Zeliade Systems

Analytical scores for stress scenarios

Zeliade Systems

Title:	Analytical scores for stress scenarios
Author:	Zeliade Systems
Number of pages:	11
First version:	03/07/2020
Current version:	03/07/2020
Revision:	1.0

Contents

In this work, inspired by the Archer-Mouy-Selmi approach [2], we present two methodologies for scoring the stress test scenarios used by CCPs for sizing their Default Funds. These methodologies can be used by risk managers to compare different sets of scenarios and could be particularly useful when evaluating the relevance of adding new scenarios to a pre-existing set.

1. Comparison of sets of Hypothetical stress scenarios

After the financial crisis of 2008, the topic of stress testing got more and more attention from the financial infrastructures environment, specifically from Central Clearing Counterparties (CCPs). Both the PFMI-IOSCO [7] and EMIR regulations [6], [4] require CCPs to specify *extreme but plausible scenarios* for the sizing of their default funds. Moreover the EMIR regulation requires a review of the same stress scenarios (at least annually according to Article 31 of EMIR RTS [4] and according to the key explanations of Principles 4, 5, and 7 of PFMI [7]). As such a risk manager may face the conundrum of assessing the benefits of adding or changing sets of scenarios. In this work we provide two methodologies giving a quantitative assessment of the advantages or disadvantages of the new set of scenarios.

1.1 Plausible Hypothetical Scenarios

A hypothetical scenario is, as the name suggests, a stylized scenario designed in order to capture a tail risk. Hypothetical stress scenarios can be designed starting from risk manager views, as for example a parallel shift of all yield curves for fixed income products, or using as a base a scenario obtained via quantitative methods, coming from the fit of a distribution, or a PCA.

By construction, the hypothetical scenarios do not refer directly to the history. So, it is a good indicator of the scenarios quality to estimate their plausibility. Moreover, the size of the hypothetical scenarios are generally calibrated separately for each risk factor, so that the plausibility of the joint scenarios does not have the same confidence level as the one dimentional quantile level. The plausibility can be estimated, for instance, by evaluating the log-likelihood of the hypothetical scenarios, using a joint-distribution calibrated on the risk factors. A joint Student distribution can be calibrated on the historical data of the risk factors, then used to estimate the likelihood of the scenario.

The problem of the design of extreme but plausible scenarios has been tackled in the literature, for example by Thomas Breuer and co-authors in [3]. They look for the scenario that gives the worst loss, under a plausibility constraint (likelihood less than a cap). While this approach is valid for portfolio management, it is not well suited for the Stress Testing context of CCPs. In their second paper [5], the authors consider a dual problem whose solution does not depend on additional dimensionality of the problem and which closely resembles the problems faced by a risk manager at a CCP. This approach is made even more explicit by Q.Archer, P.Mouy and M.Selmi (LCH), who proposed (cf. [2]) a framework for the design of extreme but plausible scenarios. Their methodology considers a linear portfolio P, so that the P&Ls from risk factor returns s is simply $P^t s$, and assumes a calibrated historical distribution of the risk factors, with density f.

The approach can be formalized by the following maximum likelihood problem:

$$\max_{P^t s \leq q} f_\theta(\mathrm{s}) \tag{1}$$

where

- s is the vector of the risk factors returns.
- P is the portfolio positions.
- f_θ is the density function of the joint-distribution of all the risk factor returns s.
- q is a cap constraint on the portfolio loss.

For instance a heavy-tail distribution, like a T- Student law can be calibrated on the risk factors returns. q can be chosen as the α-quantile of the distribution of the portfolio loss $P^t s$. For example, if s is a Gaussian vector, then $P^t s$ is Gaussian so that q can be chosen as a (one-dimensional) Gaussian quantile.

This methodology has the advantage of yielding extreme scenarios, with a loss cap, so that we are ensured to have meaningful scenarios. Moreover, if the distribution is chosen to be a standard elliptic (Student or Gaussian, with correlation Σ, null average and marginal standard deviation 1), then the problem (1) admits a simple closed formula solution:

$$S^*(P) = q \frac{\Sigma P}{{}^t P \Sigma P}. \tag{2}$$

As pointed out in [2] and [5] another advantage of this methodology is the fact that, at least in an elliptical setting, the solution of (1) is not dependent on the presence of additional risk factors not appearing in the portfolio, while the primal problem considered in [3] gives losses dependent on the introduction of risk factors unrelated to the portfolio.

1.2 The score functions

We apply now the ideas from [2] to the comparison of sets of stress scenarios.

We start by considering a single linear portfolio P (such that the P&L associated to the risk factor return s is ${}^t P s$) and a set of stress scenarios $\mathcal{S} := \{S_0, \dots, S_n\}$ for the risk factor returns s.

We can thus calculate the stress loss associated to the portfolio P as:

$$l(P) := \min_{0 \le i \le n} {}^t P S_i \,,$$

we then select the scenario $\hat{S}(P)$ among $\mathcal{S}$ that drive $l(P)$. If several drivers are found, we select the driver that maximize the density function (another criterion could be applied, such as minimizing the Mahalanobis distance ${}^t\hat{S}(P)\Sigma^{-1}\hat{S}(P)$ among the candidate drivers).

We then compute the scenario $S^*(P)$ solving

$$\sup_{S \,;\, {}^t P S \le {}^t P \hat{S}} f_\theta(S) \,, \tag{3}$$

that is to say the most plausible scenario generating a loss equal (or greater) to the worst loss obtained with the stress scenarios. Equivalently, the two scenarios $\hat{S}(P)$ and $S^*(P)$ generate the same loss $l(P)$ but $S^*(P)$ is the most plausible with respect to the distribution assumption.

We finally introduce the two score functions. The first score function measures the quality of the ratio loss to plausibility of $\hat{S}(P)$, and is given by

$$\phi_{\mathcal{S}}(P) := \frac{f_\theta\left(\hat{S}(P)\right)}{f_\theta\left(S^*(P)\right)} \in]0,1] \tag{4}$$

The higher the score, the better it is, as a high ϕ indicates that the stress scenarios contained in the set $\mathcal{S}$ are close, in a plausibility sense, to the most likely scenario inducing the same level of losses for the portfolio.

The second score is a geometrical criterion, measuring to what extent a driver is in the same direction as the optimal scenario.

$$\psi_{\mathcal{S}}(P) := \frac{\langle \hat{S}(P), S^*(P) \rangle}{\left\|\hat{S}(P)\right\| \left\|S^*(P)\right\|} \in]-1,1] \tag{5}$$

Also in this case the higher the score, the better it is, as it indicates that the "risk direction" of the portfolio is captured by the stress scenario set $\mathcal{S}$.

1.3 Applying the scores to sets of scenarios

Suppose now that we have two sets of stress scenarios: $\mathcal{S} = \{S_0, \cdots, S_n\}$,and $\mathcal{T} = \{T_0, \cdots, T_m\}$, possibly partially overlapping, and we want to evaluate the advantages of one set with respect to the other. The score functions we introduced in the previous section allows us to do it in the following way:

- Select a reference set of portfolios $P_0, \ldots, P_M$.
- Calculate the values $\phi_{\mathcal{S}}, \psi_{\mathcal{S}}(P_i)$'s and $\phi_{\mathcal{T}}, \psi_{\mathcal{T}}(P_i)$'s.
- Define a final score from those values.

The choice of the final score depends on the risk manager view. In our numerical result parts we propose two different approaches.

- **Scenario approach**, which is particularly meaningful when the set $\mathcal{T}$ is a modification of the set $\mathcal{S}$. For each stress scenario we compute the average and standard deviation of the function ψ and ϕ on the set of portfolios for which the scenario is the driver. This approach allows to have a view on which scenarios could be eventually modified, or even eliminated as being very far from optimality, either from a plausibility or geometrical points of view.
- **Portfolio approach**. We compare the score functions on each portfolio. This approach allows to better understand for which portfolios the risk is not correctly sized, and it can be used for understanding which test portfolios are not sufficiently stressed by the current set of stress scenarios.

We point out that the proposed scores should be used more as a non rejection indicator and not as an acceptance one, similarly to the Kupiec Test which says that an hypothesis can not be rejected, not that it should be accepted.

1.4 Finding the optimal scenario

For elliptical distributions, the solution of (3) can be found exactly as described in [2]. However, for the more generic meta-elliptical distributions (introduced in [1]) this is no more the case. As these are the distributions we will fit the risk factors returns on, we provide two possible alternatives for finding the most plausible scenario at a given loss.

We recall that a meta-elliptical distribution f_θ is a multidimensional distribution with elliptical copula. The setting we will consider consists of a T-Student copula with T-Student marginals, and it is consequently characterized by a parameter θ^* containing:

- the location vector μ and scale vector σ
- the correlation matrix Σ
- the vector of marginal degrees of freedom ν
- the degrees of freedom $\bar{\nu}$ of the copula(denoting also in the sequel the d-dimensional constant vector $(\bar{\nu}, \dots, \bar{\nu})$).

1.4.1 Approximate solution

The first method was proposed by Mouy et al. [2] and it is based on approximating the meta-elliptical distribution by an elliptical distribution, i.e. using the same degree of freedom for the copula and the marginals.

We start by normalizing the distribution, via the linear transformation $\tilde{s} := (s - \mu)/\sigma$, and we get the equivalent problem

$$\sup_{\tilde{s};{}^t\tilde{P}\tilde{s} \leq \tilde{q}} \tilde{f}_{\theta^*}(\tilde{s})$$

where

- $\tilde{f}_{\theta^*}(\tilde{s}) := f_{\theta^*}(s)$
- $\tilde{P} := \sigma P$
- $\tilde{q} := {}^tP(\hat{S} - \mu)$.

If the distribution $\tilde{f}_{\theta^*}$ was elliptical, the optimal scenario for the problem above would be given by (2)

$$\tilde{S}^*(P) = \tilde{q}\frac{\Sigma\tilde{P}}{{}^t\tilde{P}\Sigma\tilde{P}},$$

Transporting it back to the original problem, one has:

$$S^*(P) := \mu + \sigma\left(\tilde{q}\frac{\Sigma\tilde{P}}{{}^t\tilde{P}\Sigma\tilde{P}}\right).$$

As stated above, the approximated solution is obtained by approximating the meta-elliptical distribution with an elliptical distribution, obtaining the sub-optimal scenario

$$\bar{S}(P) := \mu + \sigma T_{\nu}^{-1} \circ T_{\bar{\nu}} \left(\tilde{q} \frac{\Sigma \tilde{P}}{{}^t\tilde{P}\Sigma\tilde{P}} \right)$$

where $T_\nu(x)$ is the vector $(T_{\nu_i}(x_i))_{1\leq i\leq d}$, T_{ν_i} being the CDF of a standard T-Student distribution with ν_i degrees of freedom.

We point out that the approximation quality is strongly linked to the "almost linearity" of the function $T_{\nu}^{-1} \circ T_{\bar{\nu}}$ around 0. In the case where $\bar{\nu}$ and ν are very different, the approximation could be poor, with significant discrepancies both in term of optimal density value and on loss constraint violation.

1.4.2 Exact numerical solution

An exact solution can also be recovered numerically using classical optimizers. In fact the target function is easy and fast to calculate and the constraint is linear.

Moreover, as the applications for which our methodologies are devised require the score calculation to be done once for all or at low frequency so using a time consuming resolution method is not an issue.

Finally, to calculate a score, it is possible to restrict the relevant portfolios to involve a low number of risk factors (e.g. spreads, involving each only 2 risk factors). The effective dimension of the exact resolution can then be lowered and the optimization made easier.

2. Numerical experiments

We thus performed our experiments on the synthetic Yields curves provided by the European Central Bank and downloadable at http://sdw.ecb.europa.eu/:

- **AAA**: synthetic curves aggregated from the AAA issuers of the EURO zone (dynamic basket).
- **ALL**: synthetic curves aggregated from all the issuers of the EURO zone (dynamic basket).

We used the pillars $6M, 1Y, 2Y, 3Y, 4Y, 5Y$ of those yield curves.

We assume the following setting:

- **Reference set of portfolios**: we consider spread portfolios of the form $(B_i, -B_j)$ where:
  - B_i, B_j are some Bonds with semi-annual coupons, with a time-to-maturity equal to one of the pillars' maturity
  - $\beta = -1$ and $\beta = -D_i/D_j$.
  - D_i, D_j are the durations of the bonds B_i and B_j.
  - we obtain $2 \times 2 \times \left(\frac{12\times 13}{2}\right) = 264$ portfolios.
- **Distribution assumption**: a meta-t distribution on the Yield rate returns.
- **Bond pricing**: we approximate the P&L for a bond to be $(\Delta Y)D$ where ΔY is the Yield rate move and D the base bond duration.

We will obtain the optimal scenarios via numerical optimization.

2.1 Low Dimensional or Full Risk setting

Should we fit a single distribution on all the risk factor at the same time, or one on each single test portfolio? While for Gaussian distribution this does not have an impact, in the case of meta-elliptical distributions that we are considering, the situation is a little bit more complicated. This is because, while the marginal distributions are fixed, the copula can vary.

From a stability point of view, our decision makes the scores dependent on the number of risk factors chosen, as the copula is fitted on each group separately. However, we believe that these additional degrees of freedom allow to better measure the risk and give a better understanding of the differences between sets of stress scenarios.

2.2 The Stress Scenarios

We consider 2 sets of stress scenarios: a base and an enriched one. As the scenario generation methodology is not the focus of this work we decided to use over-simplified and stylized sets. Moreover, this choice allows us to better highlight the contribution of our scores, as the difference between the sets havealso a clear interpretation.

Both the base and the enriched set are obtained starting from the first three components of a Principal Component Analysis performed separately on the returns of the **AAA** and **ALL** curves. The vectors are rescaled by a factor $3 \times \sigma_i$ where σ_i is the explained standard deviation associated and combined as follows:

- The **base set** $\mathcal{S}$ considers only combination of the same level, and with the same sign: ($\pm n^{th}$ component **AAA**, $\pm n^{th}$ component **ALL**), $n = 1, 2, 3$, for a total 6 possible stressed scenarios;
- The **enriched set** $\mathcal{S}'$ considers the scenarios in the base set, plus the combination given by ($\pm n^{th}$ component **AAA**, $\mp n^{th}$ component **ALL**), $n = 2, 3$, for a total 10 possible stressed scenarios.

We plot here the three drivers of the risk scenarios for the two curves. Notice that the PCA analysis provides three main directions that indeed qualitatively correspond to the shift, slope change and curvature change (displaying respectively the sign patterns $+$, $-/+$ and $+/-/+$) often used in devising hypothetical scenarios for rate curves.

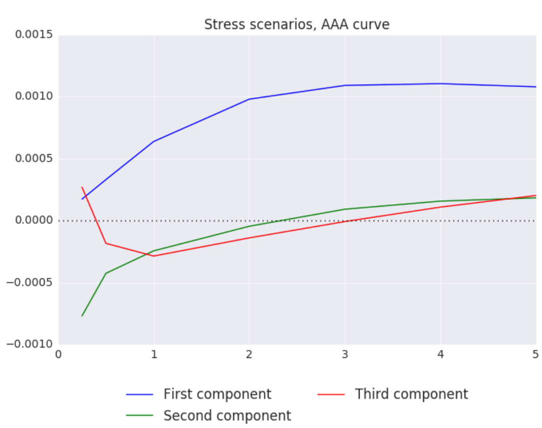

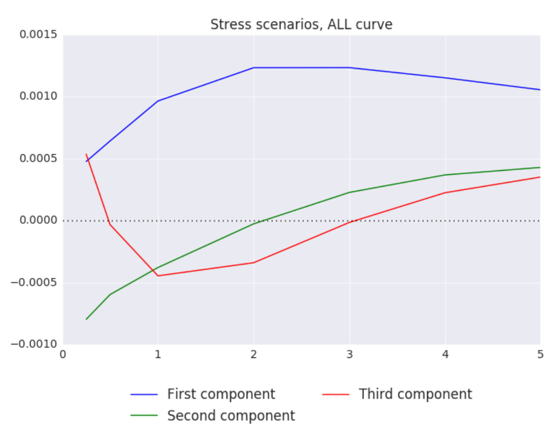

2.3 Result analysis

In the table below we compare the average and standard deviation for the loss to plausibility $\phi.(P)$ and the geometrical score $\psi.(P)$ aggregated at the driving scenario level, and considered as a whole (last line). The columns *Quantity* indicates for how many portfolios the selected scenario produces the largest losses.

	Base Scenarios					Enriched Scenarios				
Scenario	Quantity	ϕ mean	ϕ std	ψ mean	ψ std	Quantity	ϕ mean	ϕ std	ψ mean	ψ std
(+1, +1)	80	0.334	0.258	0.858	0.193	64	0.929	0.085	0.397	0.248
(-1, -1)	80	0.348	0.263	0.872	0.171	64	0.934	0.079	0.414	0.250
(+2, +2)	28	0.442	0.244	0.512	0.331	17	0.537	0.324	0.381	0.256
(-2, -2)	28	0.437	0.246	0.497	0.343	17	0.522	0.330	0.374	0.254
(+3, +3)	24	0.666	0.259	0.746	0.274	21	0.733	0.285	0.640	0.261
(-3, -3)	24	0.658	0.262	0.736	0.283	21	0.723	0.293	0.632	0.263
(+2, -2)	0	0	0	0	0	23	0.868	0.218	0.797	0.236
(-2, +2)	0	0	0	0	0	23	0.869	0.225	0.802	0.237
(+3, -3)	0	0	0	0	0	7	0.879	0.131	0.839	0.142
(-3, +3)	0	0	0	0	0	7	0.867	0.148	0.834	0.143
Total	264	0.420	0.285	0.766	0.282	264	0.530	0.301	0.833	0.243

Table 1: Comparison of the scores on the different scores obtained with the different stress scenarios sets.

We can see that the introduction of the new scenarios does not deteriorate significantly the score of the existing scenarios (it actually improves it in some cases), and that the new scenarios have average scores quite elevated.

It is up to the risk manager to decide if, according to his/her expertise, the new scenarios are acceptable or not. Particular attention should be paid in case the scores of the new scenarios are high, but the average score for some of the old scenarios has been lowered. Again, we would like to highlight that our scores are not to be intended as an acceptance tool, but more as a non rejection one.

We have compared above the score functions ψ and ϕ at a scenario level. In the figure below we compare the two scores at a portfolio level.

- **Left graph**: the functions $P \to \phi_{\mathcal{S}}(P)$ (blue), and $P \to \phi_{\mathcal{S}'}(P)$ (green). The two functions are plotted with the test portfolios in the ascending order for $P \to \phi_{\mathcal{S}}(P)$.
- **Right graph**: the functions $P \to \psi_{\mathcal{S}}(P)$ (blue), and $P \to \psi_{\mathcal{S}'}(P)$ (green). The two functions are plotted with the test portfolios in the ascending order for $P \to \psi_{\mathcal{S}}(P)$.

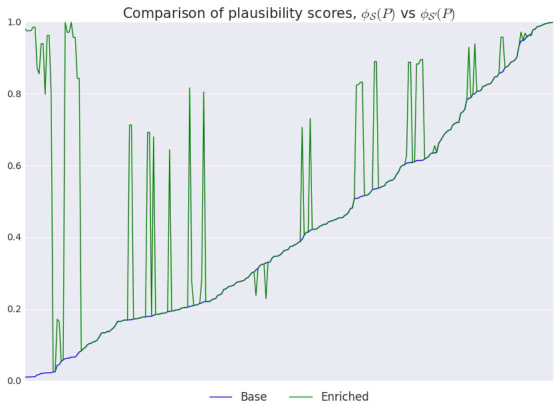

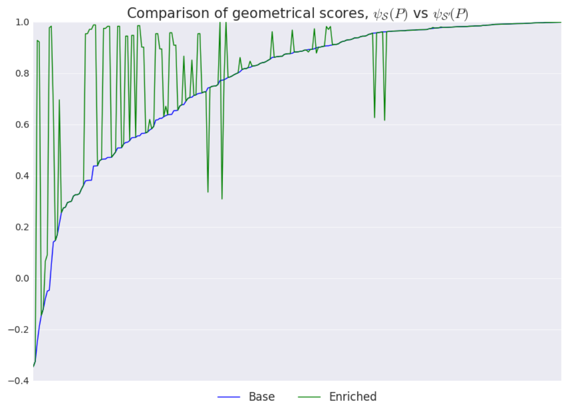

We can see that, with the exception of very few portfolios, the scores obtained by the enriched scenarios are higher than the one obtained by the base set (when the score is the same the scenario driving the stressed value is the same in the two sets).

2.4 Comparison of probability

One natural question is whether or not the new scenarios, even if obtaining better scores for the analyzed portfolios, are plausible enough.

A priori some of the new scenarios may be very close or even coincide with their corresponding most plausible scenarios, however, when compared with the scenarios previously available they may be way less plausible. A typical example could be a rescaling of one of the existing scenarios by a factor > 1. This could induce (for some test portfolios) bigger losses, but at the same time the stress scenarios would be less plausible.

A trade-off between scenario plausibility and losses may happen and it is up to the risk manager to analyze it and decide if it is acceptable or not, but it can also happen that the new drivers not only generate bigger losses but are also more probable, simply because they explore new direction with respect to the old ones and result in a more significant position with respect to the reference portfolios.

In the figure below we present the two cases:

- **Left Graph**: Long one bond **AAA** with maturity $3Y$, short one **ALL** with maturity $3Y$. In this case the new driver not only provides a higher loss but also has a higher probability.
- **Right Graph**: Long one bond **ALL** with maturity $6M$, short one **AAA** with maturity $6M$. In this case the new driver there is a trade-off higher losses lower plausibility.

The dashed lines represent the level line of the distribution on the scenarios (either driver or optimal).

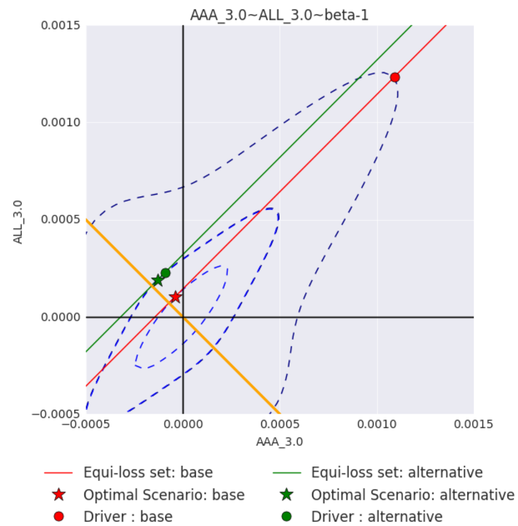

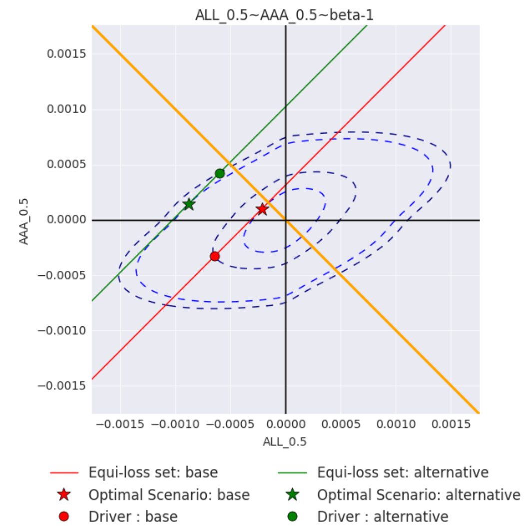

The graph on the left shows the most interesting case from a risk manager point of view. The original set of scenario did not tackle in a good way the risk associated with the specific test portfolio, as the driver is orientated almost perpedicularly with respect to the optimal scenario. The enlarged set introduces a scenario which is in a "good direction" respect to this specific test portfolio, providing a larger loss (≈ 2.2 times the original one) while being simultaneously many times more plausible (the density ratio between the two driving scenarios is around 19).

The graph on the right shows a less appealing case: the enlarged set introduces a scenario which generates bigger losses but which is, at the same time, more unlikely (even if both scores are higher in this case). The risk manager will thus have to decide, based on his expertise, if the trade off plausibility vs. risk is acceptable or if the new scenario (or set of scenarios) should not be taken in consideration.

3. Conclusions

In this work we have presented two methodologies which can help risk managers to compare sets of stress scenarios and in particular to assess the benefits of the introduction of new scenarios the existing ones. The two methodologies allow the risk manager to analyze different aspects of the stress scenarios, notably their position and relevance for the reference sets of portfolios.

The proposed methodologies have a clear and natural meaning which allows to better understand the benefit of one set of scenarios with respect to the other.

References

[1] The meta-elliptical distributions with given marginals, Fang, Hong-Bin and Fang, Kai-Tai and Kotz, Samuel, Journal of Multivariate Analysis, 2002.

[2] Extremely (un)likely: A plausibility approach to stress testing, Mouy, Pierre and Archer, Quentin and Selmi, Mohamed, Risk.net, 2017.

[3] What is a plausible Stress Scenario, Breuer, Thomas and Krenn, Gerald, Computational Intelligence: Methods and Applications, ICSC Academic Press, Canada/Switzerland, 2001.

[4] COMMISSION DELEGATED REGULATION (EU) No 153/2013 of 19 December 2012, supplementing Regulation (EU) No 48/2012 of the European Parliament and of the Council with regard to regulatory technical standards on requirements for central counterparties, European Commission, 2013.

[5] Overcoming dimensional dependence of worst case scenarios and maximum loss, Breuer, Thomas, Journal of Risk, 2008.

[6] Regulation (EU) 648/2012 of 4 July 2012 of the European Parliament and Council on OTC derivatives, central counterparties and trade repositories, European Commission, 2012.

[7] CPSS-IOSCO Principles for Financial Market Infrastructures, Bank for International Settlements and International Organization of Securities Commissions, 2012.