Dynamic Large Financial Networks via Conditional Expected Shortfalls

Giovanni Bonaccolto

School of Economics and Law, “Kore” University of Enna, Italy

Massimiliano Caporin

Department of Statistical Sciences, University of Padova, Italy

Bertrand B. Maillet

Corresponding Author

*Emlyon Business School (AIM QUANT Research Center), Paris, France
Univ. La Reunion (CEMOI), Saint-Denis, France; Variances, Paris, France*

Abstract

In this article, we first generalize the Conditional Auto-Regressive Expected Shortfall (CARES) model by introducing the loss exceedances of all (other) listed companies in the Expected Shortfall related to each firm, thus proposing the CARES-X model (where the ‘X’, as usual, stands for eXtended in the case of large-dimensional problems). Second, we construct a regularized network of US financial companies by introducing the Least Absolute Shrinkage and Selection Operator in the estimation step. Third, we also propose a calibration approach for uncovering the relevant edges between the network nodes, finding that the estimated network structure dynamically evolves through different market risk regimes. We ultimately show that knowledge of the extreme risk network links provides useful information, since the intensity of these links has strong implications on portfolio risk. Indeed, it allows us to design effective risk management mitigation allocation strategies.

Keywords: Finance, Financial networks, Portfolio analysis, Systemic risk, Expectile regression

Email addresses: giovanni.bonaccolto@unikore.it (Giovanni Bonaccolto),
massimiliano.caporin@unipd.it (Massimiliano Caporin), bmaillet@em-lyon.com
(Bertrand B. Maillet)

1. Introduction

In the aftermath of the recent subprime crisis, a significant increase of research interest in systemic risk measures and financial networks has been observed in the financial economics and econometrics literature. Both theoretical and empirical points of view were adopted and led to original operational research methodologies (e.g. Calabrese and Osmetti, 2019; Gupta et al., 2021). Indeed, the overall knowledge of the level of systemic risk associated with each single company or financial instrument reveals several advantages. In fact, it might help regulators identify the Systemically Important Financial Institutions (SIFI), and thus help to follow the evolution of their fragilities, as well as assess and gauge the stability of the entire financial system. These aspects also have implications for the diffusion of crises within the financial system, with a link to the financial contagion literature.

From a different standpoint, systemic risk aspects might also be integrated as a complement to other market risk management tools, for monitoring both single assets and the entire financial market. In particular, these characteristics of emerging systemic risk can be seen as useful for developing early warning systems designed for detecting the surge of system-wide events. Consequently, this might have implications both in terms of risk mitigation strategies for individual components of the financial market, as well as in reducing the impact of systemic events, beside contributing to the literature on risk minimization (e.g. Kremer et al., 2018).

In this context, our article aims at shedding further light on the link between asset co-movements and portfolio risk, when focusing both on extreme risk and financial connections in a large network. Our work is, indeed, at the juncture of three strands of literature.

The first focuses on extreme risk measures and, more specifically, on Expected Shortfall (ES). Indeed, ES, also called Conditional Value-at-Risk (CVaR), has garnered increasing interest in the portfolio management literature (see, among others, Barnard et al., 2018; Pac and Pinar, 2018; Ramponi and Campi, 2018). In this context, we first need to explicitly reference the Conditional Auto-Regressive Expectile (CARE) model by Taylor (2008), as this seminal article is the basis of several others, our current article included. Among the family of such articles, Engle and Manganelli (2004) first proposed to apply GARCH-type modelling to quantiles, in a so-called Conditional Auto-Regressive Value-at-Risk (CAViaR) by regression

quantile. Gerlach and Chen (2014) and Meng and Taylor (2020) generalized the approach of dynamic quantile models with the incorporation of intra-day information in conditional expectile models, estimated using a Bayesian method. Hamidi et al. (2015) also built on this approach through the use of a combination of Dynamic Auto-Regressive Expectiles (DARE) models, which is analogous to the Dynamic Auto-Regressive Quantile (DAQ) modelling approach of Gouriéroux and Jasiak (2008).

Kim and Lee (2016) investigated some of the theoretical and empirical properties of VaR and ES estimations based on Taylor’s (2008) proposal and a broad class of non-linear conditional expectile models, whilst Patton et al. (2019) and Taylor (2019) proposed various dynamic semi-parametric models for VaR and ES. Another road is taken in trying to diversify the model risk, using a combination of models as in Taylor (2020). We will hereafter characterize the extreme risks of companies using their ESs. Our first contribution will thus be in the estimation of Conditional Auto-Regressive Expected Shortfall (CARES) for various companies.

The second strand of literature is related to large financial networks. Among the various approaches currently available to build such networks, we mention the use of Granger causality (Billio et al., 2012), the spillover summarized by forecast error variance decomposition (Diebold and Yilmaz, 2009, 2012, 2014, 2015), the use of VaR (conditional or unconditional) within a static model (Hautsch et al., 2014; Adrian and Brunnermeier, 2016), the construction of composite networks monitoring different dimensions of asset interconnections (Bonaccolto et al., 2019a). Our second contribution will be the focus on the estimation of extreme risk asset interconnections. It will be based on a vector Conditional Auto-Regressive Expected Shortfall approach (called CARES-X hereafter, where the ‘X’ stands for eXtended in the case of large-dimensional problems), whilst previous works were not fully dynamic and were mainly relying on Value-at-Risk (VaR) links.

The third stream of literature relies on the impact of systemic risk on asset allocation and financial characteristics of various portfolios, gauging the importance of measuring the severity of tail events when optimizing a portfolio in an operational research approach (e.g. Maillet et al., 2015). More specifically, a particular strand of the literature related to systemic risk focuses on the implications of portfolio allocation and portfolio strategies, taking into account two different aspects: i) evaluating financial network-related and asset interdependence-related risk factors, which impact on risk factor exposures or are priced in the cross-section of assets (Billio et al., 2017); and ii) assessing whether the knowledge of systemic risk or asset interdependence (i.e. through network links) might be useful in portfolio

allocation and risk management (Clemente et al., 2019).

In other words, in the same spirit as Borri (2019) and Nguyen et al. (2020), who focused on cryptocurrencies and tail-relationships, we present herein a generalization of the approach by Hautsch et al. (2014), with estimations of financial networks focusing on dynamic ES. Indeed, Hautsch et al. (2014), when estimating the impact of other companies on the VaR of targeted companies, made use of a penalized quantile regression, thus not fully dealing with potential dynamic evolutions in the conditional quantiles. We thus rely on the approach of Hautsch et al. (2014), but using estimations of conditional ES, following Taylor (2008). Finally, we introduce penalization to uncover the relevant interconnections and exclude the irrelevant ones, making our approach feasible even in a large cross-section of assets (conditional on the availability of a sufficiently long time-series). The estimation of a financial network builds on the identification of relevant edges between its nodes, starting from the model estimates. While Hautsch et al. (2014) made a selection across companies using a unique threshold on model parameters, we make a step forward and introduce a criterion for the calibration of the threshold. We use a company-specific threshold, which is also time dependent. Such a choice, more expensive from a computational viewpoint, has the advantage of allowing us to recover a proper identification of the links between the extreme risks of the analyzed companies under different market regimes.

Starting from our methodological advancements, we move to empirical analyses, first estimating the network of companies belonging to the financial sector of the US equity market. Our analyses clearly show that the network derived from the CARES-X model, as summarized by vertex degrees and density, is coherent with market swings, both during the subprime financial crisis, as well as the following recovery. Then, we secondly exploit the informative content of the network in terms of portfolio risk and performance analysis. Resorting to naïve portfolio allocation strategies, thus mitigating weight estimation risk and uncertainty (Michaud, 1989), we highlight the impact of tail risk interconnections on the measurement of portfolio exposure to extreme events. By separating companies in terms of their network exposure, we show that it is possible to implement allocation strategies capable of mitigating the overall portfolio risk.

The tools we introduce and the empirical evidence we provide might be of relevant interest for portfolio risk managers, at the portfolio construction and allocation level, as well as for portfolio risk monitoring. In this last respect, by adapting our methodology to higher frequency returns, it will be possible to detect the surge of systemic risk events, for instance monitoring

the changes in the network density or in the vertex degree. This might have relevant implications on the derivation of early warning indicators. Finally, by combining our extreme-risk networks with approaches focusing on different dimensions of companies' systemic risk, we might improve the evaluation of composite and multilayer networks, whose stability and structure may be of interest to regulators when analyzing various market configurations across different risk regimes.

The article proceeds as follows: in Section 2, we describe our proposal for financial network estimation, also discussing the relation of our approach with competing methodologies. Section 3 describes the data and the implementation choices for our methodology, while Section 4 shows the estimated networks and analyzes their features and dynamic behavior. Section 5 discusses the use of the resulting networks within a portfolio setting and, finally, Section 6 concludes.

2. Penalized CARES-X model and network estimation

Let R_j be the return of the j^{th} financial company, whose realization at time t is denoted as $r_{j,t}$, for $j = 1, \dots, N$ and $t = 1, \dots, T$.¹ We study the links among N companies within a network that we estimate from a pessimistic perspective; namely, a network recovered from companies when they are jointly in distress. We measure the state of distress of each company through ES (i.e. the tail expectation of R_j), which is equal to:

$$ES(R_j; \theta) = \mathbb{E}[R_j | R_j \leq Q(R_j; \theta)], \quad (1)$$

where $Q(R_j; \theta)$ is the quantile of R_j at the level θ , with θ taking low values, typically in the interval $(0, 0.05]$, because we focus on the left tail of the distribution of R_j .

In our analysis, we hereafter estimate ES using the expectile regression method, following the approach introduced by Taylor (2008).² We first define the population expectile of R_j at level τ as the quantity $\mu(R_j; \tau)$

¹More precisely, our dataset includes $B > T$ daily returns for each company. We use a rolling window scheme in the empirical analysis and, then, divide the overall dataset into sub-samples of equal size, including T returns for each stock (see Section 3 for additional details). The methods described in this methodological section refer to a generic sample which spans the time interval $[1, T]$; that is, the first sub-sample resulting from the rolling window procedure. However, they apply in the same way for the other sub-samples.

²See, among others, Artzner et al. (1999), Rockafellar and Uryasev (2000) and Acerbi and Tasche (2002) for a detailed review on the properties of ES.

which minimizes the following expected loss (see, among others, Taylor, 2008; Bellini and Bernardino, 2017):

$$\mathbb{E} \left[\left| \tau - \mathbb{I}_{\{R_j < \mu(R_j; \tau)\}} \right| (R_j - \mu(R_j; \tau))^2 \right], \quad (2)$$

where $\tau \in (0, 1)$, whereas $\mathbb{I}_{\{\cdot\}}$ is an indicator function which takes the value of one, if the condition into braces is true, and zero otherwise.

Taylor (2008) shows that the ES of R_j at level θ (i.e. $ES(R_j; \theta)$) corresponds to the following function of $\mu(R_j; \tau)$:

$$ES(R_j; \theta) = \left(1 + \frac{\tau}{(1 - 2\tau)\theta} \right) \mu(R_j; \tau) - \frac{\tau}{(1 - 2\tau)\theta} \mathbb{E}[R_j], \quad (3)$$

where $\mathbb{E}[R_j]$ is the expected value of R_j .

The relationship between $ES(R_j; \theta)$ and $\mu(R_j; \tau)$ in Equation (3) is defined for a scalar $\mu(R_j; \tau)$. Notably, the relation holds even if expectiles are conditional on a set of explanatory variables (Taylor, 2008). This property allows us to empirically estimate the τ^{th} expectile of $r_{j,t}$ (i.e. $\mu(r_{j,t}; \tau)$) from a model which also accounts for the impact of a set of covariates. Among the possible specifications for $\mu(r_{j,t}; \tau)$, we may use the Conditional Auto-Regressive Expectile (CARE) model introduced by Taylor (2008), which takes the following form:

$$\mu(r_{j,t}; \tau) = \alpha_{\tau}^{(j)} + \gamma_{1,\tau}^{(j)} \mu(r_{j,t-1}; \tau) + \gamma_{2,\tau}^{(j)} |r_{j,t-1}|. \quad (4)$$

The CARE model thus captures the persistence and dynamics of expectiles over time, sharing the properties of the Conditional Auto-Regressive VaR (denoted CAViaR) model introduced by Engle and Manganelli (2004) for regression quantiles. Including the latent and lagged expectile $\mu(r_{j,t-1}; \tau)$ smoothes the changes of $\mu(r_{j,t}; \tau)$ over time, whereas the parameter $\gamma_{2,\tau}^{(j)}$ directly links $\mu(r_{j,t}; \tau)$ to the past return $r_{j,t-1}$. Following Taylor (2008), we employ, for the estimation of the unknown parameters in Equation (4), the algorithm developed by Engle and Manganelli (2004) for the CAViaR model, but adapted to minimize the following objective function:

$$L(\alpha_{\tau}^{(j)}, \gamma_{1,\tau}^{(j)}, \gamma_{2,\tau}^{(j)}) = \frac{1}{T-1} \sum_{t=2}^T \left| \tau - \mathbb{I}_{\{r_{j,t} < \mu(r_{j,t}; \tau)\}} \right| [r_{j,t} - \mu(r_{j,t}; \tau)]^2, \quad (5)$$

where $\mu(r_{j,t}; \tau)$ is defined in Equation (4).

Taylor (2008) proposed a simple rule to empirically derive $ES(r_{j,t}; \theta)$ from Equation (4), advocating that $ES(r_{j,t}; \theta) = \mu(r_{j,t}; \tau_j^*)$. Therefore,

this rule builds on the estimation of $\tau_j^\star$; that is, the value of τ which makes the percentage of in-sample violations in Equation (5) equal to θ , as such:

$$(T-1)^{-1} \sum_{t=2}^T \mathbb{I}_{\{r_{j,t} < \mu(r_{j,t}; \tau_j^\star)\}} = \theta, \quad (6)$$

where θ is fixed *a priori*.

This rule leads to the Conditional Auto-Regressive Expected Shortfall (CARES) model of Taylor (2008), which is defined as follows:

$$ES(r_{j,t}; \theta) = \mu(r_{j,t}; \tau_j^\star) = \alpha_{\tau_j^\star}^{(j)} + \gamma_{1, \tau_j^\star}^{(j)} \mu(r_{j,t-1}; \tau_j^\star) + \gamma_{2, \tau_j^\star}^{(j)} |r_{j,t-1}|. \quad (7)$$

A simple CARES model allows us to estimate the ES of $r_{j,t}$ conditional on the past history of $r_{j,t}$, but, unfortunately, does not include the effects of other financial companies. However, we are interested in measuring the co-movements between the j^{th} company and other institutions in a state of joint distress, capturing the persistence and dynamics of $ES(r_{j,t}; \theta)$ at the same time. We then extend the model in Equation (7) by including the loss exceedances, as defined by Hautsch et al. (2014), of the other $N-1$ conditioning companies. By doing so, we link the extreme quantiles of the returns yielded by the N companies we focus on, emphasizing their co-movements during tail events.³ More precisely, let $\bar{r}_j$ and $\hat{\sigma}_{r_j}$ be the sample mean and standard deviation of $r_{j,t}$, respectively, so that $z_{j,t} = (r_{j,t} - \bar{r}_j) / \hat{\sigma}_{r_j}$ is the standardized value of $r_{j,t}$. Following Hautsch et al. (2014), we compute the loss exceedances, defined as:

$$h_{j,t} = z_{j,t} \mathbb{I}_{\{z_{j,t} < q(z_{j,t}; 0.1)\}}, \quad (8)$$

where $q(z_{j,t}; 0.1)$ is the unconditional sample quantile of $z_{j,t}$ at the probability level 0.1, for $t = 1, \dots, T$ and $j = 1, \dots, N$.

The resulting CARES-X model (where the ‘X’ stands for eXtended in the case of large-dimensional problems) then takes the following specification:

$$\begin{aligned} ES(r_{j,t}; \theta) &= \mu(r_{j,t}; \tau_j^\star) = \alpha_{\tau_j^\star}^{(j)} + \gamma_{1, \tau_j^\star}^{(j)} \mu(r_{j,t-1}; \tau_j^\star) + \gamma_{2, \tau_j^\star}^{(j)} |r_{j,t-1}| \\ &+ \boldsymbol{\psi}_{\tau_j^\star}^{(j)} \mathbf{h}'_{-j,t} + \boldsymbol{\delta}_{\tau_j^\star}^{(j)} \mathbf{w}'_t, \end{aligned} \quad (9)$$

³By employing a different methodological approach, Bonaccolto et al. (2019a,b) also analyzed the relationships among financial institutions focusing on the extreme quantiles of their returns. They highlighted the relevance of the so-called quantile-located or quantile-on-quantile effects when studying the systemic impact of financial institutions during tail events.

where $\mathbf{h}_{-j,t}$ is an $1 \times (N - 1)$ vector, obtained by removing the element $h_{j,t}$ from $\mathbf{h}_t = [h_{1,t} \cdots h_{N,t}]$, with $1 \leq j \leq N$, $\mathbf{w}_t$ is an $1 \times K$ vector of standardized control variables observed at time t , whereas τ_j^* is the value of τ which satisfies the condition in Equation (6), given the new specification of $\mu(r_{j,t}; \tau_j^*)$ in (9).⁴

Note that our CARES-X model does not include the lagged expected shortfalls of the $N - 1$ conditioning companies $\mu(r_{i,t-1}; \tau_i^*)$, with $1 \leq i \leq N$ and $i \neq j$. Indeed, we aim at estimating the contemporaneous relationships between the loss exceedances of the conditioning companies and the ES of R_j , rather than forecasting $ES(r_{j,t}; \theta)$. For practical reasons, our proposed method should be flexible for dealing with a large number of companies. Nevertheless, the number of parameters to estimate in Equation (9) increases with N and the resulting accumulation of estimation errors becomes a critical issue. Of course, we do not know *a priori* which of the variables in $\mathbf{h}_{-j,t}$ have a relevant impact on $ES(r_{j,t}; \theta)$ and, thus, we face the following *dilemma*. On the one hand, including too many regressors could imply over-fitting problems. On the other hand, we could suffer from an omitted variable bias when using a restricted subset of covariates. We deal with the curse of dimensionality using a well-known variable selection tool, namely, the Least Absolute Shrinkage and Selection Operator (LASSO) introduced by Tibshirani (1996) and widely employed in economic and financial applications (e.g. Torri et al., 2018; Pun and Wong, 2019; Cui et al., 2020).

Hence, we estimate the parameters in Equation (9) by minimizing the following loss function:

$$\begin{aligned} L(\boldsymbol{\beta}_{\tau_j^*}^{(j)}) &= \frac{1}{T-1} \sum_{t=2}^T \left| \tau_j^* - \mathbb{I}_{\{r_{j,t} < \mu(r_{j,t}; \tau_j^*)\}} \right| [r_{j,t} - \mu(r_{j,t}; \tau_j^*)]^2 \\ &+ \lambda_{\tau_j^*}^{(j)} \left\| \boldsymbol{\psi}_{\tau_j^*}^{(j)} \right\|_1, \end{aligned} \quad (10)$$

where $\mu(r_{j,t}; \tau_j^*)$ and $\boldsymbol{\beta}_{\tau_j^*}^{(j)} = [\alpha_{\tau_j^*}^{(j)}, \gamma_{1, \tau_j^*}^{(j)}, \gamma_{2, \tau_j^*}^{(j)}, \boldsymbol{\delta}_{\tau_j^*}^{(j)}, \boldsymbol{\psi}_{\tau_j^*}^{(j)}]$ are, respectively, the conditional ES and the parameters of our CARES-X model defined in Equation (9), $\lambda_{\tau_j^*}^{(j)} > 0$ is a tuning parameter which governs the intensity of the penalization, whereas $\|\cdot\|_1$ is the ℓ_1 -norm.

⁴The procedure we use to compute τ_j^* in our empirical analysis is described in detail in Section 3.

The greater $\lambda_{\tau_j^*}^{(j)}$ is, the greater the number of coefficients in $\boldsymbol{\psi}_{\tau^*}^{(j)}$ that approach zero, resulting in a sparser solution. We select the optimal value of this tuning parameter by employing a cross-validation technique which is suitable for time series (see Section 3). This method is commonly used in applied machine learning, being flexible and easy to understand and implement, providing at the same time accurate results. We highlight that, among the parameters in (10), only $\boldsymbol{\psi}_{\tau^*}^{(j)}$ is penalized. By doing so, we emphasize the relationships between the j^{th} company and the remaining $N - 1$ conditioning institutions and select, among the latter, the ones which strongly impact $ES(r_{j,t}; \theta)$. At the same time, we capture the dynamics of the response variable over time as well as potential effects arising from the equity and bond markets. Note that the penalty function in (10) can be expressed as the adaptive LASSO employed by Liao et al. (2019), by setting the weights multiplying the parameters in $\boldsymbol{\psi}_{\tau^*}^{(j)}$ equal to one, and the weights of the remaining parameters equal to zero.

We LASSO-select the i^{th} conditioning institution (for $i = 1, \dots, N$ and $i \neq j$) if the impact of its loss exceedance in Equation (9), which is quantified by the specific coefficient in $\boldsymbol{\psi}_{\tau^*}^{(j)}$ that multiplies $h_{i,t}$, is, in absolute value, greater than a given threshold η ; for the latter, we provide an empirical identification strategy in Section 3. We then build an adjacency matrix $\mathbf{A} = [a_{j,i}] \in \mathbb{R}^{N \times N}$, in which $a_{j,i} = 1$ if the loss exceedance of the i^{th} company has a strong impact (i.e. is LASSO-selected) on the ES of the j^{th} institution, and $a_{j,i} = 0$ otherwise, for $i, j = 1, \dots, N$ and $i \neq j$. We exclude self-loops, so that $a_{j,i} = 0$ if $j = i$.

The closest contribution to this study is the work of Hautsch et al. (2014). However, we differ from Hautsch et al. (2014) on several relevant points. First, the method introduced by Hautsch et al. (2014) builds on a quantile regression model. Different contributions in the literature (see, among others, Newey and Powell, 1987; Efron, 1991; Koenker, 1992, 1993; Jones, 1994; Yao and Tong, 1996; Koenker, 2005; Taylor, 2008; Yang and Zou, 2015; Bellini and Bernardino, 2017; Furno and Vistocco, 2018) highlighted more appealing statistical and computational properties of expectile regression compared to quantile regression (see Appendix A). Second, an extreme quantile is interpreted as VaR, whereas expectiles lead to ES. ES is a coherent measure of risk (Artzner et al., 1999), whereas this is not the case for VaR. The superior mathematical and financial properties of ES motivated the changes that the Basel Committee on Banking Supervision (BCBS) has adopted in the Basel III accords. In particular, the BCBS has substituted VaR with ES to define the market risk capital requirements. Our network

is then consistent with the latest regulatory standards. Other differences, explained in detail in Appendix A, concern the specification of the regression model and the method used to LASSO-select the relevant connections among the N financial institutions.

3. Data and empirical setup

Our dataset includes the daily returns of the companies which belong to the basket of index STOXX 600 Financials from January 2, 2004 to November 16, 2017 (i.e. $B = 3620$ trading days).⁵ The number of these companies is equal to $N = 104$, classified into four sectors: banking (38), insurance (26), financial services (18) and real estate (22). We report in Appendix B the list of the $N = 104$ financial institutions included in our dataset. We observe a clear volatility clustering in our data. Figure C.9 in Appendix C reflects the turmoil experienced during particular events in both the US and the European markets from January 2, 2004 to November 16, 2017. Among them, the US subprime crisis (period 2007—2009) has the strongest impact.

The estimation of the CARES-X model in Equation (9) requires a set of covariates to include in $\mathbf{w}_t$. Similar to Hautsch et al. (2014), we use nine control variables related to bond and equity markets, which are listed as follows: i) the return of the STOXX Europe 600 Index; ii) the return of the VSTOXX volatility index; iii) the change in the three-month EBF Euribor rate; iv) the change in the ECB short-term repo rate; v) the change in the EMU benchmark 10-year government bond index; vi) the change in the ICE BofAML Pan-Europe large cap corporate index; vii) the change in the ICE BofAML 10+ year Euro corporate index; viii) the change in the ICE BofAML 10+ year AAA Euro government index; and ix) the return of the five-year Europe banks sector CDS index.⁶ From a principal component analysis of the nine control variables described above, we found that the first three components explain 83.14% of the variability in the data. We then use the standardized values of these three components, in place of the original variables, as elements of $\mathbf{w}_t$. By doing so, we exploit the near totality of the information contained in the nine original series, while achieving relevant benefits in terms of computational burden and stability of the resulting estimates.

⁵The data are recovered from Thomson Reuters Datastream.

⁶The control variables are taken from Thomson Reuters Datastream.

We estimate the parameters of the CARES-X model defined in Equation (9) by setting $\theta = 0.05$ and minimizing the loss function in Equation (10). However, we do not know *a priori* τ_j^* ; that is, the value of τ which makes the percentage of in-sample violations equal to $\theta = 0.05$. In contrast, we need to select τ_j^* from a set of expectile levels τ . We also consider the possibility that the optimal value of the tuning parameter in Equation (10) varies with τ . Therefore, we define a set of 20 equally-spaced expectile levels in $\mathcal{S}_\tau = \{\tau_1, \dots, \tau_{20}\}$, where $\tau_1 = 0.001$ and $\tau_{20} = 0.05$, and minimize, for each $\tau_i \in \mathcal{S}_\tau$ and for a sufficiently large grid of $\lambda_{\tau_i}^{(j)}$ values, the following loss function:

$$\frac{1}{B-1} \sum_{t=2}^B \left| \tau_i - \mathbb{I}_{\{r_{j,t} < \mu(r_{j,t}; \tau_i)\}} \right| \cdot [r_{j,t} - \mu(r_{j,t}; \tau_i)]^2 + \lambda_{\tau_i}^{(j)} \left\| \boldsymbol{\psi}_{\tau_i}^{(j)} \right\|_1,$$

where $\mu(r_{j,t}; \tau_i) = \alpha_{\tau_i}^{(j)} + \gamma_{1,\tau_i}^{(j)} \mu(r_{j,t-1}; \tau_i) + \gamma_{2,\tau_i}^{(j)} |r_{j,t-1}| + \boldsymbol{\psi}_{\tau_i}^{(j)} \mathbf{h}'_{-j,t} + \boldsymbol{\delta}_{\tau_i}^{(j)} \mathbf{w}'_t$.⁷

For each $\tau_i \in \mathcal{S}_\tau$ we select an optimal value of $\lambda_{\tau_i}^{(j)}$, denoted as $\hat{\lambda}_{\tau_i}^{(j)}$, using a cross-validation method which is suitable to deal with time series. This method is similar to a standard 3-fold cross-validation (Hastie et al., 2009), differing only in the definition of the training and validation sets. In particular, we define three training sets, which cover the time intervals $[1, 1500]$, $[1, 2500]$ and $[1, 3400]$, respectively. The corresponding validation sets span the intervals $[1501, 1700]$, $[2501, 2700]$ and $[3401, 3600]$, respectively, so that they have a constant size of 200 observations.

We point out that this cross-validation exercise builds on training sets whose extent differs from that of the samples from which we estimate the CARES-X parameters. This is a consequence of employing expanding training sets (which have different widths, by definition) that are coherent with the time evolution of our variables of interest. Alternatively, we could use a standard k -fold cross-validation (Hastie et al., 2009) to build training sets of equal size, which coincides with the width of the estimation window. However, such a choice would induce a distortion in our analyses. In fact, by choosing purely random samples and inserting them into either validation or training sets, we would destroy the temporal dependencies which are present in our time series, altering the properties of our Conditional Auto-Regressive model. In contrast, the cross-validation method we employ preserves the chronological order in our time series and, hence, we do

⁷We focus on $\tau_i \in [0.001, 0.05]$, without using greater values, because we found that τ_j^* does not exceed the value of 0.05 in our empirical analysis, for $j = 1, \dots, N$.

not negatively affect the Auto-Regressive dynamics of the CARES-X model. Moreover, we adopt relatively larger training sets to increase the stability of the Auto-Regressive components in the CARES-X model. This evidence becomes clearer when focusing on the lowest and extreme expectile levels of the $\mathcal{S}_\tau$ set.

For each $\tau_i \in \mathcal{S}_\tau$, we choose the optimal tuning parameter, from a sequence of 100 $\lambda_{\tau_i}^{(j)}$ values. We determine the maximum value of this sequence as the smallest value of $\lambda_{\tau_i}^{(j)}$ for which all entries of $\hat{\psi}_{\tau_i}^{(j)}$ (i.e. the estimate of $\psi_{\tau_i}^{(j)}$) are equal to zero. By multiplying this maximum value by a factor of 0.0001, we obtain the minimum of the same sequence, at which the impact of LASSO is weaker, approaching the solution of a standard expectile regression without penalties. We then determine a decreasing sequence of $\lambda_{\tau_i}^{(j)}$ on the log scale from the maximum to the minimum. This procedure has been suggested by Friedman et al. (2010) and is the default choice of the ‘gcdnet’ R package. We compute the maximum of the $\lambda_{\tau_i}^{(j)}$ sequence without including the Auto-Regressive CARES components to make the algorithm computationally more efficient, without loss of accuracy. Indeed, we found that the resulting sequence fits our complete CARES-X model well: the estimates in $\hat{\psi}_{\tau_i}^{(j)}$ follow a regular path, starting from a point at which all loss exceedances are selected (i.e. LASSO has a null impact when $\lambda_{\tau_i}^{(j)}$ takes relatively low values), until the ℓ_1 -norm penalty saturates its effects, so that all coefficients in $\hat{\psi}_{\tau_i}^{(j)}$ are equal to zero.

We study the behaviour of our network over time by implementing a rolling window scheme, that we describe below. We divide our dataset into rolling sub-samples which have a constant dimension $T \times N$, where $T = 750 < B = 3620$, with a step of ten days. As a result, the first sub-sample includes the returns yielded by the N companies from the first to the 750th day, whereas the subsequent estimation windows span, respectively, the time intervals $[11, 760], \dots, [2871, 3620]$. We then obtain 288 rolling sub-samples and estimate a different network from each of them. Focusing on the first estimation window, for instance, we minimize for each $\tau_i \in \mathcal{S}_\tau$ the

following loss function:⁸

$$L\left(\boldsymbol{\beta}_{\tau_i}^{(j)}\right) = \frac{1}{T-1} \sum_{t=2}^T \left| \tau_i - \mathbb{I}_{\{r_{j,t} < \mu(r_{j,t}; \tau_i)\}} \right| \cdot [r_{j,t} - \mu(r_{j,t}; \tau_i)]^2 + \widehat{\lambda}_{\tau_i}^{(j)} \left\| \boldsymbol{\psi}_{\tau_i}^{(j)} \right\|_1, \quad (11)$$

where $\mu(r_{j,t}; \tau_i) = \alpha_{\tau_i}^{(j)} + \gamma_{1,\tau_i}^{(j)} \mu(r_{j,t-1}; \tau_i) + \gamma_{2,\tau_i}^{(j)} |r_{j,t-1}| + \boldsymbol{\psi}_{\tau_i}^{(j)} \mathbf{h}'_{-j,t} + \boldsymbol{\delta}_{\tau_i}^{(j)} \mathbf{w}'_t$ and $\boldsymbol{\beta}_{\tau_i}^{(j)} = [\alpha_{\tau_i}^{(j)}, \gamma_{1,\tau_i}^{(j)}, \gamma_{2,\tau_i}^{(j)}, \boldsymbol{\delta}_{\tau_i}^{(j)}, \boldsymbol{\psi}_{\tau_i}^{(j)}]$.⁹

We minimize the loss function defined in Equation (11) using the ‘BFGS’ method and with a maximum of 100000 iterations. Starting from $\tau_1 = 0.001$, we use the coefficients derived from the CARE model in Equation (4), estimated at the expectile level $\tau = 0.001$, as initial conditions for the parameters $\gamma_{1,\tau_1}^{(j)}$ and $\gamma_{2,\tau_1}^{(j)}$. In contrast, the initial conditions of the remaining parameters in Equation (11) are obtained from a standard expectile regression model, employing the generalized coordinate descent algorithm. When i in τ_i is greater than one (so that $\tau_i > 0.001$), we use the coefficients obtained for τ_{i-1} as initial conditions. This setup induces a faster convergence of the estimation algorithm.

We compute the in-sample violations for each pair $(\tau_i, \widehat{\lambda}_{\tau_i}^{(j)})$ in Equation (11), for $i = 1, \dots, 20$, and select two expectile levels:

$$\tau_{i,m}^{(j)} = \max \left(\tau_i \in \mathcal{S}_\tau : (T-1)^{-1} \sum_{t=2}^T \mathbb{I}_{\{r_{j,t} < \mu(r_{j,t}; \tau_i)\}} < \theta \right)$$

and

$$\tau_{i,M}^{(j)} = \min \left(\tau_i \in \mathcal{S}_\tau : (T-1)^{-1} \sum_{t=2}^T \mathbb{I}_{\{r_{j,t} < \mu(r_{j,t}; \tau_i)\}} > \theta \right),$$

which are related to the optimal tuning parameters $\widehat{\lambda}_{\tau_{i,m}}^{(j)}$ and $\widehat{\lambda}_{\tau_{i,M}}^{(j)}$, respectively. We then define a new grid of ten values of equally spaced τ in the interval $[\tau_{i,m}^{(j)}, \tau_{i,M}^{(j)}]$. For each of them, we re-estimate the parameters in

⁸To each $\tau_i \in \mathcal{S}_\tau$ corresponds a specific optimal tuning parameter $\widehat{\lambda}_{\tau_i}^{(j)}$, that we compute in a previous step from the full-sample data, using the cross-validation technique previously described.

⁹Computing an optimal tuning parameter for each of the 288 rolling estimation windows would be computationally expensive. We then prefer to select $\widehat{\lambda}_{\tau_i}^{(j)}$ for each expectile level τ_i from the full-sample data.

Equation (11), while the tuning parameter is computed from a linear interpolation between $\widehat{\lambda}_{\tau_{i,m}}^{(j)}$ and $\widehat{\lambda}_{\tau_{i,M}}^{(j)}$. We now use as initial conditions the coefficients obtained for $\tau_{i,M}^{(j)}$. With this two-step estimation approach, we select with a greater accuracy the optimal expectile level, minimizing the distance between the in-sample violations and θ .

After estimating the parameters in Equation (11), we focus on $\widehat{\psi}_{\tau_j^*}^{(j)}$, that is, the estimate of the $(N - 1)$ vector $\psi_{\tau_j^*}^{(j)}$. $\widehat{\psi}_{\tau_j^*}^{(j)}$ includes the estimated impact of $N - 1$ institutions being in a stressed state (expressed by their corresponding loss exceedances) on the ES of company j . Some of the elements in $\widehat{\psi}_{\tau_j^*}^{(j)}$ are exactly equal to zero, and this stems from the use of the LASSO penalty. Other coefficients in $\widehat{\psi}_{\tau_j^*}^{(j)}$ approach zero (but are not exactly equal to zero), pointing out weak or negligible effects; this is also due to the ℓ_1 -norm penalty in Equation (11), which leads to sparse solutions. We remind the reader that there exists a link from company i to company j (i.e. an edge in our network) if the coefficient in $\widehat{\psi}_{\tau_j^*}^{(j)}$ multiplying the loss exceedance of the former is, in absolute value, greater than a threshold η .

We need to determine a value of η that allows us to identify the relevant edges between the nodes of our network. We refer to those effects which stand out from the mass of null or negligible coefficients in $\widehat{\psi}_{\tau_j^*}^{(j)}$. For this purpose, we set $\eta^{(j)} = q\left(\left|\widehat{\psi}_{\tau_j^*}^{(j)}\right|; 0.75\right) + 2 \cdot IQR\left(\left|\widehat{\psi}_{\tau_j^*}^{(j)}\right|\right)$, where $q\left(\left|\widehat{\psi}_{\tau_j^*}^{(j)}\right|; 0.75\right)$ and $IQR\left(\left|\widehat{\psi}_{\tau_j^*}^{(j)}\right|\right)$ are, respectively, the third quartile and the interquartile range of the absolute values of the elements in $\widehat{\psi}_{\tau_j^*}^{(j)}$. Note that greater values of $\widehat{\lambda}_{\tau_i}^{(j)}$ in (11) would lead to sparser solutions, characterized by a smaller number of nonzero coefficients, which would be more isolated from a larger mass of zeros. Therefore, the resulting $\eta^{(j)}$ threshold would still pick the same relevant variables which are LASSO-selected with smaller values of $\widehat{\lambda}_{\tau_i}^{(j)}$.¹⁰

¹⁰Analyzing further the results, we found that the estimates of the $\eta^{(j)}$, for $j = 1, \dots, N$, exhibit a similar distribution across economic sectors of companies included in our dataset, which highlights the fact that the chosen threshold specification does not induce a specific bias in the network construction. We are here indebted to an anonymous referee for his/her demand of robustness checks and related suggestions.

We stress that $\eta^{(j)}$ is specific for the j^{th} response company and for a given rolling sub-sample; therefore, it takes different values for $j = 1, \dots, N$. Moreover, the $\eta^{(j)}$ values are also time dependent.¹¹ In contrast to Hautsch et al. (2014), we then do not fix *a priori* the value of $\eta^{(j)}$ to a constant value of 0.0001. By doing so, we calibrate $\eta^{(j)}$ to the specific features of the data we use to estimate each expectile regression model for each sub-sample of our rolling window scheme. Indeed, the magnitude of the estimated coefficients changes from periods of market stability to periods of financial turmoil. Therefore, a constant threshold $\eta^{(j)} = 0.0001$ might be misleading to identify the true, relevant links, which trigger contagion effects within our network during a specific market regime. We report an example in Figure 1, where we display the absolute values of the coefficients in $\hat{\psi}_{\tau_j^*}^{(j)}$ that we obtain when estimating the ES of Crdit Agricole conditional on the loss exceedances of the other companies in our dataset, by adopting the daily returns observed during the time interval $[1, 750]$; that is, the first rolling sample. We also draw two thresholds: 0.00034 (in red), computed as $\eta^{(j)} = q\left(\left|\hat{\psi}_{\tau_j^*}^{(j)}\right|; 0.75\right) + 2 \cdot IQR\left(\left|\hat{\psi}_{\tau_j^*}^{(j)}\right|\right)$, and 0.0001 (in blue). We see that only the former is able to identify links clearly different from zero.

We use $\hat{\psi}_{\tau_j^*}^{(j)}$ and $\eta^{(j)}$ to define the j^{th} row of the adjacency matrix $\mathbf{A} = [a_{j,i}] \in \mathbb{R}^{N \times N}$, which determines the structure of the network, for $j = 1, \dots, N$. We focus on unweighted networks. Therefore, the element $a_{j,i}$ takes the value of one, if the loss exceedance of the i^{th} company is LASSO-selected as relevant regressor—i.e. its corresponding coefficient in $\hat{\psi}_{\tau_j^*}^{(j)}$ is, in absolute value, greater than $\eta^{(j)}$ —to explain the ES of the j^{th} institution. We exclude self-loops; that is, $a_{j,i} = 0$ if $j = i$. Moreover, our network is directed and then $a_{j,i} \neq a_{i,j}$, for $i, j = 1, \dots, N$ and $i \neq j$.

The discussions and methods presented above apply to the network estimated from the first sub-sample of our rolling window procedure, which spans the time interval $[1, 750]$. We iteratively use the same methods to obtain the subsequent networks estimated, respectively, from the time intervals $[11, 760], \dots, [2871, 3620]$, so that we use all of the available data. As a result, we obtain 288 dynamic networks overall.

We conclude this section providing some details about the computational times required by our algorithm, using an Intel®Core i7-4710HQ@2.50GHz

¹¹The time index has been suppressed to simplify the notation.

Figure (1) Absolute values of the coefficients in $\hat{\psi}_{\tau_j^*}^{(j)}$ that we obtain when estimating the ES of Crédit Agricole conditional on the loss exceedances of the other companies in our dataset, by adopting the daily returns observed during the time interval $[1, 750]$. The red and the blue lines represent the thresholds 0.00034 and 0.00010, respectively.

(64-bit operating system) computer with six cores. The cross-validation exercise implemented on the 20 expectile levels $\tau_i \in \mathcal{S}_\tau$ requires four hours and 15 minutes, on average, for six of the $N = 104$ companies included in our dataset. In contrast, the estimation of the CARES-X model in our rolling window scheme requires, on average, 51 seconds for six companies and a single (rolling) subset. Therefore, the cross-validation exercise is more expensive in terms of computational costs. However, we can adopt different solutions to make the cross-validation development more efficient. First, we can record significant improvements by employing a sufficiently large number of cores for parallel computations. Second, we found that the optimal value of $\lambda_{\tau_i}^{(j)}$ that we compute by means of cross-validation does not sensibly change across the 20 entries of $\mathcal{S}_\tau$, as this set spans the relatively narrow interval $[0.001, 0.05]$. As a result, we could significantly reduce the computational costs by employing a lower cardinality of $\mathcal{S}_\tau$, keeping the extreme values of the interval $[0.001, 0.05]$, while incurring a minor accuracy loss. Finally, we stress that this exercise is implemented only once using the full-sample data. Therefore, it does not lead to strong concerns for real applications based on our empirical setup. In fact, we only need to update the CARES-X coefficients every 10 days, following our rolling window scheme, and this procedure is less expensive in terms of computational burden, as mentioned above.

4. Dynamic analysis of the large financial network: some evidence

The rolling window procedure described in Section 3 allows us to recover and analyze the dynamic behavior of our network time series. Indeed, our dataset spans the time interval January 2, 2004—November 11, 2017, which is characterized by different market regimes. By exploiting the informative content of our rolling analysis, we check whether and to what extent our dynamic network changes when moving from either a normal or bull market condition to a stressed state. The latter acquires a central importance in this study, because our method focuses on the relationships between different companies when they are all in distress.

Figure 2 displays estimated networks corresponding to various illustrative sub-samples: i) before crisis (January 2, 2004—November 16, 2006); ii) during crisis (January 11, 2008—November 25, 2010); iii) post-crises (March 29, 2013—February 11, 2016) and iv) the most recent bullish period in the studied sample (January 2, 2015—November 16, 2017). Figure 3 exhibits trends of the network density and reciprocity in the whole sample (January 2, 2004—November 11, 2017).¹² Figure 4 displays the corresponding vertex degree evolution, split by economic sectors (banking, insurance, financial services and real estate), over the entire sample (January 2, 2004—November 11, 2017). Figures 5 and 6 distinguish in- from out-degrees.¹³

Figure 2(a) displays the network estimated using the data from January 2, 2004 to November 16, 2006 (midpoint: June 10, 2005); this is the first network derived from the rolling window scheme. It thus reflects the relationships among the $N = 104$ companies before the outbreak of the US sub-prime crisis, whose reference event was the default of Lehman Brothers in September, 2008. By clustering the $N = 104$ into four sectors, we observe a greater concentration of links among real estate and insurance companies, whereas banks are rather isolated.

Figure 2(b) displays the results we obtain from the 106th rolling estimation window, which spans the period January 11, 2008—November 25, 2010 (midpoint: June 19, 2009). Therefore, the resulting network captures the impact of two relevant events: the US subprime crisis (period 2007—2009) and the Greek government-debt crisis, which started in the fall of 2009. The latter triggered the subsequent distress of other Eurozone countries, such as Ireland, Italy, Portugal and Spain (period 2010—2011). The relevant risks and losses suffered during this time interval are well depicted in Figure 2(b):

¹²We define the network density and reciprocity in Equations D.3 and D.4.

¹³We define the in- and out-degrees in Equations D.1 and D.2.

Figure (2) Networks estimated from different periods. Nodes are clustered according to the economic sector of the corresponding companies: orange (banking), light blue (insurance), yellow (financial services) and green (real estate).

the network becomes denser, highlighting a larger number of links among $N = 104$ companies in distress. This phenomenon is also evident in Figure 3, which displays the trend of network density over time. The density significantly increases from the second half of 2007, remaining at high values throughout the European debt crisis. We observe the peak of the network density in the 125th sub-sample of our rolling window procedure, which spans the interval October 3, 2008—August 18, 2011 (midpoint: March 12, 2010).

Figure (3) Trend of the network density (black solid line) and reciprocity (red dashed line) from January 2, 2004 to November 16, 2017.

Figure (4) Trend of the average vertex degree of each sector from January 2, 2004 to November 16, 2017.

It is also interesting to observe from Figure 3 that the proportion of mutual connections (i.e. the reciprocity) significantly decreases in April, 2007, to reach its minimum in August, 2007, taking low values throughout the entire crisis period. This evidence emphasizes the systemic relevance during a crisis period of a cluster of institutions, which impact the ES of other companies and are net contributors of risk (i.e. the incoming edges of risk (i.e. the incoming edges of these companies are smaller than their outgoing edges). This applies even more specifically to banks.

By comparing Figures 2(a) and 2(b), we observe how banks acquire a greater relevance, in terms of a larger concentration of links, during the crisis period. We also note from Figure 4 a significant growth of the average vertex degree of the banking sector at the beginning of the US subprime

Figure (5) Trend of the average vertex in-degree of each sector from January 2, 2004 to November 16, 2017.

Figure (6) Trend of the average vertex out-degree of each sector from January 2, 2004 to November 16, 2017.

crisis. However, we do not observe significant differences between the banking sector and the other sectors in terms of vertex in-degree (Figure 5). In contrast, the banking sector emerges when analyzing the vertex out-degree (Figure 6). This result emphasizes the significant impact of banks on the stability of other institutions and, then, on the entire financial system during tail events. This highlighted evidence is consistent with the findings of Billio et al. (2012), namely: banks appear to be more systemically relevant than other financial industry groups.

At the end of the crisis period analyzed above, our network again becomes sparser, as we see from Figures 2(c) and 3. This evidence points out the beginning of a new market phase, characterized by weaker relationships in terms of extreme risk, resulting in a greater stability of the system. Inter-

estingly, we observe relevant changes in terms of reciprocity as well, which records a significant growth from the 134th rolling sub-sample, the midpoint of which is July, 2010 (see Figure 3). The banking sector still records high values of vertex degree, on average (see Figure 4). The active impact of banks, quantified by the vertex out-degree is lower than before, on average, but still high compared to the other sectors, with a stable trend until November, 2017 (see Figure 6). The distance between the banking sector and the other companies in our network is more evident than before when looking at the vertex in-degree, which has a stable trend from the outbreak of the US subprime crisis to November, 2017 (see Figure 5). Therefore, banks appear now to be more exposed to the risk of suffering a negative impact from other nodes.

The network again becomes denser in the final part of our dynamic analysis (see Figure 2(d)), although not at the levels recorded during the years 2007—2011, whereas its reciprocity decreases (see Figure 3). We link these to the turmoil experienced during two important events: i) the complicated negotiations between the Greek government and its international creditors in the summer of 2015, which threatened a Greek default along with a potential exit from the monetary union; and ii) the Brexit referendum in June 2016, when the UK voted to leave the European Union. The banking and real estate sectors record the greatest portion of stressed connections (see Figures 2(d) and 4). This phenomenon is more evident in terms of in-going links (see Figure 5), emphasizing a larger risk of being affected by other institutions. By analyzing the out-going links, we also highlight the impact of the insurance sector, in addition to banks and real estate companies (see Figure 6). Summing up the results, we see that the local and global statistics analyzed above provide evidence that our method is able to capture the relationships among financial institutions in pronounced distress during tail events. As a consequence, our methodology may be a useful tool in identifying companies which have a relevant impact on the financial system during crisis periods (as detected, for instance, in Figure 8).

5. Implications for portfolio risk and performance

We now study the relevance of our method within a portfolio framework. For this purpose, we focus on portfolios including the assets of the N companies we consider (i.e. the nodes of our network). The CARES model allows us to quantify the systemic relevance of N financial institutions in the occurrence of tail events, but does not optimize a specific objective function

leading to optimal portfolio weights. However, our network provides information which potentially affects the portfolio performance and risk during a turmoil or during/after systemic events. In fact, our approach leads to the identification of two clusters of companies. The first one, denoted as RT (which stands for Right Tail), includes the systemically important institutions, which have a significant impact on a large number of nodes in our network. The second one, denoted as LT (which stands for Left Tail), includes the institutions which have the lowest impact in our network. In our study, the impact of the i^{th} company on the entire network is equal to the out-degree OD_i defined in Equation (D.2), for $i = 1, \dots, N$.

For each of the 288 sub-samples derived from our rolling window procedure, we identify the two groups LT and RT of companies by computing the 25th and 75th percentiles of the values in $\mathbf{OD} = [OD_1 \ OD_2 \ \dots \ OD_N]$. The LT set includes companies which have an impact lower than or equal to the 25th percentile of $\mathbf{OD}$. In contrast, the RT set regroups institutions which have an impact greater than or equal to the 75th percentile of $\mathbf{OD}$. We assess how the composition of the LT and RT clusters evolves over time in Figure 7. The companies listed on the left side of Figure 7 are sorted by country. The companies belonging to LT are depicted in green, whereas the ones belonging to RT are depicted in red. As explained in Section 3, we estimate 288 networks by implementing a rolling window procedure with a sample size of 750 days and a step of 10 days ahead. For improving readability, we display, in Figure 7, the configuration of LT and RT derived from 144 of these 288 networks, equally spaced between January, 2004 and November, 2017. Each date represented on the top of the graph is the midpoint of the sub-sample from which we obtain the corresponding network.

We note that the composition of LT and RT groups is quite stable over time. Furthermore, we also remark that companies belonging to the same country often have a similar role. For instance, we can refer to some of the Italian and Spanish institutions, which have belonged to the RT set almost continuously since 2010. The same is true for some of the British and French companies, which have belonged to the LT group almost continuously during the same time interval. On the one hand, investing in systemically important institutions exposes our portfolio to the risk they face in a stressed state, for instance during tail events, impacting the stability of many other companies. In fact, the existence of interconnections in our network exposes single companies to the distress of other companies, potentially increasing the magnitude of contagion at a portfolio level, with an overall greater risk. On the other hand, investing in systemically important institutions may have some advantages during bull market periods, when their good performance

has a positive impact on the stability of the system. In contrast, the companies belonging to the LT set would have a weak impact on the extreme risk of the overall portfolio during tail events, but also a lower positive contribution during bull market phases.

Building on the different behaviors of LT and RT labelled companies during bull and bear markets, we now try to understand how the identification of these two clusters has an impact in terms of portfolio risk. For this purpose, we compare five different naïve portfolio strategies. The first one is the equally-weighted portfolio (denoted EW), which is commonly used as a benchmark in the financial literature because of its remarkable performance and diversification properties; see DeMiguel et al. (2009), Duchin and Levy (2009) and Tu and Zhou (2011). In our context, we would say that EW represents the optimal solution when we believe that all the N companies in distress have the same impact on the ES of the other nodes in our network. Therefore, the N assets would take the same weight, as they provide the same contribution to the overall portfolio risk. The second strategy, denoted LL (which stands for Long on LT), selects the assets of the companies belonging to LT only, weighing them equally into a global portfolio. In contrast, the weights of the stocks not included in LT are set equal to zero to satisfy the budget constraint. This corresponds to a strategy aimed at largely reducing the possible impact of systemic events on portfolio risk. LL should be defensive during phases of market turmoil, but might be penalized in terms of performance during a bullish period.

The third investment rule (LR, i.e. Long on RT) is implemented by equally weighing the stocks of the companies included in RT, whereas the other assets take a null weight. Opposite to LL, investing in highly interconnected companies is penalized during stressed phases and aggressive in bullish periods. The fourth reference strategy (denoted LL-SR, which stands for Long Left-Short Right) opposes two different portfolios, both defined above. In particular, LL-SR takes a long position on the portfolio we obtain by implementing LL and a short position on the portfolio derived from LR. Regarding the fifth strategy (SL-LR, i.e. Short Left-Long Right), it has an opposite behavior to LL-SR. Indeed, SL-LR takes a short position in the portfolio we obtain by implementing LL and a long position in the portfolio derived from LR. As a result, LL-SR and SL-LR might be seen as a kind of (imperfect) market neutral rules, playing with the differences in risk and return between two opposite clusters of companies' stocks.

Above all, we should note here that the resulting competing portfolios for the five strategies do not directly rely on the optimization of any specific explicit objective function (for instance, a risk measure, such as the

portfolio variance, VaR or ES) and, thus, do not provide optimal weights related to any program linked to financial characteristics of the final portfolio. Consequently, on one side, the differences in their performance exclusively depend on the information retrieved from our network. On the other side, our results are not exposed to any direct source of model risk related to portfolio weights (Michaud, 1989). Moreover, naïve allocation strategies might anyway be comparable to optimized allocations; see, for instance, DeMiguel et al. (2009). Also, on an *a priori* ground, we should note that, when clustering the companies into either LT or RT, we are exposed to the risk that assets with similar features (belonging to the same group) respond in the same way to tail events, with some under-diversification potential effects. In contrast, LL-SR (and also SL-LR) includes a larger set of assets having an opposite role in our network. Moreover, the possibility of setting both long and short positions (with either LL-SR or SL-LR) allows us to offset the different reactions of LT and RT companies to tail events.

We now describe the procedure we use to build each of the competing portfolios introduced above. For each of them, and for each of the 288 sub-samples derived from the rolling window procedure described in Section 3, we compute a $1 \times N$ vector of portfolio weights $\boldsymbol{\omega}_t = [\omega_{1,t} \ \omega_{2,t} \ \cdots \ \omega_{N,t}]$, where $\omega_{i,t}$ is the weight of the i^{th} asset that we determine at time t for a given investment strategy, for $i = 1, \dots, N$ and $t = 750, 760, \dots, 3610, 3620$. For instance, except for the EW strategy (where the weights are constant at the level $1/N$), the first vector of portfolio weights (i.e. $\boldsymbol{\omega}_{750}$) is retrieved from the first network, that we estimate at time $t = 750$, using the daily returns spanning the time interval $[1, 750]$. We then compute 750 portfolio returns as $r_{p,t} = \boldsymbol{\omega}_{750} \mathbf{r}'_t$, where $\mathbf{r}_t = [r_{1,t} \ r_{2,t} \ \cdots \ r_{N,t}]$, for $t = 1, 2, \dots, 750$. Therefore, these 750 portfolio returns correspond to the output of our first network. By doing so, we do not forecast optimal portfolios for a time horizon that goes beyond the 750th day. This is then an in-sample evaluation, as the primary focus of this study is to understand whether, and to what extent, the identification of the clusters of LT and RT companies (that we exclusively determine from the first rolling network) has an impact in terms of portfolio risk and performance. We do not forecast the composition of LT and RT after the 750th day. In contrast, we observe the future compositions of LT and RT sets of companies from the subsequent rolling sub-samples and, from time to time, we try to understand if they have an impact upon the performance of the resulting portfolios.

We compute from the in-sample portfolio returns $r_{p,1}, \dots, r_{p,750}$, a set of risk and performance measures, which are described as follows: first, we compute the Expected Shortfall (ES) because it has a central role in

our study. The tail expectation of R_j defined in Equation (1) typically takes negative values when setting $\theta = 0.05$. We use in this section the tail expectation with a negative sign to assess the competing portfolios, so that a greater and positive ES points out a greater risk. Secondly, we compute the Standard Deviation (SD) of the 750 portfolio returns to have a more complete picture about the portfolio risk. Third, we also compute the Sharpe Ratio (SR) and the ratio between the mean and the ES of the 750 portfolio returns (ESR) to provide information about the risk-adjusted performance. We collect these statistics in $\mathbf{D}_1 = [ES_1 \ SD_1 \ SR_1 \ ESR_1]$, where the subscript ‘1’ points out the fact that they are computed from the first network.

We then implement the rolling window procedure with a horizon of ten days ahead. As a result, we estimate the second network and, consequently, rebalance the portfolio weights at time $t = 760$, computing additional 750 portfolio returns: $r_{p,t} = \boldsymbol{\omega}_{760} \mathbf{r}_t'$ for $t = 11, 12, \dots, 760$. Again, we summarize these new portfolio returns using the descriptive statistics described above and obtain $\mathbf{D}_2 = [ES_2 \ SD_2 \ SR_2 \ ESR_2]$. We repeat this procedure for the subsequent 286 sub-samples using a step of a 10-days-ahead grid, until we use all of the available data, obtaining the corresponding vectors $\mathbf{D}_3 = [ES_3 \ SD_3 \ SR_3 \ ESR_3]$ to $\mathbf{D}_{288} = [ES_{288} \ SD_{288} \ SR_{288} \ ESR_{288}]$. Note that the vectors $\mathbf{D}_1, \dots, \mathbf{D}_{288}$ are dynamically computed over time from different networks, which, in turn, reflect different market regimes. However, our method aims at estimating the impact of the loss exceedances of company i on the ES of company j , $\forall i \neq j$. Therefore, an analysis which focuses on a crisis period (in place of the full-sample analysis) would acquire a greater relevance in our study, because it captures the relationships among different institutions when they are in a state of joint distress.

We saw in Section 4 that the networks estimated from the rolling sub-samples numbered 60 to 133 reflect the strongest degree of distress in our analysis, due to the effects of two relevant events: the US subprime crisis and the European sovereign debt crisis. The impact of these events is evident by estimating the univariate ESs of the $N = 104$ companies included in our dataset, for each of the 288 rolling sub-samples. Figure 8 displays the average ES for each sub-sample.¹⁴ We can see from Figure 8 that the period including the sub-samples numbered 60 to 133 is characterized by the strongest risk.

¹⁴For each company and for each sub-sample, we estimate the ES at the 5% level by employing the historical simulation approach (see, among others, Hull, 2011).

Figure (8) Average Expected Shortfall (%) at the 5% level of the $N = 104$ companies included in our dataset, computed for each of the 288 sub-samples resulting from the rolling window procedure described in Section 3. The crisis period includes the sub-samples from 60 (midpoint: September 2007) to 133 (midpoint: July 2010). The horizontal dashed lines represent the average Expected Shortfall before, during and after the crisis period.

The 60th sub-sample spans the interval April, 2006—February, 2009 (midpoint: September, 2007). The 133th sub-sample includes the interval January, 2009—December, 2011 (midpoint: July, 2010). We then select as a crisis period the interval which begins with the 60th sub-sample and ends with the 133th sub-sample. As a result, we use the statistics included in the 74×4 matrix such as: $\mathbf{D}^{Crisis} = [\mathbf{D}_{60}; \mathbf{D}_{61}; \dots; \mathbf{D}_{133}]$, to evaluate the performance of the competing investment strategies during the crisis period. We check whether and to what extent the results change when focusing on the entire period January, 2004—November, 2017, which is characterized by an alternation of different market regimes. In addition, we also study the performance of the strategies focusing on a stable period. We see from Figure 8 that the period before the US subprime crisis is characterized, on average, by a lower risk.

To confirm the crisis and stable period sample choices, we run a simple three-regime switching regression on the indicators displayed in Figures 3 and 8, checking if we find evidence of switches in the mean of these variables. We find that for all the three indicators a regime is associated with the crisis period, with a start ranging from mid-May to mid-July, 2007, and the end between mid-June and mid-October, 2010, thus coherent with our choice for the crisis period. For the stability period, we observe more stable behavior in the regimes before the crisis than after the crisis, in line with our choice of

using the pre-crisis period as a stable period.¹⁵ This evidence also emerges from the network analysis reported in Section 4. In this regard, we found that the average ES displayed in Figure 8 is highly correlated with the density and reciprocity of our dynamic networks (depicted in Figure 3), as well as with the Country-Level Index of Financial Stress (CLIFS) and the Composite Indicator of Systemic Stress (CISS) provided by the European Central Bank (see Figure C.10 in Appendix C).

Therefore, we select as a stable period the interval which begins with the first sub-sample and ends with the 59th sub-sample, taking into account the statistics included in the 59×4 matrix such as: $\mathbf{D}^{Stable} = [\mathbf{D}_1; \mathbf{D}_2; \dots; \mathbf{D}_{59}]$. Finally, to contrast with stable and crisis periods, we also define a 288×4 matrix such as: $\mathbf{D}^{Full Sample} = [\mathbf{D}_1; \mathbf{D}_2; \dots; \mathbf{D}_{288}]$, which reports thus the full-sample results. We hereafter analyze each series (by column) of $\mathbf{D}^{Crisis}$, $\mathbf{D}^{Full Sample}$ and $\mathbf{D}^{Stable}$ using the following descriptive statistics: first quartile (Q1), median (MED), mean (MEAN) and third quartile (Q3). We report the results obtained from the crisis period in the left panel of Table 1.

The LL strategy outperforms EW in terms of ES, whereas LR turns out to be the riskiest strategy. Therefore, investing in the companies with the largest (weakest) impact on our network implies a greater (lower) risk. It is interesting to observe that this result holds for all the statistics reported in the left panel of Table 1(a). We motivate this result as follows: the companies belonging to RT have a relevant impact on the ES of other institutions. However, when implementing LR, we invest in the companies with the largest impact only, triggering the propagation of contagions within the portfolio. ES significantly decreases when using the strategies LL-SR and SL-LR. Therefore, the possibility of combining long and short positions within the same portfolio to offset the effects of opposite clusters leads to relevant improvements. The LL-SR outperforms the SL-LR, confirming the evidence that a lower risk is associated with long positions in companies with a weaker impact on our network. We obtain similar findings when analyzing the standard deviation (Table 1(b)). Again, LR provides the greatest risk, whereas LL outperforms both EW and LR strategies. Once again, the risk sensibly decreases when implementing LL-SR and SL-LR, which have the same performance. This is due to the fact that the variance of $-aX$, where a is a scalar and X is a given random variable, is equal to $var(-aX) = a^2 var(X) = var(aX)$. In other words, we lose the information

¹⁵The Markov Switching regressions are available upon request.

Table (1) Portfolio performance and risk during different periods

RULE	Q1	MED	MEAN	Q3	Q1	MED	MEAN	Q3	Q1	MED	MEAN	Q3
	CRISIS PERIOD				FULL-SAMPLE				STABLE PERIOD			
(a) ES (%)												
EW	5.62	5.65	5.54	5.68	2.72	3.42	3.65	4.77	1.85	2.43	2.69	2.93
LL	4.89	5.12	5.08	5.28	2.67	3.37	3.52	4.30	1.83	2.45	2.62	3.03
LR	6.17	6.42	6.36	6.68	2.85	3.61	4.03	5.49	1.99	2.53	2.89	3.05
LL-SR	1.69	1.88	1.90	2.12	0.99	1.23	1.34	1.62	0.91	1.00	1.04	1.15
SL-LR	2.01	2.31	2.29	2.56	1.01	1.27	1.47	1.89	0.89	0.98	1.09	1.07
(b) SD (%)												
EW	2.17	2.23	2.20	2.26	1.13	1.43	1.50	1.92	0.74	0.94	1.07	1.22
LL	1.95	2.06	2.04	2.12	1.15	1.40	1.46	1.75	0.77	0.97	1.07	1.26
LR	2.44	2.54	2.52	2.62	1.20	1.57	1.67	2.18	0.80	0.97	1.14	1.28
LL-SR	0.82	0.92	0.91	1.01	0.46	0.58	0.63	0.76	0.42	0.46	0.48	0.51
SL-LR	0.82	0.92	0.91	1.01	0.46	0.58	0.63	0.76	0.42	0.46	0.48	0.51
(c) SR (%)												
EW	-4.30	-3.59	-3.11	-1.79	-1.55	0.74	1.36	3.73	-0.99	4.22	4.38	11.44
LL	-3.64	-3.05	-2.52	-1.67	-1.51	1.84	2.04	3.66	-1.59	3.61	3.63	10.07
LR	-4.82	-4.17	-3.79	-2.25	-2.71	-0.86	0.20	3.03	-0.88	3.52	4.19	11.35
LL-SR	3.26	4.45	4.79	6.10	2.95	5.44	4.90	7.63	-3.65	-2.51	-0.77	-0.91
SL-LR	-6.10	-4.45	-4.79	-3.26	-7.63	-5.44	-4.90	-2.95	0.91	2.51	0.77	3.65
(d) ESR (%)												
EW	-1.69	-1.44	-1.23	-0.71	-0.63	0.30	0.61	1.59	-0.41	1.64	1.76	4.58
LL	-1.47	-1.24	-1.00	-0.68	-0.63	0.78	0.91	1.52	-0.68	1.43	1.51	4.20
LR	-1.92	-1.66	-1.48	-0.88	-1.18	-0.37	0.12	1.23	-0.37	1.36	1.70	4.57
LL-SR	1.55	2.18	2.32	3.06	1.44	2.58	2.34	3.62	-1.64	-1.10	-0.21	-0.38
SL-LR	-2.36	-1.85	-1.88	-1.32	-3.42	-2.29	-2.12	-1.25	0.43	1.20	0.54	1.72

Notes: From left to right, the table reports the following statistics: first quartile (Q1), median (MED), mean (MEAN) and third quartile (Q3), for each column of the matrices $\mathbf{D}^{Crisis}$ (left panel), $\mathbf{D}^{Full Sample}$ (central panel) and $\mathbf{D}^{Stable}$ (right panel), respectively. The columns of each matrix $\mathbf{D}^{Crisis}$, $\mathbf{D}^{Full Sample}$ and $\mathbf{D}^{Stable}$ report the series of different risk and performance indicators, that we dynamically compute for each investment strategy listed in the first column. Panel (a) reports the statistics computed on the first columns of $\mathbf{D}^{Crisis}$, $\mathbf{D}^{Full Sample}$ and $\mathbf{D}^{Stable}$, which include the series of the portfolio ESs at the 5% level. Panels (b) and (c) report the statistics computed, respectively, on the second (portfolio standard deviations) and third (portfolio Sharpe ratios) columns of $\mathbf{D}^{Crisis}$, $\mathbf{D}^{Full Sample}$ and $\mathbf{D}^{Stable}$. Panel (d) reports the statistics computed on the fourth columns of $\mathbf{D}^{Crisis}$, $\mathbf{D}^{Full Sample}$ and $\mathbf{D}^{Stable}$, which include the ratios between the average portfolio returns and ESs. EW is the equally weighted portfolio of the entire set of (N) companies. LL is the equally weighted portfolio of the companies belonging to the cluster LT. LR is the equally weighted portfolio of the companies belonging to the cluster RT. LL-SR is the strategy which takes a long position on LL and a short position on LR. SL-LR is the strategy which takes a short position on LL and a long position on LR.

about the sign of a position when computing the variance. From the left panel of Table 1(c)-(d), we can see the strategy LL-SR dominates the other investment strategies in terms of risk-adjusted performance, quantified by two different indicators: the Sharpe ratio and the ratio between the mean portfolio return and ES. This result is due to two different factors. First, the best performance of LL-SR in terms of risk (decreasing the denominators of the SR and ESR ratios). Second, LL-SR only yields, on average, a positive portfolio return.

We now check whether, and to what extent, the results change when extending the focus to the entire period January, 2004—November, 2017. Results are reported in Table 1 (central panel): LL outperforms EW and LR in terms of ES (Table 1(a)) and Standard Deviation (Table 1(b)) as before. However, the differences between LL and EW (or LR) strategies are now less evident with respect to the crisis period, as we check by comparing the left and central panels of Table 1. Likewise, the differences between LL-SR and SL-LR are less evident in the central panel of Table 1(a). By comparing the left and central panels of Table 1(a)-(b), we highlight the fact that the risk sensibly decreases when moving to the full-sample analysis for all the portfolio strategies. This result points out the fact that a longer time interval, which includes bull and bear market phases, partially mitigates the effects of the crisis period. Again, LL-SR outperforms the other investment strategies in terms of risk and risk-adjusted performance (central panel of Table 1(c)-(d)).

Finally, we analyze the results obtained from the stable period in the right panel of Table 1. The statistics reported in the right panel of Table 1(a)-(b) point out a lower risk with respect to the crisis and the full-sample results, reflecting a general stability of the market. The differences among the competing strategies in the right panel of Table 1(a)-(b) are even lower than the ones derived from the full sample analysis. These results support the evidence that the information derived from our networks has a little impact in terms of portfolio risk during stable phases of the market. Interestingly, all the investment strategies (except, however, the LL-SR strategy), provide, on average, a positive risk-adjusted performance (see the right panel of Table 1(c)-(d)). Therefore, we obtain opposite results when contrasting stable and crisis periods.

In addition to the in-sample analysis, we also set up an out-of-sample exercise for completeness, using the portfolio weights derived from each estimation window (in-sample estimates) and evaluating their fit with the out-of-sample returns of the N companies. In particular, for each estimation window ending at time t , we compute ten out-of-sample portfolio returns

as $r_{p,t+h} = \omega_t \mathbf{r}'_{t+h}$, for $h = 1, \dots, 10$. Following our rolling window scheme, we update the portfolio weights every ten days, when a new network is estimated, until we use all of the available data. As a result, for each strategy, we obtain a vector of 2870 out-of-sample returns:

$$\mathbf{r}_p^{Full\ Sample} = \left[\underbrace{r_{p,751} \cdots r_{p,760}}_{using\ \omega_{750}} \underbrace{r_{p,761} \cdots r_{p,770}}_{using\ \omega_{760}} \cdots \underbrace{r_{p,3611} \cdots r_{p,3620}}_{using\ \omega_{3610}} \right],$$

that we summarize by computing the statistics ES, SD, SR and ESR.

Similar to the in-sample analysis, and in addition to the full-sample results, we evaluate the performance of the competing strategies during the crisis and the stable periods that we defined above. For this purpose, we also calculate ES, SD, SR and ESR on the vectors $\mathbf{r}_p^{Crisis} = [r_{p,971} \cdots r_{p,1710}]$ and $\mathbf{r}_p^{Stable} = [r_{p,751} \cdots r_{p,970}]$, which correspond to the time intervals September, 2007—July, 2010 and November, 2006—September, 2007, respectively. We report the out-of-sample results in Table 2. Some of the results obtained from the in-sample analysis also hold out of sample. First, LL outperforms EW and LR in terms of ES during the crisis period. Second, ES and SD significantly decrease when using LL-SR and SL-LR. Third, the values of ES and SD increase when moving from the stable to the crisis period for all the strategies, indicating a greater risk. However, other results change out of sample. For instance, SL-LR outperforms the other strategies in terms of risk and risk-adjusted performance, being the only rule which provides positive values of SR and ESR. All the investment strategies, except SL-LR, record, on average, negative portfolio returns during the stable period, in contrast to the tendency we observed in sample. This result may be due to the fact that the series of out-of-sample portfolio returns in $\mathbf{r}_p^{Stable}$ begins in November 2006, whereas the first 750 in-sample observations, which lead to a greater stability, are now excluded, being used only in sample. Therefore, the stable period identified above in the in-sample analysis becomes shorter from an out-of-sample perspective, suffering more from the initial fears related to the US subprime crisis.

As a conclusion, the results presented in this section highlight that it is possible to derive important information from our network, allowing us to build financial portfolios which provide not only a lower risk, but also a better risk-adjusted performance. In particular, these improvements become more pronounced during tail events. Of course, the relevance of our analyses might be further highlighted by moving from naïve portfolio allocations to optimized or regularized portfolio strategies. In this respect, our proposal

Table (2) Out-of-sample results

RULE	ES	SD	SR	ESR	ES	SD	SR	ESR	ES	SD	SR	ESR
	CRISIS PERIOD				FULL-SAMPLE				STABLE PERIOD			
EW	5.77	2.33	-3.56	-1.44	3.99	1.61	-0.65	-0.26	2.84	1.10	-1.96	-0.76
LL	5.64	2.35	-4.24	-1.76	4.02	1.61	-0.77	-0.31	2.81	1.10	-3.65	-1.42
LR	6.23	2.53	-2.31	-0.94	4.25	1.74	-0.66	-0.27	2.97	1.13	-1.57	-0.60
LL-SR	2.29	1.01	-4.07	-1.79	1.57	0.68	-0.13	-0.05	1.10	0.47	-4.70	-2.02
SL-LR	2.28	1.01	4.07	1.80	1.48	0.68	0.13	0.06	0.96	0.47	4.70	2.31

Notes: From left to right, the table reports the following statistics: ES (%) at the 5% level, standard deviation (%), Sharpe ratio (%) and ratio between the average portfolio return and ES (%). These statistics are computed on the out-of-sample portfolio returns obtained, respectively, from the crisis period (September, 2007—July, 2010; left panel), the full-sample (November, 2006—November, 2017; central panel) and the stable period (November, 2006—September, 2007; right panel). EW is the equally weighted portfolio on the entire set of (N) companies. LL is the equally weighted portfolio on the companies belonging to the cluster LT. LR is the equally weighted portfolio of the companies belonging to the cluster RT. LL-SR is the strategy which takes a long position on LL and a short position on LR. SL-LR is the strategy which takes a short position on LL and a long position on LR.

might be combined with the recent findings of Giuzio and Paterlini (2019) and Bonaccolto and Paterlini (2020) to verify the benefits of diversification during distress periods.

6. Conclusion, discussion and future researches

By combining the CARE model of Taylor (2008) with the intuition put forward by Hautsch et al. (2014) for the construction of financial networks, we present a methodology for estimating the extreme risk network of financial companies quoted in the US market. As a further methodological contribution, we introduce a calibration criterion for detecting the relevant edges between the network nodes. Our empirical evidence shows the appropriateness of the proposed approach, with network evolution characterized by an increase in density and vertex degrees during periods of market distress. When exploiting the informative content of the networks within a portfolio allocation framework, we show it is possible to design risk mitigation strategies which do not impact upon portfolio profitability. Our empirical results and methodological contributions might thus be of interest for both investors and market regulators.

Our study builds on a penalized expectile regression, in which the ℓ_1 -norm penalty applied to the parameters multiplying the loss exceedances

allows us to identify the most relevant links in the resulting network. However, alternative penalty functions can be adapted to our methodological proposal. For instance, it would be interesting to examine how the network structure changes when moving from convex to non-convex penalty functions, such as the Smoothly Clipped Absolute Deviation (SCAD) as introduced by Fan and Li (2001). The properties of SCAD applied to expectile regression problems have been studied by Liao et al. (2019). It would be interesting to compare LASSO with different penalty functions not only when estimating unweighted but also weighted networks, in which we assess the intensity of the links within a large set of companies in distress. Moreover, in this study, we generalize the method introduced by Taylor (2008), by focusing on large networks in which the relevant links, which reflect a state of joint distress of the analyzed companies, are identified through the LASSO method. It would be interesting to estimate our CARES-X model by employing a different methodological approach, building, for instance, on the recent contributions of Taylor (2019) and Wang and Gerlach (2019). By following the cited approaches, we could jointly estimate VaR and ES without the requirement of the grid search of the expectile level parameter τ , with relevant advantages in terms of computational efficiency. We could think about using the LASSO method in Taylor’s (2019) Asymmetric Laplace joint loss, which belongs to the Fissler and Ziegel’s (2001) loss class. Finally, we estimate our CARES-X model by setting $\theta = 0.05$, as it is a standard choice in the portfolio management literature. But our CARES-X model is also flexible enough to be applied to other operational research approaches, dedicated to risk management for instance. In this later case, it would be interesting to assess the estimates resulting from $\theta < 0.05$, such as 0.025, that is recommended by the Basel III Accord. We include these complementary analyses and extensions in our future research agenda.

Acknowledgements

We thank Jean-Luc Prigent for positive comments and encouragement about the first drafts of this article. We are also grateful to the Editor, Emanuele Borgonovo, and to the two anonymous referees who helped us to improve the article with fair comments, as well as to the participants of the XIIth International Conference on Computational and Financial Econometrics held in Pisa, and to the participants of the seminar organized by the University of Modena and Reggio Emilia for important suggestions, as well as the Referees of the IFC2020. Resources linked to this article are available on: www.systemic-risk-hub.org. The usual disclaimer applies.

References

- Acerbi, C., & Tasche, D. (2002). Expected Shortfall: a natural coherent alternative to Value at Risk. *Economic Notes*, 31, 379–388.
- Adrian, T., & Brunnermeier, M.K. (2016). CoVaR. *American Economic Review*, 106, 1705–1741.
- Artzner, P., Delbaen, F., Eber, J., & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9, 203–228.
- Barnard, R., Pearce, K., & Trindade, A. (2018). When is tail mean estimation more efficient than tail median? Answers and implications for quantitative risk management. *Annals of Operations Research*, 262, 47–65.
- Bellini, F., & Bernardino, E.D. (2017). Risk management with expectiles. *The European Journal of Finance*, 23(6), 487–506.
- Billio, M., Caporin, M., Panzica, R., & Pelizzon, L. (2017). The impact of network connectivity on factor exposures, asset pricing and portfolio diversification. *SAFE Working Paper Series, Leibniz Institute for Financial Research SAFE*.
- Billio, M., Getmansky, M., Lo, A., & Pelizzon, L. (2012). Econometric measures of connectedness and system risk in the finance and insurance sectors. *Journal of Financial Economics*, 104, 535–559.
- Bonaccolto, G., Caporin, M., & Panzica, R. (2019a). Estimation and model-based combination of causality networks among large us banks and insurance companies. *Journal of Empirical Finance*, 54, 1–21.
- Bonaccolto, G., Caporin, M., & Paterlini, S. (2019b). Decomposing and backtesting a flexible specification for CoVaR. *Journal of Banking & Finance*, 108, 105659.
- Bonaccolto, G. and Paterlini, S. (2020). Developing new portfolio strategies by aggregation. *Annals of Operations Research* 292(2), 933–971.
- Borri, N. (2019). Conditional tail-risk in cryptocurrency markets. *Journal of Empirical Finance*, 50, 1–19.
- Calabrese, R. and Osmetti, S.A. (2019). A new approach to measure systemic risk: a bivariate copula model for dependent censored data. *European Journal of Operational Research* 279(3), 1053–1064.

- Clemente, G., Grassi, R., & Hitaj, A. (2019). Asset allocation: new evidence through network approaches. *Annals of Operations Research*, forthcoming.
- Cui, H., Rajagopalan, S., & Ward, A.R. (2020). Predicting product return volume using machine learning methods. *European Journal of Operational Research*, 281(3), 612–627.
- DeMiguel, V., Garlappi, L., Nogales, F.J., & Uppal, R. (2009). A generalized approach to portfolio optimization: improving performance by constraining portfolio norms. *Management Science*, 55(5), 798–812.
- DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal versus naïve diversification: how inefficient is the 1/N portfolio strategy? *Review of Financial Studies*, 22, 1915–1953.
- Diebold, F., & Yilmaz, K. (2009). Measuring financial asset return and volatility spillovers, with application to global equity markets. *The Economic Journal*, 119, 158–171.
- Diebold, F., & Yilmaz, K. (2012). Better to give than to receive: predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28, 57–66.
- Diebold, F., & Yilmaz, K. (2014). On the network topology of variance decompositions: measuring the connectedness of financial firms. *Journal of Econometrics*, 182, 119–134.
- Diebold, F., & Yilmaz, K. (2015). *Financial and macroeconomic connectedness: a network approach to measurement and monitoring*. Oxford University Press.
- Duchin, R., & Levy, H. (2009). Markowitz versus the talmudic portfolio diversification strategies. *Journal of Portfolio Management*, 35, 71–74.
- Efron, B. (1991). Regression percentiles using asymmetric squared error loss. *Statistica Sinica*, 1, 93–125.
- Engle, R.F., & Manganelli, S. (2004). CAViaR: Conditional Autoregressive Value at Risk by regression quantiles. *Journal of Business & Economic Statistics*, American Statistical Association, 22, 367–381.
- Fan, J. and R. Li (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456), 1348–1360.

- Fissler, T. and J.F. Ziegel (2016). Higher order elicibility and Osband’s principle. *The Annals of Statistics*, 44(4), 1680–1707.
- Friedman, J.H., Hastie, T., & Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1), 1–22.
- Furno, M., & Vistocco, D. (2018). *Quantile Regression: Estimation and Simulation*. John Wiley & Sons.
- Gerlach, R.H., & Chen, C.W.S. (2014). Bayesian Expected Shortfall forecasting incorporating the intraday range. *Journal of Financial Econometrics*, 14, 128–158.
- Giuzio, M., & Paterlini, S. (2019). Un-diversifying during crises: is it a good idea? *Computational Management Science*, 16, 401–432.
- Gouriéroux, C., & Jasiak, J. (2008). Dynamic quantile models. *Journal of Econometrics*, 147, 198–205.
- Gupta, A., Wang, R., & Lu, Y. (2021). Addressing systemic risk using contingent convertible debt – a network analysis. *European Journal of Operational Research*, 290(1), 263–277.
- Hamidi, B., Hurlin, C., Kouontchou, P., & Maillet, B. (2015). A DARE for VaR. *Finance*, 36, 7–38.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning*. Springer.
- Hautsch, N., Schaumburg, J., & Schienle, M. (2014). Financial network systemic risk contributions. *Review of Finance*, 19, 1–54.
- Hull, J. (2011). *Options, futures and other derivatives*. Pearson.
- Jones, M. (1994). Expectiles and m-quantiles are quantiles. *Statistics & Probability Letters*, 20(2), 149–153.
- Kim, M., & Lee, S. (2016). Nonlinear expectile regression with application to Value-at-Risk and Expected Shortfall estimation. *Computational Statistics & Data Analysis*, 94, 1–19.
- Koenker, R. (1992). When are expectiles percentiles? *Econometric Theory*, 8(3), 423–424.

- Koenker, R. (1993). When are expectiles percentiles? *Econometric Theory*, 9(3), 526–527.
- Koenker, R. (2005). *Quantile regression*. Number 38. Cambridge university press.
- Koenker, R., & Bassett, G. (1978). Regression quantiles. *Econometrica*, 46(1), 33–50.
- Kremer, P., Talmaciu, A., & Paterlini, S. (2018). Risk minimization in multi-factor portfolios: what is the best strategy? *Annals of Operations Research*, 266, 255–291.
- Liao, L., Park, C., & Choi, H. (2019). Penalized expectile regression: an alternative to penalized quantile regression. *Annals of the Institute of Statistical Mathematics*, 71, 409–438.
- Maillet, B., Tokpavi, S., & Vaucher, B. (2015). Global minimum variance portfolio optimisation under some model risk: a robust regression-based approach. *European Journal of Operational Research*, 244(1), 289–299.
- Meng, X., & Taylor, J.W. (2020). Estimating Value-at-Risk and Expected Shortfall using the intraday low and range. *European Journal of Operational Research*, forthcoming.
- Michaud, R. (1989). The Markowitz optimization enigma: is 'optimized' optimal? *Financial Analysts Journal*, 45, 31–42.
- Newey, W., & Powell, J. (1987). Asymmetric least squares estimation and testing. *Econometrica*, 55, 819–847.
- Nguyen, L.H., Chevapatrakul, T., & Yao, K. (2020). Investigating tail-risk dependence in the cryptocurrency markets: a LASSO quantile regression approach. *Journal of Empirical Finance*, 58, 333–355.
- Nolde, N., & Ziegel, J.F. (2017). Elicitability and backtesting: perspectives for banking regulation. *Annals of Applied Statistics*, 11, 1833–1874.
- Pac, A., & Pinar, M. (2018). On robust portfolio and naïve diversification: mixing ambiguous and unambiguous assets. *Annals of Operations Research*, 266, 223–253.
- Patton, A.J., Ziegel, J.F., & Chen, R. (2019). Dynamic semiparametric models for Expected Shortfall (and Value-at-Risk). *Journal of Econometrics*, 211, 388–413.

- Pun, C.S., & Wong, H.Y. (2019). A linear programming model for selection of sparse high-dimensional multiperiod portfolios. *European Journal of Operational Research*, 273(2), 754–771.
- Ramponi, F.A., & Campi, M.C. (2018). Expected shortfall: Heuristics and certificates. *European Journal of Operational Research*, 267(3), 1003–1013.
- Rockafellar, R., & Uryasev, S. (2000). Optimization of Conditional VaR. *Journal of Risk*, 2, 21–41.
- Taylor, J.W. (2008). Estimating Value at Risk and Expected Shortfall using expectiles. *Journal of Financial Econometrics*, 6, 231–252.
- Taylor, J.W. (2019). Forecasting Value at Risk and Expected Shortfall using a semiparametric approach based on the asymmetric Laplace distribution. *Journal of Business & Economic Statistics*, 37, 121–133.
- Taylor, J.W. (2020). Forecast combinations for Value at Risk and Expected Shortfall. *International Journal of Forecasting*, forthcoming.
- Tibshirani, R. (1996). Regression analysis and selection via the LASSO. *Journal of the Royal Statistical Society, Series B*, 58(1), 267–288.
- Torri, G., Giacometti, R., & Paterlini, S. (2018). Robust and sparse banking network estimation. *European Journal of Operational Research*, 270(1), 51–65.
- Tu, J., & Zhou, G. (2011). Markowitz meets Talmud: a combination of sophisticated and naïve diversification strategies. *Journal of Financial Economics*, 99(1), 204–215.
- Wang, C., & Gerlach, R. (2019). Semi-parametric realized nonlinear conditional Autoregressive Expectile and Expected Shortfall models. *arXiv preprint <https://arxiv.org/abs/1906.09961>*.
- Yang, Y., & Zou, H. (2015). Nonparametric multiple expectile regression via er-boost. *Journal of Statistical Computation and Simulation*, 85(7), 1442–1458.
- Yao, Q., & Tong, H. (1996). Asymmetric least squares regression estimation: a nonparametric approach. *Journal of Nonparametric Statistics*, 6(2-3), 273–292.

Online Appendix

Appendix A. Differences with related models

The closest contribution to this study is the work of Hautsch et al. (2014). The authors also estimated financial networks from a pessimistic viewpoint, using the LASSO to deal with the curse of dimensionality. In addition, they also use loss exceedances to stress the impact of the $N - 1$ conditioning institutions. Nevertheless, we differ from Hautsch et al. (2014) on several relevant points.

Hautsch et al. (2014) defined the links among N nodes using a different risk measure; that is, VaR. In particular, Hautsch et al. (2014) estimated the τ^{th} conditional quantile of $r_{j,t}$ using the following quantile regression model (Koenker and Bassett, 1978):

$$Q(r_{j,t}; \tau) = \kappa_{\tau}^{(j)} \mathbf{x}_{j,t}', \quad (\text{A.1})$$

where τ takes low values in the interval $(0, 1)$, typically $\tau \in \{0.01, 0.05\}$, so that $Q(r_{j,t}; \tau)$ represents the VaR of $r_{j,t}$ at level τ conditional on the row vector $\mathbf{x}_{j,t} = [1, \mathbf{y}_{j,t-1}, \mathbf{w}_{t-1}, \mathbf{h}_{-j,t}]$, for $j = 1, \dots, N$. The elements of $\mathbf{x}_{j,t}$ are the following: i) $\mathbf{y}_{j,t-1}$, which includes lagged specific characteristics of the j^{th} company (such as leverage, maturity mismatch, market-to-book value, market capitalization, equity return volatility) along with its lagged return $r_{j,t-1}$; ii) $\mathbf{w}_{t-1}$, which includes lagged macroeconomic state variables; and iii) $\mathbf{h}_{-j,t}$, which coincides with the vector of loss exceedances we use in our CARES-X model. The parameters in Equation (A.1) are estimated by adding an ℓ_1 -norm penalty to the loss function which characterizes the standard quantile regression, as introduced by Koenker and Bassett (1978). Therefore, there is an influence from the i^{th} company to the j^{th} institution in the resulting network, if $h_{i,t}$ is LASSO-selected as a relevant driver of $Q(r_{j,t}; \tau)$; that is, if the loss exceedance of the former has a significant impact on the VaR of the latter. According to Hautsch et al. (2014), $h_{i,t}$ is LASSO-selected if the absolute value of its corresponding coefficient is greater than 0.0001.

In contrast to the method introduced by Hautsch et al. (2014), our CARES-X model builds on an expectile regression model. Although quantile regression is more robust to outliers, different contributions in the literature highlighted several advantages of expectile regression; see, among many others, Newey and Powell (1987), Efron (1991), Koenker (1992, 1993), Jones (1994), Yao and Tong (1996), Koenker (2005), Taylor (2008), Yang and Zou (2015), Bellini and Bernardino (2017) and Furno and Vistocco (2018). Some

of them are summarized as follows: i) expectile regression is computationally simple, building on an asymmetric least squares loss, which is differentiable everywhere. In contrast, the check loss function characterizing the quantile regression model is not everywhere differentiable, so that the underlying optimization routine might require a set of restrictions affecting the computational efficiency; ii) it is possible to define conditional quantiles as a function of expectiles. Indeed, there exists a one-to-one mapping between quantiles and expectiles. As a result, we compute quantiles from expectiles, exploiting the computational advantages behind the latter; iii) expectiles have a more global dependence on the form of the distribution compared to quantiles (Koenker, 2005). Altering the shape of the upper tail of the $r_{j,t}$'s distribution does not change the quantiles of the lower tail, but it does have an impact on all the expectiles (Taylor, 2008). As a result, expectiles are more sensitive to tail events, signalling more promptly the systemic impact of a given financial company; and iv) the estimated expectile curve is typically smoother than the one derived from quantiles.

Moreover, an extreme quantile is interpreted as VaR, whereas expectiles lead to ES, which has better properties from both a mathematical and a financial viewpoint. For instance, we refer to the subadditivity property, which is consistent to the concept of financial diversification. ES is subadditive, whereas this is not the case for VaR. As a result, VaR is not a coherent measure of risk (Artzner et al., 1999). As highlighted by Acerbi and Tasche (2002), broken axioms always lead to paradoxical, wrong results, and VaR makes no exception. Besides, VaR is an estimate of the maximum loss a company can potentially incur with a given probability level over a given time period, but it does not provide any information about the magnitude of the actual loss when it is greater than VaR. In contrast, ES overcomes this limitation (Hull, 2011). The better properties of ES motivate the changes that the Basel Committee on Banking Supervision (BCBS) has adopted in the Basel III accords. In particular, the BCBS has substituted VaR with ES for the definition of market risk capital requirements. Our network is then consistent with the latest regulatory standards.

Other differences concern the specification of the regression model. In contrast to Hautsch et al. (2014), we do not use the lagged firm-specific characteristics in $\mathbf{y}_{j,t-1}$ to link $\mu(r_{j,t}; \tau)$ to the past of $r_{j,t}$. Some of these covariates (such as leverage and market-to-book value) are retrieved from balance sheet indicators, which are observed with a low frequency. Hautsch et al. (2014) made them available for each $t = 1, \dots, T$ using an interpolation, which is affected by inevitable estimation errors. In contrast, our method is more efficient, as the past of $r_{j,t}$ along with the dynamics and

persistence of expectiles over time are directly captured by a single latent variable; that is, $\mu(r_{j,t-1}; \tau_j^*)$. For similar reasons, we do not need lagged macroeconomic variables in $\mathbf{w}_{t-1}$ to capture the dynamics from $t-1$ to t , but directly estimate contemporaneous relationships between $\mathbf{w}_t$ and $ES(r_{j,t}; \theta)$. Indeed, we found that $\mathbf{w}_t$ has a more significant impact on $ES(r_{j,t}; \theta)$ when compared to $\mathbf{w}_{t-1}$.

Finally, we differ from Hautsch et al. (2014) for the way in which we LASSO-select the relevant connections among the N financial institutions. Among the covariates in Equation (9), we penalize the loss exceedances only to stress further their impact. By doing so, we filter stronger and more significant links. We also stress the fact that Hautsch et al. (2014) used a threshold of 0.0001 to LASSO-select the relevant loss exceedances. However, this threshold is arbitrarily defined and minimal differences in the absolute values of the coefficients could lead to different structures of the resulting network. In contrast, the threshold η that we use to LASSO-select the relevant loss exceedances in (9) has a statistical interpretation (see Section 3).

Appendix B. List and information of the network nodes

Table (B.3) Information on the network nodes

NUMBER	COMPANY	SECTOR	COUNTRY
1	ALLIED IRISH BANKS	BANKING	IRE
2	BANCO BPM	BANKING	ITA
3	BANK OF IRELAND GROUP	BANKING	IRE
4	BANKINTER 'R'	BANKING	SPA
5	BARCLAYS	BANKING	GBR
6	BBV.ARGENTARIA	BANKING	SPA
7	BANCO DE SABADELL	BANKING	SPA
8	BANCO SANTANDER	BANKING	SPA
9	BNP PARIBAS	BANKING	FRA
10	BPER BANCA	BANKING	ITA
11	CLOSE BROTHERS GROUP	BANKING	GBR
12	COMMERZBANK (XET)	BANKING	GER
13	CRÉDIT AGRICOLE	BANKING	FRA
14	CREDIT SUISSE GROUP N	BANKING	SWI
15	DANSKE BANK	BANKING	DEN
16	DEUTSCHE BANK (XET)	BANKING	GER
17	DNB	BANKING	NOR
18	ERSTE GROUP BANK	BANKING	AUS
19	SOCIETE GENERALE	BANKING	FRA

Continued on next page

NUMBER	COMPANY	SECTOR	COUNTRY
20	HSBC HDG.	BANKING	GBR
21	ING GROEP	BANKING	NET
22	INTESA SANPAOLO	BANKING	ITA
23	JYSKE BANK	BANKING	DEN
24	KBC GROUP	BANKING	BEL
25	KOMERCNI BANKA	BANKING	CZR
26	LLOYDS BANKING GROUP	BANKING	GBR
27	MEDIOBANCA BC.FIN	BANKING	ITA
28	NATIXIS	BANKING	FRA
29	NORDEA BANK	BANKING	SWE
30	ROYAL BANK OF SCTL.GP.	BANKING	GBR
31	SEB 'A'	BANKING	SWE
32	STANDARD CHARTERED	BANKING	GBR
33	SVENSKA HANDBKN.'A'	BANKING	SWE
34	SWEDBANK 'A'	BANKING	SWE
35	SYDBANK	BANKING	DEN
36	UNIONE DI BANCHE ITALIAN	BANKING	ITA
37	UBS GROUP	BANKING	SWI
38	UNICREDIT	BANKING	ITA
39	AEGON	INSURANCE	NET
40	AGEAS (EX-FORTIS)	INSURANCE	BEL
41	ALLIANZ (XET)	INSURANCE	GER
42	ASSICURAZIONI GENERALI	INSURANCE	ITA
43	AVIVA	INSURANCE	GBR
44	AXA	INSURANCE	FRA
45	BALOISE-HOLDING AG	INSURANCE	SWI
46	BEAZLEY	INSURANCE	GBR
47	CNP ASSURANCES	INSURANCE	FRA
48	HANNOVER RUCK. (XET)	INSURANCE	GER
49	HELVETIA HOLDING N	INSURANCE	SWI
50	HISCOX (DI)	INSURANCE	GBR
51	LEGAL & GENERAL	INSURANCE	GBR
52	MAPFRE	INSURANCE	SPA
53	MUENCHENER RUCK. (XET)	INSURANCE	GER
54	OLD MUTUAL	INSURANCE	GBR
55	PRUDENTIAL	INSURANCE	GBR
56	RSA INSURANCE GROUP	INSURANCE	GBR
57	SAMPO 'A'	INSURANCE	FIN
58	SCOR SE	INSURANCE	FRA
59	ST.JAMES'S PLACE	INSURANCE	GBR
60	STOREBRAND	INSURANCE	NOR
61	SWISS LIFE HOLDING	INSURANCE	SWI
62	SWISS RE	INSURANCE	SWI
63	UNIPOLSAI	INSURANCE	ITA
64	ZURICH INSURANCE GROUP	INSURANCE	SWI
65	3I GROUP	FINANCIAL SERVICES	GBR

Continued on next page

NUMBER	COMPANY	SECTOR	COUNTRY
66	AAREAL BANK (XET)	FINANCIAL SERVICES	GER
67	ACKERMANS & VAN HAAREN	FINANCIAL SERVICES	BEL
68	DEUTSCHE BOERSE (XET)	FINANCIAL SERVICES	GER
69	GBL NEW	FINANCIAL SERVICES	BEL
70	INDUSTRIVARDEN 'A'	FINANCIAL SERVICES	SWE
71	INTERMEDIATE CAPITAL GP.	FINANCIAL SERVICES	GBR
72	INTRUM JUSTITIA	FINANCIAL SERVICES	SWE
73	INVESTEC	FINANCIAL SERVICES	GBR
74	INVESTOR 'B'	FINANCIAL SERVICES	SWE
75	KINNEVIK 'B'	FINANCIAL SERVICES	SWE
76	LONDON STOCK EX.GROUP	FINANCIAL SERVICES	GBR
77	MAN GROUP	FINANCIAL SERVICES	GBR
78	NEX GROUP	FINANCIAL SERVICES	GBR
79	PARGESA 'B'	FINANCIAL SERVICES	SWI
80	PROVIDENT FINANCIAL	FINANCIAL SERVICES	GBR
81	SCHRODERS	FINANCIAL SERVICES	GBR
82	WENDEL	FINANCIAL SERVICES	FRA
83	BRITISH LAND	REAL ESTATE	GBR
84	CASTELLUM	REAL ESTATE	SWE
85	COFINIMMO	REAL ESTATE	BEL
86	DERWENT LONDON	REAL ESTATE	GBR
87	DEUTSCHE EUROSHP (XET)	REAL ESTATE	GER
88	FABEGE	REAL ESTATE	SWE
89	FASTIGHETS BALDER 'B'	REAL ESTATE	SWE
90	FONCIERE DES REGIONS	REAL ESTATE	FRA
91	GEICNA REIT	REAL ESTATE	FRA
92	GREAT PORTLAND ESTATES	REAL ESTATE	GBR
93	HAMMERSON	REAL ESTATE	GBR
94	ICADE REIT	REAL ESTATE	FRA
95	INTU PROPERTIES	REAL ESTATE	GBR
96	JM	REAL ESTATE	SWE
97	KLEPIERRE	REAL ESTATE	FRA
98	LAND SECURITIES GROUP	REAL ESTATE	GBR
99	LUNDBERGFORETAGEN 'B'	REAL ESTATE	SWE
100	PSP SWISS PROPERTY AG	REAL ESTATE	SWI
101	SEGRO	REAL ESTATE	GBR
102	SHAFTESBURY	REAL ESTATE	GBR
103	SWISS PRIME SITE	REAL ESTATE	SWI
104	UNIBAIL-RODAMCO SE REIT	REAL ESTATE	NET

Appendix C. Additional figures

Figure (C.9) Trend of the average percentage returns of the $N = 104$ companies included in our dataset from January 2, 2004 to November 16, 2007.

Figure (C.10) Trend of the average Country-Level Index of Financial Stress (CLIFS) of the Countries listed in the fourth column of Table B.3 (except Norway and Switzerland, because CLIFS is not available for these Countries), provided by the European Central Bank with a monthly frequency, superimposing the density, reciprocity (taken with the negative sign) and ES series displayed in Figures 3 and 8, respectively. These time series are standardized so that they are expressed in the same scale. For simplicity, we do not display here the trend of the Composite Indicator of Systemic Stress (CISS) provided by the European Central Bank, as it behaves similarly to the average CLIFS (linear correlation coefficient of 0.91).

Appendix D. Definition of degrees, density and reciprocity

We evaluate our network using both local and global measures, according to the fact that we focus on either a single node (or vertex) or the overall graph, respectively. Starting from the former, we compute the in- and the out-degree of each node; that is, the number of edges (or links) pointing in towards and out from a vertex, respectively. We estimate directed and unweighted networks with no self-loops, each of which is determined by an adjacency matrix $\mathbf{A} = [a_{j,i}] \in \mathbb{R}^{N \times N}$ (see Section 3 for additional details on the computation of $\mathbf{A}$). As a result, we calculate the in-degree of node j and the out-degree of vertex i , respectively, as follows:

$$ID_j = \sum_{i=1}^N a_{j,i}, \quad (\text{D.1})$$

$$OD_i = \sum_{j=1}^N a_{j,i}, \quad (\text{D.2})$$

for $(i, j) = 1, \dots, N$.

Likewise, we also define the degree of node j ; that is, the sum of edges, without considering their direction, incident on it, for $j = 1, \dots, N$. Moving to the global statistics, we compute the density and the reciprocity of each rolling network. The density of a graph is the frequency of realized edges relative to potential edges:

$$den = \frac{|E|}{N(N-1)}, \quad (\text{D.3})$$

where $|E|$ is the number of realized links within a given network.

The reciprocity is defined as the proportion of mutual, or reciprocated, links:

$$rec = \frac{\sum_{j,i} (\mathbf{A} \odot \mathbf{A}')_{j,i}}{\sum_{j,i} a_{j,i}}, \quad (\text{D.4})$$

where $\odot$ denotes the Hadamard product.