

An Overview of Value at Risk

Darrell Duffie and Jun Pan

Preliminary Draft: January 21, 1997

This review¹ of value at risk, or “VaR,” describes some of the basic issues involved in measuring the market risk of a financial firm’s “book,” the list of positions in various instruments that expose the firm to financial risk. While there are many sources of financial risk, we concentrate here on *market risk*, meaning the risk of unexpected changes in prices or rates. Credit risk should be viewed as one component of market risk. We nevertheless focus narrowly here on the market risk associated with changes in the prices or rates of underlying traded instruments over short time horizons. This would include, for example, the risk of changes in the spreads of publicly traded corporate and sovereign bonds, but would not include the risk of default of a counterparty on

¹Darrell Duffie and Jun Pan are at The Graduate School of Business, Stanford University, Stanford CA 94305-5015, USA. Telephone: 415-723-1976. Our email addresses are duffie@baht.stanford.edu and junpan@ecu.stanford.edu, respectively. Please do not copy without permission of the authors. This review owes much in its design and contents to many conversations with Ken Froot of Harvard Business School, Joe Langsam of Morgan Stanley, and Ken Singleton of Stanford University. They should not be held responsible for errors, and our opinions do not necessarily coincide. We are also grateful for the editorial guidance of Stephen Figlewski and financial support from the Financial Research Initiative at the Graduate School of Business, Stanford University. We are grateful for research assistance by Yaroslav Bazaily, Qiang Dai, Eugene Demler, Mark Ferguson, and Tom Varner, and for conversations with Mark Williams and Wei Shi of Bank of America; Jim Cogill, David Dougherty, Anurag Saksena, Dieter Dorp, and Fred Shin of The Royal Bank of Canada; Vince Kaminsky and Matthew Verghese of Enron Capital and Trade Resources; Robert Litterman of Goldman Sachs; Andy Morton and Dexter Senft of Lehman Brothers; Steve Benardete, Robin Brenner, Steve Brauer, Adam Duff, Craig Gustaffson, Joe Langsam, and Patrick De Saint-Aignan of Morgan Stanley and Company; Matt Page of Susquehanna Investment Group; Dan Mudge of Banker’s Trust; Arturo Estrella of The Federal Reserve Bank of New York; and Peter Glynn of Stanford University.

a long-term swap contract. The measurement and management of counterparty default risk involves a range of different modeling issues, and deserves its own treatment.²

Other forms of financial risk include liquidity risk (the risk of unexpectedly large and stressful negative cash flow over a short period) and operational risk, which includes the risk of fraud, trading errors, legal and regulatory risk, and so on. These forms of risk are considered only briefly.

This article is designed to give a fairly broad and accessible overview of VaR. We make no claims of novel research results, and we do not include a comprehensive survey of the available literature on value at risk, which is large and growing quickly.³ While we discuss some of the econometric modeling required to estimate VaR, there is no systematic attempt here to survey the associated empirical evidence.

1 Background

In managing market risk, there are related objectives:

1. Measure the extent of exposure by trade, profit center, and in various aggregates.
2. Charge each position a cost of capital appropriate to its market value and risk.
3. Allocate capital, risk limits, and other scarce resources such as accounting capital to profit centers. (This is almost the same as 2.)
4. Provide information on the firm's financial integrity and risk-management technology to contractual counterparties, regulators, auditors, rating agencies, the financial press, and others whose knowledge might improve the firm's terms of trade, or regulatory treatment and compliance.
5. Evaluate and improve the performance of profit centers, in light of the risks taken to achieve profits.
6. Protect the firm from financial distress costs.

²An example of an approach that measures market risk, including credit risk, is described in Jamshidian and Zhu [1997].

³Other surveys of the topic include Jackson, Maude, and Perraudin [1995], Linsmeier and Pearson [1996], Mori, Ohsawa, and Shimizu [1996], Littlejohn and Fry [1996] and Phelan [1995]. For empirical reviews of VaR models, see Hendricks [1995], Beder [1995], and Marshall and Siegel [1996].

These objectives serve the welfare of stakeholders in the firm, including equity owners, employees, pension-holders, and others.

We envision a financial firm operating as a collection of profit centers, each running its own book of positions in a defined market. These profit centers could be classified, for example, as “equity,” “commodity,” “fixed income,” “foreign exchange,” and so on, and perhaps further broken down within each of these groups. Of course, a single position can expose the firm to risks associated with several of these markets simultaneously. Correlations among risks argue for a unified perspective. On the other hand, the needs to assign narrow trading responsibilities and to measure performance and profitability by area of responsibility suggest some form of classification and risk analysis for each position. We will be reviewing methods to accomplish these tasks.⁴

Recent proposals for the disclosure of financial risk call for firm-wide measures of risk. A standard benchmark is the *value at risk* (“VaR”). For a given time horizon t and confidence level p , the value at risk is the loss in market value over the time horizon t that is exceeded with probability $1 - p$. For example, the Derivatives Policy Group⁵ has proposed a standard for over-the-counter derivatives broker-dealer reports to the Securities and Exchange Commission that would set a time horizon t of two weeks and a confidence level p of 99 percent, as illustrated in Figure 1. Statistically speaking, this value-at-risk measure is the “0.01 critical value” of the probability distribution of changes in market value. The Bank for International Settlements (BIS) has set p to 99 percent and t to 10 days for purposes of measuring the adequacy⁶ of bank capital, although⁷ BIS would allow limited use of the benefits of statistical diversification across

⁴Models of risk-management decision making for financial firms can be found in Froot and Stein [1995] and Merton and Perold [1993]. The Global Derivatives Study Group, G30 [1993] reviews practices and procedures, and provides a follow up survey of industry practice in Group of Thirty [1994].

⁵See Derivatives Policy Group [1995].

⁶For more on capital adequacy and VaR, see Dimson [1995], Jackson, Maude, and Perraudin [1995], and Kupiec and O’Brien [1995].

⁷See the December 12, 1996 communiqué of the Bank for International Settlements, “announcing an amendment to the Basle Committee on Banking Supervision,” from *BIS Review*, Number 209, December 12, 1995, Basle, Switzerland. See also the draft ISDA response to the Basle market risk proposal made in April, 1995, in a memo from Susan Hinko of ISDA to the Basle Market Risk Task Force, July 14, 1995. The ISDA response proposes to allow more flexibility in terms of measurement, but require that firms disclose a comparison between the value-at-risk estimated at the beginning of each period, and the ultimately realized marks to market. This would presumably lead to some discipline regarding choice of methodology. Incidentally, VaR is not the difference between the *expected*

different positions, and factors up the estimated 0.01 critical value by a multiple of 3. Many firms use an overnight value-at-risk measure for internal purposes, as opposed to the two-week standard that is commonly requested for disclosure to regulators, and the 99-percent confidence level is far from uniformly adopted. For example, J.P. Morgan discloses its daily VaR at the 95-percent level. Bankers Trust discloses its daily VaR at the 99-percent level.

One expects, in a stationary environment for risk, that a 99-percent 2-week value-at-risk is a 2-week loss that will be exceeded roughly once every four years. Clearly, then, given the over-riding goal of protecting the franchise value of the firm, one should not treat one's measure of value-at-risk, even if accurate, as the level of capital necessary to sustain the firm's risk. Value at risk is merely a benchmark for *relative* judgements, such as the risk of one desk relative to another, the risk of one portfolio relative to another, the relative impact on risk of a given trade, the modeled risk relative to the historical experience of marks to market, the risk of one volatility environment relative to another, and so on. Even if accurate, comparisons such as these are specific to the time horizon and the confidence level associated with the value-at-risk standard chosen.

Whether the VaR of a firm's portfolio of positions is a relevant measure of the risk of financial distress over a short time period depends in part on the liquidity of the portfolio of positions, and the risk of adverse extreme net cash outflows, or of severe disruptions in market liquidity. In such adverse scenarios, the firm may suffer costs that include margins on unanticipated short-term financing, opportunity costs of forgone "profitable" trades, forced balance-sheet reductions, and the market-impact costs of initiating trades at highly unfavorable spreads. Whether the net effect actually threatens the ability of the firm to continue to operate profitably depends in part on the firm's net capital. Value at risk, coupled with some measure of cash-flow at risk,⁸ is relevant in this setting because it measures the extent of potential forced reductions of the firm's capital over short time periods, at some confidence level. Clearly, however, VaR captures only one aspect of market risk, and is too narrowly defined to be used on its own as a sufficient measure of capital adequacy.

In order to measure VaR, one relies on

value and the 0.01-critical value, but rather the difference between the *current* portfolio value and the 0.01 critical value at the specified time horizon. To the error tolerance of current modeling techniques, and for short time horizons, there is not much difference in practice.

⁸By "cash-flow at risk" we mean a "worst-case", say 0.99 critical value, of net "cash" outflow over the relevant time horizon.




Figure 1: Value at Risk (DPG Standard)

1. a model of random changes in the prices of the underlying instruments (equity indices, interest rates, foreign exchange rates, and so on).
2. a model for computing the sensitivity of the prices of derivatives to the underlying prices.

In principle, key elements of these two basic sets of models are typically already in place for the purposes of pricing and hedging derivatives. One approach to market risk measurement is to integrate these models across the trading desks, and add the additional elements necessary for measuring risks of various kinds. Given the difficulty of integrating systems from diverse trading environments, however, a more common approach is a unified and independent risk-management system. In any case, the challenges are many, and include data, theoretical and empirical models, and computational methods.

The next section presents models for price risk in the underlying markets. The measurement of market risk for derivatives and derivative portfolios are then treated

in Sections 3 through 5.

As motivation of the remainder, the reader should think in terms of the following broadly defined recipe for estimating VaR:

1. Build a model for simulating changes in prices across all underlying markets, and perhaps changes in volatilities as well, over the VaR time horizon. The model could be a parameterized statistical model, for example a jump-diffusion model based on given parameters for volatilities, correlations, and tail-fatness parameters such as kurtosis. Alternatively, the model could be a “bootstrap” of historical returns, perhaps “refreshed” by recent volatility estimates.
2. Build a data-base of portfolio positions, including the contractual definitions of each derivative. Estimate the size of the “current” position in each instrument (and perhaps a model for changes in position size over the VaR time horizon, as considered in Section 2).
3. Develop a model for the revaluation of each derivative for given changes in the underlying market prices (and volatilities). On a derivative-by-derivative basis, the revaluation model could be an explicit pricing formula, a delta-based (first-order linear) approximation, a second-order (delta-and-gamma based) approximation, or an analytical approximation of a pricing formula that is “splined” for VaR purposes from several numerically-computed prices.
4. Simulate the change in market value of the portfolio, for each scenario of the underlying market returns. Independently generate a sufficient number of scenarios to estimate the desired critical values of the profit-and-loss distribution with the desired level of accuracy.

We will also consider, in Section 4, the accuracy of shortcut VaR approximation methods based on multiplication of an analytically estimated portfolio standard deviation by some scaling factor (such as 2.33 for the 0.01 critical value under an assumption of normality).

2 Price Risk

This section reviews basic models of underlying price risk. Key issues are “fat tails” and the behavior and estimation of volatilities and correlations.

2.1 The Basic Model of Return Risk

We begin by modeling the daily returns R_1, R_2, \dots on some underlying asset, say on a continuously-compounding basis. We can always write

$$R_{t+1} = \mu_t + \sigma_t \epsilon_{t+1}, \quad (2.1)$$

where

μ_t is the expectation of the return R_{t+1} , conditional on the information available at day t . (In some cases, we measure instead the “excess” expected return, that is, the extent to which the expected return exceeds the overnight borrowing rate.)

σ_t is the standard deviation of R_{t+1} , conditional on the information available at time t .

ϵ_{t+1} is a “shock” with a conditional mean of zero and a conditional standard deviation of one.

The *volatility* of the asset is the annualized standard deviation of return. The volatility at day t is therefore $\sqrt{n} \sigma_t$, where n is the number of trading days per year. (In general, the annualized volatility over a period of T days is $\sqrt{n/T}$ times the standard deviation of the total return $R_{t+1} + \dots + R_{t+T}$ over the T -day period.) “Stochastic volatility” simply means randomly changing volatility. Models for stochastic volatility are considered below.

One sometimes assumes that the shocks $\epsilon_1, \epsilon_2, \dots$ are statistically independent and have the same probability distribution, denoted “*iid*”, but both of these assumptions are questionable for most major markets.

A *plain-vanilla* model of returns is one in which μ and σ are constant parameters, and in which the shocks are “white noise,” that is, *iid* and normally distributed. This is the standard benchmark model from which we will consider deviations.

2.2 Risk-Neutral Versus Actual Value at Risk

Derivative pricing models are based on the idea that there is a way to simulate returns so that the price of a security is the expected discounted cash flow paid by the security. This distorted price behavior is called “risk-neutral.” The fact that this risk-neutral

pricing approach is consistent with efficient capital markets⁹ does not mean that investors are risk-neutral. Indeed the actual risk represented by a position typically differs from that represented in risk-neutral models.

For purposes of measuring value-at-risk at short time horizons such as a few days or weeks, however, the distinction between risk-neutral and actual price behavior turns out to be negligible for most markets. (The exceptions are markets with extremely volatile returns or severe price jumps.) This means that one can draw a significant amount of information for risk-measurement purposes from one's derivative pricing models, *provided they are correct*. Because this proviso is such a significant one, many firms do not in fact draw much risk-measurement information about the price behavior of underlying markets from their risk-neutral derivative pricing models. Rather, it is not unusual to rely on historical price data, perhaps filtered by some sort of statistical procedure. Option-implied volatilities are sometimes used to replace historical volatilities, but the goal of standard risk-measurement procedures that are independent of the influence (benign or otherwise) of the current thinking of option traders has sometimes ruled out heavy reliance on derivative-implied parameters. We shall have more to say about option-implied volatility later in this section.

The distinction between risk-neutral and actual price behavior becomes increasingly important over longer and longer time horizons. This can be important for measuring the credit exposure to default by a counterparty. One is interested in the actual, not risk-neutral, probability distribution of the market value of the position with the counterparty. For that reason alone, if not also for measuring the exposure of the firm to long-term proprietary investments, it may be valuable to have models of price risk that are not derived solely from the risk-neutral pricing of derivatives.

2.3 Fat Tails

Figure 2 shows the probability densities of two alternative shocks. The “thinner tailed” of the two is that of a normally distributed random variable. Even though the fatter tailed shock is calibrated to the same standard deviation, it implies a larger overnight VaR at high confidence levels. A standard measure of tail-fatness is kurtosis, which is $E(S_t^4)$, the expected fourth power of the shock. That means that kurtosis estimates are highly sensitive to extremely large returns! For example, while the kurtosis of

⁹See Harrison and Kreps [1979].

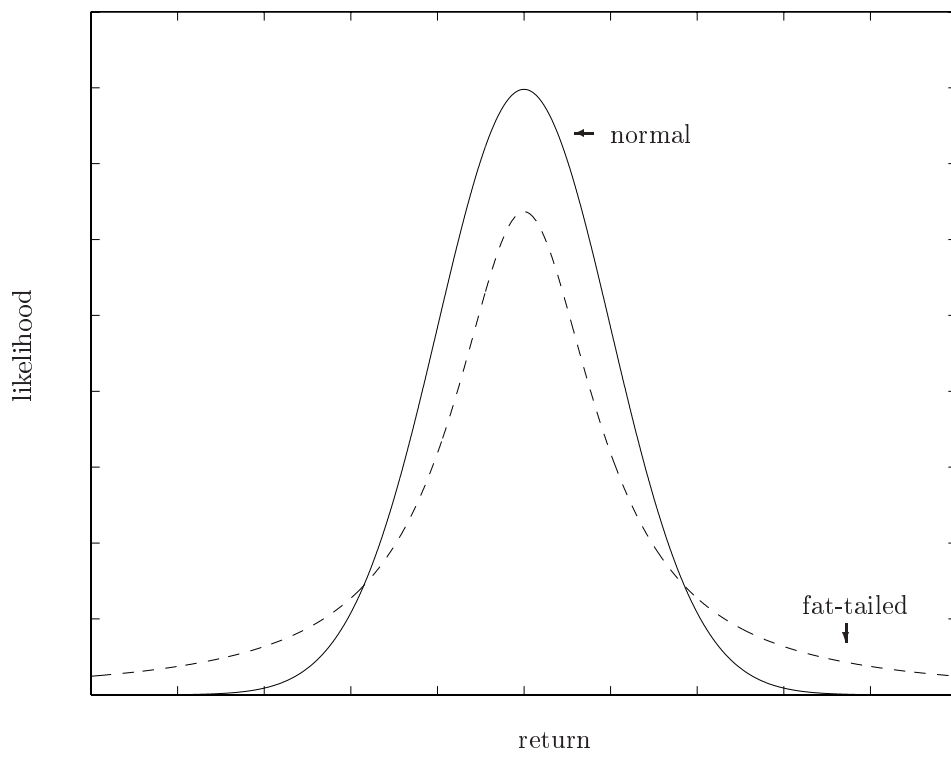


Figure 2: Normal and Fat Tails

a normally distributed shock is 3, S&P 500 daily returns for 1986 to 1996 have an extremely high sample kurtosis of 111, in large measure due to the exceptional returns associated with the market “crash” of October, 1987. The “Black-Monday” return of this crash represents a move of roughly 20 to 25 standard deviations, relative to conventional measures of volatility just prior to the crash!


If one is concerned exclusively with measuring the VaR of direct exposures to the underlying market (as opposed to certain non-linear option exposures), then a more pertinent measure of tail fatness is the number of standard deviations represented by the associated critical values of the return distribution. For example, the 0.01 critical value of the standard normal is approximately 2.33 standard deviations from the mean. By this measure, S-and-P 500 returns are not particularly fat-tailed at the 0.01 level. The 0.01 critical value for S-and-P 500 historical returns for 1986-96 is approximately 2.49 standard deviations from the mean. The 0.99 “right-tail” critical value, which is the relevant statistic for the value at risk of short positions, is only 2.25 standard deviations from the mean. As shown in Figure 3, the 0.05 and 0.95 critical values of S&P 500 returns are in fact closer to their means than would be suggested by the normal distribution. One can also see that S&P 500 returns have negative skewness, meaning roughly that large negative returns are more common than large positive returns.¹⁰

Appendix F provides, for comparison, sample statistics such as kurtosis and tail critical values for returns in a selection of equity, foreign exchanges, and commodity markets. For many markets, return shocks have fatter than normal tails, measured either by kurtosis or tail critical values at typical confidence levels. Figures 4 and 5 show that many typical underlying returns have fat tails, both right and left, at both daily and monthly time horizons. For the markets included¹¹ in Figures 4 and 5, left tails are typically fatter at the 99% confidence level, showing a predominance of negative skewness (especially for equities).

Fat tails can arise through different kinds of models, many of which can be explained

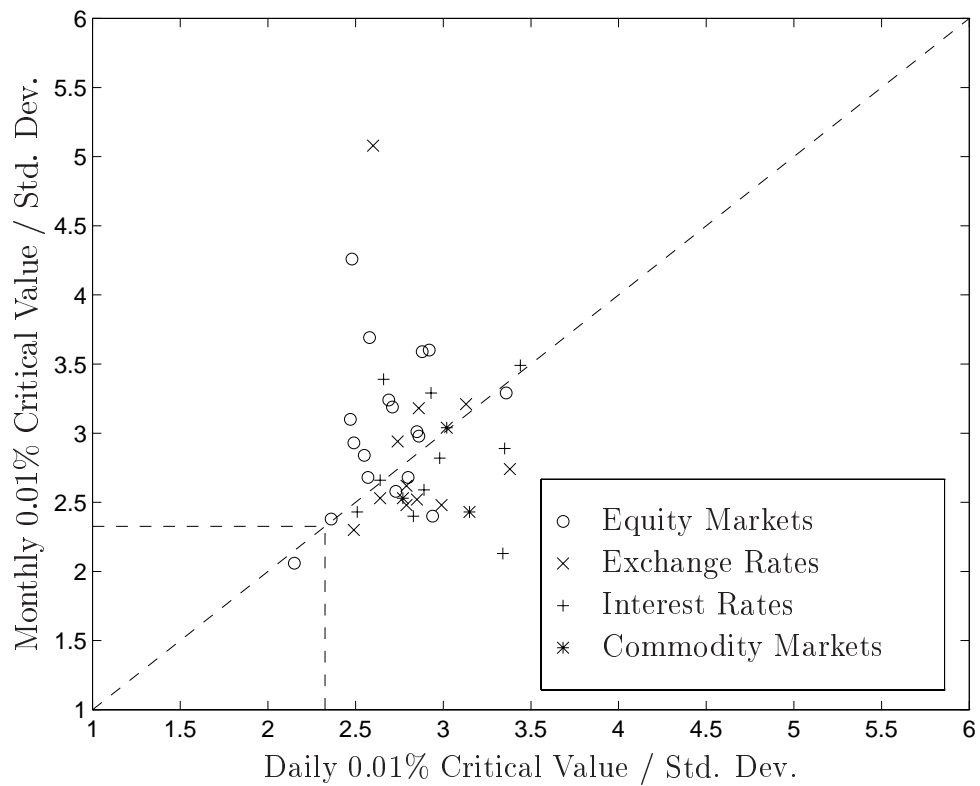
¹⁰Skewness is the expected third power of shocks.

¹¹The markets shown are those for equities, foreign currencies, and commodities shown in the table of sample return statistics in Appendix F, as well as a selection of interest rates made up of: US 3-month LIBOR, US 2-year Treasury, US 30-year Treasury, UK 3-month Bank Bills, UK overnight discount, German Mark 3-month rate, German Mark 5-year rate, French Franc 1-month rate, Swedish discount rate, Yen 1-month rate, and Yen 1-year rate. Changes in log rates are used as a proxy for returns, which is not unreasonable for short time periods provided there are not jumps.



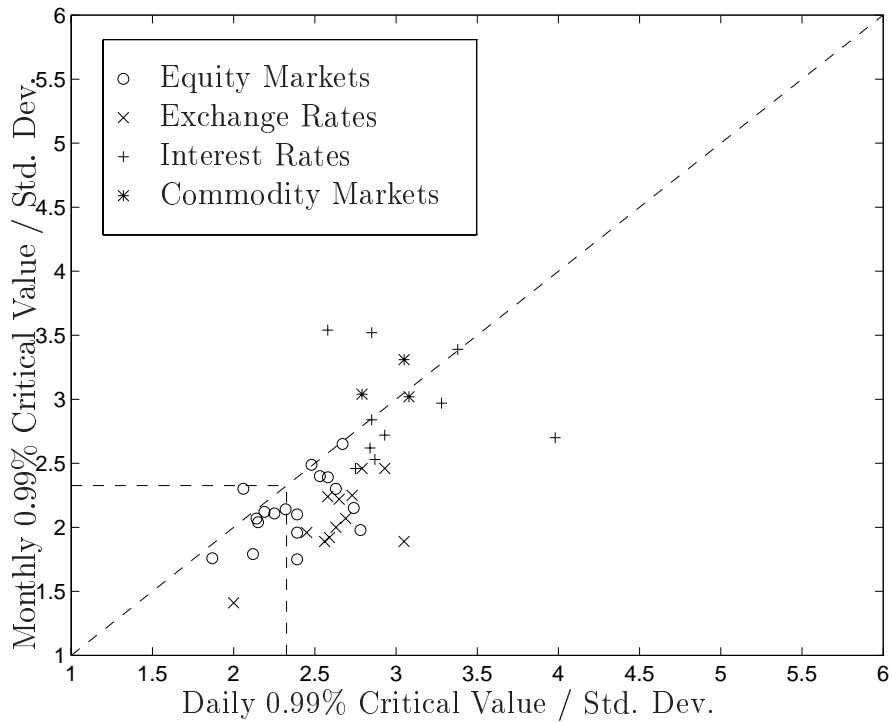
Plotted are the standard normal density (dashed line) and a frequency plot (as smoothed by the default spline supplied with Excel 3.0) of S&P 500 daily returns divided by the sample standard deviation of daily returns, for 1986-96.

Figure 3: Historical Distribution of S-and-P 500 Return Shocks



The statistics shown are the critical values of the sample distribution of daily and monthly returns, divided by the corresponding sample standard deviation.

Figure 4: Left Tail Fatness of Selected Instruments



The statistics shown are the critical values of the sample distribution of daily and monthly returns, divided by the corresponding sample standard deviation.

Figure 5: Right Tail Fatness of Selected Instruments

with the notion of “mixtures of normals.” The idea is that if one draws at random the variance that will be used to generate normal returns, then the overall result is fat tails.¹² For example, the fat-tailed density plotted in Figure 2 is that of a t -distribution, which is a mixture of normals in which the standard deviation of the normal is drawn at random from the inverted gamma-2 distribution.

While there are many possible theoretical sources of fat tails, we will be emphasizing two in particular: “jumps,” meaning significant unexpected discontinuous changes in prices, and “stochastic volatility,” meaning volatility that changes at random over time, usually with some persistence.

2.4 Jump-Diffusions

A recipe for drawing fat-tailed returns by mixing two normals is given in Appendix A. This recipe is consistent (for short time periods) with the so-called *jump-diffusion* model, whose impact on value-at-risk measurement is illustrated in Figure 6, which shows plots of the left tails of density functions for the price in two weeks of \$100 in current market value of the underlying asset, for two alternative models of price risk. Both models have *iid* shocks, a constant mean return, and a constant volatility σ of 15%. One of the models is plain vanilla (normal shocks). The price of the underlying asset therefore has a log-normal density, whose left tail is plotted in Figure 6. The other model is a jump-diffusion, which differs from the plain-vanilla model only in the distribution of shocks. For the jump-diffusion model, with an expected frequency of once per year, the daily return shock is “jumped” by adding an independent normal random variable with a standard deviation of $\nu = 10\%$. The jump arrivals process is a classical “Poisson,” independent of past shocks. The jump standard deviation of 10% is equivalent in risk to that of a plain-vanilla daily return with an annual volatility of 158%. Because the plain-vanilla and jump-diffusion models are calibrated to have the same annual volatility, and because of the relatively low expected frequency of jumps, the two models are associated with roughly the same 2-week 99% value-at-risk measures. The jump-diffusion VaR is slightly larger, at \$6.61, than the plain-vanilla VaR of \$6.45. The major implication of the jump-diffusion model for extreme loss shows up much farther out in the tail. For the jump-diffusion setting illustrated in Figure 6, one can calculate that with an expected frequency λ of roughly once every 140 years, one

¹²For early models of this, see Clark [1973].

will lose overnight at least one quarter of the value of one's position. In the comparison plain-vanilla model, one would expect to wait far longer than the age of the universe to lose as much as one quarter of the value of one's position overnight.¹³ Appendix F shows that there have been numerous daily returns during 1986-1996, across many markets, of at least 5 standard deviations in size. Under the plain-vanilla model, a 5-standard-deviation return is expected less than once per million days. Even 10-standard-deviation moves have occurred in several markets during this 10-year period, but are expected in the plain-vanilla-model less than once every 10^{23} days!

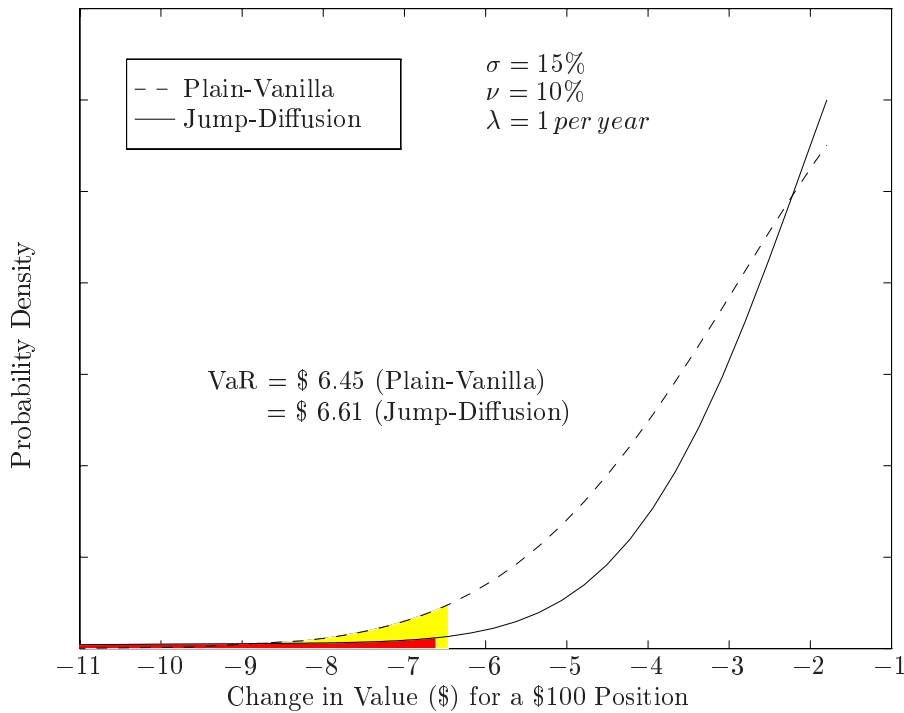


Figure 6: 2-Week 99%-VaR for Underlying Asset

Figure 7 compares the same plain-vanilla model to a jump-diffusion with 2 jumps per year, with each jump having a standard deviation of 5 percent. Again, the plain vanilla and jump-diffusion models are calibrated to the same volatility. While the 99% 2-week VaR for the underlying asset is about the same in the plain-vanilla and jump-diffusion models, the difference is somewhat larger than that shown in Figure 6. The

¹³The expected frequency of an overnight loss of this magnitude in the plain-vanilla model was verbally related to us by Mark Rubinstein.

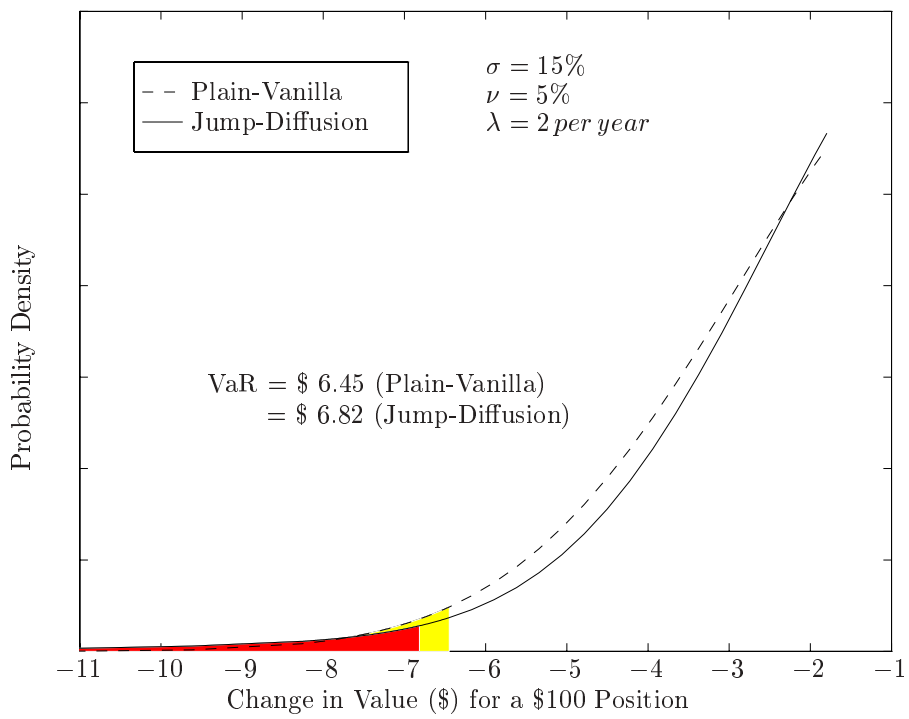


Figure 7: 2-Week 99%-VaR for Underlying Asset

implications of the jump-diffusion model for the value at risk of option positions can be more dramatic, as we shall see in Section 3.

2.5 Stochastic Volatility

The second major source of fat tails is stochastic volatility, meaning that the volatility level σ_t changes over time at random, with persistence. By *persistence*, we mean that relatively high recent volatility implies a relatively high forecast of volatility in the near future. Likewise, with persistence, recent low volatility is associated with a prediction of lower volatility in the near future. One can think of the jump-diffusion model described above as approximated in a discrete-time setting by an extreme version of a stochastic volatility model in which the volatility is random, but with no persistence; that is, each day’s volatility is drawn at random independently of the last, as in the example described in Appendix A.

Even if returns are actually drawn each day with thin tails, say normally distributed, given knowledge of that day’s volatility, we would expect to see fat tails in a frequency plot of un-normalized daily returns, because returns for different days are generated with different volatilities, the usual “mixing-of-normals” story. If this were indeed the cause of the fat tails that we see in Figures 4 and 5, we would expect to see the tail fatness in those plots to be reduced if we normalized each day’s return by an estimate of the level of the volatility σ_t for that day.

The effect of stochastic volatility on left tail fatness and negative skewness could be magnified over time by negative correlation between returns and changes in volatility, which is apparent, for example, in certain¹⁴ equity markets.

We will devote some attention to stochastic volatility models, not only because of the issue of fat tails, but also in order to address the estimation of current volatility, a key input to VaR models.

While one can envision a model for stochastic volatility in which the current level of volatility depends in a non-trivial way on the entire path that volatility has taken in the past, we will illustrate only Markovian stochastic volatility models, those of the form:

$$\sigma_t = F(\sigma_{t-1}, z_t, t), \tag{2.2}$$

¹⁴For the empirical evidence in equity markets of stochastic volatility and correlation of volatility and returns, see for example Bekaert and Wu [1997].

where F is some function in three variables and z_1, z_2, \dots is white noise. The term “Markovian” means that the probability distribution of the next period’s level of volatility depends only on the current level of volatility, and not otherwise on the path taken by volatility. This form of volatility also rules out, for reasons of simplification, dependence of the distribution of changes in volatility on other possible state variables, such as volatility in related markets and macro-economic factors, which one might actually wish to include in practice.

In principle, we would allow correlation between the volatility shock z_t and the return shock ϵ_t of (2.1), and this has important implications for risk management. For example, negative correlation implies negative skewness in the distribution of returns. So that the VaR of a long position could be more than the VaR of a short position of equal size.

There are several basic classes of the Markovian stochastic volatility model (2.2). Each of these classes has its own advantages, in terms of both empirical reasonability and tractability in an option-pricing framework. The latter is particularly important, since option valuation models may, under certain conditions, provide volatility estimates implicitly, as in the Black-Scholes setting. We will next consider some relatively simple examples.

2.5.1 Regime-Switching Volatility

A “regime-switching” model is one in which volatility behaves according to a finite-state Markov chain. For example, if one takes two possible levels, v_a and v_b , for volatility in a given period, we can take the transition probabilities of σ_t between v_a and v_b to be given by a matrix

$$\Pi = \begin{pmatrix} \Pi_{aa} & \Pi_{ab} \\ \Pi_{ba} & \Pi_{bb} \end{pmatrix}.$$

For example, if $\sigma_t = v_a$, then the conditional probability¹⁵ that $\sigma_{t+1} = v_b$ is Π_{ab} . An example, with parameters estimated¹⁶ from oil prices, is illustrated in Figure 8. One may want to allow for more than 2 states in practice. The diagonal probabilities

¹⁵This fits into our general Markovian template (2.2) by taking $F(v_a, z, t) = v_a$ for all $z \leq z_a^*$, where z_a^* is chosen so that the probability that $z_t \leq z_a^*$ is Π_{aa} , by taking $F(v_a, z, t) = v_b$ whenever $z > z_a^*$, and likewise for $F(v_b, z, t)$.

¹⁶This and the other energy volatility estimates reported below are from Duffie and Gray [1995]. For more extensive treatment of regime-switching models of volatility, see Gray [1993] and Hamilton [1990].

Volatility of Oil




Figure 8: Regime-Switching Volatility Estimates for Light Crude Oil

Π_{aa} and Π_{bb} of the regime-switching model can be treated as measures of volatility persistence.

2.5.2 Auto-Regressive Volatility

A standard Markovian model of stochastic volatility is given by the *log-auto-regressive* model:

$$\log \sigma_t^2 = \alpha + \gamma \log \sigma_{t-1}^2 + \kappa z_t, \quad (2.3)$$

where α , γ , and κ are constants.¹⁷ Volatility persistence is captured by the coefficient γ . A value of γ near zero implies low persistence, while a value near 1 implies high persistence. We always assume that $-1 < \gamma < 1$, for otherwise volatility is “explosive.”

The term structure of volatility is the schedule of annualized volatility of return, by the time-horizon over which the return is calculated. For the stochastic volatility model (2.2), in the case of independent shocks to returns and volatility,¹⁸ the term structure of conditional volatility is

$$\begin{aligned}\bar{v}_{t,T} &\equiv \sqrt{\frac{\text{var}_t(R_{t+1} + \dots + R_T)}{T-t}} \\ &= \sqrt{\frac{\bar{\sigma}^2}{T-t} \sum_{k=0}^{T-t-1} \sigma_t^{2\gamma^k} \exp\left(\frac{-\alpha\gamma^k}{1-\gamma} - \frac{\kappa^2\gamma^{2k}}{2(1-\gamma^2)}\right)},\end{aligned}\tag{2.4}$$

where

$$\bar{\sigma}^2 = \exp\left(\frac{\alpha}{1-\gamma} + \frac{1}{2} \frac{\kappa^2}{1-\gamma^2}\right)$$

is the steady-state¹⁹ mean of σ_t^2 .

For the case of non-zero correlation between volatility and shocks, one can obtain explicit calculations for the term structure of volatility in the case of normally distributed shocks, but the calculation is more complicated.²⁰ Allowing this correlation is empirically quite important.

¹⁷From (2.1), with constant mean returns, we may write $\log(R_t - \mu)^2 = \log\sigma_{t-1}^2 + \log S_t^2$. Harvey, Ruiz, and Shepard [1992] and Harvey and Shepard [1993] have shown that one can estimate the log auto-regressive model coefficients by quasi-maximum likelihood, which is indeed consistent under certain technical restrictions. Taking $\log S_t^2$ to be normally distributed, this would be a standard setup for Kalman filtering of volatility. In such a setting, we would have access to standard methods for estimating volatility given the coefficients α , γ , and κ , and for estimating these coefficients by maximum likelihood. See, for example, Brockwell and Davis [1991] for the consistency of the estimators in this setting.

¹⁸This calculation is repeated here from Heynen and Kat [1993].

¹⁹That is, $\bar{\sigma}^2 = \lim_t E(\sigma_t^2)$.

²⁰Kalman filtering can be applied in full generality here to get the joint distribution of return shocks conditional on the path of volatility. With joint normality, all second moments of the conditional distribution of return shocks are deterministic. At this point, one applies the law of iterated expectations to get the term volatility as a linear combination of the second moments of the log-normal stochastic volatilities, which is also explicit. The same calculation leads to an analytic solution for option prices in this setting, extending the Hull-White model to the case of volatility that is not independent of shock returns. See Willard [1996].

2.5.3 Garch

Many modelers have turned to ARCH (autoregressive conditional heteroscedasticity) models of volatility proposed by Engle [1982], and the related GARCH and EGARCH formulations, because they capture volatility persistence in simple and flexible ways. For example, the GARCH²¹ model of stochastic volatility proposed by Bollerslev [1986] assumes that

$$\sigma_t^2 = \alpha + \beta(R_t - \mu)^2 + \gamma\sigma_{t-1}^2,$$

where²² α , β , and γ are positive constants. Here, γ is the key persistence parameter: A high γ implies a high carryover effect of past to future volatility, while a low γ implies a heavily damped dependence on past volatility.

One can estimate the parameters α , β , and γ from returns data. For example, estimated GARCH parameters associated with crude oil have maximum likelihood estimates (with t statistics in parentheses) from recent data²³ given by

$$\sigma_t^2 = \underset{(2.95)}{0.155} + \underset{(9.31)}{0.292} (R_t - \mu)^2 + \underset{(30.1)}{0.724} \sigma_{t-1}^2.$$

The estimated persistence parameter for daily volatility is 0.724.

Under the “non-explosivity” condition $\delta \equiv \beta + \gamma < 1$, the steady-state volatility²⁴ is $\bar{\sigma} = \sqrt{\alpha/(1 - \delta)}$. One can show that the term structure of volatility associated with the GARCH model is

$$\bar{v}_{t,T} = \sqrt{(T - t)\bar{\sigma}^2 + (\sigma_{t+1}^2 - \bar{\sigma}^2) \frac{1 - \delta^{T-t}}{1 - \delta}}.$$

A potential disadvantage of the GARCH model, noting that the impact of the current return R_t on σ_{t+1}^2 is quadratic, is that a day of exceptionally large absolute returns can cause instability in parameter estimation, and from this “overshooting” in forecasted volatility. For example, with any reasonable degree of persistence, a market crash or “jump” could imply an inappropriately sustained major impact on forecasted volatility.²⁵

²¹This is known more precisely as the “GARCH(1,1)” model. For specifics and generalizations, as well as a review of the ARCH literature in finance, see Bollerslev, Chou, and Kroner [1992].

²²The GARCH model is in the class (2.2) of Markov models since we can write $\sigma_t = F(\sigma_{t-1}, z_t) = [\alpha + \beta\sigma_{t-1}^2 z_t + \gamma\sigma_{t-1}^2]^{1/2}$, where $z_t = \epsilon_t$ is white noise.

²³See Duffie and Gray (1995).

²⁴This is $\lim_{T \rightarrow \infty} E(\sigma_t^2)$. The non-explosivity condition fails for the parameter estimates given for crude oil.

²⁵Sakata and White [1996] have therefore suggested “high-breakdown point” estimators in this sort

2.5.4 Egarch

A potentially more flexible model of persistence is the exponential Garch, or “EGARCH” model proposed by Nelson [1991], which takes the form²⁶

$$\log \sigma_t^2 = \alpha + \gamma \log \sigma_{t-1}^2 + \beta_1 \left(\frac{R_t - \mu}{\sigma_{t-1}} \right) + \beta_2 \left(\left| \frac{R_t - \mu}{\sigma_{t-1}} \right| - \sqrt{\frac{2}{\pi}} \right).$$

The term structure of volatility implied by the EGARCH model is

$$\bar{v}_{t,T} = \sqrt{\sum_{k=0}^{T-t-1} C_k \sigma_t^{2\gamma^k}},$$

where C_k is a relatively complicated constant given, for example, by Heynen and Kat [1993]. Nelson [1990] has shown that the EGARCH model and the log-auto-regressive model (2.2) converge with decreasing period length, and appropriate normalization of coefficients, to the same model.

2.5.5 Cross-Market Garch

One can often infer volatility-related information for one market from changes in the volatility of returns in another. A simple model that accounts for cross-market inference is the multivariate GARCH model. For example, a simple 2-market version of this model takes

$$\begin{pmatrix} \sigma_{a,t}^2 \\ \sigma_{ab,t} \\ \sigma_{b,t}^2 \end{pmatrix} = \alpha + \beta \begin{pmatrix} R_{a,t}^2 \\ R_{a,t}R_{b,t} \\ R_{b,t}^2 \end{pmatrix} + \gamma \begin{pmatrix} \sigma_{a,t-1}^2 \\ \sigma_{ab,t-1} \\ \sigma_{b,t-1}^2 \end{pmatrix},$$

where

- $R_{a,t}$ is the return in market a at time t
- $R_{b,t}$ is the return in market b at time t
- $\sigma_{a,t-1}$ is the conditional volatility of $R_{a,t}$
- $\sigma_{b,t-1}$ is the conditional volatility of $R_{b,t}$
- $\sigma_{ab,t-1}$ is the conditional covariance between $R_{a,t}$ and $R_{b,t}$

of environment, and give example estimates for S&P 500 returns.

²⁶The term $\sqrt{2/\pi}$ is equal to $E_t[|(R_t - \mu)/\sigma_{t-1}|]$.

- α is a vector with 3 elements
- β is a 3×3 matrix
- γ is a 3×3 matrix.

With β and γ assumed to be diagonal for simplicity, a maximum-likelihood estimate for the bivariate GARCH model for heating oil (a) and crude oil (b) is given by

$$\begin{bmatrix} \sigma_{a,t}^2 \\ \sigma_{ab,t} \\ \sigma_{b,t}^2 \end{bmatrix} = \begin{bmatrix} .23963 \\ (1.9830) \\ .11408 \\ (1.6360) \\ .083939 \\ (1.5507) \end{bmatrix} + \begin{bmatrix} .15663 & 0 & 0 \\ (4.8101) & & \\ 0 & 1.3227 & 0 \\ & (2.5760) & \\ 0 & 0 & .13509 \\ & & (2.0763) \end{bmatrix} \begin{bmatrix} R_{a,t}^2 \\ R_{a,t}R_{b,t} \\ R_{b,t}^2 \end{bmatrix} \\ + \begin{bmatrix} .81675 & 0 & 0 \\ (19.076) & & \\ 0 & .84643 & 0 \\ & (17.969) & \\ 0 & 0 & .86455 \\ & & (18.609) \end{bmatrix} \begin{bmatrix} \sigma_{a,t-1}^2 \\ \sigma_{ab,t-1} \\ \sigma_{b,t-1}^2 \end{bmatrix},$$

with t -statistics shown in parentheses.

One notes the differences between the univariate and multivariate GARCH parameters for crude oil (alone). In principle, cross-market information can only improve the quality of the model if the multivariate model is appropriate.

2.6 Term Structures of Tail-Fatness and Volatility

Like volatility, tail-fatness, as measured for example by kurtosis, has a term structure according to the time horizon over which the total return is calculated. In the plain-vanilla model, the term structures of both volatility and tail-fatness are flat. In general, the term structures of tail-fatness and volatility have shapes that depend markedly on the source of tail-fatness. Here are several cases to consider.

1. **Jumps** Consider the case of constant mean and volatility, and *iid* shocks with fat tails. (This could be, for example, a jump-diffusion setting.) In this case, the term structure of volatility is flat. As illustrated in Figure 9, the central limit theorem tells us that averaging *iid* variables leads to a normally distributed variable.²⁷ We therefore expect that the term structure of tail fatness for the jump-diffusion model underlying Figure 6 to be declining, when measured by kurtosis. This is borne out in Figure 10. For example, while the 1986-96 sample daily return kurtosis for the S&P 500 index is 111, at the monthly level, the sample kurtosis for this period is 16.5 (estimated on an overlapping basis). If we were to measure tail fatness by the number of standard deviations to a particular critical value, such as the 0.01 critical value, however, the term structure of tail fatness would first increase and then eventually decline to the normal level of 2.326, as illustrated in Figure 11. At the 0.01 critical level, for typical market parameters such as those shown in Figures 6 and 7, the likelihood of a jump on a given day is smaller than 0.01, so the impact of jumps on critical values of the distribution shows up much farther out in the tail than at the 0.01 critical value. At an expected frequency of 2 jumps per year, we would expect the 0.01-critical value to be more seriously affected by jumps at a time horizon of a few weeks.

2. **Stochastic Volatility** Suppose we have constant mean returns and *iid* normal shocks, with stochastic volatility that is independent of the shocks. The term structure of volatility can have essentially any shape, depending on the time-series properties of $\sigma_t, \sigma_{t+1}, \dots$. For example, under an autoregressive model (2.2) of stochastic volatility, the term structure of volatility (2.4) approaches an asymptote from above or from below, as illustrated in Figure 12, depending on whether the initial volatility σ_t is above or below the stationary level. This plot is based on a theoretical stochastic volatility model (2.2), using as the parameters the maximum-likelihood estimates $\alpha = -5.4$, $\gamma = 0.38$, and $\kappa = 1.82$ for this model fitted to the Hang Seng Index by Heynen and Kat [1993]. The three initial levels shown are the steady-state mean volatility implied by the model (B), one standard deviation of the steady-state distribution above the mean (A), and one

²⁷The theory of large deviations, outlined in Appendix B for a different application, can be used to address the speed of convergence to normal tails. For special cases, such as our simple jump-diffusions, the calculations are easy. Figure 9 plots the densities of *t*-distributed variables with the indicated degrees of freedom. The case of “ $t = \infty$ ” is standard normal.


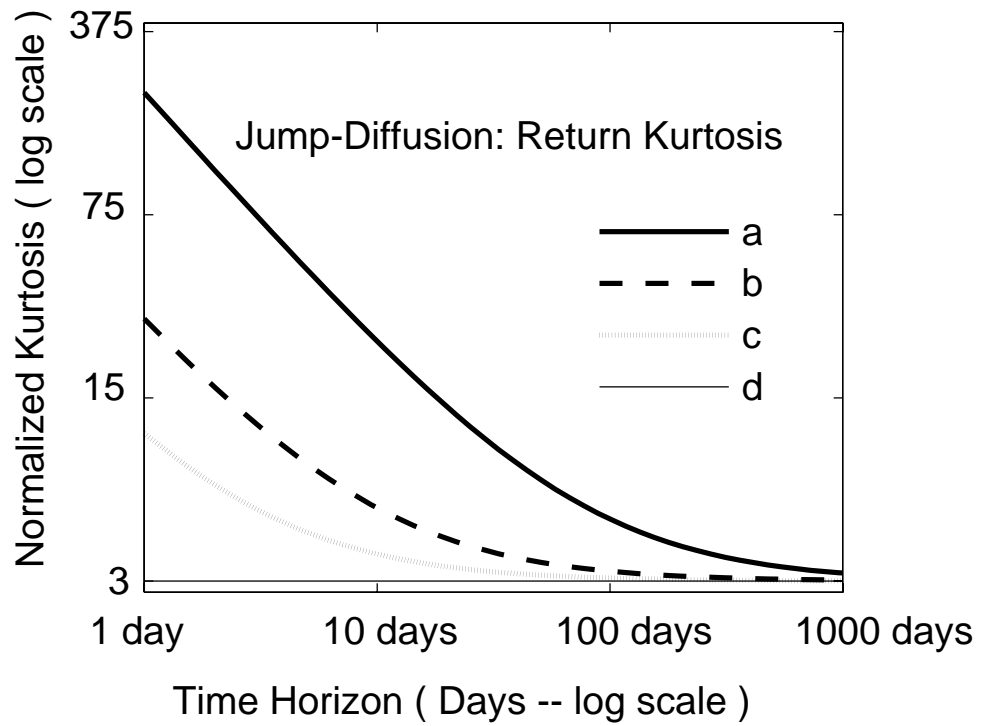


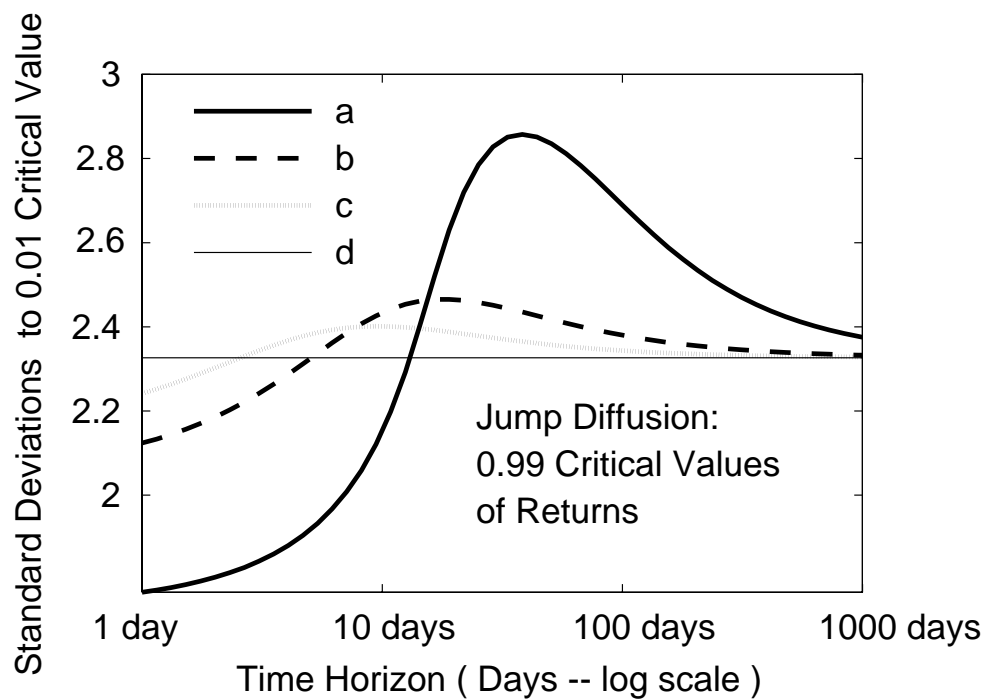
Figure 9: Tail-Thinning Effect of the Central Limit Theorem



Kurtosis of return is shown for the following cases:

- (a) $\sigma = 15\%$, $\lambda = 1.0$, $\nu = 10\%$; (b) $\sigma = 15\%$, $\lambda = 2.0$, $\nu = 5\%$
(c) $\sigma = 15\%$, $\lambda = 3.0$, $\nu = 3.33\%$; (d) plain-vanilla with $\sigma = 15\%$.

Figure 10: Term Structure of Kurtosis for the Jump-Diffusion Model



99% critical value is shown for the following cases:
 (a) $\sigma = 15\%$, $\lambda = 1.0$, $\nu = 10\%$; (b) $\sigma = 15\%$, $\lambda = 2.0$, $\nu = 5\%$
 (c) $\sigma = 15\%$, $\lambda = 3.0$, $\nu = 3.33\%$; (d) plain-vanilla with $\sigma = 15\%$.


Figure 11: Term Structure of 0.99 Critical Values of the Jump-Diffusion Model

standard deviation of the steady-state distribution below the mean (C). Starting from the steady-state mean level of volatility, the term structure of kurtosis is increasing and then eventually decreasing back to normal, as illustrated²⁸ for case “B” in Figure 13. This “hump-shaped” term structure of tail fatness arises from the effect of taking mixtures of normals with different variances drawn from the stochastic volatility model, which initially increases the term structure of tail fatness. The tail fatness ultimately must decline to standard normal, as indicated in Figure 13 by virtue of the central limit theorem.²⁹ For typical VaR time horizons, however, the term structure of kurtosis is increasing from the standard normal level of 3, as shown in Figure 14. This plot is based on three different theoretical stochastic volatility models, using as the parameters the maximum-likelihood estimates for the British Pound (A), which has extremely high mean reversion of volatility and extremely high volatility of volatility, the Hang-Seng Index (B), which has more moderate mean reversion and volatility of volatility, and the S&P 500 Index (C), which is yet more moderate.³⁰ Uncertainty about the initial level of volatility would cause some variation from this story, and effectively increase the initial level of kurtosis, as illustrated for the case “A” of random initial volatility, shown in Figure 13, for which the initial volatility is drawn from the steady-state distribution implied by the estimated parameters. A caution is in order: We can guess that the presence of jumps would result in a relatively severe mis-specification bias for estimators of the stochastic volatility model (2.2). For example, a jump would appear in the estimates in the form of a high volatility of volatility and a high mean-reversion of volatility. The presence of both jumps and stochastic volatility is anticipated for these three markets. Evidence for both jumps and stochastic volatility (modeled in the form of a GARCH) is presented by Jorion [1989].

²⁸This plot is based on a theoretical stochastic volatility model (2.2), using as the parameters the maximum-likelihood estimates $\alpha = -8.8$, $\gamma = 0.18$, and $\kappa = 3.5$ for this model fitted to the dollar price of the British Pound by Heynen and Kat [1993].


²⁹We are grateful to Ken Froot for pointing this out. We can rely on the fact that, over time intervals of “large” length, the volatilities at the beginning and end of the intervals are “essentially” independent, in the sense of the central limit theorem for recurrent Markov processes.

³⁰These parameter estimates are given above for the Hang-Seng Index and the Pound, and for the S&P 500 are $\alpha = -0.51$, $\gamma = 0.94$, and $\kappa = 0.055$ fitted by Heynen and Kat [1993].




A: High Initial Deterministic Volatility
 B: Steady-State Average Initial Deterministic Volatility
 C: Low Initial Deterministic Volatility.

Figure 12: Term Structure of Volatility (Hang Seng Index - Estimated)



A: Steady-State Random Initial Volatility
 B: Deterministic Initial Volatility

Figure 13: Long-Run Kurtosis of Stochastic-Volatility Model



A – British Pound, B – Hang Seng Index, C – S&P 500 Index

Figure 14: Estimated Term Structure of Kurtosis for Stochastic Volatility

3. **Mean Reversion** Suppose we have constant daily volatility and *iid* normal shocks, but we have *mean reversion*. For example, let $\mu_t = \alpha(R^* - \bar{R}_{t-1})$, where $\alpha > 0$ is a coefficient that “dampens” cumulative total return $\bar{R}_t = R_1 + \dots + R_t$ to a long-run mean R^* . This model, which introduces negative autocorrelation in returns, would be consistent, roughly, with the behavior explained by Froot [1993] and O’Connell [1996] of foreign exchange rates over very long time horizons. For this model, the term structure of volatility is declining to an asymptote, while the term structure of tail fatness is flat.

2.7 Estimating Current Volatility

A key to measuring VaR is obtaining an estimate of the current volatility σ_t for each underlying market. Various methods could be considered. The previous sub-section offers a sample of stochastic volatility models that can, in principle, be estimated from historical data. Along with parameter estimates, one obtains at each time period an estimate of the current underlying volatility. See Hamilton [1994]. Other conventional estimators for current volatility are described below.

2.7.1 Historical Volatility

The *historical volatility* $\hat{\sigma}_{t,T}$ implied by returns R_t, R_{t+1}, \dots, R_T is the usual naive volatility estimate

$$\hat{\sigma}_{t,T}^2 = \frac{1}{T-t} \sum_{s=t+1}^T (R_s - \hat{\mu}_{t,T})^2,$$

where $\hat{\mu}_{t,T} = (R_{t+1} + \dots + R_T)/(T-t)$. In a plain-vanilla setting, this (maximum-likelihood) estimator of the constant volatility parameter σ is optimal, in the usual statistical sense. If the plain-vanilla model of returns applies at arbitrarily fine data frequency (with suitable adjustment of μ and σ for period length), then one can learn the volatility parameter within an arbitrarily short time interval³¹ from the historical volatility estimator. Empirically, however, returns at exceptionally high frequency have statistical properties that are heavily dependent on institutional properties of the

³¹Literally, $\lim_{T \rightarrow \infty} \hat{\sigma}_{t,T} = \sigma$ almost surely, and since an arbitrary number of observations of returns is assumed to be possible within an arbitrarily small time interval, this limit can be achieved in an arbitrarily small amount of calendar time.

market that are of less importance over longer time periods.³²

For essentially every major market, historical volatility data strongly indicate that the constant-volatility model does not apply. For example, the rolling 180-day historical volatility estimates shown in Figure 15, for a major Taiwan equity index, appear to indicate that volatility is changing in some persistent manner over time.³³ Incidentally, in the presence of jumps we would expect to see large upward “jumps” in the 180-day rolling historical volatility, at the time of a jump in the return, coupled with a downward jump in the rolling volatility precisely 180 days later, which suggests caution in the use of rolling volatility as an estimator for actual volatility.

Exponential weighting of data can be incorporated in order to place more emphasis on more recent history in estimating volatility. This amounts to a restrictive case of the GARCH model, and is the standard adopted by J.P. Morgan for its RiskMetrics volatility estimates. (See Phelan [1995]).


2.7.2 Black-Scholes Implied Volatility

In the plain-vanilla setting, it is well known that the price of an option at time t , say a European call, is given explicitly by the famous Black and Scholes [1973] formula $C_t = C^{BS}(P_t, \sigma, \tau, K, r)$, given the underlying price P_t , the strike price K , the time τ to expiration, the continuously compounding constant interest rate r , and the volatility σ . It is also well known that this formula is strictly increasing in σ , as shown in Figure 16, so that, from the option price C_t , one may theoretically infer without error

³²When estimating σ , in certain markets one can also take special advantage of additional financial price data, such as the high and low prices for the period, as shown by Garman and Klass [1980], Parkinson [1980], and Rogers and Satchell [1991].

³³Of course, even in the constant-volatility setting, one expects the historical volatility estimate to vary over time, sometimes dramatically, merely from random variation in prices. (This is sometimes called “sampling error.”) One can perform various tests to ascertain whether changes in historical volatility are “so large” as to cause one to reject the constant volatility hypothesis at a given confidence level. For example, under the constant volatility hypothesis, the ratio $F_{a,b} = \hat{\sigma}_{t(a),T(a)}^2 / \hat{\sigma}_{t(b),T(b)}^2$ of squared historical volatilities over non-overlapping time intervals has the F distribution (with degrees of freedom given by the respective lengths of the two time intervals). From standard tables of the F distribution one can then test the constant-volatility hypothesis, rejecting it at, say, the 95-percent confidence level, if $F_{a,b}$ is larger than the associated critical F statistic. (One should take care not to select the time intervals in question in light of one’s impression, based on observing prices, that volatility apparently differs between the two periods. This would introduce selection bias that makes such classical tests unreliable.)

Taiwan Weighted (TW): 180-Day Historical Volatility



Source: Datastream
Daily Excess Returns

Figure 15: Rolling Volatility for Taiwan Equity Index

the volatility parameter $\sigma = \sigma^{BS}(C_t, P_t, \tau, K, r)$. The function $\sigma^{BS}(\cdot)$ is known³⁴ as the Black-Scholes implied volatility. While no explicit formula for σ^{BS} is available, one can compute implied volatilities readily with simple numerical routines.³⁵




Figure 16: Black-Scholes “Price of Volatility”

In many (but not all) markets, option-implied volatility is a more reliable method of forecasting future volatility than any of the standard statistical methods that have been based only on historical return data. (For the empirical evidence, see Canina and Figlewski [1993], Campa and Chang [1995], Day and C.Lewis [1992], Jorion [1995], Lamoureux and Lastrapes [1993], and Scott [1992].) Of course, some markets have no reliable options data!

Because we believe that volatility is changing over time, one should account for this in one’s option-pricing model before estimating the volatility implied by option prices. For example, Rubinstein [1994], Dupire [1992], Dupire [1994], and Derman and Kani [1994] have explored variations of the volatility model

$$\sigma_t = F(P_t, t), \tag{2.5}$$

³⁴This idea goes back at least to Beckers [1981].

³⁵For these and many other details on the Black-Scholes model and extensions, one may refer to Cox and Rubinstein [1985], Stoll and Whaley [1993], and Hull [1993], among many other sources.

where P_t is the price at time t of the underlying asset, for some continuous function F that is chosen so as to match the modeled prices of traded options with the prices for these options that one observes in the market. This is sometimes called the *implied-tree* approach.³⁶

2.7.3 Option-Implied Stochastic Volatility

One can also build option valuation models that are based on stochastic volatility, and obtain a further generalization of the notion of implied volatility. For instance, a common special case of the stochastic volatility models of Hull and White [1987], Scott [1987], and Wiggins [1987] assumes that, after switching to risk-neutral probabilities, we have independent shocks to returns and volatility. With this (in the usual limiting sense of the Black-Scholes model for “small” time periods) one obtains the stochastic-volatility option-pricing formula

$$C_t = C^{SV}(P_t, \sigma_t, t, T, K, r) \equiv E^* \left[C^{BS}(P_t, v_{t,T}, T - t, K, r) \right], \quad (2.6)$$

where

$$v_{t,T} = \sqrt{\frac{1}{T-t}(\sigma_t^2 + \dots + \sigma_{T-1}^2)} \quad (2.7)$$

is the root-mean-squared term volatility, $C^{BS}(\cdot)$ is the Black-Scholes formula, and E^* denotes risk-neutral expectation at time t . This calculation follows from the fact that, if volatility is time-varying but deterministic, then one can substitute $v_{t,T}$ in place of the usual constant volatility coefficient to get the correct option price $C^{BS}(P_t, v_{t,T}, T - t, K, r)$ from the Black-Scholes model.³⁷ With the above independence assumption, one can simply average this modified Black-Scholes formula over all possible (probability-weighted) realizations of $v_{t,T}$ to get the result (2.6).

For at-the-money options (specifically, options struck at the forward price of the underlying market), the Black-Scholes option pricing formula is, for practical purposes, essentially linear in the volatility parameter, as illustrated in Figure 16. In the “Hull-White” setting of independent stochastic volatility, the naive Black-Scholes implied volatility for at-the-money options is therefore an effective (albeit risk-neutralized) forecast of the root-mean-squared term volatility $v_{t,T}$ associated with the expiration date of

³⁶See Jackwerth and Rubinstein [1996] for generalizations and some empirical evidence.

³⁷This was noted by Johnson and Shanno [1987].

the option. On top of any risk-premium³⁸ associated with stochastic volatility, correlation between volatility shocks and return shocks causes a bias in Black-Scholes implied volatility as an estimator of the expectation of the root-mean-squared volatility $v_{t,T}$. (This bias can be corrected; see for example Willard [1996].) The root-mean-squared volatility $v_{t,T}$ is itself larger than annualized average volatility $(\sigma_t + \dots + \sigma_{T-1})/\sqrt{(T-t)}$ over the period before expiration, because of convexity effect of squaring in (2.7) and Jensen’s Inequality.

The impact on Black-Scholes implied volatilities of randomness in volatility is more severe for away-from-the-money options than for at-the-money options. A precise mathematical statement of this is rather complicated. One can see the effect, however, through the plots in Figure 16 of the Black-Scholes formula with respect to volatility against the exercise price. For near-the-money options, the plot is roughly linear. For well-out-of-the-money options, the plot is convex. A “smile” in plots of implied volatilities against exercise price thus follows from (2.6), Jensen’s inequality, and random variation in $v_{t,T}$. We can learn something about the degree of randomness in volatility from the degree of convexity of the implied-vol schedule.³⁹

It may be useful to model volatility that is both stochastic, as well as dependent on the price of the underlying asset. For example, we may wish to replace the univariate Markovian stochastic-volatility model with

$$\sigma_t = F(\sigma_{t-1}, P_t, z_t, t),$$

so that one combines the stochastic-volatility approach with the “implied tree” approach of Rubinstein, Dupire, and Derman-Kani. To our knowledge, this combined model has not yet been explored in any systematic way.

2.7.4 Day-of-the-week and other seasonal volatility effects

Among other determinants of volatility are “seasonality” effects. For example, there are day-of-the-week effects in volatility that reflect institutional market features, including the desire of market makers to close out their positions over weekends. One can

³⁸See, for example, Heston [1993] for an equilibrium model of the risk premium in stochastic volatility.

³⁹We can also learn about correlation between returns and changes in volatility from the degree of “tilt” in the smile curve. See, for example, Willard [1996]. For econometric models that exploit option prices to estimate the stochastic behavior of volatility, see Pastorello, Renault, and Touzi [1993] and Renault and Touzi [1992].

“correct” for this sort of “seasonality,” for example by estimating volatility separately for each day of the week.

For another example, the seasons of the year play an important role in the volatilities of energy products. For instance, the demand for heating oil depends on winter weather patterns, which are determined in the winter. The demand for gasoline is greater, and shows greater variability, in the summer, and gasoline prices therefore tend to show greater variability during the summer months.

2.8 Skewness

Skewness is a measure of the degree to which positive deviations from mean are larger than negative deviations from mean, as measured by the expected third power of these deviations. For example, equity returns are typically negatively skewed, as show in in Appendix F. If one holds long positions, then negative skewness is a source of concern for value at risk, as it implies that large negative returns are more likely, in the sense of skewness, than large positive returns.

If skewness in returns is caused by skewness in shocks alone, and if one’s model of returns is otherwise plain vanilla, we would expect the skewness to become “diversified away” over time, through the effect of the central limit theorem, as illustrated in Figure 17 for positively skewed shocks.⁴⁰ In this case, that is, the term structure of skewness would show a reversion to zero skewness over longer and longer time horizons. If, on the other hand, skewness is caused, or exacerbated, by correlation between shocks and changes in volatility (negative correlation for negative skewness), then we would not see the effect of the central limit theorem shown in Figure 17.

2.9 Correlations

A complete model of price risk requires not only models for mean returns, volatilities, and the distribution of shocks for each underlying market, but also models for the relationships across markets among these variables. For example, a primary cross-market piece of information is the conditional correlation at time t between the shocks in markets i and j . Campa and Chang [1995] address the relative ability to forecast

⁴⁰Plotted in Figure 17 are the densities of $V(n)/n$ for various n , where $V(n)$ is the sum of n independent squared normals. That is $V \sim \chi_n^2$. By the central limit theorem, the density of $\sqrt{n}V(n)$ converges to that of a normal.




Figure 17: Skewness Correcting Effect of Diversification

correlation of various approaches, including the use of the implied volatilities of cross-market options.

In order to measure value-at-risk over longer time horizons, in addition to the conditional return correlations one would also depend critically on one's assumptions about correlations across markets between changes in volatilities.

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.

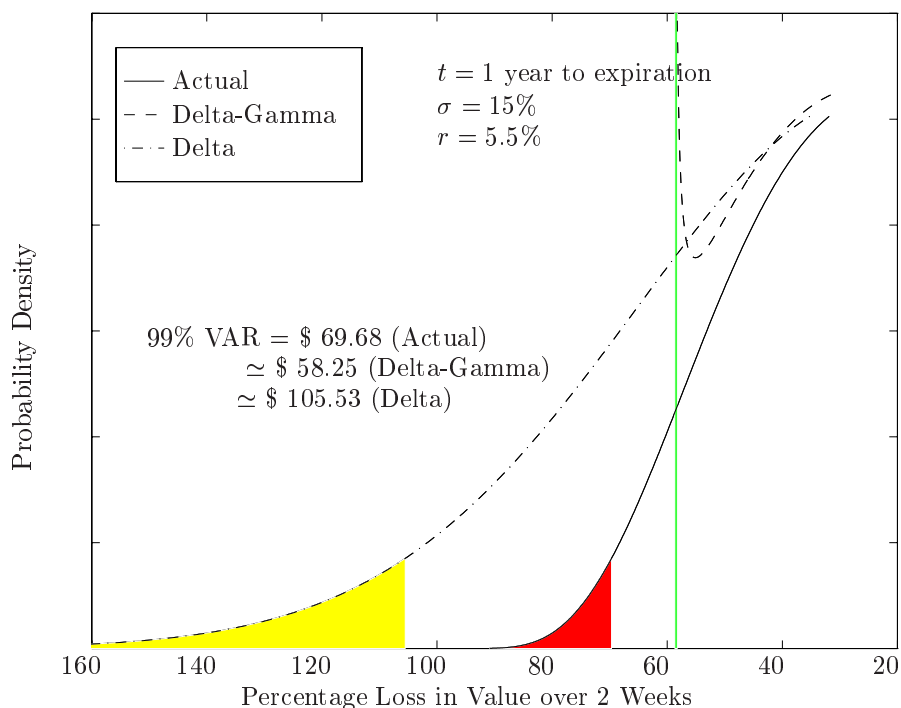


Figure 19: 2-Week Loss on 20% Out-Of-Money Put (Plain-Vanilla Returns)

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.

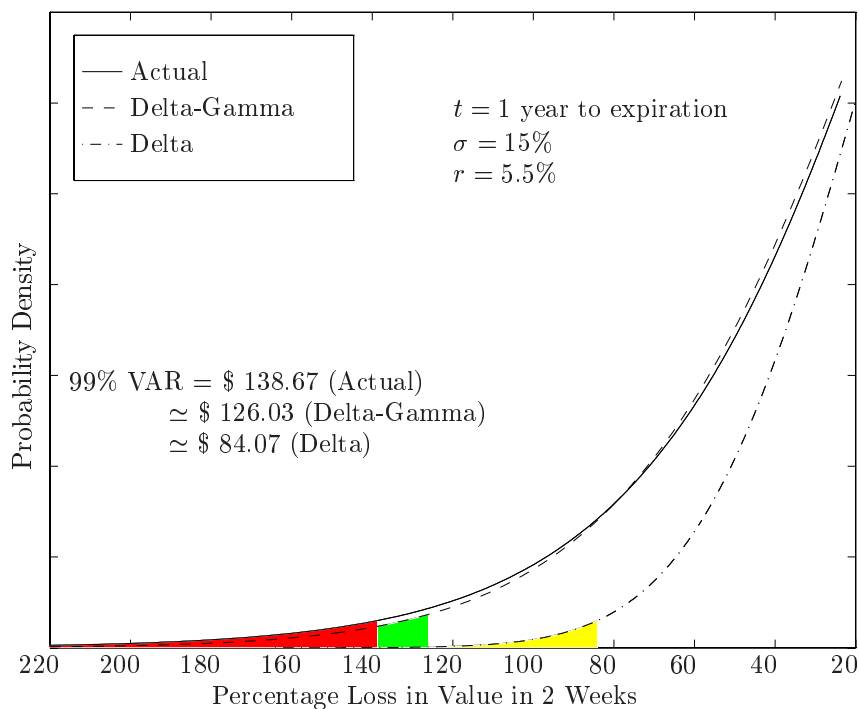


Figure 20: 2-Week Loss on Short 20% Out-Of-Money Put (Plain-Vanilla Returns)

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].

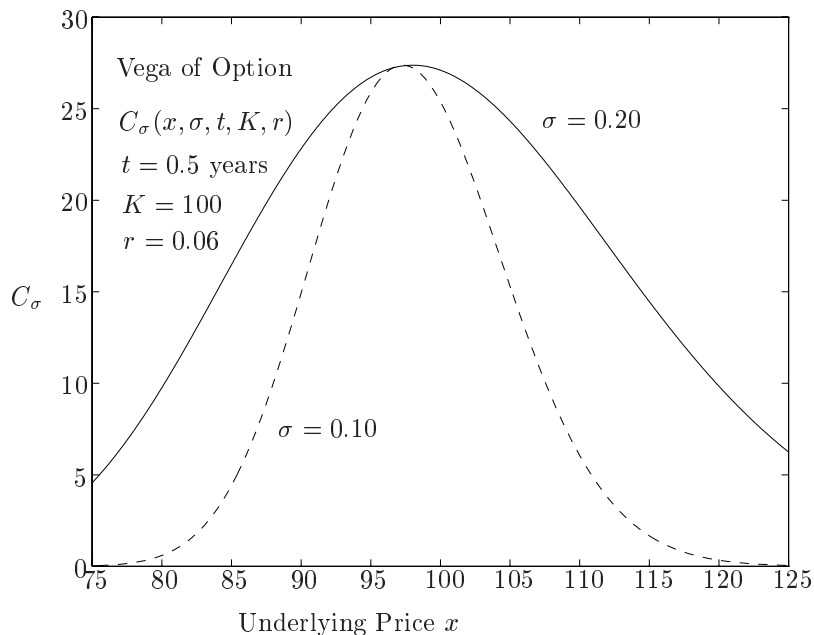


Figure 23: Sensitivity of Call Price to Volatility

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.

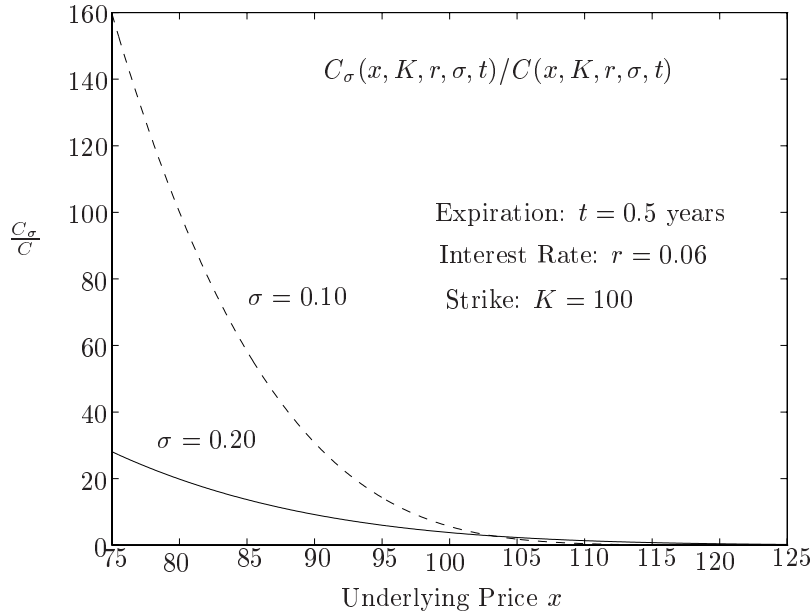


Figure 24: Relative Sensitivity of Call Price to Volatility

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

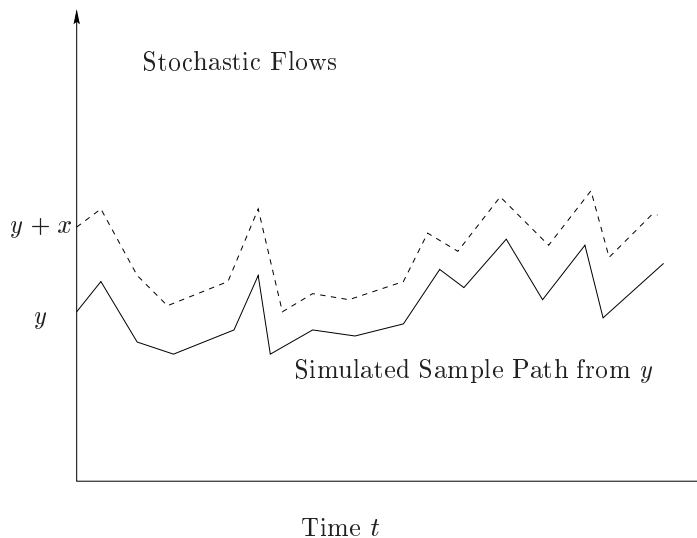


Figure 25: Using Stochastic Flows to Avoid Re-Simulation

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.

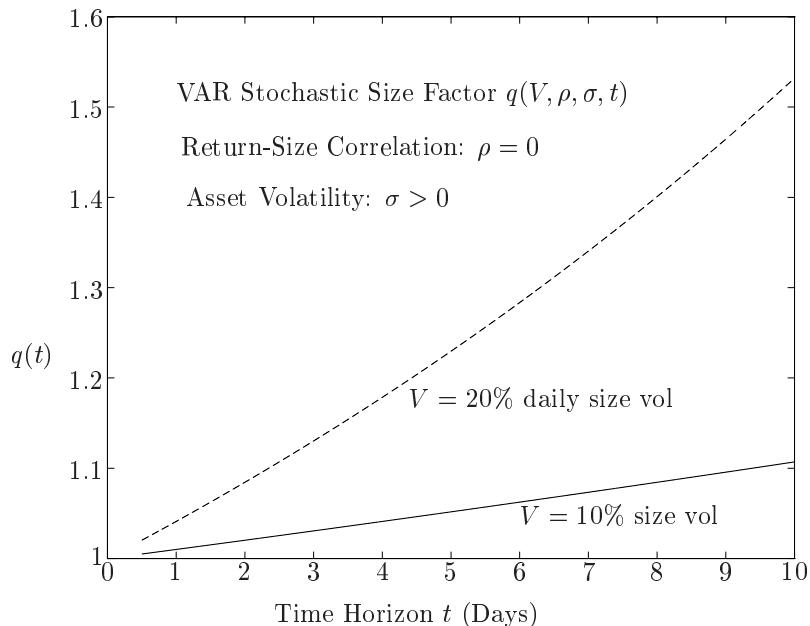


Figure 26: Impact of Position Size Volatility on Value at Risk

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.

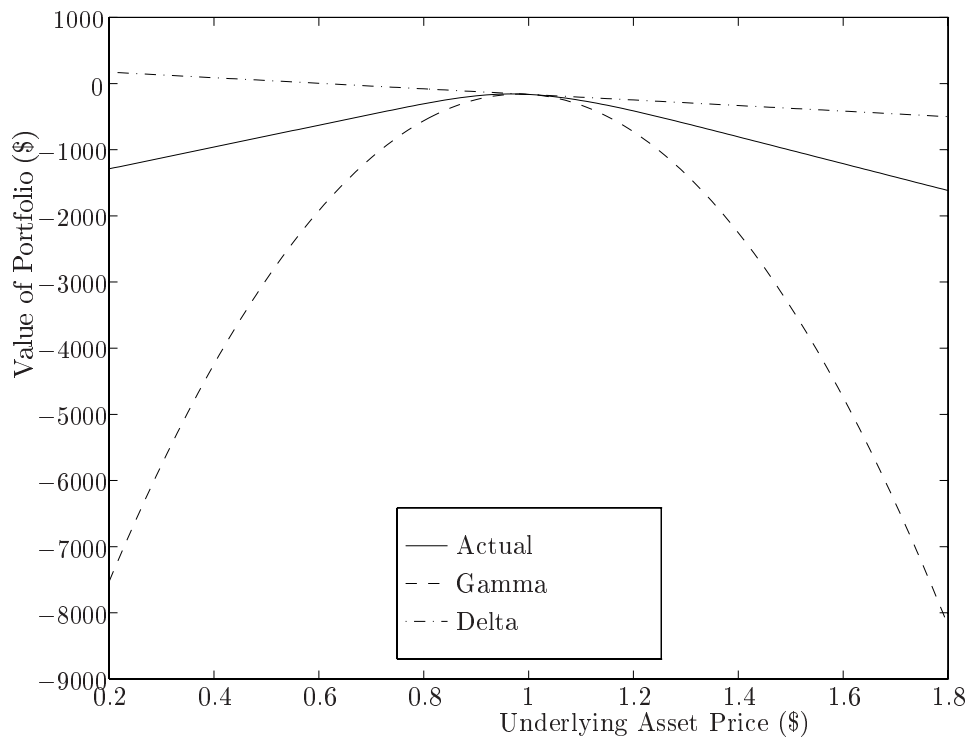


Figure 27: Value Approximations for a Random Portfolio of 10,996 options

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].

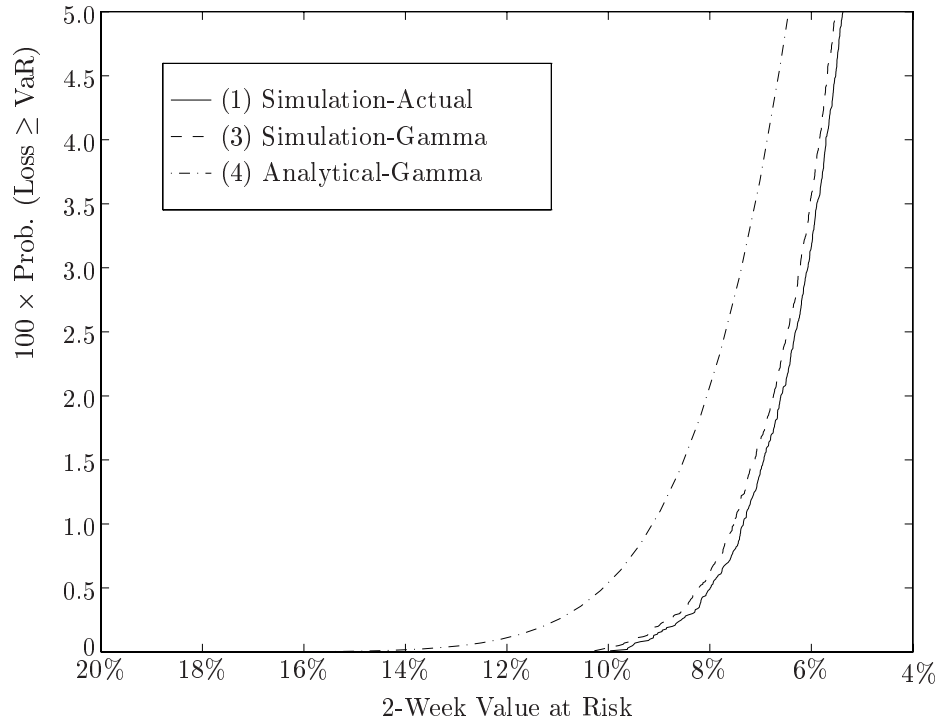


Figure 29: Value at Risk of Long Option Portfolio – Plain-Vanilla Model

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.

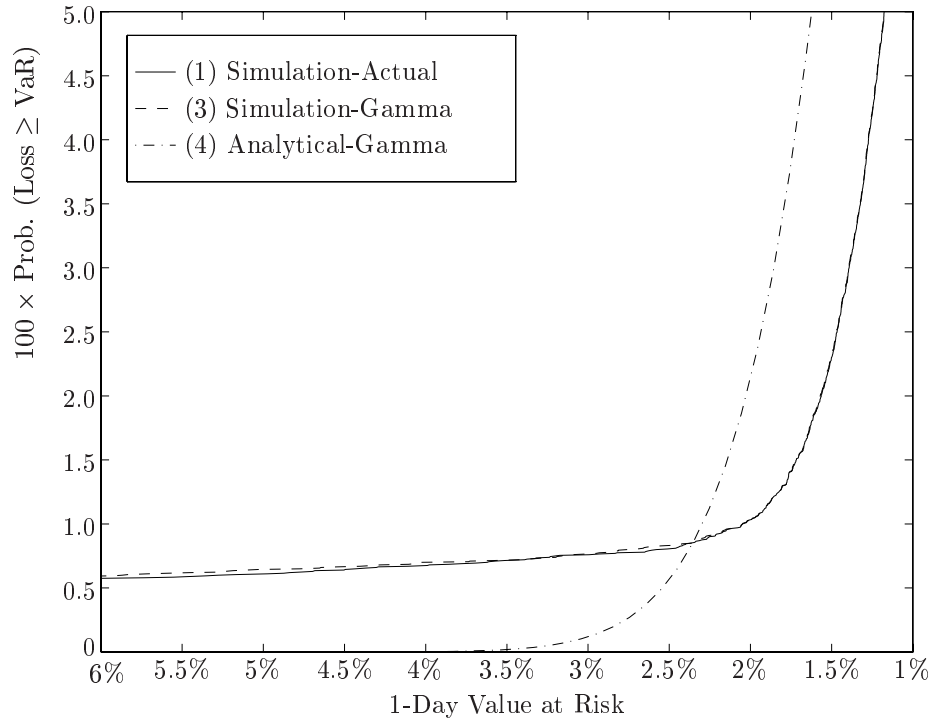


Figure 30: Value at Risk for Short Option Portfolio – Jump-Diffusion Model 3

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].

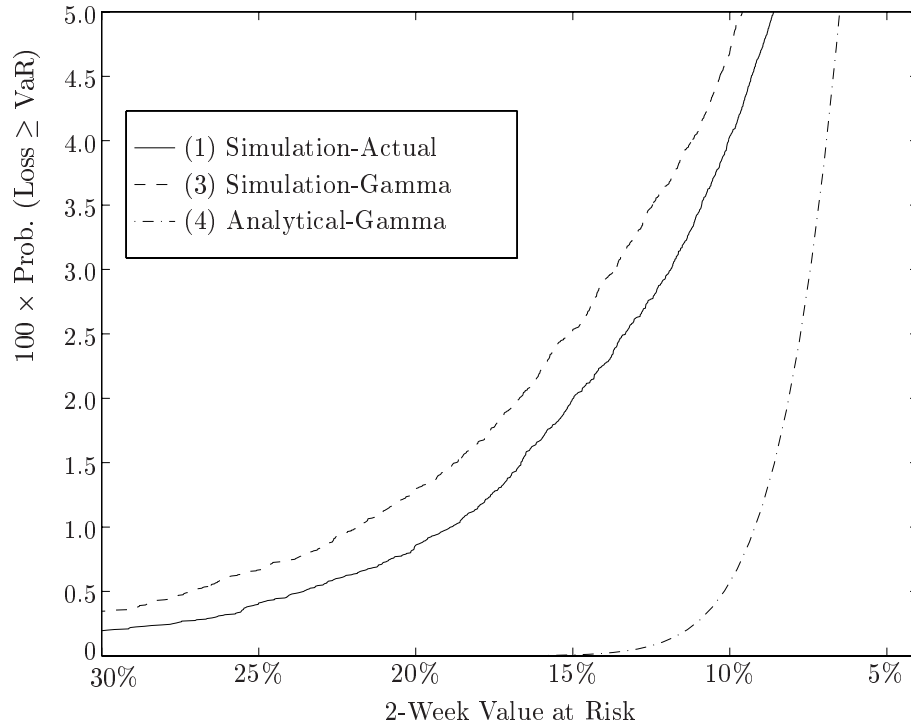


Figure 31: Value at Risk for Short Option Portfolio – Jump-Diffusion Model 3

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.

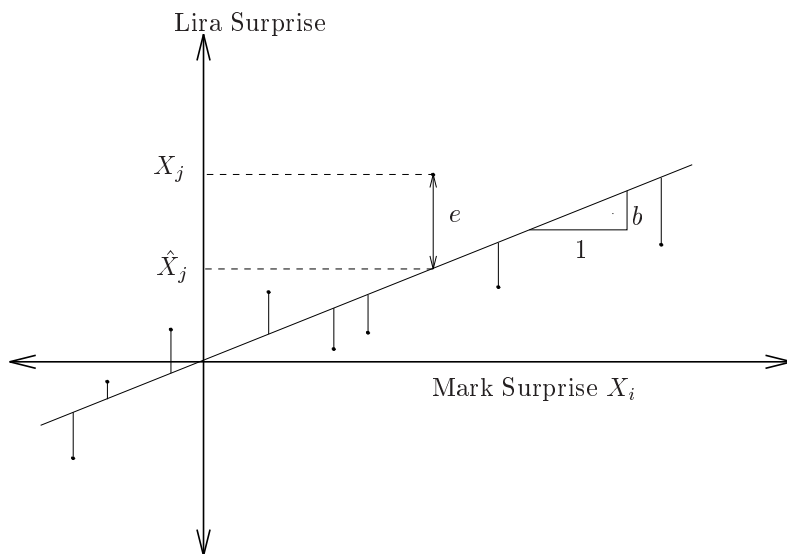


Figure 34: Regressing Lira “Surprises” on Mark “Surprises”

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.

Appendices

A. Simulating Fat Tailed Distributions

Suppose one wants to simulate a random variable of zero mean and unit variance, but with a given degree of tail fatness (fourth moment). Sticking to the more-or-less “bell-curved” shapes for the probability density of returns (and ignoring skewness), one could adopt the following approach,⁶⁵ based on the idea that a random variable has fat tails if it can be expressed as a random mixture of normal random variables of different variances.

First draw a random variable Y whose outcomes are 1 and 0, with respective probabilities p and $1 - p$. Independently, draw a standard normal random variable Z . Let α and β be the standard deviations of the two normals to be “mixed.” If the outcome of Y is 1, let $X = \alpha Z$. If the outcome of Y is 0, let $X = \beta Z$. We want to choose α and β so that the variance of X is 1. We have $\text{var}(X) = p\alpha^2 + (1 - p)\beta^2 = 1$, so that

$$\beta = \sqrt{\frac{1 - p\alpha^2}{1 - p}}.$$

Now we can choose p and α to achieve a given kurtosis or 0.99 critical value. The kurtosis of X is $E(X^4) = 3(p\alpha^4 + (1 - p)\beta^4)$.

B. How Many Scenarios is Enough?

This appendix shows how to compute an answer to the following sort of question, which can be used to decide how many scenarios is “sufficient” for measuring the percentile measure for the loss on a given portfolio of positions.

“Suppose an event occurs with probability p . With k independent trials, what is the likelihood $\pi(k)$ that we estimate p to be δ or larger?”

For the case of risk management, the event of concern is whether losses are no greater than some critical level. The danger to be avoided is over-estimation of the

⁶⁵This approach was suggested by Robert Litterman of Goldman Sachs at a meeting in March 1996 of the Financial Research Initiative at Stanford University.

probability of this event, for it would leave a firm's risk manager with undue confidence regarding the firm's risk. For example, The key "error likelihood" $\pi(k)$ depends of course on the number k of scenarios simulated. As k goes to infinity, the law of large numbers implies that $\pi(k)$ goes to zero. For example, we can show the following. If $p = 0.95$ and $\delta = 0.975$, then regardless of the distribution of the underlying market values of positions in the book, $\pi(k) \leq e^{-0.008k}$. With 1000 scenarios, for instance, the error probability is less than 3.5 parts per 10,000.

We obtain this upper bound on $\pi(k)$ along with the following general result, which allows us to derive error probabilities for other cases of estimated percentile and assumed true percentiles than $p = 0.95$ and $\delta = 0.975$. For example, with $p = 0.95$ and $\delta = 0.975$, the number of scenarios necessary to keep the error probability below 0.01 is approximately 4600.

In order to state the general result, we suppose that Y_1, Y_2, \dots , is an independently and identically distributed (*i.i.d.*) sequence of random variables, with $E(Y_i) = \mu$. We know that $\hat{\mu}(k) = (Y_1 + \dots + Y_k)/k \rightarrow \mu$ almost surely, but at what rate? We let g denote the moment-generating function of Y_i , that is,

$$g(\theta) = E[\exp(\theta Y_i)].$$

Large Deviations Theorem. Under mild regularity,⁶⁶

$$P[\hat{\mu}(k) \geq \delta] \leq e^{-k\gamma(\theta)},$$

where $\gamma(\theta) = \delta\theta - \log[g(\theta)]$.

We can get an optimal upper bound of this form by maximizing $\gamma(\theta)$ with respect to θ . Under purely technical⁶⁷ conditions, the solution θ^* provides an upper bound $\exp[-k\gamma(\theta^*)]$ that, asymptotically with k , cannot be improved.

In our application, we suppose that X is the random variable whose percentiles are of interest. We let Y_i be a "binomial random variable" (that is, a "Bernoulli trial"), with an outcome of 1 if the i -th simulated outcome of X is above the cutoff percentile level, and zero otherwise. The probability that $Y_i = 1$ is some number p , the true quantile score for this cutoff, which is 0.95 in the above example. We let $\hat{p}(k) = (Y_1 + \dots + Y_k)/k$, be the estimate of p . We are checking to see how likely it is that our estimate is larger than δ , which is 0.975 in the above example.

⁶⁶For details, see Durrett [1991].

⁶⁷Again, see Durrett [1991].

Optimizing on θ , we have

$$P(\hat{p}(k) \geq \delta) \sim \exp(-k\Gamma),$$

where

$$\Gamma = \delta \log \delta + (1 - \delta) \log(1 - \delta) - \delta \log p - (1 - \delta) \log(1 - p).$$

With $p = 0.95$ and $\delta = 0.975$, $\Gamma = 0.008$. We can now solve the equation $\exp(-\Gamma \times k) = c$ for k . For a confidence of $c = 0.99$, we see that

$$k = -\frac{1}{\Gamma} \log(c) = 576 \text{ simulations.}$$

That is, the probability that $\hat{p}(k) \geq 0.975$ is roughly $\exp(-576 \times 0.008) = 0.01$. For $\delta = 0.96$, in order to achieve 99-percent confidence, $k = 4600$ simulations are suggested.

C. Multi-Factor Scenarios

Continuing the discussion of scenario analysis begun in Section 5, we could consider the expected change in the Canadian Government forward curve conditioned on a given move in the U.S. forward curve. For illustration, we could suppose that the risk factors associated with the U.S. forward rate curve are unexpected movements of X_1, \dots, X_m basis points at each of k respective maturities, and that the scenario outcome for the forward curve shift is the vector $x = (x_1, \dots, x_m)$ of basis point changes at the respective maturities. For example, x could be the forward curve shift vector associated with the first principle component of U.S. forward curve changes.

For some other given risk factor X_k , say the unexpected change in the Canadian 5-year forward rate, we are interested in computing the expected change in X_k given the outcome $X_1 = x_1, X_2 = x_2, \dots, X_m = x_m$. Assuming joint normality of the rates, we have

$$E(X_k | X_1, X_2, X_3, \dots, X_m) = (X_1, \dots, X_m)^\top A^{-1}q,$$

where $A = \text{cov}(X_1, \dots, X_m)$ is the $k \times k$ covariance matrix of X_1, \dots, X_k , and q is the vector whose i -th element is $\text{cov}(X_i, X_m)$. Evaluating this conditional expectation at the scenario shift $(X_1, \dots, X_m) = (x_1, \dots, x_m)$, we have the desired result. One can do this for each Canadian rate to get the expected response of the Canadian forward curve. An approximation of the re-valuation of Canadian fixed-income products at this shift can be done on a delta basis. (For straight bonds, this is an easy calculation.)

Of course, our focus in this example on Canadian and U.S. forward rates is simply for illustration. We could have used equity returns, foreign exchange rates, or other risk factors. The only necessary information is the covariance matrix for the risk factors, and the scenario of concern.

D. Tail-Fatness of Jump-Diffusion Models

The calculations for tail-fatness of the jump-diffusion model considered in Section 2 are shown below for reference only.

D.1 Critical Values

We consider the return on the asset that undergoes a jump diffusion

$$\begin{aligned} S_t &= S_0 \exp(\alpha t + X_t) \\ X_t &= \beta B_t + \sum_{k=0}^{N(t)} \nu Z_k, \end{aligned} \tag{1}$$

where B is a standard Brownian motion and $N(t)$ is the number of jumps that occur by time t . Each jump νZ_k is normally distributed with mean zero and standard deviation ν . The arrival rate of jumps (Poisson) is λ . Then the total volatility σ is defined by

$$\sigma^2 = \beta^2 + \lambda\nu^2. \tag{2}$$

We are interested in the critical value, at confidence p and time horizon t , that is, $C_{p,t} = \{x : P(X_t \leq x) = p\}$. Using the law of iterated expectations and conditioning on the number of jumps,

$$\begin{aligned} P(X_t \leq x) &= \sum_{k=0}^{\infty} p(k, t) P(X_t \leq x | N(t) = k) \\ &= \sum_{k=0}^{\infty} p(k, t) N(0, \sqrt{\beta^2 t + k\nu^2}, x), \end{aligned} \tag{3}$$

where $p(k, t)$ is the Poisson probability of k arrivals within t units of time, given by

$$p(k, t) = \frac{(\lambda t)^k e^{-\lambda t}}{k!} \tag{4}$$

and $N(a, s, x)$ is the probability that a random variable that is distributed normally with mean 0 and standard deviation s has an outcome less than x .

D.2 Kurtosis

The normalized kurtosis is defined by

$$K_t = \frac{E(X_t^4)}{[Var(X_t)]^2} = \frac{E(X_t^4)}{\sigma^4 t^2}. \quad (5)$$

After tedious calculation, the numerator can be shown to be

$$\begin{aligned} E(X_t^4) &= E[(\beta B_t + \sum_{k=0}^{N(t)} Z_k)^4] \\ &= E[3(\beta^2 t + N(t)\nu^2)^2] = 3[\beta^4 t^2 + (\lambda + \lambda^2 t)\nu^4 t + 2\nu^2 \beta^2 \lambda t^2]. \end{aligned} \quad (6)$$

E. Option Portfolio Value-at-Risk

Table 3 is a summary of the moments of the simulated distribution of the “short” option portfolio described in Section 4.5, based on different models for the underlying. Specifically, in Jump-diffusion Model 1, the diffusion part is simulated using the RiskMetrics covariance matrix for July 29, 1996, while each asset jumps independently with poisson arrival of intensity $\lambda = 1$ per year. The standard deviation of jumps size for each asset is taken to be half of the corresponding RiskMetrics standard deviation. So, in this example, the total covariance matrix is not matched to that of RiskMetrics. Jump-diffusion Models 2 and 3, however, are parameterized in such a way that the total covariance matrix of the underlying assets is matched to that of RiskMetrics. Model 2 has one-fifth of its total covariance coming from jumps, and four-fifths from diffusion, while in Model 3 half of the total covariance comes from jumps, and half from diffusion. The moments for the long portfolio are of exactly the same magnitude, except for sign reversals for odd moments (mean and skewness).

Table 4 shows the estimated 1% and 0.4% values at risk (critical values) of the “predominantly short” option portfolio, designated “Portfolio S,” over time horizons of one day and two weeks. The portfolio is normalized to an initial market value of -100 dollars. The “predominantly long” option portfolio, designated “Portfolio L,” has a total market value of $\$100$. Of course, the left tail of the distribution of value changes for Portfolio L can be obtained from the right tail of Portfolio S. Table 5 shows the estimated 1% and 0.4% critical values for Portfolio L over one day or two weeks, estimated by the various approximation methods described in Section 4.

Table 3: Moments of the Simulated Distribution

Time Span	Model	Method	mean	s-dev	skewness	kurtosis
Overnight	pl. vanilla	Actual	-100.02	0.97	-0.15	3.06
		Gamma	-100.02	0.98	-0.15	3.07
		Delta	-99.88	0.97	-0.03	3.04
	jump-duffsion (Model 1)	Actual	-100.02	1.02	-0.15	3.72
		Gamma	-100.03	1.04	-0.25	4.73
		Delta	-99.86	1.02	-0.07	4.47
	jump-duffsion (Model 2)	Actual	-100.00	1.10	-8.65	233.85
		Gamma	-100.01	1.17	-11.44	343.38
		Delta	-99.87	0.98	-1.94	59.17
	jump-duffsion (Model 3)	Actual	-100.01	1.18	-7.43	103.24
		Gamma	-100.02	1.26	-8.38	122.42
		Delta	-99.87	0.99	1.02	76.31
2 Weeks	pl. vanilla	Actual	-100.22	3.76	-0.48	3.40
		Gamma	-100.28	3.90	-0.51	3.48
		Delta	-98.18	3.73	-0.06	3.03
	jump-duffsion (Model 1)	Actual	-100.23	3.96	-0.42	3.45
		Gamma	-100.39	4.21	-0.52	3.89
		Delta	-97.90	3.91	-0.04	3.17
	jump-duffsion (Model 2)	Actual	-100.26	4.37	-2.35	18.01
		Gamma	-100.38	4.87	-3.47	34.89
		Delta	-98.22	3.81	-0.10	6.29
	jump-duffsion (Model 3)	Actual	-100.27	4.69	-2.65	16.07
		Gamma	-100.42	5.25	-3.20	21.74
		Delta	-98.26	3.79	-0.17	7.99

F. Sample Statistics for Daily Returns

For reference purposes, we record in the table below some sample statistics for daily returns for the period 1986 to 1996 for a selection of equity indices, foreign currencies, and commodities. The statistics for foreign equity returns are in local currency terms. The raw price data were obtained from Datastream.

Table 4: Critical Values of the “Short Option” Portfolio

Model	Method		Overnight		2 Weeks	
			1%	0.4%	1%	0.4%
pl. vanilla	Analytical	Delta	2.27	3.27	8.70	12.54
		Gamma	2.28	3.28	9.12	13.15
	Simulation	Actual	2.41	2.80	10.48	12.18
		Gamma	2.42	2.81	10.91	12.77
		Delta	2.18	2.51	7.10	8.42
jump-diffusion (Model 1)	Analytical	Delta	2.19	3.16	8.42	12.14
		Gamma	2.20	3.17	8.81	12.70
	Simulation	Actual	2.51	3.05	10.58	12.44
		Gamma	2.60	3.10	11.63	14.03
		Delta	2.25	2.70	6.97	8.54
jump-diffusion (Model 2)	Analytical	Delta	2.28	3.29	8.74	12.59
		Gamma	2.29	3.30	9.16	13.20
	Simulation	Actual	2.15	2.53	15.57	23.14
		Gamma	2.15	2.54	17.80	26.81
		Delta	1.93	2.24	7.51	11.30
jump-diffusion (Model 3)	Analytical	Delta	2.29	3.30	8.77	12.64
		Gamma	2.30	3.31	9.19	13.25
	Simulation	Actual	2.05	8.83	18.89	25.01
		Gamma	2.05	9.54	21.85	28.57
		Delta	1.63	2.60	9.68	13.55

Shown are the annualized sample standard deviation (volatility), the sample skewness, sample normalized kurtosis, the number of days on which the return was more than 10 sample standard deviations below the mean, the number of days on which the return was more than 5 sample standard deviations below the mean, the number of days on which the return was more than 5 sample standard deviations above the mean, the number of days on which the return was more than 10 sample standard deviations above the mean, the number of standard deviations to the 0.4 percent critical value of the sample distribution, and the number of standard deviations to the 99.6 percent

Table 5: Critical Values of the “Long Option” Portfolio

Model	Method		Overnight		2 Weeks	
			1%	0.4%	1%	0.4%
pl. vanilla	Analytical	Delta	2.27	3.27	8.70	12.54
		Gamma	2.28	3.28	9.12	13.15
	Simulation	Actual	2.18	2.45	7.34	8.17
		Gamma	2.19	2.45	7.53	8.39
		Delta	2.41	2.64	10.49	11.49
jump-diffusion (Model 1)	Analytical	Delta	2.19	3.16	8.42	12.14
		Gamma	2.20	3.17	8.81	12.70
	Simulation	Actual	2.29	2.61	7.89	8.86
		Gamma	2.30	2.70	8.17	9.29
		Delta	2.56	2.93	11.18	12.60
jump-diffusion (Model 2)	Analytical	Delta	2.28	3.29	8.74	12.59
		Gamma	2.29	3.30	9.16	13.20
	Simulation	Actual	1.99	2.25	7.16	7.92
		Gamma	1.99	2.26	7.37	8.18
		Delta	2.20	2.51	11.02	14.11
jump-diffusion (Model 3)	Analytical	Delta	2.29	3.30	8.77	12.64
		Gamma	2.30	3.31	9.19	13.25
	Simulation	Actual	1.64	1.92	6.50	7.52
		Gamma	1.65	1.93	6.70	7.81
		Delta	1.87	2.37	12.52	16.20

critical value of the sample distribution.

Table 6: Sample Return Statistics for Selected Markets

Daily Return Statistics 1986-1996

Name	Vol. (Annual)	Skew	Kurtosis	< - 10 sd	< - 5 sd	> 5 sd	> 10 sd	0.4%	99.6%
S&P 500	15.9%	-4.8	110.7	1	5	3	0	-3.64	2.76
NASDAQ	15.2%	-5.1	109.7	2	2	2	0	-3.57	2.62
NYSE All Share	14.7%	-5.2	121.4	1	5	2	0	-3.54	2.68
Mexico Bolsa (Pesos)	26.3%	-0.2	7.8	0	3	4	0	-3.59	3.75
Mexico Bolsa (US\$)	32.0%	0.0	14.9	0	6	4	0	-3.98	3.44
FTSE 100	15.0%	-1.7	28.6	2	4	3	0	-3.37	2.94
FTSE All Share	13.6%	-1.9	29.1	2	5	4	0	-3.94	2.94
German DAX 30 Perf.	19.7%	-0.9	15.6	1	7	8	0	-4.25	3.37
France DS Mkt.	17.5%	-0.9	13.1	0	7	4	0	-3.72	3.14
France CAC 40	19.4%	-0.6	10.9	0	6	4	0	-3.37	3.17
Sweden Veckans Aff.	18.4%	-0.4	12.1	0	8	9	0	-4.76	3.96
Milan B.C.I.	20.3%	-0.9	13.3	1	5	1	0	-3.58	3.25
Swiss Perf.	15.6%	-2.6	32.0	3	8	3	0	-4.65	3.71
Australia All Ord.	17.1%	-7.8	198.3	1	7	3	0	-3.54	2.34
Nikkei 500	20.4%	-0.4	34.8	2	4	6	1	-3.70	3.50
Hang Seng	26.7%	-6.4	143.6	2	6	1	0	-4.08	2.93
Bangkok S.E.T.	25.0%	-0.6	9.8	0	9	5	0	-4.86	3.59
Taiwan Weighted	36.1%	-0.1	5.0	0	0	1	0	-3.08	3.20
US\$: English Pound	11.0%	-0.2	5.9	0	1	1	0	-3.60	3.20
US\$: Mexican New Peso	18.9%	-8.1	217.5	4	12	8	1	-5.15	3.20
US\$: German Mark (DM)	11.3%	-0.1	5.3	0	0	1	0	-3.40	3.40
US\$: French Franc (FF)	10.9%	0.0	5.7	0	0	2	0	-3.50	3.45
US\$: Swedish Krone (SK)	10.5%	-0.6	10.9	1	3	4	0	-3.50	3.30
US\$: Italian Lira	11.3%	-0.6	8.6	0	4	0	0	-3.60	3.30
US\$: Swiss Franc (SF)	12.6%	0.0	4.9	0	0	1	0	-3.20	3.30
US\$: Australian Dollar	9.4%	-0.7	8.1	0	4	1	0	-4.30	2.75
US\$: Japanese Yen	11.1%	0.4	8.1	0	3	5	0	-3.10	3.68
US\$: Hong Kong Dollar	0.8%	-0.6	17.4	0	11	10	0	-5.40	4.30
US\$: Thai Baht	2.4%	0.7	33.8	1	9	13	2	-4.80	4.95
US\$: Canadian Dollar	4.4%	-0.1	7.2	0	3	6	0	-3.80	3.25
Gold (First Nearby)	13.4%	-0.7	11.4	0	4	3	0	-4.24	3.44
Oil (First Nearby)	38.7%	-0.8	21.8	1	7	7	0	-4.14	4.07
Oil (Sixth Nearby)	27.5%	-0.6	15.6	1	3	3	0	-3.82	3.82

References

- Bassi, F., P. Embrechts, and M. Kafetzaki (1996). A Survival Kit on Quantile Estimation. Working paper, Department of Mathematics, ETHZ, CH-8092 Zürich, Switzerland.
- Beckers, S. (1981). Standard Deviations Implied in Option Process as Predictors of Future Stock Price Variability. *Journal of Banking and Finance* 5, 363–382.
- Beder, T. S. (1995, September–October). VAR: Seductive but Dangerous. *Financial Analysts Journal*, 12–24.
- Bekaert, G. and G. Wu (1997). Asymmetric Volatility and Risk in Equity Markets. Working paper, Graduate School of Business, Stanford University.
- Black, F. and M. Scholes (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* 81, 637–659.
- Bollerslev, T. (1986). Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics* 31, 307–327.
- Bollerslev, T., R. Chou, and K. Kroner (1992). ARCH Modeling in Finance: A Review of Theory and Empirical Evidence. *Journal of Econometrics* 52, 5–59.
- Brockwell, P. and R. Davis (1991). *Time Series: Theory and Methods* (Second ed.). New York: Springer-Verlag.
- Butler, J. S. and B. Schachter (1996). Improving Value-at-Risk Estimates by Combining Kernel Estimation with Historical Simulation. Working paper, Comptroller of the Currency, Washington DC.
- Caffisch, R. and W. Morokoff (1996). Valuation of Mortgage Backed Securities Using the Quasi-Monte Carlo Method. Working paper, Department of Mathematics, University of California, Los Angeles.
- Campa, J. and P. K. Chang (1995). The Forecasting Ability of Correlations Implied in Foreign Exchange Options. Working paper, Stern School of Business, New York University.
- Canina, L. and S. Figlewski (1993). The Informational Content of Implied Volatility. *Review of Financial Studies* 6(3), 659–681.
- Clark, P. (1973). A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices. *Econometrica* 41, 135–155.
- Cox, J. and M. Rubinstein (1985). *Options Markets*. Englewood Cliffs: Prentice-Hall.
- Day, T. and C. Lewis (1992). Stock Market Volatility and the Information Content of Stock Index Options. *Journal of Econometrics* 52, 267–287.

- Derivatives Policy Group (1995, March). A Framework for Voluntary Oversight of the OTC Derivatives Activities of Securities Firm Affiliates to Promote Confidence and Stability in Financial Markets. New York.
- Derman, E. and I. Kani (1994, February). Riding on the Smile. *Risk*, 32–39.
- Dimson, E. (1995). CAD Fails Stress Test. *Risk* 8(12).
- Duffie, D. and S. Gray (1995). Volatility in Energy Prices. In *Managing Energy Risk*, Chapter 2. Risk and Enron.
- Dupire, B. (1992). Arbitrage Pricing with Stochastic Volatility. Working paper, Société Générale, Paris.
- Dupire, B. (1994, January). Pricing with a Smile. *Risk*, 18–20.
- Durrett, R. (1991). *Probability: Theory and Examples*. Wadsworth.
- Engle, R. (1982). Autoregressive Conditional Heteroskedasticity with Estimates of The Variance of United Kingdom Inflation. *Econometrica* 50, 987–1008.
- Estrella, A. (1994). Taylor, Black, and Scholes: Series Approximations and Risk Management Pitfalls. Working paper, Federal Reserve Bank of New York.
- Estrella, A., D. Hendricks, J. Kambhu, S. Shin, and S. Walter (1994). The Price Risk of Options Positions: Measurement and Capital Requirements. *Federal Reserve Bank of New York Quarterly Review*, 27–43.
- Fournie, E., J. Lebuchoux, and N. Touzi (1996). Small Noise Expansion and Importance Sampling. Working paper, Ecole Normale Supérieure and Université de Paris IX Dauphine.
- Froot, K. and J. Stein (1995). Risk Management; Capital Budgeting and Capital Structure Policy for Financial Institutions: An Integrated Approach. Working paper, Harvard Business School.
- Froot, K. A. (1993). Currency Hedging Over Long Horizons. Working paper, Harvard and NBER.
- Garman, M. and M. Klass (1980). On the Estimation of Security Volatilities from Historical Data. *Journal of Business* 53, 67–78.
- Global Derivatives Study Group, G30 (1993, July). Derivatives: Practices and Principles.
- Gray, S. (1993). Regime-Switching Models of Volatility. Working paper, Graduate School of Business, Stanford University.
- Group of Thirty (1994, December). Derivatives: Practices and Principles — Follow-up Surveys of Industry Practice. Washington DC.

- Grundy, B. D. and Z. Wiener (1996). The Analysis of VAR, Deltas and State Prices: A New Approach. Working paper, Rodney L. White Center for Financial Research, The Wharton School.
- Hamilton, J. (1990). Analysis of Time Series Subject to Changes in Regime. *Journal of Econometrics* 45, 39–70.
- Hamilton, J. (1994). *Time Series Analysis*. Princeton University Press.
- Harrison, M. and D. Kreps (1979). Martingales and Arbitrage in Multiperiod Securities Markets. *Journal of Economic Theory* 20, 381–408.
- Harvey, A., E. Ruiz, and N. Shepard (1992). Multivariate Stochastic Variance Models. Working Paper LSE Financial Markets Group Discussion Paper, London School of Economics.
- Harvey, A. and N. Shepard (1993). The Econometrics of Stochastic Volatility. Working Paper LSE Financial Markets Group Discussion Paper Number 166, London School of Economics.
- Hendricks, D. (1995). Evaluation of Value-at-Risk Models Using Historical Data. Working paper, Federal Reserve Bank of New York.
- Heston, S. (1993). A Closed-Form Solution for Options with Stochastic Volatility, with Applications to Bond and Currency Options. *Review of Financial Studies* 6, 327–344.
- Heynen, R. and H. Kat (1993). Volatility Prediction: A Comparison of the Stochastic Volatility, GARCH(1,1) and EGARCH(1,1) Model. Working paper, Department of Operations Research, Erasmus University.
- Hull, J. (1993). *Options, Futures, and Other Derivative Securities, Second Edition*. Englewood Cliffs: Prentice-Hall.
- Hull, J. and A. White (1987). The Pricing of Options on Assets with Stochastic Volatilities. *Journal of Finance* 42, 281–300.
- Jackson, P., D. J. Maude, and W. Perraudin (1995). Capital Requirements and Value-at-Risk Analysis. Working paper, Bank of England.
- Jackwerth, J. and M. Rubinstein (1996). Implied Binomial Trees. *Journal of Finance* 49, 771–818.
- Jamshidian, F. and Y. Zhu (1997). Scenario Simulation: Theory and Methodology. *Finance and Stochastics* 1 (1), 43–68.
- Johnson, H. and H. Shanno (1987). The Pricing of Options when the Variance is Changing. *Journal of Financial and Quantitative Analysis* 22, 143–151.
- Jorion, P. (1989). On Jump Processes in the Foreign Exchange and Stock Markets. *Review of Financial Studies* 4, 427–445.

- Jorion, P. (1995). Predicting Volatility in the Foreign Exchange Market. *Journal of Finance* 50(2), 507–528.
- Kunita, H. (1990). *Stochastic Flows and Stochastic Differential Equations*. Cambridge University Press.
- Kupiec, R. H. and J. M. O'Brien (1995). The Use of Bank Trading Risk Models for Regulatory Capital Purposes. Working paper, Division of Research and Statistics, Division of Monetary Affairs, Federal Reserve Board, Washington, D.C.
- Lamoureux, C. and W. Lastrapes (1993). Forecasting Stock-Return Variance: Toward an Understanding of Stochastic Implied Volatilities. *Review of Financial Studies* 6, 293–326.
- Linsmeier, T. J. and N. D. Pearson (1996). Risk Management: An Introduction to Value at Risk. Working paper, University of Illinois at Urbana-Champaign.
- Littlejohn, G. and A. Fry (1996). *Systemic Risk Facing the World's Financial Institutions*. FT Financial Publishing, London.
- Marshall, C. and M. Siegel (1996). Value at Risk: Implementing a Risk Measurement Standard. Working paper.
- Merton, R. and A. Perold (1993). Theory of Risk Capital in Financial Firms. *Journal of Applied Corporate Finance* 5, 16–32.
- Mori, A., M. Ohsawa, and T. Shimizu (1996). Calculation of Value at Risk and Risk/Return Simulation. Working paper, Institute for Monetary and Economic Studies, Bank of Japan.
- Nelson, D. (1990). ARCH Models as Diffusion Approximations. *Journal of Econometrics* 45, 7–39.
- Nelson, D. (1991). Conditional Heteroskedasticity in Asset Returns: A New Approach. *Econometrica* 59, 347–370.
- O'Connell, P. (1996). The Overvaluation of Purchasing Power Parity. Working paper, Department of Economics, Harvard University. forthcoming in *Journal of International Economics*.
- Page, M. and E. Feng (1995, February). Measuring the Risk of a Portfolio of Options on Several Exchange Rates. Working paper, Susquehanna Investment Group, Philadelphia.
- Parkinson, M. (1980). The Extreme Value Method for Estimating the Variance of The Rate of Return. *Journal of Business* 46, 434–452.
- Pastorello, S., E. Renault, and N. Touzi (1993). Statistical Inference for Random Variance Option Pricing. Working paper, Université de Toulouse, France.

- Phelan, M. J. (1995). Probability and Statistics Applied to the Practice of Financial Risk Management: the Case of JP Morgan's RiskMetrics. Working paper, The Wharton Financial Institution Center.
- Renault, E. and N. Touzi (1992). Stochastic Volatility Models: Statistical Inference from Implied Volatility. Working paper, Université de Toulouse, France.
- Rogers, C. and L. Satchell (1991). Estimating Variance from High, Low, and Closing Prices. *Annals of Applied Probability* 1, 504–512.
- Rubinstein, M. (1994). Implied Binomial Trees. Working paper, Haas School of Business, University of California, Berkeley.
- Sakata, S. and H. White (1996). High Breakdown Point Conditional Dispersion Estimation with Application to S&P 500 Daily Returns Volatility. Working paper, University of California, San Diego.
- Schoenmakers, J. and A. Heemink (1996). Fast Valuation of Financial Derivatives. Working paper, Delft University of Technology, The Netherlands.
- Scott, L. (1987). Option Pricing when the Variance Changes Randomly: Theory, Estimation, and an Application. *Journal of Financial and Quantitative Analysis* 22, 419–438.
- Scott, L. (1992). The Information Content of Prices in Derivative Security Markets. *IMF Staff Papers* 39, 596–625.
- Stoll, H. and R. Whaley (1993). *Futures and Options: Theory and Applications*. Cincinnati: Southwestern.
- Wiggins, J. (1987). Option Values Under Stochastic Volatility: Theory and Empirical Estimates. *Journal of Financial Economics* 19, 351–372.
- Willard, G. (1996). Calculating Prices and Sensitivities for Path-Dependent Derivative Securities in Multi-Factor Models. Working paper, Department of Finance, Washington University, St. Louis.