n - \frac{\alpha}{\beta} \sum_{i=1}^n (1 - e^{-\beta(t_n - t_i)}) + \sum_{i=1}^n \ln(\lambda + \alpha R_i),$$

⁶If $\{t_i\}_{i=1, \dots, n}$ represents the random sequence of increasing event times $0 < t_1 < \dots < t_n$ associated with the point process, then $N_t \doteq \sum_{i \geq 1} \mathbf{1}_{t_i \leq t}$ defines the right continuous counting function. In what follows we will refer equivalently to the process and its counting function.

	Assicurazioni Generali	Intesa Sanpaolo
λ (min^{-1})	$(2.1 \pm 0.2) \times 10^{-3}$ ***	$(2.5 \pm 0.2) \times 10^{-3}$ ***
α (min^{-1})	$(3.1 \pm 1.3) \times 10^{-2}$ *	$(5.9 \pm 2.1) \times 10^{-2}$ **
β (min^{-1})	$(2.5 \pm 0.9) \times 10^{-1}$ **	$(4.3 \pm 1.2) \times 10^{-1}$ ***

Table 5: List of the parameters of the one dimensional Hawkes processes. Significance codes: $p_{\text{value}} < 0.001$ ‘***’, $0.001 \leq p_{\text{value}} < 0.01$ ‘**’, $0.01 \leq p_{\text{value}} < 0.05$ ‘*’, and $p_{\text{value}} \geq 0.05$ ‘ ’.

where the R function satisfies the recursion $R_i = e^{-\beta(t_i - t_{i-1})}(1 + R_{i-1})$ for $i \geq 2$ and $R_1 = 0$. In Ref. [32] it is shown that the stationarity of the process is guaranteed when $\int_0^\infty \nu(s)ds < 1$, which in our case reduces to the requirement $\alpha/\beta < 1$. If stationarity holds and under the further constraint that the base intensity is constant, the expected number of jumps in an arbitrary time interval of length T is given by $\lambda T/(1 - \alpha/\beta)$. The latter observation will be useful for the calibration of the factor models that we will discuss in Section 7 of this paper. The theoretical characterization of the maximum likelihood estimator of the Hawkes process parameters that we employ in this paper has been developed in [51] and [52], while the simulation algorithm we use is based on the procedure discussed in [53], which directly derives from the Shedler-Lewis thinning algorithm, [54].




Figure 6: MJ probability test under Hawkes null for the assets Assicurazioni Generali (left) and Intesa Sanpaolo (right).

We estimate the Hawkes processes on the univariate series of jumps of the investigated stocks. In Table 5 we report the parameters values with the associated errors and significance values for Generali and Intesa Sanpaolo. All the values are statistically significant. It is worth noting that the value of β for both assets confirms the empirical evidence enlightened by the bootstrap analysis. Indeed, the kernel relaxation times, given by β^{-1} , are in full agreement with clustering scales of few minutes identified with the bootstrap.

We then test the Hawkes model in its ability of reproducing the MJ probability at different time scales. The left and right plots in Figure 6 are obtained after the calibration of a one dimensional Hawkes process on the jumps time series of the assets Assicurazioni Generali and Intesa Sanpaolo, respectively. The figure shows that we can not reject the null both at 1% and at 5% significance levels (obtained from $N_{MC} = 10^4$ Monte Carlo simulations). The result holds both for short window lengths (as evidenced for Generali up to half an hour) and longer

horizons (up to one day for Intesa Sanpaolo).

We therefore conclude that univariate Hawkes processes are able to capture the empirical dynamics and time clustering of jumps for individual stocks.

6.2 Bivariate case




Figure 7: CJ probability test under independent Hawkes null for the pair of Italian assets Generali - Intesa Sanpaolo.

Are independent Hawkes processes able to describe the empirical behaviour of CJ probabilities? In order to answer this question, we compute the estimator (7) on each pair of assets and in Figure 7 we show an example of the results from the test computed on the pair Generali - Intesa Sanpaolo. The figure shows that the independent Hawkes process miserably fails in describing the CJ probability at all time scales. This is clearly due, at least in part, to a lack of coupling between the two processes. In order to capture such a dependence we calibrate a bivariate Hawkes process.

A K -dimensional Hawkes process is a linear self-exciting process defined by the multivariate intensity $\mathbf{I}(t) = (I^1(t), \dots, I^K(t))'$, where the k -type intensity with an exponential kernel of order one is given by

$$I^k(t) = \lambda^k(t) + \sum_{m=1}^K \sum_{t_i^m < t} \alpha_{km} e^{-\beta_{km}(t-t_i^m)}.$$

All the parameters which appear in the above expression are strictly positive, and the stationarity of the process is guaranteed if the spectral radius of the matrix $\Gamma = \left(\frac{\alpha_{km}}{\beta_{km}} \right)_{k,m=1,\dots,K}$ is strictly smaller than one. The parameters α_{kk} are responsible for the self-exciting property of the point process, while the remaining $K-1$ α_{km} 's capture the cross exciting effect of a jump in the stock m on the process of the asset k . When $K=2$ the number of free parameters is equal to ten, and the maximization of the likelihood becomes less trivial than the univariate case. To optimize the likelihood, we initialize the parameters with the values suggested by the one dimensional calibration, and we constrain their value to remain strictly positive. We preliminarily perform $N_{Ann} = 100$ searches of the maximum with the simulated annealing algorithm, then, with the optimal candidate supplied by the stochastic search, we initialize the deterministic search via conjugate gradient.

λ^1 (min^{-1})	$(2.0 \pm 0.1) \times 10^{-3}$ ***	λ^2 (min^{-1})	$(2.3 \pm 0.2) \times 10^{-3}$ ***
α_{11} (min^{-1})	$(1.6 \pm 0.9) \times 10^{-2}$ *	α_{21} (min^{-1})	$(4.5 \pm 0.1) \times 10^{-4}$ ***
α_{12} (min^{-1})	$(14 \pm 0.1) \times 10^{-4}$ ***	α_{22} (min^{-1})	$(3.4 \pm 1.3) \times 10^{-2}$ **
β_{11} (min^{-1})	$(3.7 \pm 1.6) \times 10^{-1}$ *	β_{21} (min^{-1})	$(7.0 \pm 18) \times 10^{-1}$
β_{12} (min^{-1})	$(3.2 \pm 3.4) \times 10^{-1}$	β_{22} (min^{-1})	$(4.9 \pm 1.6) \times 10^{-1}$ ***

Table 6: List of the parameters of the bivariate Hawkes process for ¹ Generali and ² Intesa Sanpaolo. Significance codes: $p_{\text{value}} < 0.001$ ‘***’, $0.001 \leq p_{\text{value}} < 0.01$ ‘**’, $0.01 \leq p_{\text{value}} < 0.05$ ‘*’, and $p_{\text{value}} \geq 0.05$ ‘’.




Figure 8: CJ probability test under bivariate Hawkes null for Generali - Intesa Sanpaolo.

In Table 6 we present the results of the calibration of a bi-dimensional Hawkes on the joint jumps process of Generali-Intesa. We then use Monte Carlo simulations to compute confidence bands for CJ probability as shown in Figure 8. It is clear from the figure that bivariate exponential Hawkes fails to describe the cross sectional clustering of jumps observed in real data. This might be due to several facts: the value of the cross exciting constants α_{12} and α_{21} is at least of one order of magnitude smaller than that of the constants α_{11} and α_{22} ; moreover, the time constants β_{12} and β_{21} associated to the cross excitations are poorly statistically significant. One possible reason could be the non reliability of the parameters since the optimization procedure is performed in a high dimensional space on a strongly non linear objective function. However, in this respect, we perform numerous experiments in a Monte Carlo framework. One test is the simulation of a bivariate Hawkes process, whose parameters are chosen equal to the values of Table 6. On each random realization, we apply the above optimization scheme and record the optimal values. With a statistics of 10^3 values per parameter, we measure no significant bias induced by the multidimensional maximization procedure. Moreover, we simulate 10^3 independent pairs of univariate Hawkes of length 44440 minutes, whose parameters are given in Table 5. We then perform the ten dimensional optimization on each copy, and find a confirmation of the tendency of the cross-exciting constants to significantly decrease. In light of these numerical evidences, we tend to marginalize the role played by the optimization scheme. Finally, we believe that the very reason why Hawkes processes fail to describe the cross-sectional dependence is that they are not designed by construction to capture synchronous effects, which in the case under our consideration seem to dominate the dependence between jumps.

7 A factor model approach to jumps

In this section we abandon the idea of an N dimensional point process, and we develop a different approach, which we will refer to as a jump factor model. The intuition which drives the modelling is that a stock jumps both because it is triggered by jumps of a market common factor and because of an idiosyncratic term.

The first issue that we want to clarify is if, even in a Poisson framework, the factor mechanism is able to describe the empirical dependence structure. Then we will propose a scalable model, which ideally should be effective and also sufficiently simple and robust.

7.1 Bivariate Poisson factor model

In order to illustrate the main idea, we start with a toy model which enlightens the role of the market factor and clarifies the mechanism which generates the dependence structure. However, this is an unrealistic model, where the behaviour of the assets is completely determined by the evolution of the factor process.

We assume that there is one unobserved market factor point process. When the factor jumps, the stock S_1 jumps with probability p_1 and the stock S_2 jumps with probability p_2 . If the factor does not jump, neither the first nor the second stocks can jump; obviously the converse is not true. If the market factor is described by a Poisson process of intensity λ_F , then the expected number of factor's jumps over the period T is given by $\lambda_F T$. In order to fix the free parameters of the model, λ_F , p_1 , and p_2 , we use the following relations

$$\begin{aligned} p_1 \lambda_F T &= n_1, \\ p_2 \lambda_F T &= n_2, \\ p_1 p_2 \lambda_F T &= n_{12}, \end{aligned} \tag{9}$$

where n_1 and n_2 are the realized number of jumps of the stock S_1 and S_2 , respectively, while n_{12} represents the realized number of cojumps among S_1 and S_2 within the one minute sampling interval. The first two equations thus require that the expected number of jumps of the single assets matches the realized values; the latter relation guarantees that, on average, the realized number of cojumps among the assets corresponds to the theoretical expectation. In terms of the counting functions N_t^1 and N_t^2 , which describe the jumps processes of the stocks, we have $n_1 = N_T^1$ and $n_2 = N_T^2$. The system of Equations (9) can be inverted and provides a direct way to express the parameters of the model in terms of the observable quantities T , n_1 , n_2 , and n_{12} :

$$\lambda_F = \frac{n_1 n_2}{n_{12}} \frac{1}{T}, \quad p_1 = \frac{n_{12}}{n_2}, \quad \text{and} \quad p_2 = \frac{n_{12}}{n_1}.$$

In Figure 9 we show the result of the test of the cojumps estimator $\hat{p}_w^{\text{CJ}}$ against a null represented by the Poisson factor model that we have just described. The time period that we consider corresponds to the usual interval of 88 days, during which Generali and Intesa Sanpaolo jump 103 and 127 times, respectively, while the realized number of cojumps n_{12} is equal to 26. From these values we obtain a Poisson intensity equal to $\lambda_F = 1.1 \times 10^{-2}$ per minute, and probabilities $p_1 = 0.20$ and $p_2 = 0.25$. The confidence intervals for the null are estimated drawing $N_{MC} = 10^4$ paths from the Poisson factor process, and then thinning each realization with Bernoulli variables of probability p_1 and p_2 .

The results provided by this simple model are very satisfactory and should convince the reader that the proposed mechanism is adequate to capture the cross dependence between jumps. However, the model is quite unrealistic in several respects: (i) the assets can not jump independently of the factor, (ii) the Poisson nature of the process leads to a severe




Figure 9: CJ probability test under Poisson factor null for Generali - Intesa Sanpaolo.

underestimation of the realized number of multiple jumps of the same stock, and (iii) the approach has not a straightforward extension to an arbitrary number of assets. In order to amend all these drawbacks, in the next section we will discuss a more general and flexible model, rooted on the factor idea, but able to achieve a higher level of realism and scalability.

7.2 N dimensional Hawkes factor model with idiosyncratic components

We now consider a set of N assets. With respect to the previous model, we need to introduce explicitly a proxy of the common factor process. The naive idea of a proxy of the factor based on those events when a cojump among all the N stocks happens realistically does not work. In the data sample of twenty stocks that we consider in the current study, we experience just a single event when all the stocks jump simultaneously. The proxy that we propose is characterized by a counting function which increases by one unit whenever we detect a cojump which involves a group of assets whose number J conflicts with the null of independence. We therefore need to identify the threshold value $\bar{J}$ compatible with independence. In general it will be time dependent, $\bar{J}_t$, and will vary between one and N . In order to detect $\bar{J}$ we propose a rigorous statistical methodology. However, we preliminarily want to convince the reader that the range of variability of $\bar{J}$ for the case under study $N = 20$ is quite narrow. First of all, we can exclude the case $\bar{J} = 1$, since otherwise we would consider every jump that occurs on the market as potentially systemic. On the other side, if we fix $\bar{J}$, by definition then we would consider the events which involve $J < \bar{J}$ jumps as likely occurrences. As shown by the bootstrap experiment in Section 4.2, the probability associated with events involving more than two stocks is extremely low, and this would fix $\bar{J} = 3$. However, the null represented by the bootstrapped data is too extreme, since it does not take into account any effect of self excitation. Heuristically, we conclude that $\bar{J} \simeq 4$.

We now present a rigorous procedure that supports the above heuristic argument. To each stock in our sample we associate a finite number of event times $\{t_i^s\}$, where $s = 1, \dots, N$ and $i = 1, \dots, n_s$ label the assets and the asset specific jumps, respectively. In equivalent terms, each asset is characterized by a counting function N_t^s , and $n_s = N_T^s$ is the realized number of jumps of the stock s . According to the procedure of jumps' identification, the t_i^s 's are measured in minutes and take only integer values running from one to T . A counting process N_t^s is

associated with an intensity function I_t^s defined by

$$\mathbb{P}(N_t^s \text{ has a jump in } [t, t + \Delta t] | \mathcal{F}_t) = I_t^s \Delta t,$$

for $\Delta t \rightarrow 0^+$, where $\mathbb{P}$ stands for the probability and $\mathcal{F}_t$ corresponds to the history of the process up to present time t . We then assume that the counting processes which describe our assets correspond to N independent univariate Hawkes processes and estimate the parameters which characterize the intensities I_t^s for $s = 1, \dots, N$ via maximum likelihood. We fix Δt equal to one minute and via Equation (8), for each $t = 1, \dots, T$, we can compute the vector of probabilities $\mathbf{p}_t = (p_t^1, \dots, p_t^N)' = (I_t^1 \Delta t, \dots, I_t^N \Delta t)'$. We test if the number of jumps that we observe at time t is compatible with the cross independence among the processes. Under the null hypothesis, the discrete probability of the event $J_t = j$ reads

$$\mathbb{P}(J_t = j) = \sum_{1 \leq l_1 < \dots < l_j \leq N} p_t^{l_1} \dots p_t^{l_j} \prod_{k \in \{1, \dots, N\} \setminus \{l_1, \dots, l_j\}} (1 - p_t^k).$$

Since we repeat the test T times, we adjust the significance level with the Bonferroni's correction. If at time t^F we reject the null, we attribute the event to a systemic shock, and we remove it from the set $\{t_i^s\}$. The procedure has to be iterated as many times as required in order to remove all the systemic jumps.

In Table 7 we detail the results of the identification procedure of the systemic factor performed on our data set for the significance level 1%. We report the values of the event times $\{t_i^F\}$ identified at each iteration, and the associated threshold values $\bar{J}_{t_i^F} = \min_j \{j : \mathbb{P}(J_{t_i^F} > j) \leq 0.01/T\}$. It is worth noticing that the typical value of the time varying threshold is close to four, as anticipated. Moreover, in order to better clarify our procedure, we comment more extensively the case of the event times $\{30036, 30037, 30038\}$ corresponding to the dates May 5th 2012 $\{13.01, 13.02, 13.03\}$, when we measured three events which had involved 20, 6, and 3 assets, respectively. The systemic shock which happened at 13.01 is an extreme event which triggers all the assets. During the first iteration of the algorithm, it is correctly identified, however his shock propagates via the self-exciting kernels to the minutes in the near future and affects the intensities of all the N point processes. For this reason, the event happening at 13.02, when six assets cojumped, is not recognized as extreme. When the first iteration is completed and the event at 13.01 is removed, we recalibrate the N Hawkes processes and proceed with a second iteration. This time the event at 13.02 is successfully identified, but again its effect may condition the time points which follow, as confirmed by the third iteration.

The advantages of the above procedure are manifold. Specifically, the set of events $\{t_i^F\}$ with $i = 1, \dots, n_F$ identifies the proxy of the common factor, but not less important each reduced set $\{t_{i'}^s\} \subseteq \{t_i^s\}$ for $s = 1, \dots, N$ and $i' = 1, \dots, n_{i'}$ corresponds in a natural way to the idiosyncratic component. The latter result represents a second major improvement with respect to the bivariate model. Moreover, the entire procedure is easy scalable, since it does not depend critically on the dimension of the portfolio. Last, but not least, we are not a priori fixing the nature of the point processes which describe the factor and the idiosyncratic components. We can apply the log-likelihood approach described in 6 to the univariate sequence of the event times, fix the value of the parameters of the factor, λ_F , α_F and β_F , and of the idiosyncratic components, λ_s , α_s and β_s , for $s = 1, \dots, N$, estimate the associated p -values, and, eventually, if it is the case, reject the Hawkes process in favour of the Poisson description.

To sum up, the multivariate model describes the extreme events which occur in a portfolio of stocks as a superposition of systemic shocks, propagating to the s -th asset with probability $\tilde{p}_s$, and jumps specific of the single assets. In order to estimate the vector of probabilities

Iteration	$\{t_i^F\}$										$\bar{J}_{t_i^F}$									
1	1191	1326	1589	2351	4442	4716	5398	5640	6594	6714	3	4	4	3	4	3	3	3	3	3
	7625	7880	8262	9232	9506	13461	14074	15211	15730	15755	3	4	3	3	3	3	3	3	3	3
	15756	15922	16331	16555	17790	18511	18687	19928	20011	20411	4	3	3	3	3	3	3	3	3	3
	20506	20539	20547	21397	21402	21596	22446	24597	25056	25231	3	3	4	3	4	3	3	3	3	3
	25671	26002	26889	27030	27891	28439	28607	29403	30036	30043	3	3	3	3	3	3	3	3	5	
	30081	30518	31136	31591	31594	32161	32237	32606	33065	33066	3	3	3	3	4	3	4	4	3	6
	33319	33634	34303	34336	35566	36110	38779	39100	39101	39195	3	3	3	3	3	3	3	3	4	3
	39779	40080	40461	40536	40657	40664	40856	41831	42006	42995	3	3	3	3	3	4	3	3	3	3
	43033	43167	43212	43224	43265	43620	43761	43762			3	4	3	4	3	3	4	4		
2	1409	3687	6597	8148	30037	30520	31138	33211	34362	40658	3	3	3	3	3	3	3	3	3	3
	43213										3									
3	30038										3									

Table 7: Event times and thresholds identified by the multiple hypothesis test with significances 1%. We comment on events in bold in the main text.

$\tilde{\mathbf{p}} = (\tilde{p}_1, \dots, \tilde{p}_N)'$, we replace the system of equations given by (9), with the new relations

$$\begin{aligned} \tilde{p}_1 \frac{\lambda_F}{1 - \alpha_F/\beta_F} T &= n_1 - n_{1'}, \\ &\vdots \\ \tilde{p}_N \frac{\lambda_F}{1 - \alpha_F/\beta_F} T &= n_N - n_{N'}, \end{aligned}$$

where $\lambda^F T / (1 - \alpha_F/\beta_F)$ corresponds the average number of factor's jumps. The above equations force the expected number of cojumps among the factor and the s -th stock to balance the realized number of shocks.




Figure 10: From the top left clockwise: MJ probability test under N factor model null for the asset Generali; CJ probability test for the pairs Generali-Mediobanca, Generali-Banca Popolare Milano, and Generali-Intesa Sanpaolo.

We test our model on the data set of twenty stocks analysed in the first part of this paper, and we show in Figure 10 the results of the tests performed over the assets Generali, Mediobanca, Banca Popolare di Milano, and Intesa Sanpaolo. The plots confirm the ability to capture both the self and the cross dependence among jumps, which is a remarkable feature of a factor model easy to implement and calibrate on the data. Moreover, by construction, the model is genuinely scalable, and this fact makes it a viable alternative to more sophisticated but computationally complex models. In Table 8 we present the value of the parameters of the common factor, with the associated standard errors and p -values. From the measured confidence for the scale and decay length we can reject the null hypothesis of a Poisson model. In Table 9 we report the remaining parameters for the idiosyncratic components.

It is important to notice that we have performed a Monte Carlo experiment by drawing $N_{MC} = 10^3$ scenarios with the values given in Tables 8 and 9 and we have then re-estimated the parameters on each copy. This experiment has shown that the entire procedure is quite

	Hawkes factor
λ_F	$(2.0 \pm 0.2) \times 10^{-3}$, p -value: < 0.001
α_F	$(4.9 \pm 1.9) \times 10^{-2}$, p -value: 0.0105
β_F	$(3.3 \pm 1.1) \times 10^{-1}$, p -value: 0.0021

Table 8: List of the parameters of the common factor process.

robust, but the comparison with the process N_t^F , which is known in the artificial Monte Carlo framework, evidences a systematic underestimation of the number of jumps of the common factor. The mis-identification is due to the fact that even when the factor jumps, there is always a small, but finite, probability that only a small number of assets below the detectability threshold jumps too. In this case, our methodology does not detect a systemic event. The mis-identification probability depends on the true values of $(\tilde{p}_1, \dots, \tilde{p}_N)$, and for those fixed in the Monte Carlo experiment, it amounts approximately to 7.5%. Nonetheless, the bias reduces when the number of assets increases, and tends to zero for $N \rightarrow \infty$.

8 Conclusions and Perspectives

The detection techniques that we develop in Sections 2 and 3 show that a large number of jumps is present in financial time series. Even though the identification process of the extreme events suffers from some dependence on the details of the detection method, we believe that the idea of intersecting the different methodologies partially amends this drawback. Relying on the correct identification of jumps, we find that, as far as individual stocks are concerned, jumps are clearly not described by a Poisson process. The evidence of time clustering can be accounted for and modelled by means of linear self-exciting Hawkes processes. Moving to a cross-sectional perspective, we identify a significant number of systemic events, especially simultaneous cross jumps, that can not be reduced to a purely random effect. We have provided quantitative arguments against the idea of modelling this effect in terms of multidimensional Hawkes processes. The simultaneity of events is not captured by this class of point processes, and the increase of dimensionality of the parameters space associated with the multivariate model is discouraging. In Section 7.2 we propose a one factor model which is able to describe the main features that characterize the departure from a random behaviour of jumps, namely, the time clustering of jumps on individual stocks, the large number of simultaneous systemic jumps, and the time lagged cross excitation between different stocks.

This work opens interesting perspectives for future research. It would be of great interest to see whether these results would be confirmed by repeating the analysis on a wider collection of assets. Promising directions could be enlarging the number of investigated stocks, extending the analysis to other markets, or also considering different but related securities, such as equities, futures, options and all derivative products. Another extension of this work could be the study of the properties of jumps for a given group of securities in different periods, to assess for example whether changes in the regulations of a market have an impact on the happening of systemic events.

It would also be interesting to understand to what extent the jump detection procedure that we use in this work is sensitive to the choice of building returns from last executed prices. Other possible choices for a comparison could be employing ranges in place of returns or building the

ISIN	$\tilde{p}_s$	$\lambda_s (\times 10^{-3}), p\text{-value}$	$\alpha_s (\times 10^{-2}), p\text{-value}$	$\beta_s (\times 10^{-1}), p\text{-value}$
IT00000062072	0.31	$1.4 \pm 0.2, (*)$	$2.6 \pm 1.4, 0.067$	$2.2 \pm 1.0, 0.030$
IT00000062957	0.12	$1.1 \pm 0.2, (*)$	$1.1 \pm 1.3, 0.401$	$2.3 \pm 2.6, 0.386$
IT00000064482	0.10	$2.2 \pm 0.2, (*)$	$3.7 \pm 1.6, 0.020$	$2.6 \pm 1.1, 0.014$
IT00000068525	0.24	$1.2 \pm 0.2, (*)$	$0.9 \pm 0.4, 0.020$	$0.3 \pm 0.1, 0.021$
IT00000072618	0.48	$1.6 \pm 0.2, (*)$	$0.6 \pm 0.6, 0.277$	$1.0 \pm 1.1, 0.365$
IT0001063210	0.08	$1.0 \pm 0.2, (*)$	$0.6 \pm 0.5, 0.174$	$0.6 \pm 0.4, 0.125$
IT0001334587	0.17	$2.9 \pm 0.3, (*)$	$6.2 \pm 1.8, 0.001$	$3.0 \pm 0.8, (*)$
IT0001976403	0.31	$1.7 \pm 0.2, (*)$	$2.8 \pm 1.2, 0.015$	$1.7 \pm 0.6, 0.006$
IT0003128367	0.39	$2.6 \pm 0.3, (*)$	$2.4 \pm 0.8, 0.004$	$1.2 \pm 0.4, 0.002$
IT0003132476	0.37	$2.2 \pm 0.2, (*)$	$1.3 \pm 0.5, 0.016$	$0.9 \pm 0.3, 0.012$
IT0003487029	0.18	$1.0 \pm 0.2, (*)$	$5.1 \pm 2.8, 0.069$	$3.9 \pm 1.7, 0.019$
IT0003497168	0.31	$1.8 \pm 0.2, (*)$	$0.7 \pm 0.4, 0.126$	$0.4 \pm 0.3, 0.232$
IT0003856405	0.11	$1.6 \pm 0.2, (*)$	$0.8 \pm 0.4, 0.087$	$0.7 \pm 0.3, 0.058$
IT0004176001	0.14	$0.9 \pm 0.2, (*)$	$0.7 \pm 0.3, 0.024$	$0.2 \pm 0.1, 0.018$
IT0004231566	0.16	$1.7 \pm 0.2, (*)$	$1.8 \pm 0.9, 0.056$	$1.5 \pm 0.7, 0.031$
IT0004623051	0.15	$1.8 \pm 0.2, (*)$	$2.4 \pm 0.9, 0.008$	$1.0 \pm 0.4, 0.006$
IT0004644743	0.21	$1.4 \pm 0.2, (*)$	$3.6 \pm 1.5, 0.014$	$1.9 \pm 0.7, 0.003$
IT0004781412	0.37	$1.4 \pm 0.2, (*)$	$1.0 \pm 0.4, 0.025$	$0.5 \pm 0.2, 0.007$
LU0156801721	0.17	$0.8 \pm 0.1, (*)$	$1.0 \pm 0.7, 0.147$	$0.7 \pm 0.4, 0.092$
NL0000226223	0.22	$1.0 \pm 0.2, (*)$	$3.2 \pm 1.6, 0.048$	$1.2 \pm 0.7, 0.073$

Table 9: List of the parameters of the idiosyncratic components; $(*) < 0.001$.

latter from the mid price. The latter possibility deserves particular attention, since the mid price is always defined and there would not be any of the problems we face with the missing observations of executed price. In this context, deeper research could be made on the frequency of data sampling, an issue we never discuss in the present study.

Moving from a descriptive point of view to one that investigates the origin of the behaviour of jumps and cojumps, the research direction we consider more promising is the study of the order book in proximity of extreme systemic events, that is, cojumps involving a large number of assets. The perspective to take should be to explore the cross direction at a fixed time, more than the time direction for a single asset. The bid-ask spread, the depth of the book, the asymmetry of buy and sell volumes are all quantities whose dynamics may reveal interesting features related to systemic events.

The high level of synchronization between the jumps of different assets that we empirically observe in our data calls for possible explanations. Even if our paper is mainly methodological, we believe that some comments are needed in order to explain this fact observed in (modern) financial markets. Financial markets are becoming increasingly interconnected at a high speed due to several reasons, first of all the increased automation of the trading process and of the information processing. High frequency trading strategies, statistical arbitrageurs, hedging strategies could be partly responsible of the large number of cojumps we observe in our sample. Certainly, we believe, a proper modelling of the jump process in a systemic context is important for regulators and for investors in order to assess in a reliable way the level of risk of a market or of a large portfolio of assets.

Acknowledgements

GB and LMC acknowledge the support of the Scuola Normale Superiore grant GR12CALCAG “Entropia di Shannon ed efficienza informativa dei mercati finanziari”. FL acknowledges partial support from grant SNS11LILLB “Price formation, agents heterogeneity, and market efficiency”. FC, FL, and SM acknowledge partial support of the European Union Seventh Framework Programme FP7/2007-2013 under grant agreement CRISIS-ICT-2011-288501. MT acknowledges Flavio Baronti and Francesco Pantaleo for technical support. Finally, the authors wish to thank Enrico Melchioni for the constant encouragement.

A Data Cleaning

A.1 Outlier and split/merge detection

To detect and remove anomalous values (often referred to as outliers) that may be present in recorded data, we use the cleaning algorithm proposed by Brownlees and Gallo in [43]. The algorithm identifies the price records at tick-by-tick level which are too distant from a mean value calculated in their neighbourhood. More precisely, a price observation p_i is regarded as anomalous and removed if

$$|p_i - \bar{p}_i(k)| \geq c s_i(k) + \gamma,$$

where $\bar{p}_i(k)$ and $s_i(k)$ denote respectively the δ -trimmed sample mean and sample standard deviation of the closest k observations around i , c is a constant which amplifies the standard deviation and γ is a granularity parameter useful in the case of k equal prices producing a zero variance. We take $k = 60$, $\delta = 10\%$, $c = 3$, $\gamma = 0.05$.

A *forward stock split* (or, simply, a *stock split*) is an operation consisting in an increase of the number of shares of a company and in a simultaneous adjustment of the price so that the market capitalization of the company remains the same. A *reverse stock split* (also called a




Figure 11: The distributions of intertrade times for stocks Fiat (left panel) and Telecom Italia (right panel).

stock merge) is the opposite operation, leading to a decrease of the number of shares and to a corresponding increase of the price. A split is *m-for-n* if m new shares are released for every n old ones, with a simultaneous price change from p to pn/m .

A.2 Volatility auctions

According to the rules stated in *Rules of the markets organised and managed by Borsa Italiana S.p.A.* and *Instructions accompanying the Rules of the markets organised and managed by Borsa Italiana S.p.A.* ([44, 45]), whenever a stock's price gets too far from a reference value, Borsa Italiana is obliged to start a volatility auction phase. During this phase, which has a duration between 10 and 11 minutes, trades are suspended and a new reference value for the price is sought. Possibly, if no valid price is reached during the volatility auction phase, another such phase starts immediately after the first one and so on. We treat all intertrade times of at least ten minutes as volatility auction periods (the distributions of intertrade times shown in Figure 11 suggest this is correct) and consider as *not available* the returns at the sampling times falling in these periods. In Figure 11 we report the empirical distributions of the intertrade times for the assets Fiat and Telecom Italia. In both histograms we can easily recognize the presence of volatility auctions, which manifests in terms of the peaks appearing on the far right tail in correspondence of intertrade times close to the multiples of eleven minutes.

A.3 Intraday pattern

We filter out the intraday volatility pattern from returns by means of a simple model with intraday volatility factors. Returns at intraday time t are rescaled by a factor λ_t , which is calculated as the average, over all days, of adjusted absolute returns at time t . More precisely, if $\tilde{r}_{d,t}$ is the raw return of day d and intraday time t , we define the rescaled return

$$r_{d,t} = \frac{\tilde{r}_{d,t}}{\lambda_t},$$

where

$$\lambda_t = \frac{1}{N_{\text{days}}} \sum_{d'} \frac{|r_{d',t}|}{s_{d'}},$$

with N_{days} indicating the number of days in the sample and $s_{d'}$ the standard deviation of absolute intraday returns of day d' .




Figure 12: The intraday volatility pattern of Monte dei Paschi di Siena.




Figure 13: Autocorrelation of absolute returns of Monte dei Paschi di Siena before and after filtering the intraday volatility pattern.

We show in Figure 12 the intraday volatility pattern (that is, the factors λ_t) for Monte dei Paschi di Siena. We notice that, although the profile is somewhat noisy because of the relatively small number of days on which the average pattern is calculated, a clear change in average intraday volatility is identifiable around time 15:30, that is, when the New York Stock Exchange opens. In Figure 13 we show the effectiveness of the procedure of removing the seasonality at daily level, through the plot of the autocorrelations of absolute intraday returns.

A.4 Volatility proxy

For estimating the local volatility, we use the *realized absolute variation* and the *realized bipower variation* [55, 47]. Let the logarithmic prices $p(t)$ be generated by a process

$$dp(t) = \mu(t) dt + \sigma(t) dW(t), \quad (10)$$

with $\mu(t)$ a finite variation process, $\sigma(t)$ a càdlàg volatility process, $W(t)$ a standard Brownian motion. Let the interval $[0, t]$ be divided into subintervals of the same length δ and denote by r_i the return $p(i\delta) - p((i-1)\delta)$. Then the following probability limits hold for the *realized absolute variation* and the *realized bipower variation*:

$$p - \lim_{\delta \searrow 0} \delta^{\frac{1}{2}} \sum_{i=1}^{\lfloor t/\delta \rfloor} |r_i| = \mu_1 \int_0^t \sigma(s) ds,$$

$$p - \lim_{\delta \searrow 0} \sum_{i=1}^{\lfloor t/\delta \rfloor - 1} |r_i| |r_{i+1}| = \mu_1^2 \int_0^t \sigma^2(s) ds, \quad (11)$$

where $\mu_1 = E(|u|) = \sqrt{\frac{2}{\pi}} \simeq 0.797885$, $u \sim \mathcal{N}(0, 1)$. The asymptotic result (11) continues to hold when a jump component is added to the continuous process (10). However, for finite δ , price jumps are a source of bias in estimation of volatility through the realized bipower variation. A solution to this problem is to use the *threshold bipower variation* (see [7]), which takes into account only returns smaller than a certain threshold. We follow this idea and estimate volatility through returns which are not identified as jumps, that is, returns whose absolute value is not larger than 4 times the local volatility. Using exponentially weighted moving averages instead of flat ones, our estimators are recursively defined by the equations

$$\hat{\sigma}_{\text{abs},t} = \mu_1^{-1} \alpha |r_{t'}| + (1 - \alpha) \hat{\sigma}_{\text{abs},t-1},$$

where t' is such that $t' \leq t - 1$, $\frac{|r_{t'}|}{\hat{\sigma}_{\text{abs},t'}} \leq 4$ and $\frac{|r_\tau|}{\hat{\sigma}_{\text{abs},\tau}} > 4$ for each $t' < \tau < t$, and

$$\hat{\sigma}_{\text{bv},t}^2 = \mu_1^{-2} \alpha |r_{t''}| |r_{t'}| + (1 - \alpha) \hat{\sigma}_{\text{bv},t-1}^2,$$

where t'' and t' are such that $t'' < t' \leq t - 1$, $\frac{|r_{t''}|}{\hat{\sigma}_{\text{bv},t''}} \leq 4$, $\frac{|r_{t'}|}{\hat{\sigma}_{\text{bv},t'}} \leq 4$ and $\frac{|r_\tau|}{\hat{\sigma}_{\text{bv},\tau}} > 4$ for each $t'' < \tau < t'$ and for each $t' < \tau < t$. The value taken for the parameter of the exponentially weighted moving average is $\alpha = \frac{2}{M+1}$, with $M = 60$, corresponding to a characteristic time of 30.5 minutes.

References

- [1] Y. Aït-Sahalia, “Disentangling diffusion from jumps,” *Journal of Financial Economics*, vol. 74, pp. 487–528, 2004.
- [2] O. E. Barndorff-Nielsen and N. Shephard, “Econometrics of testing for jumps in financial economics using bipower variation,” *Journal of Financial Econometrics*, vol. 4, pp. 1–30, 2006.
- [3] G. Jiang and R. Oomen, “Testing for jumps when asset prices are observed with noise—a “swap variance” approach,” *Journal of Econometrics*, vol. 144, pp. 352–370, 2008.
- [4] S. Lee and P. Mykland, “Jumps in financial markets: A new nonparametric test and jump dynamics,” *Review of Financial studies*, vol. 21, p. 2535, 2008.
- [5] Y. Aït-Sahalia and J. Jacod, “Testing for jumps in a discretely observed process,” *Annals of Statistics*, vol. 37, pp. 184–222, 2009.
- [6] S. Lee and J. Hannig, “Detecting jumps from Lévy jump diffusion processes,” *Journal of Financial Economics*, vol. 96, pp. 271–290, 2010.
- [7] F. Corsi, D. Pirino, and R. Renò, “Threshold bipower variation and the impact of jumps on volatility forecasting,” *Journal of Econometrics*, vol. 159, pp. 276–288, 2010.
- [8] T. Bollerslev, T. H. Law, and G. Tauchen, “Risk, jumps and diversification,” *Journal of Econometrics*, vol. 144, pp. 234–256, 2008.
- [9] J. Maheu and T. McCurdy, “News arrival, jump dynamics and volatility components for individual stock returns,” *Journal of Finance*, vol. 59, pp. 755–793, 2004.

- [10] T. Bollerslev, U. Kretschmer, C. Pigorsch, and G. Tauchen, “A Discrete-Time Model for Daily S&P500 Returns and Realized Variations: Jumps and Leverage Effects,” *Journal of Econometrics*, vol. 150, pp. 151–166, 2009.
- [11] T. Andersen, T. Bollerslev, and D. Dobrev, “No-arbitrage semi-martingale restrictions for continuous-time volatility models subject to leverage effects, jumps and iid noise: Theory and testable distributional implications,” *Journal of Econometrics*, vol. 138, pp. 125–180, 2007.
- [12] S. Das and R. Uppal, “Systemic Risk and International Portfolio Choice,” *The Journal of Finance*, vol. 59, pp. 2809–2834, 2004.
- [13] J. Wright and H. Zhou, “Bond risk premia and realized jump volatility.” Working paper, Federal Reserve Board, 2007.
- [14] D. Duffie, J. Pan, and K. Singleton, “Transform analysis and asset pricing for affine jump-diffusions,” *Econometrica*, vol. 68, pp. 1343–1376, 2000.
- [15] B. Eraker, M. Johannes, and N. Polson, “The impact of jumps in equity index volatility and returns,” *Journal of Finance*, vol. 58, pp. 1269–1300, 2003.
- [16] B. Eraker, “Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices,” *The Journal of Finance*, vol. 59, pp. 1367–1404, 2004.
- [17] T. Bollerslev, V. Todorov, and S. Li, “Jump tails, extreme dependencies, and the distribution of stock returns,” *Journal of Econometrics*, 2012.
- [18] T. Bollerslev and V. Todorov, “Tails, fears, and risk premia,” *The Journal of Finance*, vol. 66, pp. 2165–2211, 2011.
- [19] R. Cont and P. Tankov, *Financial Modelling with Jump Processes*. Chapman & Hall - CRC, 2004.
- [20] O. E. Barndorff-Nielsen and N. Shephard, “Variation, jumps, market frictions and high frequency data in financial econometrics,” in *Advances in Economics and Econometrics. Theory and Applications, Ninth World Congress*, pp. 328–372, Cambridge University Press, 2007.
- [21] M. Dungey, M. McKenzie, and L. Smith, “Empirical evidence on jumps in the term structure of the us treasury market,” *Journal of Empirical Finance*, vol. 16, pp. 430–445, 2009.
- [22] J. Lahaye, S. Laurent, and C. Neely, “Jumps, cojumps and macro announcements,” *Journal of Applied Econometrics*, vol. 26, pp. 893–921, 2011.
- [23] M. Dungey and L. Hvozdyk, “Cojumping: Evidence from the us treasury bond and futures markets,” *Journal of Banking & Finance*, 2012.
- [24] D. Gilder, M. Shackleton, and S. Taylor, “Cojumps in stock prices: Empirical evidence,” 2012. Available at SSRN 2028443.
- [25] F. Gobbi and C. Mancini, “Diffusion covariation and co-jumps in bidimensional asset price processes with stochastic volatility and infinite activity lévy jumps,” in *SPIE Fourth International Symposium on Fluctuations and Noise*, pp. 660111–660111, International Society for Optics and Photonics, 2007.

- [26] J. Jacod and V. Todorov, “Testing for common arrivals of jumps for discretely observed multidimensional processes,” *The Annals of Statistics*, vol. 37, pp. 1792–1838, 2009.
- [27] Y. Aït-Sahalia, J. Cacho-Diaz, and R. Laeven, “Modeling financial contagion using mutually exciting jump processes,” tech. rep., National Bureau of Economic Research, 2012.
- [28] *Findings regarding the market events of May 6, 2010. Report of the staffs of the CFTC and SEC to the joint advisory committee on emerging regulatory issues.* 2010. <http://www.sec.gov/news/studies/2010/marketevents-report.pdf>.
- [29] A. A. Kirilenko, A. S. Kyle, M. Samadi, and T. Tuzun, “The flash crash: The impact of high frequency trading on an electronic market,” 2011. Available at SSRN 1686004.
- [30] *Flash Crash Analysis Continuing Developments.* <http://www.nanex.net>.
- [31] N. Johnson, G. Zhao, E. Hunsader, J. Meng, A. Ravindar, S. Carran, and B. Tivnan, “Financial black swans driven by ultrafast machine ecology,” 2012. Available at SSRN 2003874.
- [32] A. G. Hawkes, “Spectra of some self-exciting and mutually exciting point processes,” *Biometrika*, vol. 58, pp. 83–90, 1971.
- [33] D. Vere-Jones, “Stochastic models for earthquake occurrence,” *Journal of the Royal Statistical Society, Series B*, vol. 32, pp. 1–62, 1970.
- [34] D. Vere-Jones and T. Ozaki, “Some examples of statistical inference applied to earthquake data,” *Annals of the Institute of Statistical Mathematics*, vol. 34, pp. 189–207, 1982.
- [35] Y. Ogata, “Statistical models for earthquake occurrences and residual analysis for point processes,” *Journal of the American Statistical Association*, vol. 83, pp. 9–27, 1988.
- [36] C. G. Bowsher, “Modelling security markets in continuous time: Intensity based, multivariate point process models,” *Journal of Econometrics*, vol. 141, pp. 876–912, 2007.
- [37] L. Bauwens and N. Hautsch, “Modelling financial high frequency data using point processes,” in *Handbook of Financial Time Series* (T. Mikosch, J.-P. Kreiss, R. A. Davis, and T. G. Andersen, eds.), pp. 953–979, Springer, Berlin, 2009.
- [38] I. Muni Toke, ““Market making” in an order book model and its impact on the spread,” in *Econophysics of Order-Driven Markets* (F. Abergel, B. K. Chakrabarti, A. Chakraborti, and M. Mitra, eds.), pp. 49–64, Springer-Verlag, Milan, 2011.
- [39] I. Muni Toke and F. Pomponio, “Modelling trades-through in a limit order book using Hawkes processes,” *Economics: The Open-Access, Open-Assessment E-Journal*, vol. 6, pp. 1–23, 2012.
- [40] E. Bacry, S. Delattre, M. Hoffmann, and J. F. Muzy, “Modelling microstructure noise with mutually exciting point processes,” *Quantitative Finance*, vol. 13, pp. 65–77, 2013.
- [41] D. Filimonov and D. Sornette, “Quantifying reflexivity in financial markets: Toward a prediction of flash crashes,” *Physical Review E*, vol. 85, pp. 056108–1–9, 2012.
- [42] Y. Aït-Sahalia and T. R. Hurd, “Portfolio choice in markets with contagion,” 2012. Available at Arxiv.org:1210.1598[q-fin.PM].

- [43] C. T. Brownlee and G. M. Gallo, “Financial econometric analysis at ultra-high frequency: Data handling concerns,” *Computational Statistics & Data Analysis*, vol. 51, pp. 2232–2245, 2006.
- [44] *Rules of the markets organised and managed by Borsa Italiana S.p.A.* Borsa Italiana S.p.A., 2012. <http://www.borsaitaliana.it/borsaitaliana/regolamenti/regolamenti/regolamentoborsa-istruzionalregolamento.en.htm>.
- [45] *Instructions accompanying the Rules of the markets organised and managed by Borsa Italiana S.p.A.* Borsa Italiana S.p.A., 2012. <http://www.borsaitaliana.it/borsaitaliana/regolamenti/regolamenti/regolamentoborsa-istruzionalregolamento.en.htm>.
- [46] A. Joulin, A. Lefevre, D. Grunberg, and J.-P. Bouchaud, “Stock price jumps: news and volume play a minor role,” *Wilmott Magazine*, vol. 46, 2008.
- [47] O. E. Barndorff-Nielsen and N. Shephard, “Power and bipower variation with stochastic volatility and jumps,” *Journal of Financial Econometrics*, vol. 2, pp. 1–48, 2004.
- [48] A.-M. Dumitru and G. Urga, “Jumps in financial assets: A comparison between non-parametric jump tests,” *Journal of Business and Economic Statistics*, *Forthcoming*, 2011. Available at SSRN: <http://ssrn.com/abstract=1944221>.
- [49] J. Fan and Y. Wang, “Multi-scale jumps and volatility analysis for high-frequency financial data,” *Journal of the American Statistical Association*, vol. 102, pp. 1349–1362, 2007.
- [50] D. J. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes Volume I: Elementary Theory and Methods*. Springer, Heidelberg, 2003.
- [51] Y. Ogata, “The asymptotic behaviour of maximum likelihood estimates for stationary point processes,” *Annals of the Institute of Statistical Mathematics*, vol. 30, pp. 243–261, 1978.
- [52] T. Ozaki, “Maximum likelihood estimation of Hawkes self-exciting point processes,” *Annals of the Institute of Statistical Mathematics*, vol. 31, pp. 145–155, 1979.
- [53] Y. Ogata, “On Lewis’ simulation method for point processes,” *IEEE Transactions of Information Theory*, vol. 27, pp. 23–31, 1981.
- [54] P. A. W. Lewis and G. S. Shedler, “Simulation of nonhomogeneous Poisson processes with log linear rate function,” *Biometrika*, vol. 63, pp. 501–506, 1976.
- [55] T. G. Andersen and T. Bollerslev, “Intraday periodicity and volatility persistence in financial markets,” *Journal of Empirical Finance*, vol. 4, pp. 115–158, 1997.