

Saddlepoint Approximations for Hawkes Jump-Diffusion Processes with an Application to Risk Management*

Yacine Aït-Sahalia
Department of Economics, Princeton University
and
Roger J. A. Laeven
Department of Quantitative Economics, University of Amsterdam

May 18, 2025

Abstract

We propose a statistical model based on Hawkes processes in which large financial losses can arise in close succession serially as well as cross-sectionally. We derive in closed form saddlepoint approximations to the tails of profit and loss distributions, both marginal and joint, and use them to construct explicit risk measure formulae that account for the fact that a given financial institution's losses make it more likely that that institution will experience further losses, and that other financial institutions will experience losses as well. These closed-form risk measures can be used for comparative statics, parameter calibration, and setting capital requirements and potential systemic risk charges.

Keywords: Mutually exciting processes; Jumps; Saddlepoint approximation; Expansions in small time; Risk measures; Systemic risk.

*We are very grateful to the Editor, Associate Editor and referee for their thoughtful comments and suggestions that have substantially improved the paper. We are also grateful to seminar and conference participants at Ca' Foscari University of Venice, National Bank of Belgium, Tinbergen Institute, University of Amsterdam, University of Leuven, UNSW Sydney, Weierstrass Institute Berlin, the Financial Risks International Forum on Systemic Risk in Paris, the Measuring Risk Conference in Princeton, the Quantitative Methods in Statistics, Biostatistics and Actuarial Sciences Conference in Louvain-la-Neuve, the VIth International Conference on Applied Mathematics, Modeling and Computational Science in Waterloo, and the 2025 Workshop on Numerical Methods for Finance and Insurance in Milan for their comments and suggestions. This research was funded in part by the Netherlands Organization for Scientific Research (NWO) under grants Vidi-2009 and Vici-2020 (Laeven).

1 Introduction

Large losses at financial institutions tend to occur in close succession to one another, and are rarely confined to a single firm.¹ The reasons for this are multiple: unlike other industries where firms generally benefit from difficulties at a competitor firm, large losses at one bank generate negative externalities for other banks due to interbank loan and other exposure, decline in the value of collateral due to asset sales under duress by the failing bank, direct trading losses resulting from copycat behavior, the resulting disappearance of liquidity, among other factors. Furthermore, once started, the avalanche of losses is both self-feeding, as other firms stop trading with the one in difficulty due to loss of confidence, lines of credit dry up, margin and other collateral demands increase; and contagious, as capital withdraws everywhere when investors' overall risk appetite declines.

Figure 1 illustrates this phenomenon using the daily P&L reported by three global banks. We observe clusters of large losses in time (over days) and space (across banks). Perhaps most noticeable is the cluster that is concentrated in the last three months of 2008, which corresponds to the global financial crisis, but other less salient clusters occur as well.

Although well established empirically, and well understood theoretically, this pattern is nevertheless not fully accounted for by the assumptions embedded in many statistical models used for risk management. Yet, failure to account for these effects can lead to a large underestimation of the overall risk faced by any single bank and by the financial system as a whole, as was made plainly and painfully obvious by the global financial crisis.

Different models and measures to ascertain the contribution of an individual firm to systemic risk, or in reverse the impact of a systemic crisis on an individual firm's risk, have been developed in recent years.² Many existing models are based on continuous processes,

¹See, e.g., Reinhart & Rogoff (2009) and, more recently, Jiang et al. (2023).

²Huang et al. (2012) use Credit Default Swaps of individual banks to measure the price of insuring against systemic financial distress. Conditional autoregressive Value-at-Risk (Engle & Manganelli (2004) and extended to a multivariate setting by White et al. (2015)) models the evolution of the quantile over time using an autoregressive process. Conditional Value-at-Risk (Adrian & Brunnermeier (2016)) measures




Figure 1: Three Banks' Daily Profit and Loss (P&L).

Note: This figure plots the daily profit and loss (P&L) data for three financial institutions: Bank of America (BoA), Deutsche Bank (DB) and Royal Bank of Canada (RBC). The P&L's are measured in Millions of Dollars (BoA), Millions of Euros (DB), and Millions of Canadian Dollars (RBC). The data are obtained by a graph extraction tool, as in Pérignon et al. (2008) and Pérignon & Smith (2010), which reads data printed in the form of P&L plots as published in the annual reports of these banks. The sample period is January 2, 2001 to December 30, 2011. The data are not reported by the banks in table form and the three banks stopped publishing their P&L plots after 2011; no other global bank appears to reveal this information in its annual report at all, while banking regulators do not release the data they collect in their supervisory capacity. So although limited, these data provide a rare window into three global banks' daily P&L, and, most importantly for our purposes, their co-movements.

or make distributional assumptions that are implicitly continuous. Existing jump models typically employ Lévy processes, which have independent increments; they can capture fat tails and cross-sectional dependence, but cannot give rise to the clusters of losses, both across time and firms, that are the hallmark of banking crises.

Yet, risk management is not concerned with capturing the fine day-to-day variations or the VaR of the financial system conditional on firm i being in distress; a bank's contribution to systemic risk can be measured by the effect of an increase in its tail risk on the VaR of the entire system (Hautsch et al. (2015)); see also the Granger-causality based measure of Billio et al. (2012), the tail risk measure of Bienvenüe & Robert (2016), Marginal Expected Shortfall (Acharya et al. (2012), Acharya et al. (2017)), Conditional Expected Shortfall (Mainik & Schaanning (2014)) and SRISK (Brownlees & Engle (2017)).

responding to a single isolated shock whose impact can be absorbed, but with the amplifications via feedback loops taking place both over time and across firms, potentially leading all the way to a systemic failure. So, in this paper, we propose a statistical model designed to produce these two effects together, and use it to derive measures of risk. Different from existing measures of risk, which are primarily derived from cross-sectional dependencies, risk measures computed from our model capture not only jump clustering in the cross-sectional dimension but also jump propagation in the time-series dimension. We start with a statistical P&L model for banks using Hawkes³ self- and cross-exciting jump processes. The time-series and cross-sectional excitation in Hawkes jumps provides a natural match for the self-feeding (over time) and contagious (across firms) nature of banking crises. The Hawkes process is only one part of the statistical P&L model, which also contains a drift term and a continuous Brownian component.

The paper makes two distinct contributions. First, we approximate, in closed form, the tails of the P&L distribution implied by a Hawkes model, both marginally for a single bank and jointly for multiple banks taken together. We employ for this purpose saddlepoint expansions⁴ that we compute fully in closed form for the transition function of the process, from a base that is non-Gaussian, in light of the fat tails implied by the jump model.⁵

³Hawkes processes (Hawkes (1971)) were originally proposed to model epidemics. They have also been used to model earthquake occurrences (Ogata & Akaike (1982) and Brillinger (1988)). More recently, Hawkes processes have been employed in various financial settings, e.g., market microstructure (Bowsher (2007), Bacry et al. (2013), Jaisson & Rosenbaum (2015)), high threshold exceedances (Chavez-Demoulin et al. (2005)), stock market crashes and financial contagion (Aït-Sahalia et al. (2015)), defaults in credit derivatives (Errais et al. (2010) and Azizpour et al. (2018)), systemic credit risk and credit contagion (Aït-Sahalia et al. (2014)), and insurance (Swishchuk et al. (2021)). See also the surveys by Bacry et al. (2015) and Hawkes (2022) and the references therein.

⁴See Daniels (1954), Lugannani & Rice (1980), Daniels (1987), Wood et al. (1993), Reid (1988), Field & Ronchetti (1990). Mykland (1999) used the saddlepoint method to improve large deviation approximations of likelihood ratio statistics based on n observations.

⁵More recent contributions on Gaussian saddlepoint approximations of characteristics of random variables and vectors include Gordy (2002), Martin (2006), Kolassa & Li (2010), Huang & Oosterlee (2011), Kim & Kim (2017), and Broda et al. (2018); Aït-Sahalia & Yu (2006) develop saddlepoint approximations for the transition density of discretely sampled Markov processes, including jump-diffusions with a non-Gaussian base; Glasserman & Kim (2009) consider Gaussian saddlepoint approximations to expectations (prices) of affine jump-diffusion models; comprehensive reviews of existing results are in the textbook treatments by Butler (2007) and Kwok & Zheng (2018).

Compared to the classical use of the saddlepoint method, where one approximates the density of a statistic computed from n observations, we approximate the transition density of a (complex) stochastic process involving Hawkes jumps over a sampling interval Δ . Using operator methods and the infinitesimal Markov generator of the model, we derive closed-form conditional and unconditional cumulants and other characteristics of the model, such as closed-form saddlepoints, which serve as inputs for the saddlepoint expansions.

Next, we use the saddlepoint expansions to construct formulae for risk measures, such as Value-at-Risk and Expected Shortfall, based on the marginal P&L distributions implied by the model, as well as systemic risk measures, exploiting the model's joint P&L distributions. Deriving the tails of banks' joint P&L distributions does not require the assumption of an external copula structure to capture the dependence of potential losses: the dependence in the joint tails is implied by the model. Furthermore, because the model already accounts for the negative externality that one bank's losses imposes onto itself (self-excitation) and onto others (cross-excitation), there is no need to design new risk measures to capture these effects. Finally, because we calculate joint tail approximations in closed form, the remaining numerical effort in computing the risk measures in the model is minimal. Hence, our explicit risk measure formulae are directly amenable to comparative statics analysis, parameter calibration, and capital requirement calculations. For example, one may decompose risk measures computed from our model into a part due to self-excitation and another part due to cross-excitation and assess their effects on potential systemic risk charges.

As an illustration, we apply the model to banks' daily P&L. We compute the saddlepoint approximations for the univariate and multivariate joint tails over the typical short horizon adopted in Value-at-Risk calculations; compute the risk measures; and conduct a comparative statics analysis. We also condition upon the event that one of the banks is in a stress state and assess the impact of this event on the risk measurement of the other bank's P&L. In doing so, we assess the cross-excitation effects from one bank to another.

The rest of this paper is organized as follows: In Section 2, we present our statistical model. Section 3 discusses tail risk measurement. In Section 4, we develop closed-form characterizations of the tails and risk measures of our model, both in the univariate and the bivariate case. Section 5 provides an illustrative example. Section 6 discusses statistical inference. Conclusions are in Section 7. All technical derivations and proofs are in an online companion provided as Supplementary Material: in Section A of the Supplement, we describe operator methods in the presence of mutual excitation; in Section B, we provide the cumulants and other characteristics, such as moment functions, in closed form; Section C contains general expressions of the saddlepoint approximations; Section D provides detailed technical derivations of our saddlepoint approximations and Monte Carlo simulation results; and Section E discusses eigenvector centrality to summarize patterns of mutual excitation.

2 Profit and Loss Model Based on Mutual Excitation

Our statistical model stipulates that a bank’s profits and losses are driven by semimartingale dynamics with a drift term, a continuous Brownian component, and a jump term. In the classical jump-diffusion model in finance, dating back to Merton (1976), jumps are typically assumed to be Poissonian or more generally Lévy processes. By definition, Lévy processes specify independent increments in the time series dimension, so cannot give rise to the clustering of jumps that is apparent in Figure 1. By contrast, we model the jumps using mutually exciting Hawkes processes (Hawkes (1971)), which introduces the possibility that losses at a bank increase the likelihood of future losses both for itself and for other banks.

Jump intensities of a multivariate Hawkes process, just like those in a Cox (or doubly stochastic Poisson) process, are stochastic processes. But whereas in a Cox process the paths of the jump intensities are not affected by the realized paths of the point process they drive, jump intensities in a mutually exciting jump process explicitly depend on the

paths of the associated point process, inducing a feedback effect. Specifically, following the occurrence of a jump in one of the marginal point processes, jump intensities increase (possibly to a different degree) in all point processes, making future jumps more likely.

2.1 Banks' Self- and Cross-Excitation

We fix a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in [0, T]}, \mathbb{P})$ and denote by $N_{i,t}$, $i = 1, \dots, m$, point processes, one for each of the m banks or financial institutions. These processes are defined through their respective intensity processes, $\lambda_{i,t}$, $i = 1, \dots, m$, dictating the \mathcal{F}_t -conditional mean jump rate per unit of time. Jump intensities follow the dynamics

$$\lambda_{i,t} = \lambda_{i,\infty} + \sum_{j=1}^m \int_{-\infty}^t \phi_{i,j}(Z_{j,s}) g_{i,j}(t-s) dN_{j,s}, \quad i = 1, \dots, m, \quad (1)$$

so $\lambda_{i,t}$ increases by $\phi_{i,j}(Z_{j,s}) g_{i,j}(t-s)$ after each jump $dN_{j,s}$, $s \in (-\infty, t)$, $i, j = 1, \dots, m$. The jump decay functions $g_{i,j}$ are an essential ingredient of a Hawkes model. To this, we add a *jump amplification* function, $\phi_{i,j}(Z_{j,t})$, so that the magnitude of a jump in firm j 's P&L, $Z_{j,t}$, and not just the fact that a jump occurred, can affect the jump intensity of firm i , $\lambda_{i,t}$.

The practical tractability of the mutually exciting jump process is critically dependent on the parameterization of the intensity process $\lambda_{i,t}$. We consider the simplest case in which

$$g_{i,j}(t-s) = \beta_{i,j} \exp(-\alpha_i(t-s)), \quad s < t, \quad i, j = 1, \dots, m, \quad (2)$$

where $\alpha_i > 0$, $\beta_{i,j} \geq 0$ for all $i, j = 1, \dots, m$ are constant parameters. Under (1)–(2), a jump in one of the marginal point processes induces the jump intensities to ramp up, after which the intensities decay back exponentially, at speed α_i , towards a level $\lambda_{i,\infty}$. The instantaneous impact of a shock in firm j onto the jump intensity of firm i , $\lambda_{i,t}$, is

dictated by the product of the bank-pair-specific jump amplification function, $\phi_{i,j}(Z_{j,t})$, and the bank-pair-specific parameter $\beta_{i,j}$. Under exponential decay (2), the jump intensities feature mean-reverting dynamics since, from (1), we have that

$$d\lambda_{i,t} = \alpha_i (\lambda_{i,\infty} - \lambda_{i,t}) dt + \sum_{j=1}^m \beta_{i,j} \phi_{i,j}(Z_{j,t}) dN_{j,t}. \quad (3)$$

Importantly, assuming exponential decay makes the couple $(\mathbf{N}_t, \boldsymbol{\Lambda}_t)$ consisting of the vectors of point and intensity processes a Markov process.

The jump part of the model generates the desired features of jump clustering and propagation in time (from past losses) and in space (from one bank's losses to another bank's). The propagation is probabilistic: jump propagation does not occur with probability one nor instantaneously. Rather, the likelihood of observing further jumps in the near future following an initial shock increases under mutual excitation. The model allows for asymmetric patterns of jump clustering and jump propagation since the matrix $\boldsymbol{\beta} = [\beta_{i,j}]_{i,j=1,\dots,m}$ need not be symmetric. We illustrate the jump component of the P&L model by plotting in Figure 2 the filtered jump intensities for three global banks, exploiting the P&L data displayed in Figure 1.

2.2 Excitrality and Systemically Important Financial Institutions

The model may be applied to identify systemically important financial institutions (SIFIs), i.e., which bank(s) pose(s) the largest systemic threat to the financial system. Identifying which banks are SIFIs constitutes a key step towards setting adequate capital requirements. Bank size measured by assets is an important determinant, but due to the patterns of loss transmission some banks may play a key role in the propagation of losses to the system beyond what their size alone would suggest. Our model is designed to capture this risk and identify these banks.

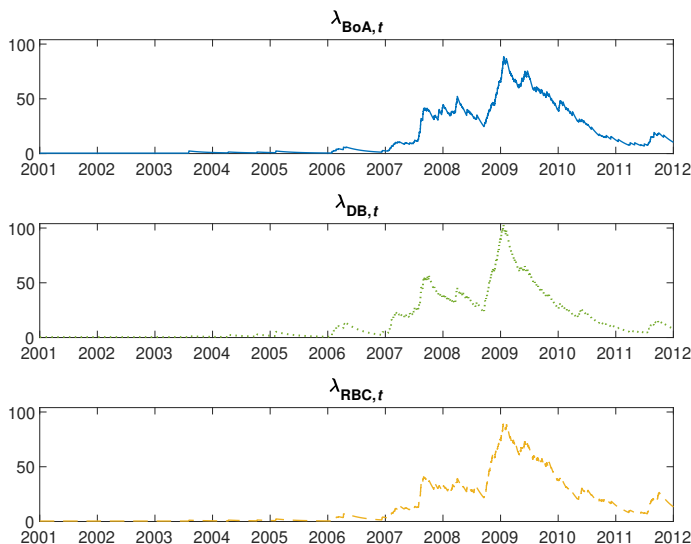


Figure 2: Illustration of the Jump Intensities Driving Three Banks’ Daily P&L.

Note: This figure plots the jump intensities for three global banks induced by the mutually exciting jump model and filtered from the P&L data displayed in Figure 1. We use the 2.5% and 97.5% quantiles of the P&L series as thresholds to decide whether a jump has occurred and next use Eqn. (3) to determine the filtered jump intensity values for the three banks. As the graph extraction tool used to obtain the data in Figure 1 inevitably has its limits regarding the accuracy of the extracted series, we use the P&L data only for illustration purposes; in particular, we use calibrated, equal self-excitation and equal cross-excitation parameters, as the P&L data are too coarse for parameter estimation. The filtered intensities clearly exhibit designated periods with clusters of jumps in time and space.

The parameters $\beta_{i,i}$ and $\beta_{i,j}$, $i, j = 1, \dots, m$, $i \neq j$, determine the extent of self- and cross-excitation of losses.⁶ Jointly with the other parameters of the jump component of our model, they capture key information about the systemic importance of the banks under scrutiny and their susceptibility to systemic risk.

The roles of the individual parameters in our model admit an easy interpretation. A larger $\beta_{i,i}$ implies a more pronounced feedback effect of a bank’s losses onto itself; a larger $\beta_{i,j}$ means the jump intensity of bank i is more affected by bank j ’s losses. A smaller α_i leads to a more persistent effect of losses in any bank onto the jump intensity of bank i . Finally, the larger λ_i , the larger the average level of bank i ’s jump intensity.

To summarize the roles of the parameters — the matrix β and the vectors α and Λ (with

⁶In the nomenclature of graphs and networks, the mutual excitation matrix β constitutes a “weighted adjacency matrix of a directed graph with self-loops”; see Section E of the Supplement.

components $\lambda_i := \mathbb{E}[\lambda_{i,t}]$ — into a single indicator of systemic importance per bank, we consider an eigendecomposition, factorizing a suitable matrix of the model parameters into its eigenvalues and eigenvectors. The directions associated with the dominant eigenvalue, as represented by the associated eigenvector, provide an informative summary of the excitation patterns among the banks that generate systemic risk. After normalization such that the elements sum to unity, the resulting indicator of systemic importance will be referred to as a measure of *excitrality*. It can be viewed as a measure of eigenvector centrality, i.e., excitation-eigenvector centrality, suitably adapted to our P&L model.⁷

Formally, we define the excitrality $\bar{\mathbf{e}}$ on the basis of the matrix $\mathcal{E} := \zeta \mathbf{\Gamma} \text{diag}(\mathbf{\Lambda})$, with ζ and $\mathbf{\Gamma}$ defined in Section A.1 and $\text{diag}(\mathbf{\Lambda}) = \sum_{i=1}^n \mathbf{U}_i \mathbf{\Lambda} \mathbf{u}_i$, where \mathbf{U}_i is an $m \times m$ -matrix with (i, i) -th element equal to one and all other elements equal to zero, and where \mathbf{u}_i is an $1 \times m$ -vector with $(1, i)$ -th element equal to one and all other elements equal to zero. This is a natural choice in view of the fact that $\mathbf{\Lambda} = \mathbf{\Lambda}_\infty + \zeta \mathbf{\Gamma} \mathbf{\Lambda}$. Under exponential decay (2), this matrix takes the simple form

$$\mathcal{E} = \begin{pmatrix} \frac{\lambda_1 \beta_{1,1}}{\alpha_1} & \dots & \frac{\lambda_m \beta_{1,m}}{\alpha_1} \\ \vdots & \ddots & \vdots \\ \frac{\lambda_1 \beta_{m,1}}{\alpha_m} & \dots & \frac{\lambda_m \beta_{m,m}}{\alpha_m} \end{pmatrix}. \quad (4)$$

To compute $\bar{\mathbf{e}}$ we first determine the non-normalized excitrality \mathbf{e} as the eigenvector associated with the dominant eigenvalue ℓ_{\max} of \mathcal{E}' : $\ell_{\max} \mathbf{e} = \mathcal{E}' \mathbf{e}$. (We assume $\mathcal{E}, \mathcal{E}'$ to be irreducible.)⁸ Next, we normalize the elements of \mathbf{e} such that they aggregate to one to obtain our excitrality indicator $\bar{\mathbf{e}}$. To illustrate, we compute in Table 1 the excitralities for several sets of model parameters.

⁷In the literature on graphs and networks, eigenvector centrality constitutes an important measure of centrality; see Section E of the Supplement.

⁸The dominant eigenvector of the “backward operator” \mathcal{E} would provide an indicator of “susceptibility”.

		\bar{e}_i	$\beta_{i,1}$	$\beta_{i,2}$	α_i	λ_i
<i>(i)</i>	Bank 1	0.50	0.8	0.2	3	2
	Bank 2	0.50	0.2	0.8	3	2
<i>(ii)</i>	Bank 1	0.67	0.8	0.2	3	2
	Bank 2	0.33	0.4	0.6	3	2
<i>(iii)</i>	Bank 1	0.77	0.8	0.2	3	2
	Bank 2	0.23	1.6	0.6	3	2
<i>(iv)</i>	Bank 1	0.84	1.6	0.2	3	2
	Bank 2	0.16	0.4	0.6	3	2
<i>(v)</i>	Bank 1	0.63	0.8	0.2	3	2
	Bank 2	0.37	0.2	0.8	2	2
<i>(vi)</i>	Bank 1	0.82	0.8	0.2	3	2
	Bank 2	0.18	0.2	0.8	3	1

Table 1: Excitrality and Systemically Important Banks

Note: This table displays the excitrality vector \bar{e} for two banks under different sets of model parameters. Quite naturally, when the parameters are symmetric among the two banks, both banks are equally systemically important (panel *(i)*). When the self- and cross-excitation parameters of bank 1 are larger than that of bank 2, bank 1 becomes more systemically important than bank 2, resulting in an increase in the first element and a corresponding decrease in the second element of the excitrality vector (see *(ii)*). A further increase in the cross-excitation effect from the first bank to the second bank makes the first bank more systemically important (see *(iii)*). An increase in the self-excitation parameter $\beta_{1,1}$ has an even more pronounced effect on the measure of excitrality (see *(iv)*). If shocks affecting bank 1 are less persistent than shocks affecting bank 2, bank 1 is more systemically important (see *(v)*). Finally, a decrease in the average level of bank 2's jump intensity makes that bank less systemically important (see *(vi)*).

2.3 Complete Model

In addition to the mutually exciting jump component, our full statistical model of financial P&L contains a drift term and a continuous Brownian component designed to capture the normal fluctuations. Let $Y_{i,t}$ denote the market value of the trading book of financial institution i and let $dY_{i,t}$ denote its continuously measured P&L. We assume it is driven

by the following dynamics:⁹

$$dY_{i,t} = \mu_i dt + \sigma_i dW_{i,t} + Z_{i,t} dN_{i,t}, \quad i = 1, \dots, m. \quad (5)$$

Here, $\mu_i \in \mathbb{R}$ is the drift; $\sigma_i > 0$ is the volatility; $\mathbf{W}_t := (W_{1,t}, \dots, W_{m,t})'$ is an m -dimensional standard Brownian motion with linear correlation coefficients $\rho_{i,j}$, $-1 \leq \rho_{i,j} \leq 1$, $i, j = 1, \dots, m$;¹⁰ $\mathbf{Z}_t := (Z_{1,t}, \dots, Z_{m,t})'$ is an m -vector of jump magnitudes with cumulative distribution function (cdf) F_{Z_i} supported on $(-\infty, \infty)$,¹¹ and assumed to be cross-sectionally and serially independent; and $\mathbf{N}_t := (N_{1,t}, \dots, N_{m,t})'$ is the vector of mutually exciting processes described previously. The vectors of Brownian motions \mathbf{W} on the one hand and jump magnitudes \mathbf{Z} and jump processes \mathbf{N} on the other hand are assumed to be mutually independent.

So far we have left the jump amplification function, $\phi_{i,j}(\cdot)$, and the jump magnitude distribution F_{Z_j} unspecified.¹² In light of the empirical evidence, a realistic model should allow for asymmetries in the jump magnitude distribution, so that the losses (i.e., negative

⁹The drift μ_i and volatility σ_i in the base model specification (5) are constant parameters. In Section A of the Supplement, Eqns. (A.6)–(A.7), we extend (5) with stochastic volatility, using an instantaneous variance process of the Heston type to allow for a leverage effect. In Section B, we provide the explicit expression for the corresponding cgf $\bar{K}(\Delta, u)$. It is visible from the expression that the leverage parameter ρ^V enters only at the higher order Δ^2 where, further, it is associated with u^3 . As a result, the leverage parameter ρ^V will not enter in the tail saddlepoint approximation at the leading order. As far as the volatility side is concerned, the tail saddlepoint approximation only features the stochastic volatility *level*. This is not surprising since the impact of the volatility component far in the tails is outweighed by that of the jump component. While the volatility level contributes to making the left tail of the distribution larger at the leading order, its effect is overall small compared to the contribution of jumps to that tail.

¹⁰The $m \times m$ -dimensional variance-covariance matrix Σ with elements $\Sigma_{i,j} = \rho_{i,j}\sigma_i\sigma_j$ is a symmetric positive semi-definite matrix and assumed to be non-singular.

¹¹One may replace Z_i by $\log(1 + \tilde{Z}_i)$. Then, the jump component in the differential of $\exp(Y_{i,t})$ becomes $\tilde{Z}_i \exp(Y_{i,t}) dN_{i,t}$. For ease of exposition, we write Z_i rather than $\log(1 + \tilde{Z}_i)$.

¹²In principle, our model and the associated analytical tools that we develop do not require specific assumptions about the jump amplification function and the jump magnitude distribution at this stage. In particular, the general results in the Supplement will be provided for a generic jump amplification function as functions of generic moments of $\phi_{i,j}(Z_j)$ and cross-products with Z_j . However, when it comes to computing explicit saddlepoint approximations, we will parameterize $\phi_{i,j}(\cdot)$ and the (moments of the) jump magnitude distribution. When doing so, we take into account the following requirements: (0) normalization: $\phi_{i,j}(0) = 0$; (i) positivity: $\phi_{i,j}(z) \geq 0$ for all $|z| \geq 0$; (ii) stationarity: achieved by restricting $\mathbb{E}[\phi_{i,j}(Z_{j,s})] = \zeta$, $i, j = 1, \dots, m$, where we normalize $\zeta \equiv 1$; (iii) non-decreasingness of the amplification: $\phi'_{i,j}(z) \geq 0$, $z > 0$, and $\phi'_{i,j}(z) \leq 0$, $z < 0$; and (iv) tractability.




Figure 3: Jump Amplification Function, $\phi_j(z)$.

Note: This figure plots an example of the jump amplification function, $\phi_j(z)$, as a function of the jump size z . The jump amplification function in this example is asymmetric, making the impact of the occurrence of negative jumps (losses) on the jump intensity more pronounced than the impact of positive jumps (gains). The parameter values are $c_{j,-} = c_{j,+}$ so $\chi_j = 1$ (see (A.13)), and $\xi_{j,-} = 0.05$, $\xi_{j,+} = 0.01$, $p_j = 0.3$, $1/\gamma_{j,-} = 100$, $1/\gamma_{j,+} = 50$.

jumps) lead to more excitation than profits (i.e., positive jumps) and that rare large losses are more likely than rare large profits. We consider the following specification (suppressing in the remainder of this subsection the subscript i for ease of notation):

$$\phi_j(z) = \begin{cases} c_{j,-} (1 - \exp(-\xi_{j,-}(-z))), & -\infty < z \leq 0; \\ c_{j,+} (1 - \exp(-\xi_{j,+}z)), & 0 < z < \infty; \end{cases} \quad (6)$$

with $c_{j,-}, c_{j,+}, \xi_{j,-}, \xi_{j,+} > 0$ denoting constant parameters. Furthermore, we will assume that Z_j has cdf given by

$$F_{Z_j}(x) = \begin{cases} p_j \exp(-\gamma_{j,-}x), & -\infty < x \leq 0; \\ p_j + (1 - p_j) (1 - \exp(-\gamma_{j,+}x)), & 0 < x < \infty; \end{cases} \quad (7)$$

with $\gamma_{j,-}, \gamma_{j,+} > 0$ and $0 \leq p_j \leq 1$, $j = 1, \dots, m$, denoting constant parameters. The jump amplification function $\phi_j(z)$ is illustrated in Figure 3.

3 Risk Measures

3.1 Univariate Risk Measures

A main application of the model consists in computing various standard risk measures¹³ such as Value-at-Risk (VaR), Expected Shortfall (ES), etc., which all depend on the left tail of a bank’s P&L distribution, in the context of the model described in Section 2. The key aspect of our model is that it incorporates the negative externality that one bank’s losses impose onto itself as well as other banks in the computation of a risk measure.

To this end, we need to evaluate the model-implied left tails of the probability density and cumulative distribution functions of a bank’s P&L, of which all law-invariant univariate and marginal risk measures are functionals. As we will see, the model of Section 2 leads to substantially larger tails, hence risk measures and capital requirements, than traditional alternatives that do not incorporate self- and mutual excitation of losses.¹⁴

Henceforth, let $X_{i,t,\Delta} := -(Y_{i,t+\Delta} - Y_{i,t})$ denote bank i ’s (sign-changed) P&L over a time interval of length Δ . We begin by considering the univariate risk of a single bank, hence we omit the index i in this subsection. We denote the cdf $F_{X_{t,\Delta}}$ by $F_{X_{t,\Delta}}(x) = \mathbb{P}[X_{t,\Delta} \leq x] \equiv P(\Delta, x)$. Furthermore, we let $p(\Delta, x)$ be the probability density function (pdf) associated with $P(\Delta, x)$ and let $\bar{P}(\Delta, x)$ be the decumulative distribution (i.e., survival) function (ddf). Next, we denote by $\text{VaR}(\Delta, p)$ the Value-at-Risk of the P&L measured over the

¹³See, e.g., Föllmer & Schied (2011), Chapter 4, and Laeven & Stajic (2013) for a detailed description of VaR and ES, classes of risk measures that do and do not encompass them, and their decision-theoretic foundations. Different tail probabilities translate, for example, into different VaR values and consequently different capital requirements, which are typically set as fixed multiples of VaR; see, e.g., the standard on minimum capital requirements for market risk, as published by the Basel Committee on Banking Supervision in 2019, and integrated into the consolidated Basel framework.

¹⁴The additional level of capital required to guard against mutual excitation risk can be useful information for banking regulators as they assess capital requirements needed in a more realistic risk scenario. Furthermore, risk management standards have evolved in response to the financial crisis of 2007–09, including the requirement that banks compute “Stressed-Value-at-Risk” (SVaR) as their usual VaR, but estimated using historical data from a continuous 12-month period of significant financial stress. Risk measures computed within our model are designed to incorporate the likelihood of financial stress as a natural rather than unexpected event: no separate estimation period is required.

time interval Δ and evaluated at probability level p . It is the p -quantile of the sign-changed P&L's cdf defined as

$$\text{VaR}(\Delta, p) = F_{X_{t,\Delta}}^{-1}(p) = \inf \{x \in \mathbb{R} \mid P(\Delta, x) \geq p\}, \quad p \in (0, 1), \quad (8)$$

(with the convention that $\inf\{\emptyset\} = \infty$) so VaR follows from inversion of the cdf $P(\Delta, x)$. VaR has been the standard measure of risk in the banking industry since the 1996 amendment to the Basel Accord of the Basel Committee on Banking Supervision.¹⁵ VaR is adequate when shortfall is to be controlled at a certain probability level, but does not capture the magnitude of the loss beyond the fact that it has exceeded the threshold.

An alternative to VaR that measures the extent of the loss is given by the expected shortfall, ES,¹⁶ defined as

$$\text{ES}(\Delta, p) = \frac{1}{1-p} \int_p^1 \text{VaR}(\Delta, q) \, dq, \quad (9)$$

for $p \in (0, 1)$. For a given probability level p , ES exceeds VaR; indeed, if the cdf is continuous, then $\text{ES}(\Delta, p) = \mathbb{E}[X_{t,\Delta} \mid X_{t,\Delta} > \text{VaR}(\Delta, p)]$, and $\mathbb{E}[X_{t,\Delta} \mid X_{t,\Delta} > x] = x + \mathbb{E}[(X_{t,\Delta} - x)^+] / \mathbb{P}[X_{t,\Delta} > x] = \mathbb{E}[X_{t,\Delta} \mathbb{1}_{\{X_{t,\Delta} > x\}}] / \bar{P}(\Delta, x)$. Henceforth, we denote

$$\mathcal{L}(\Delta, x) \equiv \mathbb{E}[X_{t,\Delta} \mathbb{1}_{\{X_{t,\Delta} > x\}}] \quad (10)$$

and will provide below saddlepoint approximations to $p(\Delta, x)$, $\bar{P}(\Delta, x)$ and $\mathcal{L}(\Delta, x)$.

¹⁵It is also used in the insurance, pension and other industries, at a variety of time horizons Δ (such as 1 or 10 days) and probability levels (such as 95% and 99%), to measure and regulate a wide variety of risks.

¹⁶ES is also known as Average-Value-at-Risk and, under slightly different definitions that agree for continuous distributions, as Conditional Value-at-Risk and Conditional Tail Expectation. It accounts for the magnitude of the loss since $\text{ES}(\Delta, p) = \text{VaR}(\Delta, p) + \frac{1}{1-p} \mathbb{E}[(X_{t,\Delta} - \text{VaR}(\Delta, p))^+]$.

3.2 Measures of Joint and Systemic Risk

Of greater concern to regulators than the isolated failure of a single bank is the joint risk of two (or more) banks getting in trouble together, or in close succession. Our model provides a parsimonious way of quantifying this risk. Self-excitation captures the negative externality that a bank's losses impose onto itself (measured by $\beta_{i,i}$), whereas cross-excitation captures the negative externality from one bank to another, as its own losses make it more likely that other banks will experience a loss as well (measured by $\beta_{i,j}$, $i \neq j$). As we will see, the self- and cross-excitation parameters have asymmetric effects on the tails of the joint P&L distribution.¹⁷

Let $X_{1,t,\Delta}$ and $X_{2,t,\Delta}$ denote the sign-changed P&L's of two banks over a time interval of length Δ . Their joint cdf and associated pdf are denoted by $P(\Delta, x_1, x_2) = \mathbb{P}[X_{1,t,\Delta} \leq x_1, X_{2,t,\Delta} \leq x_2]$ and $p(\Delta, x_1, x_2)$, respectively. We also denote by $p(\Delta, x_2 | x_1) = p(\Delta, x_1, x_2)/p(\Delta, x_1)$ the conditional pdf of $X_{2,t,\Delta}$ at x_2 conditionally upon $X_{1,t,\Delta} = x_1$. We finally denote by $P(\Delta, x_2 | x_1) = \mathbb{P}[X_{2,t,\Delta} \leq x_2 | X_{1,t,\Delta} = x_1]$ the corresponding conditional cdf and by $\bar{P}(\Delta, x_2 | x_1)$ the conditional ddf.

All law-invariant joint and systemic risk measures are functionals of these probability distributions, for which we will provide explicit approximations. Imagine, for example, a VaR calculation by a regulator monitoring two banks, and concerned with the joint probability of meeting individual loss thresholds L_1 and L_2 , i.e., $\mathbb{P}[X_{1,t,\Delta} \leq L_1, X_{2,t,\Delta} \leq L_2]$. The dependence is generated by the model without the need for an external assumption concerning joint losses in the form of a copula. As in the univariate model, for sufficiently large loss thresholds, the result is driven by the jump component of the model. Furthermore,

¹⁷For an individual bank to apply the model for its own risk calculations, such bank would need to estimate the model's parameters, including those that capture the impact that other banks' losses have on itself. An individual bank is unlikely to have access to data on P&L at its competitors, which only the regulators possess. So banking regulators would most likely have to specify what parameters $\beta_{i,j}$ are to be used by a given bank. This is not different from the current practice, where the regulators validate the internal models employed by individual banks for their risk calculations. Setting the values of these parameters by the regulators does not reveal proprietary P&L data.

the *conditional* VaR,

$$\inf \{x_2 \in \mathbb{R} \mid P(\Delta, x_2 \mid x_1) \geq p\}, \quad p \in (0, 1), \quad (11)$$

measures the risk of bank 2 given that bank 1 is experiencing losses at level x_1 over the same time interval of length Δ . It is the (generalized) inverse of $P(\Delta, x_2 \mid x_1)$, just like in the univariate and unconditional marginal cases. The difference between conditional risk measures and their unconditional marginal counterparts, for example with VaR, $\inf \{x_2 \in \mathbb{R} \mid P(\Delta, x_2 \mid x_1) \geq p\} - \inf \{x_2 \in \mathbb{R} \mid P(\Delta, x_2) \geq p\}$, can be used as a measure of the contribution of other banks' losses or systemic risk to a given bank's risk. It measures the extent to which losses in another bank, or the system as a whole, incrementally impact the risk of a given bank. Similarly, the difference between the joint tail distribution of losses and the product of the marginal tail distribution of losses captures the extent to which mutual excitation induces dependence between the occurrence of the two losses forcing a departure from the temporal and spatial independence of individual losses.

4 Model-Implied Tail Saddlepoint Expansions

As the previous section made clear, every risk measure, marginal or joint, requires an evaluation of the (joint) left tail of the P&L distribution. Analytic computations for continuous-time models with jumps¹⁸ are already complex in the baseline Lévy case, and generally not available in closed form; the Hawkes structure of our model adds yet another degree of complexity. To solve this problem, we proceed by deriving the relevant cumulants of the model in closed form as expansions in Δ , and then use these cumulants to construct fully explicit saddlepoint approximations; since the model involves jumps, we start the

¹⁸With realistic stochastic volatility parameters, a no-jumps, diffusion-only model cannot generate a tail large enough to be consistent with the observed evidence on P&L. So incorporating a jump component into the model is necessary.

saddlepoint approximations at a Bernoulli rather than a Gaussian base.

More specifically, our approach can be succinctly summarized as follows.¹⁹ We start by analyzing the cumulant generating function (cgf) of the bivariate random vector $\mathbf{X}_{t,\Delta} = (X_{1,t,\Delta}, X_{2,t,\Delta})$ representing bank 1 and bank 2's P&L, defined as

$$K(\Delta, u_1, u_2) = \log \mathbb{E} [\exp (u_1 X_{1,t,\Delta} + u_2 X_{2,t,\Delta})], \quad \mathbf{u} = (u_1, u_2) \in \mathbb{R}^2. \quad (12)$$

Once we have derived the cgf, the pdf $p(\Delta, x_1, x_2)$ of $\mathbf{X}_{t,\Delta}$ follows by Fourier inversion:²⁰

$$p(\Delta, x_1, x_2) = (2\pi)^{-2} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \exp (K(\Delta, iu_1, iu_2) - (iu_1 x_1 + iu_2 x_2)) du_1 du_2,$$

which, in turn, underlies the saddlepoint approximation, by suitably choosing the path of integration.

In our model, banks' P&L depend upon latent state variables $\boldsymbol{\varsigma}_t$.²¹ Consider a function of the form $\psi(\Delta, \mathbf{y}_1, \mathbf{y}_0, \boldsymbol{\varsigma}_1, \boldsymbol{\varsigma}_0)$, where the subscripts 1 and 0 refer to two dates separated by a time interval Δ . Recall that $Y_{i,t}$ denotes bank i 's market value of the trading book at time t , and the bank's P&L $X_{i,t,\Delta}$ is the sign-changed increment of $Y_{i,t}$ over the time interval Δ . Deriving the cgf requires evaluating $\mathbb{E} [\psi(\Delta, \mathbf{Y}_1, \mathbf{Y}_0, \boldsymbol{\varsigma}_1, \boldsymbol{\varsigma}_0)]$ for a specific choice of ψ , namely $\psi(\Delta, \mathbf{y}_1, \mathbf{y}_0, \boldsymbol{\varsigma}_1, \boldsymbol{\varsigma}_0) = \exp (\mathbf{u}'(\mathbf{y}_1 - \mathbf{y}_0))$ for a fixed value of \mathbf{u} . To do so, we use the explicit expression of the infinitesimal Markov generator \mathcal{A} of the model of Section 2, which allows us to expand the conditional expectation of ψ as a Taylor series up to order

¹⁹Full details of the derivations are contained in Sections A and B of the Supplement, where we discuss operator methods and provide cumulants for the model, and in Sections C and D of the Supplement, where we provide the general and explicit expressions of the saddlepoint approximations.

²⁰Hence,

$$p(\Delta, x_1, x_2) = (2\pi i)^{-2} \int_{\hat{u}_1 - i\infty}^{\hat{u}_1 + i\infty} \int_{\hat{u}_2 - i\infty}^{\hat{u}_2 + i\infty} \exp (K(\Delta, u_1, u_2) - (u_1 x_1 + u_2 x_2)) du_1 du_2.$$

A suitable choice of the vector (\hat{u}_1, \hat{u}_2) dictating the path of integration is the key to the saddlepoint approximation.

²¹Indeed, $\boldsymbol{\varsigma}_t = \boldsymbol{\Lambda}_t$, and in a with stochastic volatility extended version of our model, $\boldsymbol{\varsigma}_t = (\mathbf{V}_t, \boldsymbol{\Lambda}_t)$, where \mathbf{V}_t is the P&L's stochastic variance and $\boldsymbol{\Lambda}_t$ its stochastic jump intensity.

J in Δ as follows:

$$\begin{aligned} \mathbb{E}_{\mathbf{Y}_1, \boldsymbol{\varsigma}_1} [\psi(\Delta, \mathbf{Y}_1, \mathbf{Y}_0, \boldsymbol{\varsigma}_1, \boldsymbol{\varsigma}_0) | \mathbf{Y}_0, \boldsymbol{\varsigma}_0] &= \sum_{j=0}^J \frac{\Delta^j}{j!} (\mathcal{A}^j \cdot \psi) (0, \mathbf{Y}_0, \mathbf{Y}_0, \boldsymbol{\varsigma}_0, \boldsymbol{\varsigma}_0) \\ &+ O_p(\Delta^{J+1}), \end{aligned} \quad (13)$$

where subscripts in $\mathbb{E}_{\mathbf{Y}_1, \boldsymbol{\varsigma}_1}$ indicate the random variables that the expected value operates on, and $\mathcal{A}^j \cdot \psi$ is defined recursively by $\mathcal{A}^j \cdot \psi = \mathcal{A} \cdot (\mathcal{A}^{j-1} \cdot \psi)$ for all $j \geq 1$. The iterates $\mathcal{A}^j \cdot \psi$, hence all the terms in (13), are evaluated in closed form. Thus, we can obtain the conditional expectation of ψ , using the full state vector including its latent components. To evaluate (12), we need to integrate out the latent state variables:

$$\begin{aligned} \mathbb{E} [\psi(\Delta, \mathbf{Y}_1, \mathbf{Y}_0, \boldsymbol{\varsigma}_1, \boldsymbol{\varsigma}_0)] &= \mathbb{E}_{\mathbf{Y}_1, \mathbf{Y}_0, \boldsymbol{\varsigma}_1, \boldsymbol{\varsigma}_0} [\mathbb{E}_{\mathbf{Y}_1, \boldsymbol{\varsigma}_1} [\psi(\Delta, \mathbf{Y}_1, \mathbf{Y}_0, \boldsymbol{\varsigma}_1, \boldsymbol{\varsigma}_0) | \mathbf{Y}_0, \boldsymbol{\varsigma}_0]] \\ &= \sum_{j=0}^J \frac{\Delta^j}{j!} \mathbb{E}_{\mathbf{Y}_0, \boldsymbol{\varsigma}_0} [(\mathcal{A}^j \cdot \psi) (0, \mathbf{Y}_0, \mathbf{Y}_0, \boldsymbol{\varsigma}_0, \boldsymbol{\varsigma}_0)] + O(\Delta^{J+1}), \end{aligned}$$

and the last step in the necessary calculations involves computing unconditional expectations with respect to the stationary law of the state variables $(\mathbf{Y}_0, \boldsymbol{\varsigma}_0)$.

With explicit expressions for the cgf at hand, we determine the saddlepoints $(\hat{u}_1, \hat{u}_2) = (\hat{u}_1(x_1, x_2), \hat{u}_2(x_1, x_2))$ as the solution to the system of equations $\partial K(\Delta, u_1, u_2) / \partial u_1 = x_1$, $\partial K(\Delta, u_1, u_2) / \partial u_2 = x_2$.²² We also determine $\hat{u}_0 = \hat{u}_0(x_1)$ and $\hat{u}_{u_2} = \hat{u}_{u_2}(x_1)$ as the solutions to the equations $\partial K(\Delta, u_1, 0) / \partial u_1 = x_1$ and $\partial K(\Delta, u_1, u_2) / \partial u_1 = x_1$, respectively. Similarly, for a suitably selected (i.c. Bernoulli) base cgf $K_0(\Delta, w)$, we compute the associated saddlepoint $\hat{w} = \hat{w}(x_2)$ as the solution to the equation $\partial K_0(\Delta, w) / \partial w = x_2$. We compute fully explicit expressions for these saddlepoints, and their reciprocals required for the saddlepoint approximation to the conditional ddf, using Taylor series expansions in Δ .

²²Note that $K(\Delta, u_1, u_2)$ is convex in both variables ensuring the existence of a solution to this set of two equations.

Upon collecting terms, this leads to expansions in Δ of the following saddlepoint approximation for the conditional pdf (the marginal pdf and ddf, conditional ddf and joint pdf can be handled analogously, see the Supplement):

$$p^{(0)}(\Delta, x_2|x_1) = f_0(\Delta, x_2)e^{(K(\Delta, \hat{u}_2, \hat{u}_2) - \hat{u}_2 x_1) - (K(\Delta, \hat{u}_0, 0) - \hat{u}_0 x_1) - \hat{u}_2 x_2 - (K_0(\Delta, \hat{w}) - \hat{w} x_2)} \\ \times \left(\frac{\partial^2 K_0(\Delta, \hat{w})}{\partial w^2} \right)^{1/2} \left(\frac{\partial^2 K(\Delta, \hat{u}_0, 0)}{\partial u_1^2} \right)^{1/2} (|K''(\Delta, \hat{u}_1, \hat{u}_2)|)^{-1/2}, \quad (14)$$

where f_0 is the base pdf, K'' denotes the Hessian of K and $|\cdot|$ denotes the determinant.²³

We finally use these expressions to obtain explicit expressions for univariate risk measures such as VaR, ES and measures of joint and systemic risk as described in Section 3.

4.1 Tail Saddlepoint Expansion: Univariate Case

We start with the univariate case. For ease of exposition, we focus here on the pure jump case, in which $\mu = \sigma = 0$:^{24,25}

Theorem 1. *The saddlepoint approximation to the ddf of the Hawkes jump model (3), (5)–(7) with $m = 1$ and $p = 1$ is given at order Δ^2 by the following expression:*

$$\begin{aligned} \bar{P}^{(0)}(\Delta, x) &= \exp\left(\frac{\gamma x}{2} \left(\lambda + \frac{\beta q_1(\alpha q_2 - \beta q_1)}{\xi q_2(\alpha - \beta)}\right) \Delta\right) \\ &\times \left(\lambda \Delta \exp(-\gamma x) + \lambda \Delta \gamma \exp(-\gamma x)\right. \\ &\times (2q_2 \lambda(\alpha - \beta) (2\gamma \xi^2(\alpha(3\beta + \lambda) - \beta(2\beta + \lambda)) + 2\xi^3(2\alpha(\beta + \lambda) - \beta(\beta + 2\lambda)) \\ &\quad \left. - \gamma^3 \beta(\alpha - \beta)) + (3q_1^2 \beta^2 (q_2 \alpha - q_1 \beta)^2 + \lambda^2 \xi^2 q_2^2 (\alpha - \beta)^2) \gamma x\right) \\ &\times \left. \frac{1}{16\gamma \lambda \xi^2 q_2^2 (\alpha - \beta)^2} \Delta\right) + o(\Delta^2), \end{aligned} \quad (15)$$

²³This approximation may be optimized further with an alternative, judicious choice of x_2 in (14) so that the exponential term in the approximation vanishes.

²⁴The general results are available from the authors upon request.

²⁵When $p = \{0, 1\}$ in (7), we simply write γ and ξ instead of $\{\gamma_+, \gamma_-\}$ and $\{\xi_+, \xi_-\}$.

with $q_1 = \gamma + \xi$ and $q_2 = \gamma + 2\xi$.

We next consider saddlepoint approximations to ES:

Corollary 1. *The saddlepoint approximation to $\mathcal{L}(\Delta, x)$ defined in (10) in the univariate Hawkes jump model of Theorem 1 takes the following form:*

$$\mathcal{L}^{(0)}(\Delta, x) = \frac{\lambda}{\gamma} \Delta \bar{P}^{(0)}(\Delta, x) + \frac{x - (\lambda/\gamma)\Delta}{\hat{u}} p^{(0)}(\Delta, x),$$

with $\bar{P}^{(0)}(\Delta, x)$, \hat{u} and $p^{(0)}(\Delta, x)$ the saddlepoint approximation to the ddf, the saddlepoint, and the saddlepoint approximation to the pdf, given by Theorem 1, Equation (D.9) and Lemma S.3, respectively. The error term is of order $o(\Delta^2)$.

The proofs of Theorem 1 and Corollary 1 are given in the Supplement. The counterparts of Theorem 1 and Corollary 1 for the marginal tails in the bivariate Hawkes jump model are contained in Theorem S.11 and Corollary S.2 in the Supplement.

The saddlepoint approximations above are fully explicit. This is important, among other reasons, for calibration or estimation purposes of the risk measures, and to reveal the roles played by the various model parameters. For example, the derivative (i.e., sensitivity) of the saddlepoint approximation to the ddf given in Theorem 1, taken with respect to the self-excitation parameter β and evaluated at $\beta \equiv 0$, is given by the following expression at the leading order in Δ :

$$\left. \frac{\partial \bar{P}^{(0)}(\Delta, x)}{\partial \beta} \right|_{\beta=0} = \exp(-\gamma x) \frac{\lambda_\infty}{\alpha} \Delta + O(\Delta^2);$$

so, at the leading order, an increase in the degree of self-excitation leads to a fatter tail for the bank's P&L distribution.²⁶

²⁶In this comparative statics calculation, we consider λ_∞ (and not λ) to be fixed and given. Then we obtain a first-order (in Δ) effect. If, by contrast, we would consider λ to be fixed (so λ_∞ reduces when β increases to maintain equal expectations), then we would obtain a second-order (in Δ) effect. Note that when λ is fixed, we should not expect a global comparative statics result. In particular, two distributions

In Section D.5 of the Supplement, we analyze the accuracy of the saddlepoint approximations, by comparing them to Monte Carlo simulated estimates. We first specialize the P&L model to the Poissonian case, which admits a closed-form expression for the pdf, and separately compare the tail probabilities obtained from the saddlepoint approximations and from simulations to those following from the closed-form expression. The results in the Poissonian case show that the saddlepoint approximations provide more accurate estimates of the tail probabilities than the Monte Carlo simulations, even with $> 10^7$ simulated increments. Next, we consider the Hawkes case, which does not admit a closed-form expression for the pdf. We observe a suitable match between the saddlepoint approximations and the (noisy) Monte Carlo simulated estimates of the tail probabilities.

We also establish in Lemma S.2 an analog of the relative error property (see, e.g., Field & Ronchetti (1990) in the classical saddlepoint setting in which one approximates the density of a statistic computed from n observations) for the pdf of the Hawkes model.

4.2 Tail Saddlepoint Expansion: Bivariate Case

We now consider two banks together. Bank 2's P&L is then subject not only to the adverse effect of its own recent losses (as captured by $\beta_{2,2}$) but also to the adverse effects of recent losses at bank 1 (as captured by $\beta_{2,1}$), both of which make it more likely in the model that bank 2 will experience further losses. While a risk manager at bank 1 may not be directly concerned about the negative externality his or her losses impose on bank 2, a banking regulator would be concerned about risks to the financial system as a whole and therefore want to account for this cross-excitation.

As in the univariate case, we start by computing closed-form expansions in Δ of the cumulants of the model, and use those to form fully explicit saddlepoint approximations with equal expectations cannot have globally ordered tail probabilities. But for sufficiently large x , we find that the second-order effect is always positive.

from a joint Bernoulli base. For ease of exposition, we restrict attention to the special case in which $\mu_i = \sigma_i = 0$ and $\phi_{i,j}(Z_j) \equiv 1$:

Theorem 2. *The saddlepoint approximation to the conditional ddf of the Hawkes jump model (3), (5)–(7) with $m = 2$ and $p_j = 1$, $j = 1, 2$, is given at order Δ^2 by:*

$$\begin{aligned}
& \bar{P}^{(0)}(\Delta, x_2|x_1) \\
&= \exp\left(\frac{f\sqrt{\gamma_1\lambda_1\gamma_2\lambda_2x_1x_2} + \gamma_1\lambda_1x_2(c + \gamma_2^2\lambda_2^2)}{2\gamma_1\lambda_1\gamma_2\lambda_2}\Delta\right) \\
&\quad \times \left(\lambda_2\Delta \exp(-\gamma_2x_2) + \lambda_2\Delta\gamma_2 \exp(-\gamma_2x_2)\right. \\
&\quad \times \left(\frac{f(\sqrt{\gamma_1\lambda_1x_1} + \sqrt{\gamma_2\lambda_2x_2})}{8\gamma_1\lambda_1\gamma_2^2\lambda_2}\Delta^{1/2} + \left(\frac{6f\sqrt{\gamma_1^5\lambda_1^5\gamma_2\lambda_2x_1x_2}}{16\gamma_1^3\lambda_1^3\gamma_2^3\lambda_2x_2}\right.\right. \\
&\quad \left.\left.+ \frac{x_2(8c\gamma_1^3\lambda_1^3 + \gamma_2^2\lambda_2(-3b^2x_1 + \gamma_1^2\lambda_1^2(-2d + \gamma_1\lambda_1\lambda_2(\gamma_2x_2 + 4))))}{16\gamma_1^3\lambda_1^3\gamma_2^3\lambda_2x_2}\right)\Delta\right) \\
&\quad \left. + o(\Delta^2),\right. \tag{16}
\end{aligned}$$

where

$$\begin{aligned}
b &= \gamma_1^2(\nu_{1,1} + (\beta_{1,1} - \lambda_1)\lambda_1), & c &= \gamma_2^2(\nu_{2,2} + (\beta_{2,2} - \lambda_2)\lambda_2), \\
d &= -\gamma_1(2\nu_{1,1} + 2\nu_{1,2} + \beta_{1,2}\lambda_2 + \lambda_1(2\beta_{1,1} + \beta_{2,1} - 2(\lambda_1 + \lambda_2))), \\
f &= \gamma_1\gamma_2(2\nu_{1,2} + \beta_{2,1}\lambda_1 + (\beta_{1,2} - 2\lambda_1)\lambda_2),
\end{aligned}$$

with $\nu_{i,j}$, $i, j = 1, 2$, defined in Section B.2.

Corollary 2. *We have the following saddlepoint approximation for the bivariate pdf of the bivariate Hawkes jump model of Theorem 2:*

$$p^{(0)}(\Delta, x_1, x_2) = p^{(0)}(\Delta, x_2|x_1)p^{(0)}(\Delta, x_1), \tag{17}$$

with $p^{(0)}(\Delta, x_2|x_1)$ given in Lemma S.5 and $p^{(0)}(\Delta, x_1)$ given in Lemma S.4.

Given the results in Section C.2 of the Supplement, the proofs of Theorem 2 and

Lemma S.5 follow along similar lines as the proofs of Theorem 1 and Lemma S.3 (and the proofs of Theorem S.11 and Lemma S.4).

5 Application to Risk Management

We now combine the ingredients of the previous two sections. To assess the impact of self- and cross-excitation on the marginal and joint risk measures, we contrast three bivariate jump models: a simple Bernoulli jump model in which each financial institution faces at most one jump with cdf F_{Z_i} in an interval of length Δ , occurring with jump probability $\lambda_i\Delta$, and with the institutions being mutually independent; the Poissonian jump model to which our model reduces when all excitation parameters are set to zero, i.e., $\beta_{i,j} \equiv 0$, $i, j = 1, 2$; and our unrestricted mutually exciting jump model. The Bernoulli model serves as the base for the saddlepoint expansions while the Poissonian model is a special case of the Hawkes model, so we can specialize the results of Section 4 to cover all three models.

The key difference between the mutually exciting model and the Bernoulli and Poissonian models is that in the latter two models the fact that the first bank's P&L jumped predicts no increase in the probability that the second bank's P&L will jump too in the same time frame. By contrast, in our mutually exciting model based on Hawkes processes, the fact that the first bank's P&L jumped leads to a substantial increase in the probability of the second bank's P&L jumping, by a factor $\beta_{2,1}$. Recall that cross-excitation can be asymmetric with, for example, cross-excitation in the reverse direction, i.e., $\beta_{1,2}$, smaller than $\beta_{2,1}$. In addition to the cross-excitation as measured by $\beta_{2,1}$ and $\beta_{1,2}$, we also allow for the presence of self-excitation represented by $\beta_{1,1}$ and $\beta_{2,2}$.

These differences are apparent when comparing the joint pdf of the two banks' P&L's implied by the three models. The results are reported in Figure 4 in the form of contour plots of the joint right tails, or upper right quadrant, of the resulting distributions of (sign-




Figure 4: Joint pdf comparison.

Note: This figure displays contour plots of the joint pdf for three models: a Bernoulli (upper left panel), a Poissonian (upper right panel) and a bivariate mutually exciting Hawkes (lower left panel) model over a 1-day time horizon. For the Poissonian and Hawkes models, we employ the saddlepoint approximation (17). The benchmark parameters are $\alpha_1 = \alpha_2 = 1.5$, $\beta_{1,1} = \beta_{2,2} = 1.0$, $\beta_{1,2} = \beta_{2,1} = 0.4$, $\lambda_1 = \lambda_2 = 20$, $\gamma_1 = \gamma_2 = 1/100$. In the lower right panel, we consider an increase in the cross-excitation parameter $\beta_{2,1}$ from 0.4 to 1.4 and a decrease in the self-excitation parameter $\beta_{1,1}$ from 1.0 to 0.2, compared to the lower left panel.

changed) P&L's $(X_{1,t,\Delta}, X_{2,t,\Delta})$. To examine the magnitude of the excitation effect, we compare contour plots showing the P&L distributions over 1-day horizons obtained under the three models. The probability of observing successive jumps over the same small time period decreases rapidly with the number of jumps in the Bernoulli and Poissonian jump models. Under the Poissonian model, the independence over time imposed by the Lévy structure implies that multiple jumps in a single P&L series are unlikely to be recorded

over a short time period when on average a limited number of jumps occurs in a year. Furthermore, both models imply cross-sectional independence. Thus, it is very unlikely that more than a single jump will be recorded under these two models. By contrast, the introduction of mutual excitation substantially increases the probability that multiple jumps will be recorded in the same time period compared to the two simpler models: once a jump occurs, more are likely to come, both for the bank where the first loss occurred but also for other banks. The lower panel plot shows this clearly, with a noticeably larger (twofold) mass in the joint right tail.

In the lower right panel of Figure 4, we examine the impact of an increase in the cross-excitation parameter $\beta_{2,1}$, along with a decrease in the self-excitation parameter $\beta_{1,1}$ on the joint pdf. In this case, the asymmetry between the cross-excitation effects manifests itself in the higher probability of observing jumps in the second institution when compared to the first institution. As a result, the corresponding right tail is further amplified and the contour plot becomes asymmetrical.

Compared to independent bank risks, mutual excitation introduces the possibility of systemic risk by raising the probability that multiple institutions or assets will jump over the same relatively short 1-day time period over which the VaR is evaluated, and that multiple such jumps will occur — in other words, that a domino effect with systemic consequences can take place. The marginal contribution of each parameter in our model to a risk measure can be used to assess a possible additional capital charge to reflect the negative (self- or cross-) contagion externality.

5.1 Univariate Tails and Risk Measures

Since VaR is the inverse of the cdf, saddlepoint approximations to the VaR are readily obtained from the saddlepoint approximations to the cdf given in Section 4.1, specifically Theorem 1, and similarly for ES, using Corollary 1.

In the univariate case, Hawkes jumps are distinguished from Poisson jumps (i.e., $\beta = 0$) by the effect of self-excitation (i.e., $\beta > 0$). With realistic parameter values, self-excitation acts quickly to increase the jump intensity in response to past jumps, and the tail of the bank's P&L distribution will be affected. A Bernoulli or Poissonian jump component gives rise to a larger tail compared to a no-jumps, diffusion-only model. However, it is highly unlikely with Bernoulli and Poissonian jumps under realistic parameter values to observe more than a single jump over a 1- or 10-day horizon, which then results in computing VaR and setting capital reserves to cover only a one-jump event over that period. Yet, the empirical evidence strongly suggests that losses tend to occur in quick succession. Hawkes jumps make it possible, and even likely if the self-excitation parameter β is sufficiently high, to record successive jumps in a short interval of time, and this impacts the magnitude of the tail of the bank's P&L distribution and ultimately its risk measures.

To illustrate, we evaluate the impact of the self-excitation parameter β on the univariate VaR of a bank. Figure 5 displays the VaR of a bank's daily P&L at the 99.7% probability level in the univariate self-exciting jump model as a function of the self-excitation parameter β . We observe that the VaR increases sharply with the extent of self-excitation.

We next move to the bivariate model, but analyze in this subsection just the univariate marginal tail within the bivariate model. In Figure 6, we plot the unconditional marginal VaR at the 99.7% probability level of bank 1's daily P&L in the bivariate mutually exciting jump model as a function of the self- and cross-excitation parameters $\beta_{1,1}$ and $\beta_{1,2}$. The figure demonstrates two important features. First, as is intuitively appealing, the extent of self-excitation has a slightly more pronounced impact on the marginal VaR than the extent of cross-excitation, hence the slight asymmetry in the contour lines. Second, and most importantly, both the self-exciting and the cross-exciting phenomenon induce a sharp, mutually reinforcing, increase in the marginal VaR measure when compared to the Poissonian model, which occurs as a special case when $\beta_{i,i} = \beta_{i,j} = 0$ for all (i, j) .

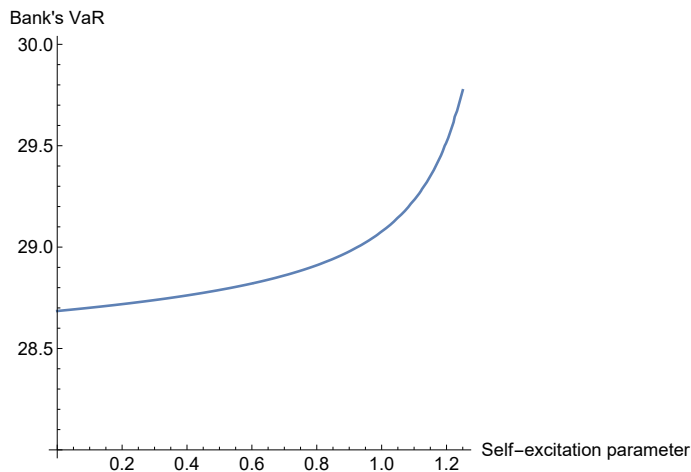


Figure 5: Univariate VaR.

Note: This figure plots the Value-at-Risk of a bank’s daily P&L at the 99.7% probability level in the univariate self-exciting jump model as a function of the self-excitation parameter β . We employ the saddlepoint approximation to the ddf in the univariate P&L model (given in (15) and extended to allow for a non-zero value of the volatility parameter σ), invert the expression, and display the resulting quantile as a function of the parameter β . The benchmark parameters are $\alpha = 1.5$, $\lambda = 1.0$, $\xi = 10$, $\gamma = 1/100$, $\mu = 0$, $\sigma = 100$.

5.2 Bivariate Tails and Risk Measures

Empirically, losses in the banking sector tend to be clustered not only in a given bank but also across banks.²⁷ Our model can generate this effect. Bivariate conditional probabilities, bivariate conditional VaR’s and contour plots of the joint probability density are obtained from (inverting) the respective saddlepoint approximations; see Lemma S.5, Theorem 2 and Corollary 2.

In the left panel of Figure 7 we display the VaR of bank 2’s daily P&L at the 99.7% probability level in the bivariate mutually exciting model conditionally upon a loss in bank 1’s P&L account, as a function of the magnitude of bank 1’s loss and of the cross-excitation parameter $\beta_{2,1}$. The figure shows that when the cross-excitation effect is small, as in the Poissonian model where it is fully absent, a loss in the other bank’s P&L account has a muted effect on the bank’s conditional VaR. However, when the cross-excitation effect

²⁷Recall the references in footnote 1.

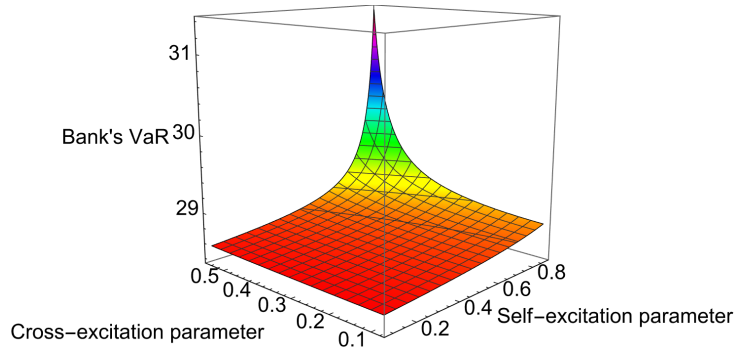


Figure 6: Unconditional marginal VaR in the bivariate model.

Note: This figure plots the unconditional marginal Value-at-Risk of bank 1's daily P&L at the 99.7% probability level in the bivariate mutually exciting jump model as a function of the self- and cross-excitation parameters $\beta_{1,1}$ and $\beta_{1,2}$. We employ the saddlepoint approximation to the marginal ddf in the bivariate P&L model (D.10), invert the expression, and display the resulting quantile as a function of the parameter vector $(\beta_{1,1}, \beta_{1,2})$. The benchmark parameters are $\alpha_1 = \alpha_2 = 1.5$, $\beta_{1,1} = \beta_{2,2}$, $\beta_{1,2} = \beta_{2,1}$, $\lambda_1 = \lambda_2 = 1.0$, $\gamma_1 = \gamma_2 = 1/100$.

becomes more pronounced, the conditional VaR of a bank becomes more and more affected by the loss of the other bank. And the effect increases with the magnitude of the other bank's loss. The sharp increase in the bank's VaR shown in the figure demonstrates that computing a bank's VaR by assuming its losses occur in a Poissonian fashion (i.e., in isolation) can lead to a severe understatement of the actual risk faced by the bank.

In the right panel of Figure 7 we plot the conditional VaR of bank 2's daily P&L as a function of bank 1's loss and of the self-excitation parameter $\beta_{2,2}$. Upon comparing the right panel of Figure 7 to the left panel, we see that both the time-series (self-excitation) and the cross-sectional dimensions (cross-excitation) are relevant when assessing risk: irrespective of the other bank's loss, self-excitation still induces a sharp increase in the conditional VaR, which is consistent with the pattern in Figure 6.

6 Model Inference

Before concluding, we briefly outline how statistical inference can be conducted. The method we have employed to compute the saddlepoint expansion and tail measures consists




Figure 7: Conditional VaR: Impact of Another Bank's Losses on a Bank's VaR

Note: The left panel in this figure plots the conditional Value-at-Risk of bank 2's daily P&L at the 99.7% probability level in the bivariate mutually exciting jump model as a function of the loss y_1 in bank 1's P&L account and of the cross-excitation parameter $\beta_{2,1}$. We employ the saddlepoint approximation to the conditional ddf in the bivariate P&L model (16), invert the expression, and display the resulting conditional quantile as a function of y_1 and $\beta_{2,1}$. The benchmark parameters are $\alpha_1 = \alpha_2 = 1.5$, $\beta_{1,1} = \beta_{2,2} = 1.0$, $\beta_{1,2} = \beta_{2,1}$, $\lambda_1 = \lambda_2 = 1.0$, $\gamma_1 = \gamma_2 = 1/100$. The right panel in this figure plots the conditional Value-at-Risk of bank 2's daily P&L at the 99.7% probability level in the bivariate mutually exciting jump model as a function of the loss y_1 in bank 1's P&L account and of the self-excitation parameter $\beta_{2,2}$. We employ the saddlepoint approximation to the conditional ddf in the bivariate P&L model (16), invert the expression, and display the resulting conditional quantile as a function of y_1 and $\beta_{2,2}$. The benchmark parameters are $\alpha_1 = \alpha_2 = 1.5$, $\beta_{1,1} = \beta_{2,2}$, $\beta_{1,2} = \beta_{2,1} = 1.0$, $\lambda_1 = \lambda_2 = 1.0$, $\gamma_1 = \gamma_2 = 1/100$.

in evaluating conditional expectations of functions of the state vector as series in the time interval Δ : recall Eqn. (13). The same method, applied to other choices of the function ψ than the one that gives rise to the cgf, can be used to compute generalized moment functions for discrete time observations of the banks' P&L generated by the model of Section 2.3.

For example, one can compute (conditional) means, variances, covariances between the P&Ls, as well as higher order moments,²⁸ or more general martingale estimating functions (see Bibby & Sørensen (1995)) over a discrete sampling interval Δ , and derive M-estimators of the parameters of the model; these estimators can be robustified as in La Vecchia & Trojani (2010). It is further possible to allow for random time intervals between observations (as long as they do not Granger-cause the P&Ls), as in Aït-Sahalia & Mykland (2004).

²⁸Explicit expressions are given in Section B of the Supplement, exploiting Section A.5. Specifically, Sections B.1 and B.2 provide the first four moments and the auto- and cross-covariance functions of the regular and squared processes in closed form, for the univariate and bivariate statistical P&L models.

7 Conclusions

We study a natural statistical model for banks' P&L based on Hawkes processes, and characterize its tails. A notable feature of the model, and hence of the risk measures derived from it, is that it generates dependence in banks' P&L in both the cross-sectional and the time-series dimensions, consistent with the empirical and theoretical facts that a bank's losses impose negative externalities onto itself and other banks.

Employing a model that can generate these dependencies in time and space is key to designing effective risk management strategies, and we develop risk measure analytics for this model based on saddlepoint approximations. The formulae derived are closed form, so the effort in computing the risk measures is minimal, no numerical routine is required, and the risk measure formulae may readily be used for comparative statics, parameter calibration, and capital requirement calculations.

The model can also be used for more realistic scenario generation and stress testing in a regulatory framework. Scenarios generated from the model do not just account for the possibility of isolated big losses every now and then, but also allow for, and even make likely, the occurrence of multiple large losses in close succession and across-the-board, so that financial institutions can get in trouble together and such risk can be quantified.

SUPPLEMENTARY MATERIAL

Appendix: Operator methods in the presence of mutual excitation; cumulants and other characteristics in closed form; general expressions of the saddlepoint approximations; technical derivations of our saddlepoint approximations for the univariate case, the marginal tail in the bivariate model and the bivariate case, and simulation results; preliminaries for eigenvector centrality in spectral graph theory. (.pdf file)

References

- Acharya, V. V., Engle, R. & Richardson, M. (2012), ‘Capital shortfall: A new approach to ranking and regulating systemic risks’, *American Economic Review* **30**, 2–47.
- Acharya, V. V., Pedersen, L. H., Philippon, T. & Richardson, M. (2017), ‘Measuring systemic risk’, *Review of Financial Studies* **102**, 59–64.
- Adrian, T. & Brunnermeier, M. K. (2016), ‘Co VaR’, *American Economic Review* **106**, 1705–1741.
- Aït-Sahalia, Y., Cacho-Diaz, J. & Laeven, R. J. (2015), ‘Modeling financial contagion using mutually exciting jump processes’, *Journal of Financial Economics* **117**, 585–606.
- Aït-Sahalia, Y., Laeven, R. J. & Pelizzon, L. (2014), ‘Mutual excitation in Eurozone sovereign CDS’, *Journal of Econometrics* **183**, 151–167.
- Aït-Sahalia, Y. & Mykland, P. A. (2004), ‘Estimators of diffusions with randomly spaced discrete observations: A general theory’, *Annals of Statistics* **32**, 2186–2222.
- Aït-Sahalia, Y. & Yu, J. (2006), ‘Saddlepoint approximations for continuous-time Markov processes’, *Journal of Econometrics* **134**, 507–551.
- Azizpour, S., Giesecke, K. & Schwenkler, G. (2018), ‘Exploring the sources of default clustering’, *Journal of Financial Economics* **129**, 154–183.
- Bacry, E., Delattre, S., Hoffmann, M. & Muzy, J. F. (2013), ‘Modelling microstructure noise with mutually exciting point processes’, *Quantitative Finance* **13**, 65–77.
- Bacry, E., Mastromatteo, I. & Muzy, J. F. (2015), ‘Hawkes processes in finance’, *Market Microstructure and Liquidity* **1**, 1–59.
- Bibby, B. M. & Sørensen, M. S. (1995), ‘Martingale estimation functions for discretely observed diffusion processes’, *Bernoulli* **1**, 17–39.
- Bienvenue, A. & Robert, C. Y. (2016), ‘Systemic tail risk distribution’, *Annals of Economics and Statistics* **123/124**, 29–52.
- Billio, M., Getmansky, M., Lo, A. W. & Pelizzon, L. (2012), ‘Econometric measures of connectedness and systemic risk in the finance and insurance sectors’, *Journal of Financial Economics* **104**, 535–559.
- Bowsher, C. G. (2007), ‘Modelling security market events in continuous time: Intensity based, multivariate point process models’, *Journal of Econometrics* **141**, 876–912.

- Brillinger, D. R. (1988), ‘Some statistical methods for random process data from seismology and neurophysiology’, *Annals of Statistics* **16**, 1–54.
- Broda, S. A., Krause, J. & Paoletta, M. S. (2018), ‘Approximating expected shortfall for heavy-tailed distributions’, *Econometrics and Statistics* **8**, 184–203.
- Brownlees, C. T. & Engle, R. F. (2017), ‘SRISK: A conditional capital shortfall measure of systemic risk’, *The Review of Financial Studies* **30**, 48–79.
- Butler, R. W. (2007), *Saddlepoint Approximations with Applications*, Cambridge University Press, Cambridge, UK.
- Chavez-Demoulin, V., Davison, A. C. & McNeil, A. J. (2005), ‘A point process approach to Value-at-Risk estimation’, *Quantitative Finance* **5**(2), 227–234.
- Daniels, H. (1954), ‘Saddlepoint approximations in statistics’, *Annals of Mathematical Statistics* **25**, 631–650.
- Daniels, H. (1987), ‘Tail probability approximations’, *International Statistical Review* **55**, 37–48.
- Engle, R. F. & Manganelli, S. (2004), ‘CAViaR: Conditional autoregressive value at risk by regression quantiles’, *Journal of Business & Economic Statistics* **22**, 367–381.
- Errais, E., Giesecke, K. & Goldberg, L. R. (2010), ‘Affine point processes and portfolio credit risk’, *SIAM Journal on Financial Mathematics* **1**, 642–665.
- Field, C. A. & Ronchetti, E. (1990), *Small Sample Aymptotics*, Lecture Notes - Monograph Series, Institute of Mathematical Statistics, Hayward, CA.
- Föllmer, H. & Schied, A. (2011), *Stochastic Finance*, 3rd edn, De Gruyter, Berlin.
- Glasserman, P. & Kim, K.-K. (2009), ‘Saddlepoint approximations for affine jump-diffusion models’, *Journal of Economic Dynamics and Control* **33**, 15–36.
- Gordy, M. B. (2002), ‘Saddlepoint approximation of CreditRisk⁺’, *Journal of Banking & Finance* **26**, 1335–1353.
- Hautsch, N., Schaumburg, J. & Schienle, M. (2015), ‘Financial network systemic risk contributions’, *Review of Finance* **19**, 685–738.
- Hawkes, A. G. (1971), ‘Spectra of some self-exciting and mutually exciting point processes’, *Biometrika* **58**, 83–90.

- Hawkes, A. G. (2022), ‘Hawkes jump-diffusions and finance: A brief history and review’, *The European Journal of Finance* **28**, 627–641.
- Huang, X. & Oosterlee, C. W. (2011), ‘Saddlepoint approximations for expectations and an application to CDO pricing’, *SIAM Journal on Financial Mathematics* **2**, 692–714.
- Huang, X., Zhou, H. & Zhu, H. (2012), ‘Systemic risk contributions’, *Journal of Financial Services Research* **42**, 55–83.
- Jaisson, T. & Rosenbaum, M. (2015), ‘Limit theorems for nearly unstable Hawkes processes’, *Annals of Applied Probability* **25**, 600–631.
- Jiang, E., Matvos, G., Piskorski, T. & Seru, A. (2023), Monetary tightening and U.S. bank fragility in 2023: Mark-to-market losses and uninsured depositor runs?, Technical report, NBER.
- Kim, S. & Kim, K.-K. (2017), ‘Saddlepoint methods for conditional expectations with applications to risk management’, *Bernoulli* **23**, 1481–1517.
- Kolassa, J. E. & Li, J. (2010), ‘Multivariate saddlepoint approximations in tail probability and conditional inference’, *Bernoulli* **16**, 1191–1207.
- Kwok, Y. K. & Zheng, W. (2018), *Saddlepoint Approximation Methods in Financial Engineering*, Springer, Cham, Switzerland.
- La Vecchia, D. & Trojani, F. (2010), ‘Infinitesimal robustness for diffusions’, *Journal of the American Statistical Association* **105**, 703–712.
- Laeven, R. J. & Stajda, M. (2013), ‘Entropy coherent and entropy convex measures of risk’, *Mathematics of Operations Research* **38**, 265–293.
- Lugannani, R. & Rice, S. (1980), ‘Saddlepoint approximation for the distribution of the sum of independent random variables’, *Advances in Applied Probability* **12**, 475–490.
- Mainik, G. & Schaanning, E. (2014), ‘On dependence consistency of CoVaR and some other systemic risk measures’, *Statistics & Risk Modeling* **31**(1), 49–77.
- Martin, R. (2006), ‘The saddlepoint method and portfolio optionalities’, *Risk Magazine* **19**, 93–95.
- Merton, R. C. (1976), ‘Option pricing when underlying stock returns are discontinuous’, *Journal of Financial Economics* **3**, 125–144.

- Mykland, P. A. (1999), ‘Bartlett identities and large deviations in likelihood theory’, *Annals of Statistics* **27**, 1105–1117.
- Ogata, Y. & Akaike, H. (1982), ‘On linear intensity models for mixed doubly stochastic Poisson and self-exciting point processes’, *Journal of the Royal Statistical Society, B* **44**, 102–107.
- Pérignon, C., Deng, Z. Y. & Wang, Z. J. (2008), ‘Do banks overstate their Value-at-Risk?’, *Journal of Banking & Finance* **32**(5), 783–794.
- Pérignon, C. & Smith, D. R. (2010), ‘Diversification and Value-at-Risk’, *Journal of Banking & Finance* **34**(1), 55–66.
- Reid, N. (1988), ‘Saddlepoint methods and statistical inference’, *Statistical Science* **3**(2), 213–227.
- Reinhart, C. M. & Rogoff, K. S. (2009), *This Time Is Different: Eight Centuries of Financial Folly*, Princeton University Press, Princeton, NJ.
- Swishchuk, A., Zagst, R. & Zeller, G. (2021), ‘Hawkes processes in insurance: Risk model, application to empirical data and optimal investment’, *Insurance: Mathematics and Economics* **101**, 107–124.
- White, H., Kim, T.-H. & Manganelli, S. (2015), ‘VAR for VaR: Measuring tail dependence using multivariate regression quantiles’, *Journal of Econometrics* **187**, 169–188.
- Wood, A. T. A., Booth, J. G. & Butler, R. W. (1993), ‘Saddlepoint approximations to the CDF of some statistics with nonnormal limit distributions’, *Journal of the American Statistical Association* **88**, 680–686.