Full Probabilistic Control for Direct & Robust, Generalized & Targeted Stressing of the Correlation Matrix (Even When Eigenvalues are Empirically Challenging)

J.D. Opdyke Head of Enterprise Risk Analytics, Allstate

QuantMinds International, November 2-6, 2020 Hamburg, Germany



© J.D. Opdyke 1 of 45

J.D. Opdyke is Head of Enterprise Risk Analytics, VP-Financial Risk and Return Management, at Allstate.

The views presented herein are solely those of the author and do not necessarily reflect the views of Allstate or any of its subsidiaries or legal entities.

J.D. Opdyke, the sole author of the work contained herein, presents this paper in the public domain with the permission of, and gratitude to, Mark Prindiville, PhD, CRO – Allstate.

The author would also like to thank Ming Dai, PhD for reviewing an earlier draft of this paper, as well as Dan Kirsner, PhD for replicating many of the empirical results and for catching errors in two related peer-reviewed papers. Any errors herein are my own.



Guiding Principle / Driver:

"Measurement is the first step that leads to control and eventually to improvement. If you can't measure something, you can't understand it. If you can't understand it, you can't control it. If you can't control it, you can't improve it." (emphasis added)

- H.J. Harrington

This presentation is all about increasing the accuracy, precision, robustness, and comprehensiveness of measuring portfolio risk for seamless, transparent portfolio risk management. Absent useable and scientifically defensible measurement, there simply is no meaningful risk management.





Paper Abstract

In practice, risk measurement of the majority of enterprise-level portfolios, and even many investment portfolios, requires stressing the correlation matrix directly, rather than (solely) stressing its underlying variables, due to 1. data paucity or incomplete time series, 2. matrices where at least some of the values are based on subject-matter expertise, and/or 3. the need to specify and test the effects of correlation values nowhere close to those reflected in historical data (i.e. the majority of extreme scenarios used in forward-looking stress testing). Surprisingly few papers in the literature address this common, real-world situation, and their approaches arguably are either adhoc, lack solid statistical underpinnings, do not allow for direct, probability-based stressing, and/or remain nonrobust as they fail under empirically challenging conditions (e.g. near-zero eigenvalues resulting from the need to first enforce positive definiteness of the matrix). We borrow from recent advances in the literature for generating random correlation matrices (based on the identity matrix) to design a method that both mitigates and eliminates these drawbacks when directly stressing real-world correlation matrices (other than the identity matrix). Our approach can be used for both generalized and targeted stressing. The former perturbs the entire correlation matrix, which can be used to account for difficult-to-model or difficult-to-anticipate second and third order effects of extreme scenarios, as well as providing much needed percentiles of the distribution of the entire matrix. Targeted stressing, on the other hand, allows for particular correlation values to be changed by specified amounts based directly on the probability of observing such changes due to the event/scenario. And both generalized and targeted stressing can be performed concurrently, based on the same proposed approach, which provides full probabilistic control while automatically enforcing positive definiteness. We demonstrate the method on realistic, reasonably large matrices (100x100) that have had positive definiteness enforced via Higham (2002), reflecting a common occurrence for most enterprise-level portfolios and even many investment portfolios. Although it requires numeric integration (or alternatively, fast and accurate calculations of Gauss hypergeometric functions) for all but very small matrices, our approach's runtimes are comparable to those of competing methods. Implementation is straightforward, and results robustly outperform existing methods in the literature, especially when matrices are empirically challenging (e.g. near-zero eigenvalues).

QuantMinds International

Contents

- I. Introduction and Background
 - 1. Measuring Portfolio Risk Requires Scenario/Stress Testing
 - 2. Scenario/Stress Testing Requires Directly Stressing the Correlation Matrix
- II. Directly Stressing/Perturbing An Estimated (Pearson's) Correlation Matrix
 - 1. Recent Advances Stressing/Perturbing the Identity Matrix
 - 2. Other Approaches to Directly Perturbing An Estimated Correlation Matrix
 - 3. C3: A New Distribution for Directly Perturbing An Estimated Correlation Matrix
- **III.** Using C3 for Targeted Scenario Analytics
- **IV. Using C3 for Generalized Stress Testing**
- V. Using C3 for Concurrent Targeted and Generalized Stress Testing
- **VI.** Summary and Conclusions
- **VII. Next Steps / Further Research**
- VIII. References

- Obtaining a broad, accurate, and complete understanding and measurement of a portfolio's risks* requires subjecting it to wide-ranging and sometimes extreme scenarios via 'stress testing' because:
 - historical data is limited: if not in its detail then in the breadth and scope of the scenarios it represents, and it cannot provide a comprehensive view of a portfolio's current or future risks beyond the narrowly defined ranges of (relatively recent) past experience.
 - when things go bad, they go bad together: this is especially true during times of extreme stress (e.g. pandemic + economic upheaval) when correlations have been shown to move together much more dramatically than during more 'normal' times (see Pritsker, M., 2006; BIS-BCBS, 2011**; Golub & Guo, 2012; Junior & Franca (2012); Ng et al., 2013; and So et al., 2013).
 - no data: sometimes the construction of a portfolio, by design, means that it simply has no historical data (e.g. new products without close comparables), making its stress testing that much more critical to accurate risk measurement.
 - portfolios are unique: even if textbook, perfectly complete time series data exists for all of a
 portfolio's assets/risks, the forward-looking purposes and risk appetites for a particular portfolio
 typically are very different from those of others (otherwise they would be the same portfolio); so
 comprehensively measuring its risks requires varying very specific combinations of individual
 characteristics to assess the effects of possible future scenarios relevant to that particular
 portfolio.
- Comprehensive stress testing requires not only stressing the data inputs on which the framework's parameter estimates are based, but also directly stressing the parameter estimates themselves, and this includes the estimated correlation matrix.

*<u>NOTE</u>: The methods developed in this paper are broadly applicable to the risk measurement of any portfolio, but are especially useful for cross-risk-siloed enterprise risk portfolios, and even many investment portfolios.

Ought.Minds

I.2 Stress Testing Requires Directly Stressing The Corr Matrix

- One of the most important components of a model framework estimating portfolio risk is the **dependence structure** between the marginal assets/risks of the portfolio.
- For all its limitations, one of the most widely used measures of dependence is Pearson's (linear) **correlation matrix**. As the scaled version of the variance-covariance matrix, it is foundational.
- Comprehensive stress testing requires not only stressing the data inputs on which the correlation estimates are based, but also directly stressing the correlation estimates themselves.
- Why stress the estimated correlation matrix directly?
 - Data paucity and/or incomplete time series prevents empirical estimation of many of the pairwise correlations (especially for reasonably large portfolios, e.g. a portfolio of 100 assets/risks contains 4,950 pairwise correlation values, and it is often unlikely that sufficiently complete time series exist or are available in comparable formats to estimate all of these).
 - Some correlations are, **by necessity, subject-matter expertise estimates** rather than empirical estimates (e.g. new products without close comparables); this is more common with enterprise portfolios, but can and often does occur in investment portfolios.
 - Combinations of correlation values relevant for many (sometimes extreme) scenarios are often nowhere near to anything reflected in historical data, requiring their explicit specification absent empirical support.

**<u>NOTE</u>: BIS-BCBS, 2011 – "...in order to calculate stressed VaR accurately it is also necessary to stress the correlation matrix used in all VaR methodologies. It is a repeated observation that during times of extreme volatility, such as occurs during every market crash, correlations are dramatically perturbed relative to their 'normal' historical values. In general, most correlations tend to increase during market crises, asymptotically approaching 1.0 during periods of complete meltdown, such as occurred in 1987, 1998 and 2008."



- Recent papers have advanced our understanding of, and ability to efficiently generate, the distribution of valid (positive definite) correlation matrices <u>based on the identity matrix</u>.
- This widely has been described as generating 'random correlation matrices,' but it is important to note that this distribution is centered only on the identify matrix (the matrix with off-diagonal values of zero).
- Before describing this distribution, it is important to first note that of all random symmetric matrices with values of 1.0 along the diagonal and off-diagonal values randomly uniformly drawn from between –1.0 to 1.0 (so all matrices that LOOK like correlation matrices), the ratio of those that are positive definite very rapidly converges to zero as the dimension of the matrix, p, increases (see Bohn & Hornik, 2014, and Pourahmad & Wang, 2015):

$$\Pr\left(rand "C" \sim PosDef\right) = X = \frac{\prod_{j=1}^{p-1} \left[\sqrt{\pi} \Gamma\left(\frac{j+1}{2}\right)\right]^{j}}{2^{p(p-1)/2}} < \prod_{j=1}^{p-1} \left[\frac{\sqrt{\pi}}{2}\right]^{j} = \left[\frac{\sqrt{\pi}}{2}\right]^{p(p-1)/2}; \lim_{p \to \infty} \left[X\right] = 0$$

 This means that empirically, even a relatively small matrix of p=25 has a probability of less than 0.000000000000000183 (or less than 2 in ten quadrillion) of being positive definite. Randomly generating a sufficient number of valid matrices by deleting all the invalid ones would take eons: correlation matrix sampling algorithms must efficiently generate only valid, positive definite correlation matrices.

QuantMinds International

- **To efficiently sample only positive definite matrices**, note that the Cholesky factorization of any symmetric positive definite matrix (which includes all valid correlation matrices) has rows whose squares sum to 1.0. This is another way of stating that they exist on the unit hypersphere (or technically, as used below, the unit hemisphere of dimension p). So while we cannot simply uniformly perturb individual correlation values directly and remain within the realm of positive definite correlation matrices, we CAN sample only positive definite matrices from the p-dimensional hemisphere **using polar coordinates** (actually, just the angles, as the vectors all equal 1.0). We simply need the right distribution to execute this sampling (this idea of sampling polar angles originates, at least in financial applications, with Rebonato and Jaeckel, 2000).
- Pourahmad & Wang (2015) show that the uniform distribution of positive definite matrices on the pdimensional hemisphere is proportional to the determinant of the Jacobian, which is defined in terms of the Cholesky factorization as shown below (see also Cordoba et al., 2018)

det $\left[J(U)\right] = 2^{p} \prod_{i=1}^{p-1} u_{ii}^{i}$ where U is the Cholesky factorization of correlation matrix $R = UU^{t}$

- There are a number of algorithms in the literature that perform sampling of this distribution (albeit not all using spherical angles directly):
 - The c-vines and onion methods (Lewandowski et al., 2009)
 - The restricted Wishart distribution approach of Wang et al. (2018).
 - The direct formulation method of Madar (2015)
 - The Cholesky-Metropolis method of Cordoba et al. (2018)
 - The angles distribution of Makalic and Schmidt (2018) combined with the polar representation of Rapisarda et al. (2007)

International

- We proceed with the Makalic and Schmidt (2018) distribution and sampling algorithm as it exhibits runtimes as fast as its competitors (see Cordoba et al., 2019), is straightforward to implement, has been implemented on a widely circulated R package (see "randcorr" in CRAN, 2018), and we use it to derive below a broader distribution of which it is a special case. We also describe later how sampling angles appears to be more robust, especially under empirically challenging conditions, than other approaches (e.g. sampling eigenvalues).
- Note also that **sampling the angles** corresponding to the trigonometric representation of the Cholesky factorization (shown below) in this way **allows each to be sampled independently**.
- The c(k)*sin(x)^k distribution of Makalic and Schmidt (2018) (PDF below) must be sampled based on rejection sampling using a scaled Beta distribution as the envelope because its CDF (also shown below but not derived in Makalic and Schmidt, 2018), is difficult to invert analytically to perform simpler sampling methods, like inverse transform sampling.

$$f_{X}(x) = c_{k} \cdot \sin^{k}(x), x \in (0,\pi), k = 1,2,3... \text{ where normalization constant } c_{k} = \frac{\Gamma(k/2+1)}{\sqrt{\pi}\Gamma(k/2+1/2)}$$

$$F_{X}(x;k) \sim \frac{1}{2} - c_{k} \cdot \cos(x) \cdot {}_{2}F_{1}\left[\frac{1}{2}, 1 + \frac{1-k}{2}; \frac{3}{2}; \cos^{2}(x)\right],$$

where $_{2}F_{1}[a,b;c;z] = \sum_{n}^{\infty} \frac{(a)_{n}(b)_{n}}{(c)_{n}} \cdot \frac{z^{n}}{n!}$ where $(h)_{n} = h(h+1)(h+2)\cdots(h+n-1), n \ge 1, (h)_{0} = 1$

International

• The Makalic and Schmidt (2018) sampling algorithm is shown in the following pages.

II.1 Recent Advances Stressing/Perturbing the Identity Matrix



• Importantly, note for later how the spread is a function of k, the column location in the matrix:

C

$$Stdev(X) = \sqrt{\int_{0}^{\pi} x^{2}c_{k} \sin^{k}(x)dx} - \left[\int_{0}^{\pi} xc_{k} \sin^{k}(x)dx\right]^{2}} = \sqrt{p_{=4}F_{q=3}\left(\left\{1,1,1,\frac{3+k}{2}\right\}, \left\{\frac{3}{2},2,2+k\right\};1\right) - \frac{\pi^{2}}{4}, k \ge 1 \text{ where}}{p_{q}F_{q}\left(\left\{a1,\ldots,ap\right\}, \left\{b1,\ldots,bq\right\};z\right)} = \sum_{n=0}^{\infty} \frac{(a1)_{n}\cdots(ap)_{n}}{(b1)_{n}\cdots(bq)_{n}} \frac{z^{n}}{n!}, \text{ where } (c)_{n} = c(c+1)(c+2)\cdots(c+n-1), n \ge 1, (c)_{0} = 1$$
QuantMinds
Denote a 11 of 45

II.1 Recent Advances Stressing/Perturbing the Identity Matrix

Note that "k" corresponds to the number of columns minus the column number: k = p – j, although this
is not explained in Makalic and Schmidt (2018)

let
$$f_x(x) = c_k \sin^k(x)$$
, $x \in (0,\pi)$, $k \ge 1$ where normalization constant $c_k = \frac{\Gamma(k/2+1)}{\sqrt{\pi}\Gamma(k/2+1/2)}$

note that c_k is just the inverse of integrating over the entire domain of the hemisphere, $x \in (0, \pi)$ Then sample from $f_x(x)$ as follows:

Then sample from $f_X(x)$ as follows: 1. Generate X ~ Beta $(k+1,k+1) \sim \frac{x^k (\pi - x)^k}{B(k+1,k+1)\pi^{(2k+1)}}, x \in (0,\pi), k \ge 1$

where
$$B(q,r) = \frac{\Gamma(q)\Gamma(r)}{\Gamma(q+r)}$$
 = Euler's beta function

2. Generate U ~ Uniform(0,1)

3. Accept X if
$$\frac{\ln(U)}{k} \le \ln\left(\frac{\pi^2 \sin(X)}{4X(\pi - X)}\right)$$

4. Otherwise go to 1.

The algorithm's maximum expected iterations per sample is only $\pi/(2\sqrt{2})=1.11$, consistent with our empirical implementation of it herein.

QuantMinds International

 Once the polar angles have been randomly sampled per the above, they are converted via trigonometric transformations into their corresponding Cholesky factorization (see Rapisarda et al., 2007, and Pourahmadi & Wang, 2015; see van Oest, 2019, for an alternate parameterization).

$$B = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0\\ \cos(\theta_{2,1}) & \sin(\theta_{2,1}) & 0 & \cdots & 0\\ \cos(\theta_{3,1}) & \cos(\theta_{3,2})\sin(\theta_{3,1}) & \sin(\theta_{3,2})\sin(\theta_{3,1}) & \cdots & 0\\ \cos(\theta_{4,1}) & \cos(\theta_{4,2})\sin(\theta_{4,1}) & \cos(\theta_{4,3})\sin(\theta_{4,2})\sin(\theta_{4,1}) & \cdots & 0\\ \vdots & \vdots & \vdots & \ddots & \vdots\\ \cos(\theta_{n,1}) & \cos(\theta_{n,2})\sin(\theta_{n,1}) & \cos(\theta_{n,3})\sin(\theta_{n,2})\sin(\theta_{n,1}) & \cdots & \prod_{k=1}^{n-1}\sin(\theta_{n,k}) \end{bmatrix}, \text{ where } \theta_{i,j}, i > j \text{ are the previously sampled angles.}$$

- (Note the relationship between the diagonals above, the determinant of the Jacobian defined previously, and the sin(x)^k distribution defined in Makalic and Schmidt, 2018)
- Finally, we simply multiply the Cholesky factorization by its transpose to obtain a randomly sampled correlation matrix from the space of all positive definite correlation matrices of dimension p (which are uniformly distributed on the p-dimensional hemisphere).

$$R_{sample} = BB^{t}$$

QuantMinds International

• So the relationship between the correlation values and their polar angles is one-to-one and non-directional (see (2) below from Pourahmadi and Wang, 2015):

1. sample the angles (independently) based on Makalic and Schmidt (2018)

2. apply trigonometric functions to the polar angles to obtain the Cholesky factorization

3. multiply the Cholesky factorization by its transpose to obtain the correlation matrix in reverse...

3. estimate the correlation matrix

2. obtain the Cholesky factorization of the correlation matrix

1. use inverse trigonometric functions on 2. (shown later) to obtain corresponding polar angles

$$r_{i,j} = \cos\left(\theta_{i,1}\right)\cos\left(\theta_{j,1}\right) + \prod_{k=2}^{i-1}\cos\left(\theta_{i,k}\right)\cos\left(\theta_{j,k}\right)\prod_{l=1}^{k-1}\sin\left(\theta_{i,l}\right)\sin\left(\theta_{j,l}\right) + \cos\left(\theta_{j,i}\right)\prod_{l=1}^{i-1}\sin\left(\theta_{i,l}\right)\sin\left(\theta_{j,l}\right) \text{ for } 1 \le i < j \le n$$

- Sampling from the p-dimensional hemisphere ensures, via these transformations, that the correlation matrix automatically is positive definite. As discussed later, this transformation appears to be more robust under challenging empirical conditions than other methods (e.g. sampling eigenvalues).
- This solves the problem of (efficiently) randomly sampling valid (positive definite) correlation matrices, but the sampling is only from a distribution of correlation matrices based upon the identify matrix, i.e. the matrix of correlations whose (off-diagonal) values are all zero.
- When we estimate values of a correlation matrix under real world conditions, their true, underling values are almost certainly not the identity matrix.
- How do we perform the same sampling/perturbation exercise above based on a real, estimated correlation matrix? And establish a distribution of matrices based on this observed/estimated, real-world correlation matrix?

QuantMinds International

- Very few papers suggest approaches for defining and generating a distribution for the correlation matrix based on an observed/estimated correlation matrix (as opposed to the identity matrix).
- Similar to portions of the identity-matrix literature discussed above, <u>Rapisarda et al. (2007)</u> suggest sampling polar angles directly, but they provide no guidance on what distribution to use other than to hint at the form of the transformation such a distribution might take:
 - "...it is sufficient to re-parameterize them [the angles] in terms of real numbers in $(-\infty,\infty)$ as follows:"

$$\theta_{i,j} = \frac{\pi}{2} + \arctan\left(x_{k,j}\right) \text{ for } j = 1, \dots, k-1$$

... but we have no guidance on how the x's are distributed.

• Similarly, <u>Galeeva et al. (2007)</u> propose several approaches, one of which is perturbing the polar angles directly via:

$$\hat{\theta}_{i,j} = \arctan\left(\tan\left(\theta_{i,j} + \frac{\pi}{2}\right)\left(1 + \sigma z_{i,j}\right)\right) + \frac{\pi}{2} \text{ where } z_{i,j} \sim N(0,1)$$

"with the goal of generating random angles around the base angles with some distribution which is symmetric and centered around the base-correlation $\theta_{i,j}^0$ [angle] for every i.j"

• This approach has several limitations.

QuantMinds International

- This approach of Galeeva et al. (2007) has several limitations.
 - Note first that the resulting distribution, regardless of the choice of σ , is in fact NOT symmetric about the base angle(s) as Galeeva et al. (2007) seek, but rather, is skewed. Yet based on the finite support of (0, π), we believe skewness is much more reasonable here than symmetry.
 - First, the choice of a Gaussian distribution appears arbitrary here. And this requires the empirical estimation of σ which (unnecessarily) adds estimation error. It also begs the question of what to do when historical data is limited or different across assets/risks for such an estimation.
 - We find the distribution empirically to be very sensitive to the values of σ , so if this parameter cannot be estimated robustly, or even estimated at all due to data paucity, it will render the angle distribution non-robust, mispecified, or non-existent.
 - Finally and most importantly, assuming normality also does not respect the requirement proven in the "identity-matrix" literature that the distributions of the angles depend directly and materially on their position in the matrix, that is, based on the "k" parameter in the sin(x)^k distribution, which is the number of columns minus the column number: k = p j. This is consistent with the proven requirement of proportionality with the determinant of the Jacobian of the Cholesky factorization (see Pourahmad & Wang, 2015, and Cordoba et al., 2018, among others). This makes intuitive sense as the range of possible valid values of each correlation obviously depend directly on those elsewhere in the matrix. This dependence between correlation values, expressed explicitly in the trigonometric representation of the Cholesky factorization (solving this problem: ignoring this arguably is a fatal flaw.

uantMinds

International

QuantMinds International

٠

II.2 Approaches to Directly Perturbing An Estimated Corr Matrix

- Other approaches proposed by Galeeva et al. (2007) include bootstrapping, perturbing correlation values, and perturbing eigenvalues.
- The bootstrap approach is that of <u>Fengler and Schwendner (2003)</u> and is strictly limited to the use of historical data which is in most cases way too limited in the scope and breadth of the scenarios and stresses it represents to provide comprehensive measurement of portfolio risks.
- The direct perturbation of correlation values is that of <u>Turkay et al. (2003)</u> which perturbs blocks of correlations via reordering the matrix, which leaves the rest of the matrix unchanged while preserving positive definiteness. This approach does not provide direct, probabilistic control over a single correlation, let alone arbitrary combinations of correlations, as does our derived approach, and is also arguably more unwieldy to implement, especially for large(r) portfolios.
- Finally, their proposed eigenvalue perturbations rely on exponentiation of perturbed, random normal variables and estimated variances based on historical data. Here the assumption of normality arguably is arbitrary, the estimated variances add additional estimation error. Also, our preliminary empirical testing shows that perturbing eigenvalues fails under challenging empirical conditions, e.g. when the positive definiteness of the matrix has to be enforced algorithmically (e.g. via Higham, 2002, or a similar approach) and eigenvalues are virtually zero (or at least unreliably estimated).



II.2 Approaches to Directly Perturbing An Estimated Corr Matrix

- <u>So et al. (2013)</u> use a block approach in ways similar to Turkay et al. (2003), but they provide material improvements computationally. However, positive definiteness must explicitly be enforced in the algorithm, and it is far more involved (it always involves reordering the entire matrix) compared to our approach. Finally, it does not allow for direct, probabilistic control over individual correlations, as does our approach, making it arguably less transparent, and making probabilistically-defined scenarios less directly translatable and useable for the method.
- While <u>Loland et al. (2013)</u> primarily focus on enforcing positive definiteness, their Bayesian approach appears to also provide a distribution of the observed/estimated (now positive definite) correlation matrix. Even if used for this purpose, however, it would not appear to provide direct, probabilistic control over individual correlations, as does our approach.
- <u>Ghosh et al. (2020)</u> focus on developing shrinkage estimates of correlation values using a Bayesian approach. They perturb the spherical angles of the matrix directly (as does the method we present below) to great effect computationally, but because they shrink these values towards $\pi/2$ (i.e. the identity matrix) this approach does not provide a useable prior for the angles for stress testing.
- <u>Ho (2016)</u> uses an empirical likelihood approach for targeted stressing explicitly, obtaining the nearest correlation matrix (vis-à-vis Kullback-Leibler divergence) to the stressed matrix generated by weighting the input data. The convergence algorithm that maximizes the objective function is slow (104 seconds), even for a small matrix (10x10), and even though it automatically enforces positive definiteness, the method does not provide direct probabilistic control over individual correlations (as our method presented below does) just closest matches to specified values, which arguably is a major drawback for targeted stress.

QuantMinds International

 <u>Hardin et al. (2013)</u> utilize a normalized vector of independent gaussian random variables to perturb the observed correlation matrix. While relatively straightforward to implement, a nontrivial drawback compared to our method is its reliance on minimum eigenvalues that can be reasonably estimated.
 <u>"The amount of noise that can be added to the original matrix is determined by its smallest</u> <u>eigenvalue</u>."

"We provide the user with ... a general algorithm to apply to <u>any correlation matrix for which the</u> <u>smallest eigenvalue can be reasonably estimated</u>." (emphases added)

- As discussed above, correlation matrices estimated on large portfolios often (perhaps usually) are not positive definite for a wide range of reasons, and once positive definiteness is enforced using reliable, proven methods (e.g. Higham, 2002), the smallest eigenvalue of the resulting matrix is almost always virtually zero. This makes its estimation, let alone its perturbation, difficult, if not wholly unreliable.
- Our approach shown below has been tested extensively on reasonably large (100x100 and larger) portfolios under these conditions, without any noticeable adverse effects. In other words, preliminary indications are that perturbing polar angles rather than eigenvalues may be a much more robust approach to the problem, at least for many (most) real-world, sizeable matrices.



It is important to reemphasize here that **directly stressing the estimated/observed correlation matrix is not inconsistent in any way with concurrently stressing the inputs to the matrix**. Many approaches perturb the underlying distributions to the correlation matrix (e.g. Kupiec, 1998; Ng et al., 2013; Zhang et al., 2015; Packham & Woebbeking, 2019), but our goal is to perturb the matrix directly due to all the real-world data constraints this addresses (which stressing the inputs alone does not), as discussed at the beginning of this paper. Indeed, **stressing the correlation matrix directly can (and should be) done in conjunction with stressing its inputs**. See BIS Guidance below, and Opdyke (2019) for related methods of extreme quantile estimation for these stressed inputs:

<u>BIS-BCBS, 2011</u>, pp. 28-29: "This would have the effect of lengthening the tails of the Gaussian (normal) loss distributions that underlie all standard VaR calculations. ... A more sophisticated approach might include not only linear transforms of multivariate normal risk factors but also employing 'fat-tailed' distributions to model the extreme loss events more accurately. Examples of those 'extreme value theory' distributions are the Gumbel, Generalised Pareto, Weibull, Fréchet, and the Tukey g&h distributions. ... However, in order to calculate stressed VaR accurately it is also necessary to stress the correlation matrix used in all VaR methodologies. It is a repeated observation that during times of extreme volatility, such as occurs during every market crash, correlations are dramatically perturbed relative to their 'normal' historical values. In general, most correlations tend to increase during market crises, asymptotically approaching 1.0 during periods of complete meltdown, such as occurred in 1987, 1998 and 2008. ...Certain methods that could be meaningful in this context can be identified in the earlier literature on stress testing. Employing fat-tailed distributions for the risk factors and replacing the standard correlation matrix with a stressed one are two examples." (emphasis added)

QuantMinds International

٠

- Goals of our approach include obtaining a distribution of spherical angles that:
 - Is (median-) centered on the observed angle values
 - Is unimodal, continuously differentiable, and <u>sampled angles are independent</u> (of each other)
 - Is increasingly skewed as the observed angle values approach the bounds of support on $(0, \pi)$
 - Makes no distributional assumptions about the underlying input data (e.g. not necessarily normal, which yields Wishart-distributed correlation matrices)
 - Defines its spread as a function of its position in the correlation matrix, specifically as a function of its column defined by k = p – j (total #columns minus column number).
 - Respects the successive, hierarchical constraint placed on the angle distributions by their column number (recall that this is due to the nature of the determinant of the Jacobian, as shown along the diagonal of B (the Cholesky reparameterization) which are products of increasing numbers of sin functions of the angles).
- Additional goals include:
 - Straightforward implementation.
 - Use of a single, univariate distribution (with independence across angles) allowing for direct, probabilistic control over the values of the angles, and hence, the correlations.
 - This provides flexibility for real-world applications, allowing for the estimation and use of directionally conditional probabilities that are empirically based, based on sme, or both.
 - Robustness, even when eigenvalues are empirically challenging (i.e. virtually zero), as when positive definiteness must be enforced via an algorithm like Higham (2002).
- Our solution satisfies all the above criteria. We use Makalic and Schmidt (2018) as a starting point, and derive a broader distribution of which theirs is but a special case. In other words, we introduce a location parameter, allowing the distribution of the true angle values to deviate from $\pi/2$, which is their (only) center when perturbing the identity matrix.

QuantMinds International

OuantMinds

Step by Step implementation of proposed sampling algorithm to obtain the distribution of an estimated / observed Correlation Matrix:

a. Obtain the p(p-1) angles corresponding to the original correlation matrix via Rapisarda et al. (2007). $\theta_{i,j}$ below are the angles and $b_{i,j}$ are the cells in the Cholesky reparameterization matrix B shown previously above:

$$\theta_{i,1} = \arccos\left(b_{i,1}\right), i = 2, \dots, p, \text{ and } \theta_{i,j} = \arccos\left(\frac{b_{i,j}}{\prod_{k=1}^{j-1}}\sin\left(\theta_{i,k}\right)\right), 2 \le j < i \le p. *$$

- b. Systematically perturb these angles independently according to some distribution (C3 below).
- c. For each perturbation, translate each matrix of angles back into the Cholesky reparameterization matrix B shown above. j-1

$$b_{1,1} = 1, \ b_{i,1} = \cos(\theta_{i,1}), \ i = 2, \dots, p \text{ and } b_{i,j} = \frac{\prod_{k=1}^{n} \sin(\theta_{i,k}) \text{ for } i = j}{\cos(\theta_{i,j}) \cdot \prod_{k=1}^{j-1} \sin(\theta_{i,k}) \text{ for } 2 \le j \le i-1}$$

- d. Obtain a perturbed correlation matrix via: $R_{perturbed_i} = BB^t$
- e. Perform a. d. some large number of times to obtain the matrix's distribution.

*NOTE: There is an error in this formula in both Pourahmadi and Wang (2015) and Zhang et al., (2013). We have verified, both empirically and analytically, that Rapisarda et al. (2007) above is right (see 8/8/20 Email from JD Opdyke to Dan Kirsner). QuantMinds International © J.D. Opdyke 22 of 45

- Our contribution is to provide the angle distribution for b. in the steps above.
- We start with Makalic and Schmidt's (2018) distribution:

 $f_X(x) = c_k \cdot \sin^k(x), x \in (0,\pi), k = 1,2,3...$ where normalization constant $c_k = \frac{\Gamma(k/2+1)}{\sqrt{\pi}\Gamma(k/2+1/2)}$

• We transform X as follows:

$$z_{i,j} = \arctan\left(\tan\left(\theta_{i,j} - \frac{\pi}{2}\right) + \tan\left(x - \frac{\pi}{2}\right)\right) + \frac{\pi}{2}$$

where $\theta_{i,j}$ is the angle corresponding to the observed correlation value.

- This translates the range of the observed angle from (0, π) to (-∞,∞), does the same for the sampled value from Makalic and Schmidt's (2018) distribution, x, then shifts the sampled value x by θ_{i,j} and retransforms the support back to (0, π). The resulting random variable z is median-centered on the observed angle(s) because tan and arctan are monotonic transformations.
- To obtain the distribution of Z, we use the familiar relationship between the densities of transformed random variables,

$$g(z) = f(x(z)) \cdot \left| \frac{dx}{dz} \right|$$

to obtain...

QuantMinds International

$$f_{z}(z;k,\theta) \sim \frac{\Gamma(k/2+1)}{\sqrt{\pi}\Gamma(k/2+1/2)} \cdot \left(\sin\left[\arctan\left(\tan\left(z - \pi/2 \right) - \tan\left(\theta - \pi/2 \right) \right) \right] \right)^{k} \cdot \frac{\sec(z - \pi/2)^{2}}{1 + \left(\tan\left(z - \pi/2 \right) - \tan\left(\theta - \pi/2 \right) \right)^{2}},$$

where $z \in \{0,\pi\}, \ 0 < \theta < \pi, \ k = 1,2,3..., \ c_{k} = \frac{\Gamma(k/2+1)}{\sqrt{\pi}\Gamma(k/2+1/2)}$

This simplifies to "C3" (for "Cosecant, Cotangent, Cotangent")

$$f_{Z}(z;k,\theta) \sim c_{k} \cdot \left[\csc(z)\right]^{2} \left(1 + \left[\cot(z) - \cot(\theta)\right]^{2}\right)^{-1-k/2},$$

International

with a CDF of $F_{Z}(z;k,\theta) \sim \frac{1}{2} + c_{k} \cdot \left(\cot(\theta) - \cot(z)\right) \cdot {}_{2}F_{1}\left[\frac{1}{2}, 1 + \frac{k}{2}; \frac{3}{2}; -\left(\cot(z) - \cot(\theta)\right)^{2}\right],$ where ${}_{2}F_{1}[a,b;c;u] = \sum_{n}^{\infty} \frac{(a)_{n}(b)_{n}}{(c)_{n}} \cdot \frac{u^{n}}{n!}$ where $(h)_{n} = h(h+1)(h+2)\cdots(h+n-1), n \ge 1, (h)_{0} = 1$

Note the correspondence between the CDF and the identity $\arctan(u) = u \cdot {}_{2}F_{1}\left[1, \frac{1}{2}; \frac{3}{2}; -u^{2}\right]$

Makalic and Schmidt (2018) is now just a special case of C3 where the location parameter $heta=\pi/2$.



OuantMinds International

3

3

© J.D. Opdyke 25 of 45





© J.D. Opdyke 26 of 45

- Our sampling approach provides a distribution of an estimated/observed correlation matrix that
 - 1. allows us to correctly, accurately, and comprehensively account its sampling/estimation error.
 - 2. provides direct, probabilistic control over individual correlations WITHOUT violating the positive definiteness of the matrix; in other words, we have a valid methodology for changing individual correlation values based on the corresponding probabilities associated with the magnitude of the change. This allows us to answer questions like:
 - a. What are the chances of the correlation between assets X and Y increasing at least 0.2?
 - b. What is the minimum decrease in the correlation between X and Y associated with a 50% probability?
 - c. What are the chances of the correlations between X and Y, X and Z, AND Z and Q concurrently changing by at least +0.2, +0.1, and –0.5, respectively?
 - 3. allows for the **identification of percentiles (quantiles) of the entire distribution**. Larger percentiles (e.g. 99%tile) would correspond to the generalized effects of extreme scenarios (e.g. pandemic+economic upheaval).
- In other words, we provide one, consistent method for performing two types of stress, either separately or concurrently: Targeted or Scenario Stress (2. above), and Generalized Stress (3. above).
- Targeted Stress = change specific correlations, according to C3, by directionally conditional probabilities (e.g. increase/decrease an angle's CDF by 0.10 to correspond with a particular scenario (say, a 1-in-10-year event)); and Generalized Stress = perturb C3 many times to create the empirical distribution of the observed/estimated correlation matrix (and obtain its empirical percentiles).

QuantMinds International

٠

- **Targeted Stress** is changing those particular correlations, in the controlled, probabilistic manner described above, relevant to a particular scenario (e.g. pandemic+economic upheaval).
- Below we show a screen shot of an Excel workbook that implements Targeted Stress on a user-٠ defined correlation matrix (available upon request). Its small size allows us to use closed-form solutions for calculating the CDF of C3, but larger matrices require numeric integration (or fast and accurate calculation of Gauss hypergeometric functions).*
- Each of the C3 sampling steps identified above are shown explicitly in the workbook, and changes in ٠ the correlations are driven by probabilities provided by the user (e.g. a 1-in-10 year change (90th%tile of the annual distribution, or probability=0.10) induces the change in a particular correlation).
- These probabilities can be empirically estimated, based on subject matter expertise exclusively, or some combination of the two. As long as the changes in correlation are directionally correct vis-à-vis the scenario, the probabilities will deterministically change the correlations accordingly.**
- Note that Ng et al. (2013) and Yu et al. (2014) identify changes in "peripheral" correlations due to changes in "core" correlations, but when taking a polar angle view of the problem as herein, such correlation 'dependencies' already are captured by the trigonometric transformation of the Cholesky factor; in other words, Targeted Stressing will necessarily alter correlations not explicitly altered by the user, as it must to maintain positive definiteness (see our attached Excel workbook to observe this interactively). Our approach in using C3 goes a step beyond by explicitly distinguishing between targeted and generalized stressing, and allowing either or both to be implemented concurrently.

*NOTE: We have used numeric integration here on larger matrices (over 100x100) successfully with major statistical programming packages (SAS® and R), but much larger matrices (p>1,000) may require symbolic programming languages, like Mathematica® or Maple, to calculate the CDF values (either via numeric integration or fast and accurate calculation of Gauss hypergeometric functions). **NOTE: Note that, per Zhang et al. (2015) and Lu et al. (2020), correlations decrease monotonically in their corresponding angles. **OugntMinds** OuantMinds International International

- Note that taking the spherical angle approach to this problem gives us more control over the ٠ entire matrix when specifying changes to specific correlations: the fact that the 'dependencies' between correlations are explicitly determined by the trigonometric transformation of the Cholesky factor, rather than as an indirect consequence of simulated inputs as in Ng et al. (2013) and Yu et al. (2014), means that we have some direct control over which other combinations of correlations get altered as the result of our scenario-specific targeted changes, due to one fact: we can permute the rows and the columns of any correlation matrix without affecting the distributions of either the individual correlations OR the entire matrix.*
 - *The distribution of any correlation value per Forrester and Zhang (2018) per Pourahmadi and Wang (2015) (\leftarrow same result, different notation \rightarrow) $(1 r_{i,i}^2)^{[2(k-1)+p]}$

$$r_{i,j} \propto \left(1 - r_{i,j}^2\right)^{\left[k+p/2-1\right]}, \ i > j = Beta\left(k + \frac{p}{2}, k + \frac{p}{2}\right) \text{ distribution on } \left(-1, 1\right) \qquad \qquad \overline{Beta\left(2k + p - 1, \frac{1}{2}\right)}$$

- And the distribution of the ended define mean $f(r) = c_p(k) \left[\left(\prod_{j=1}^p \prod_{l=1}^{j-1} \left[\sin(\theta_{j,l}) \right]^2 \right)^k = c_p(k) \left[\det(R) \right]^k, \ j = 1, \dots, p-1, \ i > j$ where $c_p(k) = \prod_{j=1}^{n-1} \left[\frac{\Gamma\left(\frac{2k+j}{2}+1\right)}{\sqrt{\pi}\Gamma\left(\frac{2k+j+1}{2}\right)} \right]^j$ is the normalization constant. And the distribution of the entire matrix per Pourahmadi & Wang (2015) (see also Forr. & Zhang 2018)
- So if changing one correlation value of a cell in the matrix alters many others (because of the ٠ trigonometric functions applied to the Cholesky factorization) which we don't want to alter, we can simply reorder the matrix so that only that correlation (or at least fewer other correlations) change as a result of changing the 'core' scenario correlation(s).
- Again, this can be seen interactively in the Excel workbook. **OuantMinds** International © J.D. Opdyke 29 of 45

III. Using C3 for Targeted Scenario Analytics



NOTE: This Excel workbook implementation is available from the author upon request.

QuantMinds International



© J.D. Opdyke 30 of 45

- Now we turn to Generalized Stress, wherein we
 - 1. generate a distribution of the correlation matrix around its observed values using C3,
 - 2. identify a desired percentile and corresponding quantile of the distribution (e.g. VaR99), &
 - 3. **systematically increment** the entire observed correlation matrix **to match this percentile value** (while automatically enforcing positive definiteness).
- Risk of the entire portfolio only can be accurately measured when taking into account its dependence structure, which allows for measurement of its diversification benefit. This is the difference between a portfolio of perfectly dependent assets/risks one wherein the values/risks of p assets move perfectly in tandem and one wherein assets/risks sometimes move in opposite directions. The former provides NO benefit from diversification, because the assets all are identical, while the later provides a diversification benefit whose size depends on the degree to which asset values/risks move in opposite directions (as well as the relative sizes of these assets).
- Herein we are using the correlation matrix as the dependence structure, the dimensionality of which is reduced to a univariate distribution of 'capital' (the actual \$ associated with the marginal loss/returns distributions) so that diversification benefit can be quantified. The strong assumption of multivariate ellipticity of our marginal distributions (of assets/risks) allows us to use the simple metric provided by the 'square root rule' to do this (see Solvency II). For the <u>comparative</u> purposes of this paper, this assumption is largely ancillary and does not change the primary findings or specific results of this paper. But a quantification of portfolio risk and diversification benefit of a real world portfolio in an <u>absolute</u> sense requires a model framework with more realistic assumptions, such as those provided by various forms of distributionally flexible copulas.

• So we quantify 'capital' and measure diversification benefit using 'the square root rule' shown below:

Capital = $\sqrt{VRV^t}$ where V = vector of p VaRs and R = correlation matrix of dimension p x p

- In addition to our choice of metric for measuring capital and diversification benefit, our chosen risk metric – Value-at-Risk ("VaR"), which is simply the quantile of each marginal loss/returns distribution (and ultimately of the capital distribution) – is ancillary to the main findings of this paper. Use of other risk metrics (Expected Shortfall, Expectiles, etc.) will provide similar results on a relative basis.
- Once we have generated a distribution of, say, 100,000 correlation matrices around the observed correlation matrix via perturbation of C3, and converted these to 100,000 capital values via the square root rule, we identify the quantile associated with the (conditional) 99% tile of this distribution: in other words, the (conditional) VaR99 (this is 'conditional' on the observed matrix, i.e. the 99% tile BEYOND the estimated, base matrix).
- However, there are many correlation matrices that provide this same value of VaR99: for example, if we have only p=2, the solution providing VaR99=X Capital is an entire ellipse (under the above assumptions) of correlation matrices; so if p = 100, the solution is a 100-dimensional hyper-ellipse.
- The point of generalized stressing is that we have no priors regarding specific pairs of correlations where stress may manifest. This is helpful in catching difficult-to-anticipate and/or difficult-toquantify second and third order effects of a large, multivariate, impactful scenario (e.g. pandemic+economic upheaval). But this also actually gives us direction as to which correlation matrix to select that satisfies VaR99 = X Capital.

 Once we have the value of capital at VaR99, we can simply increase the entire observed correlation matrix by increasing each cell the same percentage of its value towards an upper bound of 1.0. In other words,

[1] $R^* = R + (1 - R)\alpha$ where *R* is a valid (positive definite) correlation matrix and $0 \le \alpha < 1$

- When R* yields a value of capital equal to the VaR99, we select that R*.
- α is incremented using any convergence algorithm (we use bisection, which is guaranteed to converge, and is not materially slower here than more complicated approaches).
- Note that in [1] and related formulae below, (1 R) indicates the 'unit matrix' minus R, where the unit matrix has values of 1.0 in every cell: "1" here does NOT represent the identify matrix, which has values of 1.0 only on the diagonal and zeros in every other cell.
- This approach [1] is useful, straightforward, and appropriate for a number of reasons:



 $[1] \quad R^* = R + (1-R)\alpha,$

- This approach [1] is useful, straightforward, and appropriate for a number of reasons:
 - This approach makes sense conceptually: we are increasing all correlations by a percentage of the distance (α) between their present value and the upper bound of 1.0. By design, we are stressing the entire matrix similarly, because we have no priors regarding the relative sizes of the increases of any pairwise correlations; this is reserved for the particulars of targeted stressing, which would employ hyper-stress on particular correlation values, and can (and conservatively, should) be implemented in conjunction with this generalized stressing of the entire matrix to aid in capturing overlooked or underestimated, yet critical second/third order effects that are difficult to anticipate and/or quantify.
 - This approach respects the required bounds of a valid correlation matrix $(-1 \le r_{i,j} \le 1 \text{ for all } i, j)$
 - This approach respects the requirement of unit diagonals $(r_{i,i} = 1 \text{ for all } i)$
 - Importantly, this approach automatically respects the requirement of Positive Definiteness:

if *R* is positive definite and $0 \le \alpha < 1$, then R^* is positive definite. See next page for proof.

QuantMinds International

IV. Using C3 for Generalized Stress Testing

for [1] $R^* = R + (1 - R)\alpha$, if *R* is positive definite and $0 \le \alpha < 1$, then R^* is positive definite. PROOF: $R + (1 - R)\alpha = R + \alpha - \alpha R = R(1 - \alpha) + \alpha$ Term A Term B

A. As $0 \le \alpha < 1$, $R(1-\alpha)$ is *R* multiplied by a real number > 0 which, according to Horn & Johnson (2012), remains positive definite (see Horn & Johnson, Obs.7.1.3, p.430)

B. And adding a constant greater than zero (i.e. α) to any positive definite matrix yields a matrix which is positive definite, as shown below.

Let
$$C = R(1-\alpha)$$
, $x \in \mathbb{R}^n \setminus \mathbf{0}$
then $x^t Cx = \sum_i \sum_j x_i x_j c_{i,j} > 0$
 $x^t (C+\alpha)x = \sum_i \sum_j x_i x_j (c_{i,j} + \alpha) = \sum_i \sum_j (x_i x_j c_{i,j} + x_i x_j \alpha) = \sum_i \sum_j x_i x_j c_{i,j} + \alpha \left(\sum_i x_i\right)^2$
since $x \in \mathbb{R}^n$, $\alpha \sum_i x_i^2 \ge 0$, so $\sum_i \sum_j x_i x_j (c_{i,j} + \alpha) \ge \sum_i \sum_j x_i x_j c_{i,j} > 0$, so
 $x^t (C+\alpha)x \ge x^t Cx > 0$
So $C + \alpha$ is positive definite, and $R(1-\alpha) + \alpha = R + (1-R)\alpha$ is positive definite.

QuantMinds International

• Just as the Excel workbook provides working examples of Targeted Stress, below we provide examples of the above described application of Generalized Stress.*

Again, $Capital = \sqrt{VRV^t}$ where V = vector of p VaRs and R = correlation matrix of dimension pxp

- Both the base correlation matrix and the distribution of the vector of VaRs will determine results.
- VaRs in large portfolios typically are (VERY) skewed. We use here 3 vectors of p=100 VaRs:
 - a constant VaR vector = \$50m for each of the 100 marginal distributions as a benchmark/baseline;
 - an exponentially distributed VaR vector (λ =1) for the mildly skewed case (e.g. largest 4 of 100 VaRs are 16.5% of the total); and
 - a Pareto (a=1.2) distributed VaR vector for the very (more typical) skewed case (e.g. top 4 of 100 VaRs are 40% of the total).
 - All VaR vectors are scaled to total \$5b.
- We use two 100x100 correlation matrices: one where correlations are drawn at random from a uniform distribution where 20% each are -0.50, -0.25, 0, 0.25, 0.50, and an unbalanced and more extreme matrix where half of those in the 0.50 bucket are increased evenly to 0.75 (5%) and to 0.85 (5%). Positive definiteness is enforced via Higham (2002), so our smallest eigenvalues are virtually 0.
- We conducted **100,000 simulations** using C3 for each VaR vector + Correlation Matrix combination. Diversification benefit (the difference from the undiversified \$5b total) is large because of the large proportion of correlations that are negative (fully 40%), but the goal was to test wide-ranging values.

*<u>NOTE</u>: The concave function in the capital calculation (square root) means that the estimates of capital will be biased downwards due to Jensen's Inequality (see Jensen, 1906). We focus here on conditional VaRs (i.e. the VaR99 of the distribution <u>beyond</u> the observed/estimated base correlation matrix, not of the entire distribution) so the difference between the base capital estimate and the conditional VaR99 will remain largely if not wholly unaffected by Jensen's inequality. To obtain an approximation of the size of this bias on an absolute basis we conducted the above simulations using the identity matrix as the base matrix (with the two different skewed VaR vectors) and compare the means of the capital distributions to the true value under the identity matrix. The bias as approximated by these differences was de minimis (always between negative \$2m and zero).



OuantMinds

International

 Even under very skewed distributions of VaRs and unbalanced correlation matrices we obtain useable and useful distributions based on perturbing our observed/estimated matrices using C3. Note that when compared to baseline capital, increases from generalized stress capital (defined as the conditional VaR99) range from about 10–20%.

VaR Distribution	Flat	Pareto	Pareto	Exponential	Exponential
Correlation Matrix	Balanced	Balanced	Unbalanced	Balanced	Unbalanced
Base Matrix	\$483	\$1,259	\$1,373	\$596	\$816
Base %tile	48.6%	50.6%	69.5%	40.8%	71.7%
Conditional 99%tile	\$573	\$1,407	\$1,511	\$717	\$919
Base - Cond. 99%tile	\$90	\$148	\$137	\$121	\$103
%Difference	18.7%	11.7%	10.0%	20.3%	12.6%
25%tile	\$461	\$1,218	\$1,299	\$578	\$759
50%tile	\$484	\$1,258	\$1,342	\$606	\$790
75%tile	\$507	\$1,297	\$1,384	\$635	\$821
90%tile	\$528	\$1,332	\$1,421	\$662	\$849
95%tile	\$541	\$1,353	\$1,444	\$678	\$866
99%tile	\$565	\$1,392	\$1,486	\$708	\$898
99.5%tile	\$573	\$1,407	\$1,501	\$720	\$909
99.9%tile	\$592	\$1,439	\$1,533	\$743	\$936

Generalized Stress Simulation Based on C3: Results



© J.D. Opdyke 38 of 45

IV. Using C3 for Generalized Stress Testing





- Note that the unbalanced correlation cases are not centered on their observed/estimated (base) correlation matrices, and even the exponential 'balanced' case is not exactly centered.
- Given the skewed natures of C3 (when $\theta \neq \pi/2$) and to a lesser degree the VaR vector, and the fact that the transformations from angles to Cholesky factor to correlations are highly nonlinear, there is no reason to expect to see such 'centering,' which is what we'd observe under linear transformations.
- However, we explore the possibility that these transformations might be closer to linear if the observed angles were the means of the C3 distributions rather than the medians.
- To implement this we empirically shifted the value of the 'observed' angle, $\theta_{i,j}$, until the mean of $z_{i,j}$ was equal to the original value of $\theta_{i,j}$; recall that

$$z_{i,j} = \arctan\left(\tan\left(\theta_{i,j} - \frac{\pi}{2}\right) + \tan\left(x - \frac{\pi}{2}\right)\right) + \frac{\pi}{2}$$

International

• In other words, we obtain $\theta_{i,j}^* = \theta_{i,j} \pm \Delta$ (via bisection) such that

$$\theta_{i,j} = \frac{\Gamma(k/2+1)}{\sqrt{\pi}\Gamma(k/2+1/2)} \int_{0}^{\pi} z \cdot \left[\csc\left(z\right)\right]^{2} \left(1 + \left[\cot\left(z\right) - \cot\left(\theta_{i,j}^{*}\right)\right]^{2}\right)^{-1-k/2} dz$$

- To provide a numerical example, if k=2 and $\theta_{i,j} = \pi/4 = 0.785398$, then $\theta_{i,j}^* = 0.6728905$.
- Preliminary results show that the capital distributions are no more centered on the base matrices (based on comparing their quantiles) when using the mean-centered C3 distributions; and the capital distributions generally remain the same (again, based on little change in their quantiles).
- Further research is focusing on analytically deriving a version of C3 centered on means.

- In addition to the correlation (capital) distributions above, **Generalized Stressing** provides us with a ٠ correlation matrix R* wherein the entire matrix is stressed to achieve the value of, say, VaR99 (or any quantile of the distribution).
- But what if we wanted to ALSO use Targeted Stress concurrently? We have a scenario where ٠ generalized stress is appropriate for conservatively capturing unanticipated or difficult-to-quantify second and third order effects, but we also know where particular areas of stress will manifest: can we do both together?
- Yes! Since the Generalized Stress matrix is positive definite, and all Targeted Stress changes to ٠ specific correlations in it will result in a positive definite matrix, we can hyper-stress those pairwise correlations for which Targeted Stress values are greater than Generalized Stress values, and the resulting matrix will remain positive definite. And both the effects of Generalized Stress and Targeted Stress will be captured in the final matrix.

For i > j, IF $\theta_{Obs_{i,j}} \le \theta_{Tar_{i,j}}$ THEN $\theta_{i,j}^* = \theta_{Tar_{i,j}}$; ELSE IF $\theta_{Gen_{i,j}} \le \theta_{Tar_{i,j}}$ THEN $\theta_{i,j}^* = \theta_{Gen_{i,j}}$; ELSE $\theta_{i,j}^* = \theta_{Tar_{i,j}}$; Now apply the trigonometric transformations to $\theta_{i,i}^*$ to obtain B^* , and $R_{combined}^* = B^* (B^*)^T$

- In other words, simply apply Targeted Stress to the Generalized matrix (and ignore targeted stress) correlations that are less than those corresponding values in the generalized matrix UNLESS the targeted stress values are less than the original observed correlations).
- This framework for scenario/stress testing the correlation matrix using C3 to perturb the polar angles ٠ associated with each correlation, either for targeted correlations, the entire correlation matrix generally, or both combined - makes this approach extremely flexible and useful in a wide range of settings requiring the accurate and comprehensive measurement of portfolio risk via stress testing. **OugntMinds OuantMinds** International International

- Accurately and comprehensively measuring portfolio risk requires directly scenario/stress testing its dependence structure, which in many cases is measured by a Pearson's correlation matrix.
- We have presented an easily implemented, flexible, robust approach to scenario/stress testing the observed/estimated correlation matrix that borrows from the 'identity matrix / random correlation' literature to develop a new distribution for the spherical angles on which individual correlations are based.
- This approach allows for the consistent yet flexible application of **one method for both targeted (extreme)** scenario testing, generalized stress testing, and both combined.
- <u>Targeted stressing</u> provides direct, probabilistic control over specific correlations in the matrix that are most relevant to a given scenario (without violating positive definiteness), and <u>generalized stressing</u> provides an entire distribution of the matrix, giving much needed percentiles (quantiles) useful for assessing system-wide effects and capturing difficult to anticipate and/or quantify second and third order effects of wideranging scenarios (e.g. pandemic+economic upheaval).
- Our approach allows us to directly and transparently answer questions such as:
 - What are the chances of the correlation between assets X and Y increasing at least 0.2?
 - What is the minimum decrease in the X,Y correlation associated with a 50% probability?
 - What are the chances of the correlations between X and Y, X and Z, AND Z and Q concurrently changing by at least +0.2, +0.1, and –0.5, respectively?
 - What is the 99%tile / VaR99 of the distribution of my correlation matrix?
- Directly sampling correlations' spherical angles appears to be more robust than competing methods, such as perturbing eigenvalues, at least under challenging empirical conditions, such as using matrices on which positive definiteness has been enforced (e.g. via Higham, 2002); this is a relatively common situation for large(r) enterprise-level, and even many investment portfolios.

QuantMinds International

OughtMinds

- Estimating (large) correlation matrices under real world conditions is challenging even under ideal empirical conditions as stability issues are endemic to all estimators, even those traditionally considered to be optimal (see Lopez de Prado, 2018).
- Further research here could explore the possibility of using C3 in a regression on the correlation estimates, where C3 is the error term; the use of previously unused covariates (with this appropriate error term) could increase the precision and robustness of the estimates, if not also their accuracy.
- This could be much more straightforward than other regression-like approaches, such as those of Zhang et al. (2015), Lee at al. (2019), Packham and Woebbeking (2019), and Lu et al. (2020), since it would solely be a direct estimation of a location parameter (that measures the strength of linear relationship).
- Potentially using C3 in this way could augment powerful existing methods for increasing the precision, accuracy, and robustness of estimates of correlation matrices, such as de-noising, de-toning, and clustering (see Lopez de Prado, 2016, 2018, 2019 and 2020). Even as is, C3 could be complementary to and useful for correlation matrices that already have had these methods applied to them.



Contact Information:

J.D. Opdyke Head of Enterprise Risk Analytics VP – Financial Risk and Measurement, ERRM Division John.Opdyke@Allstate.com 617-943-6463







VIII. References

- Bank for Intl Settlements, Basel Committee on Banking Supervision, Working Paper No 19, "Messages from the academic literature on risk measurement for the trading book," 31 January 2011.
- Bohn, W., Hornik, K., (2014), "Generating random correlation matrices by the simple rejection method: Why it does not work," Stat. & Prob.Letters, 87 (C), 27-30.
- Cordoba, I., Varando, G., Bielza, C., and Larranaga, P., (2018), "A fast Metropolis-Hastings method for generating random correlation matrices," IDEAL, pp. 117-124, part of Lecture Notes in Computer Science, Vol 11314.
- Cordoba, I., Varando, G., Bielza, C., and Larranaga, P., (2019) R implementation of Cordoba et al. (2018) in github.com/irenecrsn/rcor
- CRAN, randcorr package, Version 1.0, published 11/16/18, Schmidt, D., and Makalic, E., https://cran.r-project.org/web/packages/randcorr/index.html. ٠
- Email correspondence, 08/08/20, from JD Opdyke to Dan Kirsner. ٠
- Fengler, M., and Schwendner, P., (2003), "Correlation Risk Premia for Multi-Asset Equity Options," SFB 373 Discussion Papers 2003,10, Humboldt University of Berlin, Interdisciplinary . Research Project 373: Quantification and Simulation of Economic Processes.
- Forrester, P., and Zhang, J., (2020), "Parameterizing Correlation Matrices," Journal of Multivariate Analysis, Vol. 178.
- Galeeva, R., Hoogland, J., & Eydeland, A., (2007), "Measuring Correlation Risk," publicly available manuscript.
- Ghosh, R., Mallick, B., and Pourahmadi, M., (2020), "Bayesian Estimation of Correlation Matrices of Longitudinal Data," Bayesian Analysis, forthcoming. ٠
- Golub, A., and Guo, Z., (2012), "Correlation Stress Testing Using the Random Matrix Theory: An Empirical Implementation to the Chinese Market," presented at Marie Curie ITN Final . Conference. Financial Risk Management & Risk Reporting, April 11, 2019.
- Hardin, J., Garcia, S., and Golan, D., (2013), "A Method for Generating Realistic Correlation Matrices," The Annals of Applied Statistics, Vol. 7, No. 1, 1733-1762.
- Higham, N., (2002), "Computing the Nearest Correlation Matrix – A Problem from Finance," IMA Journal of Numerical Analysis, 22, 329-343.
- Horn, R., Johnson, C., (2012), Matrix Analysis, 2nd Ed., Cambridge University Press.
- Jensen, J. L. W. V. (1906). "Sur les fonctions convexes et les inégalités entre les valeurs moyennes," Acta Mathematica, 30 (1): 175–193. ٠
- Junior, L., and Franca, I., (2012), "Correlation of Financial Markets in Times of Crisis," Physica A: Statistical Mechanics and its Applications, 391(1-2), 187-208. .
- Kuipec, p., (1998), "Stress Testing in a Value at Risk Framework," The Journal of Derivatives Fall, 6(1) 7-24.
- Lee, K., Cho, H., Kwak, M., & Jang, E., (2019), "Estimation of Covariance Matrix of Multivariate Longitudinal Data Using Modified Cholesky and Hypersphere Decompositions," Biometrics, 1-12.
- Lewandowski, D., Kurowicka, D., and Joe, H., (2009), "Generating random correlation matrices based on vines and extended onion method," Journal of Multivariate Analysis, 100(9), 1989-2001. .
- Loland, A., Huseby, R., Hjort, N., and Frigessi, A., (2013), "Statistical Corrections of Invalid Correlation Matrices," Scandinavian Journal of Statistics, 40(4), 1-18.
- Lopez de Prado, M., (2016), "Building Diversified Portfolios that Outperform Out of Sample," The Journal of Portfolio Management, 42 (4) 59-69.
- ٠ Lopez de Prado, M., (2018), Machine Learning Asset Allocation, https://ssrn.com/abstract=3469961.
- Lopez de Prado, M., (2019), A Robust Estimator of the Efficient Frontier, https://ssrn.com/abstract=3469964. .
- Lopez de Prado, M., (2020), Machine Learning for Asset Managers, Cambridge University Press Cambridge Elements, Quantitative Finance, Cambridge, UK.
- Lu, F., Xue, L., and Wang, Z., (2019), "Triangular Angles Parameterization for the Correlation Matrix of Bivariate Longitudinal Data," J. of the Korean Statistical Society, 49:364-388. ٠
- Madar, V., (2015), "Direct formulation to Cholesky decomposition of a general nonsingular correlation matrix," Stat. & Prob. Letters, Vol. 103, pp142-147. ٠
- Makalic, E., Schmidt, D., (2018), "An efficient algorithm for sampling from sin(x)^k for generating random correlation matrices," arXiv: 1809.05212v2 [stat.CO].
- Ng, F., Li, W., and Yu, P., (2014), "A Black-Litterman Approach to Correlation Stress Testing," Quantitative Finance, 14:9, 1643-1649.
- Opdyke, J.D., (2019), "Getting Extreme VaR Right: Eliminating Convexity & Approximation Biases Under Heavy-tailed, Moderately-Sized Samples," forthcoming, presented at RiskMinds ٠ International, December 2, 2019, Amsterdam, NL.
- Packham, N., & Woebbeking, F., (2019), "A factor-model approach for correlation scenarios and correlation stress testing," J. of Banking & Finance, 101:92-103. ٠
- Pourahmadi. M., Wang, X., (2015), "Distribution of random correlation matrices: Hyperspherical parameterization of the Cholesky factor," Stat. and Probability Letters, 106, (C), 5-12. ٠
- Pritsker, M., (2006), "The hidden dangers of historical simulation," Journal of Banking & Finance, Volume 30, Issue 2, 561-582.
- Rapisarda, F., Brigo, D., & Mercurio, F., (2007), "Parameterizing Correlations: A Geometric Interpretation," IMA Journal of Management Mathematics, 18(1), 55-73. .
- Rebonato, R., and Jackel, P., (2000), "The Most General Methodology for Creating a Valid Correlation Matrix for Risk Management and Option Pricing Purposes," J. of Risk, Vol 2., No. 2., 17-27.
- So, M., and Wong, J., (2013), "Stress Testing Correlation Matrices for Risk Management," The North American Journal of Economics and Finance, 26, 310-322.
- . Solvency II, Directive 2009/138/EC, 11/25/09, p.L335-7, https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32009L0138&from=EN
- Stuart, A., and Ord, J.K., (1994), Kendall's Advanced Theory of Statistics: Volume 1, Distribution Theory, Sixth Edition. Edward Arnold, London. .
- Turkay, S., Epperlein, E., and Christofides, N., (2003), "Correlation Stress Testing for Value-at-Risk," Journal of Risk, 5(4), 75-90.
- Van Oest, R., (2019), "Unconstrained Cholesky-based Parametrization of Correlation Matrices," Communications in Statistics Simulation and Computation, DOI: . 10.1080/03610918.2019.1628271.
- Wang, Z. Wu, Y., and Chu, H., (2018), "On equivalence of the LKJ distribution and the restricted Wishart distribution." arXiv:1809.04746v1. ٠
- Yu, P., Li, W., Ng, F., (2014), "Formulating Hypothetical Scenarios in Correlation Stress Testing via a Bayesian Framework," The North Amer. J. of Econ. and Finance, Vol 27, 17-33.
- Zhang, W., Leng, C., and Tang, Y., (2015), "A Joint Modeling Approach for Longitudinal Studies," Journal of the Royal Stat. Society, Series B, 77(1), 219-238. QuantMinds International antminas

International

© J.D. Opdyke 45 of 45