



Financial contagion through space-time point processes

Giada Adelfio¹ · Arianna Agosto² · Marcello Chiodi¹ · Paolo Giudici² 

Accepted: 1 July 2020 / Published online: 8 July 2020
© The Author(s) 2020

Abstract

We propose to study the dynamics of financial contagion by means of a class of point process models employed in the modeling of seismic contagion. The proposal extends network models, recently introduced to model financial contagion, in a space-time point process perspective. The extension helps to improve the assessment of credit risk of an institution, taking into account contagion spillover effects.

Keywords Contagion models · Credit risk · Space-time point processes

JEL Classification G01 · C58 · C63

1 Introduction

The global financial crisis has led to an increasing research literature on systemic risk, with different definitions and measurement models.

According to the European Central Bank, “Systemic risk is the risk of experiencing a strong systemic event, which adversely affects a number of systemically important intermediaries or markets”. This definition introduces two key elements for the study of systemic risk: first, a trigger event, transmitted to the system as a whole, and not only to individual institutions; second, financial contagion, which is spread from each financial institution to another, in a high-order “domino” effect.

While systemic risk definitions share this broad view, systemic risk models are still quite divergent. For an introduction and a review see, for instance, Allen and Carletti (2006) and Jorion and Zhang (2009).

✉ Paolo Giudici
paolo.giudici@unipv.it

¹ Department of Economics, Management and Statistics, University of Palermo, Palermo, Italy

² Department of Economics and Management, University of Pavia, Pavia, Italy

Econometric measures of systemic risk have been proposed in particular by Adrian and Brunnermeier (2011), Acharya et al. (2012) and Brownlees and Engle (2012). On the basis of market share prices, these authors consider systemic risk as determined in the market, and calculate the quantiles of the estimated loss probability distribution of a financial institution, conditional on an extreme event in the financial market (or vice-versa). This approach is useful to establish policy thresholds aimed, in particular, at identifying the most systemic financial institutions. However, it is a bivariate approach, which allows to calculate the risk of an institution conditional on a reference market but which does not address the issue of how risks are transmitted between different institutions, in a multivariate framework.

A different stream of research considers systemic risks as exogenous, explained by causal factors. This is the approach proposed, in particular, by Ang and Longstaff (2012) and Betz et al. (2014), who explain whether the default probability of an institution depends on a set of exogenous risk sources. While powerful from an early warning perspective, causal models, similarly to bivariate ones, concentrate on single institutions rather than on the economic system as a whole and, therefore, may underestimate systemic sources of risk arising from contagion effects within the system.

Trying to address the multivariate nature of systemic risk, researchers have recently proposed financial network models, able to combine the rich structure of networks [see, e.g. Lorenz et al. (2009) and Battiston et al. (2012)] with a more parsimonious approach that can estimate contagion effects from the dependence structure among market prices. Seminal contributions in this framework are Billio et al. (2012) and Diebold and Yilmaz (2014), who derive contagion measures based on Granger-causality tests and variance decompositions. Recently, Ahelegbey et al. (2015) and Giudici and Spelta (2016) extended this methodology introducing correlation network models.

While bivariate and causal models explain how the risk of an institution is affected by the market, or by exogeneous risk factors, correlation network models explain whether the same risk depends on endogenous contagion effects, in a cross-sectional perspective. While bivariate and causal models can signal when a crisis could happen, on the basis of the exogeneous triggers, but do not describe how that crisis is transmitted, correlation network models are not predictive, but can describe how a crisis, once triggered, is transmitted through the whole system.

Our aim is to improve financial network models, allowing them to be a predictive and not only a descriptive measure of systemic risk. To this aim, we employ point process models that combine the space dimension, characteristic of cross-sectional models, with a time dimension. Doing so, we contribute to the financial networks literature, extending systemic risk models with a new model, based on point processes, that takes both time and space dynamics into account.

A point process is a random collection of points, each representing the time and space coordinates of a single event. Point processes are widely used to describe observed patterns in different fields, such as the locations of infected patients in epidemiology, the positions of human settlements in geography, and the epicentres

of earthquakes in seismology: for a review see, for example, Diggle (1983), Ripley (1977) and Schabenberger and Gotway (2005).

In all the above situations, the main research issue is to establish whether the points exhibit complete spatial randomness, so that they can be taken as realisations of a simple Poisson process, or if they are characterised by space-time aggregation. The aggregation may be due to the clustering of the points around a stochastic process, which leads to a strong interaction among points themselves.

To model events that are clustered together, self-exciting point processes, such as the Hawkes model (Hawkes and Adamopoulos 1973) or the Epidemic Type Aftershock Sequence (ETAS) model (Ogata 1988) have been proposed in the literature. These models are largely used to describe phenomena for which the occurrence of an event increases the probability of occurrence of other events, close in time and/or space.

In finance, while Hawkes type models are already in use, as discussed in Bauwens and Hautsch (2009), ETAS models are much less known. In this paper we propose to customise ETAS models to understand financial contagion. In seismic applications, the occurrence of an earthquake shock increases the likelihood of other shocks, in the close regions. In our financial context, the occurrence of “a shock” on a financial institution increases the likelihood of other shocks, in the “close” institutions.

We aim to extend and adapt ETAS models to the financial context. This involves two main methodological challenges: (i) to provide a definition of what a financial shock mean, and a related measure, that can mimic the well known magnitude scale employed for earthquakes; (ii) to provide an adequate measure of distance between institutions that can mimic the physical distance.

We contribute to the methodological literature on ETAS models addressing the above challenges.

To solve (i) we propose, to employ, as a measure of financial shock on an institution, the occurrence that the market price of its Credit Default Swap (CDS) exceeds an extreme threshold, defined by the overall Value at Risk. We remark that a CDS expresses which premium is priced by financial markets to insure the investors against the default risk of an institution. It is therefore an “objective” way to measure the occurrence of a financial shock on an institution. Assuming that the market correctly prices the overall credit risk of an institution, by means of the risk premium implicit in the CDS price, we will try to separate, in such risk premium, the idiosyncratic component, due to the individual characteristics of an institution, from the systemic component, due to contagion from the other institutions. Doing so, the importance of each institution will depend not only on its position, as in most financial network models, but also on its risk magnitude (expressed by the CDS price) and on the risk magnitude of its neighbours.

To solve (ii) we propose to employ a multidimensional scaling technique on the observed vector of CDS prices of each institution. According to the efficient market assumption, financial market prices are expected to contain all relevant information on the underlying institutions. By comparing the time series of CDS prices, we can get a measure of the perceived distance in credit risk between institutions, along T time points. We propose to project such distance on a two dimensional space, which

can be visualised and better understood than a T dimensional space. Multidimensional scaling allows to minimize the difference between the original distances and those expressed in the two dimensional space.

To summarise, our paper contains both a methodological contribution, that allows the extension of ETAS models to the financial context; and an applied contribution, as it extends financial network models in a predictive context.

The paper is organised as follows: in Sect. 2 we present the ETAS model that we propose to model financial contagion. In Sect. 3 we present the application of the model to real data that concern the largest European financial institutions. Some concluding remarks follow in Sect. 4.

2 Methodology

Let $Z = T \times S$ be a general space-time closed region, with T the time domain and S a two-dimensional space. For a generic point $\mathbf{z} = (t, \mathbf{s})$ of Z , the *conditional intensity function* (Daley and Vere-Jones 2003) is the expected frequency of events in a space-time unit region, conditional on the history \mathcal{H}_t of the point process up to time t , that is:

$$\lambda(\mathbf{z}) = \lambda(t, \mathbf{s} | \mathcal{H}_t) = \lim_{\mathcal{L}(\Delta t), \mathcal{L}(\Delta \mathbf{s}) \rightarrow 0} \frac{E[N([t, t + \Delta t] \times [\mathbf{s}, \mathbf{s} + \Delta \mathbf{s}]) | \mathcal{H}_t]}{\mathcal{L}(\Delta t)\mathcal{L}(\Delta \mathbf{s})}$$

where $\mathcal{L}(x)$ is the Lebesgue measure of x , and $\Delta t, \Delta \mathbf{s}$ are time and space increments respectively.

Conditional intensities $\lambda(\mathbf{z})$ are usually made dependent on some unknown parameters.

The parametric conditional intensity function of a branching model in a point (t, \mathbf{s}) can be defined as the sum of a term that describes the long-term variation (spontaneous activity or background) and a term that describes the short-term variation (induced activity or offsprings):

$$\lambda_\psi(t, \mathbf{s} | \mathcal{H}_t) = \mu f(\mathbf{s}) + \sum_{t_j < t} v_\theta(t - t_j, \mathbf{s} - \mathbf{s}_j) \tag{1}$$

with $\psi = (\theta, \mu)'$, the unknown parameters, describing short-term variation (θ) and long term variation (μ), with $f(\mathbf{s})$ and $\sum_{t_j < t} v_\theta(\cdot)$ the corresponding density functions.

Given $i = 1, \dots, n$ events, the spatial-temporal observations for each of them are the occurrence time t_i , the geographical coordinates $\mathbf{s}_i = (x_i, y_i)$ and the magnitude m_i , which represents the marker of the Point Process (Daley and Vere-Jones 2003). To jointly model all these observations, the ETAS specification is based on a parametric conditional intensity function that can be written as follows:

$$\lambda_\psi(t, \mathbf{s} | \mathcal{H}_t) = \mu f(\mathbf{s}) + \sum_{t_j < t} g(t - t_j | m_j) \ell(x - x_j, y - y_j) \tag{2}$$

where m_j is the magnitude of the j -th inducing event. Note that in the ETAS model

the short term component of the space-time point process is obtained as the product of a time intensity and a spatial intensity.

The long-term background component is instead usually assumed to have an intensity that is stationary in time (μ), but inhomogeneous in space ($f(\cdot)$).

The time occurrence rate of aftershocks at time t , following an event of time t_j of magnitude m_j , is then described by the following parametric density:

$$g(t - t_j | m_j) = \frac{\kappa_0 e^{(\alpha)(m_j - m_0)}}{(t - t_j + c)^p}, \quad \text{with } t > t_j \tag{3}$$

where m_0 is the completeness magnitude threshold, that is, the lower bound above which events are recorded; κ_0 is a constant which measures *aftershocks productivity*; c and p are parameters of the modified Omori’s law (Utsu 1961), with p indicating the decay rate of aftershocks in time.

Finally, the spatial rate of the ETAS model is usually described by the following density:

$$\ell(x - x_j, y - y_j) = \alpha \left\{ (x - x_j)^2 + (y - y_j)^2 + d \right\}^{-q} \tag{4}$$

which relates the occurrence rate of aftershocks to the magnitude m_j , through a parameter α that measures the influence of each sequence, and the parameters d and q which are related to related to the spatial influence of the shocks.

To estimate the parameters of the ETAS model, with the available event data, a main issue is the simultaneous estimate of all the different components of the intensity function. If the long-term density component $\mu f(\mathbf{s})$ in (2) is known, the remaining parameters ($\kappa_0, c, p, \alpha, d, q$) can be estimated by the Maximum Likelihood method, using computationally intensive numerical techniques (Ogata 1998). In particular, a procedure to maximize the log likelihood function, based on the expectation-maximization algorithm, was introduced in Veen and Schoenberg (2008). Otherwise, when the long-term spatial component $\mu f(\mathbf{s})$ in Eq. (2) is unknown, it can be obtained through kernel density estimators, using appropriate smoothing constant bandwidths (Silverman 1986).

We now introduce a method aimed at estimating the parameters of the ETAS model in the financial contagion context.

For this purpose, we assume that the intensity function $\lambda(\cdot)$ of a space-time point process, observed in the space region S and in the time interval $(T_0; T_{max})$, depends on a set of parameters ψ . Let also $\hat{\lambda}_{\psi; H_{t_k}}(\cdot)$ indicate a generic estimator of $\lambda(\cdot)$, based on observations until t_k .

Given a set of k observed values \mathbf{z}_i , the log likelihood for the ETAS process, calculated in $\hat{\lambda}_{\psi; H_{t_k}}(\cdot)$, becomes:

$$\log L(\hat{\lambda}_{\psi; H_{t_k}}(\mathbf{z}); H_{t_k}) = \sum_{i=1}^k \log \hat{\lambda}_{\psi; H_{t_k}}(\mathbf{z}_i) - \int_{T_0}^{T_{max}} \int_S \hat{\lambda}_{\psi; H_{t_k}}(\mathbf{z}) ds dt \tag{5}$$

We remark that the above likelihood cannot be obtained when the long term spatial component $\mu f(\mathbf{s})$ in Eq. (2) is unknown. To solve this problem we may estimate it

by means of non-parametric methods, such as kernel density, using a vector of smoothing parameters, representing window sizes. However, by doing so, the maximisation of the log likelihood of the point process is reached when window sizes tend to zero and the estimated intensities put mass only on the observed points, providing a useless result. To overcome this problem, we propose to employ the Forward Likelihood-Based (FLP) method developed in Chiodi and Adelfio (2011) and Adelfio and Chiodi (2015), which is based on the idea of maximising the sum of increments of the likelihood, obtained adding an observation at the time, to progressively take past information into account. The resulting approach, named ETAS-FLP, consists of a simultaneous estimation of the non-parametric and parametric components of the ETAS model, alternating standard likelihood methods (to estimate the parameters of the short-term component) with the FLP approach (to estimate the long-term component).

The new estimation strategy can be formally defined as follows. Let $\log L(\hat{\lambda}_{\psi; H_k}(\mathbf{z}); H_{t_{k+1}})$ be the log likelihood computed on the first $k + 1$ observations, applying the estimator $\hat{\lambda}_{\psi; H_k}(\mathbf{z})$ on the previous k observations.

The *predictive information* content of the first k observations on the $k + 1$ th is then defined as:

$$\delta_{k,k+1}(\psi) = \log L(\hat{\lambda}_{\psi; H_k}(\mathbf{z}); H_{t_{k+1}}) - \log L(\hat{\lambda}_{\psi; H_k}(\mathbf{z}); H_{t_k})$$

Note that this expression resembles an objective function used in out-of-sample cross-validation, but only applied to the immediately next observation.

Parameter estimates are finally obtained choosing $\hat{\psi}(H_{t_k})$ which maximises:

$$FLP_{k_1, k_2}(\hat{\psi}) \equiv \sum_{k=k_1}^{k_2} \delta_{k,k+1} \quad (6)$$

with the constants k_1 and k_2 to be appropriately chosen. In our applications we set, without loss of generality, $k_1 = \lfloor \frac{n}{2} \rfloor$ and $k_2 = n - 1$.

Chiodi and Adelfio (2011) showed that the above FLP procedure provides lower estimation errors of space-time intensity functions with respect to classical methods. Adelfio and Chiodi (2015) interpreted the FLP measure as a martingale and provided a proof of the asymptotic existence and uniqueness of the solution. They also showed applying the procedure to seismic contagion leads to rather stable results. To allow further testing and reproducibility of the methodology, the same authors implemented the algorithm in the R package `etasFLP` (Chiodi and Adelfio 2017).

The specific contribution of this paper is the extension of the ETAS-FLP approach, developed for the study of contagion effects of seismic shocks, to the study of contagion effects of financial shocks. To achieve this goal, we need to: (a) define a ‘response’ variable which defines and measures the magnitude of a financial shock; (b) replace ‘physical’ spatial coordinates, used in the seismic literature to measure contagion effects of seismic shocks, with ‘financial’ spatial coordinates that can measure contagion effects of financial shocks.

Addressing these two issues is the main methodological novelty of this paper. Concerning point a), we assume that corporate default swap (CDS) prices can be used to measure the magnitude of a financial shock. This seems reasonable, since CDS spreads correspond to the insurance premium that the market is willing to pay to insure against the default of the corresponding institution.

While the magnitude of a seismic shock is measured by the physical losses that materialise in a set of geographical sites, hit by an earthquake, the magnitude of a financial shock can be measured by the monetary losses incurred by investors in a set of specific institutions, hit by a crisis. While seismic losses are estimated through expert assessments, financial losses are estimated by the financial markets and are incorporated in market prices, such as CDS spreads.

The evolution of an earthquake is observed through a sequence of seismic shocks, correlated in time and space; similarly, we assume that the evolution of a financial crisis is observed through a sequence of CDS price variations, correlated in time and space. By assuming that the process defined by Eq. (2) is also suitable to describe contagion effects in the financial context, we can model the interconnection among different institutions, accounting for both long-term and short-term dependence.

Concerning issue b), a key challenge for extending the ETAS model from seismic to financial contagion is to get appropriate ‘financial’ spatial coordinates to replace the geographical ones. To this aim, in Sect. 3 we apply the multidimensional scaling approach [see, for instance, Torgerson (1958); Cox (2001) and Gower (1966)] which allows to define reference coordinates for the available corporate institutions in a low dimensional space, which, for simplicity and without loss of generality, will be chosen as bidimensional. In such a space, the distances between any two institutions will be proportional to the distance between the vectors of their observed CDS prices, along the whole considered time period. This way, we define an ‘endogenous’ system of coordinates that reproduces, in a bidimensional space, the distance between corporates estimated by the financial markets through the daily CDS prices.

Both in the seismic and in the financial context, statistical models provide useful tools not only related to the ‘description’ of the events, but also to the development of ‘early warning’ indicators of distress. To this aim, we need to further extend the ETAS-FLP procedure, making it able to provide predictions of the future CDS prices. This extension can be summarised by the following pseudo-code:

- Start from the estimated ETAS model, based on the FLP estimation using data up to time t .
- Forecast the intensity of the following day, that is, compute the intensity at time $t + dt$, conditional on the history up to t , $\hat{\lambda}_{\psi;H_t}(t + dt, \mathbf{s})$ (in our context dt is one day).
- Re-estimate the model based on the predicted intensity and repeat the steps for a period of n days.

A further extension is based on the following remark. To improve prediction for events occurring in the time interval $t + dt$, we should account for the effect of the

events that may occur meanwhile, conditionally on the history up to H_{t+dt} . In other words, we should consider the conditional intensity function $\lambda_\psi(\cdot|H_{t+dt})$, such that:

$$\begin{aligned} \lambda_\psi(t + dt|H_{t+dt}) &= \lambda_\psi(t + dt|H_t)Prob(N(dt) = 0) \\ &+ [\lambda_\psi(t + dt|H_t) + g(dt)]Prob(N(dt) > 0) \\ &\approx \lambda_\psi(t + dt|H_t) + g(dt)\lambda_\psi(t + dt|H_t) \end{aligned}$$

with $g(\cdot)$ defined as in Eq. (3).

Then, to predict the number of events in $t + dt$, we should consider the integral with respect to time:

$$\int_0^1 \lambda_\psi(t + dt|H_t)dt + \int_0^1 g(dt)\lambda_\psi(t + dt|H_t)dt$$

The first part of the sum above describes the predicted number of events in $t + dt$ conditionally to the history up to t , while the second part represents an additional effect due to the events that may occur between t and dt .

In Sect. 3, we derive predictions that include the previous ‘additional’ component, showing that it can indeed further improve prediction accuracy.

With the aim of highlighting the contribution of our approach, the predictions obtained by the ETAS model in the empirical application (Sect. 3) will be compared with those computed through a more traditional count model, that is a Poisson model applied to the event count time series. In particular, we consider the so-called Poisson Autoregression (PAR), whose properties and inference theory were studied by Fokianos et al. (2009). Because of its capability of capturing overdispersion and long memory effects in count time series, an extension of PAR was indeed applied by Agosto et al. (2016) to model the empirically found clustering in the time series of default counts.

The PAR model assumes that a count time series y_t , conditional on its past, follows a Poisson distribution with a linear autoregressive intensity $\lambda_t > 0$:

$$\begin{aligned} y_t|\mathcal{F}_{t-1} \sim Poisson(\lambda_t) &\Leftrightarrow P(y_t = y|\mathcal{F}_{t-1}) = \frac{\lambda_t^y \exp(-\lambda_t)}{y!} \\ \lambda_t &= \omega + \alpha y_{t-1} + \beta \lambda_{t-1} \end{aligned} \tag{7}$$

with $\omega, \alpha, \beta > 0$. Fokianos et al. (2009) showed that, under certain conditions ensuring ergodicity of the process, model (7) can be consistently estimated through a maximum likelihood approach.

3 Application

We consider corporate default swap (CDS) spread data, for a total of 108,316 observations, from 01/01/2008 to 31/12/2015, obtained from the 5 years CDS time series of the Bloomberg database. Each observation describes the daily CDS spread of 54 European financial corporations (bank and insurance companies), belonging to 11 different countries.

We first need a valid definition of what is a “financial shock”, identified by time and space coordinates and by some kind of marker that can correspond to the magnitude of an earthquake in the seismic context.

To this aim, we consider $\Delta_{CDS}(t)$, the relative time difference between two consecutive CDS spreads (the CDS return). We then define a shock event to occur at time t if $\Delta_{CDS}(t)$ exceeds a fixed threshold: the 99th percentile of the observed return distribution of all CDS returns. This identifies a total of 1,084 events, characterised by a time point and a magnitude defined by the corresponding CDS spread.

We then need to define a notion of “distance” between the considered financial companies, that can mimic the spatial distance between earthquake locations. To this aim, we propose to apply a multidimensional scaling (MDS) projection (Torgerson 1958; Cox 2001) for the 54 vectors that contain the 8-years long CDS time series of the considered financial corporations.

Multidimensional scaling methods can approximate multivariate distances among pairs of observations as distances between the corresponding points in a lower p -dimensional space. Starting from a $n \times n$ distance matrix, MDS returns an $n \times p$ configuration matrix, whose rows are the coordinates of the n points in a p -dimensional space for some $p < n$ (Mardia et al. 1979). Here we set $p = 2$. This choice does not only ease the interpretation, but also allows to capture most of the data variability. Table 2 in appendix supplementary material shows indeed that the first two coordinates explain a very large part of the CDS price variance, higher than 95% in all considered years.

In our context, we compute the Euclidean distances between the considered 54 financial corporates, using all their observed CDS daily spreads in the period 2008-2015. Through the MDS projection, we get a set of points such that the distances between them, in a bidimensional space, are approximately equal to their original distances.

Doing so we have a complete description of each observation, in terms of magnitude, time, and space location, which can be used in input to the ETAS model.

Figure 1 contains the temporal distribution of the events as well as their spatial representation, obtained from the multidimensional scaling projection. In the temporal distribution, event times, which are discrete by nature, are randomly moved by a very small amount to avoid visual overlapping.

From Fig. 1 (left) note, as expected, the clustering of many events at the burst of the 2008 financial crisis, with the maximum peak observed on September 15th, 2008. Besides this spike, many other critical events occurred in March 2008, May 2010, June 2012 and June 2014.

From Fig. 1 (right) note that the spatial distribution obtained from the MDS projection individuates three main clusters of the considered 54 financial corporates: the largest one, on the left, contains corporates of the countries identified, during the crisis, as “core” European countries (France, Germany, Netherlands, Nordic countries). Moving along the first coordinate, we observe corporates from Italy, Spain, Portugal and Ireland: the countries identified as “peripheral” countries. Last, on the right of the figure, we can observe a small cluster of financial companies from Greece, the country identified as the “most peripheral” one. Finally, Piraeus Bank,

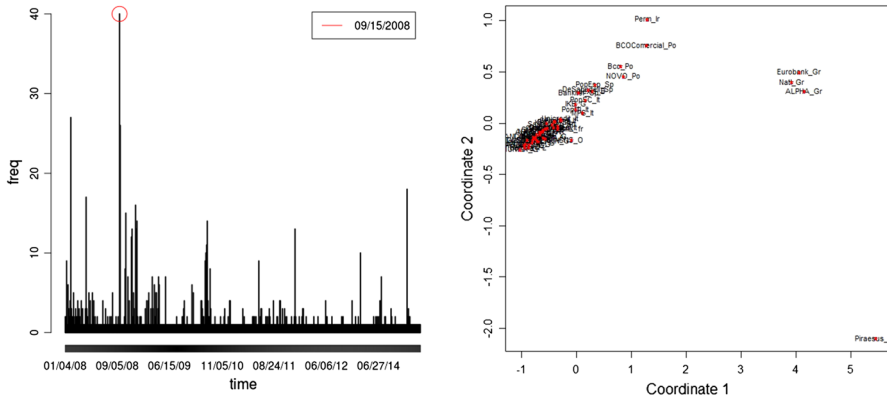


Fig. 1 Temporal (left) and spatial (right) distributions of crisis events occurring to European financial corporates, 2008–2015. The temporal distribution reports, for each considered day on the x-axis, the frequency of the crisis events, that is, the number of CDS spreads whose value exceeded the 99% percentile, calculated on the aggregated time series. The spatial distribution represents the two “financial” coordinates of each of the 52 considered corporates, obtained from the application of a MDS bidimensional projection of all the distances between the available corporates, based on all daily observations

also from Greece, for which CDS trading was even suspended, due to its very negative values, appears isolated in the bottom right part of the figure.

Overall, the spatial representation provided by Fig. 1(right) corresponds to the classification of the European countries into “core” and “peripheral”, particularly during the financial crisis, and is consistent with what employed in the financial literature.

It is worth noticing that this representation could change if considering subsamples of the CDS price time series, rather than the full 8-year sample. To give an insight of the evolution of the spatial representation over the analyzed period, in the appendix supplementary material (Table 3) we report the correlation matrix of the distances calculated in different years. The correlation values show that, after an initial instability in the first two years - corresponding to the spread of the global financial crisis - the distances between institutions become stable.

We remark that other authors have obtained similar conclusions, finding high stability of the distances between countries, except during strong crisis times. See, e.g., Avdjiev et al. (2019), who employed an exogenous measure based on Bank for International Settlements database containing capital flows between countries.

We now apply the Forward Likelihood Prediction (FLP) method to estimate the proposed ETAS model, defined in Eq. (2). The obtained estimates for the seven unknown parameters of the model are reported in Table 1.

From Table 1, note that $\mu = 0.005793$, which indicates that a financial crisis event occurs, on average, every 150 days. Furthermore, the estimate of κ_0 is positive, which indicates that the strength (‘magnitude’) of a previous event positively influence the subsequent offspring intensity. Comparing the values of p and q note that the latter is higher and, therefore, since a lower value of p correspond to a lower decay in time (rather than in space), the results establish

Table 1 Estimates of the ETAS-FLP model in Eq. 2, applied to the available CDS data

	Estimates	Standard errors
μ	0.005793	0.002442
κ_0	0.000260	0.000058
c	0.001411	0.000334
p	0.968939	0.010587
α	0.011795	0.003157
d	0.008519	0.001438
q	1.923259	0.091496

that, for the considered European corporates, time contagion is more persistent than space contagion.

In the following three subsections, we move from a descriptive to a predictive viewpoint, applying the proposed ETAS-FLP model to make predictions, in an out-of-sample perspective. In particular, in Sects. 3.1–3.3 we forecast crisis events using all the CDS spread data available at the time of prediction. Coherently with previous considerations concerning the stability of the spatial representation and to smooth initial instabilities in the joint CDS dynamics, we use constant distances estimated through our MDS procedure.

3.1 Predicting foreshock events

In the time distribution of events represented in Fig. 1, some peaks of intensity can be observed. While some of them, such as the one of September 18th, 2008, appear to be caused by previous crisis events (for which a high CDS spread value is observed) others, such as the peak on September 15th, 2008, do not seem to be induced by the effect of previous events and could rather be caused by external exogenous factors.

Using the terminology used in earthquake modelling, we refer to the former as ‘aftershock’ events (typically induced by a large intensity event) and to the latter as ‘foreshock’ events (typically occurring before a large intensity event). In this subsection we consider the prediction of foreshock events and, in the following, the prediction of aftershock events.

Figure 2 plots the one-ahead predictions obtained with our proposed estimation procedure, for the events occurring in the time interval around September 15th, 2008. Besides the observed event, the figure shows the predictions obtained with and without the additional component in the conditional intensity, which takes into account the effect of events occurring in the interval $t + dt$.

Figure 2 shows that the model does not well capture the first cluster of events, occurring on September 15th, in which as many as 40 crisis events occur: among Dutch, French, German, Italian, Portuguese and Spanish corporates. The observed large number of events is clearly higher than the predicted one.

However, note that after September 18th, in which a large shock hits the Credit Industriel CDS spread, in France (black vertical dotted line), with $\Delta_{CDS} = 1.27$, the




Fig. 2 Comparison between observed and predicted frequency of events for the events of September 2008 by the ETAS-FLP approach. The predicted lines are obtained both with and without the additional component (red and green line respectively)

model does improve considerably its predictions. This is likely due to the fact that the model has acquired more data and, therefore, makes better predictions.

To better evaluate the performance of the proposed model, and to check its validity in a more challenging exercise, we now test its predictive behaviour in a longer, weekly, out-of-sample horizon.

Figure 3 reports the one-week-ahead predictive distribution of the events, computed with the model that includes the additional component.

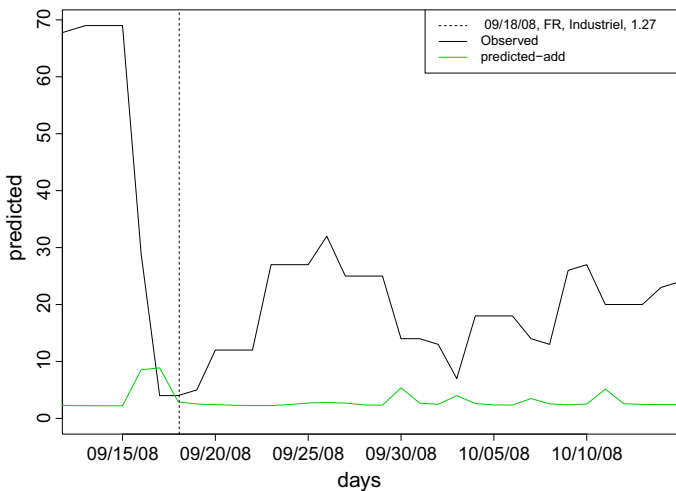


Fig. 3 Comparison between observed and predicted frequency of events by the ETAS-FLP approach with the additional component (green line) for the events of September 2008: one week ahead predictions

Figure 3 confirms the previous findings: the model does not predict well initially but then, as it learns from more data, improves.

To better understand whether our model is valid, we compare its predictions with those obtained with the PAR model defined in Eq. (7). This will allow us to verify whether the ETAS model, that contains a spatial component, improves the predictions with respect to a pure autoregressive model, that contains only a temporal component. Specifically, we carry out an out-of-sample exercise, where the observations up to 6 trading days before the September 15th default peak are used for an initial estimation of a PAR model and 35 one-day-ahead forecasts are computed. The estimation and prediction time intervals are chosen so to mimic the predictive exercise represented in Fig. 2. Figure 4 shows the predictions obtained with the PAR model.

From Fig. 4 note that the PAR model is not able to properly capture the default events cluster. Indeed, a visual comparison with Fig. 2 shows that the autoregressive specification underperforms the ETAS one. It follows that the proposed space-time ETAS model is better than a pure time autoregressive one, and a spatial component is necessary.

Having found that a spatial component is necessary, we focus the analysis on the spatial contagion effects on the predictions. To this aim, we derive the predictive space intensity, under two different conditioning sets. We consider the conditional intensity at September 15th, 2008 and at September 29th, 2008. Figure 5 shows the predictive intensities, for each spatial point, based on the data up to that time point.

In Fig. 5 (top), that refers to the conditional intensity at September 15th, we observe a clear dependence in the spatial intensity, not only between corporates in the same country, as expected, but also between financial corporates belonging to neighbouring countries (such as Spain and Italy or France and Italy).

Figure 5 (bottom) refers to the conditional intensity at September 29th. We remark that September 29th occurs some days after the ‘large shock’ in Credit Industriel and that, on the same day, 15 events are observed: 9 among Belgian, French, German and Dutch corporates, and 6 involving Italian, Spanish and Greek




Fig. 4 Comparison between observed and predicted frequency of events, for the event of September 2008: PAR model

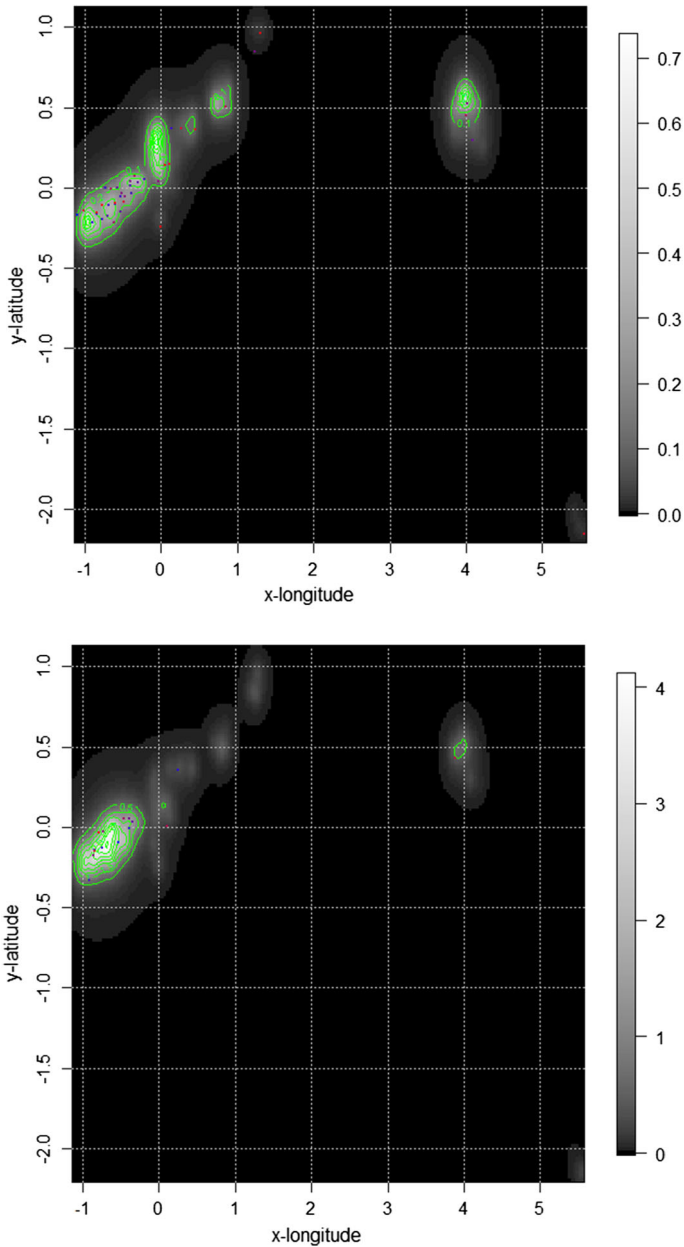


Fig. 5 Predictive spatial intensity by the ETAS-FLP approach, at September 15th, 2008 (top) and at September 29th (bottom), measured in terms of the two MDS coordinates (latitude and longitude)

companies. Figure 5 (bottom) clearly shows that the large shock in the CDS spread of Credit Industriel on September 18th, 2008 has increased the intensity in France and in the neighbouring countries, leading to a substantial shift in the main kernel of

the intensity. This is consistent with our previous finding: the model improves by learning from shock data.

3.2 Predicting aftershock events

We now consider the prediction of ‘aftershock’ events which, presumably, are caused endogenously, differently from ‘foreshock’ events that are likely caused by exogenous factors. In particular, we consider the shock occurring on March 7th, 2008, when a large CDS value in the French Groupama was observed ($\Delta_{CDS} = 6.15$). On that day several CDS variations exceeding the threshold were observed: 3 in France, 6 in Italy, 3 in Spain, 1 in Greece, 4 in Germany.

Figure 6 reports the one-day-ahead predictive distribution of the events.

From Fig. 6 note that, six days after the shock (on March, 13th) five further events occur, and they are well predicted by our model: 1 in Belgium, 1 in France, 2 in the Netherlands and 1 in Portugal. Overall, our model seem to well predict the number of crisis events. Comparing Fig. 6 with Fig. 2, it appears that our model is better suited to predict aftershock events, rather than foreshock events.

To further evaluate the performance of our model, we also consider one-week-ahead predictions. Figure 7 reports the one week ahead predictive distribution of the events.

From Fig. 7 note that the one-week-ahead predictions are worse, as expected, than the one-day ahead predictions, but better than those relative to foreshock events.

We then compare our one-day-ahead predictions with those obtained by the PAR model. The latter is firstly estimated up to 3 days before the Groupama event,




Fig. 6 Comparison between observed and predicted frequency of events for the event of March 2008 by the ETAS-FLP approach. The predicted lines are obtained both with and without the additional component (red and green line respectively)

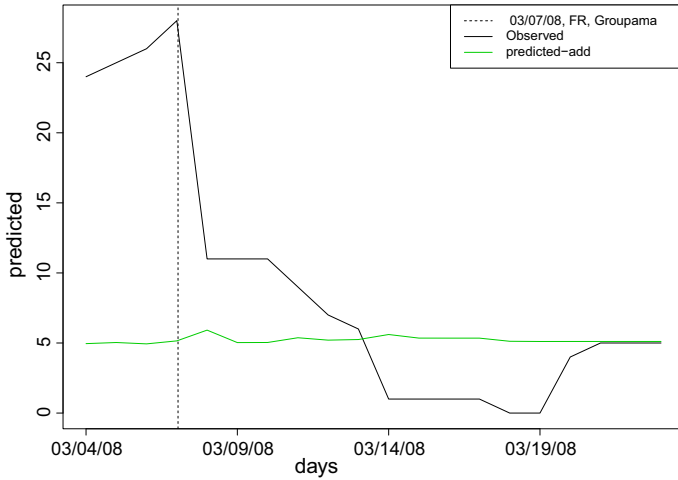


Fig. 7 Comparison between observed and predicted frequency of events by the ETAS-FLP approach with the additional component (green line) for the events in March 2008: one week ahead predictions




Fig. 8 Comparison between observed and predicted frequency of events for the event in March 2008: PAR model

providing a one-day-ahead forecast for the following 20 days. Figure 8 shows the obtained predictions.

Figure 8, compared with Fig. 6, shows that the PAR model does not well capture the default peaks, where it underperforms the ETAS model. Again, the results show that a spatial component is necessary.

As in the previous subsection, we now focus on the spatial contagion effects. We consider data until March 7th and 13th, 2008, respectively. The corresponding conditional intensities are represented, in Figs. 9 (top) and 9 (bottom), respectively.

Comparing the two plots in Fig. 9, note that the effect of the 2008 Groupama event, generating the stress area in red in Fig. 8, persists a few days after in the neighbouring region, as shown in Fig. 9 (bottom).

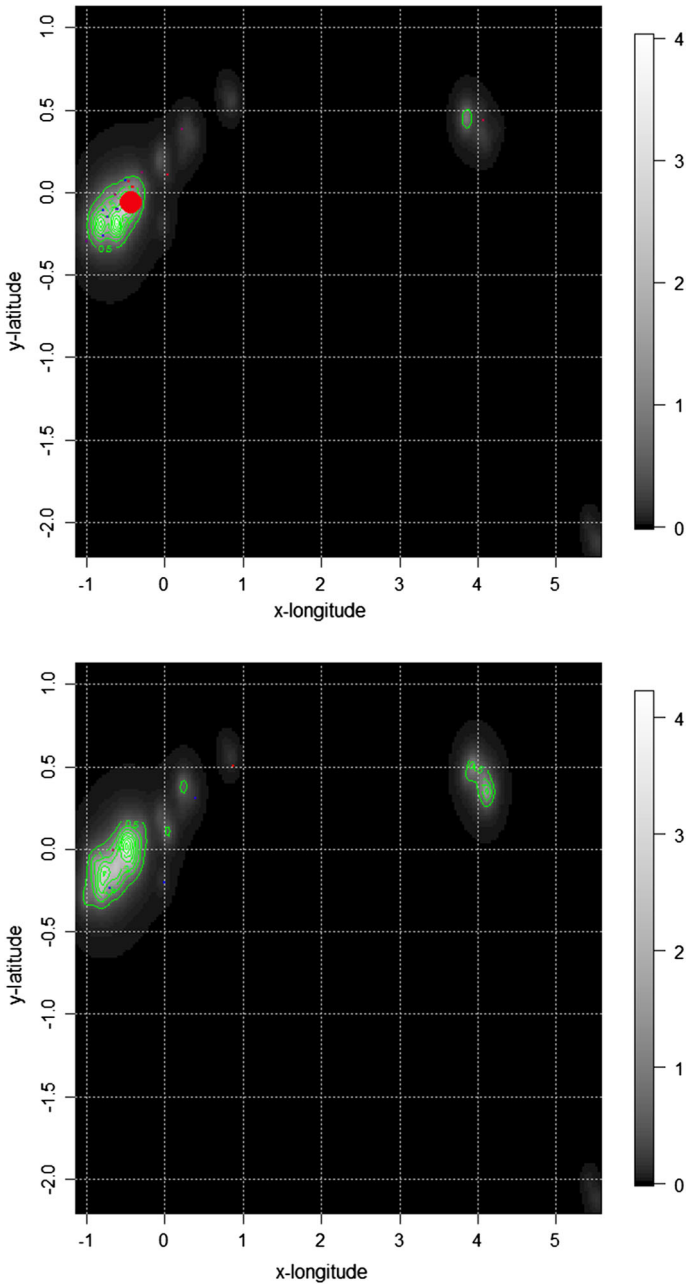


Fig. 9 Predictive spatial intensity by the ETAS-FLP approach, at March 7th (top) and March 13th (bottom), 2008, measured in terms of the two MDS coordinates (latitude and longitude)

3.3 Predicting mixed events

We finally consider the prediction of “mixed” events, that exhibit both a ‘foreshock’ and an ‘aftershock’ component. An example of this occurrence are the events that cluster around the event occurring to Groupama on June 7th, 2012. For this event, we have the ‘typical’ representation of a seismic sequence: foreshocks, mainshock and aftershocks. The mainshock is the big event on June 7th, 2012, with $\Delta_{CDS}(t) = 5.20$, preceded by foreshocks, probably induced by external factors not directly related to the strength of the main event, and followed by an aftershocks sequence, that is reflected in an increase of the offspring intensity. The one-day-ahead predictive distribution around June 7th, 2012 is reported in Fig. 10.

From Fig. 10 note that, also in this case, the proposed approach well performs in terms of prediction, especially after the occurrence of the ‘large shock’ in the Groupama CDS spread, which provides a noticeable gain of information.

To better evaluate the performance of our model, we consider, as before, one-week-ahead predictions. Figure 11 reports the one week ahead predictive distribution of the events.

From Fig. 11 note that the one-week-ahead predictions are now quite well performing, compared to the ones discussed in the previous subsections. This is consistent with the fact that the model has been fed with sufficient shock data.

We can also compare our one-day ahead predictions with those computed with a PAR model. Specifically, we can use observations up to 17 trading days before the June 7th event for an initial estimation of the PAR model and then, provide 50 one-step-ahead forecasts. Figure 12 shows the obtained predictions.




Fig. 10 Comparison between observed and predicted frequency of events for the event of June 2012 by the ETAS-FLP approach. The predicted lines are obtained both with and without the additional component (red and green line respectively)

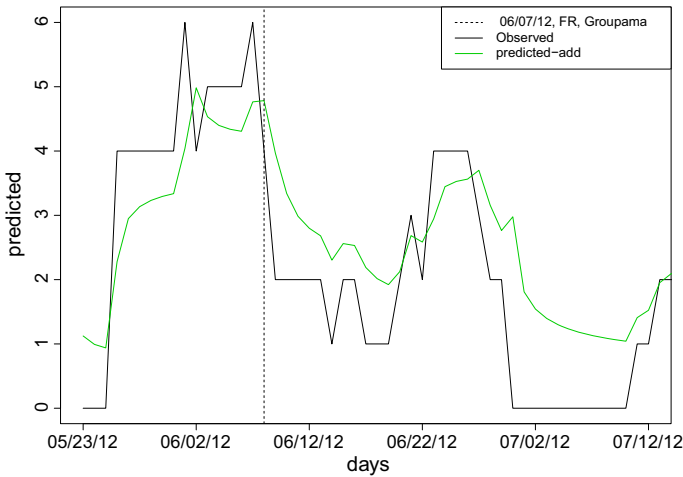


Fig. 11 Comparison between observed and predicted frequency of events by the ETAS-FLP approach with the additional component (green line) for the days in June 2012: one week ahead predictions





Fig. 12 Comparison between observed and predicted frequency of events for the days in June 2012: PAR model

By comparing Fig. 12 with Fig. 10, note that the ETAS specification again outperforms the PAR model in prediction, giving further evidence in favor of our space-time model against a pure autoregressive model.

From a spatial contagion point of view, we can examine the spatial predicted intensity, by looking at three conditional intensities: up to May 31st (one week before the event), in Fig. 13 (top); up to June 7th (the day of the big event), in Fig. 13 (middle); and up to June 28th, 2012 (twenty one days after the event), in Fig. 13 (bottom).

Comparing the three plots in Fig. 13, note that the stress activity, which emerges before the large event of June 7th, and “erupts” on that day, is widely spread among the four main clusters previously identified. However, after the event of June 7th, the spatial intensity changes, shifting mostly to the France area (Fig. 13,middle),

Fig. 13 Predicted spatial intensity by the ETAS-FLP approach at May 31st, 2012 (top), June 7th, 2012 (middle), and June 28th (bottom), measured in terms of the two MDS coordinates (latitude and longitude)



with the effect persisting even twenty one days afterwards, in the same area (Fig. 13, bottom).

4 Concluding remarks

In the paper, we have extended ETAS space-time point process, typically employed in the study of seismic contagion, to model financial contagion. To do so we have defined the occurrence of a financial shock by an extreme variation of the CDS spread of a certain institution, and we have introduced a projection method aimed at deriving “financial coordinates” of a set of institutions, that reproduces in a bidimensional space the multidimensional distances between the corresponding time series. We have also extended the ETAS-FLP estimation procedure to derive predictions which can be employed to gauge model accuracy.

Our empirical findings show that the proposed model takes advantage from the spatial modelling component, as it leads to predictions more accurate than those obtained with a pure autoregressive model. In addition, the model performs better when the events to be predicted are aftershocks, or mixed, as more shock data are available for the learning process.

We believe that our proposed model could be useful by policy makers, regulators and supervisors, as an indicator of the emergence of a financial crisis in one or more institutions, caused by contagion from other institutions. Such an indicator can be used as an “early warning” signal aimed at preventing and, therefore, mitigating in advance, the systemic consequences of financial crisis.

Future research developments include evaluating the robustness of the results, for different spatial regions and other types of institutions, in particular sovereign entities. This would require to extend the approach to include specific covariates, such as the macro economic characteristics of the considered countries.

Other developments may include the application of the proposed methodology to other types of financial risks, such as operational risks and cyber risks.

Acknowledgements Open access funding provided by Università degli Studi di Pavia within the CRUI-CARE Agreement. The Authors thank the anonymous referees, for their suggestions, which have helped improving the quality of the paper.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Funding The paper has been partially supported by the national grant of the Italian Ministry of Education University and Research (MIUR) for the PRIN-2015 program, “Complex space-time modelling and

functional analysis for probabilistic forecast of seismic events'. It has also received support from the European Union's Horizon 2020 research and innovation program "FIN-TECH: A Financial supervision and Technology compliance training programme" under the Grant Agreement No. 825215 (Topic: ICT-35-2018, Type of action: CSA), and from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 750961.

Table 2 Percentage of CDS price variance explained by the first two spatial coordinates for each sample year

Year	Explained variance
2008	0.9294
2009	1.0000
2010	0.9762
2011	0.9573
2012	0.9618
2013	0.9662
2014	0.9867
2015	0.9992

Table 3 Correlation matrix of the Euclidean distances between institutions calculated in different years based on the CDS time series

Year	2008	2009	2010	2011	2012	2013	2014	2015
2008	1.0000	0.6889	0.1644	0.1574	0.1544	0.1520	0.1449	0.1311
2009	0.6889	1.0000	0.5577	0.5755	0.5687	0.5685	0.5619	0.5289
2010	0.1644	0.5578	1.0000	0.9461	0.8818	0.8400	0.8258	0.7646
2011	0.1574	0.5755	0.9460	1.0000	0.9832	0.9650	0.9582	0.9194
2012	0.1544	0.5687	0.8818	0.9832	1.0000	0.9918	0.9862	0.9584
2013	0.1520	0.5685	0.8400	0.9650	0.9918	1.0000	0.9985	0.9828
2014	0.1449	0.5619	0.8258	0.9582	0.9862	0.9985	1.0000	0.9854
2015	0.1311	0.5289	0.7646	0.9194	0.9584	0.9828	0.9854	1.0000

Appendix: Supplementary material: spatial representation

This section presents the results of a robustness analysis of the spatial representation used in the paper.

Table 2 reports the percentage of CDS price variance explained by the first two spatial coordinates for each sample year, while Table 3 shows the correlation matrix of the distances between institutions calculated in different years.

References

- Acharya V, Engle R, Richardson M (2012) Capital shortfall: a new approach to ranking and regulating systemic risks. *Am Econ Rev Pap Proc* 102(3):59–64
- Adelfio G, Chiodi M (2015) Alternated estimation in semi-parametric space-time branching-type point processes with application to seismic catalogs. *Stoch Environ Res Risk Assess* 14:151–186
- Adrian T, Brunnermeier M (2011) Covar. NBER working paper 17454, National Bureau of Economic Research
- Agosto A, Cavaliere G, Kristensen D, Rahbek A (2016) Modeling corporate defaults: poisson autoregressions with exogenous covariates (PARX). *J Empir Finance* 38(B):640–663
- Ahelegbey D, Billio M, Casarin R (2015) Bayesian graphical models for structural vector autoregressive processes. *J Appl Econom* 31(2):357–386
- Allen F, Carletti E (2006) Credit risk transfer and contagion. *J Monet Econ* 53(1):89–111
- Ang A, Longstaff F (2012) Systemic sovereign credit risk: lessons from the US and Europe. Technical report, National Bureau of Economic Research
- Avdjiev S, Giudici P, Spelta A (2019) Measuring contagion risk in international banking. *J Financ Stab* 42:36–51
- Battiston S, Delli Gatti D, Gallegati M, Greenwald B, Stiglitz J (2012) Liasons dangereuses: Increasing connectivity risk sharing, and systemic risk. *J Econ Dyn Control* 36(8):1121–1141
- Bauwens L, Hautsch N (2009) Modelling financial high frequency data using point processes. *Handbook of financial time series*, pp 953–979
- Betz F, Oprica S, Peltonen T, Sarlin P (2014) Predicting distress in european banks. *J Bank Finance* 45(C):225–241
- Billio M, Getmansky M, Lo A, Pelizzon L (2012) Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *J Financ Econ* 104:535–559
- Brownlees C, Engle R (2012) Volatility, correlation and tails for systemic risk measurement. Technical report, New York University
- Chiodi M, Adelfio G (2011) Forward likelihood-based predictive approach for space-time processes. *Environmetrics* 22:749–757
- Chiodi M, Adelfio G (2017) etasflp: Estimation of an etas model. mixed flp (forward likelihood predictive) and ml estimation of non-parametric and parametric components of the etas model for earthquake description. R package version 1.4.1
- Cox T (2001) *Multidimensional scaling*. Chapman and Hall, London
- Daley DJ, Vere-Jones D (2003) *An introduction to the theory of point processes*, 2nd edn. Springer, New York
- Diebold F, Yilmaz K (2014) On the network topology of variance decompositions: measuring the connectedness of financial firms. *J Econom* 182:119–134
- Diggle P (1983) *Statistical analysis of spatial point patterns*. Academic Press, London
- Fokianos K, Rahbek A, Tjøstheim D (2009) Poisson autoregression. *J Am Stat Assoc* 104:1430–1439
- Giudici P, Spelta A (2016) Graphical network models for international financial flows. *J Bus Econ Stat* 34(1):128–138
- Gower JC (1966) Some distance properties of latent root and vector methods used in multivariate analysis. *Biometrika* 53:325–328
- Hawkes A, Adamopoulos L (1973) Cluster models for earthquakes-regional comparison. *Bull Int Stat Inst* 45(3):454–461
- Jorion P, Zhang G (2009) Credit contagion from counterparty risk. *J Finance* 64:2053–2087
- Lorenz J, Battiston S, Schweitzer F (2009) Systemic risk in a unifying framework for cascading processes on networks. *Eur Phys J B Condens Matter Complex Syst* 71(4):441–460
- Mardia KV, Kent JT, Bibby JM (1979) *Multivariate analysis*. Academic Press, London
- Ogata Y (1988) Statistical models for earthquake occurrences and residual analysis for point processes. *J Am Stat Assoc* 83(401):9–27
- Ogata Y (1998) Space-time point-process models for earthquake occurrences. *Ann Inst Stat Math* 50(2):379–402
- Ripley B (1977) Modelling spatial patterns. *J R Stat Soc* 39(2):172–212
- Schabenberger O, Gotway C (2005) *Statistical methods for spatial data analysis*. Chapman and Hall, London
- Silverman BW (1986) *Density estimation for statistics and data analysis*. Chapman and Hall, London

Torgerson WS (1958) Theory and methods of scaling. Wiley, New York

Utsu T (1961) A statistical study on the occurrence of aftershocks. *Geophys Mag* 30:521–605

Veen A, Schoenberg FP (2008) Estimation of space-time branching process models in seismology using an em-type algorithm. *JASA* 103(482):614–624

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.