

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.

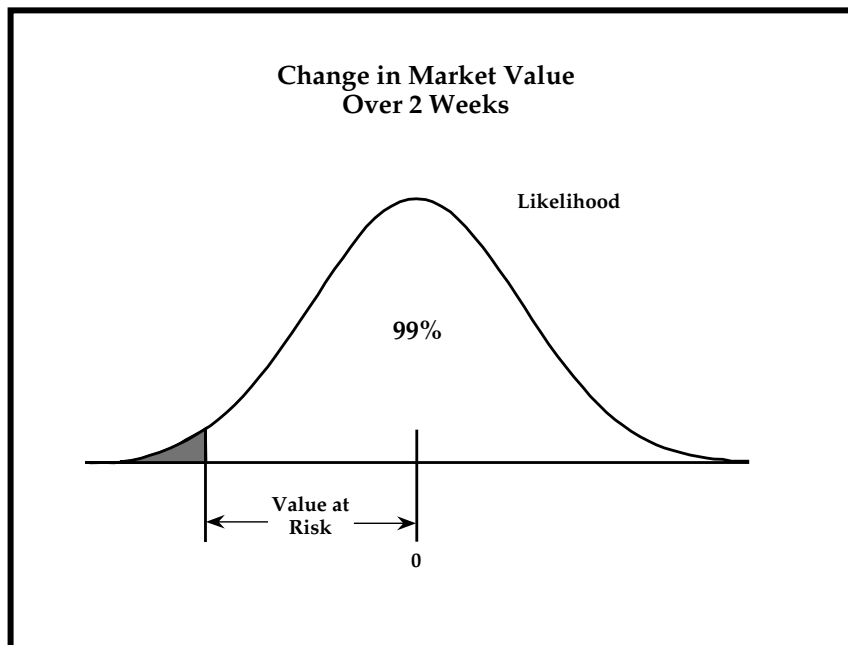


Figure 1: Value at Risk (DPG Standard)

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.



Source: Datastream
Daily Excess Returns, 10-Year Basis


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

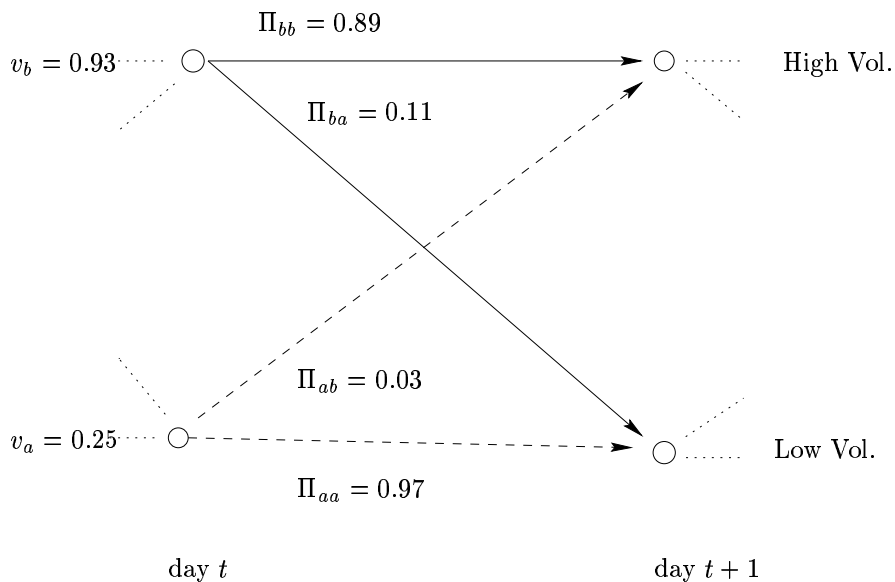


Figure 8: Regime-Switching Volatility Estimates for Light Crude Oil

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

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

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].

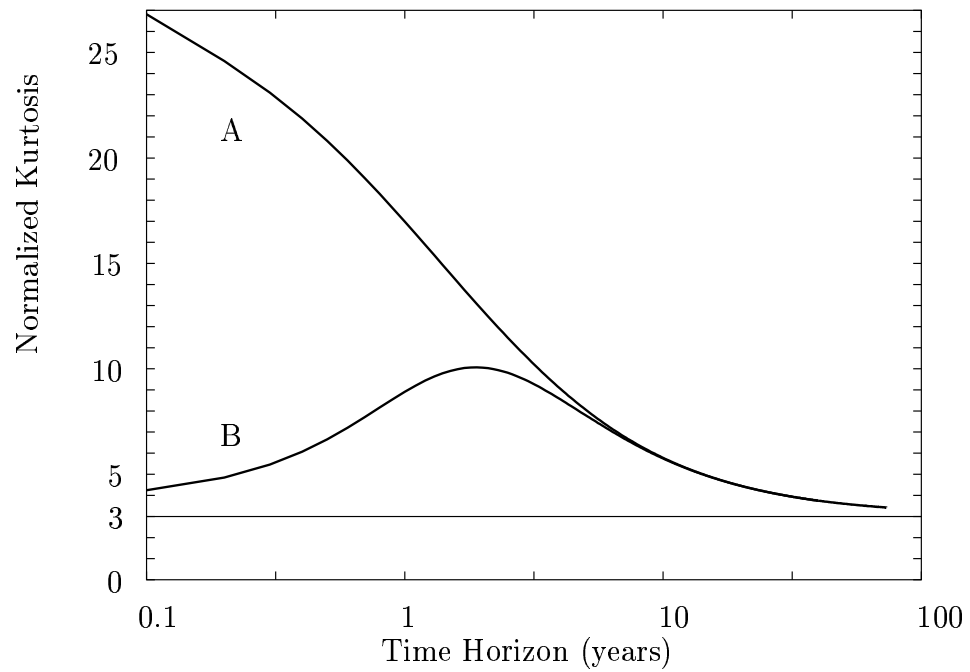
²⁹We are grateful to Ken Froot for pointing this out. We can rely on the fact that, over time intervals of “large” length, the volatilities at the beginning and end of the intervals are “essentially” independent, in the sense of the central limit theorem for recurrent Markov processes.

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].



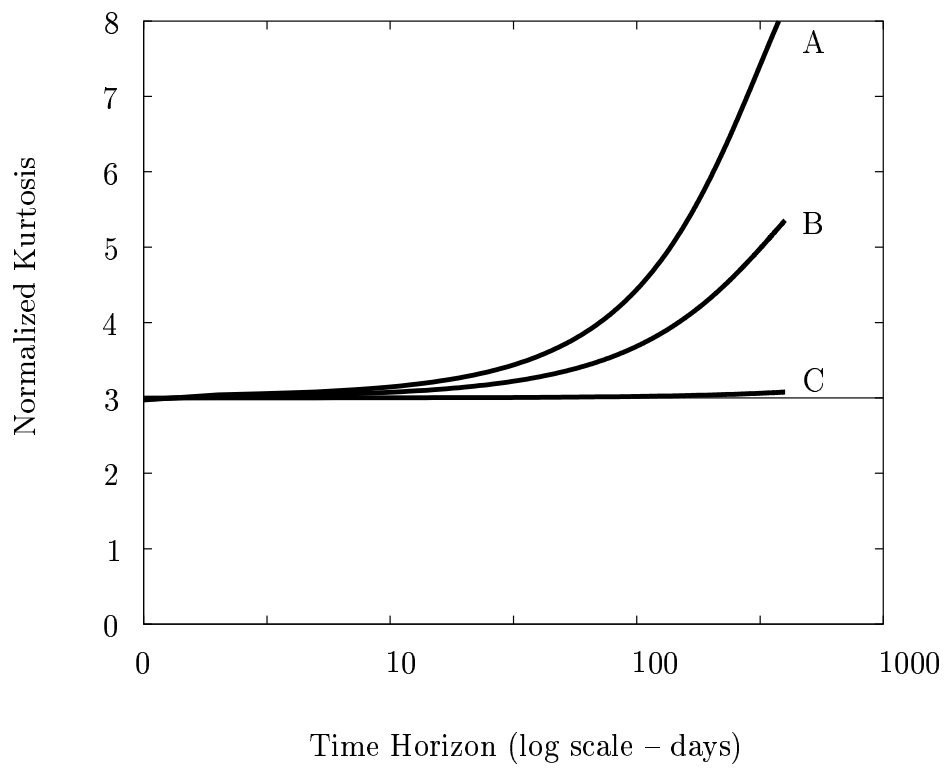
A: High Initial Deterministic Volatility
 B: Steady-State Average Initial Deterministic Volatility
 C: Low Initial Deterministic Volatility.

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.