9

Value-at-Risk

During the last few years there have been many changes in the way financial institutions evaluate risk. Improvements are continuously being sought to relate the regulatory capital that must be available to the underlying risks that a firm takes. Hence regulations have played a major role in the development of risk measurement techniques. The Basel Committee on Banking Supervision now recommend two types of models for measuring market risk on a daily basis; details of these models are given in the Basel Accord Amendment of 1996.

These internal models have become industry standards for measuring risk not only for external regulatory purposes, but also for internal risk management and control. One approach is to quantify the maximum loss over a large set of scenarios for movements in the risk factors over a certain time horizon. Another approach is to weight scenarios with probabilities and to assess the level of loss that has some low probability of being exceeded over a fixed time horizon.[1] This measure is called the portfolio value-at-risk (VaR).[2] Both approaches assume the portfolio is not managed during the fixed time horizon.

This chapter is about the VaR approach to measuring market risk, for capital requirements and for internal risk management. The first section looks at the current risky environment in financial markets and gives a brief overview of the developments in risk capital regulation during the last few years. It outlines the limitations with current regulations and discusses how they are being addressed by the new proposals from the Basel Committee that are currently under consultation (Basel 2). The theoretical and practical advantages and limitations of VaR as a measure of portfolio risk are described in §9.2. Here some of the traditional 'sensitivity-based' risk measures and some of the new alternative risk measures are also described.

The next two sections look at the VaR models that have emerged as industry standard during the last few years. The covariance VaR model is described in

[1] In our uncertain environment loss is a random variable, so it is only possible to make *probabilistic* statements about the loss from a portfolio.

[2] Or more precisely, the *market VaR* of the portfolio. Although VaR was introduced in the context of market risk, recently the context has been extended to credit VaR and operational VaR.

§9.3, with examples of its application to different types of linear portfolios. Option portfolios have too many non-linear characteristics to be measured by a covariance VaR model; instead one of the simulation methods described in §9.4 will normally be applied.

Section 9.5 describes the methods that are normally used to validate a VaR model and the impact of backtesting results on risk capital requirements. This section also deals with the sensitivity of VaR estimates to assumptions about model parameters, something that deserves careful investigation by both internal and external risk control functions. The chapter concludes by reviewing the methods by which market VaR estimates of current or potential positions can be complemented by stress testing and scenario analysis to control the impact of extreme market movements.

There are three workbooks on the CD to supplement this chapter: **covariance VaR**, **historical VaR** and **Monte Carlo VaR**. The reader can use these to specify portfolios, compute VaR and perform stress tests.

9.1 Controlling the Risk in Financial Markets

Are financial markets more risky than ever? They are certainly more volatile. One of the reasons for this is the increased ability of financial institutions to create leverage. Hedge funds can take extraordinarily highly leveraged positions because their models are supposedly designed to diversify most of the risks. New derivative products are continually being structured to allow companies and banks to increase leverage in more ingenious ways than ever.

Financial activity is unstable and risky by its very nature. As new markets are opened and new products are developed, market liquidity may be insufficient to accommodate our growing appetite for leverage. In young markets for new products sometimes it is just not possible to understand the risks completely. Even when proper pricing models have been developed they may be so new that they are only understood by 'quants'. In established markets that are better understood there can always be an event that has never been experienced before.

Volatility in itself does not imply risk. There does not seem to be a strong connection between crises in financial markets and the risk to the real economy

But volatility in itself does not imply risk. What do we mean by risk anyway? Financial risk should be perceived on three levels: the risks to individual consumers, the risks to a firm, along with its shareholders and investors, and the risks to the markets as a whole. Consider one of the most volatile financial experiences of the twentieth century — the global stock market crash of 1987. All three levels of risk were affected by this event. Firms and individuals suffered immediate and direct financial loss, and the knock-on effect of the crash undermined the fundamental stability of the world's economy. In more recent years there have been a number of other crises in financial markets around the globe, but it is mainly in

Japan that the financial market crisis has had a devastating effect on the economy. Elsewhere there does not seem to be a strong connection between crises in financial markets and the risk to the real economy. Thus the main sufferers are the investors. For example, in 1998 the Long Term Capital Management (LTCM) company, whose shareholders included some of the most successful Wall Street traders and most respected Nobel Prize winners, had debts of the order of $100 billion and extraordinarily highly leveraged positions amounting to approximately $1 trillion. Unfortunately, their mathematical models (which had been beating the markets consistently for a number of years) could not cope with the totally unexpected Russian debt crisis and did not prevent them from huge illiquid exposures in other markets. The US Federal Reserve bailed them out at a cost of $3.5 billion. But most of the losses were borne by LTCM's shareholders and investors.

During the last decade many banks and securities houses have also experienced very large losses. These have often been attributed to fraud, bad management or poor advice, and it is hoped that the institutions concerned will not repeat the same mistakes. But the scale of these individual losses should be put into perspective, when it is considered that the value of large multinational corporations could change by similar amounts in the course of a normal day's trading on the stock market.

The scale of these individual losses should be put into perspective, when it is considered that the value of large multinational corporations could change by similar amounts in the course of a normal day's trading on the stock market

The forces behind the measurement and control of financial risk are, therefore, extremely strong. There are internal forces to gain the optimal return on capital (where capital is risk-based) and to ensure the survival of a firm as a whole. There are external forces that are driven by competition, by the enormous growth in the risk management industry and by the increased volatility of financial markets, with new products that enable participants to increase leverage to very high levels. And there are regulatory forces to promote fair competition between firms, protect the solvency of financial institutions, and control systemic risk.

9.1.1 The 1988 Basel Accord and the 1996 Amendment

Three main tools are available to regulators for the measurement and control of financial risk: minimum risk capital requirements; inspections and reporting requirements; and public disclosure and market discipline. This section concerns only the first of these three *pillars of regulation*, the minimum risk capital requirements that are imposed by the regulatory body. The discussion focuses on the Basel Accord of 1988 and its 1996 Amendment which, though only formally adopted by the G10 countries, have had an enormous international influence.

The 1988 Basel Accord linked minimum capital standards to credit risk, and this was extended to market risks in the 1996 Amendment. The basic principle of the 1996 Amendment is to measure regulatory capital by a minimum solvency condition, the *Cooke ratio*. This is the ratio of eligible capital to risk-weighted assets, where eligible capital is the sum of tier 1 (core capital), tier 2 (complementary capital) and tier 3 (sub-supplementary capital), and risk-weighted assets are the sum of a credit risk capital requirement (CRR) and a market risk capital requirement (MRR).

Regulators currently favour placing transactions in the banking book, where securities attract a CRR

Credit risk requirements apply to all positions except the equity and debt positions in a trading book, foreign exchange and commodities. At the time of writing the CRR is simply calculated as a percentage of the nominal (for on-balance-sheet positions) or a credit equivalent amount (for off-balance-sheet positions).[3] Regulators currently favour placing transactions in the banking book, where securities attract a CRR (the banking book CRR is very broadly defined to compensate for the lack of MRR). In particular, the current method used to calculate the CRR provides little incentive to diversify a portfolio or to employ other risk mitigation techniques. However, the Basel 2 Accord that is currently under discussion will define new rules for calculating the CRR, for implementation by G10 banks in 2005.

The MRR applies to all on- and off-balance-sheet positions in a trading book. It is necessary to mark positions to market, which is easily done in liquid markets where bid and offer prices are readily available. Otherwise it is acceptable to mark the portfolio to a model price, where the value of a transaction is derived from the value of liquid instruments, or to make a liquidity adjustment that is based on an assessment of the bid–offer spread.

The standardized approach is to sum the MRRs of positions in four different categories or *building blocks*: equities, interest rates, foreign exchange and gold, and commodities. In each block the total MRR is the sum of general and specific risk requirements that are percentages of the net and gross exposures, respectively; these percentages depend on the building block.[4] The minimum solvency ratio that has been set in the 1988 Accord is 8%. That is, the eligible (tier 1, tier 2 and tier 3) capital of a firm must be at least 8% of the sum of the MRR and CRR. In this way the regulator protects the solvency of a firm by tying its total risk exposure to its capital base.

9.1.2 Internal Models for Calculating Market Risk Capital Requirements

The 1996 Amendment outlined an alternative approach to measuring the MRR, which is to use an internal model to determine the total loss to a firm when netted over all positions in its trading book (§9.1.2). Internal models for measuring the MRR must determine the maximum loss over 10 trading days at

[3] This percentage is determined by the type of counterparty.

[4] For example, in equities the general MRR is 8% of net exposure and the specific MRR is often 8% of gross exposure.

a 99% confidence level. They are subject to some strong qualitative and quantitative requirements. Qualitative requirements include the existence of an independent risk management function for audit and control, a rigorous and comprehensive stress testing program on the positions of the firm, and on the IT and control side the MRR model must be fully integrated with other systems. Quantitative requirements include frequent estimation of model parameters, separate assessment for the risks of linear and non-linear portfolios, rigorous model validation techniques, the use of a minimum number of risk factors and a minimum length of historic data period.

Although these models can be based entirely on scenario analysis (§9.6), many firm assess their MRR using a VaR model. VaR has been defined as *the loss (stated with a specified probability) from adverse market movements over a fixed time horizon, assuming the portfolio is not managed during this time*. So VaR is measured as a lower percentile of a distribution for theoretical profit and loss that arises from possible movements of the market risk factors over a fixed risk horizon. To see this, first note that the loss (or profit) for a portfolio that is left unmanaged over a risk horizon of h days is

$$\Delta_h P_t = P_{t+h} - P_t.$$

In other words, $\Delta_h P_t$ is the forward-looking h-day theoretical (or 'unrealized') P&L, that is, the P&L obtained by simply marking the portfolio to market today and then leaving it unchanged and marking it to market again at the risk horizon. We do not know exactly how the underlying risk factors are going to move over the next h days, but we do have some idea. For example, we might expect that historical volatilities and correlations would remain much the same. The possibilities for movements in risk factors can be summarized in a (multivariate) distribution, and this in turn will generate a distribution of $\Delta_h P_t$, as each set of possible values for the risk factors at the risk horizon are entered into the pricing models for the portfolio, weighted by their joint probabilities.

The significance level of VaR, that is, the probability that is associated with a VaR measurement, corresponds to the frequency with which a given level of loss is expected to occur. Thus a 5% 1-day VaR corresponds to a loss level that one expects to exceed, in normal market circumstances, one day in 20. And a 1% 1-day VaR is the loss level that might be seen one day in 100.[5] Now the definition of VaR above can be rephrased as follows: *the 100α% h-day VaR is that number x such that the probability of losing x, or more, over the next h days equals 100α%* — in mathematical terms,

the $100\alpha\%$ h-day VaR is that number x such that $\text{Prob}(\Delta_h P_t < -x) = \alpha$.

Sometimes we use the notation $\text{VaR}_{\alpha,h}$ to emphasize the dependency of the VaR measurement on the two parameters α, the significance level, and h, the holding period. Thus,

[5] These are also called the 95% and 99% VaR measures, but no confusion should arise.

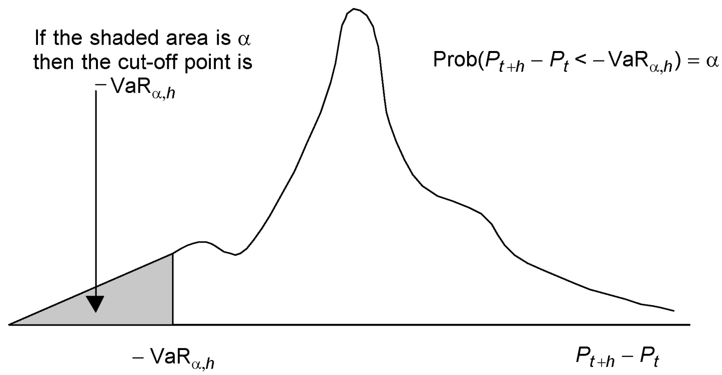

Figure 9.1 The P&L density and value-at-risk.

Table 9.1: VaR tables

Holding period, h	Significance level, α		
	0.01	0.05	0.1
1			
2		← VaR increases	
. . .		↓	
10			

$$\mathrm{Prob}(\Delta_h P_t < -\mathrm{VaR}_{\alpha,h}) = \alpha \tag{9.1}$$

is a mathematical statement that is equivalent to saying that the 100α% h-day VaR measurement x is the lower α quantile of the unrealized P&L distribution, as depicted in Figure 9.1.

100α% h-day VaR measurement x is the lower α quantile of the unrealized P&L distribution

Table 9.1 shows how VaR varies with the choice of significance level and holding period. A single VaR model will generate a whole table of VaR estimates, that will increase as the significance level increases (that is, as α decreases). From Figure 9.1 it is evident that the percentile that cuts off 100α% of the area under the density will move to the left as the area α decreases. The VaR will also increase as the holding period increases because the uncertainty in P&L will generally increase with holding period. In some cases the square-root-of-time rule is employed (§3.3) so that the h-day VaR is simply taken as $\sqrt{h}$ times the 1-day VaR. If the square-root-of-time rule is not invoked then the VaR measure is simply the lower percentile of the historical h-day P&L distribution or, when a covariance matrix is used, it is the VaR based on the h-day covariance matrix.[6]

[6] When a covariance matrix is used in the VaR model, a 1-day covariance matrix is transformed to an h-day covariance matrix by multiplying each element in the matrix by h. In linear portfolios this is equivalent to multiplying the 1-day VaR by $\sqrt{h}$; but it is not so for option portfolios, where the square-root-of-time rule should not be used. More details are given in the help sheets for the VaR models on the CD.

The regulatory MRR is a multiple of the average of the last 60 days' 1% 10-day VaR estimates when netted across the whole firm, or the previous day's VaR estimate, whichever is greater (§9.5.1). Normally VaR models produce more realistic risk capital measurements than those obtained using the standardized approach. In fact for well-diversified or well-hedged portfolios capital requirement savings may exceed 50%. VaR models have many advantages and considerable potential for internal and external risk management and control, but there are also a number of disadvantages with their use, and these are discussed in §9.2.

9.1.3 Basel 2 Proposals

There are number of problems with the current framework for measuring risk capital. For example, the measurement of MRRs is quite complex, whereas the measurement of credit risk is still quite crude and operational risk is totally ignored. And even MRRs only focus on a few short-term risks with a single criterion of limited usefulness: market risks in the banking book are ignored. There is also no proper integration of the separate measures of market and credit risk: the two measures are simply summed, as if they were perfectly correlated (§9.3.5). But, if anything, market and credit risks will be negatively correlated. For example, with a simple instrument such as a swap, one is either owed money and therefore subject to credit risk, or one owes money and therefore subject to market risk.

There is also no proper integration of the separate measures of market and credit risk: the two measures are simply summed, as if they were perfectly correlated

In 1999 and 2001 the Basel Committee prepared three consultative papers that aimed to address some of these problems. The intention of Basel 2 is to widen the scope of its regulatory framework to cope with more sophisticated institutions and products and to cover a broader range of risks. In particular, for the first time operational risk will be included in risk capital requirements. More advanced models for measuring CRRs will be introduced, that allow better credit quality differentiation and risk mitigation techniques. Another proposal that will have a great impact is that market risk capital will be required to cover positions in the banking book. Current expectations are that the Basel 2 Amendment will be implemented in stages, starting in 2005. More details can be found on www.bis.org, www.isda.org and www.bba.org.uk.

9.2 Advantages and Limitations of Value-at-Risk

Since regulators have imposed minimum capital requirements to cover market risks that are based on internal models, VaR has become the ubiquitous measure of risk. This has many advantages. VaR can be used to compare the market risks of all types of activities in the firm, and it provides a single measure that is easily understood by senior management. The VaR concept can be extended to other types of risk, notably credit risk and operational risk. It takes into account the correlations and cross-hedging between various asset categories or risk factors, and it can be calculated according to a number of different methods. VaR may also take account of specific risks by including individual equities among risk factors or including spread risk for bonds. And it may be calculated separately by building block, although without assessed correlations between building blocks, VaR measurements are simply added.

There is a danger that VaR calculations may be seen as a substitute for good risk management

There are, however, many disadvantages with the use of VaR. It does not distinguish between the different liquidities of market positions, in fact it only captures short-term risks in normal market circumstances. VaR models may be based on unwarranted assumptions, and some risks such as repo costs are ignored. The implementation costs of a fully integrated VaR system can be huge and there is a danger that VaR calculations may be seen as a substitute for good risk management. Furthermore, in the course of this chapter we shall see that VaR measures are very imprecise, because they depend on many assumptions about model parameters that may be very difficult to either support or contradict.

9.2.1 Comparison with Traditional Risk Measures

The traditional measures of risk for a fixed income portfolio that can be represented by a cash-flow map, such as a portfolio of bonds or loans, are based on sensitivity to movements in a yield curve. For example, the standard *duration* measure is a maturity-weighted average of the present values of all cash flows.[7] Another traditional measure of yield curve risk is the *present value of a basis point move* (PVBP), the change in present value of cash flows if the yield curve is shifted up by 1 basis point.[8]

The traditional risk measures of an equity portfolio are based on the sensitivities to the risk factors, that is the portfolio 'betas'.[9] However, in §8.1 the market risk of an equity portfolio was attributed to *three* sources: the variances (and covariances) of the underlying risk factors, the sensitivities to these risk factors, and the specific variance or residual risk. Therefore the traditional 'beta' risk measure relates only to the undiversifiable part of the risk: the risk that cannot be hedged away by holding a large and diversified portfolio. The beta ignores the risk arising from movements in the underlying risk factors and the specific risks of a portfolio. One of the major advantages of VaR is that it does not ignore these other two sources of risk.

---

[7] The basic measure of duration is $\Sigma(PV_t \times M_t)/\Sigma(PV_t)$. It is a measure of interest rate risk because if interest rates rise the present value of cash flows will decrease, but the income from reinvestment will increase, and the duration marks the 'break-even' point in time where the capital lost from lower cash flows has just been recovered by the increased reinvestment income.

[8] There are 10^4 basis points in 100%, thus to convert a cash flow to PVBP terms, a simple approximation is to multiply the cash flow amount by t and then divide the result by 10^4. To see this note that the interest rate sensitivity is $-\partial \ln P/\partial r$. For a single cash flow of \$1 at time t, $P = e^{-rt}$ so $-\partial \ln P/\partial r = t$.

[9] And risk measures in option portfolios have also been based purely on sensitivities. The main risk factors in options portfolios are the price and the volatility of the underlying; the main sensitivities to these risk factors, delta, gamma and vega, are defined in §2.3.3.

Traditional risk measures cannot be compared across the different activities of the firm. For example, even though they both measure first-order changes in the portfolio as a result of movements in the underlying, the duration of a bond portfolio (which is measured in months or years) cannot be compared with the delta of an options portfolio. For comparative purposes the duration or delta can be multiplied by a notional amount and a variation in the underlying risk factor (interest rates or share price). Nevertheless it is difficult to assess which activities are taking the most risk using these sensitivity-based measures, and to allocate capital accordingly. One of the main advantages of VaR is that it takes into account the volatilities and correlations of risk factors, so it is comparable across different asset classes.

Sensitivity-based risk measures have a number of limitations. They only make sense within each trading unit, and they cannot be compared across different activities to see which area has the most risk. They cannot be aggregated to give an overall exposure across all products and currencies. They do not indicate how much could be lost, either in normal circumstances or under extreme events.

9.2.2 VaR-Based Trading Limits

Long-term capital allocation between different asset classes and different activities is normally addressed using risk-adjusted performance measures (§7.2.3). On the other hand, short-term trading limits are based only on risk and not on returns: in the short term returns can be assumed to be nearly zero. Sensitivity-based trading limits have been used for some time, but more recently there have been considerable internal and regulatory forces to implement trading limits that are based on VaR. VaR provides a risk measure that focuses on the profit and loss from different activities in the firm taken together. It is therefore natural that a firm would seek to use VaR in a unified framework for allocating capital, not just to satisfy regulatory purposes, but also to allocate capital among the different activities of the firm.

However, it is a very difficult task to replace the traditional sensitivity-based trading limits with trading limits based on VaR. Not only are traders thinking in terms of sensitivities, they have the software to calculate immediately the impact of a proposed trade on their limit. In contrast, real-time VaR systems for large complex portfolios are simply not available, unless many approximations are made (§9.4.3). Therefore, if trading limits are based on VaR, the trader will need to use both VaR- and sensitivity-based systems, at least initially, so that he knows what are the implications of a trade from all perspectives.

Real-time VaR systems for large complex portfolios are simply not available, unless many approximations are made

9.2.3 Alternatives to VaR

As a measure of risk, VaR has some limitations. First, like volatility, it is affected by 'good' risk as well as 'bad' risk. Following Dembo and Freeman

Table 9.2: Good risk, bad risk, VaR and volatility

Return:	Probability of that return in portfolio *A*	Probability of that return in portfolio *B*
−10	0.01	0.01
−7.5	0.04	0.04
−5	0.05	0.25
−2.5	0.1	0.25
0	0.5	0.3
2.5	0.15	0.1
5	0.15	0.05
Expected return	0.225	−1.775
s.d.	3.124	3.128
1% VaR	−10	−10
5% VaR	− 7.5	− 7.5

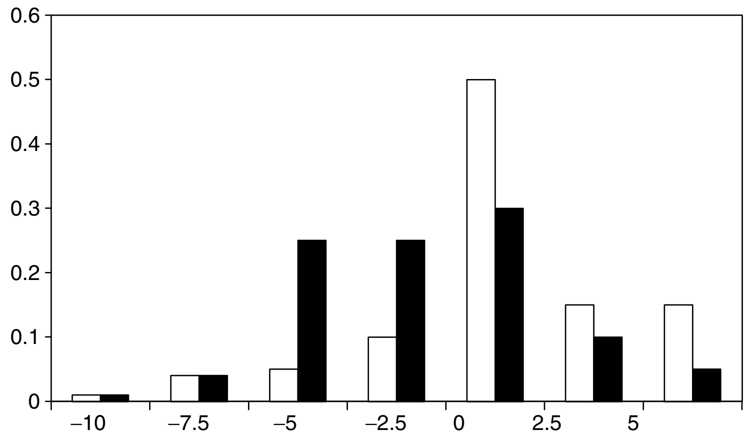

Figure 9.2 Which is more risky?

(2001), a simple example to illustrate this is given. Consider the distributions of returns for two portfolios *A* and *B* that are shown in Table 9.2 and Figure 9.2. Portfolio *A* has a greater chance of positive returns, and this is reflected in the fact that the expected return for portfolio *A* is larger than it is for portfolio *B*. However, if one only looks at the standard risk measures, these tell the same story for both portfolios. That is, the volatility (standard deviation), the 5% VaR and the 1% VaR are identical for both portfolios. Even though much of the variation in the return to portfolio *A* is concentrated on the up side, so that *A* has more 'good' risk than *B*, this is not revealed by looking at the volatility or the VaR of these portfolios.

Downside risk measures are based on the returns that fall short of a benchmark return. For example, the *semi-variance* operator introduced by Markovitz (1959) measures the variance of all returns that are less than the expected return:

$$SV = E((\min(0,\ R - E(R)))^2). \tag{9.2}$$

More generally, $E(R)$ may be replaced by any benchmark return B that can be time-varying or fixed. Dembo and Freeman (2001) advocate the use of *regret* as a measure of downside risk. The regret operator is defined as

$$\text{Regret} = -E(\min(0,\ R - B)). \tag{9.3}$$

This has the same form as the pay-off to a put option with strike equal to the benchmark return. Regret therefore has the intuitive interpretation of an insurance cost — the cost of insuring the downside risk of a portfolio. Regret will easily distinguish 'good' risk from 'bad' risk. For example, if the benchmark return $B = 0$ then the regret of portfolio A in the example above is 0.9, whereas the regret of portfolio B is 2.275.

Regret will easily distinguish 'good' risk from 'bad' risk

Artzner *et al.* (1997) have used an axiomatic approach to the problem of defining a satisfactory risk measure. They set out certain attributes that one should reasonably require of any risk measure, and call risk measures that satisfy these axioms 'coherent'. A *coherent risk measure* ρ assigns to each loss X a risk measure $\rho(X)$ such that the following conditions hold.

1. Risk is *monotonic*: if $X \geqslant Y$ then $\rho(X) \geqslant \rho(Y)$.
2. Risk is *homogeneous*: $\rho(tX) = t\rho(X)$ for $t > 0$.
3. Risk-free condition: $\rho(X + nr) = \rho(X) - n$, where r is the risk-free rate.
4. Risk is *sub-additive*: $\rho(X + Y) \leqslant \rho(X) + \rho(Y)$.

These attributes guarantee that the risk function is convex; this corresponds to risk aversion (§7.2.4).[10] The risk-free condition ensures that if an amount n of the riskless asset is added to the position then the risk measure will be reduced by n; thus capital requirements will be reduced accordingly. The last attribute is very important: it ensures that the total risk is no more than the sum of the risks of individual positions and without this there would be no incentive to diversify portfolios.

One of the great disadvantages of VaR is that it is not sub-additive

VaR is not a coherent risk measure because it does not necessarily satisfy axiom 4. The violation of axiom 4 has serious consequences for risk management, since axiom 4 allows decentralized calculation of the risks from different positions in the firm and the sum of these individual risk measures will be conservative (over-) estimate of the total risk. In a structure of limits and sub-limits for different activities or individual traders, axiom 4 implies that the respect of sub-limits will guarantee the respect of global limits. Therefore one of the great disadvantages of VaR is that it is not sub-additive.

Artzner *et al.* (1999) have introduced a new risk measure called *conditional VaR* that is a coherent risk measure, and has a simple relation to the ordinary VaR. Ordinary VaR corresponds to a lower percentile of the theoretical P&L distribution, a threshold level of loss that cuts off the lower tail of the

[10] That is, $\rho(tX + (1-t)Y) \leqslant t\rho(X) + (1-t)\rho(Y)$.

distribution. Conditional VaR is the expected loss given that the loss has exceeded the VaR threshold. That is:[11]

$$\text{Conditional VaR} = E(X|X > \text{VaR}). \tag{9.4}$$

A simple example will illustrate the difference between VaR and conditional VaR. Suppose that 100 000 simulations of the risk factors at the risk horizon are used to generate a P&L distribution such as that in Figure 9.1, and consider the 1000 largest losses that are simulated. Then the 1% VaR is the smallest of these 1000 losses and the 1% conditional VaR is the average of these 1000 losses. So conditional VaR will always be at least as great as VaR, and usually it will be greater.

The conditional VaR risk measure will be encountered again in §10.2.1, when it will be related to the mean excess loss in the peaks over threshold model. An estimate of conditional VaR may be obtained by fitting the standard parametric distribution for excesses over a threshold, that is, the generalized Pareto distribution that is described in §10.2.1.

9.3 Covariance VaR Models

The covariance VaR methodology was introduced in October 1994 by J.P. Morgan; in fact it is the methodology that underlies the RiskMetrics daily data that are available at www.riskmetrics.com. A detailed discussion of these data, which consist of three very large covariance matrices of the major risk factors in global financial markets, has been given in §7.3.

9.3.1 Basic Assumptions

In the covariance VaR methodology the only data necessary to compute the VaR of a linear portfolio is a covariance matrix of all the assets in the portfolio. One does of course need to know the portfolio composition, but the only other data necessary are the variances and covariances of the asset returns. These can be measured using any of the standard methods — usually a moving average or GARCH methodology will be employed, as explained in Part 1 of the book — and regulators recommend that at least one year of historic data be used in their construction.

The fundamental assumption is that the portfolio P&L is normally distributed. That is, if $\Delta_h P_t = P_{t+h} - P_t$ denotes the h-day unrealized P&L, we assume[12]

[11] Of course conditional VaR depends on the same parameters α, the significance level, and h, the holding period, as the corresponding VaR.

[12] The validity of such an assumption is questionable, and should be investigated using some of the tests that are explained in §10.1.

$$\Delta_h P_t \sim N(\mu_t, \sigma_t^2). \tag{9.5}$$

The $100\alpha\%$ h-day Value-at-Risk is that number $\mathrm{VaR}_{\alpha,h}$ such that $\mathrm{Prob}(\Delta_h P_t < -\mathrm{VaR}_{\alpha,h}) = \alpha$. Now, applying the standard normal transformation:

$$\mathrm{Prob}([\Delta_h P_t - \mu_t]/\sigma_t < [-\mathrm{VaR}_{\alpha,h} - \mu_t]/\sigma_t) = \alpha;$$

or, since $[\Delta_h P_t - \mu_t]/\sigma_t \sim N(0, 1)$, and denoting $[\Delta_h P_t - \mu_t]/\sigma_t$ by the standard normal variate Z_t,

$$\mathrm{Prob}(Z_t < [-\mathrm{VaR}_{\alpha,h} - \mu_t]/\sigma_t) = \alpha.$$

But for a standard normal variate Z_t,

$$\mathrm{Prob}(Z_t < -Z_\alpha) = \alpha,$$

where Z_α is the 100αth percentile of the standard normal density. Therefore,

$$[-\mathrm{VaR}_{\alpha,h} - \mu_t]/\sigma_t = -Z_\alpha;$$

written another way, we have the formula for covariance VaR,

$$\mathrm{VaR}_{\alpha,h} = Z_\alpha \sigma_t - \mu_t. \tag{9.6}$$

It has already been mentioned that VaR as a risk measure is only suitable for short-term risks, so it is normal to assume that $\mu_t = 0$. Now, Z_α is simply a constant given in the standard normal tables (1.645 for $\alpha = 0.05$, 1.96 for $\alpha = 0.025$, 2.33 for $\alpha = 0.01$, and other values may be found from the tables in the back of the book). Thus it is the volatility of the P&L, that is, σ_t, that determines the portfolio VaR.

It is the volatility of the P&L that determines the portfolio VaR in the covariance method

9.3.2 Simple Cash Portfolios

The P&L volatility is easily computed for a simple cash portfolio. Recall from §7.1.1 that the variance of a linear portfolio is a quadratic form: that is, the portfolio return has a variance $\mathbf{w}'\mathbf{V}\mathbf{w}$, where $\mathbf{w}$ is the vector of portfolio weights and $\mathbf{V}$ denotes the covariance matrix of asset returns. Similarly, the portfolio P&L has a variance $\mathbf{p}'\mathbf{V}\mathbf{p}$, where $\mathbf{p}$ is the vector of nominal amounts invested in each asset. Thus if the portfolio is represented as a linear sum of its constituent assets, the covariance method gives the portfolio VaR as (9.6) with

$$\sigma_t = (\mathbf{p}'\mathbf{V}\mathbf{p})^{1/2}.$$

For example, consider a portfolio with two assets with \$1 million invested in asset 1 and \$2 million invested in asset 2. If asset 1 has a 10-day variance of 0.01, asset 2 has a 10-day variance of 0.005 and their 10-day covariance is 0.002, then

$$\mathbf{p}'\mathbf{V}\mathbf{p} = \begin{pmatrix} 1 & 2 \end{pmatrix} \begin{pmatrix} 0.01 & 0.002 \\ 0.002 & 0.005 \end{pmatrix} \begin{pmatrix} 1 \\ 2 \end{pmatrix} = 0.038,$$

and the P&L volatility is $(\mathbf{p}'\mathbf{V}\mathbf{p})^{1/2} = 0.038^{1/2} = \$0.195$ million, so the 5% 10-day VaR is $1.645 \times 0.195 = \$0.32$ million.

Often the historic asset information is given in terms of annualized volatilities and correlations, instead of h-day variances and covariances. But the formula $\mathbf{p}'\mathbf{V}\mathbf{p}$ for the P&L variance can be written equally well in terms of volatilities and correlations. Since the covariance is the correlation multiplied by the product of the square roots of the variances we have

$$\mathbf{p}'\mathbf{V}\mathbf{p} = \mathbf{v}'\mathbf{C}\mathbf{v},$$

where $\mathbf{C}$ is the correlation matrix and $\mathbf{v}$ is the vector of positions multiplied by the square root of the corresponding asset variance. For example, the data in the above example could have been expressed in the following form: the correlation between two assets is 0.2828, asset 1 has annualized volatility of 50% and asset 2 has an annualized volatility of 35.355%. Then one may calculate the asset 10-day standard deviations as $0.5/(250/10)^{1/2} = 0.5/5 = 0.1$ and $0.35355/(250/10)^{1/2} = 0.35355/5 = 0.0707$, respectively (assuming 250 days per year). Then

$$\mathbf{v}'\mathbf{C}\mathbf{v} = \begin{pmatrix} 0.1 & 0.1414 \end{pmatrix} \begin{pmatrix} 1 & 0.2828 \\ 0.2828 & 1 \end{pmatrix} \begin{pmatrix} 0.1 \\ 0.1414 \end{pmatrix} = 0.038,$$

which is the same as $\mathbf{p}'\mathbf{V}\mathbf{p}$ in the previous example.

9.3.3 Covariance VaR with Factor Models

Many linear portfolios are too large to be represented at the asset level and are instead represented by a mapping. Commonly a factor model is used for large equity portfolios or a cash-flow map is used for fixed income portfolios. In this subsection we suppose that a large equity portfolio, with k assets, has been represented by a factor model with m risk factors. Now from §7.1.1 we know that the variance of the portfolio P&L that is due to the risk factors is $\mathbf{p}'\mathbf{B}\mathbf{V}_X\mathbf{B}'\mathbf{p}$, where $\mathbf{p}$ denotes the $n \times 1$ vector of nominal amounts invested in each asset, $\mathbf{B}$ is the $n \times k$ matrix of factor sensitivities (the (i, j)th element of $\mathbf{B}$ is the beta of the ith asset with respect to the jth risk factor) and $\mathbf{V}_X$ is the $k \times k$ risk factor covariance matrix appropriate to the risk horizon.

For a simple example, consider an equity portfolio with \$2 million invested in US and UK stocks. Suppose the net portfolio beta with respect to the FTSE 100 is 1.5 and the net portfolio beta with respect to the S&P 500 is 2. Then the vector $\mathbf{p}'\mathbf{B} = (3, 4)$. Suppose the 1-day risk factor covariance matrix for the FTSE 100 and the S&P 500 is

$$\begin{pmatrix} 0.0018 & 0.0002 \\ 0.0002 & 0.0012 \end{pmatrix} = 10^{-4} \times \begin{pmatrix} 18 & 2 \\ 2 & 12 \end{pmatrix}.$$

Then the variance of the P&L due to the risk factors is

$$10^{-4} \times \begin{pmatrix} 3 & 4 \end{pmatrix} \begin{pmatrix} 18 & 2 \\ 2 & 12 \end{pmatrix} \begin{pmatrix} 3 \\ 4 \end{pmatrix} = 0.0402.$$

Thus the 1% 1-day VaR due to the risk factors is \$2.33 million $\times \sqrt{0.0402}$ = \$0.467 million.

The specific risk of the portfolio also contributes to the P&L variance. In fact from (8.7) we have

$$\text{P\&L variance} = \mathbf{p}'\mathbf{B}\mathbf{V}_X\mathbf{B}'\mathbf{p} + \mathbf{p}'\mathbf{V}_\varepsilon\mathbf{p}, \tag{9.7}$$

where $\mathbf{V}_\varepsilon$ is the $n \times n$ specific risk covariance matrix. There are therefore two components in the covariance VaR:

$$\text{VaR due to market risk factors} = Z_\alpha(\mathbf{p}'\mathbf{B}\mathbf{V}_X\mathbf{B}'\mathbf{p})^{1/2};$$
$$\text{Specific VaR} = Z_\alpha(\mathbf{p}'\mathbf{V}_\varepsilon\mathbf{p})^{1/2}.$$

Saving the residual covariance matrix from the factor $\mathbf{V}_\varepsilon$ model therefore allows a separate VaR measure for the specific risk of the portfolio. Risk capital charges for specific risk can be substantially reduced by this method.[13]

Note that the total VaR should not be measured as the sum of these two components.[14] The total *variance* is the sum, not the total volatility, so the total VaR is obtained from (9.7) as

$$\text{Total VaR} = Z_\alpha(\mathbf{p}'\mathbf{B}\mathbf{V}_X\mathbf{B}'\mathbf{p} + \mathbf{p}'\mathbf{V}_\varepsilon\mathbf{p})^{1/2}.$$

Similar principles are used in the **covariance VaR** spreadsheet on the CD. It calculates VaR for a dollar investor in an international equity portfolio. Total VaR is disaggregated into equity VaR and FX VaR as explained in the help sheet.

9.3.4 Covariance VaR with Cash-Flow Maps

To compute the covariance VaR of a portfolio of bonds or loans, it may be represented by a cash-flow map $\mathbf{p} = (PV_1, PV_2, \ldots, PV_k)$, where PV_t denotes the present value of the income flow at time t for fixed maturity dates $t = 1, \ldots, k$. The present value data should be translated to an absolute scale, so that VaR is measured in terms of P&L, and normally we represent the cash flow amount by its sensitivity to a 1 basis point move in interest rates. Then the covariance matrix must also refer to changes in basis points.

The present value data should be translated to an absolute scale, so that VaR is measured in terms of P&L. Normally we represent the cash flow amount by its sensitivity to a 1 basis point move in interest rates. Then the covariance matrix must also refer to changes in basis points

For example, compute the 1% 1-day VaR for a cash flow of \$1 million in 3 months' time, and \$2 million in 6 months' time, when both the 3-month and the 6-month interest rates are 6%, the 3-month interest rate has an annualized daily volatility of 10%, the 6-month interest rate has an annualized daily volatility of 9% and their correlation is 0.95.

[13] There is a limit to this reduction: the specific risk charge as measured by (9.7) cannot be less than 50% of the charge that would be made under the standardized rules.

[14] In §9.3.5 we see that it is only under very special conditions that one can aggregate covariance VaR measures.

First, the cash-flow amounts must be mapped to P&L sensitivities to changes of 1 basis point. The present value of a basis point move (PVBP) is one of the traditional risk measures that were mentioned in §9.2.1. Following footnote 8, to convert the 3-month cash flow to sensitivity terms it should be divided by 4×10^4; and to convert this to present value terms it should be discounted at the annual rate of 6% (which means dividing by approximately 1.015); similarly the 6-month cash flow should be divided by 2×10^4 and divided by 1.03 to view it in (approximate) present value sensitivity terms. Thus the 3-month cash flow has the present value of a 1 basis point move of

$$\$(10^6 \times 10^{-4})/(4 \times 1.015) = \$24.63;$$

similarly the 6-month cash flow has a present value, for a 1 basis point move, of

$$\$(2 \times 10^6 \times 10^{-4})/(2 \times 1.03) = \$97.09.$$

Secondly, the covariance matrix must refer to rate changes in basis points; if the rates are currently 6% the 10% volatility means that the 3-month rate can vary by ±60 basis points over one year and the 9% volatility means that the 6-month rate can vary by ±54 basis points over one year. In fact, in basis points the 1-day variances are, assuming 250 days per year, $(0.1 \times 600)^2/250 = 14.4$ and $(0.09 \times 600)^2/250 = 11.664$, respectively. Their 1-day covariance is $0.95(14.4 \times 11.664)^{1/2} = 12.312$, so the 1-day covariance matrix in basis point terms is

$$\begin{pmatrix} 14.4 & 12.312 \\ 12.312 & 11.664 \end{pmatrix}.$$

Now, recall that the h-day variance of the portfolio P&L is $\mathbf{p}'\mathbf{V}\mathbf{p}$, where $\mathbf{V}$ is the h-day $k \times k$ covariance matrix for interest rates at maturity dates $t = 1, \ldots, k$, and that the $100\alpha\%$ h-day covariance VaR of this portfolio is simply

$$\mathrm{VaR}_{\alpha,h} = Z_\alpha(\mathbf{p}'\mathbf{V}\mathbf{p})^{1/2}.$$

Thus, in the example above, we calculate

$$\mathbf{p}'\mathbf{V}\mathbf{p} = \begin{pmatrix} 24.63 & 97.09 \end{pmatrix} \begin{pmatrix} 14.4 & 12.312 \\ 12.312 & 11.664 \end{pmatrix} \begin{pmatrix} 24.63 \\ 97.09 \end{pmatrix} = \$177\,570,$$

so the 1% 1-day VaR is $2.33 \times 177\,570^{1/2} = \$981.84$.

A standard cash-flow mapping method, where cash flows are linearly interpolated between adjacent vertices, is not appropriate for VaR models

It should be noted that a standard cash-flow mapping method, where cash flows are linearly interpolated between adjacent vertices, is not appropriate for VaR models. In fact the VaR of the original cash flow will not be the same as the VaR of the mapped cash flow if this method is used. To keep VaR constant one has to use quadratic interpolation between adjacent vertices. Consider the cash-flow X in Figure 9.3 that is between the vertices A and B. Suppose it is t_A days from vertex A and t_B days from vertex B. The linear interpolation method would be to map a proportion $p = t_B/(t_A + t_B)$ of X to A and a proportion

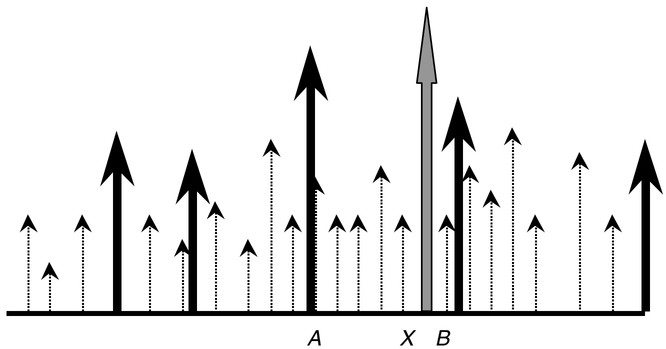

Figure 9.3 Cash-flow mapping.

$1-p = t_A/(t_A+t_B)$ to B. But the variance of the resulting mapped cash flow will not be the same as the variance of the original cash flow. A VaR-invariant cash-flow map is to map a proportion p to vertex A and $1-p$ to vertex B, where

$$V(X) = p^2 V(A) + (1-p)^2 V(B) + 2p(1-p)\text{cov}(A, B). \tag{9.8}$$

The required proportion p is found by solving the quadratic equation (9.8), using estimates for $V(A)$, $V(B)$, $V(X)$ and $\text{cov}(A, B)$. The variances of returns at vertices A and B and their covariance can be obtained from standard covariance matrices, and the variance of the original cash flow X is obtained by linear interpolation between $V(A)$ and $V(B)$ according to the number of days:

$$V(X) = [t_B/(t_A+t_B)]V(A) + [t_A/(t_A+t_B)]V(B).$$

For example, consider a cash flow $X = \$1$ million at 1 year and 200 days. The adjacent vertices are assumed to be A at 1 year and B at 2 years, with 365 days between them. Then $t_A = 200$ and $t_B = 165$, so $t_A/(t_A+t_B) = 0.548$ and $t_B/\ (t_A+t_B) = 0.452$; so with linear interpolation \$0.452 million would be mapped to the 1-year maturity and \$0.548 million would be mapped to the 2-year maturity. However, this mapping would change the VaR of the cash flow. Now suppose that $V(A) = 4 \times 10^{-5}$ and $V(B) = 3.5 \times 10^{-5}$, and that $\text{cov}(A, B) = 3.3 \times 10^{-5}$. By linear interpolation $V(X) = 0.452 \times 4 \times 10^{-5} + 0.548 \times 3.5 \times 10^{-5} = 3.726 \times 10^{-5}$. So the proportion p of the cash flow that should be mapped to A to keep VaR constant is found by solving

$$3.72 = 4p^2 + 3.5(1-p)^2 + 6.6p(1-p)$$

or

$$0.9p^2 - 0.4p - 0.22 = 0$$

giving $p = -0.3198$ or 0.7643. Taking the positive root shows that a VaR invariant cash flow map is to map \$0.7643 million to A and \$0.2357 million to B. The fact that much more is mapped to A than under linear interpolation is due to the fact that A has a much higher volatility than B. On the other hand, if A had variance less than the variance of B, then the VaR invariant cash-flow map would map less to A than is mapped using the linear interpolation method.

9.3.5 Aggregation

Covariance VaR is determined by the volatility of the portfolio P&L, so it behaves just like a volatility. In particular, the rules for adding volatilities apply to the aggregation of covariance VaR measures and, just as one should not simply add volatilities to obtain a total volatility, it is not normally appropriate to add covariance VaR measures. To see why, we now derive a formula for adding the covariance VaR measures of two positions A and B. From the usual rule for the variance of a sum, the variance of the total position is given by

$$\sigma_{A+B}^2 = \sigma_A^2 + \sigma_B^2 + 2\rho\sigma_A\sigma_B,$$

or

$$\sigma_{A+B}^2 = (\sigma_A + \sigma_B)^2 - 2\sigma_A\sigma_B + 2\rho\sigma_A\sigma_B.$$

So

$$\text{Total VaR} = [(\text{VaR}_A + \text{VaR}_B)^2 - 2(1-\rho)\text{VaR}_A\text{VaR}_B]^{1/2}. \tag{9.9}$$

The total VaR is only the same as the sum of the VaRs if we assume the two positions are perfectly correlated

The total VaR is only the same as the sum of the VaRs if we assume the two positions are perfectly correlated, that is, $\rho = 1$. Note that the specific VaR in §9.3.3 should be uncorrelated with the VaR due to the market risk factors (see the comments made at the end of §8.1.1). Thus it makes little sense to add the specific VaR to the risk factor VaR. However, to err on the side of conservatism, this is exactly what regulators recommend.

This example has shown that one has to take into account the correlation, when aggregating the VaR for different portfolios. This is particularly important when there is low or negative correlation between the assets or factors, because in that case (9.9) shows that the total VaR will be much less than the sum of the individual VaRs. The interested reader should use the **covariance VaR** spreadsheet on the CD to see how equity and FX VaR are aggregated to give total VaR.

9.3.6 Advantages and Limitations

The covariance method has one clear advantage: it is very quick and simple to compute, so there is no impediment to intra-day VaR calculations and traders can perform a quick calculation to find the impact of a proposed trade on their VaR limit. However, it also has substantial disadvantages:

- It has very limited applicability, being suitable only for linear portfolios or portfolios that are assumed to be linear with respect to the risk factors.
- It assumes that the portfolio P&L distributions at any point in time are normal. This is certainly true if the returns to individual assets, or risk factors, in the portfolio are normally distributed. But often they are not.[15] If there are major differences between a covariance VaR measure and a historical simulation VaR measure for the same portfolio, it is likely that the non-normality of asset or risk factor returns distributions is a main source of error in the covariance VaR measure (Brooks and Persand, 2000).
- It assumes that all the historic data, including possibly very complex dependencies between risk factors, are captured by the covariance matrix. However, covariance matrices are very limited. Firstly, they are very difficult to estimate and forecast (Chapter 5). Secondly, correlation (and covariance) is only a linear measure of co-dependency (§1.1). And thirdly, unless large GARCH covariance matrices are available (§7.4) it is common to apply the square root of time rule to obtain 10-day covariance matrices. But this assumes that volatility and correlation are constant, which may be a gross simplification for the 10-day horizon.

This last point highlights the main problem with covariance VaR measures — that they are only as accurate as the risk model parameters used in the calculation. We have already noted, many times, how unreliable covariance matrices can be. Moreover, in the covariance equity VaR model, factor sensitivities are also used to compute the VaR, and these too are subject to large and unpredictable forecast errors (§8.2 and §4.5.1). Finally, in the covariance fixed income VaR model, the cash flow map itself can be a source of error — one can obtain quite different VaR measures depending on how the portfolio is mapped to the standard risk factors (§9.3.4).

This last point highlights the main problem with covariance VaR measures — that they are only as accurate as the risk model parameters used in the calculation

9.4 Simulation VaR Models

During the past few years the use of simulation methods for VaR analysis has become standard; simulation methods can overcome some of the problems mentioned in §9.3.6. Historical simulation, in particular, is an extremely popular method for many types of institutions. However, it is not easy to capture the path-dependent behaviour of certain types of complex assets unless Monte Carlo simulation is used.

Monte Carlo VaR measures require a covariance matrix and so their accuracy is limited by the accuracy of this matrix. They also take large amounts of computation time if positions are revalued using complex pricing models. Often full revaluation is simply not possible unless VaR calculations are

[15] In that case it is possible to adapt the covariance VaR measure to the assumption that returns are fat-tailed. In §10.3.1 it will be shown how to generalize covariance VaR so that the portfolio P&L is assumed to be generated by a mixture of normal densities. Also see the spreadsheet **normal mixture VaR** on the CD.

performed overnight, so intra-day VaR measures are often obtained using approximate pricing functions. The trade-off between speed and accuracy in the Monte Carlo VaR methods has been a focus of recent research (Glasserman *et al*., 2001).

This section outlines the basic concepts for historical VaR and Monte Carlo VaR. It explains when and how they should be applied, and outlines the advantages and limitations of each method.

9.4.1 Historical Simulation

The absence of distributional assumptions and the direct estimation of variation without the use of a covariance matrix are the main strengths of the historical VaR model

The basic idea behind historical simulation VaR is very straightforward: one simply uses real historical data to build an empirical density for the portfolio P&L. No assumption about the analytic form of this distribution is made at all, nor about the type of co-movements between assets or risk factors. It is also possible to evaluate option prices and other complex positions for various combinations of risk factors, so it is not surprising that many institutions favour this method.

Historical data on the underlying assets and risk factors for the portfolio are collected, usually on a daily basis covering several years. Regulators insist that at least a year's data be employed for internal models that are used to compute market risk capital requirements, and recommend using between 3 and 5 years of daily data. These data are used to compute the portfolio value on each day during the historic data period, keeping the current portfolio weights constant. This will include computing the value of any options or other complex positions using the pricing models.

In a linear portfolio we represent the h-day portfolio return $\Delta_h P_t / P_t$ as a weighted sum of the returns R_i to assets or risk factors, say,

$$\Delta_h P_t / P_t = w_i R_{i,t} + \ldots + w_k R_{k,t}.$$

The w_i are the portfolio weights (so they sum to 1) or the risk factor sensitivities (§8.1.2). Historical data are obtained on each R_i, and then the portfolio price changes over h days are simulated as

$$\Delta_h P_t = \Sigma(w_i P_t) R_{i,t} = \Sigma p_i R_{i,t},$$

where the p_i are the actual amounts invested in each asset, or, in the case of a factor model, the nominal risk factor sensitivities. This representation allows h-day theoretical (or 'unrealized') P&Ls for the portfolio to be simulated from historical data on $R_1, \ldots, R_k$.

With options portfolios full valuation at each point in time is desirable, but for complex products often a price approximation such as a delta–gamma–vega approximation is used. The current portfolio deltas, gammas and vegas are

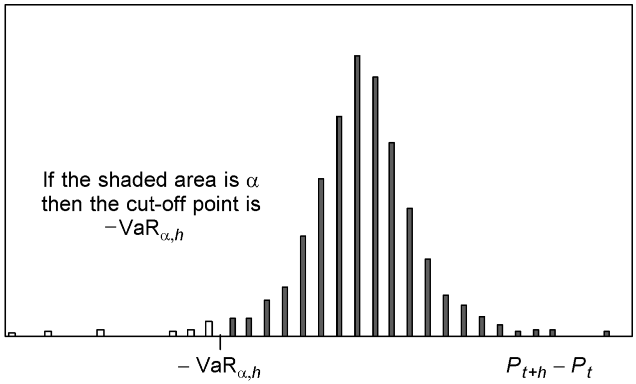

Figure 9.4 Value-at-risk from a simulated P&L density.

Table 9.3: Problems of using a long historic data period for historical VaR

Problem	Possible solution
Data on all relevant assets and risk factors may not be available.	Simulate data using a method such as principal component analysis (§6.4.2).
Full valuation using binomial trees or time-consuming numerical methods on thousands of data points may preclude the use of this method for real-time VaR calculations.	Analytic approximations may be necessary, but this introduces an additional source of error in the VaR measure.
Very long historic data periods may contain a number of extreme market events from far in the past that are not necessarily relevant to current 'normal' circumstances.	Filter out the extreme events from the historic data to obtain the VaR measure under 'normal' market circumstances that is used to calculate the MRR (but see the warning below).
The underlying assets have been trending during the historic period and this does not necessarily reflect their future performance.[16]	Mirror each upward/downward move in the data set with a downward/upward move of the same magnitude, thus doubling the size of the data set and removing any bias from trends.

The ***historical VaR*** *spreadsheet on the CD illustrates how to implement the model*

computed and applied to the historic series on the underlying assets and risk factors (§9.4.3).

The empirical h-day P&L density is obtained by building a histogram of the h-day differences $\Delta P_t = P_{t+h} - P_t$ for all t. Then the historical $\mathrm{VaR}_{\alpha,h}$ is the lower 100αth percentile of this distribution, as shown in Figure 9.4. It will be sensitive to the historic period chosen. For example, if only 300 observations are used in this empirical density, the lower 1% tail will consist of only the three largest losses. The VaR will be the smallest of these three largest losses, and the

[16] However note that the trend over 10 days will be much smaller than the variations, so any bias introduced by trending data will be very small.

conditional VaR will be the average of these three losses. Clearly neither will be very robust to changes in the data period.

From the point of view of robustness it is obviously desirable to use a very long historic data period. But this also has a number of disadvantages, which are summarized, along with possible solutions, in Table 9.3. Also, since the portfolio composition has been determined by current circumstances, how meaningful is it to evaluate the portfolio P&L using market data from far in the past? Long ago the availability, liquidity, risk and return characteristics of the various products in the portfolio will have been quite different; the current portfolio would never have been held, so what is the point in valuing it, under these circumstances?

A word of caution on the use of exceptional events that have occurred in historic data for 'normal circumstances' VaR and for 'stress circumstances' VaR. If such events are filtered out, as suggested above, one has to be very careful not to throw out the relevant part of the data. For VaR models, the relevant part is the exceptional losses. If the exceptional event occurred far in the past and the view is taken that similar circumstances are unlikely to pertain during the next 10 days, then there should be no objection to removing them from the historic data set that is used to compute the normal circumstances VaR that forms the basis of the MRR. These events can always be substituted into the data set for the historical scenarios used to calculated stress VaR measures, for setting trading limits and so forth, possibly using the statistical bootstrap so that they occur at random times and with random frequencies in a simulation of historic data. However, if the exceptional events occurred more recently then it is not acceptable to filter them out of the data for MRR computations, setting trading limits or any other VaR application.

If such events are filtered out, one has to be very careful not to throw out the relevant part of the data. For VaR models, the relevant part is the exceptional losses

As an example of recent 'exceptional' losses that should be used in normal VaR computations, consider the daily data from 2 January 1996 to 2 October 2000 on 20 US stocks that were analysed in Table 4.1. The data showed seven extremely large negative returns among 20 stocks over about 1200 days. Over a VaR period of 10 days one should conclude that the chance of an exceptional event per stock is about $(7 \times 10)/(20 \times 1200)$ or 0.29%. Therefore if all 20 stocks are taken together, the probability of an exceptional event in the portfolio over the next 10 days will be more than 5%. One could conclude that in a portfolio of 20 US stocks an 'exceptional' return is quite 'normal'.

9.4.2 Monte Carlo Simulation

*The **Monte Carlo VaR** spreadsheet is limited to 3000 simulations*

Instead of using actual historic data to build an empirical P&L distribution, it may be more convenient to simulate the movements in underlying assets and risk factors from now until some future point in time (the 'risk horizon' of the model). Taking their current values as starting points, thousands of possible values of the underlying assets and risk factors over the next h days are

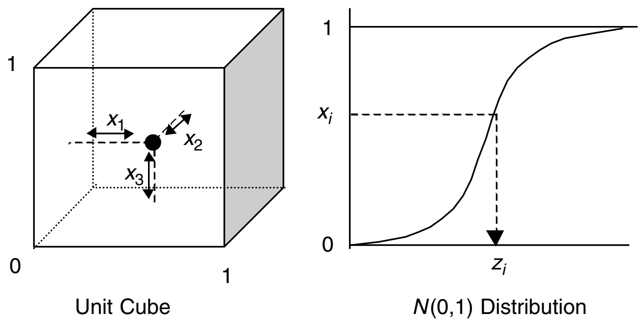

Figure 9.5 Sampling the hypercube and simulating independent $N(0, 1)$ observations.

generated using Monte Carlo methods. This very large set of scenarios is then used to obtain thousands of possible values for the portfolio in h days' time, and a histogram of the differences between these and the current portfolio value is obtained. As with the historic simulation method, the VaR measure is simply the lower percentile of this distribution.

It is necessary to generate these scenarios in a realistic manner. One should not include, for example, a scenario where the 2-year swap rate increases by 20 basis points at the same time as the 3-year swap rate decreases by 50 basis points. Not only will the volatility of an asset determine its possible future values in h days' time, one also needs to take account of the correlations between different assets in the portfolio. For this reason one usually employs an h-day covariance matrix for all the underlying assets and risk factors in the portfolio.

Monte Carlo VaR calculation can be summarized in three stages, which are now explained using the example of a portfolio that has k correlated risk factors. We denote the h-day returns to these risk factors by $R_1, \ldots, R_k$ and their covariance matrix by $\mathbf{V}$.

1. *Take a random sample on k independent standard normal variates*. First a random point in the k-dimensional unit cube is chosen, as this corresponds to a set of k independent random numbers, each between 0 and 1. The case $k = 3$ is illustrated in Figure 9.5. It shows how each random number x_i is obtained from a point in the cube and then used with the probability integral transform to obtain a random sample on the standard normal distribution. Denote this $k \times 1$ vector of independent random samples from $N(0, 1)$ by $\mathbf{z}$.[17]
2. *Use the covariance matrix to transform this sample into correlated h-day returns*. Obtain the Cholesky decomposition of $\mathbf{V}$, as explained in §7.1.2, to convert the random sample $\mathbf{z}$ to a vector of normal returns $\mathbf{r}$ that reflects the appropriate covariance structure.

[17] It is possible to use other forms of distribution, analytic or empirical, but the normal distribution is the usual one.

3. *Apply the pricing model to the simulated correlated h-day risk factor returns.* This will give one simulated h-day portfolio return, and multiplying this by the current price gives a single h-day profit or loss figure.

These three stages are repeated several thousand times, to obtain many P&L figures and thereby build an h-day P&L distribution for the portfolio. How many times will depend on the number of risk factors, that is, the dimension of the hypercube. Figure 9.6 illustrates why. If only a few hundred random points are taken from the cube it is quite possible that they will, by chance, be concentrated in one part of the cube . . . and the next time, by chance, the points could be concentrated in another area of the cube, in which case the Monte Carlo VaR measures will vary enormously each time they are recalculated.

To increase robustness of Monte Carlo VaR estimates one could, of course, take many thousands of points in the cube each time, so that sampling errors are reduced. But then the time taken for all these portfolio revaluations would be completely impractical. Jamshidian and Zhu (1997) show how principal components analysis can be used to gain enormous increases in computational efficiency in Monte Carlo VaR calculations. It is also standard to employ advanced sampling methods, such as deterministic sequences, so that the cube will be covered in a representative manner by fewer samples. But even with the aid of advanced sampling techniques or principal component analysis it will usually be necessary to employ several thousands of points in the sample.

The third stage of Monte Carlo VaR is so time-consuming that it is impossible to perform thousands of revaluations of the portfolio within a realistic time frame

Stages 1 and 2 take hardly any time, and if the third stage is also not too time-consuming there will be no impediment to using thousands of simulations to obtain accurate and robust VaR estimates. However, it is often the case that the third stage of Monte Carlo VaR is so time-consuming that it is impossible to perform thousands of revaluations of the portfolio within a realistic time frame. In that case approximate pricing functions will need to be employed, including analytic forms for complex options and Taylor approximations to portfolio value changes (§9.4.3).

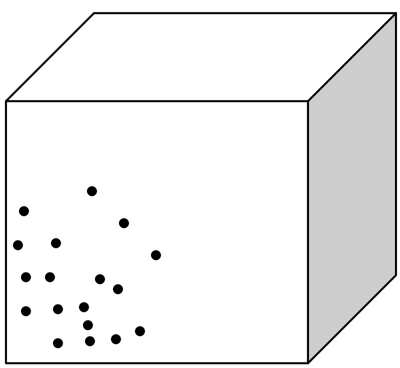

Figure 9.6 Sampling the hypercube with too few points.

The main advantages of Monte Carlo simulation are that it is widely applicable, it is able to capture path-dependent behaviour of complex products and, since simulation techniques are often already employed in the front office, it is operationally efficient to employ these models in VaR calculation also. The main disadvantages of Monte Carlo VaR are the need to use a covariance matrix, which introduces another well-known source of error, and the fact that there is considerable trade-off between speed and accuracy. Often accuracy will need to be compromised in order to complete the VaR calculations within a realistic time frame.

9.4.3 Delta–Gamma Approximations

The **Monte Carlo VaR** spreadsheet calculates the VaR for an option portfolio using full (Black–Scholes) valuation and also using a number of different Taylor approximations based on the option Greeks.

Delta approximations for options and complex instruments are, of course, very convenient. With a delta approximation the portfolio value changes are just a linear function of the underlying asset price changes. Delta approximations are simple, but the VaR obtained in this way tends to be very inaccurate, in any portfolio with gamma or convexity effects. The delta approximation is a local approximation, holding good only for small changes in the underlying price. For VaR measures we need to consider the tail of the P&L distribution, that is, the effect of large price changes. Figure 9.7 explains why the VaR approximation errors from a delta-only representation tend to be rather too large.

The *delta–gamma–theta representation* is a second-order Taylor expansion of the portfolio value change with respect to the changes in the underlying prices (denoted by the vector ΔS). That is,

$$\Delta P \approx \theta \Delta t + \boldsymbol{\delta}' \Delta \mathbf{S} + \tfrac{1}{2} \Delta \mathbf{S}' \Gamma \Delta \mathbf{S}, \qquad (9.10)$$

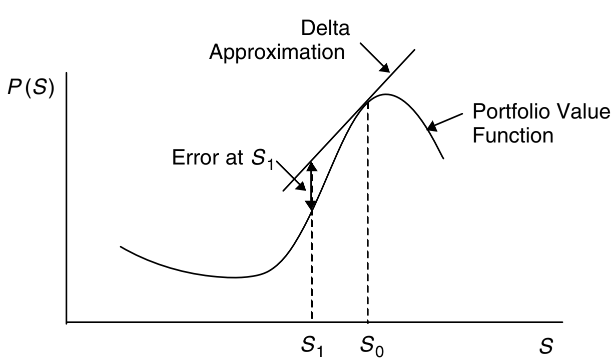

Figure 9.7 Error from delta-only approximation.

where θ is the partial derivative of the portfolio value with respect to time, $\boldsymbol{\delta}$ is a vector of first partial derivatives of the portfolio value with respect to the components of $\mathbf{S}$, and Γ is the Hessian matrix of second partial derivatives with respect to $\mathbf{S}$. Without the first term, (9.10) is the *delta–gamma representation*.

Finite differences are normally employed to compute the option sensitivities. For example, the option delta and gamma may be computed using central differences by valuing the portfolio at the current price S, giving the current value $P(S)$ and for small perturbations above and below this price, at $S + \varepsilon$ and $S - \varepsilon$ for each underlying S in the vector $\mathbf{S}$. Then the delta and gamma corresponding to S are approximated by:

$$\delta \approx [P(S + \varepsilon) - P(S - \varepsilon)]/2\varepsilon \tag{9.11}$$

and

$$\gamma \approx [P(S + \varepsilon) - 2P(S) + P(S - \varepsilon)]/\varepsilon^2. \tag{9.12}$$

The accuracy of the delta–gamma approximation can be tested in the usual way, by comparing the delta–gamma value changes with the full valuation value changes over a suitable historic period. If the approximation is not working sufficiently well it is normal to include volatility as another risk factor, in the *delta–gamma–vega–theta representation*

$$\Delta P \approx \theta\Delta t + \boldsymbol{\delta}'\Delta\mathbf{S} + \tfrac{1}{2}\Delta\mathbf{S}'\Gamma\Delta\mathbf{S} + \mathbf{v}'\Delta\boldsymbol{\sigma}, \tag{9.13}$$

where $\mathbf{v}$ is a vector of first partial derivatives of the portfolio value with respect to the volatility of each component of $\mathbf{S}$, this vector of volatilities being denoted $\boldsymbol{\sigma}$.[18] Theta and the components of the vega vector are also computed using first finite differences. The procedure is the same as with the delta components, only this time using small perturbations on time and the current value of the volatility rather than the underlying price. More details can be seen on the CD.

The delta–gamma, delta–vega and delta–gamma–vega approximations are commonly employed as an approximate pricing function when full revaluation is considered too time-consuming

The delta–gamma, delta–vega and delta–gamma–vega approximations to value changes in a portfolio have an important role to play in simulation VaR models. They are commonly employed as an approximate pricing function when full revaluation is considered too time-consuming. Glasserman *et al.* (2001) show how to use them in Monte Carlo VaR, not as a pricing approximation but as a guide to sampling, with a combination of importance sampling and stratified sampling methods. Rouvinez (1997) shows how the delta–gamma approximation can be used for an analytic approximation to the VaR for an options portfolio.

[18] This assumes that volatility changes are independent of price changes, which is only a rough approximation (§2.3).

9.5 Model Validation

VaR models have become the primary method for calculating MRR, so regulators have set rigorous guidelines for the internal backtesting of these models. The first part of this section describes the backtesting methodology that is outlined in the 1996 Amendment and the translation of backtesting results into risk capital requirements. The second part discusses why different VaR models often give quite different results for the same portfolio, highlighting the various sources of error in each method. VaR estimates can be very sensitive to the choice of historical data period and assumptions about model parameters such as the covariance matrix. Assessment of a VaR model should therefore include an analysis of the sensitivity of results to these decisions and assumptions.

9.5.1 Backtesting Methodology and Regulatory Classification

A 1% 1-day VaR measure gives the level of loss that would be exceeded in normal market circumstances one day in every 100, if the portfolio is left unmanaged. That is, the unrealized P&L, also called the theoretical P&L, of the portfolio will show a loss greater than the 1% 1-day VaR estimate, on average, one day in every 100. So if an accurate VaR model is tested over a period of 1000 days one would expect 10 losses that exceed the level of VaR if the model is accurate. However, if the model is predicting a VaR that is too low, more than 10 exceptional losses will be observed. These 'exceptional' losses are illustrated in Figure 9.8. They form the basis of the VaR model backtests that were proposed in the 1996 Amendment to the Basel Accord.

The total number of exceptional losses may be regarded as a random variable that has a binomial distribution. The probability of 'success' (an exceptional loss) is $p = 0.01$ for a 1% VaR, and the number of trials n is the number of days for the backtest, say $n = 1000$. Then the expected number of exceptional losses

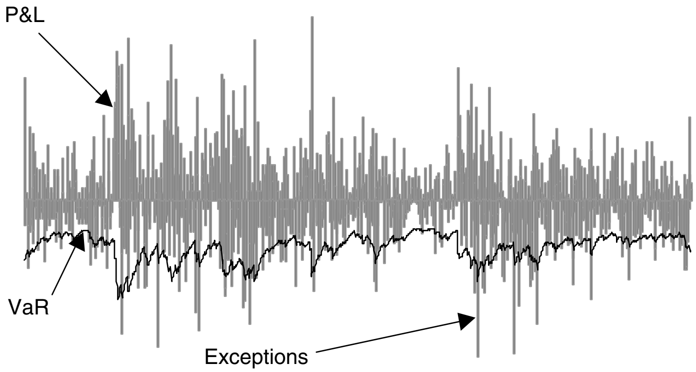

Figure 9.8 Backtests.

is $np = 10$ and the variance of the exceptional losses is $np(1-p) = 9.9$. Therefore the standard deviation is $\sqrt{9.9} = 3.146$ and, using the fact that a binomial distribution is approximately normal when n is large and p is small, a 99% confidence interval for the number of exceptional losses, assuming that the VaR model is accurate, is approximately

$$(np - Z_{0.005}\sqrt{(np(1-p))},\ np + Z_{0.005}\sqrt{(np(1-p))}$$

(see §A.2.1). In this example it becomes

$$(10 - 2.576 \times 3.146,\ 10 + 2.576 \times 3.146) = (1.896,\ 18.104).$$

That is, we are about 99% sure that we will observe between 2 and 18 exceptional losses if the VaR model is correct. Similarly, if the VaR model is accurate, we are 90% confident that between 5 and 15 exceptional losses will be observed in the backtest, because

$$(np - Z_{0.005}\sqrt{(np(1-p))},\ np + Z_{0.005}\sqrt{(np(1-p))}$$
$$= (10 - 1.645 \times 3.146,\ 10 + 1.645 \times 3.146) = (4.825,\ 15.175).$$

Whether or not actual P&L gives rise to more exceptions during backtests than theoretical P&L will depend on the nature of trading

The 1996 Amendment to the Basel accord describes the form of backtests that must be undertaken by firms wishing to use a VaR model for the calculation of MRR. Regulators recommend using the last 250 days of P&L data to backtest the 1% 1-day VaR that is predicted by an internal model. The model should be backtested against both theoretical and actual P&L.[19] Whether or not actual P&L gives rise to more exceptions during backtests than theoretical P&L will depend on the nature of trading. If the main activity is hedging one should expect fewer exceptions, but if traders are undertaking more speculative trades that increase P&L volatility, then the opposite will be observed.

For each area of operations in the firm, such as equity derivatives trading, a backtest is performed by first choosing a candidate portfolio that reflects the type of positions normally taken. The portfolio is held fixed and for each of the last 250 days the VaR is compared with the P&L of the portfolio. The number of exceptional losses over the past 250 days is then recorded.

Regulators will not necessarily adhere to these rules in a hard-and-fast fashion

From the binomial model, the standard error for a 1% VaR for a backtest on 250 days is $\sqrt{(250 \times 0.01 \times 0.99)} = \sqrt{2.475} = 1.573$, so a 90% confidence interval for the number of exceptions observed if the VaR model is accurate $(2.5 \pm 1.645 \times 1.573)$. That is, one is approximately 90% confident that no more than 5 exceptions will occur when the VaR model is accurate. Thus regulators will accept that VaR models which give up to 4 exceptional losses during backtests are performing their function with an appropriate accuracy. These models are labeled 'green zone' models and have a multiplier of $k = 3$ for

[19] Actual includes the P&L from positions taken during the day (even when they are closed out at the end of the day), fees, commissions and so forth. Theoretical P&L is the P&L that would have been obtained if the position had been left unchanged.

Table 9.4: Multipliers for the calculation of market risk requirements

Zone	Number of Exceptions	Multiplier k
Green	Up to 4	3
	5	3.04
	6	3.05
Yellow	7	3.65
	8	3.75
	9	3.85
Red	10 or more	4

the calculation of MRR.[20] If more than four exceptions are recorded the multiplier increases up to a maximum of 4 as shown in Table 9.4 (some 'red zone' models may in fact be disallowed). Regulators will not necessarily adhere to these rules in a hard-and-fast fashion. Some allowance may be made if markets have been particularly turbulent during the backtesting period, particularly if the model has performed well in previous backtests. It should be expected that the results of the backtests will depend on the data period chosen.

9.5.2 Sensitivity Analysis and Model Comparison

VaR estimates are highly sensitive to the assumptions of the VaR model and the other decisions that will need to be made. For example, covariance VaR models and Monte Carlo VaR models both require a covariance matrix, the generation of which requires many assumptions about the nature of asset or risk factor returns. If the same covariance matrix is employed to calculate the covariance VaR and the Monte Carlo VaR of a linear portfolio using both of these methods one should obtain the same result. The reader can verify that this is the case using the VaR spreadsheets on the CD. Differences will only arise if one drops the assumption of normality in one or other of the models (as in §10.3.1, for example), or if too few simulations are being employed in the Monte Carlo VaR.

Historical VaR models employ neither covariance matrices nor normality assumptions, so it is often the case that historical VaR estimates differ substantially from covariance or Monte Carlo VaR estimates. If this is found to be the case one should test whether the assumption of normality is warranted (§10.1). One should also investigate the accuracy of the covariance matrix forecasts, as far as this is possible (§5.1). If neither the covariance matrix nor the normality assumptions are thought to be a problem, then it is possible that it is the historical VaR estimate that is inaccurate, and this will most likely be due an inappropriate choice of historical period. It may be too short, so the

It is often the case that historical VaR measures differ substantially from covariance or Monte Carlo VaR measures

[20] The market risk capital charge will be either yesterday's VaR measure or k times the average of the last 60 days' VaR measures, whichever is greater.

P&L distribution is based on only a few points, in which case VaR estimates will not be robust to changes in the data. Readers should see this for themselves by choosing only a short historical period in the **historical VaR** workbook. It may be too long if there were many very extreme returns in the risk factors a long time ago which have not been repeated during the more recent past.

When attempting to validate a VaR model the risk control function will need to examine many VaR tables, along the lines of that shown in Table 9.1, where the model parameter assumptions are displayed with the table. So for each VaR table the type of covariance matrix should be clearly stated, as should the type of cash-flow maps or factor models if these are employed. In historical VaR models the historical look back period should be stated, and in Monte Carlo VaR models the number of simulations, and any advanced sampling techniques should also be mentioned.

It is the role of risk control to investigate how robust the VaR estimates are to different choices of data and parameters

It is the role of risk control to investigate how robust the VaR estimates are to different choices of data and parameters. This can be achieved using a simple visual analysis—for example, a plot of historical VaR estimates as the historical data period is increased, or a series of covariance VaR plots as certain correlations are varied (cf. Figure 5.4). Since so many assumptions will need to be made about the values of the VaR model parameters this type of sensitivity analysis is a crucial part of model development. Even if the VaR model appears to perform well in backtests, if it does not prove to be robust to small changes in parameter assumptions there is a good chance that future backtesting results will be poor, and the model will be moved into a different Basel zone.

9.6 Scenario Analysis and Stress Testing

The final part of this chapter will examine how the risk characteristics of portfolios can be assessed through scenario analysis and stress testing. Scenario analysis examines the value of a portfolio as the underlying risk factors are perturbed from their current values. Stress testing is really a part of scenario analysis, but instead of considering the sort of perturbations that are expected in normal market circumstances, one looks at the portfolio value when risk factors are moved to extreme positions.

There is much to be said for evaluating portfolios using scenario analysis. Scenario-based methods such as the mark-to-future framework (Dembo, 2000; Dembo *et al.*, 2000) do not depend on distributional assumptions and can incorporate the path-dependent behaviour of any type of security.[21] It is therefore not surprising that many banks prefer to quantify their MRR using scenario-based calculations rather than internal VaR models.

[21] More details are available from www.mark-to-future.com.

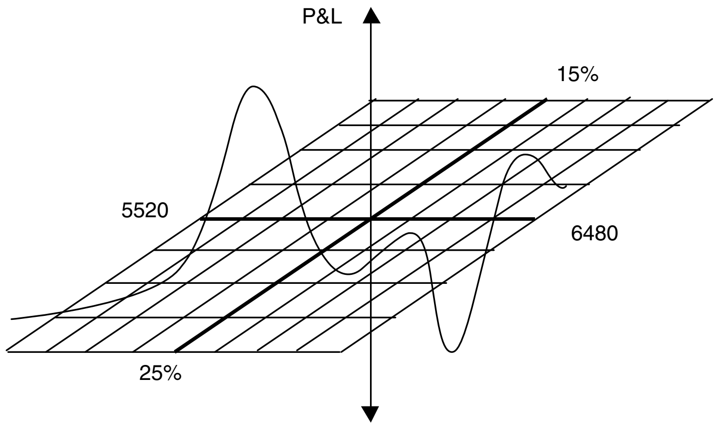

Figure 9.9 Scenario analysis.

9.6.1 Scenario Analysis

As if the building of a historic database on asset prices and risk factors that is updated in real time is not enough of a challenge, a fully functional risk management capability needs to provide a library of scenarios for use by traders and managers. The role of this library should be to allow managers to assess all risk characteristics of existing positions and to allow traders to examine the risk and return of potential trades under various scenarios. The library should consist of

A fully functional risk management capability needs to provide a library of scenarios for use by traders and managers

- covariance matrices (so that the portfolios can be stress-tested for extreme market conditions as described in §9.6.3);
- joint distributions on risk factors and volatilities (which are used to compute the expected loss of a portfolio, as described in §9.6.2);
- smile surfaces, including those that are particularly relevant to current market conditions (these can be generated using the method outlined in §6.3).

The 1996 Amendment to the Basel Accord laid down various guidelines if MRR is being assessed via scenario analysis.[22] Regulators require that portfolios are valued on a grid that is defined according to both country and asset class. The scenarios are for underlying risk factors and risk factor volatilities, so the grid may be envisaged in two dimensions, as shown in Figure 9.9.

Figure 9.9 shows the type of grid that could be used to revalue the net UK equity option positions. The underlying risk factors are the FTSE 100 index and the at-the-money FTSE 100 implied volatility. Their current values are

[22] These scenarios are much less comprehensive that those that should be considered by internal risk management.

taken as the origin of coordinates and the boundary of the grid is defined as ±8% for the value of the FTSE 100 and ±25% for ATM volatility. Thus if the FTSE index is currently at 6000, and ATM implied volatility is currently 20%, the boundaries are set at (5520, 6480) for the index and (15%, 25%) for the ATM volatility. The volatility shifts are imposed by the regulators and applied to the volatility term structure; large shifts (±30%) are taken for short maturities and smaller shifts are taken for longer maturities (e.g. ±8% for 1 year and ±3% for 5 years).

A fairly coarse grid is defined in this range, say at intervals of 100 points for the index and 1% for the implied volatility. In fact usually the grid will contain only 7 × 3 points (seven divisions for the underlying and three divisions for the implied volatility). The regulator does not impose more points but will want some justification that the maximum loss found on the grid cannot be too far from the actual maximum loss. Therefore the portfolio is revalued at each point on this grid and then potential 'hot-spots' are discovered, where the portfolio makes large losses; the grid is then refined about these points and the portfolio is revalued in these regions using the finer grid. Eventually the maximum loss of the portfolio over all possible scenarios is recorded.[23]

If all positions are cash or futures there is no need to revalue the portfolio over a grid or to account for implied volatilities, because in this case the portfolio value will be a linear function of the underlying and the maximum loss will occur at one or other extreme. Interest-rate-dependent products are treated in this way. They are banded into maturity buckets, and the net positions in each bucket are subject to yield variations from ±100 basis points for the short buckets going down to ±60 basis points for the long buckets. If the maximum loss occurs at +100 basis points for the 6–9-month bucket and −100 basis points in the 9–12-month bucket, this is a highly unlikely scenario, so there is also some netting across adjacent buckets. The rules for interest-rate-dependent products are quite complex and the interested reader should consult the current FSA Guidelines for Banking Supervision,[24] or equivalent, for further details.

This type of scenario analysis is performed for all asset types and all countries. There is netting within each asset class but not across countries, and the MRR is then the sum of all the maximum losses that are recorded from each scenario analysis.

9.6.2 Probabilistic Scenario Analysis

If a portfolio is properly hedged then the maximum loss over a scenario grid will occur at a very unlikely scenario. Suppose the maximum loss occurs when

[23] Useful software for performing this type of analysis is available from www.fea.com.

[24] See www.fsa.gov.uk/pubs/supervisor.

the underlying asset price is virtually unchanged at the same time as its implied volatility increases by 5%. This is very unlikely to happen: the discussion in §2.3 shows that there is often a strong negative correlation between prices and volatilities, particularly in equity markets. However, one could reduce the MRR by hedging the portfolio against this scenario. This would be costly, and of little practical importance, but it might be worth putting on the hedge if the MRR will be reduced considerably. This is one example where regulatory policy does not promote good risk management practice.

This is one example where regulatory policy does not promote good risk management practice

Although it is necessary to follow regulators' rules for market risk capital calculations, internal risk management may want to use a more sophisticated approach when assessing the risk characteristics of a portfolio. For one thing, a broader range of movements in both the underlying and the implied volatilities could be applied. Secondly, it would be more informative to compute not the maximum loss, irrespective of how likely it is to occur, but the expected loss of a portfolio, using realistic assumptions about the joint distribution of its risk factors.

It would be more informative to compute the expected loss of a portfolio, using realistic assumptions about the joint distribution of its risk factors

A simple model that employs historic data to generate a joint distribution for movements in an underlying asset and movements in its implied volatility was described in §2.3.4. A joint distribution of risk factor movements, such as those shown in Figures 2.9 and 2.10, can be used to obtain a more realistic assessment of the portfolio's risk characteristics. The expected loss of the portfolio can be obtained by multiplying the profit or loss by the relevant probability and then summing over the entire scenario range. In addition, one can compute the loss standard deviation, to indicate the accuracy of the assessment of expected loss.

9.6.3 Stress-Testing Portfolios

The 1996 Amendment also gave guidelines for the rigorous and comprehensive stress-testing of portfolios. This should include stress testing portfolios for:

- the repeat of an historic event, such as the global equity Black Monday crash in 1987;
- a breakdown in volatility and correlation that is associated with a stress event;
- changes in liquidity accompanying the stress event.

All three types of stress test can be performed using a covariance matrix. In the covariance and Monte Carlo VaR models all that is required is to replace the current covariance matrix by a covariance matrix that reflects the stress test. Even when the VaR model uses historic simulation, it is possible to perform the stress tests using a covariance matrix. Following Duffie and Pan (1997), let $\mathbf{r}$ denote the vector of returns that is used in the historical VaR model, and compute its equally weighted covariance matrix $\mathbf{V}$. Suppose that $\mathbf{W}$ is a stress

Even when the VaR model uses historic simulation, it is possible to perform the stress tests using a covariance matrix covariance matrix (of one of the types defined below), and let **C** and **D** respectively denote the Cholesky decompositions of **V** and **W**. Now transform the historic returns series **r** into another vector of returns $\mathbf{r}^* = \mathbf{D}\mathbf{C}^{-1}\mathbf{r}$. These returns will reflect the conditions of **W** rather than **V** and the stress test is performed using **r*** in the historical VaR model in place of the actual historic returns **r**.

What type of covariance matrix should be used in a stress test? If the test is against an extreme event that actually occurred, then it is simply the covariance matrix that pertained at that time. Historic data on all assets and risk factors around the time of extreme events will have to be stored so that historic covariance matrices that include data on these extreme events can be generated.

If the stress test is for a change in liquidity accompanying the stress event then the stressed matrix can be a new covariance matrix that reflects changes in holding periods associated with different asset classes. For example, consider a covariance matrix for positions in Eastern European equity markets. The risk factors that define the covariance matrix will be the relevant equity indices and the relevant exchange rates. Following an extreme event in the markets, suppose it is possible to liquidate the equity positions locally within 5 days without much effect on the bid–ask spread, but that the foreign exchange market becomes very illiquid and that it will take 25 days to hedge the currency exposure. Then an approximate stress covariance matrix can be obtained by multiplying all elements in the equity block and the equity–FX block of the matrix by 5, but the elements in the FX block by 25. Of course, it is not certain that the resulting covariance matrix will be positive semi-definite. The only way to ensure this is to use the same multiplication factor for all elements, which is in effect modelling liquidity deterioration using an h-day VaR measure for reasonably large h.

For a breakdown of volatilities and correlations it is easier to use the decomposition of a covariance matrix **V** into the product

$$\mathbf{V} = \mathbf{DCD},$$

where **C** is the correlation matrix and **D** is the diagonal matrix with standard deviations along the diagonal. Stress tests can therefore be performed by perturbing the volatilities separately from the correlations, and it is only the changes in the correlation matrix that will affect the positive definiteness of **V**.[25] Readers will observe this when using the stress VaR settings in the VaR spreadsheets on the CD. One is therefore free to change any volatilities to any (positive) level during stress tests, and the resulting VaR measures will always be non-negative. But some changes in the correlation matrix would be disallowed, and so it is important to check for positive definiteness (§7.1.3).

[25] **V** is positive semi-definite if and only if $\mathbf{x}'\mathbf{V}\mathbf{x} \geqslant 0$ for all $\mathbf{x} \neq \mathbf{0}$. But $\mathbf{x}'\mathbf{V}\mathbf{x} = \mathbf{y}'\mathbf{C}\mathbf{y}$ where $\mathbf{y} = \mathbf{D}\mathbf{x}$, and, since **D** is diagonal with positive elements, $\mathbf{y} \neq \mathbf{0}$ if and only if $\mathbf{x} \neq \mathbf{0}$. So **V** is positive semi-definite if and only if **C** is positive semi-definite.

Principal components analysis has natural applications to both scenario analysis and stress testing portfolios

It is worth mentioning that principal components analysis has natural applications to both scenario analysis and stress testing portfolios (Jamshidian and Zhu, 1997). Recall from §6.1 that a principal components representation can be written $\mathbf{X} = \mathbf{P}\mathbf{W}'$ where $\mathbf{X}$ is a set of standardized returns, $\mathbf{P}$ is a matrix of principal components and $\mathbf{W}$ is the matrix of factor weights in (6.2). Both $\mathbf{X}$ and $\mathbf{P}$ have columns that represent time series (the columns of $\mathbf{P}$ are the principal components) but $\mathbf{W}$ is a matrix of constants, which captures the correlation in the system. Stress-testing correlations can therefore be performed by changing the factor weights matrix $\mathbf{W}$. The advantages of this method include the ability to stress test for correlation breakdown without being required to use a covariance matrix.

Frye (1998) shows how shifts in the first few principal components are useful for the scenario analysis of interest-rate-dependent products. Yield curves lend themselves particularly well to principal component analysis, and the first principal component has the interpretation of a parallel shift in the yield curve (§6.2.1). A scenario that increases the first principal component by the equivalent of 100 basis points, for example, is therefore a computationally efficient method for evaluating the effect of parallel shift scenarios. Similarly, the second component represents a yield curve tilt and the third the curvature, so one can capture very complex scenarios on yield curve movements using just a few scenarios on the first three principal components.