Financial Contagion as a Self-Exciting Point Process[⋆]

Luca Persia[a,b,*]

[a]*Institute of Wealth and Asset Management, Zurich University of Applied Sciences, Zurich, Switzerland*
[b]*Institute of Computing, Università della Svizzera italiana, Lugano, Switzerland*

Abstract

Financial crises rarely manifest as isolated shocks; rather, they cluster over time and propagate across markets. Building on this idea, this paper develops an event-based framework for systemic risk and financial contagion in which extreme market movements are treated as arrivals in a self-exciting point process. We construct daily stress events from left tail return exceedances across U.S. sector ETFs and study whether their arrival mechanism is compatible with conditionally independent shocks or exhibits endogenous amplification and contagion. We estimate univariate and multivariate Hawkes processes over two crisis regimes with different dominant mechanisms, the Global Financial Crisis (GFC) and the COVID-19 episode. The results show that the GFC displays stronger self-excitation, a higher spectral radius of the contagion matrix, and more persistent sectoral amplification, especially within the financial sector. By contrast, the COVID-19 episode is characterized by a sharper but more transient intensity spike and a lower degree of endogenous persistence. To assess whether the estimated processes are economically meaningful, we compare them with the VIX as an external benchmark for market-wide uncertainty. In both windows, the multivariate Hawkes model tracks the VIX closely and performs competitively relative to a DCC-GARCH benchmark. Overall, the results support a macro-financial interpretation of crises as self-exciting systems and illustrate how Hawkes models can provide operational monitoring tools for financial stability, including measures of endogeneity, contagion, and persistence.

Keywords: Hawkes process, Financial contagion, Financial stability, System dynamics

⋆Preprint version, June 2026.

*Corresponding author
Email address: `pess@zhaw.ch; luca.persia@usi.ch` (Luca Persia)

1. Introduction

Financial crises can trigger periods of systemic stress in which difficulties in one part of the system can lead to broad market imbalances and financing distress. Typically, two narratives are used to explain why crises occur. The first emphasizes rare and largely exogenous shocks, often discussed as "black swans" in the sense of Taleb (2007). In this view, turmoil is triggered by external unexpected events, such as pandemics, sudden geopolitical escalations, or other major disruptions. Financial regulators often frame systemic resilience in these terms, focusing on how financial infrastructures and intermediaries cope with low probability and high impact shocks. The second narrative emphasizes that fragility can accumulate endogenously within the system, as analyzed extensively from Sornette (2003) and Bouchaud (2024). This narrative focuses on the fact that liquidity spirals and fire sales can create feedback loops in which small shocks are amplified and propagate across markets and institutions (Brunnermeier and Pedersen, 2009; Adrian and Shin, 2010; Allen and Gale, 2000). This latter view is close to the broader concept of self-organized criticality in complex systems, where interactions among components drive the aggregate state toward a marginally stable regime, allowing small perturbations to generate large cascades (Bak, 1996). From a systemic risk perspective, the distinction is important because it changes what should be monitored. If crises are mostly exogenous, the relevant object is exposure to external shocks; if crises are largely endogenous, the relevant object is the internal propagation and persistence.

Most empirical research on systemic risk and financial contagion relies on statistical models for time series of returns and related observables, with the main focus on volatility dynamics and co-movements across assets and markets. Standard approaches include models such as GARCH and stochastic volatility, dynamic conditional correlation (DCC) specifications, dynamic copulas, and connectedness measures based on variance decompositions or balance-sheet linkages (Engle, 2002; Bollerslev, 1986; Diebold and Yilmaz, 2012; Forbes and Rigobon, 2002). These tools are well suited to capturing persistence and spillovers in second moments, and they fit naturally with a monitoring practice that treats stress as a continuous latent state evolving over time. In parallel, policy institutions often summarize systemic conditions through composite stress indices. A prominent example are the Composite Indicator of Systemic Stress (CISS) developed at the European Central Bank (Holló et al., 2012) and the St. Louis Fed Financial Stress Index (STLFSI), widely used as measures of financial condition within the Federal Reserve System (FED). These indicators are valuable for contemporaneous monitoring, but they are not designed to specify a data-generating process for crisis dynamics. In particular, they do not model the discrete cascade of large adverse events that tends to characterize crises, and they do not directly map today's shock into a short-term prediction about the probability of further shocks. We see this as a practical limitation, as many macroprudential issues today concern the propagation of risk rather than the level of a stress index. We may want to consider different questions: *How long is high stress likely to persist after a shock? How much does one episode of stress increase the short-term probability of further episodes? To what extent are crises driven by external event versus self-reinforcing feedback?* These are questions about the timing and clustering of extreme events, and are difficult to answer using continuous value stress indicators or volatility models, unless we introduce additional layers of assumptions about how extremes arise from the underlying time series.

A different perspective that we propose in this paper is an event-driven view of systemic risk and contagion by modeling the occurrence of extreme losses as a point process. Rather than working directly with the full return series, we detect large adverse moves and treat them as stress events. The key empirical question is whether the sequence of such events is compatible with a memoryless mechanism, or whether it exhibits self-excitation. In different terms, if crises were largely a sequence of independent events, a Poisson process would be a reasonable starting point; if instead crises contain a relevant endogenous and self-exciting component, we should observe temporal clustering that cannot be explained by deterministic time variation alone. Self-exciting point process, or Hawkes processes, introduced by Hawkes (1971), provide a probabilistic approach to model this mechanism. They formalize a trivial intuition: after a shock, stress does not reset instantly to normal, but remains elevated and decays over time. In that sense, Hawkes processes can offer a statistical representation of propagation and persistence in financial markets.

The same modeling logic is well established in other fields where aftershocks are common. In seismology, seismic sequences are commonly modeled as self-exciting processes in which each event temporarily increases the probability of subsequent events; epidemic-type aftershock sequence (ETAS) models are a standard example (e.g., see Ogata (1988) and Zhuang (2012)). Similar ideas related to point processes have been used to describe contagion and clustering in epidemiology, crime, and information diffusion (Mohler et al., 2011; Mohler, 2014; Zhao et al., 2015; Browning et al., 2021). In quantitative finance, point processes are also widely used. Bacry et al. (2015) provides a comprehensive review of the most general applications of Hawkes processes in finance, ranging from market microstructure and high-frequency trading (Large, 2007) to risk management and insurance (Dassios and Zhao, 2011), to asset valuation (e.g., see Hainaut (2016) and Liu and Zhu (2019)) and price dynamics beyond geometric Brownian motion (Hawkes, 2018). Several papers have also used related ideas to model contagion in financial jumps and clustered extremes. Aït-Sahalia et al. (2015) explored mutually exciting jump intensities in a jump diffusion framework to capture how price jumps in one market can trigger jumps elsewhere. Errais et al. (2010) have employed self- and cross-exciting intensities to capture clustering and spillovers in default events.

The contribution of our paper lies in how Hawkes processes are used for systemic risk analysis. We exploit two of their properties that are particularly informative. The first is a direct measure of endogeneity via the branching ratio, which quantifies the expected number of events triggered by previous events rather than by independent shocks. This is a clear way to quantify how self-sustaining a crisis episode is. The second is the multivariate extension, in which each market segment has its own intensity and events in one segment can increase the intensity in another. The excitation coefficients form a contagion matrix that can be interpreted as a directed network of stress propagation, providing a reduced-form representation of contagion channels that complements correlation measures. Lastly, we provide an empirical contribution by testing this framework to compare two crisis regimes with different dominant mechanisms. The first is the Global Financial Crisis (GFC), which can be interpreted as a crisis mainly characterized by endogenous financial amplification. The second is the COVID-19 episode, which can be interpreted as a largely exogenous shock that nevertheless generated temporary endogenous propagation in financial markets. This comparison connects directly to the conceptual distinction introduced above: crises may originate outside the financial system, but the persistence and severity of their market impact depend on the internal feedback structure of the system.

We implement the model on U.S. sector ETFs. For each crisis window, we test goodness-of-fit using time-rescaling diagnostics and Kolmogorov–Smirnov tests on the transformed inter-event gaps. We also compare the model estimated intensity to the VIX index, which is widely used as a market measure of expected volatility and risk sentiment.

This diagnostic assesses whether event-driven intensities capture the peaks and persistence of systemic stress. As a time series benchmark, we use a DCC-GARCH, since it is a standard model for volatility and correlation and represents a widely used alternative view of market stress (Engle, 2002).

The remainder of the paper is structured as follows. Section 2 develops the theoretical framework, introducing point process models for financial stress events, the marked multivariate Hawkes specification, and the corresponding estimation and stability properties. Section 3 details the construction of the empirical dataset and applies the model to the GFC and the COVID-19 episode, comparing clustering, contagion, and systemic persistence through the branching matrix and intensity dynamics, and benchmarking the resulting stress signals against a DCC-GARCH specification. Finally, Section 4 discusses policy relevance and limitations, emphasizing the role of Hawkes modeling for financial stability and outlining directions for future research.

2. Methodology

This section sets up the point process framework used throughout the paper. The core idea is to treat extreme market moves as event arrivals in continuous time, considering one trading day as the time unit.

2.1. Point processes

We model the arrival of extreme events in a market as a point process on $[0, \infty)$. Let $N(t)$ denote the cumulative number of events observed up to time t. We call $\{N(t) : t \geq 0\}$ a *counting process* if:

1. $N(t) \in \mathbb{N}_0 = \{0, 1, 2, \dots\}$ for all $t \geq 0$ and $N(0) = 0$;
2. $N(t)$ has right-continuous sample paths and jumps of size at most $+1$ [1].
3. for every finite horizon T, $N(T)$ is almost surely finite.

The (random) jump times $0 < t_1 < t_2 < \dots$ at which $N(t)$ increases by one are the event arrival times; the point process can be represented equivalently by $\{t_n\}_{n \geq 1}$ or by the counting path $\{N(t)\}_{t \geq 0}$.

We write $\mathcal{H}_{t^-}$ for the history (filtration) generated by the process up to just before time t, i.e., the information contained in $\{t_n : t_n < t\}$ (equivalently, in $\{N(s) : 0 \leq s < t\}$). Conditioning on $\mathcal{H}_{t^-}$ distinguishes models where arrivals are independent of the past from models where past events change the short-run arrival rate. A basic property of counting processes concerns increments over disjoint intervals. For times $0 \leq t_1 < t_2 < \cdots < t_n$, consider the increments

$$N(t_2) - N(t_1),\ N(t_3) - N(t_2),\ \dots,\ N(t_n) - N(t_{n-1}).$$

We say that N has independent increments if these increments are independent for every such choice of times. We say that N has stationary increments if the distribution of $N(t+h) - N(t)$ depends only on $h > 0$ and not on t. In point processes, the central object to model is the *conditional intensity*, defined as

$$\lambda(t \mid \mathcal{H}_{t^-}) := \lim_{\Delta t \to 0^+} \frac{\mathbb{E}[N(t + \Delta t) - N(t) \mid \mathcal{H}_{t^-}]}{\Delta t},$$

[1] This is the standard simple-process convention. In the empirical multivariate application, multiple sector events can occur on the same trading day; these ties are handled by the implementation convention described below.

whenever the limit exists. Intuitively, $\lambda(t \mid \mathcal{H}_{t^-})$ is the instantaneous expected arrival rate at time t , given the information just before t.

2.2. Homogeneous Poisson process

A first benchmark for event arrivals is the homogeneous Poisson process.

Definition 1 (Homogeneous Poisson process). *A counting process $\{N(t) : t \geq 0\}$ is a homogeneous Poisson process with rate $\lambda > 0$ if:*

1. *for every interval $I = (s,t] \subset [0,\infty)$ with length $|I| = t - s$,*
$$N(I) := N(t) - N(s) \sim \text{Poisson}(\lambda|I|);$$
2. *for any finite collection of disjoint intervals $I_1, \dots, I_n$, the random variables $N(I_1), \dots, N(I_n)$ are independent.*

For each fixed $t \geq 0$ this implies

$$\mathbb{P}\big(N(t) = n\big) = \frac{(\lambda t)^n}{n!}\, \mathrm{e}^{-\lambda t}, \qquad n = 0, 1, 2, \dots.$$

Both $\mathbb{E}[N(t)]$ and $\mathrm{Var}(N(t))$ equal λt, meaning that the variance-to-mean ratio is exactly one. Deviations from one indicate under-dispersion (< 1) or over-dispersion (> 1), which is often a first sign that independent arrivals are too restrictive for financial stress data. Let t_n denote the time of the n-th event,

$$t_n := \inf\{t \geq 0 : N(t) \geq n\}, \qquad n = 1, 2, \dots,$$

and define the inter-arrival times $Z_n := t_n - t_{n-1}$ with $t_0 := 0$. For a homogeneous Poisson process,

$$Z_n \sim \text{Exp}(\lambda), \qquad \text{i.i.d.}$$

so the waiting time to the next event is memoryless. In different terms, conditional on the present, the distribution of the next waiting time does not depend on how long has passed since the last arrival.

2.3. Non-homogeneous Poisson process

The homogeneous Poisson model assumes a constant arrival rate, which is rarely realistic in financial markets, where calm and turbulent periods alternate. A natural extension is the non-homogeneous, or inhomogeneous, Poisson process, in which the intensity varies deterministically over time.

Definition 2 (Non-homogeneous Poisson process). *A counting process $\{N(t) : t \geq 0\}$ is an inhomogeneous Poisson process with (deterministic) intensity function $\lambda(t) \geq 0$ if:*

1. *for any interval $[s,t] \subset [0,\infty)$,*
$$N(t) - N(s) \sim \text{Poisson}\bigg(\int_s^t \lambda(u)\, du\bigg);$$
2. *increments over disjoint intervals are independent.*

The inhomogeneous Poisson process still enforces independent increments, but allows the background arrival rate to change over time through $\lambda(t)$. In our context, it serves as a baseline for deterministic variation in event rates and is specified with a log-linear deterministic intensity,

$$\lambda(t) = \mathrm{e}^{\delta_0 + \delta_1 t}, \tag{1}$$

where t is measured in trading days from the beginning of the observation window.

2.4. Self-exciting point process

When modeling financial events over time, extreme movements rarely occur in isolation; Poisson processes assumptions are therefore too restrictive, as it is also natural to expect that the intensity of a point process will remain fairly high for a certain time range after the main event occurs, and eventually decay following a certain behavior over time. This endogenous clustering in the macro-financial context is precisely what Poisson models exclude. A Hawkes process captures this by letting past events raise the future intensity. In the univariate case the conditional intensity is specified as

$$\lambda(t \mid \mathcal{H}_{t^-}) = \mu + \sum_{t_n < t} \phi(t - t_n), \tag{2}$$

where $\mu \geq 0$ is the baseline, or exogenous, intensity and $\phi(u) \geq 0$ is the excitation kernel. $u = t - t_n > 0$ denotes the elapsed time since a past event at time t_n, so $\phi(u)$ describes how the impact of that event decays as the lag u increases. Setting $\phi \equiv 0$ collapses (2) to a Poisson process. A common and analytically convenient choice is to choose an exponential kernel, expressed in the branching parameterization:

$$\phi(u) = \eta\, \beta\, \mathrm{e}^{-\beta u} \mathbf{1}_{\{u>0\}}, \tag{3}$$

with $\eta > 0$ and $\beta > 0$. Under (3), each event causes an instantaneous jump of size $\eta\beta$ in the intensity, and this excess decays exponentially at rate β. The ratio β^{-1} controls the memory time-scale, while the integral of the kernel directly yields the *branching ratio*:

$$\int_0^\infty \phi(u)\, du = \int_0^\infty \eta\, \beta\, \mathrm{e}^{-\beta u}\, du = \eta.$$

In the cluster representation of Hawkes processes, η equals the expected number of direct offspring events triggered by a typical event (Hawkes and Oakes, 1974). The subcritical regime is $\eta < 1$, under which the process admits a stationary version with finite mean intensity and event clusters die out on average (Hawkes and Oakes, 1974; Daley and Vere-Jones, 2003). At $\eta = 1$ the system is critical, while for $\eta > 1$ it is supercritical and cascades can grow explosively with positive probability. η can indeed be interpreted as a summary measure of market endogeneity: as suggested by Filimonov and Sornette (2012), values near one indicate that observed stress is dominated by internal feedback, and small exogenous triggers can generate long term clusters and reflexivity mechanism in financial markets.

2.5. Multivariate Hawkes and financial contagion

In an interconnected macro-financial system, what matters for systemic risk is not only how often each market experiences extreme events, but also how shocks propagate across markets; a multivariate Hawkes process provides a compact reduced-form description of such propagation. In a multivariate specification, each market has its own conditional intensity, and past events in any market can increase the near-term intensity of others through an excitation matrix. In this sense, the excitation matrix can be read as a directed weighted network encoding spillover channels, a viewpoint widely used in financial applications of Hawkes processes (Bacry et al., 2015).

2.5.1. Intensity

Let each observed event be represented by a triplet (t_k, d_k, r_k) for $k = 1, \dots, N$, where $t_k \in (0, T]$ is time, $d_k \in \{1, \dots, d\}$ is the market in which the event occurs, and $r_k \geq 0$ is a raw mark measuring severity. In our framework, raw marks are log-transformed and normalized by the overall sample mean of the transformed marks:

$$m_k = \frac{\log(1 + r_k)}{N^{-1} \sum_{\ell=1}^{N} \log(1 + r_\ell)}. \tag{4}$$

For likelihood and estimation it is convenient to work with the equivalent component-wise representation: for each type $i \in \{1, \dots, d\}$ let $\{t_{i,n}\}_{n=1}^{N_i}$ be the event times of type i in $(0, T]$ and let $\{m_{i,n}\}_{n=1}^{N_i}$ be the corresponding transformed marks. This is simply a reindexing of the same marked event stream, where each pooled event (t_k, d_k, m_k) corresponds to one pair $(t_{d_k,n}, m_{d_k,n})$.

Let $\lambda_i(t \mid \mathcal{H}_{t^-})$ denote the conditional intensity of an event of type i at time t, given the history $\mathcal{H}_{t^-}$. We specify a marked multivariate Hawkes model with exponential kernels as

$$\lambda_i(t \mid \mathcal{H}_{t^-}) = \mu_i + \sum_{j=1}^{d} \sum_{n:\, t_{j,n} < t} m_{j,n}\, B_{j,i}\, \beta_i\, \mathrm{e}^{-\beta_i (t - t_{j,n})}, \qquad i = 1, \dots, d, \tag{5}$$

where $\mu_i > 0$ is the baseline, or exogenous, event rate of market i, and the double sum is the endogenous part such that each past event of type j adds an aftershock component to market i. The coefficient $B_{j,i} \geq 0$ is the branching coefficient from market j to market i, while $\beta_i > 0$ controls how quickly this effect decays in market i. The mark $m_{j,n}$ scales the jump size, so that more severe events create proportionally stronger excitation.

2.5.2. Branching contagion matrix and stability

The coefficients $B_{j,i}$ in (5) are written directly in branching matrix form. To see this, consider a unit mark event of type j occurring at time $t_{j,n}$. At a later time $t > t_{j,n}$, its effect on the type-i intensity depends on the lag $u = t - t_{j,n} > 0$. The corresponding type-j to type-i triggering kernel is

$$\phi_{j \to i}(u) = B_{j,i}\, \beta_i\, \mathrm{e}^{-\beta_i u} \mathbf{1}_{\{u > 0\}}.$$

In the branching, or cluster, interpretation of Hawkes processes, $\phi_{j \to i}(u)\, du$ is the expected number of direct type-i offspring events generated by one unit mark type-j event at lags between u and $u + du$. Integrating over all future lags gives

$$\int_0^\infty \phi_{j \to i}(u)\, du = \int_0^\infty B_{j,i}\, \beta_i\, \mathrm{e}^{-\beta_i u}\, du = B_{j,i}.$$

Thus $B_{j,i}$ is the expected number of direct type-i events generated by one unit mark type-j event. Collecting these coefficients defines the branching, or contagion, matrix

$$B = \begin{pmatrix} B_{1,1} & B_{1,2} & \cdots & B_{1,d} \\ B_{2,1} & B_{2,2} & \cdots & B_{2,d} \\ \vdots & \vdots & \ddots & \vdots \\ B_{d,1} & B_{d,2} & \cdots & B_{d,d} \end{pmatrix}.$$

Diagonal entries $B_{i,i}$ quantify self-excitation within market i, whereas off-diagonal entries $B_{j,i}$ for $j \neq i$ quantify cross-excitation, or contagion, from market j to market i (Embrechts et al., 2011; Bacry et al., 2015).

When severity marks enter the intensity multiplicatively, events no longer have the same reproductive strength. A realized event (j,n) with mark $m_{j,n}$ generates expected direct type-i offspring equal to $m_{j,n}\,B_{j,i}$. For stability, however, what matters is the average reproductive strength across source-market events rather than the strength of a single realized event. Treating marks as random variables, or in practice summarizing them by their sample means, the relevant effective marked offspring matrix is

$$B^{\mathrm{mark}}_{j,i} = \mathbb{E}[m \mid j]\,B_{j,i}, \quad B^{\mathrm{mark}} = \mathrm{diag}\big(\mathbb{E}[m \mid 1], \ldots, \mathbb{E}[m \mid d]\big)B.$$

This matrix captures the average amplification induced by the observed marks and is the effective offspring matrix used to assess stability in the marked specification. The standard subcriticality condition is

$$\rho\big(B^{\mathrm{mark}}\big) < 1,$$

where $\rho(\cdot)$ denotes the spectral radius (Embrechts et al., 2011). As in the univariate case, values close to one indicate near-critical behavior, where feedback effects are strong enough for event clusters to become large and persistent.

2.5.3. Estimation

We estimate the model on the observation window $[0,T]$. Let $\theta := (\mu, B, \beta)$ denote the parameter vector. Following Daley and Vere-Jones (2003), for a multivariate point process with conditional intensities, the likelihood on $[0,T]$ is

$$\mathcal{L}(\theta) = \exp\left\{-\sum_{i=1}^{d}\int_0^T \lambda_i(t \mid \mathcal{H}_{t^-})\,dt\right\}\prod_{k=1}^{N}\lambda_{d_k}\left(t_k \mid \mathcal{H}_{t_k^-}\right). \tag{6}$$

The corresponding log-likelihood is

$$\ell(\theta) = \sum_{k=1}^{N}\log\lambda_{d_k}(t_k \mid \mathcal{H}_{t_k^-}) - \sum_{i=1}^{d}\int_0^T \lambda_i(t \mid \mathcal{H}_{t^-})\,dt. \tag{7}$$

A fully specified marked point process would also specify a conditional mark density $f_i^*(m \mid t, \mathcal{H}_{t^-})$ and would add the term $\sum_{k=1}^{N}\log f_{d_k}^*(m_k \mid t_k, \mathcal{H}_{t_k^-})$ to (7) (Embrechts et al., 2011). In this paper, marks are deterministic transformations of returns and are treated as fixed covariates. They therefore introduce no additional parameters beyond θ. We estimate the conditional-on-marks likelihood, with marks entering only through the conditional intensities in (5).

With the exponential kernel in (5), the compensator has a closed form. For each component i,

$$\int_0^T \lambda_i(t)\,dt = \mu_i T + \sum_{k=1}^{N} m_k\,B_{d_k,i}\,\left(1 - \mathrm{e}^{-\beta_i(T-t_k)}\right). \tag{8}$$

At an event time t_k of type d_k, the pre-event intensity is

$$\lambda_{d_k}(t_k \mid \mathcal{H}_{t_k^-}) = \mu_{d_k} + \sum_{\ell:\, t_\ell < t_k} m_\ell\,B_{d_\ell,d_k}\,\beta_{d_k}\,\mathrm{e}^{-\beta_{d_k}(t_k - t_\ell)}. \tag{9}$$

Substituting (8) and (9) into (7) gives

$$\ell(\mu, B, \beta) = \sum_{k=1}^{N} \log\left(\mu_{d_k} + \sum_{\ell:\, t_\ell < t_k} m_\ell\, B_{d_\ell, d_k}\, \beta_{d_k}\, \mathrm{e}^{-\beta_{d_k}(t_k - t_\ell)}\right) - \sum_{i=1}^{d} \mu_i T$$
$$- \sum_{i=1}^{d} \sum_{k=1}^{N} m_k\, B_{d_k, i}\, \left(1 - \mathrm{e}^{-\beta_i (T - t_k)}\right). \tag{10}$$

The numerical estimator is obtained by minimizing the regularized negative log-likelihood:

$$\mathcal{J}(\theta) = -\ell(\theta) + \lambda_B \|B\|_F^2 + \lambda_\beta \|\log \beta\|_2^2 + \lambda_s \sum_{i=1}^{d} \left[\sum_{j=1}^{d} B_{j,i}^{\text{mark}} - \bar{\rho}\right]_+^2. \tag{11}$$

The first two terms are ridge regularization terms on the branching matrix and on the logarithm of the decay rates. The last term is a stability barrier based on the column sums of B^{mark}. Since B^{mark} is nonnegative, bounding these column sums below one is a sufficient condition for subcriticality. The optimization is carried out in MATLAB using `fminunc` with analytical gradients. We ensure that μ_i, $B_{j,i}$, and β_i remain positive by carrying out the optimization in terms of their logarithms.

2.6. Goodness-of-fit

After estimation, we evaluate fit using the time-rescaling theorem (Daley and Vere-Jones, 2003). Define the compensator for each market i,

$$\Lambda_i(t) := \int_0^t \lambda_i(s)\, ds.$$

For the multivariate diagnostic used in the empirical analysis, we consider the superposed compensator, defined as

$$\Lambda_\bullet(t) := \sum_{i=1}^{d} \Lambda_i(t),$$

and the ordered pooled event times $\{t_1, t_2, \dots\}$, with the convention $t_0 := 0$. Under correct specification, the compensator increments

$$\Delta\Lambda_n := \Lambda_\bullet(t_n) - \Lambda_\bullet(t_{n-1}) \tag{12}$$

are i.i.d. Exp(1). For the Poisson and univariate Hawkes benchmarks, the same construction is applied to the corresponding univariate compensator. We assess this in two complementary ways: (i) qualitatively via Quantile–Quantile plots of empirical $\Delta\Lambda_n$ against theoretical Exp(1) quantiles, and (ii) quantitatively via one-sample Kolmogorov–Smirnov (KS) tests. For the KS test, the null hypothesis is $H_0 : \Delta\Lambda_n \overset{\text{i.i.d.}}{\sim} \text{Exp}(1)$, while the alternative hypothesis is $H_1 : \Delta\Lambda_n \not\sim \text{Exp}(1)$, which indicates a violation of the time-rescaling implication and therefore evidence of model misspecification.

2.7. Systemic risk stress interpretation

We further evaluate whether the estimated Hawkes intensities produce a sensible daily stress signal when compared to the CBOE Volatility Index (VIX)[2]. Let V_t denote the daily VIX level on business days over each crisis window. The model-implied stress score is defined as $S_t = \lambda(t)$ for the univariate Hawkes, and $S_t = \sum_{i=1}^{d} \lambda_i(t)$ for the multivariate Hawkes. Because S_t is measured in expected events per day, while VIX is reported in annualized volatility points, we do not compare levels directly. Instead, we use an in-sample linear calibration after standardization:

$$V_t = c_0 + c_1\, z(S_t) + \varepsilon_t, \quad z(S_t) := \frac{S_t - \overline{S}}{\mathrm{sd}(S)}. \tag{13}$$

The coefficients (c_0, c_1) are estimated by ordinary least squares (OLS) over the crisis window, and fitted values $\widehat{V}_t$ are computed. Model performance is summarized by the coefficient of determination (R^2), the correlation between V_t and $\widehat{V}_t$, and the root mean squared error ($\mathrm{RMSE} := \sqrt{\frac{1}{n}\sum_{t=1}^{n}(V_t - \widehat{V}_t)^2}$).

3. Empirical application

This section applies the framework to two crisis episodes. The first is the Global Financial Crisis (GFC) of 2008, which we interpret as a crisis dominated by endogenous amplification. The second episode is the COVID-19 crisis of 2020, which we interpret as a crisis triggered by a largely exogenous shock, not generated within the financial system itself, but by the sudden arrival of a global health emergency and the associated policy restrictions. For each episode, we study three objects: goodness-of-fit diagnostics based on the time-rescaling theorem, the branching (contagion) matrix as a reduced-form contagion network, and the decomposition of total intensity into market equity sectors, in order to highlight the main drivers of the intensity.

3.1. Sample and event construction

We work with the daily log-returns of the primary S&P 500 Select Sector SPDR Exchange-Traded Funds (ETFs). Specifically, our sample comprises the Materials Select Sector SPDR Fund (XLB), the Energy Select Sector SPDR Fund (XLE), the Financial Select Sector SPDR Fund (XLF), the Industrial Select Sector SPDR Fund (XLI), the Technology Select Sector SPDR Fund (XLK), the Utilities Select Sector SPDR Fund (XLU), the Health Care Select Sector SPDR Fund (XLV), and the Consumer Discretionary Select Sector SPDR Fund (XLY). The daily price and return data for these assets are retrieved from the Center for Research in Security Prices (CRSP) database.[3]. This broad set allows us to investigate amplification and contagion within and across sectors. We define a first window to cover the GFC from `2007-06-01` to `2011-01-01`, a period that spans the buildup, acute phase, and early normalization of the crisis. We define a second window to cover the COVID-19 shock from `2019-06-01` to `2024-01-01`, encompassing the onset of the pandemic, the sharp market correction in March 2020, and the recovery phase.

[2]The VIX is an annualized implied-volatility index derived from S&P 500 option prices and is commonly used as a proxy for market-wide uncertainty or stress. It reflects variance risk pricing and risk-neutral expectations.

[3]Data is sourced from the CRSP Daily Stock File (DSF), accessed via Wharton Research Data Services (WRDS).

Window	Sector	Ticker	No. events	Start	End	Mean mark	Max mark	Mean gap
GFC	Materials	XLB	17	2007-12-17	2010-04-27	0.8449	2.4259	53.875
	Energy	XLE	22	2007-07-26	2010-06-01	1.7799	6.4763	49.571
	Financials	XLF	39	2007-11-01	2009-09-01	2.0714	7.3609	17.632
	Industrials	XLI	11	2008-04-11	2010-08-11	0.5277	1.0878	85.200
	Technology	XLK	17	2007-11-08	2010-06-29	0.6592	1.7901	60.250
	Utilities	XLU	23	2007-06-07	2010-06-22	0.4388	1.0717	50.500
	Health Care	XLV	15	2007-08-09	2010-10-19	0.4643	1.8521	83.357
	Consumer Discretionary	XLY	15	2008-04-30	2010-07-16	0.4730	1.8062	57.643
COVID-19	Materials	XLB	9	2020-02-25	2023-10-18	0.6240	1.8392	166.380
	Energy	XLE	40	2020-03-06	2023-05-02	1.6777	12.7990	29.538
	Financials	XLF	14	2019-08-14	2023-03-17	1.3352	8.5872	100.850
	Industrials	XLI	10	2019-07-17	2023-04-04	0.7223	2.8610	150.780
	Technology	XLK	23	2019-08-05	2022-12-15	0.8569	2.4103	55.818
	Utilities	XLU	39	2019-06-26	2023-12-20	0.8569	4.5493	43.105
	Health Care	XLV	16	2019-10-08	2023-11-09	0.5048	2.6747	99.533
	Consumer Discretionary	XLY	27	2021-02-25	2023-10-19	0.7746	3.0238	37.154

Table 1: Event-sample summary statistics for the GFC and COVID windows across U.S. sector ETFs. Mean gap is measured in trading days.

Extreme events are detected using left-tail return exceedances. For each sector and each crisis window, an event is recorded when the daily log-return falls below a static threshold equal to the empirical 5th percentile within that window. Each event is assigned a nonnegative severity mark based on the exceedance magnitude, transformed and normalized as described in Section 2. Table 1 reports event-level summary statistics for all sectors in both crisis windows. For the multivariate Hawkes estimation, sectors with fewer than 10 events are excluded from the corresponding fit. This rule affects the COVID-19 window, where XLB has only 9 events and is therefore reported in the summary table but excluded from the multivariate process estimation.

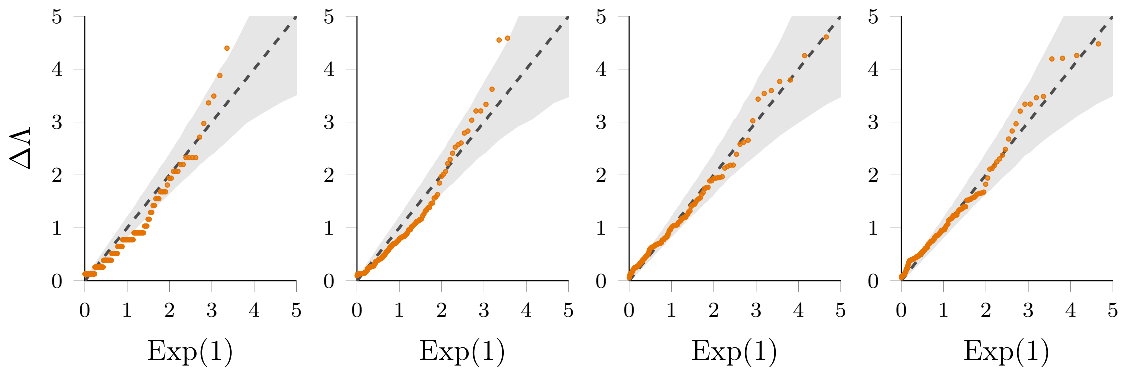

Figure 1: Goodness-of-fit diagnostics for the GFC window. The figure reports Quantile–Quantile plots of the transformed increments $\Delta\Lambda$ against the Exp(1) distribution. From left to right, the figure reports: Homogeneous Poisson, Inhomogeneous Poisson, Univariate Hawkes, Multivariate Hawkes. Under correct specification, the transformed increments should approximately follow a unit exponential distribution.

3.2. Global Financial Crisis (2008)

The GFC window is used as the benchmark case for endogenous financial amplification. In figure 1 we report the time-rescaling diagnostics for the GFC window. The homogeneous Poisson benchmark is strongly rejected, with a KS statistic of 0.158 and a p-value below 0.001; this indicates that a model with independent arrivals is not adequate for the timing of GFC stress events. Allowing for deterministic time variation improves the fit: the inhomogeneous Poisson model reduces the KS statistic to 0.109. However, the null is still rejected at the 5% level, which indicates that deterministic variation alone does not fully account for the observed clustering.

The best diagnostic performance is obtained by the univariate Hawkes model; the KS statistic is the lowest among the four specifications, 0.103, and the corresponding p-value is 0.0667, meaning that univerate Hawkes transformed residuals are not rejected at the 5% level and are consistent with the unit-exponential benchmark. Conversely, the multivariate Hawkes residuals produce a weaker diagnostic outcome, with a KS statistic of 0.141 and a p-value of 0.0035. This suggests that the multivariate model, while useful for decomposing persistence and cross-sector propagation, leaves residual structure when all sectoral residuals are pooled into a single superposed diagnostic.

Source $\to$ Target	XLB	XLE	XLF	XLI	XLK	XLU	XLV	XLY
XLB	0.04	0.02	0.02	0.01	0.14	0.00	0.22	0.19
XLE	0.05	0.07	0.06	0.01	0.05	0.06	0.08	0.05
XLF	0.08	0.12	0.81	0.01	0.06	0.04	0.02	0.01
XLI	0.00	0.00	0.00	0.10	0.00	0.01	0.00	0.01
XLK	0.03	0.36	0.00	0.00	0.08	0.07	0.01	0.07
XLU	0.01	0.04	0.00	0.00	0.00	0.39	0.10	0.00
XLV	0.03	0.02	0.00	0.00	0.01	0.03	0.03	0.01
XLY	0.01	0.00	0.04	0.12	0.05	0.01	0.00	0.06
$\rho(\widehat{B}^{\text{mark}}) = 0.840$.								

Table 2: Contagion matrix for the multivariate Hawkes model during the GFC window. Rows denote source sectors and columns denote target sectors.

Table 2 reports the (marked) contagion matrix for the GFC window. The most relevant feature is the large diagonal coefficient for XLF, with $\widehat{B}^{\text{mark}}_{\text{XLF}\to\text{XLF}} = 0.81$. This is the largest entry in the matrix and indicates strong within-sector persistence in financial stress events. The result is consistent with the economic narrative of the GFC: the deterioration of mortgage-related assets in 2007 was reinforced by funding stress, declining confidence in financial institution balance sheets, deleveraging, and liquidity pressure. In this sense, the high XLF diagonal coefficient provides reduced-form evidence of persistent stress within the financial sector. XLU shows moderate persistence, with $\widehat{B}^{\text{mark}}_{\text{XLU}\to\text{XLU}} = 0.39$, while all other sectors display a weaker feedback mechanism. The estimates therefore suggest that endogenous amplification during the GFC was not evenly distributed across sectors; it was concentrated primarily in financials, while most non-financial sectors were more exposed to incoming stress than to persistent internal feedback.

The off-diagonal coefficients provide additional information on how stress moved across sectors. The strongest off-diagonal coefficient is from technology to energy, with $\widehat{B}^{\text{mark}}_{\text{XLK}\to\text{XLE}} = 0.36$. Materials also transmit to health care and consumer discretionary, with coefficients of 0.22 and 0.19, respectively, and to technology with a coefficient of 0.14. Financial sector events also contribute to subsequent energy stress, with $\widehat{B}^{\text{mark}}_{\text{XLF}\to\text{XLE}} = 0.12$. These values suggest that the GFC did not propagate only through a single financial channel; we observe that sectoral stress also moved through broader market linkages, showing a contagion-like mechanism.

The estimated spectral radius is $\rho(\widehat{B}^{\text{mark}}) = 0.840$, indicating that the estimated process is subcritical but close to the stability boundary. This is an important result for the GFC interpretation, since the model does not describe an explosive process, but it does describe a highly persistent one. Overall, our results provides evidence of relevance of the endogenous component of the Hawkes intensity, with the financial sector self-excitation playing the central role.

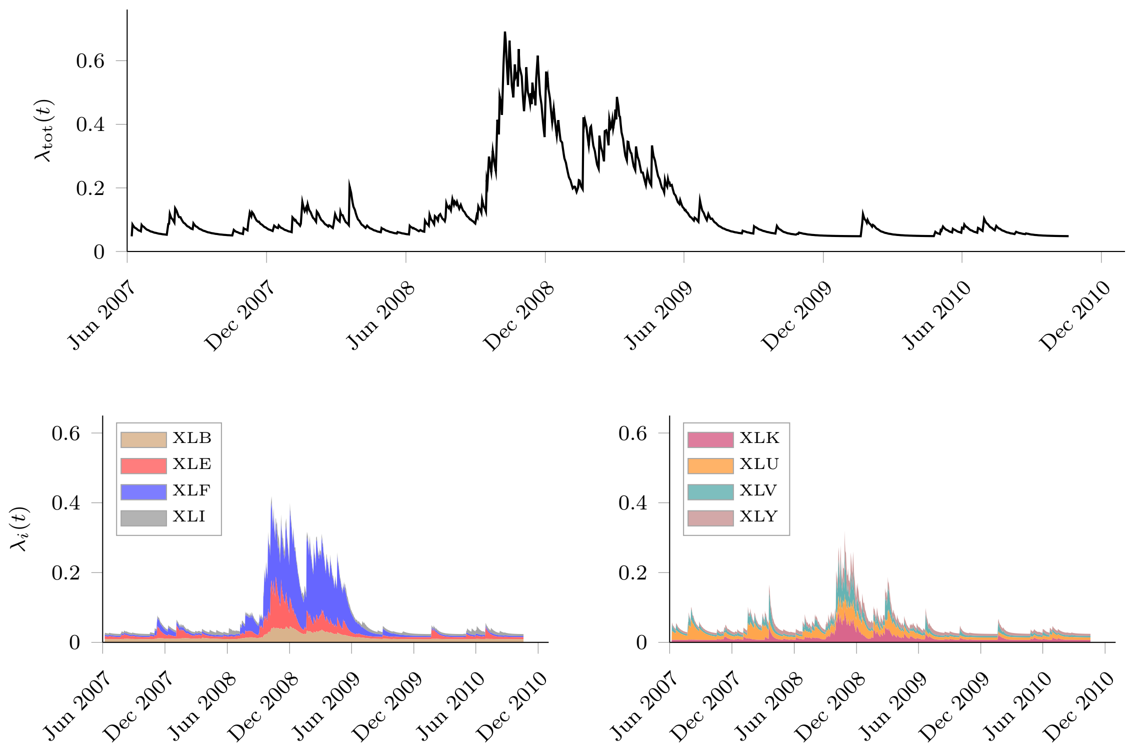

Figure 2: Hawkes intensity during the GFC (June 2007–December 2010). The top panel reports the aggregate systemic intensity $\lambda_{\text{tot}}(t) = \sum_{i=1}^{d} \lambda_i(t)$. The bottom panels show the sectoral decomposition of the fitted multivariate Hawkes intensity.

Figure 2 reports the fitted intensity path for the GFC window. This object is useful because it translates the estimated Hawkes model into a time-varying measure of event risk and can be used to directly monitor the contribution of each sector over time. The aggregate intensity $\lambda_{\text{tot}}(t)$ shows when the model assigns a higher probability of observing new stress events, while the sectoral panels show which parts of the system account for that increase. In the GFC application, the aggregate intensity rises during the central phase of the crisis and remains elevated for a persistent period. This pattern indicates that stress events are not absorbed immediately after they occur, but rather, past events continue to contribute to the conditional intensity for some time, generating a sequence of elevated event risk consistent with the critical spectral radius $\rho(\widehat{B}^{\text{mark}}) = 0.840$

3.3. COVID-19 crisis (2020)

The COVID-19 window can be interpreted as a contrasting crisis episode compared to the GFC, since the initial trigger was largely outside the financial system. The pandemic and the associated restrictions changed expectations about economic activity, earnings, liquidity needs, and policy responses over a very short period. Also in this case, we report the time-rescaling diagnostics in Figure 3. The homogeneous Poisson benchmark is rejected at the 5% level, with a KS statistic of 0.117 and a p-value of 0.0190. For the inhomogeneous Poisson model, we get similar diagnostic result, with a KS statistic of 0.116 and a p-value of 0.0198; allowing for a deterministic trend in the arrival rate therefore does not materially improve the residual fit in this window.

The univariate Hawkes model reduces the KS statistic to 0.088 and is not rejected by the KS test, with a p-value of 0.1393. However, in this case, the best diagnostic performance is obtained by the multivariate superposed Hawkes residuals, which have the

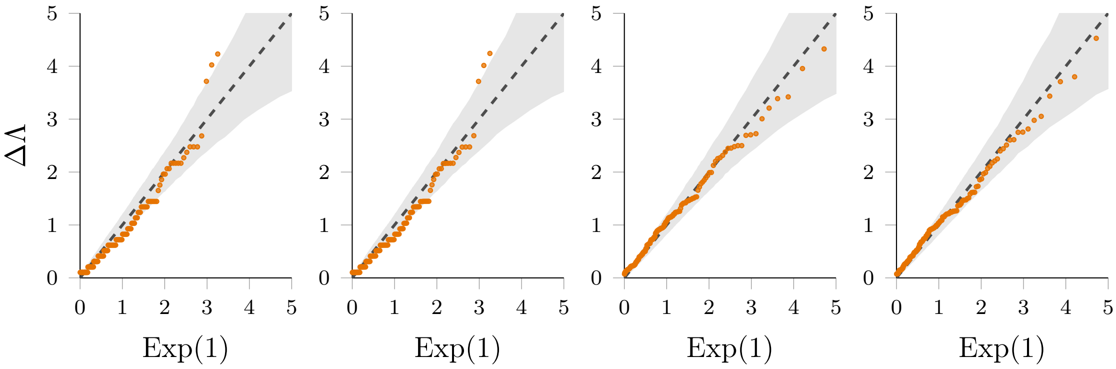

Figure 3: Goodness-of-fit diagnostics for the COVID window. The figure reports Quantile-Quantile plots of the transformed increments $\Delta\Lambda$ against the Exp(1) distribution. From left to right, the figure reports: Homogeneous Poisson, Inhomogeneous Poisson, Univariate Hawkes, Multivariate Hawkes. Under correct specification, the transformed increments should approximately follow a unit exponential distribution.

lowest KS statistic, 0.074, and the highest p-value, 0.3086. This result indicates that the events related to the COVID-19 crisis are better described once dependence is introduced, and that the multivariate structure absorbs additional residual clustering relative to the aggregate univariate specification.

Table 3 presents the (marked) contagion matrix for the COVID period. The spectral radius is $\rho(\widehat{B}^{\mathrm{mark}}) = 0.576$, substantially lower than the value obtained for the GFC. The process is still subcritical, but at the same time less persistent than in the GFC window, indicating a crisis pattern characterized by a large initial shock, followed by a market transmission mechanism that, while present, remained comparatively contained.

Source $\to$ Target	XLE	XLF	XLI	XLK	XLU	XLV	XLY
XLE	0.42	0.06	0.00	0.01	0.14	0.01	0.01
XLF	0.08	0.17	0.05	0.01	0.20	0.04	0.01
XLI	0.13	0.03	0.17	0.01	0.00	0.03	0.00
XLK	0.12	0.06	0.06	0.21	0.08	0.00	0.00
XLU	0.03	0.01	0.00	0.02	0.20	0.03	0.01
XLV	0.06	0.00	0.05	0.00	0.00	0.04	0.02
XLY	0.10	0.01	0.01	0.00	0.01	0.07	0.56
$\rho(\widehat{B}^{\mathrm{mark}}) = 0.576$.							

Table 3: Contagion matrix for the multivariate Hawkes model during the COVID-19 window. Rows denote source sectors and columns denote target sectors.

The largest within-sector coefficient is for consumer discretionary, with $\widehat{B}^{\mathrm{mark}}_{\mathrm{XLY}\to\mathrm{XLY}} = 0.56$. We can interpret this results directly looking at the timeline of events. In particular, it is consistent with the direct exposure of the sector to lockdowns, mobility restrictions, changes in household spending, and the distinction between firms harmed by restrictions and firms benefiting from stay-at-home demand. The energy sector also shows relevant within-sector persistence, with $\widehat{B}^{\mathrm{mark}}_{\mathrm{XLE}\to\mathrm{XLE}} = 0.42$. In the first few months of the COVID-19 crisis, we observed sudden decline in mobility and production that reduced expected oil demand, but at the same time we observed a relevant oil-market turmoil in March and April 2020 affecting the supply side; therefore, this results seems to be economically plausible for this window, if we consider this double effect in energy markets through both demand and price channels. As an intuitively expected result, the financial sector behaves differently from the GFC case. The diagonal coefficient for XLF is $\widehat{B}^{\text{mark}}_{\text{XLF}\to\text{XLF}} = 0.17$, far below the GFC value. This is one of the main empirical differences between the two crises, and finds also a trivial interpretation.

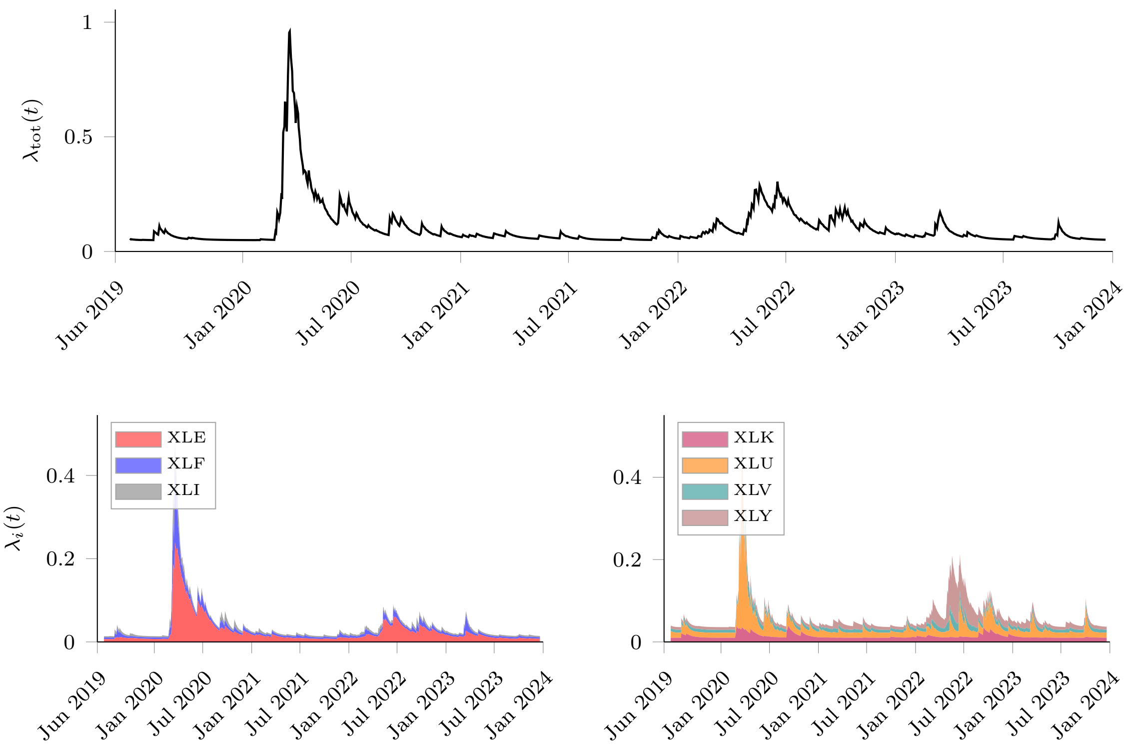

Figure 4: Hawkes intensity during the COVID window (June 2019–January 2024). The top panel reports the aggregate systemic intensity $\lambda_{\text{tot}}(t) = \sum_{i=1}^{d} \lambda_i(t)$. The bottom panels show the sectoral decomposition of the fitted multivariate Hawkes intensity. Note: XLB is excluded because it does not satisfy the minimum event threshold.

During COVID-19, financial stress was relevant, but at the same time the financial sector was not the main internal source of persistence stress in the market; rather, the shock originated from the real economy and public health restrictions, and the financial response was also shaped by rapid policy support and liquidity interventions. Looking at the off-diagonal estimates, the strongest cross-section propagation is from financials to utilities, with $\widehat{B}^{\text{mark}}_{\text{XLF}\to\text{XLU}} = 0.20$. Energy also transmits to utilities, with $\widehat{B}^{\text{mark}}_{\text{XLE}\to\text{XLU}} = 0.14$. These links suggest that defensive and "rate sensitive" sectors were not insulated from market stress; interestingly, we observe that even sectors usually viewed as relatively stable can become connected to broader liquidity and repricing dynamics during a sudden system-wide market shock. Industrials transmit to energy with $\widehat{B}^{\text{mark}}_{\text{XLI}\to\text{XLE}} = 0.13$, technology transmits to energy with $\widehat{B}^{\text{mark}}_{\text{XLK}\to\text{XLE}} = 0.12$, and consumer discretionary transmits to energy with $\widehat{B}^{\text{mark}}_{\text{XLY}\to\text{XLE}} = 0.10$. We believe this pattern is coherent with the nature of the pandemic shock; in particular, expectations about production, transport, consumption, and aggregate demand were revised at the same time, and these revisions had direct implications for energy demand and energy sector valuations.

Figure 4 reports the fitted intensity path for the COVID window. As in the GFC application, the top panel shows the total conditional intensity, while the lower panels show how much each sector contributes to that total at different points in time. The aggregate intensity increases sharply around the pandemic stress period. We note that, compared to the GFC window, the rise is more concentrated in time and the normalization is faster; this is consistent with a much stronger initial magnitude of the shock but, at the same time, also with the lower spectral radius reported in Table 3.

3.4. Baseline comparison

The calibration exercise in Table 4 compares how well the Hawkes model implied stress signals track the market implied VIX index during each crisis.

Model	R^2 **(%)**	**Correlation**	**RMSE**	β **(Slope)**	t**-statistic**
GFC					
Univariate Hawkes	71.16	0.844	6.290	9.874	41.17
Multivariate Marked Hawkes	86.31	0.929	4.334	10.874	65.81
DCC-GARCH(1,1)	85.62	0.925	4.442	10.831	63.96
COVID-19					
Univariate Hawkes	59.20	0.769	5.395	6.494	31.04
Multivariate Marked Hawkes	76.13	0.873	4.126	7.364	46.02
DCC-GARCH(1,1)	73.99	0.860	4.308	7.260	43.46

Table 4: In-sample calibration performance to the VIX index. The table reports the OLS regression statistics where the dependent variable is the daily VIX level and the independent variables are the standardized z-scores of the respective model intensities.

Across both episodes, the multivariate marked Hawkes stress signal attains the highest in-sample R^2 and correlation with VIX and the lowest RMSE among the specifications considered. During the GFC, the multivariate Hawkes calibration delivers an R^2 of 86.3%, slightly above the DCC-GARCH(1,1) benchmark and substantially above the univariate Hawkes specification. During the COVID-19, the ranking is similar, although fit levels are lower overall: the multivariate Hawkes stress signal attains an R^2 of 76.1%. The weaker performance of the univariate Hawkes model in the pandemic window is consistent with the importance of cross-sector propagation during a broad market-wide adjustment. We interpret these results as an in-sample calibration of the methodology; particularly, they indicate that event-based Hawkes intensities contain information about systemic stress that is complementary to standard volatility-based benchmarks. Additional visual comparisons between the estimated Hawkes intensities, the VIX index, and the DCC-GARCH(1,1) annualised volatility are reported in Appendix 4.2.

4. Discussion

In this paper, we developed and illustrated an event-driven approach to systemic stress based on self-exciting point processes. The central idea is that crises are often experienced as clusters of extreme losses that propagate across markets and persist after major market shocks. Hawkes processes provide a probabilistic approach to model this mechanism. In particular, estimation of Hawkes processes provide (i) conditional intensities that can be interpreted as near-term hazard rates for extreme events, (ii) a branching (contagion) matrix that summarizes propagation channels, and (iii) stability diagnostics that summarize proximity to cascade regimes.

In the empirical application, we compared two crisis episodes with different dominant mechanisms, namely the Global Financial Crisis and the COVID-19 crisis. In both windows, the Poisson benchmarks provide weaker time-rescaling diagnostics than the Hawkes specifications, indicating that independent arrivals are too restrictive to model crisis events over time. If we relax this assumption, for the GFC window the univariate Hawkes model provides the strongest aggregate timing fit, while the multivariate model is mainly informative as a sectoral propagation object. Conversely, in the COVID-19 window, the multivariate superposed diagnostic performs best, suggesting that the cross-sector structure helps absorb the clustering generated by the rapid market-wide adjustment. The contagion matrices show that the two crises have different feedback structures. In the GFC window, the marked spectral radius is high, $\rho(\widehat{B}^{\mathrm{mark}}) = 0.840$, and the largest matrix entry is the financial sector diagonal coefficient, $\widehat{B}^{\mathrm{mark}}_{\mathrm{XLF}\to\mathrm{XLF}} = 0.81$. This points to a persistent financial sector component in the crisis dynamics, and results are consistent with the sequence of the crisis, in which mortgage losses, funding stress, balance-sheet pressure, and deleveraging repeatedly affected financial institutions before the crisis spread more broadly across markets. The COVID-19 window has a different profile. The estimated spectral radius is lower, $\rho(\widehat{B}^{\mathrm{mark}}) = 0.576$, still indicating a clearly subcritical process, but with persistence particularly less centered on financials and more visible in sectors directly exposed to the real economy shock. In this case, consumer discretionary has the largest diagonal coefficient, followed by the energy sector, while the financial sector diagonal coefficient is much smaller than in the GFC window. This is consistent with the nature of the COVID shock, which affected mobility, consumption, production, energy demand, and expectations about activity over a short period. For both crisis, we report the intensity decomposition as a useful monitoring object for policy analysis, showing how different components contribute to the aggregate intensity over time.

Additionally, we examine how the calibrated Hawkes intensities reproduce VIX dynamics following a linear calibration, and we benchmark their performance against a DCC-GARCH specification. For the crisis episodes considered in this study, we observe that Hawkes models perform competitively with the DCC-GARCH benchmark. As such, they serve as a useful complement to traditional volatility models, while capturing a distinct yet relevant crisis-related feature: the transmission of tail-event risk over time and across different segments of the market.

Overall, we illustrate how existing Hawkes process tools can be applied in a coherent way to model clustered crisis events, extract interpretable measures of persistence and contagion, compare different crisis regimes, and support analyses of systemic amplification in financial systems.

4.1. Limitations

Several limitations matter for interpretation, especially in a macro-financial stability context.

The multivariate model assumes constant μ_i in the baseline specification, implying time-homogeneous exogenous risk within each crisis window. In reality, macroeconomic conditions, policy regimes, and structural changes shift the background hazard of stress events. A constant baseline can therefore mechanically attribute slow-moving exogenous variation to endogenous excitation. We also note that the threshold choice and daily discretization rule introduce modeling artifacts that should be tested for robustness. Thresholds determine which return movements become events, and therefore affect both event counts and estimated branching coefficients. Similarly, point estimates of (μ, B, β) and $\widehat{B}^{\mathrm{mark}}$ are informative, but macro-prudential interpretation ultimately requires uncertainty bands and stability checks across thresholds, samples, and crisis windows.

4.2. Future research

The most natural extensions follow directly from the highlighted limitations and from the financial-stability motivation of this study.

A key step is to replace constant μ_i with covariate-linked baselines $\mu_i(t)$ that reflect macro-financial conditions. Conceptually, this addresses the "black swan vs. internal fragility" distinction emphasized in Section 1 by allowing macro variables to move the baseline, while the endogenous kernel quantify propagation. Furthermore, a distinctive advantage of the Hawkes representation is that it supports generative scenario analysis. Given θ, one can simulate event cascades under counterfactual parameter changes and summarize systemic outcomes, such as the expected number of tail events, the expected time spent above critical stress thresholds, and cross-market propagation probabilities. From a policy perspective, this enables a mapping from interventions to propagation parameters, and we believe can be useful for stress-testing exercises for central banks or similar institutions. Finally, crisis dynamics may require multi-scale kernels and broader multivariate systems, including rates, credit, foreign exchange and commodity markets.

Acknowledgements

The authors thank Aleksandar Krajisnik, whose Bachelor thesis, *An Empirical Analysis of Systemic Risk with Self-Exciting Point Processes*, submitted at the Zurich University of Applied Sciences, School of Management and Law, Department of Banking, Finance, Insurance, in Spring Semester 2025, provided the initial basis for this research project. The authors also acknowledge Università della Svizzera italiana (USI) for providing institutional access to Wharton Research Data Services (WRDS), which supported with data collection. The authors are also grateful to the colleagues and seminar participants who discussed earlier versions of this work and provided valuable comments and suggestions.

References

Adrian, T., Shin, H.S., 2010. Liquidity and leverage. Journal of Financial Intermediation 19, 418–437. doi:10.1016/j.jfi.2008.12.002.

Aït-Sahalia, Y., Cacho-Diaz, J., Laeven, R.J.A., 2015. Modeling financial contagion using mutually exciting jump processes. Journal of Financial Economics 117, 585–606. doi:10.1016/j.jfineco.2015.03.002.

Allen, F., Gale, D., 2000. Financial contagion. Journal of Political Economy 108, 1–33. doi:10.1086/262109.

Bacry, E., Mastromatteo, I., Muzy, J., 2015. Hawkes processes in finance. Market Microstructure and Liquidity 1, 1550005. doi:10.1142/S2382626615500057.

Bak, P., 1996. How Nature Works: The Science of Self-Organized Criticality. Copernicus, New York, NY. doi:10.1007/978-1-4757-5426-1.

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 31, 307–327. doi:10.1016/0304-4076(86)90063-1.

Bouchaud, J., 2024. The self-organized criticality paradigm in economics & finance. `arXiv:2407.10284`. arXiv preprint.

Browning, R., Sulem, D., Mengersen, K., Rivoirard, V., Rousseau, J., 2021. Simple discrete-time self-exciting models can describe complex dynamic processes: A case study of COVID-19. PLOS ONE 16, e0250015. doi:10.1371/journal.pone.0250015.

Brunnermeier, M.K., Pedersen, L.H., 2009. Market liquidity and funding liquidity. The Review of Financial Studies 22, 2201–2238. doi:10.1093/rfs/hhn098.

Daley, D.J., Vere-Jones, D., 2003. An Introduction to the Theory of Point Processes. Volume I: Elementary Theory and Methods. Probability and Its Applications. 2 ed., Springer, New York, NY. doi:10.1007/b97277.

Dassios, A., Zhao, H., 2011. A dynamic contagion process. Advances in Applied Probability 43, 814–846. doi:10.1239/aap/1316792671.

Diebold, F.X., Yilmaz, K., 2012. Better to give than to receive: Predictive directional measurement of volatility spillovers. International Journal of Forecasting 28, 57–66. doi:10.1016/j.ijforecast.2011.02.006.

Embrechts, P., Liniger, T., Lin, L., 2011. Multivariate hawkes processes: An application to financial data. Journal of Applied Probability 48, 367–378. doi:10.1017/S0021900200099344.

Engle, R., 2002. Dynamic conditional correlation: A simple class of multivariate GARCH models. Journal of Business & Economic Statistics 20, 339–350. doi:10.1198/073500102288618487.

Errais, E., Giesecke, K., Goldberg, L.R., 2010. Affine point processes and portfolio credit risk. SIAM Journal on Financial Mathematics 1, 642–665. doi:10.1137/090771272.

Filimonov, V., Sornette, D., 2012. Quantifying reflexivity in financial markets: Toward a prediction of flash crashes. Physical Review E 85, 056108. doi:10.1103/PhysRevE.85.056108.

Forbes, K.J., Rigobon, R., 2002. No contagion, only interdependence: Measuring stock market comovements. The Journal of Finance 57, 2223–2261. doi:10.1111/1540-6261.00442.

Hainaut, D., 2016. A bivariate hawkes process for interest rate modeling. Economic Modelling 57, 180–196. doi:10.1016/j.econmod.2016.04.016.

Hawkes, A.G., 1971. Spectra of some self-exciting and mutually exciting point processes. Biometrika 58, 83–90. doi:10.1093/biomet/58.1.83.

Hawkes, A.G., 2018. Hawkes processes and their applications to finance: A review. Quantitative Finance 18, 193–198. doi:10.1080/14697688.2017.1403131.

Hawkes, A.G., Oakes, D., 1974. A cluster process representation of a self-exciting process. Journal of Applied Probability 11, 493–503. doi:10.2307/3212693.

Holló, D., Kremer, M., Lo Duca, M., 2012. CISS – A Composite Indicator of Systemic Stress in the Financial System. Working Paper Series 1426. European Central Bank. Frankfurt am Main, Germany.

Large, J., 2007. Measuring the resiliency of an electronic limit order book. Journal of Financial Markets 10, 1–25. doi:10.1016/j.finmar.2006.09.001.

Liu, W., Zhu, S., 2019. Pricing variance swaps under the hawkes jump-diffusion process. Journal of Futures Markets 39, 635–655. doi:10.1002/fut.21997.

Mohler, G., 2014. Marked point process hotspot maps for homicide and gun crime prediction in chicago. International Journal of Forecasting 30, 491–497. doi:10.1016/j.ijforecast.2014.01.004.

Mohler, G.O., Short, M.B., Brantingham, P.J., Schoenberg, F.P., Tita, G.E., 2011. Self-exciting point process modeling of crime. Journal of the American Statistical Association 106, 100–108. doi:10.1198/jasa.2011.ap09546.

Ogata, Y., 1988. Statistical models for earthquake occurrences and residual analysis for point processes. Journal of the American Statistical Association 83, 9–27. doi:10.1080/01621459.1988.10478560.

Sornette, D., 2003. Why Stock Markets Crash: Critical Events in Complex Financial Systems. Princeton University Press, Princeton, NJ and Oxford.

Taleb, N.N., 2007. The Black Swan: The Impact of the Highly Improbable. Random House, New York, NY.

Zhao, Q., Erdogdu, M.A., He, H.Y., Rajaraman, A., Leskovec, J., 2015. SEISMIC: A self-exciting point process model for predicting tweet popularity, in: Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, New York, NY, USA. pp. 1513–1522. doi:10.1145/2783258.2783401.

Zhuang, J., 2012. Long-term earthquake forecasts based on the epidemic-type aftershock sequence (ETAS) model for short-term clustering. Research in Geophysics 2, e8. doi:10.4081/rg.2012.e8.

Appendix

This section presents additional plots relating to the comparative analysis of the proposed measure of systemic risk and financial contagion based on the Hawkes models and the DCC-GARCH(1,1).

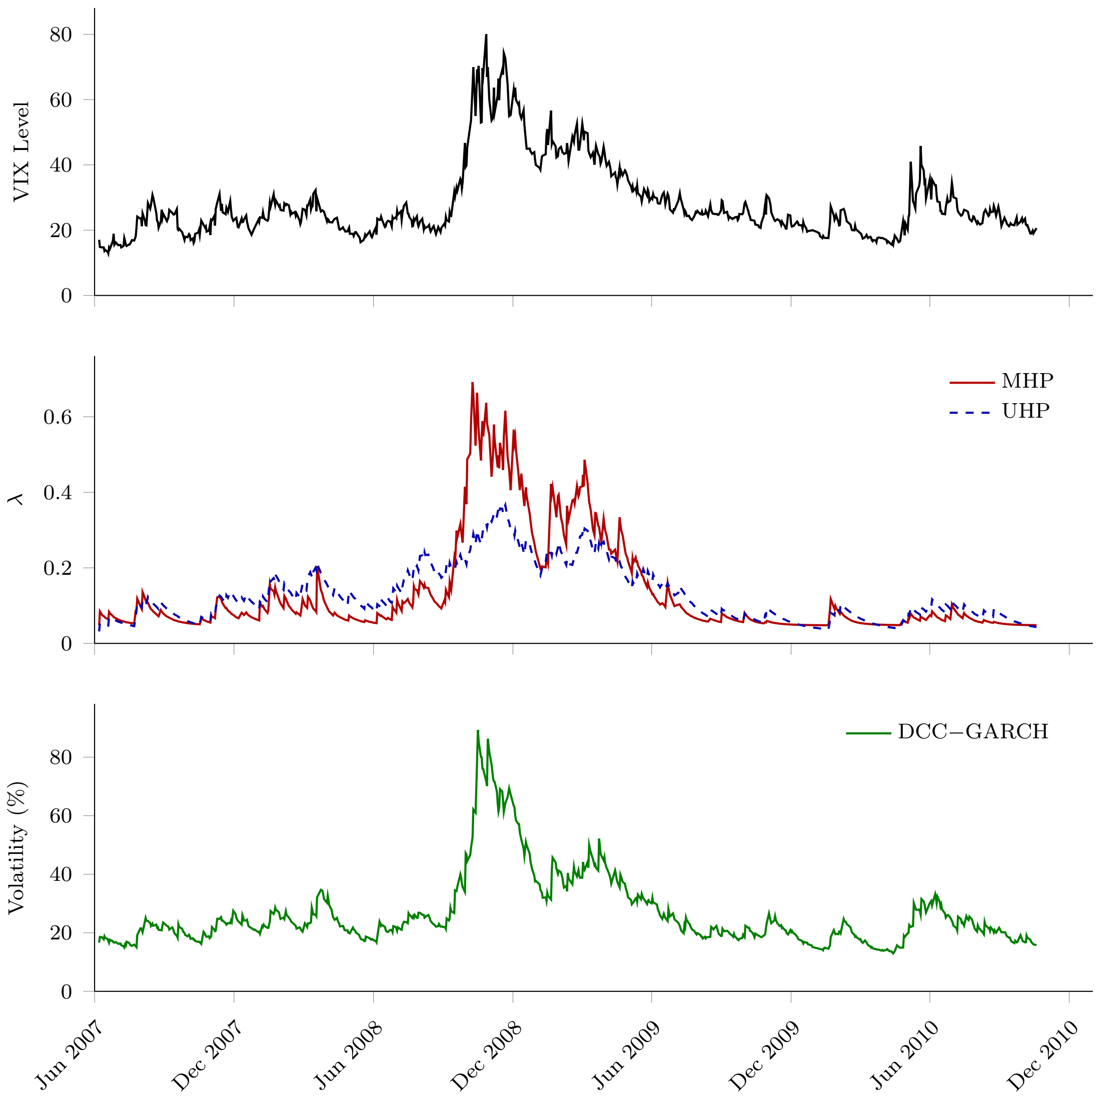

Figure 5: Visual comparison between the VIX index, the intensities estimated by Hawkes models (λ) and the annualised volatility of the DCC-GARCH model for the GFC window.

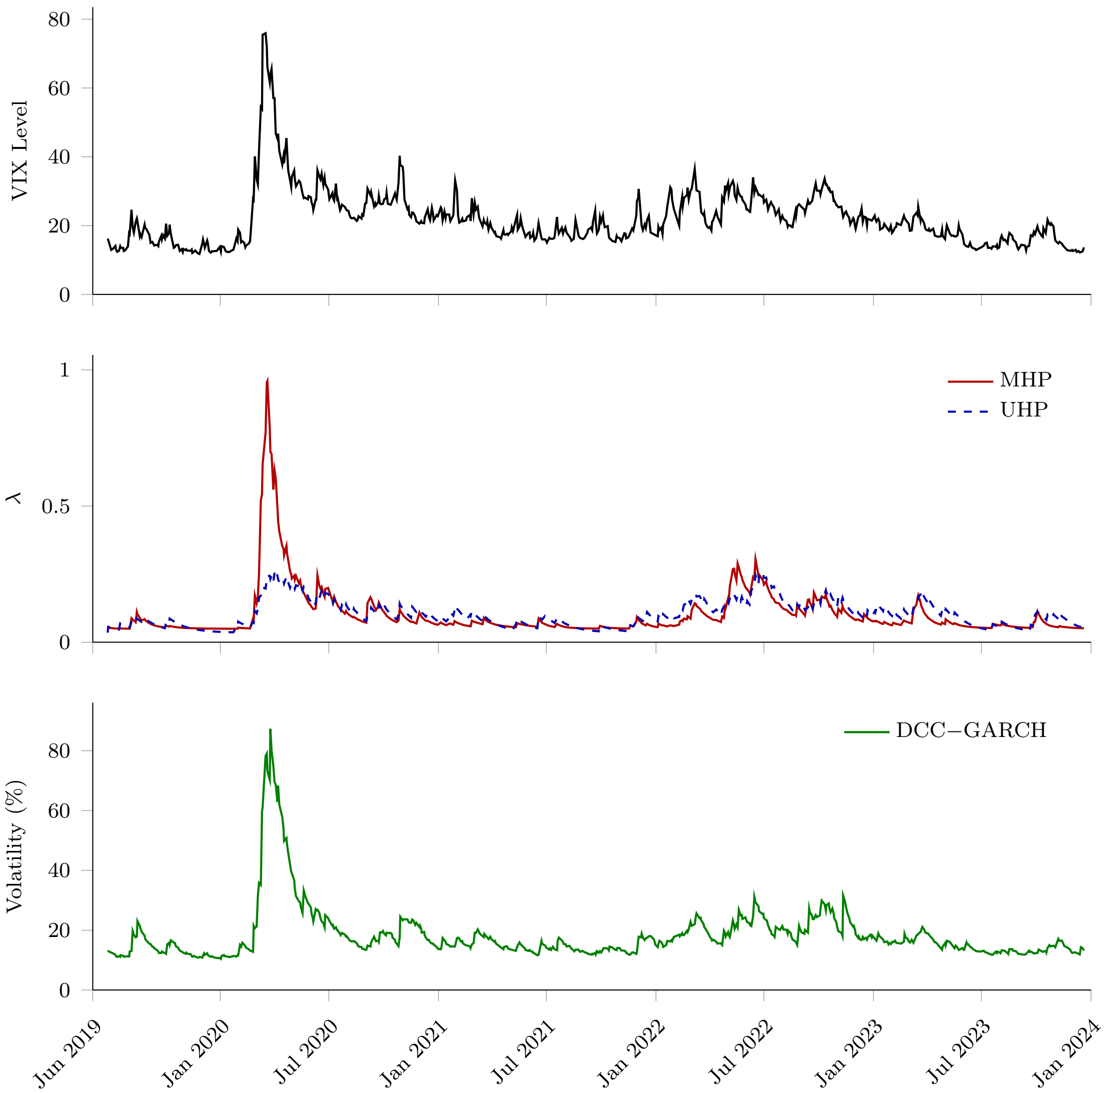

Figure 6: Visual comparison between the VIX index, the intensities estimated by Hawkes models (λ) and the annualised volatility of the DCC-GARCH model for the COVID-19 window.