CHAPTER 1
Credit Risk

Credit risk emerged as a significant risk management issue during the 1990s. In increasingly competitive markets, banks and securities houses began taking on greater credit risk from this period onwards. While the concept of 'credit risk' is as old as banking itself, it seems that only recently the nature and extent of it has increased dramatically. For example, consider the following developments:

- credit spreads tightened during the late 1990s and again during 2002–2007, to the point where blue-chip companies such as General Electric, British Telecom and Shell were being offered syndicated loans for as little as 10–12 basis points (bps) over LIBOR. To maintain margin, or the increased return on capital, banks increased lending to lower rated corporates, thereby increasing their credit risk both overall and as a share of overall risk;
- investors were finding fewer opportunities in interest rate and currency markets, and therefore moved towards yield enhancement through extending and trading credit across lower rated and emerging market assets;
- the rapid expansion of high yield and emerging market sectors, again lower rated assets, increased the magnitude of credit risk for investors and the banks that held and traded such assets.

The growth in credit risk exposure would naturally be expected to lead to more sophisticated risk management techniques than those employed hitherto. It was accompanied, however, by a rise in the level of corporate defaults and consequently higher losses due to credit deterioration, which led to a rigorous test of banks' risk management systems and procedures. It also led to a demand for the type of product that resulted in the credit derivative market.

The development of the credit derivatives market, and hence the subsequent introduction of structured credit products, was a response to the rising importance attached to credit risk management. For this reason, it is worthwhile beginning this book with a look at credit risk, credit ratings, default and credit risk measurement. In this chapter we will look at the concept of credit risk, before considering the main way that it is measured in banks and financial institutions, using the technique known as credit *value-at-risk*. We also introduce credit risk measurement methodology. First, though, we look at the incidence of corporate defaults during the 1990s and during 2007–2008.

CORPORATE DEFAULT

During the second half of the 1990s and into the new century, credit risk and credit risk management have been topical issues in the financial markets industry. Viewed statistically, 1999 onwards appear to be years of excessive corporate default, when compared with the market experience in the previous two decades. This is vividly illustrated in Figure 1.1, which shows the monetary value of corporate defaults for the period 1980–2002. The average size of corporate bond defaults also rose significantly, as we show in Figure 1.2. Adjusted for inflation, the average size of default in 2002 was over five times that for the entire period 1980–2002.

The excessive levels of corporate defaults provided confirmation that banks and bond investors needed to focus closely on credit risk management. They did this using a two-pronged approach, by concentrating on risk measurement and risk hedging. The former used so-called value-at-risk

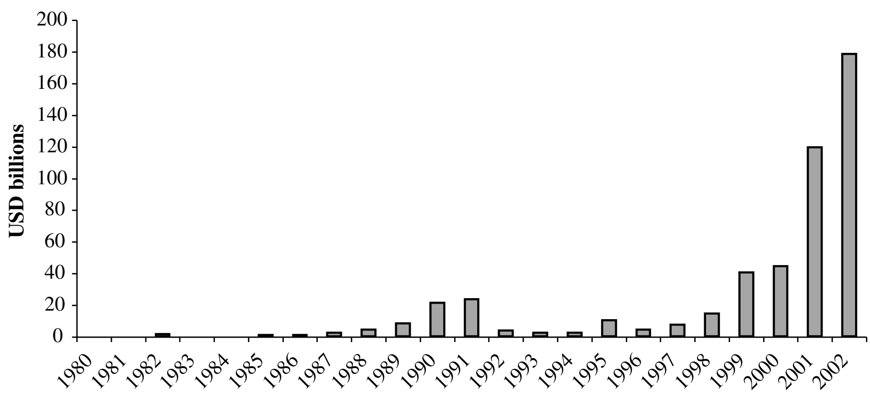

FIGURE 1.1 Global corporate defaults, 1980–2002.
Sources: S&P, CFSB. Used with permission.

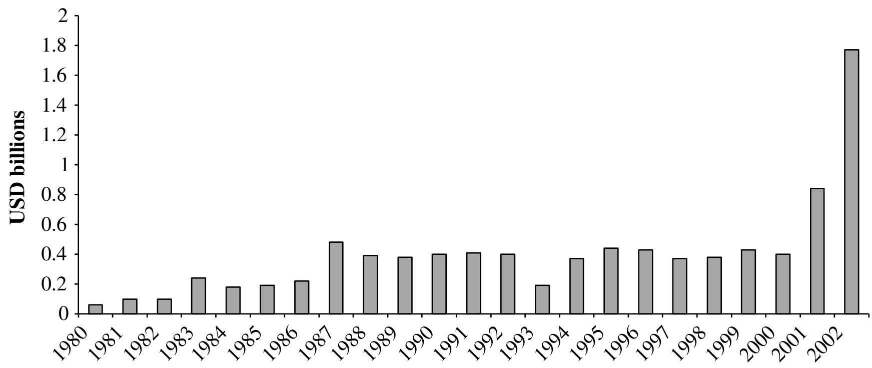

FIGURE 1.2 Average size of corporate bond defaults.
Sources: Moody's, CFSB. Used with permission.

techniques, introduced earlier in the 1990s for market risk measurement, while the latter was accomplished with credit derivatives.

CREDIT RISK

There are two main types of credit risk that a portfolio of assets, or a position in a single asset, is exposed to. These are credit default risk and credit spread risk.

Credit Default Risk

This is the risk that an issuer of debt (obligor) is unable to meet its financial obligations. This is known as *default*. There is also the case of technical default, which is used to describe a company that has not honoured its interest payments on a loan for (typically) three months or more, but has not reached a stage of bankruptcy or administration. Where an obligor defaults, a lender generally incurs a loss equal to the amount owed by the obligor less any recovery amount that the firm recovers as a result of foreclosure, liquidation or restructuring of the defaulted obligor. This recovery amount is usually expressed as a percentage of the total amount and is known as the *recovery rate*. All portfolios with credit exposure exhibit credit default risk.

The measure of a firm's credit default risk is given by its *credit rating*. The three main credit rating agencies are Moody's, Standard & Poor's and Fitch Ratings. These institutions undertake qualitative and quantitative analysis of borrowers and formally rate the borrower after their analysis. The issues considered in the analysis include:

TABLE 1.1 Long-term bond credit ratings.

Fitch	Moody's	S&P	Summary description
Investment grade—High creditworthiness			
AAA	Aaa	AAA	Gilt-edged, prime, maximum safety, lowest risk
AA+	Aa1	AA+	
AA	Aa2	AA	High-grade, high-credit quality
AA–	Aa3	AA–	
A+	A1	A+	
A	A2	A	Upper medium grade
A–	A3	A–	
BBB+	Baa1	BBB+	
BBB	Baa2	BBB	Lower medium grade
BBB–	Baa3	BBB–	
Speculative—Lower creditworthiness			
BB+	Ba1	BB+	
BB	Ba2	BB	Low grade; speculative
BB–	Ba3	BB–	
B+	B1		
B	B2	B	Highly speculative
B–	B3		
Predominantly speculative, substantial risk or in default			
CCC+		CCC+	
CCC	Caa	CCC	Substantial risk, in poor standing
CC	Ca	CC	May be in default, very speculative
C	C	C	Extremely speculative
		CI	Income bonds – no interest being paid
DDD			
DD			Default
D		D	

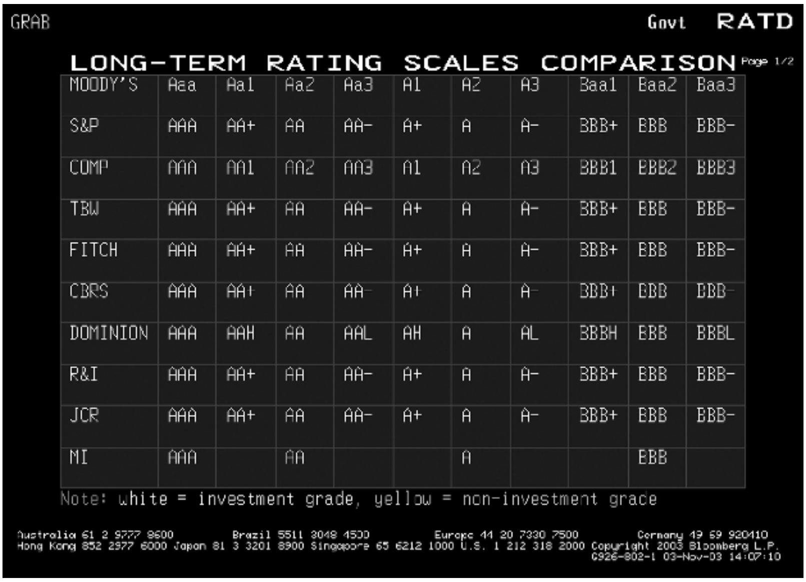

GRAB Govt RATD

LONG-TERM RATING SCALES COMPARISON Page 1/2

MOODY'S	Aaa	Aa1	Aa2	Aa3	A1	A2	A3	Baa1	Baa2	Baa3
S&P	AAA	AA+	AA	AA-	A+	A	A-	BBB+	BBB	BBB-
COMP	AAA	AA1	AA2	AA3	A1	A2	A3	BBB1	BBB2	BBB3
TBW	AAA	AA+	AA	AA-	A+	A	A-	BBB+	BBB	BBB-
FITCH	AAA	AA+	AA	AA-	A+	A	A-	BBB+	BBB	BBB-
CBRS	AAA	AA+	AA	AA-	A+	A	A-	BBB+	BBB	BBB-
DOMINION	AAA	AAH	AA	AAL	AH	A	AL	BBBH	BBB	BBBL
R&I	AAA	AA+	AA	AA-	A+	A	A-	BBB+	BBB	BBB-
JCR	AAA	AA+	AA	AA-	A+	A	A-	BBB+	BBB	BBB-
MI	AAA		AA			A			BBB	

Note: white = investment grade, yellow = non-investment grade

Australia 61 2 9777 8600 Brazil 5511 3048 4500 Europe 44 20 7330 7500 Germany 49 69 920410
Hong Kong 852 2977 6000 Japan 81 3 3201 8900 Singapore 65 6212 1000 U.S. 1 212 318 2000 Copyright 2003 Bloomberg L.P.
G926-802-1 03-Nov-03 14:07:10

FIGURE 1.3 Bloomberg screen RATD, long-term credit ratings.

- the financial position of the firm itself, for example, its balance sheet position and anticipated cash flows and revenues;
- other firm-specific issues such as the quality of management and succession planning;
- an assessment of the firm's ability to meet scheduled interest and principal payments, both in its domestic and in foreign currencies;
- the outlook for the industry as a whole, and competition within it, together with general assessments of the domestic economy.

The range of credit ratings awarded by the three rating agencies is shown at Table 1.1. Ratings can also be seen on Bloomberg page RATD, shown at Figure 1.3. We discuss credit ratings again shortly.

Credit Spread Risk

Credit spread is the excess premium, over and above government or risk-free risk, required by the market for taking on a certain assumed credit exposure. For example, Figure 1.4 shows the credit spreads in May 2009 for US dollar corporate bonds with different credit ratings (AAA, A and BBB). The benchmark is the on-the-run or active US Treasury issue for the

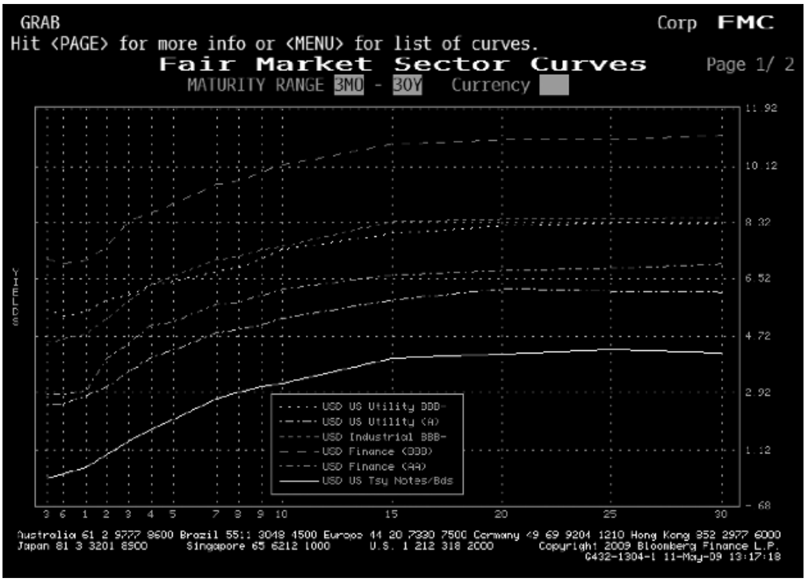

FIGURE 1.4 US dollar bond yield curves, May 2009.
© Bloomberg L.P. Used with permission. Visit www.bloomberg.com

given maturity. Note that the higher the credit rating, the smaller the credit spread; note also the higher yields on financial names. Credit spread risk is the risk of financial loss resulting from changes in the level of credit spreads used in the marking-to-market of a product. It is exhibited by a portfolio for which the credit spread is traded and marked-to-market. Changes in observed credit spreads affect the value of the portfolio and can lead to losses for investors.

CREDIT RATINGS

The risks associated with holding a fixed interest debt instrument are closely connected with the ability of the issuer to maintain the regular coupon payments as well as redeem the debt on maturity. Essentially, the *credit risk* is the main risk of holding a bond. Only the highest quality government debt and a small amount of supra-national debt are considered to be entirely free of credit risk. Therefore, at any time, the yield on a bond reflects investors' views on the ability of the issuer to meet its liabilities as set out in the bond's terms and conditions. A delay in paying a cash liability as it becomes due is known as technical default and is a cause for extreme concern for investors–failure to pay will result in the matter going to court as investors seek to recover their funds. In order to determine the ability of an issuer to meet its obligations for a particular debt issue, for the entire life of the issue, judgmental analysis of the issuer's financial strength and business prospects is required. There are a number of factors that must be considered, and larger banks, fund managers and corporates carry out their own *credit analysis* of individual borrowers' bond issues. The market also makes considerable use of formal *credit ratings* that are assigned to individual bond issues by a formal credit rating agency.

A credit rating is a formal opinion given by a rating agency, of the *credit risk* for investors holding a particular issue of debt securities. Ratings are given to public issues of debt securities by any type of entity, including governments, banks and corporates. They are also given to short-term debt such as commercial paper as well as medium-term notes and long-term debt such as bonds.

The specific factors that are considered by a ratings agency, and the methodology used in conducting the analysis, differ slightly among the individual ratings agencies. Although in many cases the ratings assigned to a particular issue by different agencies are the same, they occasionally differ and in these instances investors usually seek to determine what aspect of an issuer is given more weight in an analysis by which individual agency. Note that a credit rating is not a recommendation to buy (or equally, sell) a particular bond, nor is it a comment on market expectations. Credit analysis does take into account general market and economic conditions, but the overall point of credit analysis is to consider the financial health of the issuer and its ability to meet the obligations of the specific issue being rated. Credit ratings play a large part in the decision-making of investors, and also have a significant impact on the interest rates payable by borrowers. This become an issue of contention during the 2007–08 crash.

Purpose of Credit Ratings

Investors in securities accept the risk that the issuer may default on coupon payments or fail to repay the principal in full on the maturity date. Generally, credit risk is greater for securities with a long maturity, as there is a longer period for the issuer potentially to default. For example, if company issues 10-year bonds, investors cannot be certain that the company will still exist in ten years' time. It may have failed and gone into liquidation some time before that. That said, there is also risk attached to short-dated debt securities; indeed, there have been instances of default by issuers of commercial paper, which is a very short-term instrument.

The prospectus or offer document for an issue provides investors with some information about the issuer so that some credit analysis can be performed on the issuer before the bonds are placed on the market. The information in the offer document enables investors to perform their own credit analysis by studying this information before deciding whether or not to invest in the bonds. Credit assessments take up time, however, and also require the specialist skills of credit analysts. Large institutional investors employ specialists to carry out credit analysis; however, often it is too costly and time-consuming to assess every issuer in every debt market. Therefore investors commonly employ two other methods when making a decision on the credit risk of debt securities:

- name recognition;
- formal credit ratings.

Name recognition is when the investor relies on the good name and reputation of the issuer and accepts that the issuer is of such good financial standing, or sufficient financial standing, that a default on interest and principal payments is unlikely. An investor may feel this way about companies such as Microsoft or British Petroleum. However, the collapse of Barings Bank in 1995 suggested to many investors that it may not be wise to rely on name recognition alone in today's marketplace. The tradition and reputation behind the Barings name allowed the bank to borrow at sub-LIBOR interest rates in the money markets, which put it on a par with highest quality clearing banks in terms of credit rating. The Barings case illustrated that name recognition needs to be augmented by other methods to reduce the risk of loss due to unforeseen events. Credit ratings are increasingly used to make investment decisions about corporate or lesser developed government debt.

FORMAL CREDIT RATINGS

On receipt of a formal request, the credit rating agencies carry out a rating exercise on a specific issue of debt capital. The request for a rating comes from the organisation planning the issue of bonds. Although ratings are provided for the benefit of investors, the issuer must bear the cost. However, it is in the issuer's interest to request a rating as it raises the profile of the bonds, and investors may refuse to buy paper that is not accompanied with a recognised rating. Although the rating exercise involves a credit analysis of the issuer, the rating is applied to a specific debt issue. This means that, in theory, the credit rating is applied not to an organisation itself, but to specific debt securities that the organisation has issued or is planning to issue. In practice, it is common for the market to refer to the creditworthiness of organisations in terms of the rating of their debt. A highly rated company such as Rabobank is therefore referred to as a 'triple-A rated' company, although it is the bank's debt issues that are rated as triple-A.

The rating for an issue is kept under review and if the credit quality of the issuer declines or improves, the rating will be changed accordingly. An agency may announce in advance that it is reviewing a particular credit rating, and may go further and state that the review is a precursor to a possible downgrade or upgrade. This announcement is referred to as putting the issue under *credit watch*. The outcome of a credit watch is, in most cases, likely to be a rating downgrade; however, the review may reaffirm the current rating or possibly upgrade it. During the credit watch phase the agency will advise investors to use the current rating with caution. When an agency announces that an issue is under credit watch, the price of the bonds will fall in the market as investors look to sell out of their holdings. This upward movement in yield will be more pronounced if an actual downgrade results. For example, in October 2008 the government of Ireland was placed under credit watch (and subsequently lost its AAA credit rating); as a result, there was an immediate and sharp sell-off in Irish government Eurobonds, before the rating agencies had announced the actual results of their credit review.

RATINGS CHANGES OVER TIME

Ratings Transition Matrix

We have noted that the rating agencies constantly review the credit quality of firms they have rated. As may be expected, the credit rating of many companies will fluctuate over time as they experience changes in their corporate wellbeing. As a guide to the change in credit rating that might be expected over a 1-year period, Moody's and S&P publish historical transition matrices, which provide average rating transition probabilities for each class of rating. An example is shown at Table 1.2, which is Moody's 1-year ratings transition matrix for 2002. These results are obtained from a sample of a large number of firms over many years. In Table 1.2, the first column shows the initial rating and the first row the final rating. For instance, the probability of an A-rated company being downgraded to Baa in one year is 4.63%. The probability of the A-rated company defaulting in this year is 0.00%.

There are some inconsistencies in the ratings transition table and this is explained by Moody's as resulting from scarcity of data for some ratings categories. For instance, an Aa-rated company has a 0.02% probability of

TABLE 1.2 Moody's one-year rating transition matrix, 2002.

	Aaa	Aa	A	Baa	Ba	B	Caa	Default
Aaa	93.40%	5.94%	0.64%	0.00%	0.02%	0.00%	0.00%	0.00%
Aa	1.61%	90.55%	7.46%	0.26%	0.09%	0.01%	0.00%	0.02%
A	0.07%	2.28%	92.44%	4.63%	0.45%	0.12%	0.01%	0.00%
Baa	0.05%	0.26%	5.51%	88.48%	4.76%	0.71%	0.08%	0.15%
Ba	0.02%	0.05%	0.42%	5.16%	86.91%	5.91%	0.24%	1.29%
B	0.00%	0.04%	0.13%	0.54%	6.35%	84.22%	1.91%	6.81%
Caa	0.00%	0.00%	0.00%	0.62%	2.05%	4.08%	69.20%	24.06%

being in default at year-end, which is higher than the supposedly lower rated A-rated company. Such results must be treated with caution. The conclusion from Table 1.2 is that the most likely outcome at year-end is that the company rating remains the same. It may be that a 1-year time horizon provides little real value; hence, the rating agencies also publish transition matrices for longer periods, such as five and ten years.

We might expect an increased level of default as we move lower down the credit ratings scale. This is borne out in Figure 1.5, which is a reproduction of data published by Moody's. It shows 1-year average cumulative default rates by credit rating category, for the period 1985–2000. We see that the average 1-year default rate rises from zero for the highest rated Aaa, to 15.7% for the B3 rating category. However, investors generally attach little value to 1-year results. Figure 1.6 shows average cumulative default rates for 5- and 10-year time horizons, for the same period covered in Figure 1.5. This repeats the results shown in Figure 1.5, with higher default rates associated with lower credit ratings.

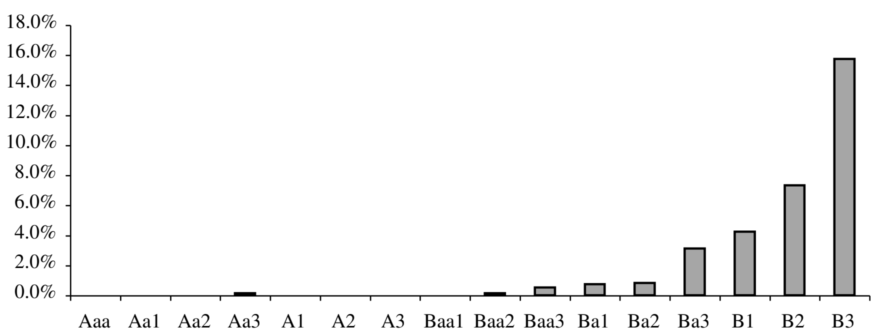

FIGURE 1.5 One-year cumulative default rates, 1985–2000.
Source: Moody's. Reproduced with permission.

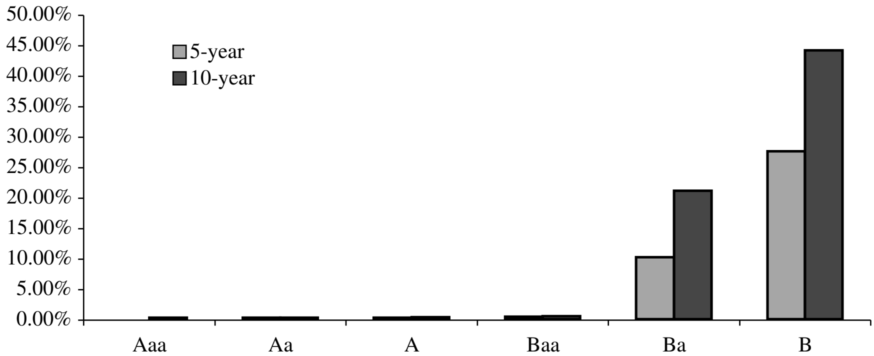

FIGURE 1.6 Five- and 10-year average cumulative default rates, 1985–2000.
Source: Moody's. Reproduced with permission.

Structured Finance Rating Transitions

In March 2009 Moody's published a *Special Comment* on ratings transitions for structured finance securities for the period 1983–2008, which covered the impact of the 2007–08 financial crisis.[1] The extent of the impact was evident from the statistics observed; for example the 12-month rate for downgrades in the global structured finance market reached an historical high of 35.5% in 2008, from a figure of 7.4% in 2007, while the upgrade rate reduced from 2.2% to 0.7%.[2] The primary driver of these results was the poor performance of the most recent structured finance deals, which had sourced the most recent (and poor quality) US sub-prime mortgage assets.

Figure 1.7 shows the global structured finance and global corporate finance 12-month rating transition matrices for the year 2008 and also for the period 1984–2008. Although structured finance assets performed worse than straight corporate bonds during 2008, it is noteworthy that over the longer term the former outperformed the latter.

Further illustration of the impact of the crisis on the structured finance market is given at Table 1.3, which is a summary of rating transition trends. Again, over the longer term structured finance assets have outperformed corporate assets, as shown in the table.

[1] This report is entitled *Structured Finance Rating Transitions: 1983–2008* and is dated March 2009. It is the seventh such annual report from Moody's. It may be obtained from www.moodys.com. Structured finance securities include asset-backed and mortgage-backed securities (ABS and MBS) as well as collateralised debt obligations (CDO), which are introduced in Chapter 12.

[2] Source: Moody's Investors Service.

	Structured Finance in 2008						
	Aaa	**Aa**	**A**	**Baa**	**Ba**	**B**	**Caa and below**
Aaa	73.89%	7.23%	6.31%	5.32%	2.84%	1.74%	2.66%
Aa	1.00%	55.51%	7.29%	5.68%	4.83%	7.98%	17.71%
A	0.27%	0.92%	58.86%	7.72%	4.78%	6.39%	21.07%
Baa	0.10%	0.05%	0.82%	55.42%	5.47%	6.26%	31.88%
Ba	0.05%	0.02%	0.05%	0.67%	54.67%	3.81%	40.74%
B				0.09%	0.21%	45.65%	54.04%
Caa and below						0.13%	99.87%

	Structured Finance: 1984-2008 average over 12-month horizon						
	Aaa	**Aa**	**A**	**Baa**	**Ba**	**B**	**Caa and below**
Aaa	97.79%	0.76%	0.53%	0.37%	0.19%	0.14%	0.21%
Aa	5.27%	87.19%	2.14%	1.12%	0.80%	1.72%	1.77%
A	1.10%	3.26%	85.61%	3.28%	1.39%	2.02%	3.34%
Baa	0.37%	0.47%	2.46%	83.17%	3.46%	2.92%	7.14%
Ba	0.15%	0.07%	0.45%	2.46%	82.33%	3.56%	10.98%
B	0.07%	0.04%	0.08%	0.34%	1.95%	83.63%	13.89%
Caa and below	0.03%			0.07%	0.08%	0.51%	99.30%

	Corporate Finance in 2008						
	Aaa	**Aa**	**A**	**Baa**	**Ba**	**B**	**Caa and below**
Aaa	95.85%	4.15%					
Aa	4.43%	91.25%	4.12%	0.10%		0.10%	
A		10.02%	87.10%	2.69%	0.06%		0.13%
Baa		0.18%	7.30%	88.63%	3.60%	0.28%	
Ba			0.18%	8.06%	83.70%	7.33%	0.73%
B	0.10%			0.19%	6.67%	83.60%	9.44%
Caa and below						15.12%	84.88%

	Corporate Finance: 1984-2008 average over 12-month horizon						
	Aaa	**Aa**	**A**	**Baa**	**Ba**	**B**	**Caa and below**
Aaa	92.76%	6.97%	0.26%		0.02%		
Aa	1.26%	91.45%	6.95%	0.27%	0.05%	0.02%	0.01%
A	0.07%	3.01%	90.91%	5.30%	0.55%	0.11%	0.04%
Baa	0.05%	0.21%	5.37%	88.33%	4.53%	1.00%	0.51%
Ba	0.01%	0.06%	0.43%	6.48%	81.47%	9.56%	2.00%
B	0.01%	0.05%	0.18%	0.40%	6.16%	81.72%	11.47%
Caa and below		0.03%	0.04%	0.19%	0.67%	11.44%	87.63%

FIGURE 1.7 Global structured finance and global corporate finance 12-month rating transition matrices.
Source: Moody's Investors Service. Reproduced with permission.

TABLE 1.3 Summary of rating transition trends, 1983–2008.

	Structured finance		Corporate finance	
	2008	1984-2008	2008	1984-2008
Downgrade rate	35.50%	6.25%	18.22%	13.47%
Upgrade rate	0.69%	2.24%	4.64%	9.86%
Average number of notches downgraded	8.30	6.99	1.64	1.78
Average number of notches upgraded	2.12	2.37	1.34	1.49

Source: Moody's Investors Service. Reproduced with permission.

It is apparent that straight corporate credit ratings are actually much less stable than ratings on structured finance securities, but that when rating changes do take place, the scale of the change is, on average, lower for the corporate ratings. In other words, the statistics appear to confirm that the higher return (on average) of structured finance securities compared to corporate bonds of the same rating reflects their higher risk, which manifests itself during extreme bear market situations.

Corporate Recovery Rates

When a corporate obligor experiences bankruptcy or enters into liquidation or administration, it defaults on its loans. However, this does not mean that all the firm's creditors will lose everything. At the end of the administration process, the firm's creditors will typically receive a portion of their outstanding loans, a *recovery* amount.[3] The percentage of the original loan that is received back is known as the *recovery rate*, which is defined as the percentage of par value that is returned to the creditor.

The seniority of a loan strongly influences the level of the recovery rate. Table 1.4 shows recovery rates for varying levels of loan seniority in 2002 as published by Moody's. The standard deviation for each recovery rate reported is high, which illustrates the dispersion around the mean and reflects widely varying recovery rates even within the same level of seniority. It is clear that the more senior a loan or a bond is, the higher the recovery rate it will have in the event of default.

Recovery rate is a key parameter used to price credit derivatives (see Chapter 6). During the 2007–2009 global recession recovery rates were lower for many industrial sectors and this impacted the settlement value of

[3] This recovery may be received in the form of other assets, such as securities or physical plant, instead of cash.

TABLE 1.4 Moody's recovery rates for varying levels of loan seniority, 2002.

Recovery rates according to loan seniority		
Seniority	*Mean*	*Standard deviation*
Senior secured bank loans	60.70%	26.31%
Senior secured	55.83%	25.41%
Senior unsecured	52.13%	25.12%
Senior subordinated	39.45%	24.79%
Subordinated	33.81%	21.25%
Junior subordinated	18.51%	11.26%
Preference shares	8.26%	10.45%

many CDS contracts. A PowerPoint slide pack in this book's CD-R discusses recovery rate as a factor in CDS evaluation.

Credit risk is measured using the value-at-risk (VaR) technique. This was first introduced as a market risk measurement tool, and subsequently applied to credit risk. Therefore in the next section we introduce the basics of the VaR methodology, which we require for an understanding of Credit VaR.

VALUE-AT-RISK (VaR)

The introduction of VaR as an accepted methodology for quantifying market risk is part of the evolution of risk management. The application of VaR has been extended from its initial use in securities houses to commercial banks and corporates, following its introduction in October 1994 when J.P. Morgan launched *RiskMetrics*™ free over the Internet.

VaR is a measure of the worst expected loss that a firm may suffer over a period of time that has been specified by the user, under normal market conditions and a specified level of confidence. This measure may be obtained in a number of ways, using a statistical model or by computer simulation. It is defined as follows:

VaR is a measure of market risk. It is the maximum loss that can occur with X% confidence over a holding period of n *days.*

Hence VaR is the expected loss of a portfolio over a specified time period for a set level of probability. For example, if a daily VaR is stated as £100,000 to a 95% level of confidence, this means that during the day there is a only a 5% chance that the loss the next day will be *greater* than £100,000. VaR measures the potential loss in market value of a portfolio using estimated volatility and correlation. The 'correlation' referred to is the correlation that exists between the market prices of different instruments in a bank's portfolio. VaR is calculated within a given confidence interval, typically 95% or 99%. It seeks to measure the possible losses from a position or portfolio under 'normal' circumstances. The definition of normality is critical and is essentially a statistical concept that varies by firm and by risk management system. Put simply, however, the most commonly used VaR models assume that the prices of assets in the financial markets follow a normal distribution. To implement VaR, all of a firm's positions data must be gathered into one centralised database. Once this is complete, the overall risk has to be calculated by aggregating the risks from individual instruments across the entire portfolio. The potential move in each instrument (that is, each risk factor) has to be inferred from past daily price movements over a given observation period. For regulatory purposes this period is at least one year. Hence the data on which VaR estimates are based should capture all relevant daily market moves over the previous year.

The main assumption underpinning VaR—and which in turn were shown to be its major weakness—is that the distribution of future price and rate changes will follow past variations. Therefore the potential portfolio loss calculations for VaR are worked out using distributions from historic price data in the observation period.

VaR is therefore a measure of a bank's risk exposure—it is a tool for measuring market risk exposure. There is no one VaR number for a single portfolio, because different methodologies used for calculating VaR produce different results. The VaR number captures only those risks that can be measured in quantitative terms, it does not capture risk exposures such as operational risk, liquidity risk, regulatory risk or sovereign risk.

Assumption of Normality

A distribution is described as *normal* if there is a high probability that any observation from the population sample will have a value that is close to the mean, and a low probability of having a value that is far from the mean. The normal distribution curve is used by many VaR models, which assume that asset returns follow a normal pattern. A VaR model uses the normal curve to estimate the losses that an institution may suffer over a given time period. Normal distribution tables show the probability of a particular observation moving a certain distance from the mean.

If we look along a normal distribution table, we see that at 21.645 standard deviations, the probability is 5%. This means that there is a 5%

probability that an observation will be at least 1.645 standard deviations below the mean. This level is used in many VaR models.

Calculation Methods

There are three common methods for calculating VaR. They are:

- the variance–covariance (or *correlation* or *parametric* method);
- historical simulation;
- Monte Carlo simulation.

Variance–Covariance Method

This method assumes the returns on risk factors are normally distributed, the correlations between risk factors are constant and the delta (or price sensitivity to changes in a risk factor) of each portfolio constituent is constant. Using the correlation method, the volatility of each risk factor is extracted from the historical observation period. Historical data on investment returns are therefore required. The potential effect of each component of the portfolio on the overall portfolio value is then worked out from the component's delta (with respect to a particular risk factor) and that risk factor's volatility.

There are different methods of calculating the relevant risk factor volatilities and correlations. Two alternatives are:

- simple *historic volatility:* this is the most straightforward method, but the effects of a large one-off market move can significantly distort volatilities over the required forecasting period. For example, if using 30-day historic volatility, a market shock will stay in the volatility figure for 30 days until it drops out of the sample range and correspondingly causes a sharp drop in (historic) volatility 30 days *after* the event. This is because each past observation is equally weighted in the volatility calculation;
- to weight past observations unequally: this is done to give more weight to recent observations so that large jumps in volatility are not caused by events that occurred some time ago. One method is to use exponentially weighted moving averages.

Historical Simulation Method

The historical simulation method for calculating VaR is the simplest method and avoids some of the pitfalls of the correlation method. Specifically, the three main assumptions behind correlation (normally distributed returns, constant correlations, constant deltas) are not needed in this case. For historical simulation, the model calculates potential losses using actual historical returns in the risk factors and so captures the non-normal distribution of risk factor returns. This means rare events and crashes can be included in the results. As the risk factor returns used for revaluing the portfolio are actual past movements, the correlations in the calculation are also actual past correlations. They capture the dynamic nature of correlation as well as scenarios when the usual correlation relationships break down.

Monte Carlo Simulation Method

The third method, Monte Carlo simulation is more flexible than the previous two. As with historical simulation, Monte Carlo simulation allows the risk manager to use actual historical distributions for risk factor returns rather than having to assume normal returns. A large number of randomly generated simulations are run forward in time using volatility and correlation estimates chosen by the risk manager. Each simulation will be different, but in total the simulations will aggregate to the chosen statistical parameters (that is, historical distributions and volatility and correlation estimates). This method is more realistic than the previous two models and therefore is more likely to estimate VaR more accurately. However, its implementation requires powerful computers and there is also a trade-off in that the time required to perform calculations is longer.

The level of confidence in the VaR estimation process is selected by the number of standard deviations of variance applied to the probability distribution. A standard deviation selection of 1.645 provides a 95% confidence level (in a 1-tailed test) that the potential estimated price movement will not be more than a given amount based on the correlation of market factors to the position's price sensitivity.

EXPLAINING VALUE-AT-RISK

Correlation

Measures of correlation between variables are important to fund managers who are interested in reducing their risk exposure through diversifying their portfolio. Correlation is a measure of the degree to which a value of one variable is related to the value of another. The correlation coefficient is a single number that compares the strengths and directions of the movements in two instruments' values. The sign of the coefficient determines the relative directions that the instruments move in, while its value determines the strength of the relative movements. The value of the coefficient ranges from 21 to 11, depending on the nature of the relationship. So if, for example, the value of the correlation is 0.5, this means that one instrument moves in the same direction by half of the amount that the other instrument moves. A value of zero means that the instruments are uncorrelated, and their movements are independent of each other.

Correlation is a key element of many VaR models, including parametric models. It is particularly important in the measurement of the variance (hence, volatility) of a portfolio. If we take the simplest example, a portfolio containing just two assets, (1.1) below gives the volatility of the portfolio based on the volatility of each instrument in the portfolio (x and y) and their correlation with one another.

$$V_{port} = \sqrt{x^2 + y^2 + 2xy \cdot \rho(xy)} \tag{1.1}$$

where

- x is the volatility of asset x
- y is the volatility of asset y
- ρ is the correlation between assets x and y.

The correlation coefficient between two assets uses the covariance between the assets in its calculation. The standard formula for covariance is shown at (1.2):

$$\text{Cov} = \frac{\sum_{i=1}^{n} (x_i - \bar{x})(y_i - \bar{y})}{(n-1)} \tag{1.2}$$

where the sum of the distance of each value x and y from the mean is divided by the number of observations minus one. The covariance calculation enables us to calculate the correlation coefficient, shown as (1.3):

$$r = \text{Cov}\frac{(1,2)}{\sigma_1 \sigma_2} \tag{1.3}$$

where σ is the standard deviation of each asset.

Equation (1.1) may be modified to cover more than two instruments. In practice, correlations are usually estimated on the basis of past historical observations. This is an important consideration in the construction and analysis of a portfolio, as the associated risks depend to an extent on the correlation between its constituents.

From a portfolio perspective, a positive correlation increases risk. If the returns on two or more instruments in a portfolio are positively correlated, strong movements in either direction are likely to occur at the same time. The overall distribution of returns is wider and flatter, as there are higher joint probabilities associated with extreme values (both gains and losses). A negative correlation indicates that the assets are likely to move in opposite directions, thus reducing risk.

It has been argued that in extreme situations, such as market crashes or large-scale market corrections, correlations cease to have any relevance, because all assets are moving in the same direction. However, most market scenarios using correlations to reduce the risk of a portfolio is considered satisfactory practice, and the VaR number for a diversified portfolio is lower than that for an undiversified portfolio.

Simple VaR Calculation

To calculate the VaR for a single asset, we calculate the standard deviation of its returns, using either its historical volatility or implied volatility. If a 95% confidence level is required, meaning we wish to have 5% of the observations in the left-hand tail of the normal distribution, this means that the observations in that area are 1.645 standard deviations away from the mean. Let us consider the following statistical data for a government bond, calculated using one year's historical observations:

Nominal:	£10 million
Price:	£100
Average return:	7.35%
Standard deviation:	1.99%

The VaR at the 95% confidence level is 1.645 × 3 × 0.0199 or 0.032736. The portfolio has a market value of £10 million, so the VaR of the portfolio is 0.032736 × 3 × 10,000,000 or £327,360. Therefore this figure is the maximum loss that the portfolio may sustain over one year for 95% of the time.

We may extend this analysis to a 2-asset portfolio. In a 2-asset portfolio, we stated at (1.1) that there is a relationship that enables us to calculate the volatility of a 2-asset portfolio. This expression is used to calculate the VaR, and is shown at (1.4):

$$Var_{port} = \sqrt{w_1^2 s_1^2 + w_2^2 s_2^2 + 2 w_1 w_2 \sigma_1 \sigma_2 r_{1,2}} \quad (1.4)$$

where

w_1 is the weighting of the first asset
w_2 is the weighting of the second asset
σ_1 is the standard deviation or volatility of the first asset
σ_2 is the standard deviation or volatility of the second asset
$r_{1,2}$ is the correlation coefficient between the two assets.

In a 2-asset portfolio, the undiversified VaR is the weighted average of the individual standard deviations. The diversified VaR, which takes into account the correlation between the assets, is the square root of the variance of the portfolio. In practice, banks calculate both diversified and undiversified VaR. The diversified VaR measure is used to set trading limits, while the larger undiversified VaR measure is used to gauge an idea of the bank's risk exposure in the event of a significant correction or market crash. This is because in a crash situation, liquidity dries up as market participants all attempt to sell off their assets. This means that the correlation relationship between assets ceases to have any impact on a book, as all assets move in the same direction. Under this scenario, it is more effective to use an undiversified VaR measure.

Although the description given here is very simple, nevertheless it explains the nature of the VaR measure. VaR is essentially the calculation of the standard deviation of a portfolio, which is used as an indicator of the volatility of that portfolio. A portfolio exhibiting high volatility has a high VaR number. An observer may then conclude that the portfolio has a high probability of making losses. Risk managers and traders may use the VaR measure to help them to allocate capital to more efficient sectors of the bank, as the return on capital can now be measured in terms of return on risk capital. Regulators may use the VaR number as a guide to the capital adequacy levels that they believe the bank requires.

VARIANCE–COVARIANCE VALUE-AT-RISK

Calculation of variance–covariance VaR

In the previous section, we illustrated how VaR can be calculated for a 2-asset portfolio. Here we illustrate how this is done using matrices.

Let us consider the following hypothetical portfolio, invested in two assets, as shown in Table 1.5. The standard deviation of each asset has been calculated on historical observation of asset returns. Note that *returns* are returns of asset prices, rather than the prices themselves. They are calculated from the actual prices by taking the ratio of closing prices. The returns

TABLE 1.5 Two-asset portfolio.

Assets	Bond 1	Bond 2	
Standard deviation	11.83%	17.65%	
Portfolio weighting	60%	40%	
Correlation coefficient			0.647
Portfolio value			£10,000,000
Variance			0.016506998
Standard deviation			12.848%
95% c.i. standard deviations			1.644853
Value-at-Risk			0.211349136
Value-at-Risk£			£2,113,491

are then calculated as the logarithm of the price relativities. The mean and standard deviation of the returns are then calculated using standard statistical formulae. This then gives the standard deviation of daily price relativities, which is converted to an annual figure by multiplying it by the square root of the number of days in a year, usually taken to be 250.

The standard equation (shown as (1.4)) is used to calculate the variance of the portfolio, using the individual asset standard deviations and the asset weightings. The VaR of the book is the square root of the variance. Multiplying this figure by the current value of the portfolio gives us the portfolio VaR, which is £2,113,491.

The *RiskMetrics*™ VaR methodology uses matrices to obtain the same results that we have shown here. This is because once a portfolio starts to contain multiple assets, the method we described above becomes inappropriate. Matrices allow us to calculate VaR for a portfolio containing many hundreds of assets, which would require assessment of the volatility of each asset and correlations of each asset to all the others in the portfolio. We can demonstrate how the parametric methodology uses variance and correlation matrices to calculate the variance, and hence standard deviation, of a portfolio. The matrices are shown in Choudhry (1999).

The advantages of the variance–covariance methodology are that:

- it is simple to apply, and straightforward to explain;
- datasets for its use are immediately available.

The drawbacks of the approach are that it assumes stable correlations and measures only linear risk, it also places excessive reliance on the normal distribution, and returns in the market are widely believed to have ‘fatter tails’ than a true to normal distribution. This phenomenon is known as *leptokurtosis*; that is, the non-normal distribution of outcomes. Another disadvantage is that the process requires *mapping*. To construct a weighting portfolio for the *RiskMetrics*™ tool, cash flows from financial instruments are mapped into precise maturity points, known as grid points. However, in most cases assets do not fit into neat grid points, and complex instruments cannot be broken down accurately into cash flows. The mapping process makes assumptions that frequently do not hold in practice.

Nevertheless the variance–covariance method is still popular in the market, and is frequently the first VaR method installed at a bank.

Historical VaR Methodology

The historical approach to VaR is a relatively simple calculation, and it is also easy to implement and explain. To implement it, a bank requires a database record of its past profit/loss figures for the total portfolio. The required confidence interval is then applied to this record, to obtain a cut-off of the worst-case scenario. For example, to calculate the VaR at a 95% confidence level, the 5th percentile value for the historical data is taken, and this is the VaR number. For a 99% confidence level measure, the 1% percentile is taken. The advantage of the historical method is that it uses the actual market data that a bank has recorded (unlike *RiskMetrics*™, for example, for which the volatility and correlations are not actual values, but estimated values calculated from average figures over a period of time, usually the last five years), and so produces a reasonably accurate figure. Its main weakness is that as it is reliant on actual historical data built up over a period of time—generally, at least one year's data is required to make the calculation meaningful. Therefore it is not suitable for portfolios with asset weightings that frequently change, as another set of data is necessary before a VaR number can be calculated.

In order to overcome this drawback, banks use a method known as *historical simulation*. This calculates VaR for the current portfolio weighting, using the historical data for the securities in the current portfolio. To calculate historical simulation VaR for our hypothetical portfolio considered earlier, comprising 60% of bond 1 and 40% of bond 2, we require the closing prices for both assets over the specified previous period (usually three or five years); we then calculate the value of the portfolio for each day in the period assuming constant weightings.

Simulation Methodology

The most complex calculations use computer simulations to estimate VaR. The most common is the Monte Carlo method. To calculate VaR using a Monte Carlo approach, a computer simulation is run in order to generate a number of random scenarios, which are then used to estimate the portfolio

VaR. The method is probably the most realistic, if we accept that market returns follow a similar 'random walk' pattern. However, Monte Carlo simulation is best suited to trading books containing large option portfolios, where price behaviour is not captured very well with the *RiskMetrics*™ methodology. The main disadvantage of the simulation methodology is that it is time-consuming and uses a substantial amount of computer resources.

A Monte Carlo simulation generates simulated future prices, and it may be used to value an option as well as for VaR applications. When used for valuation, a range of possible asset prices are generated and these are used to assess what intrinsic value the option will have at those asset prices. The present value of the option is then calculated from these possible intrinsic values. Generating simulated prices, although designed to mimic a 'random walk', cannot be completely random because asset prices, although not a pure normal distribution, are not completely random either. The simulation model is usually set to generate very few extreme prices. Strictly speaking, it is asset price *returns* that follow a normal distribution, or rather a *lognormal* distribution. Monte Carlo simulation may also be used to simulate other scenarios; for example, the effect on option 'greeks' for a given change in volatility, or any other parameters. The scenario concept may be applied to calculating VaR as well. For example, if 50,000 simulations of an option price are generated, the 95th lowest value in the simulation will be the VaR at the 95% confidence level. The correlation between assets is accounted for by altering the random selection programme to reflect relationships.

CREDIT VALUE-AT-RISK

Credit risk VaR methodologies take a portfolio approach to credit risk analysis. This means that:

- credit risks of each obligor across the portfolio are re-stated on an equivalent basis and aggregated in order to be treated consistently, regardless of the underlying asset class;
- correlations of credit quality movements across obligors are taken into account.

This allows portfolio effects—that is, the benefits of diversification and risks of concentration—to be quantified.

The portfolio risk of an exposure is determined by four factors:

- size of the exposure;
- maturity of the exposure;

- probability of default of the obligor;
- systemic or concentration risk of the obligor.

Credit VaR, like market risk VaR, considers (credit) risk in a mark-to-market framework. It arises from changes in value due to credit events; that is, changes in obligor credit quality including defaults, upgrades and downgrades.

Nevertheless, credit risk is different in nature to market risk. Typically, market return distributions are assumed to be relatively symmetrical and approximated by normal distributions. In credit portfolios, value changes are relatively small as a result of minor up/downgrades, but can be substantial upon default. This remote probability of large losses produces skewed distributions with heavy downside tails that differ from the more normally distributed returns assumed for market VaR models. This is shown in Figure 1.8.

This difference in risk profiles does not prevent us from assessing risk on a comparable basis. Analytical method market VaR models consider a time horizon and estimate VaR across a distribution of estimated market outcomes. Credit VaR models similarly look to a horizon and construct a distribution of value given different estimated credit outcomes.

When modelling credit risk the two main measures of risk are:

- distribution of loss: obtaining distributions of loss that may arise from the current portfolio. This considers the question of what the expected loss is for a given confidence level;

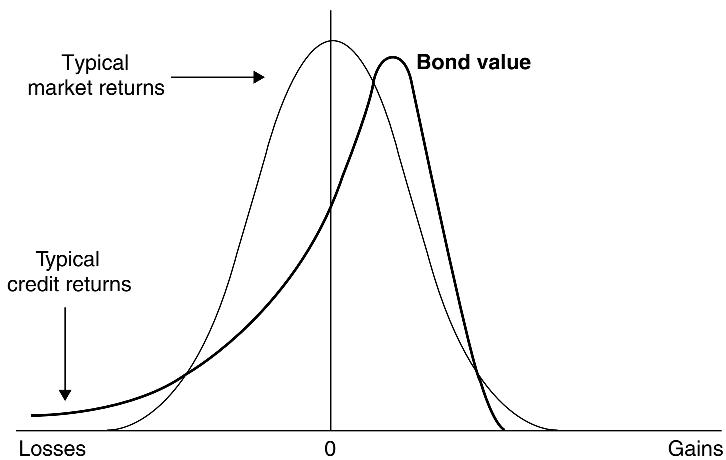

FIGURE 1.8 Comparison of distribution of market returns and credit returns.

- identifying extreme or catastrophic outcomes: this is addressed through the use of scenario analysis and concentration limits.

To simplify modelling, no assumptions are made about the causes of default. Mathematical techniques used in the insurance industry are used to model the event of an obligor default.

Time Horizon

The choice of time horizon will not be shorter than the timeframe over which risk-mitigating actions can be taken. Essentially there are two alternatives:

- a constant time horizon such as one year;
- a hold-to-maturity time horizon.

The constant time horizon is similar to the *CreditMetrics* approach and also to that used for market risk measures. It is more suitable for trading desks. The hold-to-maturity approach is used by institutions such as insurance companies and fund managers.

Data Inputs

Modelling credit risk requires certain data inputs; these include the following:

- credit exposures;
- obligor default rates;
- obligor default rate volatilities;
- recovery rates.

These data requirements present some difficulties. There is a lack of comprehensive default and correlation data and assumptions need to be made at certain times. The most accessible data are compiled by the credit ratings agencies such as Moody's and Standard & Poor's.

We now consider one methodology used for measuring credit VaR, the *CreditMetrics*™ model.

CREDITMETRICS™

CreditMetrics™ was J.P. Morgan's portfolio model for analysing credit risk, providing an estimate of VaR due to credit events caused by upgrades,

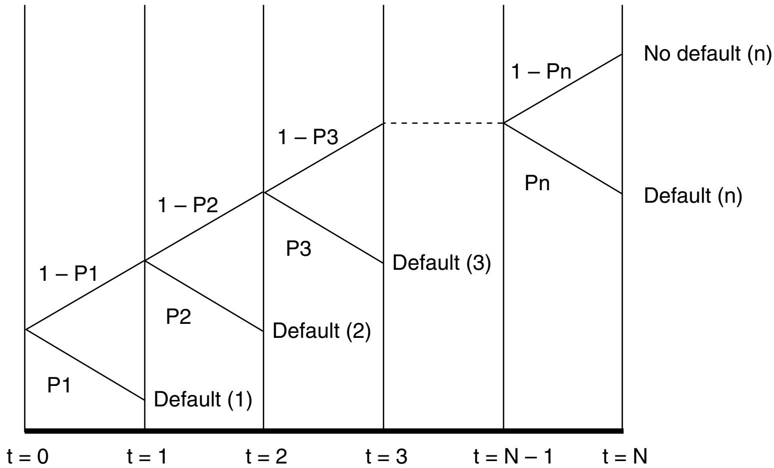

FIGURE 1.9 A binomial model of credit risk.
Source: J.P. Morgan; *RiskMetrics*™ technical document, 1997. Used with permission.

downgrades and default. A software package known as *CreditManager* is available that allows users to implement the *CreditMetrics*™ methodology.[4]

Methodology

There are two main frameworks in use for quantifying credit risk. One approach considers only two scenarios—default and no default. This model constructs a binomial tree of default versus no default outcomes until maturity. This approach is shown in Figure 1.9.

The other approach, sometimes called the RAROC (risk adjusted return on capital) approach holds that risk is the observed volatility of corporate bond values within each credit rating category, maturity band and industry grouping. The idea is to track a benchmark corporate bond (or index) that has observable pricing. The resulting estimate of volatility of value is then used as a proxy for the volatility of the exposure (or portfolio) under analysis.

The *CreditMetrics*™ methodology sits between these two approaches. The model estimates portfolio VaR at the risk horizon due to credit events that include upgrades and downgrades, rather than just defaults. Thus it adopts a mark-to-market framework. As shown in Figure 1.10, bonds within each credit rating category have volatility of value due to day-to-day

[4] The department in J.P. Morgan that developed *CreditMetrics*™ was transformed into a separate corporate entity, known as *RiskMetrics*™, during 1998.

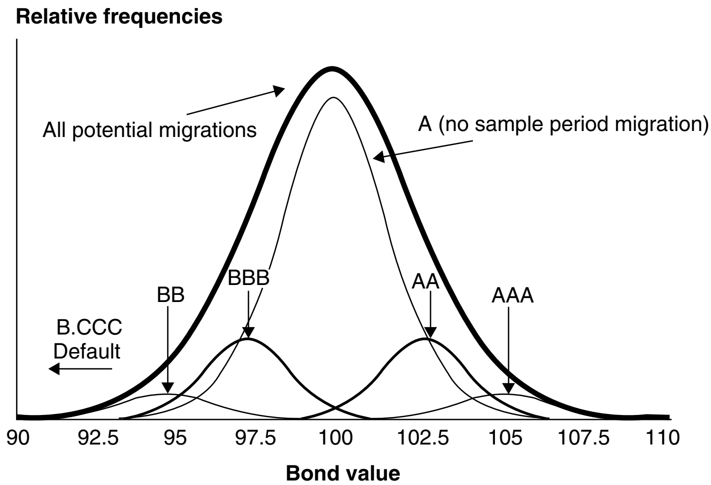

FIGURE 1.10 Distribution of credit returns by rating.
Source: J.P. Morgan; *RiskMetrics*™ technical document, 1997. Used with permission.

credit spread fluctuations. The exhibit shows the loss distributions for bonds of varying credit quality. *CreditMetrics*™ assumes that all credit migrations have been realised, weighting each by a migration likelihood.

Time Horizon

CreditMetrics™ adopts a 1-year risk horizon. The justification given in its technical document[5] is that this is because much academic and credit agency data are stated on an annual basis. This is a convenient convention similar to the use of annualised interest rates in the money markets. The risk horizon is adequate as long as it is not shorter than the time required to perform risk-mitigating actions. Users must therefore adopt their risk management and risk adjustments procedures with this in mind.

The steps involved in *CreditMetrics*™ measurement methodology are shown in Figure 1.11, described by J.P. Morgan as its analytical 'roadmap'.

The elements in each step are:

Exposures User portfolio
Market volatilities
Exposure distributions

[5] J.P. Morgan, Introduction to *CreditMetrics*™, J.P. Morgan & Co., 1997.

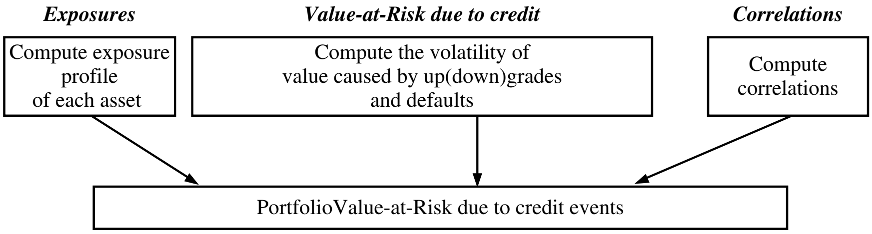

FIGURE 1.11 Analytics road map for *CreditMetrics*™.
Source: J.P. Morgan; *RiskMetrics*™ technical document, 1997. Used with permission.

VaR Due to Credit Events Credit rating
Credit spreads
Rating change likelihood
Recovery rate in default
Present value bond revaluation
Standard deviation of value due to credit quality changes

Correlations Ratings series
Models (for example, correlations)
Joint credit rating changes

Calculating the Credit VaR

CreditMetrics™ methodology assesses individual and portfolio credit VaR in three steps:

Step 1: it establishes the exposure profile of each obligor in a portfolio.
Step 2: it computes the volatility in value of each instrument caused by possible upgrade, downgrade and default.
Step 3: taking into account correlations between each of these events, it combines the volatility of the individual instruments to give an aggregate portfolio risk.

Step 1 Exposure Profiles *CreditMetrics*™ incorporates the exposure of instruments such as bonds (fixed or floating-rate) as well as other loan commitments and off-balance sheet instruments such as swaps. The exposure is stated on an equivalent basis for all products. Products covered include:

- receivables (or trade credit);
- bonds and loans;
- loan commitments;
- letters of credit;
- market-driven instruments.

Step 2 Volatility of Each Exposure from Up/Downgrades and Defaults The levels of likelihood are attributed to each possible credit event of upgrade, downgrade and default. The probability that an obligor will change over a given time horizon to another credit rating is calculated. Each change (migration) results in an estimated change in value (derived from credit spread data and in default, recovery rates). Each value outcome is weighted by its likelihood to create a distribution of value across each credit state, from which each asset's expected value and volatility (standard deviation) of value are calculated.

There are three steps to calculating the volatility of value in a credit exposure:

- the senior unsecured credit rating of the issuer determines the chance of either defaulting or migrating to any other possible credit quality state in the risk horizon;
- revaluation at the risk time horizon can be by either (1) the seniority of the exposure, which determines its recovery rate in case of default or (2) the forward zero-coupon curve (spot curve) for each credit rating category, which determines the revaluation upon up/downgrade;
- the probabilities from the two steps above are combined to calculate volatility of value due to credit quality changes.

Step 3 Correlations Individual value distributions for each exposure are combined to give a portfolio result. To calculate the portfolio value from the volatility of individual asset values requires estimates of correlation in credit quality changes. *CreditMetrics*™ allows for different approaches to estimating correlations including a simple constant correlation. This is because of frequent difficulty in obtaining directly observed credit quality correlations from historical data.

Our discussion of credit risk, and the VaR methodology for measuring such risk, is useful background for the following chapters. We are now in a position to consider the main instruments used to manage and trade credit risk exposure.

REFERENCES

Choudhry, M., *An Introduction to Value-at-Risk*, London: Securities Institute (Services) Publishing, 1999.

Duffie, D. and Singleton, K., *Credit Risk*, Princeton: Princeton University Press, 2003.

Kimber, A., *Credit Risk: From Transaction to Portfolio Management*, Oxford: Elsevier, 2004.