

Technical note on the Financial Shock Simulator (FSS)¹

1. Introduction

The Financial Shock Simulator (FSS) is a state-of-the-art tool used by the European Central Bank (ECB) to calibrate financial shocks for adverse scenarios as part of its stress test design. It has been used in the calibration of financial shocks for the European Banking Authority (EBA), European Insurance and Occupational Pensions Authority (EIOPA) and European Securities and Market Authority (ESMA) scenarios since 2014, although the tool has evolved significantly over time, especially in its practical application. This tool is regularly employed for internal and external policy analysis, such as for the impact assessments in the Financial Stability Review (Bassanin, Rancoita and Ojea Ferreiro, 2019; Rancoita and Gross, 2015; Rancoita and Gross, 2017).

The FSS is able to capture correlations in the extreme tails of financial returns' distributions and identify spillover effects across securities and jurisdictions. Furthermore, it is a tool that can combine a large number of time series (3000+), allowing for the calibration of very granular scenarios. These features have made it possible to cope with the steep increase in the granularity of scenarios that has occurred since 2014 (see, for example, the market risk scenario in the EBA's EU-wide stress test, the ESMA central counterparties (CCPs) and money market funds (MMF) scenarios, and the EIOPA insurance and institutions for occupational retirement provision (IORP) scenarios).

The FSS is based on a multivariate copula approach which calibrates the shocks and builds on the concepts of Conditional Expected Returns and Conditional Expected Shortfall. This approach takes into account all the outcomes in the entire tail of the distribution and not just the value at a certain percentile, which is the case for Value at Risk (*VaR*) (see Acerbi and Tasche, 2002). For this reason, it is particularly suitable for designing extreme scenarios, such as those used in stress testing.

Given the complexity of the database, the collection and processing of the input data are two essential steps to ensure the robustness of the results. This is because (i) data are retrieved from several sources and need to be combined and sorted according to quality, and (ii) some series are highly

¹ Prepared by E. Rancoita and J. Ojea Ferreiro.

irregular (for example some securities are not regularly traded, and others could have been suspended, etc.).

The FSS is built in a flexible manner, which enables the financial shocks to be calibrated purely non-parametrically or by assuming a Gaussian distribution. The major advantage of the non-parametric FSS is that it is able to produce a realistic quantitative assessment of the co-movement in the tail of the distributions, as financial variables are usually log-normally distributed (e.g. stock prices decline more rapidly in crisis times than they increase in booms) and usually exhibit higher dependence during crisis periods. The parametric copula can be more informative when looking to achieve quantitative benchmarks in cases in which the length of the data is short or the market liquidity of the instrument is low. However, the results of the parametric FSS should be analysed taking into consideration the underlying Gaussian assumption, which implies that results tend to be more “optimistic”. Real financial data have higher tail dependence (i.e. they exhibit a stronger cross-sectional correlation in extreme scenarios) and are leptokurtic (the extreme scenarios occur more often than what would be expected under the assumption of Gaussian distribution). Consequently, it provides a fallback solution in the absence of adequate data for the non-parametric copula and serves as a floor for the non-parametric results. It is also useful for identifying outliers in the series that normal filters do not identify.

2. FSS methodology: a copula-based approach

Let us define the financial security triggering the scenario as x_m and the remaining securities as x_i , where $i = 1, \dots, N$. For simplicity of notation, the subscript t indicating the time is ignored. The FSS is based on the joint distribution of the financial variables and its decomposition using a copula approach. The copula decomposition of the joint distribution provides high flexibility for modelling complex patterns, such as asymmetric dependence or strong joint tail behaviour. Sklar’s theorem (Sklar, 1959) states that a multivariate cumulative distribution function can be expressed as a combination of the marginal distributions and a copula that contains the co-movement between variables, i.e.:

$$F(x_i, x_m) = C(F_i(x_i), F_m(x_m)) \quad (2.1)$$

where x_j refers to a stationary economic variable with distribution function $F_j(\dots)$, $j = \{i, m\}$ and $C(\dots)$ is the copula function.

The conditional cumulative distribution for variable i , given that the triggering variable m is at its p 100th quantile, can be derived from equation (2.1) as:

$$F(x_i | x_m = F_m^{-1}(p)) = C(F_i(x_i) | F_m(x_m) = p) \quad (2.2)$$

where $C(x_i|x_m = F_m^{-1}(p)) = \frac{\partial C(F_i(x_i), F_m(x_m))}{\partial F_m(x_m)} \Big|_{F_m(x_m)=p}$. $F(x_i|x_m = F_m^{-1}(p))$ is the conditional distribution function for each variable i given that the triggering variable m is in its p 100% worst case scenario. **The conditional cumulative distribution in equation (2.2) is the cornerstone concept using the copula method to quantitatively measure and assess systemic risk and contagion spillovers** (see Adrian and Brunnermeier, 2016; Brownlees and Engle, 2016; Girardi and Ergün, 2013; Acharya, Engle, and Richardson, 2012).

We assess a set of risk measures that could give us an idea about the resilience of the economy in the event of a stress scenario defined by the narrative. In the event of a crisis we would like to know what the average behaviour of the prices of the different securities would be, i.e. conditional mean return, but also what the worst reaction would be given a high confidence level, i.e., Conditional Value-at-Risk, as well as the mean return beyond the Conditional Value-at-Risk, Conditional Expected Shortfall. All of these metrics are explained briefly below.

Conditional Mean Return (CMR). The conditional mean return of variable i , conditional on the triggering variable m being below its p 100th quantile, is obtained using Bayes' theorem as:

$$E(x_i|x_m \leq F_m^{-1}(p)) = \frac{1}{p} \int_0^1 F_i^{-1}(u) C(p|F_{i(x_i)} = u) du \quad (2.3)$$

Equation (2.3) is the main metric provided by the FSS tool and employed for the adverse scenarios. It indicates the mean value of variable i when the triggering variable m is below its p 100% worst case scenario. It is worth noting that $E(x_i|x_i \leq F_i^{-1}(p))$ is the Expected Shortfall for the triggering variable.

Conditional Value-at-Risk (CoVaR). The *CoVaR* expresses the α 100% worst case scenario for variable i conditioned on a stressed scenario for the triggering variable m . The triggering variable m is considered to be stressed when it is below its p 100th quantile, i.e. $P(x_i < CoVaR_{i|m}(\alpha, p) | x_m < F_m^{-1}(p)) = \alpha$.

This conditional quantile in copula terms would be:

$$\frac{C(F_i < CoVaR_{i|m}(\alpha, p), p)}{p} = \alpha$$

Hence, the CoVaR would be the quantile q 100th of variable i , i.e. $CoVaR_{i|m} = F_i^{-1}(q^*)$, such that:

$$\int_0^{q^*} C(p|F_{i(x_i)} = u) du = \alpha p \quad (3.2)$$

Note that unlike the unconditional VaR, here we are merging individual and joint risk, the effect of contagion on the assets' risk measure owing to a shock to the triggering variable m can be assessed as $CoVaR_{i|m} - VaR_i$, where $VaR_i = F_i^{-1}(\alpha)$.

Conditional Expected Shortfall (CoES). The CoES indicates the mean value of variable i given that i is below its CoVaR, i.e.:

$$E(x_i | x_i < CoVaR) = \frac{1}{q^*} \int_0^{q^*} F_i^{-1}(u) C(p | F_i(x_i) = u) du \quad (2.2)$$

where $q^* = F_i(CoVaR_{i|m})$.

These additional measures increase the knowledge of the quantitative effects of financial shocks on economic indicators. In terms of severity, intuitively, the conditional mean should be less severe than the CoVaR, which should be less severe than the CoES.

Figure 1: Isodensities of the joint distribution

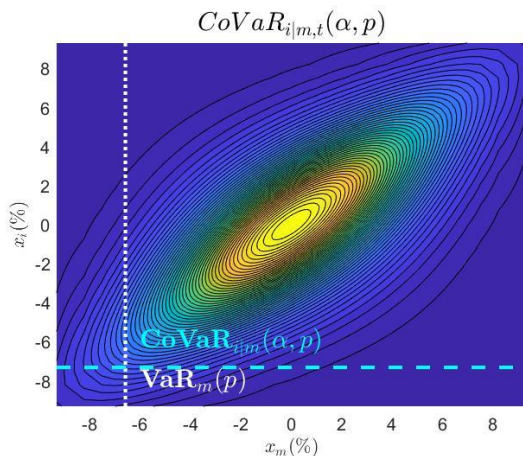
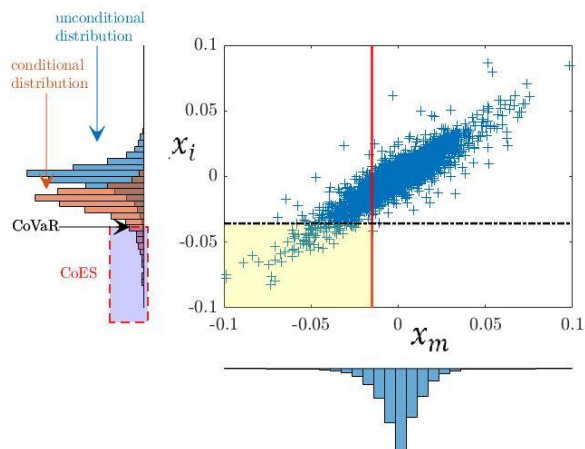


Figure 1 presents an example of how the conditional mean, the CoVaR and the CoES are computed in the probability space. Each line in Figure 1 indicates a set of scenarios that has the same probability of occurrence, i.e. the so-called isodensities. The lighter the colour of the line, the more likely it is to observe the realisations of the stochastic variables.

The white line denotes the $p100$ th quantile for the triggering variable m , i.e. $F_m^{-1}(p) = VaR_m$. The conditional mean for variable i is the average return along the y-axis considering only the probabilities of occurrence on the left of the white line. The $CoVaR_{i|m}$ would be the value that corresponds to the quantile $\alpha100^{th}$ of variable i given the set of scenarios to the left of the threshold VaR_m . Finally, the $CoES_{i|m}$ is the mean return of variable i on the left of the VaR_m and below the $CoVaR_{i|m}$. Figure 2 provides the same example looking at the realisations of a sample of 7000 pairwise observations, where the blue bars indicate the histogram of the marginal observations, the red line indicates the VaR of the triggering variable, and the orange bars indicate the histogram of the conditioned observations of the variable i given that variable m is on the left of the red line. CoVaR is the VaR measured in the histogram with orange bars while the CoES is the mean return measured with the orange bars within the purple section of the histogram.

The $CoVaR_{i|m}$ would be the value that corresponds to the quantile $\alpha100^{th}$ of variable i given the set of scenarios to the left of the threshold VaR_m . Finally, the $CoES_{i|m}$ is the mean return of variable i on the left of the VaR_m and below the $CoVaR_{i|m}$. Figure 2 provides the same example looking at the realisations of a sample of 7000 pairwise observations, where the blue bars indicate the histogram of the marginal observations, the red line indicates the VaR of the triggering variable, and the orange bars indicate the histogram of the conditioned observations of the variable i given that variable m is on the left of the red line. CoVaR is the VaR measured in the histogram with orange bars while the CoES is the mean return measured with the orange bars within the purple section of the histogram.

Figure 2: Histogram and scatter plot for bivariate data



3. Database and data preparation

3.1 Input data

The database spans a long sample of daily data including more than 3000 securities across several types of asset classes and jurisdictions, starting at the beginning of 2000 up to the most recent date and including periods of financial distress such as the 2008 financial crisis and the European sovereign debt crisis. However, subsamples could be selected for specific case studies. Several types of securities are included in the database, including equity indices, government bond yields, swap rates, implied volatilities, exchange rates, dividend yields, liquidity proxies. Annex 1: includes a subsample of the main input securities of the toolbox. As can be observed, the data are very granular. For instance, corporate yields are not only classified by different rating classes but also by type of firm (financial, non-financial) or by country. In addition to allowing for joint shocks to a given market, the FSS also makes it possible to build shocks by country or region which is helpful for simulating a European sovereign debt crisis, for example. The toolbox also distinguishes between different maturities for government yields or different obligor grades for corporate yields, which allows for the possibility of assessing a tailor-made scenario that generates a conditioning scenario for an segment of the market considered more vulnerable.

3.2 Data preparation

The FSS mostly operates on the distribution of daily returns as non-stationary time series are made stationary using standard approaches, e.g. first difference for swap rates or log-returns for equity quotes. In addition, the toolbox also makes it possible to work with the distribution of spreads

series, which is particularly useful for European government yields, for which the stress scenarios usually correspond to a widening of spreads against German bonds.

The quality of the input data is the main impediment to achieving reliable results. For this reason, the FSS toolbox performs a cleansing procedure before starting the simulation process. The toolbox applies several approaches to identify outliers and deviations from the historical path² so that possibly miswritten values can be removed while retaining the extreme values that are supported by real data behaviour. For example cleaning out extreme deviations combined with graphical tools such as charts, plots and surface graphs allows for different criteria to assess the data for a particular security or economic period. Particular attention was devoted to detecting data errors for government yields and swap rates as their presence can significantly modify the surface of the term structure across different maturities. In this case, the FSS compares the interest rates and yields across several available sources of market data and chooses the best proxy by merging several data sources, considering the possible differences between them in terms of levels or data quality.³

Once the data has been cleansed, the toolbox allows for several interpolation methods – linear, cubic or spline – to generate the most complete set of time series possible. This means that for yields and interest rates a double interpolation is performed, both across maturities and over time, thus avoiding non-smooth yield/swap rate curves in the results that arise as a consequence of certain maturities not being traded very often.

4. The FSS in practice

This section presents more technical details about how the FSS toolbox works in practical application.

4.1 Non-parametric FSS based on Monte Carlo simulations

The non-parametric application of the FSS allows the asymmetric distribution of returns and its skewness for negative values to be taken into consideration. The calibration of the deviation from the starting point level is then calculated as follows:

1. The observations of the joint distribution of all the daily rates of change are bootstrapped via Monte Carlo simulation, where the daily rates of change are assumed to be independently and identically distributed (iid) over time (see also Adrian and Brunnermeier, 2016). The new sample of daily rates of change of all securities reflects the daily cross-sectional correlation between variables, which is the key feature for capturing the response to the different scenarios. It is worth

² For instance, the tool makes it possible to choose between the elimination of the first, fifth and tenth percentiles of the distributions or to eliminate the outliers based on the distance from the trimmed or Winsorised mean. However, these cleaning procedures are not usually used for very extreme scenarios.

³ Manual corrections are also introduced following visual checks.

noting that this daily correlation obtained by bootstrapping simulation is conditioned on the data sample. The selection of the data sample allows abstraction from exogenous shocks that might affect the calibration but cannot be identified by this model. Using other econometric approaches would make it possible to identify the effect of other exogenous shocks and to abstract the responses from the impact of changes in the monetary policy, for example. In the copula model selecting a sample period for the calibration of some shocks allows abstraction from other exogenous variations. For example, selecting a period with a monetary policy stance that is constant and at a level similar to the one assumed in the scenario would allow abstraction from the endogenous effects of a change in monetary policy and, therefore, a coherent scenario calibration.

2. To provide the shock over a time horizon longer than one day, the toolbox allows for two options: (i) a linear extrapolation over the desired horizon, and (ii) the simulation of the returns over the desired horizon bootstrapped from daily returns.

The first option is the least time-consuming approach, although it could present imprecise results as the time span becomes longer (normally for a scenario horizon longer than 60 business days). In addition, this assumption simplifies the simulation process in periods where volatility and correlation across variables are relatively stable. However, this assumption might be questioned when we have structural changes and shifts in dependence across variables.

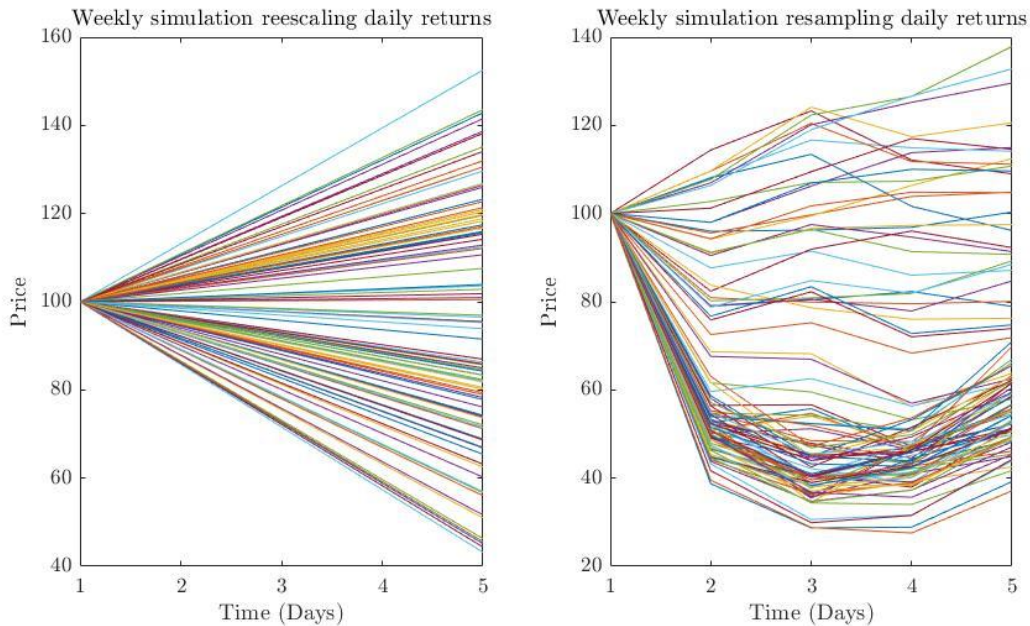
In the second case, returns are simulated over the intermediate horizon and then summed to obtain a path for the returns over the entire horizon before being rescaled⁴ to the length of the overall horizon. In the latter case, returns are simulated over the entire desired horizon and then summed. These procedures are computationally intensive.

Figure 3 shows a simulation example where the price of an asset is assumed to be log-normal. In the left-hand chart 100 paths of the final price of the asset after five days are illustrated based on the empirical daily returns and their extrapolation over one week. In the right-hand chart, the five-day paths are constructed by first extracting daily returns for each day, then compounding the returns over five days, gathering dependence shift and changes in the relationships over the sample period. Note that the simulation of $t+k$ where $k=1, \dots, 5$ is performed in a cross-sectional way, so the same day is chosen across different types of assets, ensuring consistency in the simulation results across the different assets classes. Repeating this simulation a very large number of times via bootstrapping (in the example 100 times, but actual simulations are always performed over 10,000 times) ensures that choosing very extreme returns for five days, which

⁴ The rescaling is carried out by dividing the returns obtained over the intermediate horizon by the square root of the number of days in the intermediate horizon and multiplying them by the square root of the number of days in the entire horizon.

would cause the final results to go beyond the maxima and minima of the original time series, is very unlikely (see Annex 3.)

Figure 3 : Simulation of weekly prices using daily data obtained by extrapolation and by bootstrapping daily returns




3. The triggering variable m is constructed as the weighted average of one or more variables that the narrative of the scenario defines as an exogenous shock that triggers the entire scenario, i.e., $x_m = \frac{\sum_i x_i \omega_i}{\sum_i \omega_i}$ where x_m is the triggering variable, x_i is the return of the individual security i where $i=1, \dots, N$, and ω_i is the weight assigned to that variable. The weight indicates the relative importance of a factor in triggering the scenario. For instance, to replicate the 2010 European sovereign debt crisis, the weights of the government bond yields of the Southern European countries should be higher than those of other jurisdictions. These features make it possible to obtain a scenario that is tailor-made to the narrative.
4. The observations below the p 100% worst case scenario of the triggering variable m are selected, i.e. $x_m \leq F_m^{-1}(p)$.
5. The response of the other financial variables in the sample is estimated by assessing the main returns in the simulations where the observations of the triggering variable were selected according to the previous point, i.e. $E(x_i | x_m < F_m^{-1}(p))$. Figure 4 presents an example of the returns from the other financial variables given a stress scenario for the triggering variable. The stressed

scenario occurs when the triggering variable (blue line) crosses into the red area, which would be reflected in a certain return in the response variable (green line).

6. The resulting deviation from the starting point level is given on a daily basis so it should be either rescaled or the simulated process repeated to obtain quarterly or yearly values, as explained in Step 2.

Figure 4: Financial shock simulator methodology, a visual example.



Section 5 presents an example of how the non-parametric copula works for a shock in a financial distress scenario in which equity quotes experience an extreme downward movement.

4.2 Parametric FSS based on Gaussian distribution

The FSS includes also a parametric copula model to facilitate the evaluation of the plausibility of the outcomes of the FSS and to cope with series in which there are substantial missing values.

The parametric FSS provides a quantitative floor for the assessment of the shock impacts. This increases the quantitative sensitivity to the values obtained by the non-parametrical FSS, acting as a complementary tool in the shock simulation process. Discrepancies between the non-parametric and the parametric FSS results can be also very informative, shedding light on input quality problems or complex features in the data structure.

Currently, only Gaussian distribution is considered for the joint behaviour of variables in the parametric version of the FSS. This distribution has closed-form formulas for the conditional mean, *CoVaR* and *CoES*, providing a fast quantitative assessment without relying on integration methods. A fat-tail feature, i.e. excess of kurtosis, can be observed to some degree in the unconditional Gaussian distribution owing to the use of conditional variance and covariance (see Cherubini, Gobbi, and Mulinacci, 2016). The use of a time-varying correlation also allows for some tail-dependence feature to be captured in the unconditional distribution.

The conditional variance and covariance is computed following an exponential weighted moving average (EWMA) model. The main reason for using this specification is the simplicity of relying only on a parameter λ that weighs the importance of new observations, contributing to the higher flexibility of the model. The λ parameter is between 0 and 1, where $\lambda=1$ would give us the unconditional variance and covariance, while the closer λ is to zero, the greater the importance given to more recent observations. A default value is $\lambda=0.94$, as suggested by RiskMetrics, which gives a balanced result between the historical observations and the relevance of new information. A robustness check on the results can be constructed using a λ between 0.9 and 1, for instance.

The mean is assumed to be constant if daily data, not weekly data, are used. Data are assumed to follow a first-order autoregressive model in this case, an option available in the FSS toolbox that changes the extrapolation of the results and can be also useful as a robustness check. To sum up, the most complex model choice in the parametric FSS, AR(1)-EWMA, has four parameters (two for the conditional mean, one for the conditional variance and one for the joint distribution). In this way, and without increasing either model risk or computational time, we find a balance between the flexibility of the model and its complexity. It is important to highlight that the results of the parametric FSS tool are complementary to those obtained from the non-parametric FSS tool and in no case are they substitute for them. Furthermore, the parametrical FSS makes some strong distribution assumptions that cannot be overlooked in a sensible analysis of the shock response of the financial indicators.

The next sections present the equations for the model and the closed formulae for the risk measures according to the assumptions made.

4.2.1 Model assumptions

4.2.1.1. Conditional mean

The conditional mean for the transformed variable x_i is:

$$x_{i,t} = \mu_{i,t} + \epsilon_{i,t} \quad \text{for } t=1, \dots, T \text{ and } \forall i$$

where $\mu_{i,t} = \phi_{i,0}$ if the data is daily and $\mu_{i,t} = \phi_{i,0} + \phi_{i,1}x_{i,t-1}$ if we are dealing with weekly data. The parameters are estimated using a quasi-maximum likelihood because it is quicker in computational terms.

$\epsilon_{i,t}$ is the innovation term, which presents heteroscedasticity, i.e.:

$$\epsilon_{i,t} = \xi_i \sigma_{i,t}$$

where $\sigma_{i,t}$ is the conditional standard deviation and $\xi_{i,t}$ is distributed as a standard Gaussian variable.

To extrapolate the standard deviation $\sigma_{i,t}$ to get an annual value using daily data would be $\sigma_{i,t:t+252} = \sqrt{252}\sigma_{i,t}$,

while for the weekly data the formula would be $\sigma_{i,t:t+52} = \sigma_{i,t} \left(52 + 2 \frac{\phi_{i,1}}{(1-\phi_{i,1})^2} [51(1-\phi_{i,1}) - \phi_{i,1}(1-\phi_{i,1}^{51})] \right)^{\frac{1}{2}}$ under the assumption that a year has 252 business days and 52 business weeks.

To extrapolate the mean $\mu_{i,t}$ to get an annual value using daily data would be $\mu_{i,t:t+252} = 252\phi_{i,0}$, while for

the weekly data the formula would be $\mu_{i,t:t+52} = \frac{\phi_{i,0}}{1-\phi_{i,1}} \left(52 - \frac{\phi_{i,1}-\phi_{i,1}^{53}}{1-\phi_{i,1}} \right) + \frac{r_{i,t}}{1-\phi_{i,1}} (\phi_{i,1} - \phi_{i,1}^{53})$.

4.2.1.2. Conditional variance and correlation

An EWMA model is used to model the heteroscedasticity in the data.

$$\sigma_{i,t}^2 = \epsilon_{i,t-1}^2(1-\lambda) + \sigma_{i,t-1}^2\lambda \quad \text{for } t=1,\dots,T \text{ and } \forall i$$

To model the cross-sectional dependence, the conditional covariance is computed as

$$\sigma_{ij,t} = \epsilon_{i,t}\epsilon_{j,t}(1-\lambda) + \sigma_{ij,t-1}\lambda \quad \text{for } t=1,\dots,T \text{ and } \forall i, j \neq i$$

where $\sigma_{i,0}^2 = E(\epsilon_i^2)$ and $\sigma_{ij,0} = E(\epsilon_i\epsilon_j)$.

The conditional correlation is obtained as

$$\rho_{ij,t} = \frac{\sigma_{ij,t}}{\sigma_{i,t}\sigma_{j,t}}$$

4.2.2 Closed-form solutions for the risk measures under Gaussian distribution

The joint model can be expressed in matrix terms as:

$$\begin{pmatrix} x_{m,t} \\ x_{i,t} \end{pmatrix} = \begin{pmatrix} \mu_{m,t} \\ \mu_{i,t} \end{pmatrix} + \begin{pmatrix} \sigma_{m,t} & 0 \\ 0 & \sigma_{i,t} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ \rho_{im,t} & \sqrt{1-\rho_{im,t}^2} \end{pmatrix} \begin{pmatrix} \xi_m \\ \xi_i \end{pmatrix} \quad (4.1)$$

where $\mu_{k,t}$, $\sigma_{k,t}$ and $\rho_{im,t}$ are the parameters from the previous section and $\begin{pmatrix} \xi_m \\ \xi_i \end{pmatrix}$ is a 2x1 vector of normal independent variables.

The conditional mean return of variable i is

$$E(x_i | x_m \leq F_m^{-1}(p)) = \mu_{i,t} - \frac{\sigma_{i,t}\rho_{im,t}\varphi(\Phi^{-1}(p))}{p}$$

where φ is the probability density function and Φ^{-1} is the inverse Gaussian distribution probability.

The CoVaR of variable i is

$$CoVaR_{i|m,t}(p, \alpha) = \mu_{i,t} - \sigma_{i,t} \frac{\phi_{im,t} \varphi(\Phi^{-1}(p))}{p} + \sqrt{1 - \rho_{im,t}^2} \Phi^{-1}(\alpha)$$

The *CoES* of variable i is

$$CoES_{i|m,t}(p, \alpha) = \mu_{i,t} - \sigma_{i,t} \left(\frac{\sqrt{1 - \rho_{im,t}^2} \varphi(\Phi^{-1}(\alpha))}{\alpha} + \frac{\phi_{im,t} \varphi(\Phi^{-1}(p))}{p} \right)$$

Annex 2 contains the proofs for these closed-form solutions.

5. Practical examples and use for very granular scenarios

5.1 Equity price shocks

As an example, a scenario with a joint financial distress for European stock markets is considered. Following the bootstrapping procedure in Step 1 (see Section 4.1) we have a large number of observations for the variables that gather the cross-sectional correlation. The “conditioning” variable is then constructed as a weighted average of the daily returns of stock prices. The triggering variable includes only domestic stock indices, such as the CAC 40 or the IBEX35, for example, and equal weights are assumed.

Once we have obtained the triggering variable, we check the results that fall below its p 100th lowest value, where p defines the degree of stress for the scenario. Stock prices normally decline in a downturn so the lower the value of p , the greater the scenario severity.

The values below the threshold given by the p 100th percentile belong to the simulations that are within the financial distress scenario. For those simulations, we assess the mean daily rates of change for the response variables, e.g. the stock returns in the different European countries. By doing this assessment we are gathering the higher-moment features of the joint distribution, i.e. skewness and kurtosis, which captures the joint co-movement in the extreme scenarios.

Table 1 shows the overall shocks for the five largest economies in the euro area triggered by a shock to the entire European stock market. We can observe that both for the parametric FSS and the non-parametric FSS the conditional mean and the *VaR* must be, by definition, lower in absolute values than the values of *CoVaR*. This is because in *CoVaR* we are not simply considering an adverse scenario for the domestic stock indices, as in *VaR*, or an adverse scenario for the entire European stock market, as in the conditional mean, we are computing the 10% worst return for the domestic indices given a bearish scenario across the entire European stock market. Similarly, the *CoES* must be higher in absolute terms than the *CoVaR* because we calculate the mean return for the domestic indices if those returns are worse than their *CoVaR*. Indeed the reasoning is the same as for the unconditional measure, i.e. *VaR* and *ES*.

With regard to the comparison between methodologies within the same measure, the values obtained from the non-parametric copula (NPC) are generally higher in absolute terms than those coming from the parametric copula (PC). The *VaR* output depicts a slightly higher *VaR* under the PC than under the NPC, which may be due to the effect of the conditional variance. Nevertheless, as we consider lower quantiles for the risk measure, the effects of skewness and kurtosis in the marginal distribution are visible, producing higher values for the *ES* under the NPC than under the PC. When we assess the *CoVaR* or the *CoES*, apart from the effect of the higher-moment features which causes the distribution to exhibit an asymmetric and fat-tail behaviour, the joint tail dependence increases the correlation on the extreme scenarios. This is because the difference between the *CoES* and the *ES* is much greater when computed under the NPC than it is when computed under the PC.

Table 1 : Non-parametric copula (NPC) and parametric copula (PC) results for a shock in the European stock market

	Conditional mean return		VaR		CoVaR (A)		CoVaR (B)		ES		CoES (A)		CoES (B)	
	NPC	PC	NPC	PC	NPC	PC	NPC	PC	NPC	PC	NPC	PC	NPC	PC
p=0.1														
DE (DAX)	-32	-25	-22	-24	-41	-15	-53	-34	-36	-33	-52	-19	-69	-38
ES (IBEX)	-34	-26	-23	-26	-43	-22	-60	-37	-39	-36	-58	-26	-74	-41
FR (CAC)	-34	-26	-22	-25	-38	-19	-57	-34	-37	-34	-52	-22	-69	-37
IT (MIB)	-36	-28	-26	-28	-45	-26	-63	-40	-42	-38	-64	-30	-78	-44
NL (AEX)	-30	-23	-20	-22	-38	-17	-51	-30	-34	-31	-49	-21	-65	-33

Note: This table shows the results assessed with $p=0.1$ and, for the PC, a lambda value equal to 0.94.

CoVaR(A) and CoES(A) show the conditional measures where the joint probability of the scenario is 10% while CoVaR(B) and CoES(B) has a conditional probability of 10%, i.e. the joint probability is 1%. Changes in dependence could be measured as CoES(B)-ES but the probability of observing CoES(B) is lower than ES. On the other side, losses are equally probable in CoES(A) and ES.

Annex 1: Sample of the dataset

VARIABLE	SUBGROUP	COUNTRY	EXTRA LABEL
CBS	F, NF	AT, BE, CY, DE, DK, ES, FI, FR, HU, IE, IT, LU, NL, PT, RO, SE, U2, UK, EU, CZ, SK, EA	A, AA, AAA, BBB
CDS	CDX, Itraxx	US, ASIA, EU, AT, AU, BE, BR, CH, CN, DE, DK, ES, FR, GR, IE, IN, IT, JP, NL, NO, PT, RU, SE	OVERALL, HVOL, SENFIN, XOVER, SUBFIN
GYD		AT, BE, BG, CN, CZ, DE, DK, EA, ES, FI, FR, GR, HR, HU, IE, IT, JP, KR, LT, PL, PT, RO, SE, SI, SK, UK, ZA, LV, UA, US, CH, NL, IS, RU, AG, AU, BR, CA, CL, CO, HK, ID, IN, LU, MT, MX, MY, NO, NZ, PE, PH, PK, SG, TH, TR, VN, CY	10Y, 10Y, 15Y, 1M, 1Y, 20Y, 25Y, 2Y, 30Y, 3M, 3Y, 4Y, 5Y, 6M, 6Y, 7Y, 8Y, 9Y
SWAP		CA, CH, CL, CN, CZ, DK, EA, EU, HR, HU, IN, JP, MX, NO, PL, RO, SE, SG, UK, AU, CO, HK, KR, MY, RU, TH, TR, US, ZA, ID, PE	10Y, 1M, 1Y, 20Y, 2Y, 30Y, 3M, 3Y, 5Y, 7Y, 25Y, 35Y, 40Y, 50Y
DY		US, ASIA, EU, AT, AU, BE, BR, CH, CN, DE, DK, ES, FR, GR, IE, IN, IT, JP, NL, NO, PT, RU, SE, DY	
ESX		GOLD, BRENT, REIT, COM, HF, EME, WO, WTI, EU, CH, Global, SILB, US, PE, MM, FREIGHT, AG, AT, AU, B2, BE, BG, BR, CA, CL, CN, CO, CY, CZ, DE, DK, DKK, EE, ES, F2, FERT, FI, FR, GR, HK, HR, HU, ID, IE, IL, IN, IS, IT, JP, KR, LT, LU, LV, MT, MX, MY, NGAS, NL, NO, NZ, PL, PT, RO, RS, RU, SE, SG, SI, SK, TH, TR, UA, UK, ZA	
FX		AUDUSD, EURCHF, EURCZK, EURGBP, EURHRK, EURHUF, EURPLN, EURRON, EURRSD, EURRUB, EURTRY, EURUAH, EURUSD, GBPUSD, NZDUSD, USDARS, USDBGN, USDBRL, USDCAD, USDCHF, USDCLP, USDCNY, USDCOP, USDDKK, USDHKD, USDHRK, USDHUF, USDIDR, USDINR, USDJPY, USDKRW, USDMXN, USDMYR, USDNOK, USDPEN, USDPHP, USDPKR, USDPLN, USDRUB, USDSEK, USDSGD, USDTHB, USDTRY, USDUAH, USDZAR, USDRON	

Notes: CBS – Corporate debt; DY – Dividend yield; GYD – Government yield; ESX – stock quote

Annex 2: Mathematical proofs of the closed formulas for the conditional mean, the CoVaR and the CoES

First, the closed formula for the Expected Shortfall is proved. The other proofs are based on this.

$$\begin{aligned} ES_{m,t-1}(p) &= \frac{1}{p} \int_0^p -\mu_{m,t} - \sigma_{m,t} \Phi^{-1}(s) \, ds \\ &= -\mu_{m,t} - \frac{\sigma_{m,t}}{p} \int_0^p \Phi^{-1}(s) \, ds. \end{aligned}$$

Define a change of variable $s = \Phi(r)$, then $ds = \phi(r)dr$

$$\begin{aligned} \int_{-\infty}^{\Phi^{-1}(p)} r \phi(r) \, dr &= \int_{-\infty}^{\Phi^{-1}(p)} \frac{r}{\sqrt{2\pi}} \exp(-r^2/2) \, dr \\ &= \frac{1}{\sqrt{2\pi}} [-\exp(-r^2/2)]_{-\infty}^{\Phi^{-1}(p)} \\ &= -\phi(\Phi^{-1}(p)). \end{aligned}$$

Given equation (4.1), the distribution for assets i given the distribution of the financial shock m is

$$r_{i,t}|r_{m,t} \sim N\left(\mu_{i,t} + \frac{\sigma_{i,t}\rho_t}{\sigma_{m,t}}(r_{m,t} - \mu_{m,t}), \sqrt{1 - \rho_t^2}\sigma_{i,t}\right)$$

If the realisation of the shock is expressed in terms of quantiles, i.e. $r_{m,t} = \Phi^{-1}(q)\sigma_{m,t} + \mu_{m,t}$, the mean value of i given that the financial variable is below its p 100% worst case scenario is

$$E_{t-1}(r_{i,t}|r_{m,t} < VaR_{m,t}(p)) = \mu_{i,t} + \sigma_{i,t}\rho_t \underbrace{\frac{\int_0^p \Phi^{-1}(q) dq}{p}}_{E\left(\frac{r_{m,t} - \mu_{m,t}}{\sigma_{m,t}} | r_{m,t} < VaR_{m,t}(p)\right)}$$

With the proof of the closed formula of the Expected Shortfall, the conditional mean is obtained.

The CoVaR can be derived by obtaining the quantile of the conditioned distribution of i given that the benchmark financial variable m is below its $VaR_m(p)$. Finally, the CoES can be obtained by assessing the integral of the CoVaR.

Annex 3: Comparison between the copula results and the historical distribution

The following table compares the results of a simulation of a distress scenario with historical data.

The scenario (last column) is run over the sample 01/01/2004-31/12/2018 assuming a joint decline of the EU stock prices, where the conditioning variable is the joint distribution of stock price returns in the EU and the response for each country is given by the *CoES* of each domestic stock index assuming that the joint probability for the scenario is 1/1000. The time horizon of the scenario is five days and it is calculated by bootstrapping each day over the five-day horizon and then computing the cumulative returns over the five-day period, a process which is repeated 100,000 times. The scenario results are compared with some statics of the original series. (maximum, minimum, second maximum, second minimum, percentile 0.1%).

The period 2004-2018 contains about 3500 observations and therefore the shock should be comparable to the historical minimum of the distribution. However, we also want to illustrate that repeating the bootstrapping such a large number of times means that the scenario results do reflect the historical properties of the series and that the resampling does not inflate the results.

	2-day					5-day					Shocks scenario 5 days
	Min	2nd min	Max	2nd max	Percentile 0.1%	Min	2nd min	Max	2nd max	Percentile 0.1%	prob 0.1%
AT	-13.1%	-12.9%	19.6%	18.6%	-12.6%	-28.9%	-18.9%	27.4%	22.1%	-18.4%	-17%
BE	-12.1%	-10.9%	9.5%	9.1%	-10.5%	-23.0%	-18.2%	17.7%	17.2%	-13.7%	-13%
BG	-17.6%	-12.9%	13.1%	10.4%	-12.0%	-29.4%	-24.6%	22.1%	17.5%	-22.8%	-15%
CY	-20.0%	-17.8%	29.5%	25.0%	-16.7%	-24.6%	-22.7%	44.8%	44.6%	-21.3%	-18%
CZ	-17.8%	-15.6%	22.1%	13.2%	-14.0%	-26.3%	-19.0%	28.8%	16.9%	-16.8%	-20%
DE	-11.1%	-9.6%	14.4%	12.3%	-8.9%	-21.6%	-15.5%	16.1%	16.0%	-14.0%	-14%
EE	-11.0%	-10.3%	14.5%	13.7%	-9.0%	-18.3%	-16.1%	17.3%	16.5%	-14.5%	-12%
ES	-14.0%	-12.6%	13.6%	11.6%	-10.1%	-21.2%	-19.5%	23.0%	15.3%	-14.7%	-18%
FI	-10.8%	-10.2%	11.4%	11.2%	-8.8%	-16.5%	-15.2%	16.1%	14.2%	-13.8%	-16%
FR	-12.3%	-10.8%	14.2%	11.4%	-8.6%	-22.2%	-13.8%	18.5%	15.0%	-13.1%	-15%
GR	-15.9%	-14.7%	16.4%	15.7%	-13.2%	-20.2%	-19.2%	22.6%	20.7%	-18.4%	-20%
HR	-12.2%	-10.7%	25.0%	10.9%	-10.4%	-25.5%	-19.2%	14.8%	14.7%	-17.9%	-14%
HU	-19.5%	-13.9%	24.6%	16.5%	-11.0%	-24.9%	-23.6%	32.0%	25.5%	-19.0%	-19%
IE	-16.9%	-16.3%	13.1%	10.6%	-14.0%	-27.2%	-20.5%	22.0%	17.9%	-17.2%	-20%
IT	-15.9%	-11.7%	15.6%	11.5%	-10.3%	-21.6%	-15.6%	24.0%	13.3%	-14.2%	-19%
LT	-16.0%	-15.5%	20.0%	14.6%	-10.8%	-25.5%	-18.8%	28.2%	24.3%	-16.6%	-14%
LU	-15.1%	-14.0%	10.1%	10.0%	-12.6%	-23.9%	-20.2%	18.8%	15.0%	-18.3%	-17%
LV	-13.2%	-12.2%	14.4%	12.0%	-10.9%	-20.1%	-14.4%	20.7%	18.8%	-13.7%	-13%
MT	-7.3%	-6.4%	8.3%	6.4%	-5.9%	-8.9%	-7.9%	11.4%	11.1%	-7.0%	-6%
NL	-12.8%	-12.4%	10.4%	10.3%	-10.1%	-25.0%	-14.8%	22.3%	15.1%	-13.9%	-16%
NO	-14.7%	-12.7%	11.3%	11.0%	-12.3%	-20.9%	-18.2%	25.7%	21.8%	-16.0%	-16%
PL	-9.7%	-9.7%	8.8%	8.5%	-7.9%	-16.1%	-15.7%	14.4%	13.3%	-14.0%	-14%
PT	-11.2%	-11.0%	15.1%	9.6%	-8.5%	-18.6%	-17.2%	19.9%	13.3%	-14.8%	-16%
RO	-15.8%	-13.7%	20.0%	13.3%	-13.1%	-27.0%	-25.0%	18.5%	16.2%	-19.0%	-19%
SE	-9.4%	-8.8%	12.6%	11.7%	-8.5%	-20.2%	-13.7%	22.5%	16.0%	-13.3%	-13%
SI	-11.5%	-11.2%	15.1%	10.7%	-10.4%	-17.8%	-17.5%	16.8%	12.1%	-16.6%	-14%
SK	-13.8%	-10.4%	13.9%	12.6%	-9.5%	-14.1%	-13.8%	14.8%	14.8%	-12.5%	-15%
UK	-24.0%	-11.9%	35.6%	10.1%	-10.6%	-21.8%	-21.5%	35.8%	18.9%	-14.6%	-22%
CH	-14.1%	-12.0%	17.1%	9.6%	-9.6%	-22.3%	-13.8%	15.6%	14.1%	-13.0%	-14%

6. Bibliography

- Acerbi, C. and Tasche, D. (2002), "On the coherence of expected shortfall", *Journal of Banking & Finance*, Vol. 27, No 7, pp. 1487-1503.
- Acharya, V., Engle, R. and Richardson, M. (2012), "Capital shortfall: A new approach to ranking and regulating systemic risks", *American Economic Review*, Vol. 102, No 3, pp. 59-64.
- Adrian, T. and Brunnermeier, M. K. (2016), "CoVaR", *American Economic Review*, Vol. 106, No 7, pp. 1705-1741.
- Bassanin, M., Rancoita, E. and Ojea Ferreiro, J. (2019), "Design of scenarios for the calibration of macroprudential policies", forthcoming.
- Brownlees, C. and Engle, R. (2016), "SRISK: A conditional capital shortfall measure of systemic risk", *The Review of Financial Studies*, Vol. 30, No 1, pp. 48-79.
- Cherubini, U., Gobbi, F. and Mulinacci, S. (2016), "The dynamics of Economic Variables", in Cherubini, U., Gobbi, F. and Mulinacci, S., (eds.), *Convolution Copula Econometrics*, Springer International Publishing, Basel, pp. 4-5.
- Girardi, G. and Ergün, T. (2013), "Systemic risk measurement: Multivariate GARCH estimation of CoVaR", *Journal of Banking & Finance*, Vol. 37, No 8, pp. 3169-3180.
- Rancoita, E. and Gross, M. (2015), *Scenarios for the European Insurance and Occupational Pensions Authority's EU-wide pension fund stress test in 2015*, European Systemic Risk Board, Frankfurt am Main, March.
- Rancoita, E., and Gross, M. (2017), *Adverse scenario for the European Insurance and Occupational Pensions Authority's EU-wide pension fund stress test in 2017*, European Systemic Risk Board, Frankfurt am Main, March.
- Sklar, M. (1959), "Fonctions de repartition an dimensions et leurs marges", *Publications de l'Institut Statistique de l'Université de Paris*, Vol. 8, pp. 229-231.