Simulations with exact means and covariances

Attilio Meucci presents a simple method to generate scenarios from multivariate elliptical distributions with given sample means and covariances, and shows an application to the risk management of a book of options

To perform risk and portfolio management, we must represent the distribution of the risk factors that affect the market. The most flexible approach is in terms of scenarios and their probabilities, which includes historical scenarios, pure Monte Carlo and importance sampling (see Glasserman, 2004).

Here, we present a simple method to generate scenarios from elliptical distributions with given sample means and covariances. This is very important in applications such as mean-variance portfolio optimisation, which are heavily affected by incorrect representations of the first two moments.

The same problem has been tackled by, among others, Wedderburn (1975), Cheng (1985), Li (1992) and Alexander, Ledermann & Ledermann (2008). However, these approaches require handling matrices or loops of the same size as the number of scenarios. This quickly becomes intractable for large Monte Carlo simulations. Instead, our method is a multivariate generalisation of the intuitive shift/rescaling that appears in, for example, Boyle, Broadie & Glasserman (1995). This method amounts to solving a matrix Riccati equation independent of the number of scenarios.

Methodology

Consider a multivariate normal distribution:

$$\mathbf{X} \sim N(\mathbf{m}, \mathbf{S}) \tag{1}$$

where $\mathbf{m}$ is an arbitrary expected value and $\mathbf{S}$ is an arbitrary covariance matrix. Consider the representation of this distribution in terms of the probability-scenario pairs $(p_j, \mathbf{x}_j)$, $j = 1, \ldots, J$. Our aim is to ensure that:

$$\hat{\mathbf{m}}_x = \mathbf{m}, \quad \hat{\mathbf{S}}_x \equiv \mathbf{S} \tag{2}$$

where:

$$\hat{\mathbf{m}}_x \equiv \sum_{j=1}^{J} p_j \mathbf{x}_j, \quad \hat{\mathbf{S}}_x \equiv \sum_{j=1}^{J} p_j \mathbf{x}_j \mathbf{x}'_j - \hat{\mathbf{m}}_x \hat{\mathbf{m}}'_x \tag{3}$$

denote the sample mean and sample covariance of $(p_j, \mathbf{x}_j)$. To do so, one can either constrain the probabilities p_j, or the scenarios $\mathbf{x}_j$. The former approach, pursued in, for example, Avellaneda (1999), D'Amico, Fusai & Tagliani (2003), Glasserman & Yu (2005) and Meucci (2008), is very flexible, but for large-dimensional markets it becomes computationally challenging. Here, we choose the second route, which relies on the affine equivariance of the elliptical distributions.

First, we produce an auxiliary set of scenarios:

$$\left(\tilde{p}_j, \tilde{\mathbf{y}}_j\right), \quad j = 1, \ldots, \frac{J}{2} \tag{4}$$

from the distribution $N(\mathbf{0}, \mathbf{S})$. Then we complement these scenarios with their opposite:

$$\left(\tilde{p}_j, \tilde{\mathbf{y}}_j\right) \equiv \begin{cases} \left(\tilde{p}_j / 2, \tilde{\mathbf{y}}_j\right) & \text{if } j \leq J/2 \\ \left(\tilde{p}_{j-\frac{J}{2}} / 2, -\tilde{\mathbf{y}}_{j-\frac{J}{2}}\right) & \text{if } j > J/2 \end{cases} \tag{5}$$

These antithetic variables still represent the distribution $N(\mathbf{0}, \mathbf{S})$, but they are more efficient (see Boyle, Broadie & Glasserman, 1995) and they satisfy the zero-mean condition $\hat{\mathbf{m}}_{\tilde{y}} \equiv \mathbf{0}$.

Next, we apply a linear transformation to the scenarios $\tilde{\mathbf{y}}_j$, which again preserves normality:

$$\mathbf{y}_j \equiv \mathbf{B}\tilde{\mathbf{y}}_j, \quad j = 1, \ldots, J \tag{6}$$

For any choice of the invertible matrix $\mathbf{B}$, the sample mean is null: $\hat{\mathbf{m}}_y \equiv \mathbf{0}$. To determine $\mathbf{B}$, we impose that the sample covariance $\hat{\mathbf{S}}_y$ matches the desired covariance $\mathbf{S}$. Using the affine equivariance of the sample covariance (see, for example, Meucci, 2005), we obtain a matrix Riccati equation:

$$\mathbf{S} \equiv \mathbf{B}\hat{\mathbf{S}}_{\tilde{y}}\mathbf{B}, \quad \mathbf{B} \equiv \mathbf{B}' \tag{7}$$

To solve this equation we follow Petkov, Christov & Konstantinov (1991). First, we define the Hamiltonian matrix:

$$\mathbf{H} \equiv \begin{pmatrix} \mathbf{0} & -\hat{\mathbf{S}}_{\tilde{y}} \\ -\mathbf{S} & \mathbf{0} \end{pmatrix} \tag{8}$$

1 Sample and population moments coincide in our approach

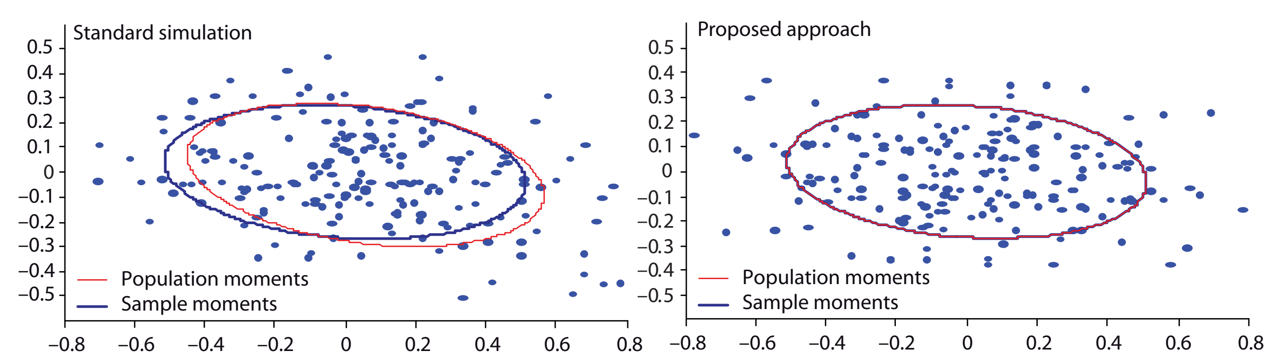

2 Profit and loss of a highly non-linear long-short option portfolio

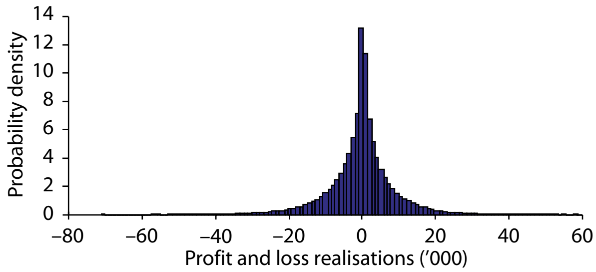

3 VAR computations in a highly non-linear long-short option portfolio

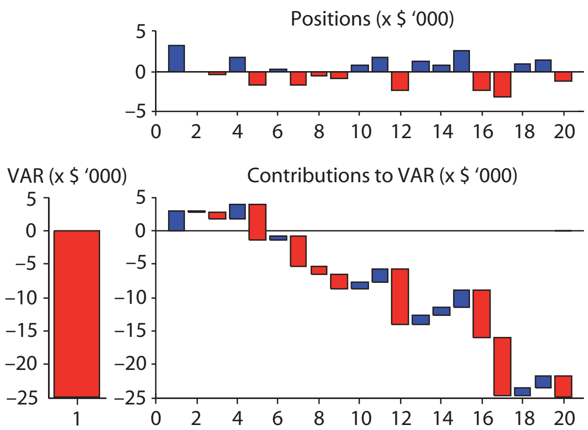

Next we perform its Schur decomposition:

$$\mathbf{H} \equiv \mathbf{UTU}' \tag{9}$$

where $\mathbf{UU}' \equiv \mathbf{I}$ and $\mathbf{T}$ is upper triangular with the eigenvalues of $\mathbf{H}$ on the diagonal sorted in such a way that the first N have negative real part and the remaining N have positive real part; the terms in this decomposition are similar in nature to principal components and are calculated by standard software packages (see, for example, Anderson *et al*, 1999). Then the solution of the Riccati equation (7) reads:

$$\mathbf{B} \equiv \mathbf{U}_{LL}\mathbf{U}_{UL}^{-1} \tag{10}$$

where $\mathbf{U}_{UL}$ is the upper left $N \times N$ block of $\mathbf{U}$ and $\mathbf{U}_{LL}$ is the lower left $N \times N$ block of $\mathbf{U}$.

With the solution (10) we can perform the affine transformation (6) and finally generate the desired scenarios:

$$\mathbf{x}_j \equiv \mathbf{m} + \mathbf{y}_j, \quad j=1,\dots,J \tag{11}$$

which satisfy (2), see figure 1, where, as in Meucci (2005), we represent the first two moments of a distribution in terms of an ellipsoid.

Note that the steps (4)–(11) only require a few fractions of a second to run even for large problems. Refer to *www.symmys.com* ⇒ *Teaching* ⇒ *MATLAB* for a fully functional implementation.

The present methodology is based on affine transformations as well as on the affine equivariance of the sample mean and covariance. Therefore, it extends straightforwardly to general elliptical distributions, such as the *t*.

Applications

To illustrate the ubiquitous nature of normal simulations, here we apply our methodology to calculate the value-at-risk and its decomposition in a book of I plain vanilla call options (see also Meucci, 2008). We denote current time as T. We notice the price $P_{T+\tau}$ of a call option at the investment horizon can be written in the format:

$$P_{T+\tau} \equiv P\left(X_{T,y}, X_{T,\sigma}; \mathcal{I}_T\right) \tag{12}$$

where $(X_{T,y}, X_{T,\sigma})$ are risk factors and $\mathcal{I}_T$ represents currently available information. Indeed, consider the Black-Scholes pricing formula:

$$C_{BS}\left(y,\sigma,K,t\right) \equiv y\Phi\left(d_1\right) - Ke^{-rt}\Phi\left(d_2\right) \tag{13}$$

where Φ is the cumulative distribution function of the standard normal distribution and:

$$d_1 \equiv \frac{-\ln\left(K/y\right) + t\left(r + \sigma^2/2\right)}{\sqrt{\sigma^2 t}} \tag{14}$$

$$d_2 \equiv \frac{-\ln\left(K/y\right) + t\left(r - \sigma^2/2\right)}{\sqrt{\sigma^2 t}} \tag{15}$$

Then:

$$P_{T+\tau} = C_{BS}\left(y_T e^{X_{T,y}}, h\left(y_T e^{X_{T,y}}, \sigma_T e^{X_{T,\sigma}}, K, T-\tau\right); K, T-\tau, r\right) \tag{16}$$

In this expression y_T is the current value and $X_{T,y} \equiv \ln(y_{T+\tau}/y_T)$ is the log-change of the underlying; σ_T is the current value and $X_{T,\sigma} \equiv \ln(\sigma_{T+\tau}/\sigma_T)$ is the log-change in $(T - \tau)$-expiry, at-the-money implied volatility; and h is a skew/smile map:

$$h\left(y,\sigma;K,T\right) \equiv \sigma + \alpha\frac{\ln\left(y/K\right)}{\sqrt{T}} + \beta\left(\frac{\ln\left(y/K\right)}{\sqrt{T}}\right)^2 \tag{17}$$

for coefficients α and β, which depend on the underlying and are fitted empirically, as in Malz (1997). Clearly, (16) is in the format (12).

Consider a portfolio represented by the vector $\mathbf{w}$, whose generic ith entry is the number of contracts in the respective call. The profit and loss then reads:

$$H_{\mathbf{w}} \equiv \sum_{i=1}^{I} w_i\left(P_i\left(X_{T,y}^{(i)}, X_{T,\sigma}^{(i)}; \mathcal{I}_T\right) - p_{i,T}\right) \tag{18}$$

where $p_{i,T}$ denotes the currently traded price of the ith call. In order to calculate the VAR, we need the distribution of the profit and loss (18). To obtain the latter, we need the joint distribution of all the sources of risk $(X_{T,y}^{(i)}, X_{T,\sigma}^{(i)})$ in the portfolio. We realise that the sources of risk are approximately invariants – their joint distribution is independent and identical across time, and thus it does not depend on the specific time. We model this joint distribution with a multivariate normal copula with non-parametric marginals:

$$\begin{pmatrix} X_y^{(1)} \\ X_\sigma^{(1)} \\ \vdots \\ X_y^{(I)} \\ X_\sigma^{(I)} \end{pmatrix} \overset{d}{=} \begin{pmatrix} \hat{F}_{X_y^{(1)}}^{-1}\left(\Phi\left(Z_1^{(1)}\right)\right) \\ \hat{F}_{X_\sigma^{(1)}}^{-1}\left(\Phi\left(Z_2^{(1)}\right)\right) \\ \vdots \\ \hat{F}_{X_y^{(I)}}^{-1}\left(\Phi\left(Z_1^{(I)}\right)\right) \\ \hat{F}_{X_\sigma^{(I)}}^{-1}\left(\Phi\left(Z_2^{(I)}\right)\right) \end{pmatrix} \tag{19}$$

In this expression:

$$\mathbf{Z} \sim N\left(\mathbf{0}, \hat{\mathbf{C}}\right) \tag{20}$$

for a suitably estimated correlation matrix $\hat{\mathbf{C}}$, and $\hat{F}_X$ denotes a suitable estimate of the marginal distribution of X. In particular, we estimate these marginals by a non-parametric kernel smoothing of the historical data. Then we use the cumulative density functions $\hat{F}_X$ to invert (19) for each entry in our joint time series of log price changes $X_{t,y}$ and the log volatility changes $X_{t,\sigma}$:

$$Z_t^{(i)} \equiv \Phi^{-1}\left(\hat{F}_{X^{(i)}}\left(X_t^{(i)}\right)\right), \quad t = 1,\ldots,T \tag{21}$$

where this simplified notation applies to both y and σ. Next, we fit the correlation (20) to the time series (21).

Now we can simulate the profit and loss distribution (18). First, we use our recipe to draw Monte Carlo scenarios from (20) in such a way that the sample mean is zero and the sample covariance exactly matches the estimated correlation matrix $\hat{\mathbf{C}}$. More precisely, we generate a $J \times 2I$ panel $\mathcal{Y}$ of J normal simulations with sample mean a $2I$ vector of zeros and sample covariance the $2I \times 2I$ identity matrix, and map it into a $J \times 2I$ panel $\mathcal{Z}$ distributed as (20) using the Cholesky decomposition of $\hat{\mathbf{C}}$. Then we map those simulations into factor realisations using (19). Since the expression of the inverse cumulative density function is not available analytically, we perform a linear interpolation of the cumulative density function, as in Meucci (2006). Next, those simulations are fed into the pricing functions that appear in (18), thereby generating a $J \times I$ panel $\mathcal{H}$ of joint profit and loss scenarios for the I options at the investment horizon. The portfolio profit and loss (18) is then represented by the simulations vector $\mathcal{H}_{\mathbf{w}} \equiv \mathcal{H}\mathbf{w}$. In figure 2, we report this distribution in an example of a portfolio long-short 20 options.

To calculate the VAR and its contributions from the different securities in the portfolio, first we express the former in terms of the latter:

$$VAR \equiv \sum_{i=1}^{I} w_i \frac{\partial VAR}{\partial w_i} \tag{22}$$

where the partial derivatives that appear in (22) can be expressed conveniently as in Hallerbach (2003) and Gourieroux, Laurent & Scaillet (2000):

$$\frac{\partial VAR}{\partial \mathbf{w}} \equiv \mathbf{p} - \mathbb{E}\left\{\mathbf{P} \middle| H_{\mathbf{w}} \equiv -VAR\right\} \tag{23}$$

where $\mathbf{p}$ denote the current prices, which are known, and $\mathbf{P}$ the horizon prices, which are a random vector, as they appear in (18).

Then the expectations in (23) are approximated numerically as in Mausser (2003) (see also Epperlein & Smillie, 2006, and Meucci *et al*, 2007):

$$\frac{\partial VAR}{\partial \mathbf{w}} \approx -\mathbf{k}'\tilde{\mathcal{H}} \tag{24}$$

In this expression, $\tilde{\mathcal{H}}$ is a $J \times I$ panel whose generic ith column is the ith column of the options profit and loss panel $\mathcal{H}$, sorted as the order statistics of the J-dimensional vector of the portfolio losses $-\mathcal{H}\mathbf{w}$, and $\mathbf{k}$ is a Gaussian smoothing kernel that peaks around the rescaled confidence level cJ. Finally, (22) yields the contributions from each option as well as the total VAR. The total VAR number then follows from (22).

In figure 3, we report the total VAR in our example of a portfolio long-short 20 options, as well as its decomposition in terms of the contributions from each call.

Risk managers can now proceed to stress test the correlation $\hat{\mathbf{C}}$ using the Cholesky decomposition of the new stress-test matrix and the $J \times 2I$ panel $\mathcal{Y}$ of uncorrelated standard normal simulations in the above process. Then they can analyse the impact of the stress test on a risk report such as figure 3, confident that the stress-test assumptions will be faithfully reflected in the simulations. ■

Attilio Meucci is head of research at Bloomberg ALPHA, Portfolio Analytics and Risk, in New York. He is grateful to an anonymous referee for helpful feedback. Email: a.meucci@bloomberg.net

References

Alexander C, D Ledermann and W Ledermann, 2008
ROM simulation: a new approach to simulation using random orthogonal matrices
Working paper

Anderson E, Z Bai, C Bischof, S Blackford, J Demmel, J Dongarra, J Du Croz, A Greenbaum, S Hammarling, A McKenney and D Sorensen, 1999
LAPACK user's guide
SIAM, third edition

Avellaneda M, 1999
Minimum-entropy calibration of asset-pricing models
International Journal of Theoretical and Applied Finance 1, pages 447–472

Boyle P, M Broadie and P Glasserman, 1995
Recent advances in simulation for security pricing
Proceedings of the 1995 Winter Simulation Conference, pages 212–219

Cheng R, 1985
Generation of multivariate normal samples with given sample mean and covariance matrix
Journal of Statistical Computation and Simulation 21, pages 39–49

D'Amico M, G Fusai and A Tagliani, 2003
Valuation of exotic options using moments
Operational Research 2, pages 157–186

Epperlein E and A Smillie, 2006
Cracking VAR with kernels
Risk August, pages 70–74

Glasserman P, 2004
Monte Carlo methods in financial engineering
Springer

Glasserman P and B Yu, 2005
Large sample properties of weighted Monte Carlo estimators
Operations Research 53, pages 298–312

Gourieroux C, J-P Laurent and O Scaillet, 2000
Sensitivity analysis of values at risk
Journal of Empirical Finance 7, pages 225–245

Hallerbach W, 2003
Decomposing portfolio value-at-risk: a general analysis
Journal of Risk 5, pages 1–18

Li K, 1992
Generation of random matrices with orthonormal columns and multivariate normal variates with given sample mean and covariance
Journal of Statistical Computation and Simulation 43, pages 11–18

Malz A, 1997
Option-implied probability distributions and currency excess returns
Federal Reserve Bank of New York – Staff Reports

Mausser H, 2003
Calculating quantile-based risk analytics with L-estimators
Journal of Risk Finance 4(3), pages 61–74

Meucci A, 2005
Risk and asset allocation
Springer

Meucci A, 2006
Beyond Black-Litterman in practice: a five-step recipe to input views on non-normal markets
Risk September, pages 114–119

Meucci A, 2008
Fully flexible views: theory and practice
Risk October, pages 97–102, available at *symmys.com > Research > Working Papers*

Meucci A, Y Gan, A Lazanas and B Phelps, 2007
A portfolio mangers guide to Lehman Brothers tail risk model
Lehman Brothers Publications

Petkov P, N Christov and M Konstantinov, 1991
Computational methods for linear control systems
Prentice Hall

Wedderburn R, 1975
Random rotations and multivariate normal simulation
Unpublished working paper