Monte Carlo within a day

In November, we published a paper showing one route to calculating intra-day value-at-risk. Here, Juan Cárdenas, Emmanuel Fruchard, Jean-François Picron, Cecilia Reyes, Kristen Walters and Weiming Yang present an alternative approach

This article presents an innovative approach to measuring intra-day VAR that combines the use of a robust parametric technique, gamma VAR[1], with Monte Carlo simulation to capitalise on the respective strengths of these models. The simulation is optimised by using parametric VAR results to limit the required number of portfolio revaluations to those random scenarios that are statistically relevant given the Greek-estimated profit and loss distribution, and as a variance reduction tool to minimise the standard Monte Carlo error term. As the results presented here will show, these techniques, combined with portfolio and market risk factor compression, significantly enhance the performance and precision of the Monte Carlo engine. Although VAR alone, no matter how sophisticated the model, is not sufficient to capture effectively all possible market moves, it is an invaluable intra-day tool for risk managers.

Estimating the loss tail of Monte Carlo

A naive or "brute force" Monte Carlo simulation for VAR performs full trade revaluations of an entire portfolio under thousands of random scenarios to approximate the profit and loss distribution. The trade revaluation process is computationally intensive, thereby preventing Monte Carlo from being a viable tool to measure risk during the trading day. Given that the goal in VAR is to obtain the "loss tail" (eg, a 5% loss tail for a 95% confidence level) of the distribution, it is reasonable to perform only full trade revaluations for those scenarios that will result in large losses for a portfolio. Our approach uses the gamma VAR distribution as a tool for determining which scenarios will result in losses in the tail of the Monte Carlo distribution and discarding irrelevant scenarios using the following step-by-step process:

☐ **Parametric VAR calculation.** First, our model calculates the first- and second-order derivatives with respect to all relevant portfolio risk factors, which will result in a vector of deltas and vegas and a matrix of gammas. Next, instead of approximating the profit and loss function (eg, by using partial simulations or its first few moments), the characteristic function is determined analytically and the full profit and loss distribution is recovered using a fast Fourier transform.

☐ **Random rate path generation.** This involves generating pseudo-random scenarios for the portfolio risk factors (interest rates, bond and equity prices, foreign exchange rates, volatilities, etc). Risk factors may either represent individual market rates or statistically based shift scenarios[2] calculated using principal component analysis (PCA). Employing PCA to compress the universe of market risk factors into their salient principal components enhances performance. The definition of the risk factors is a critical input to the analysis. Recognising that during crises historical measures of volatilities and correlations break down, it is very important to supplement VAR with event risk analysis.

☐ **Modelling of error distribution.** For each random scenario, the change in portfolio mark-to-market is approximated using the Greeks. The true profit and loss based on full revaluation is calculated for those scenarios where the Greek-based profit and loss falls below a user-specified upper bound (see figure 1). The specified upper bound will logically reside to the right of the parametric loss tail to reflect possible error in the parametric VAR model. Next, the distribution of the error between the Greek-based and full revaluation for the scenarios whose Greek-based profit and loss falls below the upper bound is modelled.

☐ **Upper bound adjustment.** Based on the calculated error distribution, the upper bound may be adjusted further to prevent any relevant scenarios from being inappropriately discarded in the Monte Carlo process. Figure 1 shows the error distribution and the adjusted upper bound used to determine relevant random scenarios. In addition, a minimum number of tail scenarios may be specified to ensure sufficient sampling in the tail region.

☐ **Monte Carlo VAR calculation.** Finally, the full valuation profits and losses for the scenarios that fall below the adjusted upper bound are ordered and the Monte Carlo VAR is obtained by assessing the α-percentile based on the original number of scenarios generated, where $(1 - \alpha)$ is the confidence level.

Variance reduction

Monte Carlo results explicitly include a calculated error term. Besides determining which randomly generated scenarios should be valued, the Greek approximation of the profit and loss distribution can also be used to reduce the variance of the estimators. Our method uses two variance reduction techniques, control variate and stratified sampling, both of which are based on the gamma VAR distribution, to reduce the Monte Carlo error term. By reducing the error term, risk managers may obtain a higher degree of precision with Monte Carlo based on the same number of random scenarios or improved performance by reducing the number of required scenarios and obtaining the same specified error.

1. Loss tail optimisation

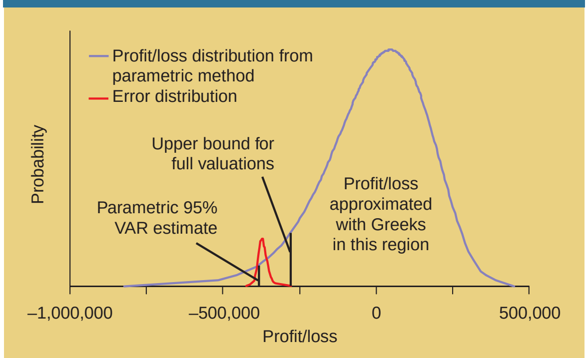

[1] *See* Risk *October 1997, pages 72–75*

[2] *For example, rather than modelling the processes of a set of zero-coupon rates that comprise a yield curve, it is possible to concisely represent its process using statistically relevant shift scenarios without significantly losing information, eg, parallel, steepening and curvature shifts*

2. Control variate case

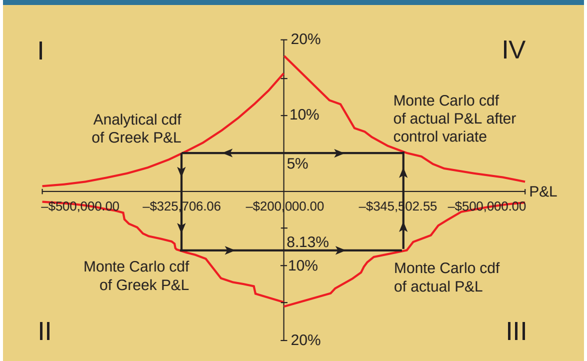

3. Brute force Monte Carlo v. control variate

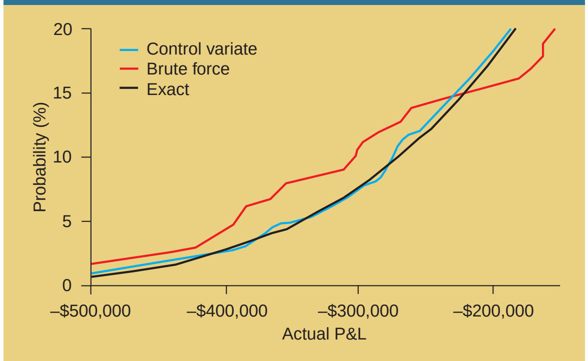

4. Cumulative distribution function using stratified sampling

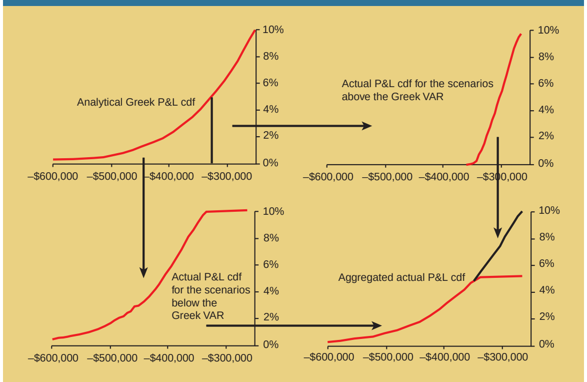

5. Parametric v. Monte Carlo VAR

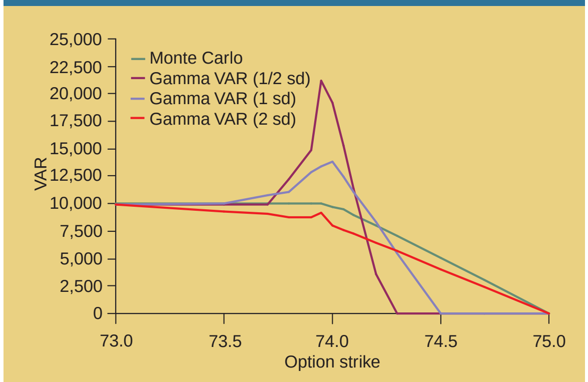

Control variate case

Given that, in most cases, the gamma VAR distribution will be highly correlated with the true distribution being estimated via Monte Carlo simulation, it may be used as a control for the true profit and loss distribution as shown in figure 2. In the first quadrant (I), the Greek VAR is obtained using the analytical approximation, ie, gamma VAR, of the profit and loss distribution. In the second quadrant (II), the Monte Carlo estimate of the Greek VAR's percentile is obtained by performing a partial simulation, ie, calculating the Greek profit and loss for randomly generated scenarios. In the third quadrant (III), the true profit and loss corresponding to the percentile found above is calculated. The actual profit and loss results are subsequently mapped in the fourth quadrant (IV) as corresponding to the VAR percentile. By using a control, the variance of the estimator is reduced by a factor, $2(1 - \rho)$, where ρ is the correlation. For most test cases, using the full gamma matrix for the calculation of the parametric VAR results in a correlation above 90%, making the simulation more than five times faster.

One of the most attractive features of the control variate is that it can be easily applied to the full distribution, yielding a much smoother curve than brute force (ie, full) Monte Carlo. Figure 3 shows a comparison of brute force Monte Carlo and control variate for 100 scenarios and a correlation of 95%. The curve labelled "exact" was obtained by running 50,000 scenarios and using the control variate.

Stratified sampling

Stratified sampling can be viewed as a "divide and conquer" variance reduction technique in which the following steps are performed (see figure 4):

- ☐ The pseudo-randomly generated scenarios are divided into two independent sets: one containing the scenarios that result in Greek-approximated profit and loss falling below the parametric VAR, and the other set containing those scenarios that yield Greek-approximated profit and loss falling above the parametric VAR.
- ☐ Each subset is sampled proportionally to its probability measure. For example, if a 95% VAR is desired, 5% of the scenarios should be below the parametric VAR and 95% should be above it. The true profit and loss distribution is calculated for each set separately using full valuation.
- ☐ Since the two sets are independent by construction, the overall Monte Carlo distribution is simply a weighted sum of the two sub-distributions.

Portfolio mapping

Mapping a portfolio into fewer equivalent trades is an additional feature incorporated in the design of our Monte Carlo engine. Equivalence, in this sense, means that the sensitivities of the mapped trades and the original portfolio to all the risk factors are equal. This technique can yield significant improvements in performance and may be employed for linear (eg, interest rate swaps) as well as non-linear products (eg, plain vanilla caps and floors). A large swap portfolio of 20,000 trades, for example, can be mapped to less than 100 trades.

Parametric VAR in Monte Carlo

Market practitioners frequently cite the need for a full Monte Carlo simulation based on the assumption that parametric VAR models give inaccurate results for short-dated at-the-money options whose Greeks are unstable and hence unreliable to use to predict profit and loss. While this is true when using analytic Greeks, the error is often of a significantly smaller magnitude when the Greeks are calculated empirically using reasonably sized perturbations of the risk factors. This key observation formed the basis of our development of a parametric Monte Carlo approach and is illustrated in figure 5 by a "worst-case" example of a bond option with one day to expiry. For this example, the VAR is calculated for various strikes and with Greeks obtained globally with step sizes equalling ½, one and two standard deviations (ie, 3-, 6- and 12-basis-point shifts, respectively).[3] The results show that the parametric VAR is inaccurate when the shift is small, because the gamma surges to an artificially high level. On the other hand, when the shift size is set to a number of standard deviations consistent with the confidence interval used in the VAR calculation, the gap narrows to a small fraction of the correct result. Therefore, the few close-to-the-money options that exist in every large portfolio will not reduce the effectiveness of the parametric VAR optimisation.

Monte Carlo results

To illustrate the beneficial effects on performance and precision of VAR estimate provided by our Monte Carlo approach, VAR results for a 97.5% confidence level and one-day time horizon are generated for a representative portfolio of a multinational dealer with 500 trades in 24 currencies (see tables A and B). Including numerous currencies and risk factors is relevant for any performance benchmark given that the size of the variance/covariance matrix and corresponding processing time for Monte Carlo expand with the number of relevant risk factors. This portfolio contains 100 risk factors with options comprising 40% of the trades (including forex barriers, Asian options and 100 American bond options).

The fact that this sample portfolio contains exotic options with non-monotonic payouts and exhibits significant convexity is crucial to demonstrating the power of the parametric Monte Carlo method given that, for these instrument types, parametric methods used in isolation are not sufficient VAR estimators. As the tables show, our Monte Carlo method handles these trades (as well as instruments priced with lattices) remarkably well with intermediate events (eg, barrier crossings) incorporated into the modelling process.[4] The same performance is observed after mapping for portfolios of any size (eg, 50,000 swaps) with the same number of options.

VAR was generated on a single Sparc Ultra-10 workstation with a 512-megabyte memory. Performance times could be substantially reduced in a distributed processing environment with multiple servers, which is the desired configuration for generating global VAR. Our Monte Carlo engine was designed based on an object model to support this technology structure.

The exact VAR for the portfolio shown in the tables was calculated by generating 50,000 pseudo-random rate paths with variance reduction; the resulting VAR was 66,075,770. Figure 6 compares the exact VAR results with the gamma VAR estimate as well as our optimised Monte Carlo VAR, with and without employing variance reduction. These results readily demonstrate the improvement in precision of results obtained by generating VAR using our Monte Carlo method.

Conclusion and implementation issues

The results displayed above clearly validate the use of parametric VAR with Monte Carlo simulation and demonstrate the performance gains obtained with the combined approach. The parametric Monte Carlo VAR alone provides a tenfold increase in performance over brute force Monte Carlo. With the variance reduction, an additional five to 20 times boost in performance is obtained. There will be instances (eg, when the Greeks are highly unstable) where the parametric approximation is not reasonable to use together with the Monte Carlo simulation. The ability to configure the VAR calculation process based on portfolio risk profiles and current market conditions is crucial to the design and practical implementation of the parametric Monte Carlo method. With our methodology, risk managers have complete flexibility to determine the extent to which performance optimisations are used when generating VAR results.

Intra-day, risk managers may take advantage of all applicable optimi-

A. Performance run time of Monte Carlo engine

VAR precision	5%	3%
Total number of scenarios	700	1900
Brute force Monte Carlo engine (no optimisation)	1hr, 44 min	4hrs, 31 min
Optimised Monte Carlo VAR	13 min	24 min

This shows the processing time of our Monte Carlo engine for the sample portfolio given differing VAR precisions (and associated required scenarios). Our experience indicates that 5% precision, which is considered more than adequate, will generally result in less than 1,000 random scenarios. As shown, the effect on performance of employing the parametric Monte Carlo VAR technique instead of brute force Monte Carlo is quite dramatic

B. Monte Carlo VAR results with variance reduction

Total number of scenarios	700		1,900	
Results	VAR	Std error	VAR	Std error
No variance reduction	65,122,890	1,543,517	64,931,509	1,221,515
Stratified sampling	66,200,240	695,671	65,754,475	484,224
Gamma VAR control variate	66,371,881	628,933	66,085,661	362,389

This shows the results of VAR and estimated error before and after applying variance reduction techniques. The computation time for variance reduction is negligible compared with total time of Monte Carlo engine and, as shown, the variance reduction is substantial. For most cases, the reductions are more than 55%, which may be interpreted as a fivefold-plus performance advantage after applying the variance reduction technique

6. Error and time v. number of scenarios

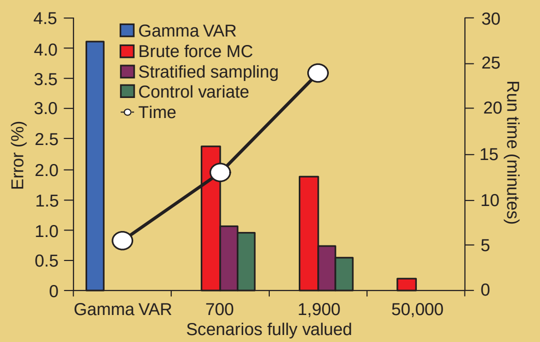

[3] For this example, Greeks were calculated using finite difference methods. Local and global Greek calculations refer to the respective amplitude of the market shift employed to calculate the Greeks. It is well known that using minute "local" shifts in market rates can yield gammas that approach infinity for at-the-money options at expiry. Thus, it is reasonable to increase the magnitude of the shift to reflect expected market moves, which will yield more reliable gammas. In this case, the option mark-to-market was recalculated based on the specified standard deviation change in market rates, ie, a 1/2 standard deviation move, representing a 3-basis-point shift in market rates

[4] As an alternative to ageing barrier trades on several dates from the initial date to the horizon date, a one-step simulation is performed where the simulated mark-to-market value of the barrier option on the final date takes into account the probability that the underlying asset has crossed the barrier over the simulation period for each scenario. Thus correlations across barrier crossing events are taken into consideration

Appendix: variance reduction

Brute force Monte Carlo

For a given level x_a of actual profit and loss, the percentile is estimated using a sum of indicator functions:

$$\overline{F_a}(x_a) = \frac{\sum_{i=1}^{N} I(a_i \le x_a)}{N} \tag{1}$$

where a_i is the actual profit and loss of scenario i. For a sufficiently large number of scenarios, this estimator will be approximately normally distributed with a mean equal to the true percentile and a variance approximately equal to:

$$\overline{\sigma^2}\left(\overline{F_a}(x_a)\right) = \frac{\overline{F_a}(x_a)\left(1 - \overline{F_a}(x_a)\right)}{N-1} \tag{2}$$

Note that to find a confidence interval for the VAR, we need to map the confidence interval on the percentile back to the profit and loss space using the cumulative distribution function.

Control variate

For given levels x_g of Greek profit and loss and x_a of actual profit and loss, the following is an unbiased estimator of the percentile of x_a:

$$\overline{F_{cv}}(x_a) = \overline{F_a}(x_a) - \left(\overline{F_g}(x_g) - F_g(x_g)\right) \tag{3}$$

where a bar denotes a Monte Carlo approximation, as opposed to an analytical result, and the index indicates whether the percentiles and profit and loss refer to the Greek approximation, g, or the actual valuation, a. Since we are looking for the v^{th} percentile, we can replace x_g by VAR_g in equation (3) to obtain:

$$v\% = \overline{F_a}\left(\overline{VAR_a}\right) - \left(\overline{F_g}(VAR_g) - v\%\right) \Rightarrow \overline{VAR_a} = \overline{F_a}^{-1}\left(\overline{F_g}(VAR_g)\right)$$

The standard error on this estimator is simply:

$$\sigma_{cv}^2 = \sigma_a^2 + \sigma_g^2 - \rho\sigma_a\sigma_g \Rightarrow \sigma_{cv}^2 = 2\sigma_a^2(1-\rho)$$

To calculate the above variance, one can simply realise that the estimator (3) is a difference of indicator functions that has an average of zero:

$$\overline{F_{cv}}\left(\overline{VAR_a}\right) = v\% + \frac{\sum_{i=1}^{N}\left(I\left(a_i \le \overline{VAR_a}\right) - I\left(g_i \le VAR_g\right)\right)}{N}$$

$$\overline{\sigma_{cv}^2} = \frac{\sum_{i=1}^{N}\left(I\left(a_i \le \overline{VAR_a}\right) - I\left(g_i \le VAR_g\right)\right)^2}{N-1}$$

$$= \frac{\sum_{i=1}^{N}\left|I\left(a_i \le \overline{VAR_a}\right) - I\left(g_i \le VAR_g\right)\right|}{N-1} \tag{4}$$

Equation (4) indicates that the variance of the estimator is asymptotically equal to the "disagreement probability": the probability that the Greek profit and loss of a random scenario will be above the Greek VAR and its actual profit and loss below the actual VAR or vice versa.

Stratified sampling

For stratified sampling, we have a weighted sum of percentile estimates:

$$\overline{F_{ss}}(x_a) = v\overline{F_1}(x_a) + (1-v)\overline{F_2}(x_a) \tag{5}$$

The indexes refer to the sub-space in which the sampling is done: one for the scenarios that fall below the Greek VAR and two for the scenarios above the Greek VAR. The standard error of the estimator is equal to:

$$\sigma_{ss}^2 = v^2\sigma_1^2 + (1-v)^2\sigma_2^2$$

$$= v^2\frac{\overline{F_1}(x_a)\left(1-\overline{F_1}(x_a)\right)}{N_1 - 1} + (1-v)^2\frac{\overline{F_2}(x_a)\left(1-\overline{F_2}(x_a)\right)}{N_2 - 1}$$

If we use biased estimates (ie, dividing by N instead of N – 1), then the proportional sampling allows us to rewrite the above equation as:

$$\sigma_{ss}^2 = v\frac{\overline{F_1}(x_a)\left(1-\overline{F_1}(x_a)\right)}{N} + (1-v)\frac{\overline{F_2}(x_a)\left(1-\overline{F_2}(x_a)\right)}{N} \tag{6}$$

Comparing equations (5) and (6) with equations (1) and (2), we see that the variance of the stratified sampling estimator is a convex combination of the variance of two brute force Monte Carlo estimators, as illustrated in the figure.

In the figure, the red parabola represents the variance of brute force Monte Carlo as a function of the target percentile, as defined by equation (2). Point 1 represents the variance for a 95% confidence level VAR, while the two points labelled 2 are the two individual terms of the stratified sampling variance. According to equations (5) and (6), the variance will be located on the line joining those two points. The exact stratified sampling variance 3 will be located on that line, right below the brute force Monte Carlo variance, as indicated by equation (5). ■

Stratified sampling variance reduction

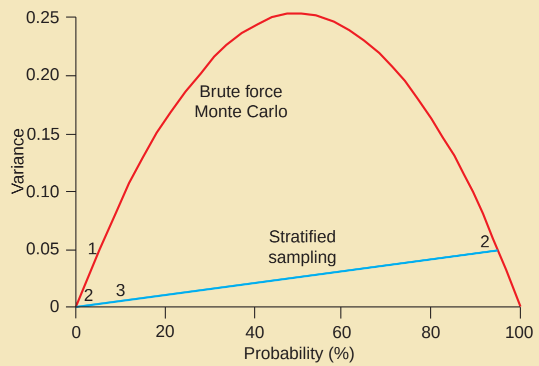

sations to generate VAR in real time and/or determine the "marginal" VAR of certain trades. At close of business day, if the full distribution of profit and loss is required (ie, not just the tail region) the Monte Carlo engine may be run in "brute force" mode involving full revaluations of the portfolio for all randomly generated rate paths. If the full distribution is not required, then the combined parametric method and Monte Carlo provides a very efficient way to estimate VAR. VAR obtained using the optimised Monte Carlo may be back-tested against brute force Monte Carlo, providing risk managers with an empirical view on the accuracy of intra-day VAR measures and the resulting ability to allocate capital and risk limits confidently and optimally during the trading day. ■

Juan Cárdenas is manager of market and credit risk, financial engineering at Summit Systems, Emmanuel Fruchard is director of financial engineering at Summit Systems, Jean-François Picron is a financial engineer at Summit Systems, Cecilia Reyes is director of quantitative risk management at ING Barings, Kristen Walters is associate director of business development at Summit Systems and Weiming Yang is application developer at Summit Systems

e-mail: juan_cardenas@summithq.com
emmanuel_fruchard@summithq.com
jean_francois_picron@summithq.com
Cecilia.Reyes@ing-barings.com
kristen_walters@summithq.com
weiming_yang@summithq.com