

3 VaR Calculations for Derivatives

This section is a brief review of delta and gamma-based VaR calculation methods for options. As we shall see, as a last resort, one can estimate VaR accurately, given enough computing resources, by Monte Carlo simulation, assuming of course that one knows the “correct” behavior of the underlying prices and has accurate derivative-pricing models. In practice, however, brute-force Monte Carlo simulation is not efficient for large portfolios, and for expositional reasons we will therefore take the delta-gamma approach seriously even for a simple option.

We will explore the “delta” and “delta-gamma” approaches for accuracy in plain-vanilla and in our simple jump-diffusion settings. It would be useful to go beyond this with an examination of the accuracy of delta-gamma-based methods with stochastic volatility and skewed return shocks of various sorts.

3.1 The Delta Approach

Suppose $f(y)$ is the price of a derivative at a particular time and at a price level y for the underlying. Assuming that f is differentiable, the *delta* (Δ) of the derivative is the slope $f'(y)$ of the graph of f at y , as depicted in Figure 18 for the case of the Black-Scholes pricing formula f of a European put option.

For small changes in the underlying price, we know from calculus that a reasonably accurate measure of the change in market value of a derivative price is obtained from the usual first-order approximation:

$$f(y + x) = f(y) + f'(y)x + \epsilon(1), \tag{3.1}$$

where $\epsilon(1)$ is the “first-order” approximation error. Thus, for small changes in the index, we could approximate the change in market value of a derivative as that of a fixed position in the underlying whose size is the delta of the derivative.

For spot or forward positions in the underlying, the delta approach is fully accurate, because the associated price function f is linear in the underlying.

The delta approximation illustrated in Figure 18 is the foundation of delta hedging: A position in the underlying asset whose size is *minus* the delta of the derivative is a hedge of changes in price of the derivative, if continually re-set as delta changes, and if the underlying price does not jump.




Figure 18: The Delta (first-order) Approximation

The VaR setting for our application of the delta approach, however, is perverse, for it is actually the *large* changes that are typically of most concern! For a given level of volatility, delta-based approximations are accurate only over short periods of time, and even then are not satisfactory⁴¹ if the underlying index may jump dramatically and unexpectedly. One can see from the convexity of option-pricing functions illustrated in Figure 18 that the delta approach over-estimates the loss on a long option position associated with any change in the underlying price. (If one had sold the option, one would under-estimate losses by the delta approach.)

The delta approach allows us to approximate the VaR of a derivative as the value-at-risk of the underlying multiplied by the delta of the derivative.⁴² Figure 19 shows,

⁴¹See Page and Feng [1995] and Estrella, Hendricks, Kambhu, Shin, and Walter [1994].

⁴²It may be more accurate to expand the first-order approximation at other points than the current price x . We use the forward price of the underlying for these calculations at the value-at-risk time horizon for these calculations, but the difference is negligible.




Figure 19: 2-Week Loss on 20% Out-Of-Money Put (Plain-Vanilla Returns)

as predicted, that the probability density function for the put price⁴³ at a time horizon of 2 weeks, shown as a solid line, has a left tail that is everywhere to the right of the density function for a delta-equivalent position in the underlying. (The option is a European put worth \$100, expiring in one year, and struck 20% out of the money. We use the plain-vanilla model for the underlying, at a volatility of 15%. The short rate and the expected rate of return on the underlying are assumed to be 5.5%.⁴⁴) In particular, the 2-week VaR (at 99% confidence) of the put is \$69.28, but is estimated by the delta approach to have a VaR of \$105.53 (representing a loss of more than the full price of the option, which is possible because the delta-approximating portfolio is a short position in the underlying.) Figure 20 shows the same VaR estimates for a short position in the same put option.

We will discuss below the more accurate “gamma” approach.

⁴³This can be calculated explicitly by the strict monotonicity of the Black-Scholes formula.

⁴⁴The short rate and expected rate of return have negligible effects on the results for this and other examples to follow.




Figure 20: 2-Week Loss on Short 20% Out-Of-Money Put (Plain-Vanilla Returns)

3.2 Impact of Jumps on Value at Risk for Options

Figure 21 illustrates the same calculations shown in Figure 19, with one change: The returns model is a jump-diffusion, with an expected frequency of $\lambda = 2$ jumps per year, and return jumps that have a standard deviation of $\nu = 5\%$. The total annualized volatility of daily returns is kept at $\sigma = 15\%$. The value-at-risk of the put has gone up from \$69.68 to \$74.09. The delta approximation is roughly as poor as it was for the plain-vanilla model. For these calculations, we are using the correct theoretical option-pricing formula,⁴⁵ the correct delta,⁴⁶ and the correct probability distribution for the

⁴⁵One can condition on the number of jumps, compute the variance of the normally distributed total return over one year associated with k jumps, use the Black-Scholes price for this case, weight by the probability of p_k of k jumps, and add up for k ranging from 1 to a point of reasonable accuracy, which is about 10 jumps.

⁴⁶The same trick used for the pricing formula works, as the derivative of a sum is the sum of the derivatives.

underlying price.⁴⁷ (We could also have done these calculations with the Black-Scholes option prices and deltas, which is incorrect. We do not expect a significant impact of this error.)




Figure 21: 2-Week Loss on 20% Out-Of-Money Put (Jump-Diffusion)

3.3 Beyond Delta to Gamma

A common resort when the first-order (that is, “delta”) approximation of a derivative revaluation is not sufficiently accurate is to move on to a second-order approximation. For smooth f , we have

$$f(y + x) = f(y) + f'(y)x + \frac{1}{2}f''(y)x^2 + \epsilon(2), \quad (3.2)$$

where the second-order error $\epsilon(2)$ is smaller, for sufficiently small x , than the first order error, as illustrated by a comparison of Figures 22 and 18.

⁴⁷Again, one conditions on the number of jumps, and adds up the k -conditional densities for the underlying return over a two-week period, and averages these densities with p_k weights. The resulting density is a weighted sum of exponentials of quadratics, which is easy to work with.




Figure 22: Delta-Gamma Hedging, second order approximation

For options, with underlying index y , we say that $f''(y)$ is the gamma (Γ) of the option. In a setting of plain-vanilla returns, both the delta and the gamma of a European option are known explicitly⁴⁸, so it is easy to apply the second-order⁴⁹ approximation (3.2) in order to get more accuracy in measuring risk exposure.

For value-at-risk calculations for the plain-vanilla returns model and plain-vanilla options, gamma methods are “optimistic” for long option positions, because the approximating parabola lies above the Black-Scholes price, as shown in Figure 22. The gamma-based value-at-risk estimate therefore *under-estimates* the actual value-at-risk.⁵⁰ We can see this in the previous two figures. Indeed the gamma-based density approxi-

⁴⁸See, for example, Cox and Rubinstein [1985].

⁴⁹One might think that even higher order accuracy can be achieved, and this is in principle correct. See Estrella, Hendricks, Kambhu, Shin, and Walter [1994]. On the other hand, the approximation error need not go to zero. See Estrella [1994].

⁵⁰This is not just a question of convexity of the option price; it is a third-derivative issue.

mations⁵¹ have a “funny tail,” corresponding to the “turn-back point” of the approximating parabola.

3.4 Gamma-Based Variance Estimates

Based on the gamma approximation, the variance of the revaluation of a derivative whose underlying is $y + X$, where X is the unexpected change, is approximated from (3.2), using the formula for the variance of a sum, by

$$\text{var}[f(y + X)] \simeq V_f(y) \equiv f'(y)^2 \text{var}(X) + \frac{1}{4} f''(y)^2 \text{var}(X^2) + f'(y) f''(y) \text{cov}(X, X^2).$$

For log-normal or normal X , these moments are known explicitly, providing a simple estimate of the risk of a position. This calculation is relatively accurate in the above settings for typical parameters. One may then approximate the value-at-risk at the 99% confidence level as $2.33\sqrt{V_f(y)}$, taking the 0.99 critical value 2.33 for the standard normal density as an estimate of the 0.99 critical value of the normalized density of the actual derivative position. The accuracy of this approximation declines with deviations from the plain-vanilla returns model, with increasing volatility, and with increasing time horizon.

3.5 Delta-Gamma Exposures of Cross-Market Derivatives

Some derivatives are based on more than one underlying. For example, a cross-rate option can be exposed to two currencies simultaneously. The delta approximation of an option exposed to two factors, say marks and yen, is to treat the position as a portfolio of two positions, Δ_i units of marks and Δ_j units of yen, where

$$\Delta_i(y_i, y_j) = \frac{\partial}{\partial y_i} f(y_i, y_j) \simeq \frac{f(y_i + x, y_j) - f(y_i, y_j)}{x},$$

and likewise for $\Delta_j(y_i, y_j)$, where $f(y_i, y_j)$ is the price of the option at the respective underlying indices y_i and y_j for marks and yen, respectively.

For a position or portfolio that is sensitive to two or more underlying indices, such as an option on a spread, in order to estimate risk to second-order accuracy, one could use the deltas and gammas with respect to each underlying. The second-order terms

⁵¹This can be calculated by the same method outlined for the delta case.

would include the “cross-gamma” of a derivative with price $f(y_i, y_j)$ at underlying prices y_i and y_j for markets i and j . The cross-gamma is defined as the derivative

$$\Gamma_{ij} = \frac{\partial^2}{\partial y_i \partial y_j} f(y_i, y_j) \simeq \frac{\Delta_i(y_i, y_j + x) - \Delta_i(y_i, y_j)}{x}.$$

For the case of $i = j$, this is the usual gamma (second derivative) of the position with respect to its underlying index.

3.6 Exposure to Volatility

For derivative positions, one may wish to include the “vega” risk associated with unexpected changes in volatility.⁵² That is, suppose the volatility parameter σ_t changes with a certain volatility of its own. The sensitivity of the option price with respect to the volatility, in the sense of first derivatives, is often called “vega.” If volatility is indeed stochastic, the Black-Scholes formula does not literally apply, although the explicit Black-Scholes vega calculation is a useful approximation of the actual vega over small time horizons.

Figure 23 illustrates the sensitivity of an option to unexpected changes in the volatility of the underlying asset. All else the same, at-the-money options are more sensitive to changes in volatility than are out-of-the-money options. This sensitivity is increasing in the initial level of the underlying volatility. Figure 24 shows, however, that per dollar of initial option premium, the sensitivity in market value to changes in volatility is greater for options that are farther out of the money, and for lower initial volatility. The distinction between the absolute and relative sensitivities of option prices to volatility arises from the fact that an option’s price declines more quickly than does its vega, as the option becomes more and more out-of-the-money and as the volatility parameter is lowered.

For example, suppose the underlying volatility is 10%. A call option struck 20% out of the money with an expiration in 6 months has a market value that almost doubles if the volatility increases unexpectedly from 10% to 11%. For another case, in which the underlying volatility is initially twice as big, the same option increases in market value by roughly 20% under the same circumstances.

Even for a major liquid currency such as the Pound or Mark, the estimated daily

⁵²See Page and Feng [1995].




Figure 23: Sensitivity of Call Price to Volatility

volatility of the volatility can at times exceed⁵³ 100% (annualized), implying non-trivial over-night exposure of an option portfolio to unexpected changes in volatility, particularly for portfolios with a significant fraction of their market value represented by out-of-the-money options on low-volatility underlying assets.

3.7 Numerical Estimation of Delta and Gamma

Other than for simple European options and certain exotics, the deltas and gammas of derivatives are not generally known explicitly. These derivatives can be estimated numerically from derivative-pricing models. For example, we can see in Figure 18 that a reasonable approximation of the delta is obtained by valuing the derivative price $f(y)$ at an underlying price y that is just below the current price, re-valuing the derivative price $f(y+x)$ at a price $y+x$ for the underlying that is just above the current price, and then computing the usual first-difference approximation

$$\Delta \simeq \frac{f(y+x) - f(y)}{x}$$

⁵³See Heynen and Kat [1993] for estimates of the volatility of volatility of certain exchange rates and equity indices. One should of course beware of mis-specification of the stochastic volatility model.




Figure 24: Relative Sensitivity of Call Price to Volatility

of the first derivative of f at $y + x/2$. With more price information, or with an estimate of gamma, one can do better than this simple method. There are related finite-difference-based approximations for gammas.

In order to evaluate the derivative prices $f(y)$ and $f(y + x)$ for this application, one may need to solve a partial differential equation, or to simulate the cash flows on the derivative at initial conditions x and $x + y$ for the underlying. This is quite computationally demanding.

Recent advances⁵⁴ in an area of stochastic calculus called “stochastic flows” allow one to exploit a single simulation of the underlying price process from y , rather than require separate simulation from y and from $y + x$. Using the single simulated path from y , one can estimate the implied path from $y + x$, as illustrated in Figure 25. There are also ways to simulate only the paths that “matter.” For example, with a put, one can condition on the event that the price of the underlying drops. See, for example, Fournie, Lebuchoux, and Touzi [1996]. We expect many new tools to emerge in this direction.⁵⁵

⁵⁴See, for example, Kunita [1990]. For an application to VaR, see Grundy and Wiener [1996].

⁵⁵For a recent example, see Schoenmakers and Heemink [1996], for a method that uses a finite-difference solution as a first step in order to speed up the second-state Monte Carlo simulation through




Figure 25: Using Stochastic Flows to Avoid Re-Simulation

3.8 Impact of Intra-Period Position-Size Volatility

It is conventional to base value-at-risk calculations on the sizes of positions at the beginning of the accounting period. (1-day and 10-day periods are common.) If one knows that the position size is expected to increase or decrease through the period, then one can approximate the effect of changing position size (assuming no correlation between changes in position size and returns) by replacing the initial position size with the square root of the mean squared position size over the accounting period. For example, it is not unusual for broker-dealers in foreign exchange to have dramatic increases in the sizes of their positions during the course of a trading day, and then to dramatically reduce their positions at the end of the trading period in order to mitigate risk over non-trading periods. If not accounted for, this could cause estimated VaR to significantly understate actual profit-and-loss risk.

Let us consider a simple example designed to explore only the effect of random variation in position size around a given mean, without considering the effect of changes in the mean itself. Suppose the underlying asset returns are plain vanilla with constant volatility σ . We suppose that the position size is a classical log-normal process with

a control-variate variance reduction. See Caffisch and Morokoff [1996] for an example of “quasi-random” Monte Carlo methods.

volatility V . We assume that the position size is constant in expectation (a “martingale”), but (in the usual sense of returns) has correlation ρ with the asset. The VaR associated with the stochastic trading strategy is increasing in ρ and, if $\rho > 0$, is increasing in V . If $\rho < 0$, the effect of stochastic position size is ambiguous.

One can show that, over a trading period of length t , the impact of stochastic position size on VaR is to multiply VaR by a factor of approximately⁵⁶

$$q(V, \rho, \sigma, t) = \frac{e^{\alpha t} - 1}{\alpha t},$$

where $\alpha = 2V^2 + 4\rho V\sigma$. The “worst case” is $\rho = 1$. Suppose $\rho = 1$ and $\sigma = 0.50$. If the standard deviation of the daily change in position size is 20 percent of its initial size, we have $V = 0.2 \times \sqrt{365} = 3.8$. For this case, a stochastic position size raises the effective volatility by a factor of approximately

$$q(3.8, 1, 0.5, t) = \frac{e^{36.8t} - 1}{36.8t}.$$

Over 1 day ($t = 1/250$), we have a stochastic-size factor of 1.1, representing roughly a 10 percent higher VaR due to stochastic position size. Over 2 weeks, however, the impact of stochastic position size in this example is a factor of 2.2. In other words, even though the position size is not changing in expected terms, if one were to treat the position size as constant over 2 weeks, the VaR would be low by a factor of 2.2. Even for small asset volatility (any non-zero σ applies) and zero correlation, we get a 2-week “bias factor” of $q(3.8, 0, \sigma, 2/52) = 1.5$ for the same position size volatility of 20 percent per day, as shown Figure 26.

In its recent disclosure documents regarding VaR, Banker’s Trust remarks on the relevance of position size volatility, although no estimates of this effect are reported.

⁵⁶From stochastic calculus, for a trading strategy funded by riskless borrowing or lending, with quantity Q_t of the underlying at time t and an asset with price X_t at time t , we have a variance of gain or loss over the period from time 0 to time t of $E[\int_0^t Q_s^2 \sigma^2 X_s^2 ds]$, neglecting the effect of expected returns on variance, which are truly negligible over typical value-at-risk horizons. Taking X and Q to be log-joint-normal with the indicated parameters leads to the stated result by calculus, ignoring $e^\epsilon - (1 + \epsilon)$ for small ϵ . We take $\sigma^2 t$ to be “small” for this purpose, but not αt . A precise calculation is easy but somewhat messier.




Figure 26: Impact of Position Size Volatility on Value at Risk

4 Portfolio VaR

Using modern portfolio methods, we could imagine a “grand unified” market-risk management model that covers all positions in all markets. In this section, we study the estimation of the VaR for the entire portfolio, accounting for diversification effects, and attempting to deal with the serious computational challenges. We will examine several numerical approaches.

4.1 Risk Factors

The firm’s portfolio of positions has a market value that could be shocked by any of a number of risk factors, such as the S&P 500 index, the 2-year U.S. Treasury rate, the WTI spot Oil price, the German Mark exchange rate, the Nikkei equity index, the 10-year Japanese Government Bond (JGB) rate, and so on. In practice, there could be several hundred, or more, such risk factors that are actually measured. We can label them X_1, \dots, X_n , treating X_i as the “surprise” component of the i -th risk factor, that is, the difference between the i -th risk factor and its expected value.

The covariances of the risk factors are key inputs.⁵⁷ The covariance C_{ij} between risk factors X_i and X_j is $\sigma_i\sigma_j\rho_{ij}$, the product of the standard deviations σ_i of X_i and σ_j of X_j with the correlation ρ_{ij} between X_i and X_j . Historically fitted correlations or standard deviations can be adjusted on the basis of option-implied information, or adjusted arbitrarily⁵⁸ for sensitivity analysis of the effects of changing covariances, as explained in Section 4.3. While the risk factors are often taken to be market rates or prices, there is no reason to exclude other forms of risk, such as certain volatilities, that are not well captured directly by prices or rates.

Suppose, to take the simplest case, that the unexpected change in market value of one's portfolio is

$$Y = \beta_1 X_1 + \beta_2 X_2 + \cdots + \beta_n X_n,$$

where β_i is the direct exposure of the portfolio to risk factor i , which we assume for the moment to be fixed over the VaR time horizon. We measure β_i as the dollar change in the market value of the portfolio in response to a unit change in risk factor i . Then the total risk (standard deviation) D of the portfolio is determined by,

$$D^2 = \sum_{i=1}^n \sum_{j=1}^n \beta_i \beta_j C_{ij}.$$

If X_1, \dots, X_n were to be treated as jointly normally distributed, then the value-at-risk, at the 99-percent level, is simply the 99-percentile change for a normally distributed random variable with standard deviation D , which is approximately $2.326D$. Suppose, for example, that we estimate a standard deviation of $D = 10$ million dollars on a daily basis. Under a normal approximation for the 0.01 critical value, this means a 99-percentile portfolio VaR of 23.26 million dollars. On a weekly basis, this is roughly $23.26 \times \sqrt{5} = 52.01$ million, and on an annual basis, this is roughly $52.01 \times \sqrt{52} = 375$ million dollars. The notion that risks may be re-scaled by the square root of the

⁵⁷For example, J.P. Morgan's RiskMetrics provides much of necessary data that could be used to construct a covariance matrix, allowing daily downloads of historical volatilities and correlations for the major currencies, equity indices, commodity and interest rates.

⁵⁸One cannot adjust covariances arbitrarily, for not every matrix is a legitimate covariance matrix. Indeed, the data provided by RiskMetrics are not literally consistent with a true covariance matrix, and simulating random variables consistent with the reported correlations can only be accomplished after adjusting the correlations. One method that works reasonably well is an eigenvalue-eigenvector decomposition of the covariance matrix, replacement of any negative values (which are in practice small relative to the largest eigenvalues) with zero, and recomposition of the covariance matrix.

time period is reasonable only if there is not significant variation in standard deviations, or correlation in price changes, over the time period in question, as discussed in Section 2, or significant non-linearities in derivative prices as functions of underlying market prices. We will examine the quality of this scaling approximation in Section 4.10.

4.2 Simulating Underlying Risk Factors

The normal distribution is, in practice, only useful as a rough approximation. Fat tails, as explained in Section 2 are common. A suggestion for simulating a fat-tailed distribution is given in Appendix A. While fat tails may be important for exposure to a single risk factor, they may be less critical for a well-diversified portfolio, because of the notion of the central limit theorem, which implies that the sum of a “large” number of independent random variables (of any probability distribution) has a probability distribution that is approximately normal, under technical regularity conditions.⁵⁹ We will soon see the quality of this analytical approximation, based on normal distributions, in an extensive example.

In any case, regardless of the shape of the probability distributions, if one can simulate X_1, \dots, X_n , then one can simulate the total unexpected change in market value, $Y = \beta_1 X_1 + \dots + \beta_n X_n$. One can then estimate the VaR as the level of loss that is exceeded by a given fraction p of simulated outcomes of Y . Appendix B discusses the issue of how many simulated scenarios is a sufficient number for reasonable accuracy. Because of sampling error, one can do better than simply using the 0.01 critical value of the simulated data to estimate the 99% VaR of the underlying distribution. See, for example, Bassi, Embrechts, and Kafetzaki [1996] and Butler and Schachter [1996] for methods that estimate “smooth” tails from sampled data.

⁵⁹See, for example, Durrett [1991]. This is not to suggest that the risk factors X_1, \dots, X_n are themselves independent, but rather that, in some cases, they may be approximately expressed in terms of independent random variables Z_1, \dots, Z_k , for some $k \leq n$. (For VaR calculations, this idea is pursued by Jamshidian and Zhu [1997].) The idea behind principal-component decomposition is an example. The question at hand is then whether the portfolio risk is sufficiently “diversified,” in terms of dependence on Z_1, \dots, Z_k , to take advantage of the principle underlying the central limit theorem.

4.3 Bootstrapped Simulation from Historical Data

In a stationary statistical environment, one can simulate underlying prices in an historically realistic manner by “bootstrapping” from historical data. For example, one can simply take a data-base of actual historical returns, unadjusted, as the source for simulated returns. This will capture the correlations, volatilities, tail fatness and skewness in returns that are actually present in the data, avoiding a need to parameterize and estimate a mathematical model, with the encumbent costs and dangers of misspecification. J.P. Morgan, for example, reports that it uses actual historical price changes to measure its VaR. For reasons of stationarity, use of historical returns is usually preferred to use of historical price changes.

On the other hand, because of significant non-stationarity, at least in terms of volatilities and correlations, one may wish to “update” the historical return distribution. For example, suppose one wishes to update the volatilities. Rather than drawing from the time series R_1, R_2, \dots of historical returns on a given asset, one could draw from the returns $\hat{R}_1, \hat{R}_2, \dots$ defined by

$$\hat{R}_i = R_i \frac{\hat{V}}{V},$$

where V is the historical volatility and \hat{V} is a recent volatility estimate, for example a near-to-expiration option-implied volatility.

Going beyond volatilities, one can update as well for recent correlation estimates. For example, suppose C is the historical covariance matrix for returns across a group of assets of concern, and \hat{C} is an updated estimate. Let R_1, R_2, \dots denote the vectors of historical returns across these markets. The historical return distribution can be updated for volatility and correlation by replacing R_t , for each past date t , with

$$\hat{R}_t = \hat{C}^{1/2} C^{-1/2} R_t, \tag{4.1}$$

where $C^{-1/2}$ is the matrix-square-root of C^{-1} , and likewise for $\hat{C}^{1/2}$. The covariance matrix associated with the modified data $\hat{R}_1, \hat{R}_2, \dots$ is then⁶⁰

$$\hat{C}^{1/2} C^{-1/2} C^{1/2} C^{-1/2} \hat{C}^{1/2} = \hat{C},$$

⁶⁰We use the fact that, for a random vector X with covariance matrix C , and for any compatible matrix A , the covariance of AX is ACA^\top .

as desired. We have not explored the implications of this linear transformation for skewness and tail behavior. As correlation estimates tend to be relatively unstable, any corrections for estimated correlation should be adopted with caution.

4.4 The Portfolio Delta Approach

Unfortunately, the exposure β_i to a given risk factor X_i is typically not constant, as assumed above. For example, if X_i is the unexpected change in the S&P 500 index and one's portfolio includes S&P 500 options, then β_i is not constant because the change in market value of the options is non-linear in X_i , as illustrated in Figure 18.

With certain types of options and other derivatives, for small changes in the underlying, the delta approach is sufficiently accurate in practice, and one could think in terms of the approximation

$$Y \simeq \Delta_1 X_1 + \cdots + \Delta_n X_n,$$

where Δ_i is the delta of the total portfolio with respect to the i -th risk factor.

For a portfolio of k different options or other derivatives on the same underlying index, with individual price functions f_1, \dots, f_k , we can compute the delta of the portfolio from the fact that

$$\frac{d}{dy} [f_1(y) + f_2(y) + \cdots + f_k(y)] = f_1'(y) + f_2'(y) + \cdots + f_k'(y). \quad (4.2)$$

That is, the delta of a sum is the sum of the deltas. One can likewise add in deltas for cross-market derivatives, as discussed in Section 3.5.

4.5 A Working Example

In order to illustrate the implications of various methods for estimating the VaR of derivatives portfolios, we will present an extensive hypothetical example. For this example, there are total of 418 underlying assets, those covered by RiskMetrics on July 29, 1996. A portfolio of plain-vanilla options on these underlying assets was simulated by Monte Carlo, with the following distribution.

- Option Type: Independently, any option is drawn with probabilities 0.5 of being a European call, and 0.5 of being a European put.

- Long or Short: Independently, an option position is long with probability 0.4 and short with probability 0.6. We will also consider a portfolio dominated by long option positions, obtained simply by reversing the signs of the quantities of all options in the portfolio.
- Maturity: Independently, the time to expiration is 1 month with probability 0.4, 3 months with probability 0.3, 6 months with probability 0.2, and 1 year with probability 0.1.
- Moneyness: Independently, a given option has a ratio m of exercise price to forward price that is log-normally distributed with mean 1 and 10% “volatility,” in the sense that $\log(m)$ has standard deviation 0.1.
- Quantity: Independently, the size Q of each option position is log-normally distributed, with $\log(Q)$ standard normal.

The 418 underlying assets can be categorized into four groups: Commodity (CM), Foreign Exchange (FX), Fixed Income (FI), and Equity (EQ). The portfolio that was randomly generated using the above parameters has a total of 10,996 options. Table 1 shows the number of options and underlying instruments in each of the four groups. Within each group, there is an equal number of options on each underlying asset. One can see in Tables 1 and 2 the approximate distribution of value and risk across the four groups. The reported standard deviations and correlations were estimated using delta approximations, and annualized, and were based on daily standard deviations and correlations for the underlying 418 assets that were calculated from RiskMetrics results on July 29, 1996.

Table 1: Descriptive Statistics of the Reference Short Option Portfolio

	CM	FX	FI	EQ	Total
Value (\$)	-4.42	-6.68	-70.68	-18.22	-100.00
Standard Deviation (\$)	7.07	4.37	9.47	9.56	18.63
Number of Instruments	34	22	340	22	418
Number of Options	612	704	7480	2200	10996

Table 2: Approximate Correlation Matrix of the Portfolio Components

	CM	FX	FI	EQ
CM	1.00	0.02	-0.05	0.10
FX	0.02	1.00	0.30	0.30
FI	-0.05	0.30	1.00	0.22
EQ	0.10	0.30	0.22	1.00

4.6 Delta and Gamma “In the Large”

As a preview of the implications of value-at-risk estimation using approximations based on deltas and gammas, we constructed a derivative portfolio on a single hypothetical underlying, consisting of all 10,996 options that were randomly generated using the algorithm described above. From the Black-Scholes formula, we can calculate the value of the portfolio as a function of the underlying asset price. For this, we assumed an underlying annualized volatility parameter of 14.6%, which was the July 29, 1996 RiskMetrics-based volatility estimate for S&P 500. Figure 27 shows a plot of this total value function, as well as its delta and delta-gamma approximations, assuming that the underlying returns are plain vanilla.⁶¹

One sees in Figure 27 the overall concavity of the payoff function of an option portfolio that is predominately short. For such a portfolio, the delta approximation is always above, and the delta-gamma approximation is always below, the actual value of the portfolio. For the version of our portfolio that is predominantly long options, the opposite conclusions apply.

Our VaR results are summarized in Appendix E. We will refer to a subset of them in the following analysis of the quality of various VaR approximations for an option portfolio.

⁶¹We have made the same calculation for the case of a jump-diffusion of the sort considered earlier in Figure 21, using the correct pricing, deltas, and gammas associated with this model. To the eye, the plots for the jump-diffusion case are virtually identical to those shown in Figure 27, and not reported.




Figure 27: Value Approximations for a Random Portfolio of 10,996 options

4.7 The Portfolio Delta-Gamma Approach

For the case of a portfolio exposed to many different underlying assets, one can compute a delta-gamma-based approximation of the market value of the entire book in terms of the deltas and gammas of the book with respect to each underlying asset and each pair of underlying assets, respectively. The (i, j) -gamma of the entire book, for any i and j , is merely the sum of the (i, j) -gammas of all individual positions. (Again, from calculus, the derivative of a sum is equal to the sum of the derivatives.)

Combining all of the within-market and across-market deltas and gammas for all positions (underlying and derivatives), the total change in value of the entire book (neglecting the time value, which is easily included and in any case is negligible for typical portfolios, for value-at-risk calculations), has the delta-gamma approximation:

$$Y(\Delta, \Gamma) = \sum_{j=1}^n \Delta_j X_j + \frac{1}{2} \sum_{j=1}^n \sum_{k=1}^n \Gamma_{jk} X_j X_k, \quad (4.3)$$

where Γ_{ij} is the (i, j) -gamma of the entire book.

From (4.3), we can compute the portfolio variance to second-order accuracy as

$$\begin{aligned} \text{var}(Y(\Delta, \Gamma)) &= \sum_j \sum_k \Delta_j \Delta_k \text{cov}(X_j, X_k) + \sum_i \sum_j \sum_k \Delta_i \Gamma_{jk} \text{cov}(X_i, X_j X_k) \\ &\quad + \frac{1}{4} \sum_i \sum_j \sum_k \sum_\ell \Gamma_{jk} \Gamma_{k\ell} \text{cov}(X_i X_j, X_j X_k). \end{aligned} \quad (4.4)$$

The first term of (4.4) may be recognized as the first-order estimate of the variance of the book discussed in the previous section. The covariance terms involving products of the form $X_i X_j$ can be computed explicitly for the case of normal or our simple jump-diffusion models of the underlying returns associated with X_1, X_2, \dots, X_n .

4.8 Sample VaR Estimates for Derivative Portfolios

For our hypothetical option portfolio, we have computed value-at-risk estimates for all combinations of the following cases:

1. short and long versions of the reference option portfolio.
2. at 1-day and at 2-week horizons.
3. at a range of confidence levels.

4. for plain-vanilla and various types of jump-diffusion return models.

We have results for each case above, for each of the following methods for estimating the VaR:

1. “Actual” – Monte Carlo simulation of all underlying asset prices, and computation of each option price for each scenario by an exact formula. We take 10,000 independent scenarios drawn with Matlab pseudo-random number generators, and use no variance-reduction methods.
2. “Delta” – Monte Carlo simulation of all underlying asset prices, and approximation of each option price for each scenario by a delta-approximation of its change in value.
3. “Gamma” – Monte Carlo simulation of all underlying asset prices, and approximation of each option price for each scenario by the delta-gamma approximation $Y(\Delta, \Gamma)$ of its change in value.
4. “Analytical-Gamma” – The explicit approximation $c(p) \times \sqrt{\text{var}(Y(\Delta, \Gamma))}$, where $c(p)$ is the p -critical value of the standard normal density (for example, 2.33 in the case of a 99% confidence VaR), and where $\text{var}(Y(\Delta, \Gamma))$ is calculated in (4.4).

For cases 1, 2 and 3, we take the p -critical value of the simulated revaluations as our VaR estimates, although kernel or other quantile-estimation methods might be preferred in practice.⁶² The quality of the “analytical-gamma” approximation (4) of VaR for our reference portfolio of options, relative to a reasonably accurate Monte Carlo estimate (1), is illustrated in Figure 28, for the case of a plain-vanilla model for the underlying returns and a 1-day horizon. For example, the actual 99% VaR is approximately 2.2% of the initial market value of the portfolio, while the analytical-gamma VaR approximation is about 2.3%. (See Table 3 of Appendix E.) The quality for a 2-week horizon (7.3% actual VaR versus 9.1% analytical-gamma VaR approximation) is shown in Figure 29, and tabulated in Appendix E. The “simulation-gamma” VaR approximation (3) is reasonably accurate for both the 1-day and 2-week horizons.




Figure 28: Value at Risk of Long Option Portfolio – Plain-Vanilla Model

4.9 Exposure to Correlated Jumps

Figure 30 shows the accuracy of gamma-based approximations for a jump-diffusion setting, calibrated to the same RiskMetrics-based covariance matrix for returns used to generate Figure 28, based on the questionable assumption that the relative sizes of volatilities and the correlations of the returns across markets are not affected by jump events. This example (the third of three jump-diffusion examples summarized in Appendix E) is designed to be extreme, in that half of the variance of the annual return of each asset is associated with the risk of a jump, with an expected arrival rate of 1 jump per year. The results for this extreme jump example provide a more dramatic comparison of the various methods, especially for the 2-week VaR, as shown in Figure 31.

Our results for cases with less extreme jumps, or jumps that are independently timed across markets, summarized in Appendix E, show a distinction from the plain-

⁶²See Bassi, Embrechts, and Kafetzaki [1996] and Butler and Schachter [1996].




Figure 29: Value at Risk of Long Option Portfolio – Plain-Vanilla Model

vanilla case that is more moderate.

4.10 Two-Week VaR by Scaling One-Day VaR

A typical short-cut to estimating the risk of a position over various time horizons is to scale by the square-root of the ratio of the time horizons. For example, a two-week (14-day) standard deviation or VaR can be approximated by scaling up a one-day standard deviation or VaR, respectively, by the factor $\sqrt{14}$. For our sample portfolio setup, this shortcut is actually quite accurate for the plain-vanilla model. The results are summarized in Appendix E. For the unrealistically extreme correlated-jump model described in the previous subsection, the shortcut method is less accurate, as shown in Figure 32 and Appendix E. There are two sources of error in this case, one being the non-linearity of the options; the other being the impact of lengthening the time horizon on the likelihood of a jump within the time horizon.




Figure 30: Value at Risk for Short Option Portfolio – Jump-Diffusion Model 3

4.11 Exposure to Volatility

For option positions, one may wish to include the “vega” risk associated with changing volatilities, as addressed for a single derivative in Section 3.5.⁶³ In principle, that means doubling the number of underlying risk factors, adding one volatility factor for each underlying market. The empirical evidence is that changes in volatility are correlated across markets, and correlated with returns within markets. That means a rather extensive addition to the list of covariances that would be estimated.

In its risk disclosure, Banker’s Trust reports that it measures and accounts for stochastic volatility risk factors in its value-at-risk reports.

4.12 Revaluation of the Book for Each Scenario

Rather than using deltas and gammas to estimate VaR to second order, one could estimate the actual value-at-risk by simulating the market value of the entire book. If

⁶³See also Page and Feng [1995].




Figure 31: Value at Risk for Short Option Portfolio – Jump-Diffusion Model 3

the prices of individual instruments are also computed by Monte Carlo, this can involve a substantial computational burden. For example, suppose that one uses 1000 scenarios to estimate the market value of the book as the expected (risk-neutral) discounted cash flow, at *given* levels of the underlying indices, X_1, \dots, X_n . One must then simulate the underlying indices themselves, say 1000 times, and then value the book for each such simulation! As illustrated in Figure 33, this implies a total of 1,000,000 simulations, each of which may involve many time periods, many indices, and many cash-flow evaluations.

An alternative is to build an approximate pricing formula for each derivative for which there is no explicit formula, such as Black-Scholes, at hand. For VaR purposes, this may be more accurate than relying on the linear or parabolic approximations that come with delta and gamma approximations, especially for certain exotic derivatives. For instance, by Monte Carlo or lattice-based calculations, one can estimate the price of a derivative at each of a small number of underlying prices, and from these fit a spline, or some other low-dimensional analytic approximation, for the derivative




Figure 32: Value at Risk of Long Option Portfolio, Jump-Diffusion Model 3

price. Then, when estimating VaR, one can draw a large number of scenarios for the underlying market price and quickly obtain an approximate revaluation of the derivative in question for each scenario.

As to how many simulations is enough for confidence in the results, a possible approach is described in Appendix B. In general, measuring risk exposure to a large and complex book of derivatives is an extremely challenging computational problem.

4.13 Testing VaR Models

Statistical tests could be applied for the validation of value-at-risk models. For example, if daily marks to market are *iid*, such tests as Kullback discrepancy or Kolmogorov-Smirnov could be used to compare the probability distribution of marks-to-market implied by the VaR model to the historical distribution of marks. The *iid* assumption, however, is unlikely to be reasonable.

A simpler test, which does not require the *iid* assumption for marks to market, is




Figure 33: Many Scenarios for Value-at-Risk Estimation

motivated by the industry practice of verifying the fraction of marks-to-market that actually exceed the VaR. For example, under the null hypothesis that the VaR model is correct, and stationarity, the fraction of days on which the 95-percent VaR for each day is exceeded by the actual mark to market for that day will converge over time to 5 percent. If the fraction differs from 5 percent by a sufficiently large margin in the available sample, one would reject the null hypothesis. The confidence level of this sort of test can be computed under simplifying assumptions. For example, suppose we have 576 days of profit-and-loss data, and only 2.5% of the days had losses in excess of the VaR estimate for that day. (Under the hypothesis that the VaR model is unbiased, the expected fraction is 5%, so perhaps the VaR model is over-estimating risk!) Assuming that the daily excess-of-VaR trials are unbiased and independent, the probability that only 2.5% percent or less of the sampled days would have losses in excess of the 5% VaR is approximately 0.01. (See Appendix B.) At typical statistical confidence levels, we would therefore reject the hypothesis that the VaR model is unbiased.

J.P. Morgan's annual report for 1995 indicates that its 95% daily VaR estimate was exceeded on 4% of the trading days in 1995. (With this, one would fail, on the basis of the test above, to reject the unbiasedness assumption for the VaR model at typical confidence levels.) This simple test may not have as much power as possible to reject

poor VaR models, as it uses relatively little of the available data. More complicated tests could assume that the distribution of marks-to-market is *iid* after some normalization. For example, one could apply a test for matching certain moments of the historical profit-and-loss distribution after daily normalization by current estimates of standard deviations. Confidence levels might be computed by Monte Carlo.

4.14 Volatility as an Alternative to VaR

Volatility is a natural measure of market risk, and one that can be measured and tested with more confidence than can VaR. Tail measures of a distribution, such as VaR, are statistically estimated with large standard errors in this setting, whereas volatility is measured with relatively smaller standard errors.

While it is useful to measure and report volatility regularly, it may be advisable to undertake periodic (or better, randomly timed, given adverse selection for trader behavior) studies of what a given level of volatility means, in terms of the likelihood of losing a level of capital that would cause significant damage to the firm's ability to operate profitably, over a time horizon that reflects the amount of time that would be needed to reduce the firm's balance sheet significantly, or to raise more capital, or both. Is a volatility of 5 percent of the firm's capital "large" or "small" relative to the threshold for distress? Value-at-risk is more to the point on this issue.

Moreover, volatility alone, as a measure of risk used for allocating position limits, would not discourage traders from adopting positions of a given standard deviation that have fat-tailed distributions. Such positions are sometimes called "peso problems" by economists, because of the fat-tailed empirical distribution of Mexican peso returns, for example, as shown for 1986–1996 in Appendix F.

4.15 Marginal Contribution to VaR of a New Position

In the plain-vanilla setting for returns, the marginal contribution to the VaR of the entire book of adding a new position, provided it is not large relative to the book, can be computed to first order with calculus to be approximately ρV , where ρ is the correlation between the new position and the rest of the book, and V is the VaR of the new position on its own. If the position is large relative to the entire book, or if plain-vanilla returns do not apply, or if the time horizon t is long, a more accurate estimate should be considered, and can be done more laboriously.

5 Scenario Exposures

For some applications, we may be concerned, for various risk-management applications, with the expected change in market value of a portfolio to a change in only one of the underlying risk factors or parameters. For example, it may be useful to know the expected change in market value of the portfolio in response to a given change in some market index, yield curve, or volatility. Most banks, for example, estimate the “PV01” of their domestic fixed-income portfolios, meaning the change in market value associated with a 100-basis point parallel shift of the yield curve.

5.1 Cross-Market Exposure Through Correlation

Even if a particular position is not a cross-market derivative, it may have cross-market exposure merely from the correlation of the underlying returns. For example, as the Lira and Mark have correlated returns, we expect a Lira position to be exposed, in expected terms, to a revaluation of the Mark.

Suppose, for example, that X_i is the unexpected change in price of the German mark. We consider the exposure of, say, 4 billion Italian lira options, whose delta with respect to the underlying Lira is 0.5. We let X_j denote the unexpected change in the Lira price. Under normality of X_i and X_j , the expected change in the Lira per unit change in the Mark can be estimated from a “regression” of the form

$$X_j = bX_i + e,$$

where b is the coefficient of the regression⁶⁴ and e is the portion of the change in the Lira that is uncorrelated with (or, equivalently, unexplained by) the change in the Mark. The regression coefficient is $b = C_{ij}/C_{ii} = \sigma_j\rho_{ij}/\sigma_i$. (For 1986-96, for weekly return data, the least-squares estimate of b is 0.84. One would not actually need to use historical regression to estimate b . Rather, one could use option-implied parameters or econometric models for correlations and volatilities.) For motivation only, the idea of estimating b through least-squares regression is illustrated in Figure 34.

For our example, to a first-order approximation, the expected exposure of the Lira

⁶⁴There is no constant “intercept” term in the regression because we are measuring unexpected changes only. Without normality, bX_i is the minimum-variance linear predictor of X_j given X_i , although it need not be the conditional expectation.




Figure 34: Regressing Lira “Surprises” on Mark “Surprises”

option portfolio to the mark is

$$\delta = 0.5 \times 4 \times b \text{ billion,}$$

the delta of the Lira option times the number of options (4 billion) times the regression coefficient b associated with Lira and Mark prices. These are readily estimated coefficients. If, for instance, the expected exposure of the Lira option portfolio to the Mark is $\delta = 0.5$ billion, then we expect to lose approximately 5 million dollars for each 1-penny change in the price of the Mark, in addition to changes in value that are uncorrelated with the Mark. This sort of scenario analysis could be useful in a discussion of the potential loss or gain of our option position in response to a piece of market news specifically regarding the German exchange rate, such as an announcement by the German central bank that is not directly related to the Lira.

In general, a trading firm may wish to estimate exposures to many scenarios. For example, one may wish to have at hand a table of scenario exposures to unit changes in each risk factor, separately. The total scenario exposure to the book X_i , the unexpected change in the Mark price, would be estimated to first order by

$$\delta_i = \sum_{j=1}^n \beta_j \frac{C_{ji}}{C_{ii}},$$

where β_j is the direct exposure to X_j on a delta-equivalent basis. For example, if $\delta_i = 0.3$ billion, then an unexpected change in the Mark price of $X_i = 0.02$ dollars would generate a total expected change in the value of the book of approximately 6 million dollars, plus some “noise” that is uncorrelated with the Mark price. In other words, insofar as the value of the firm depends on the Mark only, one could think of the total book as approximately the same as a pure 0.3 billion spot Mark position. Some of this exposure may actually be pure Mark positions, some of it may be effective Mark exposure held in other positions such as Mark derivatives, German government bonds, German equities, other European equities, and so on.

One can also estimate the portion of total risk of the book, in the sense of standard deviation, that is attributable to a given risk factor. If the volatility of risk-factor i is σ_i , the risk attributable to this risk factor is $\sigma_i\delta_i$. This attributions of risk by factor do not add up to total risk, because of the effects of diversification and correlation.

5.2 Exposure Limits

While it is natural to allocate and measure risk by trading area, there are good reasons to also measure and control total exposure to a market risk factor, including those induced by cross-market correlations. In practice, however, computational limits do not always allow for this “unified” approach, as there may simply be too many risk factors to capture all cross-market effects. Rather, Mark exposures would be measured for only a subset of positions, and indirect Mark exposures would be measured via only a subset of other risk factors.

5.3 Multiple-Factor Scenario Analysis

For purposes of scenario analysis, one may wish to take as the scenario a given change in several risk factors simultaneously. For example, with a fixed-income portfolio, one may be interested in the expected change in market value of the entire book associated with a shift of a given vector of U.S. forward rates, such as given parallel and non-parallel shifts, or some multiples of the first several principle components. This idea is worked out in Appendix C.