

A triptych approach for reverse stress testing of complex portfolios

Pascal Traccucci, Luc Dumontier, Guillaume Garchery and Benjamin Jacot present an extended reverse stress test (ERST) triptych approach with three variables: level of plausibility, level of loss and scenario. Any two of these variables can be derived, provided the third is given as input. A new version of the Levenberg-Marquardt optimisation algorithm is introduced to derive the ERST in certain complex cases

Introduction: the case of ARP portfolios

Academic theory has been mined to support the development of investment solutions containing an ever-increasing number of factors. Over the last decade, academics and practitioners have shown traditional asset classes offer limited diversification, especially in market downturns. In response, they have delved into modern portfolio theory (MPT) to identify the micro-economic factors that are the backbone of alternative risk premia (ARP) solutions. The ARP 1.0 approach combines 10–15 different long/short portfolios capturing standard investment styles such as value, carry, momentum, low risk and liquidity across a broad range of traditional asset classes. For further diversification, the ARP 2.0 approach combines up to 30 strategies by including investment banking-style premia likely to use instruments with quadratic profiles.

Many risk management frameworks cannot properly account for non-linear profiles and assess the risk of loss associated with combining an unusually high number of strategies. Specifically, historical value-at-risk is an instantaneous risk indicator and does not correspond to a clearly identified scenario; hence the need for complimentary stress tests. To build a stress-testing tool, the dataset must be simplified, and historical or predefined scenarios are used without quantifying their plausibility. Thus, parametric VAR imposes dependence on a model to benefit from an analysis framework in the form of a VAR and a sensitivity of this VAR to all the parameters of the model. This requires several numerical problems to be addressed, especially in case of quadratic profit and loss (P&L). This article presents an innovative approach: the extended reverse stress test (ERST), following on from the work of Breuer *et al* (2009) and Mouy *et al* (2017). This approach is able, with low technical costs,¹ to deliver two of three parameters, provided the third is given as input. The three parameters are scenario, level of plausibility and level of loss (see figure 1). The result is a more meaningful risk measure and one that corresponds to a clearly identified scenario.

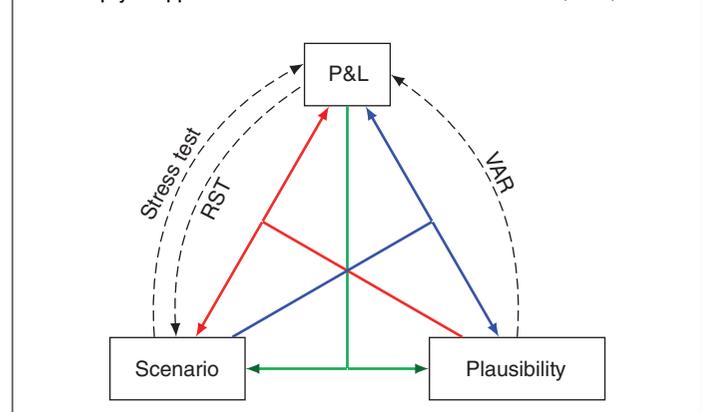
In what follows, \mathbf{S} is defined as a scenario. It is a vector with length n , which equals the number of risk factors to which the portfolio is exposed. In addition, the covariance matrix of the risk factors will be denoted by Σ .

Starting from a scenario

A scenario-driven ERST approach is suitable for a portfolio manager considering a given adverse or best-case scenario \mathbf{S}_0 . To assess the plausibility of such a scenario, the probability α_0 of a scenario being as extreme as or less extreme than \mathbf{S}_0 is computed. If α_0 is too high, $\tilde{\mathbf{S}}$, a more plausible scenario than \mathbf{S}_0 , is derived and suggested to the portfolio manager.

¹ Using an algorithm derived from the Levenberg-Marquardt one to deal with complex problems.

1 The triptych approach of the extended reverse stress test (ERST)



■ **Measuring plausibility.** The ERST relies heavily on the concept of plausibility (or likelihood) to discriminate between the scenarios generated. Multiple plausibility measures exist in the literature (Breuer *et al* 2009). In this article, plausibility is quantified in terms of the Mahalanobis distance. The latter measures the amplitude of the multivariate moves in \mathbf{S} from the mean scenario μ in units of standard deviation. It is therefore similar in a multi-dimensional space to the concept of a Z -score z or standardised variables. As a reminder:

$$z = \frac{x - \mu_X}{\sigma_X}$$

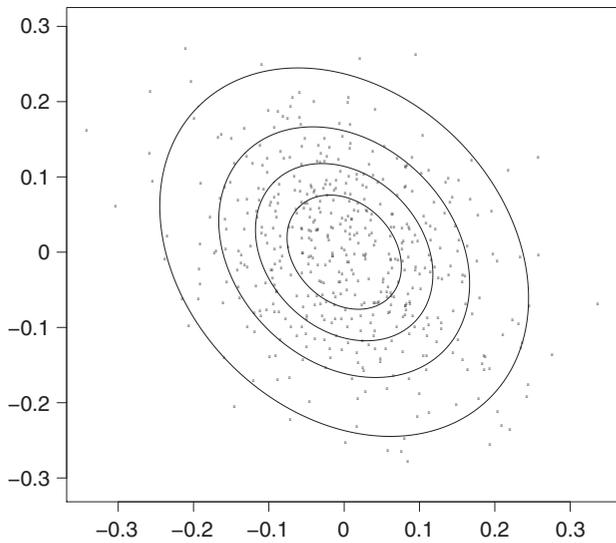
where x is the realisation of a random variable X with mean μ_X and standard deviation σ_X . The Mahalanobis distance is defined as follows:

$$\text{Maha}^2(\mathbf{S}) = (\mathbf{S} - \mu)^T \Sigma^{-1} (\mathbf{S} - \mu) \quad (1)$$

Unlike other measures, the Mahalanobis distance is both intuitive and simple to use. Its following characteristics are noteworthy:

- A low (respectively, high) Mahalanobis distance characterises a highly plausible (respectively, unlikely) scenario.
- $\text{Maha}^2(\mathbf{S}) = R^2$ is the surface of an ellipsoid of radius R . Points within the ellipsoid have a Mahalanobis distance of less than R . The further away these points are from the surface, the closer they are to the centre, and the more plausible they become.
- Assuming \mathbf{S} follows a multivariate normal distribution, $\text{Maha}^2(\mathbf{S})$ follows a $\chi^2(n)$ distribution, as proved in Studer (1997). The α quantile of a $\chi^2(n)$ density is thus the squared radius of the ellipsoid, where $\alpha\%$ of the multivariate normal scenarios \mathbf{S} remain inside. Hereafter, this ellipsoid is referred to as \mathcal{E}^α . See figure 2 as an illustration. Also, α is referred to as

2 Plausibility domains for a bivariate random variable with elliptical density



The inner (respectively, outer) ellipse corresponds to a 25% (respectively, 95%) quantile. The in-between ellipses correspond to 50% and 75% quantiles

the probability a scenario is as plausible as or more plausible than \mathbf{S} , or the probability of non-occurrence.

■ Mahalanobis distance is suited to any elliptical multivariate distribution for \mathbf{S} , which includes densities other than multivariate normal, for example, Student's t distribution. This is of primary importance as a distribution of this type is typically a better fit for historical distributions than a normal one, especially as concerns fat tails.

The plausibility of \mathbf{S}_0 can now be evaluated simply by using (1). Assuming a normal distribution, the resulting value is compared with the quantiles of a $\chi^2(n)$ to determine the probability of a scenario as extreme as or less extreme than \mathbf{S}_0 . For other elliptical distributions where the law of the Mahalanobis distance is not known, a numerical solution exists. First, the elliptical distribution that best fits \mathbf{S} is determined. Second, a Monte Carlo simulation of \mathbf{S} is performed. Then, approximate quantiles of the Mahalanobis distance are computed and the probability of non-occurrence $\alpha_{0, \text{approx}}$ of \mathbf{S}_0 can be deduced.

■ **Fitting the plausibility of a given scenario.** If α_0 or $\alpha_{0, \text{approx}}$ exceeds a given threshold α_{max} , then \mathbf{S}_0 lies outside of the admissible ellipsoid $\mathcal{E}^{\alpha_{\text{max}}}$. In this case, the closest admissible scenario $\tilde{\mathbf{S}}$ to \mathbf{S}_0 on $\mathcal{E}^{\alpha_{\text{max}}}$ is defined by uniform scaling. This definition results in minimal corrections with regard to \mathbf{S}_0 and thus adheres as closely as possible to the intuition of the portfolio manager.

For the sake of clarity, the non-constraining assumption $\boldsymbol{\mu} = \mathbf{0}$ is made. As $\tilde{\mathbf{S}} \in \mathcal{E}^{\alpha_{\text{max}}}$, it follows that $\tilde{\mathbf{S}}^T \boldsymbol{\Sigma}^{-1} \tilde{\mathbf{S}} = q^{\alpha_{\text{max}}}$, where $q^{\alpha_{\text{max}}}$ is the α_{max} quantile of the density of the squared Mahalanobis distance. This constraint leads to:

$$\tilde{\mathbf{S}} = \sqrt{q^{\alpha_{\text{max}}}} \frac{\mathbf{S}_0}{\sqrt{\mathbf{S}_0^T \boldsymbol{\Sigma}^{-1} \mathbf{S}_0}} \quad (2)$$

■ **Application.** A portfolio manager runs two long/short strategies, each based on a different spread: the first on equity indexes (S&P 500 versus Euro Stoxx 50) and the second on bonds (German Bund versus US Treasury).

Analysis of monthly prices from the previous five years (2014–19) shows low correlation ρ between spreads: $\rho \approx 0.01$. Over the same period, the spread returns have a monthly volatility of 3.3% and 1.2%, respectively.

The manager would like to know if, under these assumptions, there is a strong probability (50%, for example) the spread scenarios will incur 1.5% and 2.5% losses over one month, leading to the scenario:

$$\mathbf{S}_0 = [-1.5\%, -2.5\%] \quad (3)$$

In this example, it is assumed the spreads have either a normal or a Student's t distribution. The parameters of the elliptical distribution of reference are determined using maximum likelihood estimators derived from the historical distribution. Thus, \mathbf{S}_0 corresponds to $\alpha_0 = 91\%$ (respectively, 81%) for a normal (respectively, a Student's t) distribution. By way of comparison, the average monthly market correction observed during the fourth quarter of 2018 shows a probability of non-occurrence of approximately 77% (respectively, 69%). Therefore, the loss the manager had in mind is less plausible than expected. Setting $\alpha_{\text{max}} = 50\%$ in (2), the fitted scenario of interest for the manager is:

$$\begin{aligned} \tilde{\mathbf{S}} &= [-0.8\%, -1.3\%] \quad \text{for normal risk factors} \\ &= [-0.9\%, -1.4\%] \quad \text{for Student's } t \text{ risk factors} \end{aligned} \quad (4)$$

Thus, the fitted scenarios respect the directions intended by the portfolio manager, and only the amplitude of the shocks is changed to comply with the constraint α_{max} .

Obviously one can argue the correlation and volatility used in this example do not reflect a crisis environment where \mathbf{S}_0 occurs. An extension of this example would therefore be to stress the correlation and volatility to best reflect a financial crisis environment. This process is further explained in Traccucci *et al* (2019).

Starting from plausibility

The plausibility-driven ERST returns both the most extreme loss and a corresponding scenario for a given level of plausibility. This approach is studied in Studer (1997) and further discussed in Breuer *et al* (2009), for example. Its advantage is it returns a loss that may be compared with other existing risk measures such as VAR, which is briefly introduced in the following subsection. As shown in the rest of this section, a plausibility-driven ERST is linearly dependent on VAR for linear and some non-linear portfolios. However, this relationship is not present as a general rule for non-linear portfolios, making plausibility-driven ERST interesting and valuable. For non-linear portfolios, the approach can be seen as a continuum of VAR and expected shortfall (ES), and it sets a new paradigm for risk measurement. Some limitations do exist, however, as discussed at the end of this section.

■ **Existing VAR approach.** For a given $\alpha \in [0, 1]$, VAR_α returns the α quantile of the P&L density, indicating the P&L is not as extreme as the VAR output $\alpha\%$ of the time. The P&L density may be a historical or any other fitted density.

Taking the simple case of a linear portfolio with n risk factors and a weighting scheme $\boldsymbol{\omega}$, and assuming the risk factors are normally distributed $\mathbf{S} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$, then $\text{P\&L}(\mathbf{S}) \sim \mathcal{N}(0, \boldsymbol{\omega}^T \boldsymbol{\Sigma} \boldsymbol{\omega})$ and:

$$\text{VAR}_\alpha = -\mathcal{N}^{-1}(\alpha) \sqrt{\boldsymbol{\omega}^T \boldsymbol{\Sigma} \boldsymbol{\omega}} \quad (5)$$

for $\mathcal{N}^{-1}(\alpha)$, the α quantile of a standard normal distribution.

For a quadratic portfolio, this expression does not hold true. The distribution of $P\&L(\mathbf{S})$ may not be analytically known depending on the density function for \mathbf{S} . However, an approximate VAR can be calculated after Monte Carlo simulations of \mathbf{S} and a derivation of the probability based on the resulting P&L distribution.

With such an approach, VAR provides only one output: a loss. This does not allow a portfolio manager to dig deeper and understand where the underlying weaknesses in portfolio exposures lie.² In this respect, the plausibility-driven ERST provides a more complete result than VAR. In addition to a resulting loss, it provides a corresponding scenario, identifying the specific strengths and weaknesses of the portfolio. This allows the portfolio manager to take potential countermeasures such as hedging or portfolio adjustments. This advantage is discussed in more detail below.

■ **Problem statement.** Let α and MaxERST³ be the input level of plausibility and the output loss, respectively. The plausibility-driven ERST is then the optimisation problem:

$$\min_{\text{Maha}^2(\mathbf{S}) \leq q^\alpha} P\&L(\mathbf{S}) \quad (6)$$

In the two following sections, this problem is solved for both linear and quadratic portfolios.

■ **Application for delta-one strategies.** Here, (6) can be solved by relying on Lagrangian optimisation with Kuhn-Tucker conditions. For a given α , MaxERST and the corresponding scenario \mathbf{S}^α are:

$$\mathbf{S}^\alpha = -\sqrt{q^\alpha} \frac{\boldsymbol{\Sigma} \boldsymbol{\omega}}{\sqrt{\boldsymbol{\omega}^T \boldsymbol{\Sigma} \boldsymbol{\omega}}} \quad (7)$$

$$\text{MaxERST} = -\sqrt{q^\alpha} \sqrt{\boldsymbol{\omega}^T \boldsymbol{\Sigma} \boldsymbol{\omega}} \quad (8)$$

Figure 3 shows an application to a long/short strategy on two momentum indexes, with $\boldsymbol{\omega} = (1, -1)$. Comparing (5) and (8), MaxERST and VAR are proportional. For linear portfolios, Breuer (2006) states a similar relationship, adding that VAR and MaxERST are also proportional to the ES measure. The corresponding proof is by Sadefo-Kamdem (2004). Therefore, when \mathbf{S} is normally distributed:

$$\frac{\text{VAR}}{\mathcal{N}^{-1}(\alpha)} = \frac{\text{MaxERST}}{\sqrt{q^\alpha}} = \frac{\text{ES}}{\rho(\alpha)\alpha} = -\sqrt{\boldsymbol{\omega}^T \boldsymbol{\Sigma} \boldsymbol{\omega}} \quad (9)$$

where q^α is the α quantile of a $\chi^2(n)$ distribution and $\rho(\alpha)$ is the density of the standard normal distribution (Breuer 2006). Despite being proportional for delta-one strategies, Breuer (2006) argues MaxERST is more useful than VAR. As it is sub-additive⁴ whereas VAR is not, MaxERST has proved to be a more reliable limit system than VAR for simple non-linear portfolios such as some combinations of out-of-the-money short puts and short calls on the same underlying.

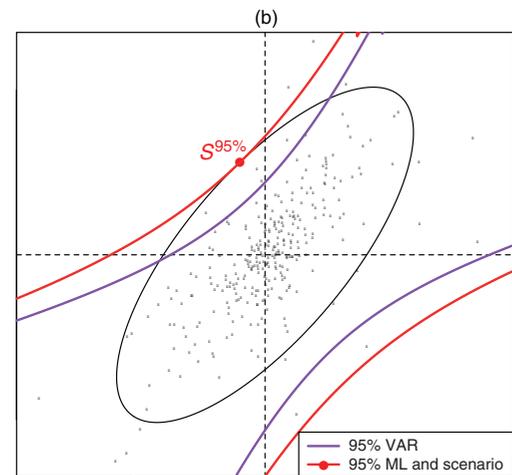
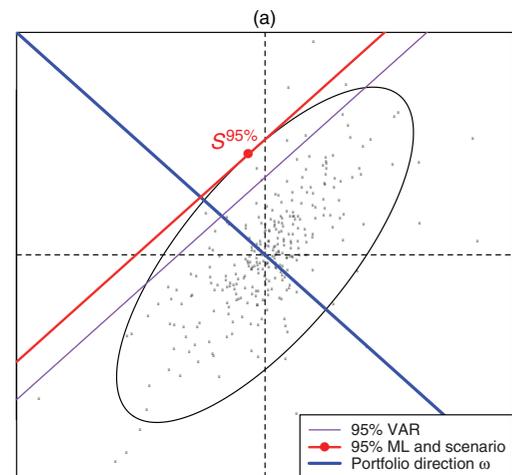
For q^α , the quantile of a $\chi^2(n)$ distribution $\lim_{n \rightarrow \infty} q^\alpha = \infty$. Therefore, if \mathbf{S} is normally distributed, the higher the number n of risk factors to which the portfolio is exposed, the more extreme MaxERST will be relative to VAR as per (9). This could create a dimensional dependency issue for irrelevant factors, as exposed in Mouy *et al* (2017) or Breuer *et al* (2009) and further discussed in the subsection titled ‘On dimensional dependency’ hereafter.

² This is possible with historical VAR but only for historical/past scenarios.

³ MaxERST was first introduced by Studer (1997) and denoted as maximum loss (ML).

⁴ By which we mean absolute losses are such that $\text{MaxERST}(\text{portfolio 1}) + \text{MaxERST}(\text{portfolio 2}) \geq \text{MaxERST}(\text{portfolios 1 + 2}) \geq 0$.

3 Plausibility-driven ERST for (a) a linear or (b) a non-linear portfolio



300 data points for two momentum indexes (in grey) are used to compute $\boldsymbol{\Sigma}$. The level of plausibility is fixed at $\alpha = 95\%$ and $\mathcal{E}^{95\%}$ is in black. Knowing the VAR and MaxERST as per (5) and (8), their corresponding iso-P&L lines are indicated

■ **Application for non-linear P&L.** For non-linear P&L, a second-order approximation is considered. Thus:

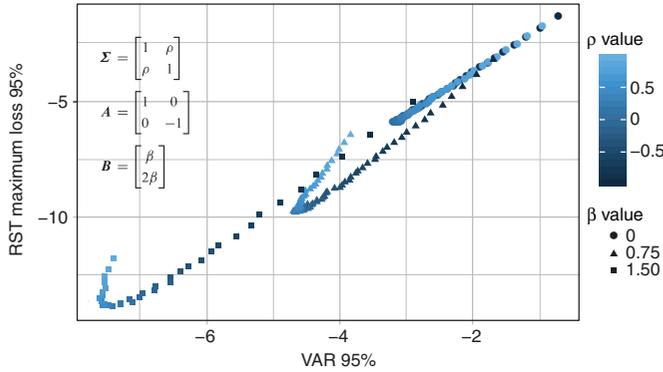
$$P\&L(\mathbf{S}) = \frac{1}{2} \mathbf{S}^T \mathbf{A} \mathbf{S} + \mathbf{B}^T \mathbf{S} \quad (10)$$

where \mathbf{A} and \mathbf{B} are the second- and first-order sensitivities of the portfolio, respectively. Second-order sensitivities being symmetric, \mathbf{A} is symmetric. With a quadratic form for P&L, the resolution of (6) is more complex. The objective function may not be convex, therefore Kuhn-Tucker conditions are irrelevant. Fortunately, there is an optimisation algorithm that can cope with this issue: the Levenberg-Marquardt algorithm. This is introduced in depth by Nocedal & Wright (1999) and applied by Studer (1997) to solve (6) with (10).

For the same two momentum indexes as in the previous subsection, results are shown in figure 3 for some \mathbf{A} and \mathbf{B} .

In addition, MaxERST is no longer linear with respect to VAR, as opposed to delta-one strategies. This result justifies the use of ERST rather than VAR for non-linear portfolios, as the approach offers added value to the portfolio manager and is a continuum of VAR. Figure 4 shows the aforementioned

4 Comparison between VAR and MaxERST (or maximum loss) outputs for $\alpha = 95\%$ and two risk factors



Two risk factors are simulated using unitary volatility for simplicity. For $\beta \neq 0$, the linear relation between VAR and MaxERST disappears. For $\beta = 1.5$, VAR does not vary much with the positive correlation, whereas MaxERST does. The latter is therefore a more interesting risk measure, as it reacts more strongly to changes in correlation

non-linearities. And yet the specific case where $B = 0$ remains linear, as proved in Traccucci *et al* (2019, appendix 1).

■ **On dimensional dependency.** As stressed in Mouy *et al* (2017) and Breuer *et al* (2009), the output of this ERST approach depends directly on the dimension of the problem, ie, the number n of risk factors under consideration. Indeed, q^α in (6) varies with n . As previously stated, q^α is, for example, the quantile of a $\chi^2(n)$ distribution for a normally distributed S .

Although this may be viewed as a source of instability, it is also positive from a portfolio management perspective: it is actually a way to account for correlation with the external yet meaningful risk factors that indirectly drive variations.

In addition, to bypass the instability caused by dimensional dependency, Rouvinez (1997) suggests replacing q^α with the Mahalanobis distance of a given scenario.

Starting from P&L

A P&L-driven ERST extends the ideas expressed by Mouy *et al* (2017) to non-linear portfolios. To this end, a new, adapted version of the Levenberg-Marquardt optimisation algorithm is defined and tested. The main advantage of such an approach as compared with starting from plausibility is to overcome the aforementioned dimensional dependency issue.

■ **Problem statement.** Given the dimensional dependency issue, it is preferable the constraint in (6) be independent of the squared Mahalanobis quantiles. Here, inverting the problem formulation works, ie, finding the scenario with optimal plausibility for a given P&L. This paves the way for the third and final approach discussed in this article. The optimisation problem becomes:

$$\min_{P\&L(S)=l} \text{Maha}^2(S) \quad (11)$$

The case for a linear P&L is discussed in Mouy *et al* (2017), but the resolution for non-linear P&L remains outstanding. The remainder of this section focuses on this.

■ **Resolution for non-linear P&L.** Rewriting (11) brings, for a loss l :

$$\min_{\frac{1}{2} S^T A S + B^T S \leq l} S^T \Sigma^{-1} S = \min_{\frac{1}{2} \hat{S}^T \hat{A} \hat{S} + \hat{B}^T \hat{S} \leq l} \|\hat{S}\|^2 \quad (12)$$

where the change of variable $\hat{S} = U^{-T} S$ is performed with U , the Cholesky decomposition matrix for Σ , and:

$$\hat{A} = U A U^T \quad (13a)$$

$$\hat{B} = U B \quad (13b)$$

Changing the variable allows the quadratic optimisation problem to work with a centred bowl rather than an ellipsoid. The problem is thus reduced to finding the closest scenario(s) \hat{S}^* to the origin and associated with the iso-loss curve of value l . This problem relates to the Levenberg-Marquardt optimisation problem used when starting from a plausibility. However, the constraint is not necessarily convex here. Therefore, a new version of the method is introduced.

It can be derived⁵ from the equivalence proved in Nocedal & Wright (1999, theorem 4.3) that \hat{S}^* is a solution to (12) if, and only if, it verifies the following conditions for λ_m , the smallest eigenvalue of \hat{A} , and a given μ :

$$(\hat{A} + \mu I) \hat{S}^* = -\hat{B} \quad (14a)$$

$$\mu \left(\frac{1}{2} \hat{S}^{*T} \hat{A} \hat{S}^* + \hat{B}^T \hat{S}^* - l \right) = 0 \quad (14b)$$

$$\mu \geq \max(0, -\lambda_m) \quad (14c)$$

The multidimensional optimisation problem (12) reduces to a scalar optimisation problem on μ under constraints (14a)–(14c). A problem of this type can be solved rapidly. As detailed below, a bisection algorithm to find the optimal μ allows for \hat{S}^* to be inferred directly.

If $\mathcal{B} = (\pi_1, \dots, \pi_n)$ is an orthonormal diagonalising basis of symmetric matrix \hat{A} , then:

$$\hat{S}^* = \sum_i \sigma_i \pi_i \quad (15)$$

$$\hat{B} = \sum_i \beta_i \pi_i \quad (16)$$

Defining $(\lambda_i)_i$, the eigenvalues of \hat{A} , $I_m = \{i, \lambda_i = \lambda_m\}$, and taking (14c) into account, (14a) expressed in \mathcal{B} becomes:

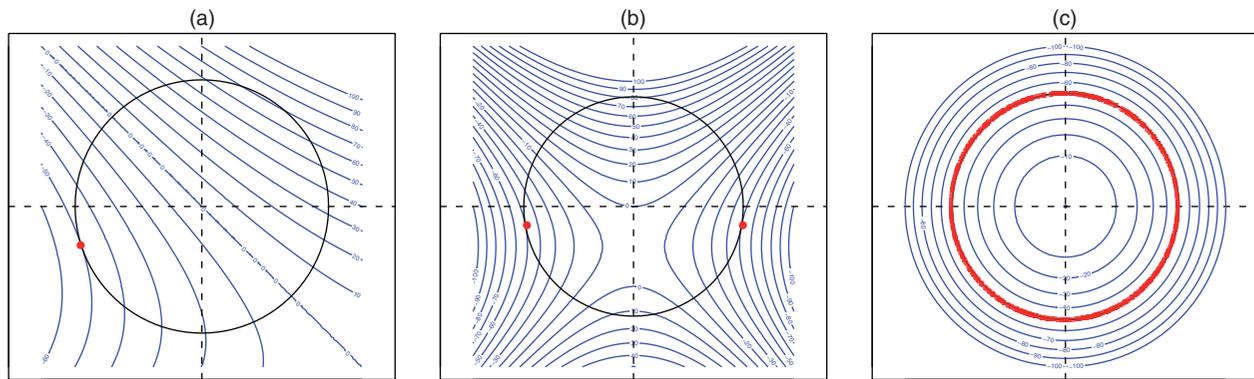
$$\sigma_i = -\frac{\beta_i}{\lambda_i + \mu} \quad \forall i \notin I_m \quad (17a)$$

$$\sigma_j = -\frac{\beta_j}{\lambda_m + \mu} \quad \forall j \in I_m \text{ if } \mu \neq -\lambda_m \quad (17b)$$

$$\sigma_j \in \mathbb{R} \quad \text{and} \quad \beta_j = 0 \quad \forall j \in I_m \text{ if } \mu = -\lambda_m \quad (17c)$$

It is thus possible to find a unique \hat{S}^* if (17b) is met, and several \hat{S}^* parameterised by $(\sigma_j)_{j \in I_m}$ if (17c) is met. This result is important, because it illustrates the different cases of existence and unicity of \hat{S}^* . In this respect, it is more complex than the functional expression obtained for the original

⁵ This article does not provide rigorous proof of this statement, resembling that in Nocedal & Wright (1999) for convex P&L. Instead, for both clarity and applicability, this article shows the statement solves all of the variations the optimisation problem (11) takes.

5 Solutions to (12) when $\hat{\mathbf{S}}$ consists of two risk factors

Depending on the P&L expression, (a) one solution, (b) two solutions or (c) an infinity of solutions are found

Levenberg-Marquardt problem analysed by Nocedal & Wright (1999) and Studer (1997).

Expressing (14b) in \mathcal{B} brings $\mu(f(\mu) - l) = 0$ with:

$$f(\mu) = \sum_i \left[\frac{\lambda_i}{2} \left(\frac{\beta_i}{\lambda_i + \mu} \right)^2 - \frac{\beta_i^2}{\lambda_i + \mu} \right] \quad \text{if } \mu \neq -\lambda_m \quad (18a)$$

$$= \sum_{i \notin I_m} \left[\frac{\lambda_i}{2} \left(\frac{\beta_i}{\lambda_i - \lambda_m} \right)^2 - \frac{\beta_i^2}{\lambda_i - \lambda_m} \right] + \frac{\lambda_m}{2} \sum_{j \in I_m} \sigma_j^2 \quad \text{if } \mu = -\lambda_m \quad (18b)$$

These different dynamics lead to the following discussion on $\hat{\mathbf{A}}$:

(1) For $\hat{\mathbf{A}}$ positive definite, $-\lambda_m < 0$ and $\mu \geq 0$ as per (14c).

(1a) If $\mu = 0$, $\hat{\mathbf{S}}^* = -\hat{\mathbf{A}}^{-1} \hat{\mathbf{B}}$ as per (14a) and:

$$\begin{aligned} \text{P\&L}(\hat{\mathbf{S}}^*) &= -\frac{1}{2} \hat{\mathbf{B}}^T \hat{\mathbf{A}}^{-1} \hat{\mathbf{B}} \\ &= -\frac{1}{2} \mathbf{B}^T \mathbf{A}^{-1} \mathbf{B}, \end{aligned}$$

which corresponds to the global minimum P&L. Such a value of μ is chosen whenever the scenario for the global minimum is more plausible than the most plausible scenario for loss l .

(1b) If $\mu \neq 0$, then the loss l is attained as per (14b). However, such a loss must be greater than the global minimum P&L. If f is continuous and increasing, a single μ corresponds to any loss and can be approximated using a bisection algorithm.

(2) For $\hat{\mathbf{A}}$ semi-positive definite, $-\lambda_m = 0$ and $\mu \geq 0$.

(2a) If $\mu = 0$, (17c) applies. As per (18b), the P&L does not vary with any σ_j , $j \in I_m$, and the corresponding risk factors become irrelevant. The dimensions of the problem are thereby reduced and it becomes similar to (1a).

(2b) If $\mu \neq 0$, then loss l is attained as per (14b). Since $\lim_{-\lambda_m^+} f = -\infty$ and $\lim_{+\infty} f = 0$, and f is still continuous and increasing, a single μ corresponds to any loss and can again be approximated using a bisection algorithm.

(3) For any other $\hat{\mathbf{A}}$, $-\lambda_m > 0$ and $\mu \geq -\lambda_m$.

(3a) If $\mu = -\lambda_m$, (17c) applies. As per (18b), the P&L still varies with σ_j , $j \in I_m$. Constraint (14b) becomes $f(\mu) = l$ and a root-finding algorithm

(such as Newton-Raphson) can determine which values σ_j , $j \in I_m$, must take. This solution may or may not be unique.

(3b) If $\mu \neq 0$, (2b) applies.

This indicates that it is only possible to solve (12) for losses ($l \leq 0$). This is actually a direct consequence of the formulation of the problem itself. Indeed, the null scenario always returns a zero-valued P&L per (10). In addition, the null scenario returns the lower boundary of the objective function in (12). Therefore, a profit input ($p > 0$) cannot be obtained as a null scenario both returns a lower value in the objective function and respects the P&L constraint. However, generating profit scenarios is of significant interest in assessing the asymmetries in portfolio P&L. Thus, for $p > 0$, (12) may be rewritten as follows:

$$\min_{-\frac{1}{2} \hat{\mathbf{S}}^T \hat{\mathbf{A}} \hat{\mathbf{S}} + \hat{\mathbf{B}}^T \hat{\mathbf{S}} \leq -p} \|\hat{\mathbf{S}}\|^2 \quad (19)$$

■ **Application to non-linear P&L.** The adapted Levenberg-Marquardt algorithm is tested on portfolios with two risk factors in figure 5. This algorithm is not sensitive to the number of risk factors, as the numerical procedure for solving (12) is reduced to analysing the function f of one variable in (18). Therefore, the numerical techniques involved are time-efficient.

Conclusion

The ERST adds value compared with traditional stress tests and risk measures such as VAR or ES, mostly because its output contains more information. This additional information can help both portfolio and risk management teams to control a portfolio's sensitivities and reallocate resources as and when needed.

Possible next steps may include:

■ A procedure for recomputing Greeks in (10) to better account for their potential instability in scenarios with high market moves.

■ A bootstrapping procedure for the covariance matrix Σ in (1). This would mitigate the error in the estimated plausibility of a scenario due to the estimation of Σ .

■ A procedure for better interpreting any ERST output scenario. An interesting starting point may be the maximum loss contribution defined by Breuer *et al* (2009).

■ A procedure to go beyond the restriction of multivariate elliptical distributions. The use of copulas as done by Mouy *et al* (2017) would serve as a starting point.

■ A procedure to stress Σ . Because of (1), the covariance matrix affects the stressing of the portfolio, which is of primary importance in risk management. In this respect, two methods to stress Σ are provided and illustrated in Traccucci *et al* (2019). They account for the risk of recorelation between supposedly independent strategies in market downturns. The

influence of the covariance matrix in generating the most plausible scenarios is also discussed. ■

Pascal Traccucci is global head of risk at La Française Group, Luc Dumontier is partner and head of factor investing at La Française Investment Solutions (LFIS), Guillaume Garchery is partner and head of quantitative research at LFIS, and Benjamin Jacot is quantitative risk manager at La Française Group in Paris. This article reflects the authors' opinions and not necessarily those of their employers. Email: p

REFERENCES

Breuer T, 2006

Providing against the worst: risk capital for worst case scenarios
Managerial Finance 32, pages 716–730

Breuer T, M Jandačka, K Rheinberger and M Summer, 2009

How to find plausible, severe, and useful stress scenarios
International Journal of Central Banking 5(3), pages 205–224

Mouy P, Q Archer and M Selmi, 2017

Extremely (un)likely: a plausibility approach to stress testing
Risk August, pages 72–77

Nocedal J and SJ Wright, 1999

Numerical Optimization, Springer Series in Operations Research
Springer

Rouvinez C, 1997

Going Greek with VaR
Risk February, pages 57–65

Sadefo-Kamdem J, 2004

Value at risk and expected shortfall for linear portfolios with mixtures of elliptically distributed risk factors
Working Paper, Université de Reims

Studer G, 1997

Maximum loss of measurement of market risk
PhD Thesis, Swiss Federal Institute of Technology

Traccucci P, L Dumontier, G Garchery and B Jacot, 2019

A triptych approach for reverse stress testing of complex portfolios
Preprint, arXiv:1906.11186 [q-fin.RM]

