# CreditMetrics™ – Technical Document

*The benchmark for understanding credit risk*

- A value-at-risk (VaR) framework applicable to *all institutions worldwide that carry credit risk in the course of their business.*
- A full portfolio view addressing credit event correlations which can identify the costs of over concentration and benefits of diversification in a mark-to-market framework.
- Results that drive: *investment decisions, risk-mitigating actions,* consistent risk-based *credit limits,* and rational *risk-based capital allocations.*

New York
April 2, 1997

*J.P. Morgan*

*Co-sponsors:*
*Bank of America*
*BZW*
*Deutsche Morgan Grenfell*
*KMV Corporation*
*Swiss Bank Corporation*
*Union Bank of Switzerland*

This *Technical Document* describes CreditMetrics™, a framework for quantifying credit risk in portfolios of traditional credit products (loans, commitments to lend, financial letters of credit), fixed income instruments, and market-driven instruments subject to counterparty default (swaps, forwards, etc.). This is the first edition of what we intend will be an ongoing refinement of credit risk methodologies.

Just as we have done with RiskMetrics™, we are making our methodology and data available for three reasons:

1. We are interested in promoting greater transparency of credit risk. Transparency is the key to effective management.

2. Our aim is to establish a benchmark for credit risk measurement. The absence of a com mon point of reference for credit risk makes it difficult to compare different approaches to and measures of credit risk. Risks are comparable only when they are measured with the same yardstick.

3. We intend to provide our clients with sound advice, including advice on managing their credit risk. We describe the CreditMetrics™ methodology as an aid to clients in under standing and evaluating that advice.

Both J.P. Morgan and our co-sponsors are committed to further the development of CreditMetrics™ as a fully transparent set of risk measurement methods. This broad sponsorship should be interpreted as a signal of our joint commitment to an open and evolving standard for credit risk measurement. We invite the participation of all parties in this continuing enterprise and look forward to receiving feedback to enhance CreditMetrics™ as a benchmark for measuring credit risk.

CreditMetrics™ is based on, but differs significantly from, the risk measurement methodology developed by J.P. Morgan for the measurement, management, and control of credit risk in its trading, arbitrage, and investment account activities. **We remind our readers that no amount of sophisticated analytics will replace experience and professional judgment in managing risks.** CreditMetrics™ is nothing more than a high-quality tool for the professional risk manager in the financial markets and is not a guarantee of specific results.

*CreditMetrics™—Technical Document*

# This book

*Authors:*

Greg M. Gupton
Morgan Guaranty Trust Company
Risk Management Research
(1-212) 648-8062
gupton_greg@jpmorgan.com

Christopher C. Finger
Morgan Guaranty Trust Company
Risk Management Research
(1-212) 648-4657
finger_christopher@jpmorgan.com

Mickey Bhatia
Morgan Guaranty Trust Company
Risk Management Research
(1-212) 648-4299
bhatia_mickey@jpmorgan.com

This is the reference document for CreditMetrics™. It is meant to serve as an introduction to the methodology and mathematics behind statistical credit risk estimation, as well as a detailed documentation of the analytics that generate the data set we provide.

This document reviews:

- the conceptual framework of our methodologies for estimating credit risk;
- the description of the obligors' credit quality characteristics, their statistical description and associated statistical models;
- the description of credit exposure types across "market-driven" instruments and the more traditional corporate finance credit products; and
- the data set that we update periodically and provide to the market for free.

In the interest of establishing a benchmark in a field with as little standardization and precise data as credit risk measurement, we have invited five leading banks, Bank of America, BZW, Deutsche Morgan Grenfell, Swiss Bank Corporation, and Union Bank of Switzerland, and a leading credit risk analytics firm, KMV Corporation, to be co-sponsors of CreditMetrics. All these firms have spent a significant amount of time working on their own credit risk management issues, and we are pleased to have received their input and support in the development of CreditMetrics. With their sponsorship we hope to send one clear and consistent message to the marketplace in an area with little clarity to date.

We have also had many fruitful dialogues with professionals from Central Banks, regulators, competitors, and academics. We are grateful for their insights, help, and encouragement. Of course, all remaining errors and omissions are solely our responsibility.

### How is this related to RiskMetrics™?

We developed CreditMetrics to be as good a methodology for capturing counterparty default risk as the available data quality would allow. Although we never mandated during this development that CreditMetrics must resemble RiskMetrics, the outcome has yielded philosophically similar models. One major difference in the models was driven by the difference in the available data. In RiskMetrics, we have an abundance of daily liquid pricing data on which to construct a model of conditional volatility. In CreditMetrics, we have relatively sparse and infrequently priced data on which to construct a model of unconditional volatility.

### What is different about CreditMetrics?

Unlike market risks where daily liquid price observations allow a direct calculation of value-at-risk (VaR), CreditMetrics seeks to *construct* what it cannot directly *observe*: the volatility of value due to credit quality changes. This constructive approach makes CreditMetrics less an exercise in fitting distributions to observed price data, and more an exercise in proposing models which explain the changes in credit related instruments.

And as we will mention many times in this document, the models which best describe credit risk do not rely on the assumption that returns are normally distributed, marking a significant departure from the RiskMetrics framework.

In the end, we seek to balance the best of all sources of information in a model which looks across broad historical data rather than only recent market moves and across the full range of credit quality migration — upgrades and downgrades — rather than just default.

Our framework can be described in the diagram below. The many sources of information may give an impression of complexity. However, we give a step-by-step introduction in the first four chapters of this book which should be accessible to all readers.

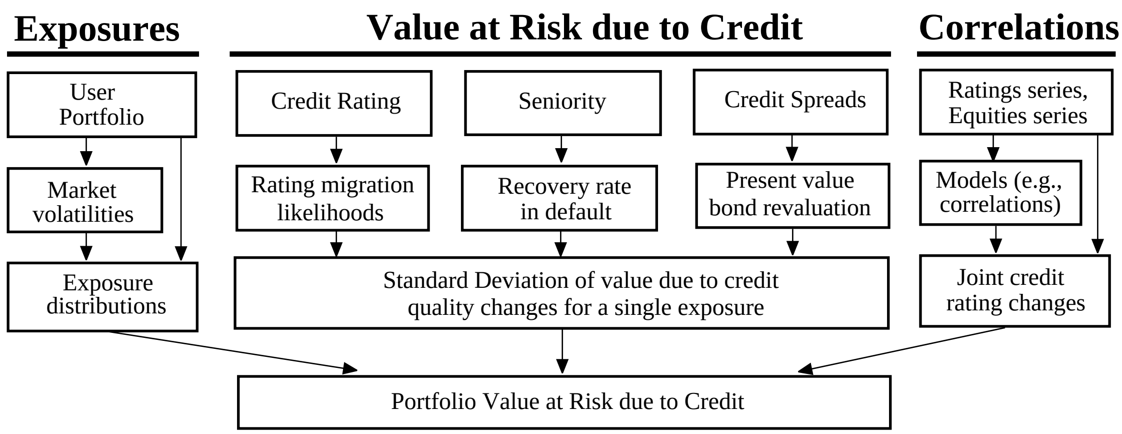


One of our fundamental techniques is *migration analysis*, that is, the study of changes in the credit quality of names through time. Morgan developed transition matrices for this purpose as early as 1987. We have since built upon a broad literature of work which applies migration analysis to credit risk evaluation. The first publication of transition matrices was in 1991 by both Professor Edward Altman of New York University and separately by Lucas & Lonski of Moody's Investors Service. They have since been published regularly (see Moody's Carty & Lieberman [96a][1] and Standard & Poor's *Creditweek* [15-Apr-96]) and are also calculated by firms such as KMV.

**Are RiskMetrics and CreditMetrics comparable?**

Yes, in brief, RiskMetrics looks to a horizon and estimates the *value-at-risk* across a distribution of historically estimated realizations. Likewise, CreditMetrics looks to a horizon and constructs a distribution of historically estimated credit outcomes (rating migrations including potentially default). Each credit quality migration is weighted by its likelihood (transition matrix analysis). Each outcome has an estimate of change in value (given by either credit spreads or studies of recovery rates in default). We then aggregate volatilities across the portfolio, applying estimates of correlation. Thus, although the relevant time horizon is usually longer for credit risk, with CreditMetrics we compute credit risk on a comparable basis with market risk.

[1] Bracketed numbers refer to year of publication.

### What CreditMetrics is not

We have sought to add value to the market's understanding of credit risk estimation, not by replicating what others have done before, but rather by filling in what we believe is lacking. Most prior work has been on the estimation of the relative likelihoods of default for individual firms; Moody's and S&P have long done this and many others have started to do so. We have designed CreditMetrics to accept as an input any assessment of default probability[2] which results in firms being classified into discrete groups (such as rating categories), each with a defined default probability. It is important to realize, however, that these assessments are only inputs to CreditMetrics, and not the final output.

We wish to estimate the *volatility of value* due to changes in credit quality, not just the *expected loss*. In our view, as important as default likelihood estimation is, it is only one link in the long chain of modeling and estimation that is necessary to fully assess credit risk (volatility) within a portfolio. Just as a chain is only as strong as its weakest link, it is also important to diligently address: (i) uncertainty of exposure such as is found in swaps and forwards, (ii) residual value estimates and their *uncertainties*, and (iii) credit quality *correlations* across the portfolio.

### How is this document organized?

One need not read and fully understand the details of this entire document to understand CreditMetrics. This document is organized into three parts that address subjects of particular interest to our diverse readers.

**Part I** **Risk Measurement Framework**
This section is for the general practitioner. We provide a practicable framework of how to think about credit risk, how to apply that thinking in practice, and how to interpret the results. We begin with an example of a single bond and then add more variation and detail. By example, we apply our framework across different exposures and across a portfolio.

**Part II** **Model Parameters**
Although this section occasionally refers to advanced statistical analysis, there is content accessible to all readers. We first review the current academic context within which we developed our credit risk framework. We review the statistical assumptions needed to describe discrete credit events; their mean expectations, volatilities, and correlations. We then look at how these credit statistics can be estimated to describe what happened in the past and what can be projected in the future.

**Part III** **Applications**
We discuss two implementations of our portfolio framework for estimating the *volatility of value due to credit quality changes*. The first is an analytic calculation of the mean and standard deviation of value changes. The second is a simulation approach which estimates the full distribution of value changes. These both embody the same modeling framework and

[2] These assessments may be agency debt ratings, a user's internal ratings, the output of a statistical default prediction model, or any other approach.

produce comparable results. We also discuss how the results can be used in portfolio management, limit setting, and economic capital allocation.

### Future plans

We expect to update this *Technical Document* regularly. We intend to further develop our methodology, data and software implementation as we receive client and academic comments.

CreditMetrics has been developed by the Risk Management Research Group at J.P. Morgan. Special mention must go to Greg M. Gupton who conceived of this project and has been working on developing the CreditMetrics approach at JPMorgan for the last four years. We welcome any suggestions to enhance the methodology and adapt it further to the changing needs of the market. We encourage academic studies and are prepared to supply data for well-structured projects.

### Acknowledgments

We would like to thank our co-sponsors for their input and support in the writing and editing of this document. In particular, we thank the KMV Corporation, which has been a pioneer in developing portfolio approaches to credit risk, and whose work has influenced many of the methods presented here.

We thank numerous individuals at J.P. Morgan who participated in this project, as well as professionals at other banks and academic institutions who offered input at various levels. Also, this document could not have been produced without the contributions of our consulting editor, Margaret Dunkle. We apologize for any omissions.

# Table of Contents

## *Part I Risk Measurement Framework*


| | | |
|---|---|---|
| **Chapter 1.** | **Introduction to CreditMetrics** | **5** |
| 1.1 | The portfolio context of credit | 5 |
| 1.2 | Types of risks modeled | 8 |
| 1.3 | Modeling the distribution of portfolio value | 8 |
| 1.4 | Different credit risk measures | 15 |
| 1.5 | Exposure type differences | 17 |
| 1.6 | Data issues | 20 |
| 1.7 | Advanced modeling features | 21 |
| **Chapter 2.** | **Stand-alone risk calculation** | **23** |
| 2.1 | Overview: Risk for a stand-alone exposure | 23 |
| 2.2 | Step #1: Credit rating migration | 24 |
| 2.3 | Step #2: Valuation | 26 |
| 2.4 | Step #3: Credit risk estimation | 28 |
| 2.5 | Choosing a time horizon | 31 |
| **Chapter 3.** | **Portfolio risk calculation** | **35** |
| 3.1 | Joint probabilities | 36 |
| 3.2 | Portfolio credit risk | 38 |
| 3.3 | Marginal risk | 40 |
| **Chapter 4.** | **Differing exposure types** | **41** |
| 4.1 | Receivables | 42 |
| 4.2 | Bonds and loans | 43 |
| 4.3 | Loan commitments | 43 |
| 4.4 | Financial letters of credit (LCs) | 46 |
| 4.5 | Market-driven instruments | 47 |


## *Part II Model Parameters*


| | | |
|---|---|---|
| **Chapter 5.** | **Overview of credit risk literature** | **57** |
| 5.1 | Expected losses | 57 |
| 5.2 | Unexpected losses | 60 |
| 5.3 | A portfolio view | 63 |
| **Chapter 6.** | **Default and credit quality migration** | **65** |
| 6.1 | Default | 65 |
| 6.2 | Credit quality migration | 66 |
| 6.3 | Historical tabulation | 67 |
| 6.4 | Long-term behavior | 70 |
| **Chapter 7.** | **Recovery rates** | **77** |
| 7.1 | Estimating recovery rates | 77 |
| 7.2 | Distribution of recovery rate | 80 |
| **Chapter 8.** | **Credit quality correlations** | **81** |
| 8.1 | Finding evidence of default correlation | 81 |

8.2 Direct estimation of joint credit moves 83
8.3 Estimating credit quality correlations through bond spreads 84
8.4 Asset value model 85
8.5 Estimating asset correlations 92

## *Part III Applications*

**Chapter 9. Analytic portfolio calculation 107**
9.1 Three-asset portfolio 107
9.2 Marginal standard deviation 110

**Chapter 10. Simulation 113**
10.1 Scenario generation 113
10.2 Portfolio valuation 116
10.3 Summarizing the results 117

**Chapter 11. Portfolio example 121**
11.1 The example portfolio 121
11.2 Simulation results 122
11.3 Assessing precision 125
11.4 Marginal risk measures 129

**Chapter 12. Application of model outputs 133**
12.1 Prioritizing risk reduction actions 133
12.2 Credit risk limits 135
12.3 Economic capital assessment 138
12.4 Summary 140

## *Appendices*

Appendix A. Analytic standard deviation calculation 147
Appendix B. Precision of simulation-based estimates 149
Appendix C. Derivation of the product of N random variables 153
Appendix D. Derivation of risk across mutually exclusive outcomes 155
Appendix E. Derivation of the correlation of two binomials 157
Appendix F. Inferring default correlations from default volatilities 159
Appendix G. International bankruptcy code summary 161
Appendix H. Model inputs 163
Appendix I. Indices used for asset correlations 166

## *Reference*

Glossary of terms 173
Bibliography 183
Index 191

# Part II

# Model Parameters

# Overview of Part II

We have seen in the previous section the general overview, scope and type of results of CreditMetrics. Now we will give more detail to the main modeling parameters used in the CreditMetrics calculation: our sources of data, how we use the data to estimate parameters and why we have made some of the modeling choices we did. There is no single step in the methodology that is particularly difficult; there are simply a lot of steps. We devote a chapter to each major parameter and have tried to present each chapter as a topic which can be read on its own. Although we encourage the reader to study all chapters, reading only a particular chapter of interest is also possible.

Part II is organized into four chapters providing a detailed description of the major parameters within the CreditMetrics framework for quantifying credit risks. Our intent has been to make this description sufficiently detailed so that a practitioner can independently implement this model. This section is organized as follows:

- **Chapter 5: Overview of credit risk literature.** To better place our efforts within the context of prior research in the credit risk quantification field, we give a brief overview of some of the relevant literature.

- **Chapter 6: Default and credit quality migration.** We present an underlying *model of the firm* within which we integrate the process of firm default and, more generally, credit quality migrations. We argue that default is just a special case of a more general process of credit quality migration.

- **Chapter 7: Recovery rates.** Since changes in value are – naturally – greatest in the state of default, our overall measure of credit risk is sensitive to the estimation of recovery rates. We also model the uncertainty of recovery rates.

- **Chapter 8: Credit quality correlations.** The portfolio view of any risk requires an estimation of – most generally – joint movement. In practice, this often means estimating correlation parameters. CreditMetrics requires the joint likelihood of credit quality movements between obligors. Since the observation of credit events are often rare or of poor quality, it is difficult to further estimate their correlations of credit quality moves. We show that the results of several different data sources corroborate each other and might be used to estimate credit quality correlations.

# Chapter 5. Overview of credit risk literature

One of our explicit goals is to stimulate broad discussion and further research towards a better understanding of quantitative credit risk estimation within a full portfolio context. We have sought to make CreditMetrics as competent as is possible within an objective and workable framework. However, we are certain that it will improve with comments from the broad community of researchers.

Extensive previous work has been done towards developing methodologies for estimating different aspects of credit risk. In this chapter, we give a brief survey of the academic literature so that our effort with CreditMetrics can be put in context and so that researchers can more easily compare our approach to others. We group the previous academic research on credit risk estimation within three broad categories:

- estimating particular individual parameters such as expected default frequencies or expected recovery rate in the event of default;
- estimating volatility of value (often termed *unexpected* losses) with the assumption of bond market level diversification; and
- estimating volatility of value within the context of a specific portfolio that is not perfectly diversified.

Also, there have been several papers on credit *pricing*, starting with Merton [74], which discuss debt value as a result of firm risk estimation in an option-theoretic framework. There is more recent work in this area which has focused on incorporating corporate bond yield spreads in valuation models, see Ginzburg, Maloney & Willner [93], Jarrow, Lando & Turnbull [96] and Das & Tufano [96]. For CreditMetrics, we have chosen to focus on the risk assessment side rather than focus on the pricing side.

## 5.1 Expected losses

Expected losses are driven by the expected probability of default and the expected recovery rate in default. We cover recovery rate expectations in much more detail in *Chapter 7* and so will devote this discussion to the expected default likelihood. The problem of estimating the chance of counterparty default has been so difficult that many systems devote all their efforts to this alone. Certainly, if the underlying estimates of default likelihood are poor, then a risk management system is unlikely to make up for this deficiency in its other parts. We will discuss three approaches that are used in practice:

- the accounting analytic approach which is the method used by most rating agencies;
- statistical methods which encompass quite a few varieties; and
- the option-theoretic approach which is a common academic paradigm for default.

We emphasize that CreditMetrics is not another rating service. We assume that exposures input into CreditMetrics will already have been labeled into discrete rating categories as to their credit quality by some outside provider.

As we discuss in *Chapter 6,* a transition matrix for use by CreditMetrics can be fit to any categorical rating system which has historical data. Indeed, we would argue that each credit scoring system should be fit with its own transition matrix. For some users with their own internal rating systems, this will be a necessary first step before applying CreditMetrics to their portfolios. If these systems have limited historical data sets available, then an estimation algorithm that expresses long-term behavior may be desirable (see *Section 6.4*).

### *5.1.1 Accounting analytic approach*

Perhaps the most widely applied approach for estimating firm specific credit quality is fundamental analysis with the use of financial ratios. Such *accounting analytic* methods focus on leverage and coverage measures, coupled with an analysis of the quality and stability of the firm's earnings and cash flows. A good statement of this approach is in Standard and Poor's *Debt Rating Criteria*.[1] These raw quantitative measures are then tempered by the judgment and experience of an industry specialist. This broad description is generally the approach of the major debt rating agencies. This approach yields discrete ordinal groups (e.g., alphabetic ratings) which label firms by credit quality.

We are aware of at least 35 credit rating services worldwide. Also, it is common that financial institutions will maintain their own in-house credit rating expertise. However, letter (or numerical) rating categories by themselves only give an ordinal ranking of the default likelihoods. A quantitative credit risk model such as CreditMetrics cannot utilize ratings without additional information. Each credit rating label must have a statistical meaning such as a specific default probability (e.g., 0.45% over a one-year horizon).

*Table 5.1*
**Moody's corporate bond average cumulative default rates (%)**

| Years | 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|---|
| Aaa | 0.00 | 0.00 | 0.00 | 0.07 | 0.23 |
| Aa1 | 0.00 | 0.00 | 0.00 | 0.31 | 0.31 |
| Aa2 | 0.00 | 0.00 | 0.09 | 0.29 | 0.65 |
| Aa3 | 0.09 | 0.15 | 0.27 | 0.42 | 0.60 |
| A1 | 0.00 | 0.04 | 0.49 | 0.79 | 1.01 |
| A2 | 0.00 | 0.04 | 0.21 | 0.57 | 0.88 |
| A3 | 0.00 | 0.20 | 0.37 | 0.52 | 0.61 |
| Baa1 | 0.06 | 0.39 | 0.79 | 1.17 | 1.53 |
| Baa2 | 0.06 | 0.26 | 0.35 | 1.07 | 1.70 |
| Baa3 | 0.45 | 1.06 | 1.80 | 2.87 | 3.69 |
| Ba1 | 0.85 | 2.68 | 4.46 | 7.03 | 9.52 |
| Ba2 | 0.73 | 3.37 | 6.47 | 9.43 | 12.28 |
| Ba3 | 3.12 | 8.09 | 13.49 | 18.55 | 23.15 |
| B1 | 4.50 | 10.90 | 17.33 | 23.44 | 29.05 |
| B2 | 8.75 | 15.18 | 22.10 | 27.95 | 31.86 |
| B3 | 13.49 | 21.86 | 27.84 | 32.08 | 36.10 |

*Source: Carty & Lieberman [96a]*
*— Moody's Investors Service*

The two major U.S. agencies, S&P and Moody's, have published historical default likelihoods for their letter rating categories. An example from Moody's is shown in *Table 5.1*.

There have been many studies of the historical default frequency of corporate publicly rated bonds. These include Altman [92], [88], [87], Altman & Bencivenga [95], Altman & Haldeman [92], Altman & Nammacher [85], Asquith, Mullins & Wolff [89], Carty & Lieberman [96a] and S&P CreditWeek [96]. These studies are indispensable, and it is important to highlight some important points from them:

- the evolution and change in the original issue high yield bond market is unique in its history and future high yield bond issuance will be different;
- most of the default history is tagged to U.S. domestic issuers who are large enough to have at least an S&P or Moody's rating; and
- the definition of "default" has itself evolved (e.g., it now typically includes "distressed exchanges").

Thus, use of these data must be accompanied by a working knowledge of how they were generated and what they represent.

[1] See: http://www.ratings.standardpoor.com/criteria/index.htm

On a more macro-economic level, researchers have found that aggregate default likelihood is correlated with measures of the business and credit cycle. For example, Fons [91] correlates aggregate defaults to GDP, while Jónsson & Fridson [96] examine also corporate profits, manufacturing hours, money supply, etc.

### *5.1.2 Statistical prediction of default likelihood*

There is a large body of more statistically focused work devoted to building credit quality estimation models, which seek to predict future default. One can identify three basic approaches to estimating default likelihood: qualitative dependent variable models, discriminant analysis, and neural networks. All of these approaches are strictly quantitative and will at least yield a ranking of anticipated default likelihoods and often can be tuned to yield an estimate of default likelihood.

Linear *discriminant analysis* applies a classification model to categorize which firms have defaulted versus which firms survived. In this approach, a historical sample is compiled of firms which defaulted with a matched sample of similar firms that did not default. Then, the statistical estimation approach is applied to identify which variables (and in which combination) can best classify firms into either group. The best example of this approach is Edward Altman's Z-scores; first developed in 1968 and now offered commercially as Zeta Services Inc. This approach yields a continuous numerical score based on a linear function of the relevant firm variables, which – with additional processing – can be mapped to default likelihoods.

The academic literature is full of alternative techniques ranging from principal components analysis, self-organizing feature maps, logistic regression, probit/logit analysis and hierarchical classification models. All of these methods can be shown to have some ability to distinguish high from low default likelihoods firms. Authors who compare the predictive strength of these diverse techniques include Alici [95], Altman, Marco & Varetto [93], and Episcopos, Pericli & Hu [95].

The application of *neural network* techniques to credit scoring include Dutta & Shekhar [88], Kerling [95], and Tyree & Long [94]. The popular press reports commercial applications of neural networks to large volume credit decisions such as credit card authorizations, but there do not appear to be commercial application yet of these neural network techniques for large corporate credits.

### *5.1.3 Option-theoretic approach*

The *option-theoretic* approach was proposed by Fisher Black and Myron Scholes in the context of option pricing, and subsequently developed by Black, Cox, Ingersoll, and most notably, Robert Merton. In this view, a firm has a market value which evolves randomly through time as new information about future prospects of the firm become known. Default occurs when the value of the firm falls so low that the firm's assets are worth less than its obligations. This approach has served as an academic paradigm for default risk, but it is also used as a basis for default risk estimation. The leading commercial exemplar of this approach is KMV. In general, this method yields a continuous numeric value such as the number of standard deviations to the threshold of default, which – with additional processing – can be mapped to default likelihoods

### *5.1.4 Migration analysis*

*Chart 5.1*
**Credit migration**

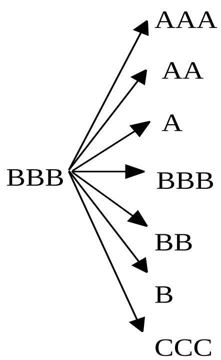


Understanding the potential range of outcomes that are possible is fundamental to risk assessment. As illustrated in *Chart 5.1*, knowing today's credit rating allows us to estimate from history the possible pattern of behaviors in the coming period. More specifically, if an obligor is BBB today, then chances are the obligor will be BBB in one year's time; but it may be up(down)graded. *Table 5.2* shows that, for instance, 86.93% of the time a BBB-rated obligor will remain a BBB, but there is a 5.30% chance that a BBB will downgrade to a BB in one year.

*Table 5.2*
**Credit quality migration likelihoods for a BBB in one year**

| | AAA | AA | A | BBB | BB | B | CCC | Default |
|---|---|---|---|---|---|---|---|---|
| BBB | 0.02% | 0.33% | 5.95% | 86.93% | 5.30% | 1.17% | 0.12% | 0.18% |

One of our fundamental techniques is *migration analysis*. Morgan developed transition matrices for our own use as early as 1987. We have since built upon a broad literature of work which applies migration analysis to credit risk evaluation. The first publication of transition matrices was in 1991 by both Professor Edward Altman of New York University and separately by Lucas & Lonski of Moody's Investors Service. They have since been published regularly (see Moody's Carty & Lieberman [96a] and Standard & Poor's *Creditweek* [15-Apr-96]) and can be calculated by firms such as KMV.

There have been studies of their predictive power and stationarity (Altman & Kao [91] and [92]). More recently, several practitioners (see Austin [92], Meyer [95], and Smith & Lawrence [95]) have used migration analysis to better estimate an accounting-based *allowance for loan and lease losses* (what we would term *expected default losses*). Also, these tools have been used to both estimate (Crabbe [95]) and even potentially improve Lucas [95b]) holding period returns. Finally, academics have constructed arbitrage free credit pricing models (see Ginzburg, Maloney and Willner [93], Jarrow, Lando & Turnbull [96] and Das & Tufano [96]). In CreditMetrics, we extend this literature by showing how to calculate the volatility of value due to credit quality changes (i.e., the potential magnitude of *unexpected* losses) rather than just expected losses.

## 5.2 Unexpected losses

The volatility of losses, commonly termed *unexpected* losses, has proven to be generally much more difficult to estimate than expected losses. Since it is so difficult to explicitly address correlations there have been a number of examples where practitioners take one of two approaches. First, they have applied methods which are statistically easy by addressing either the special case of correlations all equaling zero (perfectly uncorrelated) or correlations all equaling one (perfectly positively correlated). Neither of these is realistic.

Second, they have taken a middle road and assumed that their specific portfolio will have the same correlation effects as some index portfolio. The index portfolio can either be the total credit market ("full" diversification) or a sector index. Thus, the hope would be that statistics drawn from observing the index of debt might be applied through analogy to the specific portfolio. The institution's portfolio would be assumed to have the same correlations and profile of composition as the overall credit markets. These approaches can be grouped into two categories which we discuss in turn:

- historical default volatility; and
- volatility of holding period returns.

Although these may yield some estimate of general portfolio risk, they both suffer from an inability to do meaningful marginal analysis. These techniques would not allow the examination of marginal risk brought by adding some specific proposed transaction. There would also be no guide to know which specific names contribute disproportionate risk to the portfolio.

### *5.2.1 Historical default volatility*

Historical default volatility is available from public studies: see for example *Table 5.3,* which is taken from Carty & Lieberman [96a]. There are several hypotheses to explain why default rates would be volatile:

- defaults are simply random events and the number of firms in the credit markets is not large enough to smooth random variation;
- the volume of high yield bond issuance across years is uneven; and
- the business cycle sees more firms default during downturns versus growth phases.

All three hypotheses are likely to have some truth for the corporate credit markets.

*Table 5.3*
**Volatility of historical default rates by rating category**

| | **Default rate standard deviations (%)** | |
|---|---|---|
| **Credit rating** | **One-year** | **Ten-year** |
| Aaa | 0.0 | 0.0 |
| Aa | 0.1 | 0.9 |
| A | 0.1 | 0.7 |
| Baa | 0.3 | 1.8 |
| Ba | 1.4 | 3.4 |
| B | 4.8 | 5.6 |

*Source: Carty &Lieberman [96a] — Moody's Investors Service*

The problem with trying to understand the volatility of individual exposures in this fashion is that it must be viewed within a portfolio.

### *5.2.2 Volatility of holding period returns*

The volatility of default events is only one component of credit risk. Thus, it may also be useful to examine the volatilities of total holding period returns. A number of academic studies have performed this exercise. For corporate bonds, there are two studies by Bennett, Esser & Roth [93] and Wagner [96]. For commercial loans, there are studies by Asarnow [96] and Asarnow & Marker [95].

Once the historical return volatility is estimated – perhaps grouped by credit rating, maturity bucket, and industry/sector – some practitioners have applied them to analogous exposures in the credit portfolio. In this approach, portfolio diversification is addressed only to the extent that the portfolio under analysis is assuming to be analogous to the credit market universe. Again, there is the obvious problem of diversification differences. But there are also three practical concerns with this approach:

- historical returns are likely to poorly sample returns given credit quality migrations (including defaults) which are low-frequency but important[2];
- the data as it has been collected would require a standard deviation estimate over a sample size of less than 30 and so the standard error of the estimate is large; and
- the studies listed have commingled all sources of volatility – including interest rate fluctuations – rather than just volatility in value due to credit quality changes.

This general approach is sometimes termed the RAROC approach. Implementations vary, but the idea is to track a benchmark corporate bond (or index) which has liquidity and observable pricing. The resulting estimate of volatility of value is then used to proxy for the volatility of some exposure (or portfolio) under analysis.

Potential problems with this approach arise because of its relative inefficiency in estimating infrequent events such as up(down)grades and defaults. Observing some benchmark bond in this fashion over, say, the last year, will yield one of two qualitative results. First, the benchmark bond will neither upgraded nor downgraded and the resulting observed volatility will be (relatively) small. Second, the benchmark bond will have *realized* some credit quality migration and the resulting observed volatility will be (relatively) large.

This process of observing volatility should be unbiased over many trials. However, the estimation error is potentially high due to the infrequent but meaningful impact of credit quality migrations on value. Our approach in CreditMetrics uses long term estimate of migration likelihood rather than observation within some recent sample period and so should avoid this problem.

Consider *Chart 5.2* below. Bonds within each credit rating category can be said also to have volatility of value due to day-to-day credit spread fluctuations. The RAROC approach seeks to measure these fluctuations, but will also sometimes *realize* a potentially large move due to a credit rating migration. Our approach is probabilistic. CreditMetrics assumes that all migrations might have been realized and each is weighted by the likelihoods of migration which we argue is best estimated using long term data.

[2] The credit quality migration and revaluation mechanism in CreditMetrics gives a weight to remote but possible credit quality migrations according to their long-term historical frequency without regard to how a short-term (perhaps one year) sampling of bond prices would – or would not – have observed theses.

*Chart 5.2*
**Construction of volatility across credit quality categories**

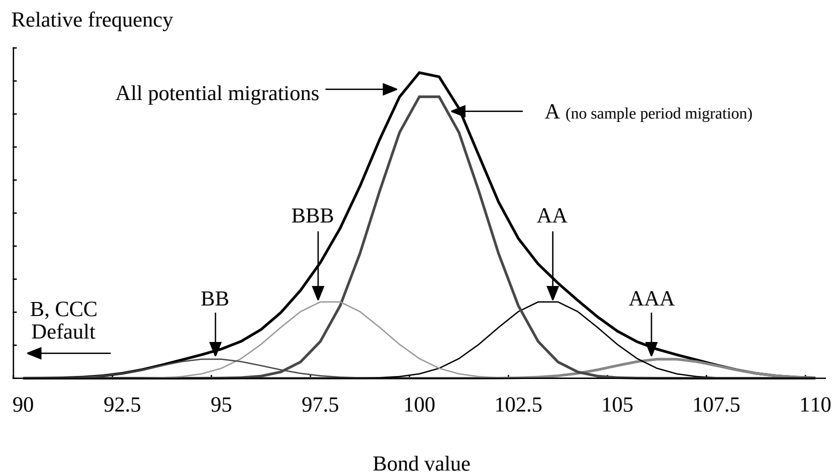


## 5.3 A portfolio view

Any analysis of a group of exposures could be called a portfolio analysis. We use the term here to mean a Markowitz-type analysis where the total risk of a portfolio is measured by explicit consideration of the relationships between individual risks and exposure amounts in a variance-covariance framework. This type of analysis was originated by Harry Markowitz, and has subsequently gone through considerable development, primarily in application to equity portfolios.

A growing number of major institutions estimate the portfolio effects of credit risk in a Markowitz-type framework. However, most institutions still rely on an intuitive assessment as to what level of over concentration to any one area may lead to problems. Thus, bank lenders, for instance, typically set exposure limits against several types of portfolio concentrations, such as industrial sector, geographical location, product type, etc. Lacking the guidance of a model, these groupings tend to be subjective rather than statistical.

For example, industrial sectors are generally defined by aggregating four-digit Standard Industrial Classification (SIC) codes into 60 or fewer groupings. This implies that the banker is assuming that credit quality correlations are higher within an industry or sector and lower between industries or sectors. It is not clear from the data that this is necessarily true. Although this is likely true for *commodity process* industries like oil refining and wood/paper manufacture, we believe it would be less true for *proprietary technology* industries like pharmaceuticals and computer software.

Modern portfolio theory is commonly applied to market risk. The volatilities and correlations necessary to calculate portfolio market volatility are generally readily measurable. In contrast, there has been relatively little academic literature on the problem of measuring diversification or over-concentration within a credit portfolio. To do this requires an understanding of credit quality correlations between obligors.

So, if we were interested in modeling the coincidence of *just* defaults, we might follow Stevenson & Fadil [95]. They constructed 33 industry indices of default experience as listed in Dun & Bradstreet's *Business Failure Record*. The correlation between these indices was their industry level estimate default correlation. While this approach is fine in concept, it suffers from the infrequency of defaults over which to correlate.

To get around this problem, another approach is to construct indices of, not just defaulted firms, but default *likelihoods* of all firms. We know of two services which publish quantitatively estimated default likelihood statistics across thousands of firms: KMV Corporation and Zeta Services. Gollinger & Morgan [93] used time series of default likelihoods (Zeta-Scores™ published by Zeta Services) to estimate default correlations across 42 industry indices. Neither of these studies has been realized in a practicable implementation.

In contrast to these academic suggestions, there is a practicable framework which is a commercial offering by KMV Corporation. In brief, they estimate the value of a firm's debt within the option theoretic framework first described in Merton [74]. Both expected default frequencies (EDFs) and correlations of default expectation are addressed within a consistent – and academically accepted – model-of-the-firm.

The approach practiced by KMV is to look to equity price series as a starting point to understanding the volatility of a firm's underlying (unlevered) asset value moves. Asset value moves can be taken to be approximately normally distributed. These asset values can in turn be mapped ordinally (one-to-one) to credit quality measure, as illustrated in *Chart 3.3*. An assumption of bivariate normality between firms' asset value moves then allows credit quality correlations to be estimated from equity prices series. This is the model on which we have constructed the equity-based correlation estimation in *Chapter 8*. J.P. Morgan has talked with KMV for at least four years on this approach to correlation and we are grateful for their input.

# Chapter 6. Default and credit quality migration

A fundamental source of risk is that the *credit quality* of an obligor may change over the risk horizon. “Credit quality” is commonly used to refer to only the relative chance of default. As we show here, however, CreditMetrics makes use of an extended definition that includes also the volatility of up(down)grades. In this chapter we do the following:

- detail our model-of-the-firm which relates changes in underlying firm value to the event of credit distress;
- generalize this model to incorporate up(down)grades in credit quality;
- discuss the historical tabulation of transition matrices by different providers;
- discuss anticipated long-term behavior of transition matrices; and
- detail an approach to estimate transition probabilities which is sensitive to both the historical tabulation and anticipated long-term behavior.

## 6.1 Default

As discussed in the previous chapter, credit rating systems typically assign an alphabetic or numeric label to rating categories. By itself, this only gives an ordinal ranking of the default likelihoods across the categories. A quantitative framework, such as CreditMetrics, must give meaning to each rating category by linking it with a default probability.[1]

In the academic research, even the definition of the default event has evolved over time. Up to 1989, it was common to look for only missed interest or principal payments (see Altman [87]). Since then, starting with Asquith, Mullins & Wolff [89], researchers realized that distressed exchanges can play an important role in default statistics. Also default rates can be materially different depending upon the population under study. If rates are tabulated for the first few years of newly issued, then the default rate will be much lower than if the population broadly includes all extant debt.

### *6.1.1 Defining credit distress*

For our purposes in CreditMetrics, we look to the following characteristics when we speak of the likelihood of credit distress:

- default rates which have been tabulated weighted by obligors rather than weighted by number of issues or dollars of issuance;
- default rates which have been tabulated broadly upon all obligors rather than just those with recent debt issuance; and

---

[1] Rating agencies commonly also include a judgment for differing recovery rates in their subordinated and structured debt rating. For instance, although senior and subordinated debt to a firm will encounter what we term “credit distress” at the exact same time, the anticipated recovery rate for subordinated is lower and thus it is given a lower rating. It is the senior rating that we look to as the most indicative of credit distress likelihood.

- default rates which are tabulated by senior rating categories (subordinated ratings include recovery rate differences, which are separate from the *likelihood* of default).

This last point is worth elaborating. We utilize credit ratings as an indication of the chance of default and credit rating migration likelihood. However, there are clearly differences in rating – to different debt of the same firm – between senior and subordinated classes. The rating agencies assign lower ratings to subordinated debt in recognition of differences in anticipated recovery rate in default. It is certainly true that senior debt obligations may be satisfied in full during bankruptcy procedures while subordinated debt is paid off only partially. In this circumstance we would say that the firm – and so *all* its debts – encountered *credit distress* even though only the subordinated class realized a *default*. Thus we take the senior credit rating as most indicative of the chance of a firm encountering *credit distress*.

### *6.1.2 Fitting probabilities of default with a transition matrix*

Based on historical default studies from both Moody's and S&P credit rating systems, we have transition matrices which include historically estimated one-year default rates. These are included as part of the dataset for CreditMetrics. Of course, there are many rating agencies beyond S&P and Moody's. There are two ways of using alternative credit rating systems depending upon what historical information is available.

- If individual rating histories are available, then tabulating a transition matrix would give first direct estimate of the transition likelihoods including default.

- If all that is available are cumulative default histories by rating category,[2] then the transition matrix which "best replicates" this history can be estimated.

In the absence of historical information, perhaps a one-to-one correspondence could be made to established rating systems based on each credit category's rating criteria.

## 6.2 Credit quality migration

Credit rating migrations can be thought of as an extension of our model of firm defaults discussed in *Section 3* and illustrated again in *Chart 6.1*. We say that a firm has some underlying value – the value of its assets – and changes in this value suggest changes in credit quality. Certainly it is the case that equity prices drop precipitously as a firm moves towards bankruptcy. If we take the default likelihood as given by the credit rating of the firm, then we can work backwards to the "threshold" in asset value that delimits default. This is treated more formally in *Section 8.4*.

Likewise, just as our firm default model uses the default likelihood to place a threshold below which a firm is deemed to be in default, so also do the rating migration probabilities define thresholds beyond which the firm would be deemed to up(down)grade from its current credit rating. The data which drive this model are the default likelihood and credit rating migration likelihoods for each credit rating. We can compactly represent these rating migration probabilities using a transition matrix model (e.g., *Table 6.2*).

[2] Moody's terms these aggregated groupings "cohorts" and S&P terms them "static pools."

In essence, a transition matrix is nothing more than a square table of probabilities. These probabilities give the likelihood of migrating to any possible rating category (or perhaps default) one period from now given the obligor's credit rating today.

*Chart 6.1*
**Model of firm value and migration**

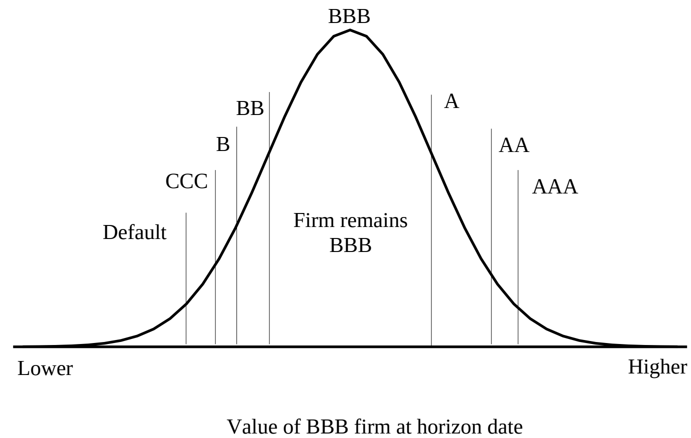


Many practical events (e.g., calls, enforced collateral provisions, spread resets) can be triggered by a rating change. These actions can directly affect the realized value within each credit rating category. For instance, a *pricing grid* – which predetermines a credit spread schedule given changes in credit rating – can reduce the volatility of value across up(down)grades.[3] Thus, we find it very convenient to explicitly incorporate awareness of rating migrations into our risk models.

## 6.3 Historical tabulation

We can tabulate historical credit rating migration probabilities by looking at time series of credit ratings over many firms. This technique is both powerful and limited. It is powerful in that we can freely model different volatilities of credit quality migration conditioned on the current credit standing. Said another way, each row in the transition matrix describes a volatility of credit rating changes that is unique to that row's initial credit rating. This is clearly an advantage since migration volatilities can vary widely between initial credit rating categories. There are, however, two assumption that we make about transition matrices. They are:

1. We assume that all firms tagged with the "correct" rating label. By this we mean that the rating agencies' are diligent in consistently applying credit rating standards across industries and countries (i.e., a "Baa" means the same for a U.S. electric utility as it does for a French bank). Of course, there is no reason that transition matrices could not be tabulated more specifically to reflect potential differences in the historical migration likelihoods of industries or countries. One caveat to this refinement might be the greater "noise" introduced by the smaller sample sizes.

[3] The securitised form of this structure is called a Credit-Sensitive Note (CSN) and is discussed in more detail in Das & Tufano [96].

2. We assume that all firms tagged with a given rating label will act alike. By this we mean that the full spectrum of credit migration likelihoods – not just the default likelihood – is similar for each firm assigned to a particular credit rating.

There are several sources of transition matrices, each specific to a particular credit rating service.[4] We advocate maintaining this correspondence even though it is common for practitioners to use shorthand assumptions, e.g., Moody's Baa is "just like" S&P's BBB, etc. Here we list three of these sources: Moody's, S&P, and KMV. Each is shown for the major credit rating categories – transition matrices which cover the minor (+/-) credit rating are also available, but are not shown here.

### *6.3.1 Moody's Investors Service transition matrix*

Moody's utilizes a data set of 26 years' worth of credit rating migrations over the issuers that they cover. These issuers are predominantly U.S.-based firms, but are including more and more international firms. The transition matrix is tabulated upon issuers conditioned on those issuers continuing to be rated at the end of the year. Thus there is no concern with having to adjust for a *no-longer-rated* "rating."

*Table 6.1*
**Moody's Investors Service: One-year transition matrix**

| **Initial Rating** | **Rating at year-end (%)** | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Aaa | Aa | A | Baa | Ba | B | Caa | Default |
| Aaa | 93.40 | 5.94 | 0.64 | 0 | 0.02 | 0 | 0 | 0 |
| Aa | 1.61 | 90.55 | 7.46 | 0.26 | 0.09 | 0.01 | 0 | 0.02 |
| A | 0.07 | 2.28 | 92.44 | 4.63 | 0.45 | 0.12 | 0.01 | 0 |
| 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 | 0.04 | 0.13 | 0.54 | 6.35 | 84.22 | 1.91 | 6.81 |
| Caa | 0 | 0 | 0 | 0.62 | 2.05 | 4.08 | 69.20 | 24.06 |

*Source: Lea Carty of Moody's Investors Service*

### *6.3.2 Standard & Poor's transition matrix*

It happens that the transition matrix published by Standard & Poor's includes a *no-longer-rated* "rating," and so we pause to discuss this issue. The majority of these withdrawals of a rating occur when a firm's only outstanding issue is paid off or its debt issuance program matures. Yet our assumption is that CreditMetrics will be applied to obligations with a known maturity. So there should be no N.R. category in application.

Thus, it makes sense to eliminate the N.R. category and gross-up the remaining percentages in some appropriate fashion. We do this as follows. Since S&P describes that they track bankruptcies even after a rating is withdrawn, the default probabilities are already fully tabulated. We believe that there is no systematic reason correlated with credit rat-

[4] KMV is not a credit rating service. They quantitatively estimate Expected Default Frequencies (EDF) which are continuous values rather than categorical labels using an option theoretic approach.

ing stating which would explain rating removals. We thus adjust all remaining migration probabilities on a *pro rata* basis as shown in *Table 6.2* below:

*Table 6.2*
**Standard & Poor's one-year transition matrix – adjusted for removal of N.R.**

| **Initial Rating** | **Rating at year-end (%)** | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | AAA | AA | A | BBB | BB | B | CCC | Default |
| AAA | 90.81 | 8.33 | 0.68 | 0.06 | 0.12 | 0 | 0 | 0 |
| AA | 0.70 | 90.65 | 7.79 | 0.64 | 0.06 | 0.14 | 0.02 | 0 |
| A | 0.09 | 2.27 | 91.05 | 5.52 | 0.74 | 0.26 | 0.01 | 0.06 |
| BBB | 0.02 | 0.33 | 5.95 | 86.93 | 5.30 | 1.17 | 0.12 | 0.18 |
| BB | 0.03 | 0.14 | 0.67 | 7.73 | 80.53 | 8.84 | 1.00 | 1.06 |
| B | 0 | 0.11 | 0.24 | 0.43 | 6.48 | 83.46 | 4.07 | 5.20 |
| CCC | 0.22 | 0 | 0.22 | 1.30 | 2.38 | 11.24 | 64.86 | 19.79 |

*Source: Standard & Poor's CreditWeek April 15, 1996*

Both of these tables are included in the CreditMetrics data set.

### *6.3.3 KMV Corporation transition matrix*

Both of the above transition matrices were tabulated by credit rating agencies. In contrast, the sample transition matrix shown in Table 6.3 was constructed from KMV EDFs (expected default frequency) for non-financial companies in the US using data from January 1990 through September 1995. Each month, the rating group based on the EDF of each company for that month was compared against the rating group it was in 12 months hence, based on its EDF at that date. This gave a single migration. There are an average of 4,780 companies in the sample each month, resulting in a total of 329,803 migration observations. Firms that disappeared from the sample were allocated into the rating categories proportionately to the population. Rating group #8 signifies default, which is treated as a a terminal event for the firm.

The purpose of this sample is to show how an alternative approach such as EDFs can be utilized to generate a transition matrix. EDFs are default probabilities measured on a continuous scale of 0.02% to 20.0%, but grouped into discrete "rating" ranges for application in CreditMetrics.

*Table 6.3*
**KMV one-year transition matrices as tabulated from expected default frequencies (EDFs)**

| Initial Rating | Rating at Year-end (%) | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 1 (AAA) | 2 (AA) | 3 (A) | 4 (BBB) | 5 (BB) | 6 (B) | 7 (CCC) | 8 (Default) |
| 1 (AAA) | 66.26 | 22.22 | 7.37 | 2.45 | 0.86 | 0.67 | 0.14 | 0.02 |
| 2 (AA) | 21.66 | 43.04 | 25.83 | 6.56 | 1.99 | 0.68 | 0.20 | 0.04 |
| 3 (A) | 2.76 | 20.34 | 44.19 | 22.94 | 7.42 | 1.97 | 0.28 | 0.10 |
| 4 (BBB) | 0.30 | 2.80 | 22.63 | 42.54 | 23.52 | 6.95 | 1.00 | 0.26 |
| 5 (BB) | 0.08 | 0.24 | 3.69 | 22.93 | 44.41 | 24.53 | 3.41 | 0.71 |
| 6 (B) | 0.01 | 0.05 | 0.39 | 3.48 | 20.47 | 53.00 | 20.58 | 2.01 |
| 7 (CCC) | 0.00 | 0.01 | 0.09 | 0.26 | 1.79 | 17.77 | 69.94 | 10.13 |

*Source: KMV Corporation*

*Table 6.3* is presented as an example and will not be included in the CreditMetrics data set. Subscribers to KMV's Expected Default Frequencies utilize a measure of default probability that is on a continues scale rather than discrete groupings offered by a credit rating agency.

Both KMV and we ourselves advocate that each credit rating (or expected default frequency) be addressed by a transition matrix tailored to that system. For this reason, the example KMV transition matrix shown here will not be part of the CreditMetrics data set. Only subscribers to KMV's expected default frequency (EDF) data would be users of such a transition matrix and so KMV will be offering it as part of that subscription.

Although it would be fine to have some issuers within a portfolio evaluated with one service (i.e., financials evaluated by IBCA) and other issuers evaluated by another service (i.e., corporates and industrials by Moody's, say), it would be inappropriate to mix systems (i.e., S&P ratings applied to Moody's transition matrix).

## 6.4 Long-term behavior

In estimating transition matrices, there are a number of desirable properties that one wants a transition to matrix to have, but which does not always follow from straightforward compilation of the historical data. In general, it is good practice to impose at least some of the desirable properties on the historical data in the form of estimation constraints.

The nature and extent of the problems encountered will be a function of the particular rating system, the number of grades considered, and the amount of historical data available. The following discussion uses S&P ratings as the basis for explaining these issues and how they can be addressed.

Historical tabulation is worthwhile in its own right. However, as with almost any type of sampling, it represents a limited amount of observation with sampling error. In addition to what we have historically observed, we also have strong expectations about credit rating migrations. For instance, over sufficient time we expect that any inconsistencies in rank order across credit ratings will disappear. By *rank order*, we mean a consistent progression in one direction such as default likelihoods always increasing – never then decreasing – as we move from high quality ratings to lower quality ratings. We list three potential short-term sampling error concerns here:

- Output cumulative default likelihoods should not violate proper rank order. For instance, *Table 6.4* below shows that AAAs have defaulted more often at the 10-year horizon than have AAs.

- Limited historical observation yields "granularity" in estimates. For instance, the AAA row in *Table 6.2* above is supported by 1,658 firm-years worth of observation. This is enough to yield a "resolution" of 0.06% (i.e., only probabilities in increments of 0.06% – or 1/1658 – are possible).

- This lack of resolution may erroneously suggest that some probabilities are identically zero. For instance, if there were truly a 0.01% chance of AAA default, then we would have to watch for another 80 years before there would be a 50% chance of tabulating a non-zero AAA default probability.

There are other potential problems with historical sampling such as the business cycle and regime shifts (e.g., the restructuring of the high-yield market in the 1980's). But these will not be addressed here.

*Table 6.4*
**Average cumulative default rates (%)**

| **Term** | **1** | **2** | **3** | **4** | **5 ...** | **7 ...** | **10 ...** | **15** |
|---|---|---|---|---|---|---|---|---|
| AAA | 0.00 | 0.00 | 0.07 | 0.15 | 0.24 ... | 0.66 ... | 1.40 ... | 1.40 |
| AA | 0.00 | 0.02 | 0.12 | 0.25 | 0.43 ... | 0.89 ... | 1.29 ... | 1.48 |
| A | 0.06 | 0.16 | 0.27 | 0.44 | 0.67 ... | 1.12 ... | 2.17 ... | 3.00 |
| BBB | 0.18 | 0.44 | 0.72 | 1.27 | 1.78 ... | 2.99 ... | 4.34 ... | 4.70 |
| BB | 1.06 | 3.48 | 6.12 | 8.68 | 10.97 ... | 14.46 ... | 17.73 ... | 19.91 |
| B | 5.20 | 11.00 | 15.95 | 19.40 | 21.88 ... | 25.14 ... | 29.02 ... | 30.65 |
| CCC | 19.79 | 26.92 | 31.63 | 35.97 | 40.15 ... | 42.64 ... | 45.10 ... | 45.10 |

*Source: S&P CreditWeek, Apr. 15, 1996*

### *6.4.1 Replicate historical cumulative default rates*

The major rating agencies have published tables of cumulative default likelihood over holding periods as long as 20 years – reported in annual increments. If we ignore for the moment the issue of autocorrelation, then it is generally true that *"there exists some annual transition matrix which best replicates (in a least squares sense) this default history."* Said another way, we can always work backwards from a cumulative default table to an implied transition matrix. *Table 6.4* illustrates part of a cumulative default probability table published by Moody's.

Cumulative default rate tables like this can be fit fairly closely by a single transition matrix.[5] Thus, it is apparently true that defaults over time are closely approximated by a transition matrix model.[6] This is an important result. It demonstrates that the statistical behavior of credit rating migrations can be captured through a transition matrix model. CreditMetrics uses a transition matrix to model credit rating migrations not only because it is intuitive but also because it is an extremely powerful statistical tool.

Below we show a transition matrix that has been created using *nothing but* a least squares fit to the cumulative default rates in *Table 6.4*. At this point, we are most interested in showing that: (i) such a matrix can be derived and (ii) that the process of defaults is closely replicated by a Markov process. (We make no claim that *Table 6.5* is a faithful replication of the historically tabulated *Table 6.2*.)

*Table 6.5*
**Imputed transition matrix which best replicates default rates**

| **Initial Rating** | **Rating at year end (%)** | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **AAA** | **AA** | **A** | **BBB** | **BB** | **B** | **CCC** | **Default** |
| AAA | 43.78 | 53.42 | 1.65 | 0.71 | 0.29 | 0.11 | 0.02 | 0.01 |
| AA | 0.60 | 90.60 | 6.20 | 1.45 | 0.93 | 0.16 | 0.04 | 0.01 |
| A | 0.22 | 2.84 | 92.97 | 3.12 | 0.56 | 0.14 | 0.07 | 0.07 |
| BBB | 2.67 | 3.29 | 12.77 | 75.30 | 5.07 | 0.60 | 0.14 | 0.17 |
| BB | 0.19 | 3.58 | 8.28 | 9.97 | 55.20 | 17.17 | 4.53 | 1.08 |
| B | 0.12 | 0.50 | 20.69 | 1.05 | 0.25 | 55.40 | 17.05 | 4.95 |
| CCC | 0.04 | 0.11 | 6.28 | 0.30 | 0.12 | 41.53 | 32.46 | 19.15 |

For comparison to *Table 6.4*, we show below in *Table 6.6* the cumulative default rates which result from this transition matrix. Again, the most important point is that *Table 6.4* and *Table 6.6* are quite close; thus the Markov process is a reasonable modeling tool. The median difference between them is 0.16% with a maximum error of 2.13%.

This "best fit" Markov process has yielded the side benefit of resolving non-intuitive rank order violations in its resulting cumulative default rates. For instance, our problem of AAA's having a 10 year default rate that was *greater* than AA's is now gone. This behavior – of non-crossing default likelihoods – is a feature that we would expect given very long sampling histories.

[5] Empirically, a transition matrix fit is not as good for cumulative default rates of *newly issued* debt (as opposed to the total debt population) due to a "seasoning" effect where sub-investment grades have an unusually low default likelihood in the first few years. This "seasoning" problem has not been apparent for bank facilities.

[6] A transition matrix model is an example of a *Markov Process*. A Markov Process is a state-space model which allows the next progression to be determined only by the current state and not information of previous states.

*Table 6.6*
**Resulting cumulative default rates from imputed transition matrix (%)**

| Term | 1 | 2 | 3 | 4 | 5 | ... | 7 | ... | 10 | ... | 15 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| AAA | 0.01 | 0.04 | 0.09 | 0.18 | 0.31 | ... | 0.66 | ... | 1.37 | ... | 2.81 |
| AA | 0.01 | 0.06 | 0.15 | 0.27 | 0.44 | ... | 0.85 | ... | 1.63 | ... | 3.12 |
| A | 0.07 | 0.17 | 0.30 | 0.46 | 0.65 | ... | 1.11 | ... | 1.94 | ... | 3.50 |
| BBB | 0.17 | 0.41 | 0.78 | 1.25 | 1.79 | ... | 2.95 | ... | 4.60 | ... | 6.83 |
| BB | 1.08 | 3.41 | 6.14 | 8.76 | 11.05 | ... | 14.53 | ... | 17.71 | ... | 20.39 |
| B | 4.95 | 10.97 | 15.75 | 19.33 | 21.98 | ... | 25.46 | ... | 28.19 | ... | 30.35 |
| CCC | 19.15 | 27.43 | 32.63 | 36.32 | 39.01 | ... | 42.49 | ... | 45.14 | ... | 47.05 |

### *6.4.2 Monotonicity (non-crossing) barrier likelihoods*

Cumulative default rates are just a special case of what we term “barrier” likelihoods. In general, we can ask, “what is the cumulative rate of crossing any given level of credit quality?” For instance, if we managed a portfolio which was not allowed to invest in sub-investment grade bonds, then we might be interested in the likelihood of any credit quality migrations which were to or across the BB rating barrier. The cumulative probabilities for crossing the “BB barrier” using the transition matrix in *Table 6.5* are as shown in *Table 6.7*. Notice that monotinicity (rank order) is violated for single-As.

*Table 6.7*
**“BB barrier” probabilities calculated from *Table 6.6* matrix (%)**

| Term | 1 | 2 | 3 | 4 | 5 | ... | 7 | ... | 10 | ... | 15 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **AAA** | 0.46 | 1.40 | 2.54 | 3.80 | 5.09 | ... | 7.74 | ... | 11.71 | ... | 18.13 |
| **AA** | 1.25 | 2.54 | 3.85 | 5.17 | 6.51 | ... | 9.17 | ... | 13.12 | ... | 19.47 |
| **A** | 0.91 | 2.00 | 3.20 | 4.49 | 5.82 | ... | 8.57 | ... | 12.69 | ... | 19.29 |
| **BBB** | 6.57 | 11.66 | 15.69 | 18.93 | 21.60 | ... | 25.78 | ... | 30.40 | ... | 36.25 |

Just as we would expect very long-term historical observation to resolve violations of non-intuitive cumulative default rank order, we should expect resolution of barrier rank ordering. This table above shows that our imputed transition matrix violates this anticipated long-term behavior.

We can now replay the least squares fit we performed when we produced *Table 6.4* with the added constraint that all possible barrier probabilities must also be in rank order. *Table 6.8* shows these same BB barrier probabilities with our new fit. (In fact, there are six non-default “barriers” for seven rating categories and our fitting algorithm addressed them all.)

*Table 6.8*
**“BB barrier” probabilities calculated from *Table 6.6* matrix (%)**

| Term | 1 | 2 | 3 | 4 | 5 | ... | 7 | ... | 10 | ... | 15 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **AAA** | 0.39 | 1.09 | 1.98 | 3.01 | 4.12 | ... | 6.52 | ... | 10.37 | ... | 16.97 |
| **AA** | 1.07 | 2.19 | 3.36 | 4.57 | 5.82 | ... | 8.39 | ... | 12.36 | ... | 19.01 |
| **A** | 1.13 | 2.42 | 3.82 | 5.29 | 6.80 | ... | 9.88 | ... | 14.48 | ... | 21.73 |
| **BBB** | 5.88 | 10.72 | 14.77 | 18.18 | 21.11 | ... | 25.89 | ... | 31.34 | ... | 38.13 |

This refinement was achieved with minimal change in the transition matrix’s fit to the cumulative default rates. The differences in predicted cumulative default rates averages only 0.06% (median is 0.02%) between the two fitted transition matrices. For comparison with *Table 6.5*, we show this new fit of our imputed transition matrix.

*Table 6.9*
**Imputed transition matrix with default rate rank order constraint**

| | ---Rating at year end (%)--- | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| **Initial Rating** | **AAA** | **AA** | **A** | **BBB** | **BB** | **B** | **CCC** | **Default** |
| **AAA** | 58.57 | 39.02 | 1.42 | 0.63 | 0.18 | 0.14 | 0.03 | 0.01 |
| **AA** | 0.71 | 89.45 | 7.47 | 1.39 | 0.72 | 0.18 | 0.05 | 0.02 |
| **A** | 0.25 | 3.83 | 91.15 | 3.73 | 0.77 | 0.14 | 0.07 | 0.06 |
| **BBB** | 2.07 | 2.26 | 10.03 | 80.29 | 4.53 | 0.50 | 0.15 | 0.18 |
| **BB** | 0.15 | 3.57 | 7.84 | 10.38 | 55.91 | 16.18 | 4.91 | 1.06 |
| **B** | 0.14 | 0.62 | 19.21 | 2.44 | 0.55 | 54.87 | 17.24 | 4.94 |
| **CCC** | 0.04 | 0.14 | 5.85 | 0.77 | 0.33 | 41.10 | 32.65 | 19.14 |

Perhaps the difference between *Table 6.5* and *Table 6.9* is that the weight of probabilities are generally moved towards the upper-left to lower-right diagonal. Also, without directly trying, we are moving towards a better approximation of the historical transition matrix shown in *Table 6.2*.

### *6.4.3 Steady state profile matches debt market profile*

Another desirable dimension of “fit” for a transition matrix is for it to exhibit a long-term steady state that approximates the observed profile of the overall credit markets. By this we mean that – among those firms which do not default – there will be some distribution of their credit quality across the available credit rating categories. To represent the rating profile across the bond market, we have taken the following data (*Table 6.10*) from Standard & Poor’s *CreditWeek* April 15, 1996.

*Table 6.10*
**Estimate of debt market profile across credit rating categories**

| S&P 1996 | AAA | AA | A | BBB | BB | B | CCC |
|---|---|---|---|---|---|---|---|
| **Count** | 85 | 200 | 487 | 275 | 231 | 87 | 13 |
| **Proportion** | 6.2% | 14.5% | 35.3% | 20.0% | 16.8% | 6.3% | 0.9% |

Mathematically, our transition matrix Markov process will have two long-term properties (i.e., more than 100 periods). First, since default is an *absorbing* state, eventually all firms will default. Second, since the initial state has geometrically less influence on future states, the profile of non-defaulted firms will converge to some steady state regardless of the firm's initial rating.

As the chart below shows, our fitting algorithm can achieve a closer approximation of the anticipated long-term steady state. The transition matrix in *Table 6.9* shows too strong a tendency to migrate towards single-A. Once we add an incentive to fit the anticipated steady state, we see that a more balanced profile is achieved.

*Chart 6.2*
**Achieving a closer fit to the long-term steady state profile**

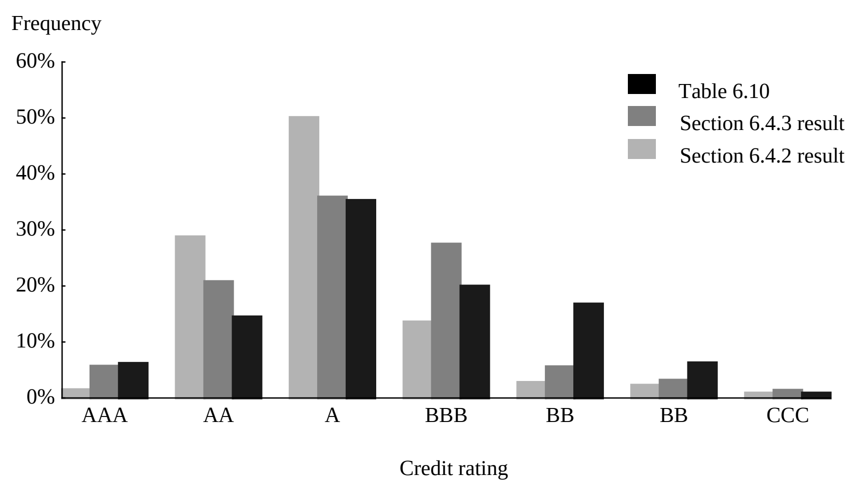


This additional soft constraint was accomplished with a negligible effect on the matrix's ability to replicate cumulative default likelihoods – and monotonicity in the barrier was still fully realized. Also, without directly trying, we are moving towards a better approximation of the historical transition matrix shown in *Table 6.2*.

### *6.4.4 Monotonicity (smoothly changing) transition likelihoods*

Though it is certainly not a requirement of a transition matrix, our expectation is that there is a certain rank ordering the likelihood of migrations as follows:

1. Better ratings should never have a higher chance of default;
2. The chance of migration should become less as the migration distance (in rating notches) becomes greater; and

3. The chance of migrating to a given rating should be greater for more closely adjacent rating categories.

As an example, we will refer to *Table 6.10*. Since the default likelihoods ascend smoothly there is no violation of #1. However, since the chance that a single-B would migrate to a single-A is greater than either a migration to BBB or BB, there is a "violation" of #2. Also, since single-B has a greater chance of migrating to single-A than does an initial BB or BBB, there is a "violation" of rule #3. The reader can find other probabilities in this table which are not monotonic in our definition.

As before, we could add the soft constraint that our fitting algorithm should endeavor to mitigate these non-rank orderings of probabilities as it seeks to replicate the cumulative default likelihoods. However, as we discuss next, there is one last source of data that we should use in best estimating our transition matrix – an historically tabulated transition matrix. Any fitting algorithm that addresses smooth transition likelihoods would have to revisit these same probabilities when it includes knowledge of the historically tabulated transition matrix. So we address them both together below.

### *6.4.5 Match historically tabulated transition matrix*

Standard & Poor's historically tabulated transition matrix was shown above in *Table 6.2*. Up to now we have discussed some of the characteristics of transition matrices and methods of addressing these. Now we will bring all this together in *Table 6.11* to give an estimate of a one-year transition matrix which is rooted in the historical data and is also sensitive to our expectation of long-term behavior.

*Table 6.11*
**Achieving a closer fit to the long-term steady state profile**

| **Initial Rating** | **Rating at year end (%)** | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **AAA** | **AA** | **A** | **BBB** | **BB** | **B** | **CCC** | **Default** |
| AAA | 87.74 | 10.93 | 0.45 | 0.63 | 0.12 | 0.10 | 0.02 | 0.02 |
| AA | 0.84 | 88.23 | 7.47 | 2.16 | 1.11 | 0.13 | 0.05 | 0.02 |
| A | 0.27 | 1.59 | 89.05 | 7.40 | 1.48 | 0.13 | 0.06 | 0.03 |
| BBB | 1.84 | 1.89 | 5.00 | 84.21 | 6.51 | 0.32 | 0.16 | 0.07 |
| BB | 0.08 | 2.91 | 3.29 | 5.53 | 74.68 | 8.05 | 4.14 | 1.32 |
| B | 0.21 | 0.36 | 9.25 | 8.29 | 2.31 | 63.89 | 10.13 | 5.58 |
| CCC | 0.06 | 0.25 | 1.85 | 2.06 | 12.34 | 24.86 | 39.97 | 18.60 |

This transition matrix is meant to be close to the historically tabulated probabilities while being adjusted somewhat to better approximate the long-term behaviors we have discussed in this section. From a risk estimation standpoint we see that there are now small but non-zero probabilities of default imputed for AAAs and AAs.

# Chapter 7. Recovery rates

Residual value estimation in the event of default is inherently difficult. At the time when a banker makes a loan or an investor buys a bond, it is in the belief not that the obligor will go bankrupt but that the instrument will outperform. So it can be especially difficult to imagine what the obligor's position will be in the unlikely event of default. Will it be a catastrophe which leaves no value to recover, or will it be a regrettable but well behaved wrapping up which affects only shareholders but leaves debt holders whole?

It is in the remote chance of an outright default that a credit instrument will realize its greatest potential loss. Across a typical portfolio, most of the credit risk will be attributable to default events. Investment grade credits will have relatively more of their volatility attributable to credit spread moves versus sub-investment grade credits, which will be primarily driven by potential default events. However, a typical portfolio will have a mixture of each, with most of the portfolio risk coming from the sub-investment grades. So the magnitude of any recovery rate in default is important to model diligently.

The academic literature in our bibliography focuses primarily on U.S. defaults post October 1, 1979 – the effective date of the 1978 Bankruptcy Reform Act. However, the general finding that recovery rates are highly uncertain with a distribution that can be modeled is applicable internationally.

In this chapter, we will discuss not only the estimation of mean expected recovery rate in default, but the important problem of addressing the wide uncertainty of recovery rate experience. This chapter is organized as follows:

- estimating recovery rate distributions, their mean and standard deviation, by seniority level and exposure type; and

- fitting a full distribution to recovery rate statistics while preserving the required 0% to 100% bounds.

We have seen much effort devoted to estimating recovery rates based on: (i) seniority ranking of debt, (ii) instrument type or use, (iii) credit rating *X*-years before default, and (iv) size and/or industry of the obligor. But the most striking feature of any historical recovery data is its wide uncertainty. Any worthwhile credit risk model must be able to incorporate recovery rate uncertainty in order to fully capture the volatility of value attributable to credit. However, once we contemplate volatility of recoveries, we must also address any potential correlation of recoveries across a portfolio. In this section, we estimate any potential correlation of recoveries across the book.

### 7.1 Estimating recovery rates

There are many practical problems in estimating recovery rates of debt in the event of default. Often there is no market from which to observe objective valuations, and if there are market prices available they will necessarily be within a highly illiquid market. Even if these issues are resolved there is the question of whether it is best to estimate values: (i) immediately upon announcement of default, (ii) after some reasonable period for information to become available – perhaps a month, or (iii) after a full settlement has been reached – which can take years.

Since there have been academic studies, see Eberhart & Sweeney [92], which conclude that the bond market efficiently prices future realized liquidation values, we take comfort in those studies which poll/estimate market valuations about one month after the announcement of default. This certainly is the value which an active investor would face whether or not he chose to hold his position after the default event.[1]

We look to the following independent studies for use in CreditMetrics. These studies refine their estimates of recovery rate according to seniority type among bonds. Among bank facilities (e.g., loans, commitments, letter of credit) the studies have viewed these as a separate "seniority" class. It is clear from the data that the historical loan recovery rates have been higher than recovery rates for senior bonds. It is not clear whether this is attributable to differences in relationship, use of borrowing or security.

### *7.1.1 Recovery rates of bonds*

For corporate bonds, we have two primary studies of recovery rate which arrive at similar estimates (see Carty & Lieberman [96a] and Altman & Kishore [96]). For bond recoveries we can look primarily to Moody's 1996 study of recovery rates by seniority class. This study has the largest sample of defaulted bond that we know of. *Table 7.1* is a partial representation of Table 5 from Moody's Investors Service, which shows statistics for defaulted bond prices (1/1/70 through 12/31/1995).

*Table 7.1*
**Recovery statistics by seniority class**
*Par (face value) is $100.00.*

| | Carty & Lieberman [96a] | | | Altman & Kishore [96] | | |
|---|---|---|---|---|---|---|
| **Seniority Class** | **Number** | **Average** | **Std. Dev.** | **Number** | **Average** | **Std. Dev.** |
| Senior Secured | 115 | $53.80 | $26.86 | 85 | $57.89 | $22.99 |
| Senior Unsecured | 278 | $51.13 | $25.45 | 221 | $47.65 | $26.71 |
| Senior Subordinated | 196 | $38.52 | $23.81 | 177 | $34.38 | $25.08 |
| Subordinated | 226 | $32.74 | $20.18 | 214 | $31.34 | $22.42 |
| Junior Subordinated | 9 | $17.09 | $10.90 | — | — | — |

As this table shows, the subordinated classes are appreciably different from one another in their recovery realizations. In contrast, the difference between secured versus unsecured senior debt is not statistically significant. It is likely that there is a self-selection effect here. There is a greater chance for security to be requested in the cases where an underlying firm has questionable hard assets from which to salvage value in the event of default.

There is no public study we are aware of that seeks to isolate the effects of different levels of security controlling for the asset quality of the obligor firm. It becomes then a

[1] There are two studies, see Swank & Root [95] and Ward & Griepentrog [93], that report high average holding period returns for debt held between the default announcement and the ultimate bankruptcy resolution. These studies also note the high average uncertainty of returns and thus the market's risk pricing efficiency.

practical problem for the risk manager to judge on a bond-by-bond basis what adjustment is best made to recovery rate estimates for different levels of security.[2]

### *7.1.2 Recovery rate of bank facilities*

For bank facilities, we again have two primary studies of recovery rate which arrive at similar estimates see (Asarnow & Edwards [95] and Carty & Lieberman [96b]). A&E track 831 commercial and industrial loan defaults plus 89 structured loans while C&L track 58 defaults of loans with Moody's credit ratings. Both studies treat bank facilities as essentially a seniority class of their own – with this being senior to all public bond seniority classes.

Moody's reports a 71% mean and 77% median recovery rate which is within sampling error of Asarnow & Edwards 65.21% mean and 78.79% median recovery rates. So these two studies are different by no more than 5%.

*Chart 7.1* below is reproduced from A&E, and we have used it to estimate the standard deviation of recovery rates, of 32.7%, which is beyond the information reported by A&E.

*Chart 7.1*
**Distribution of bank facility recoveries**

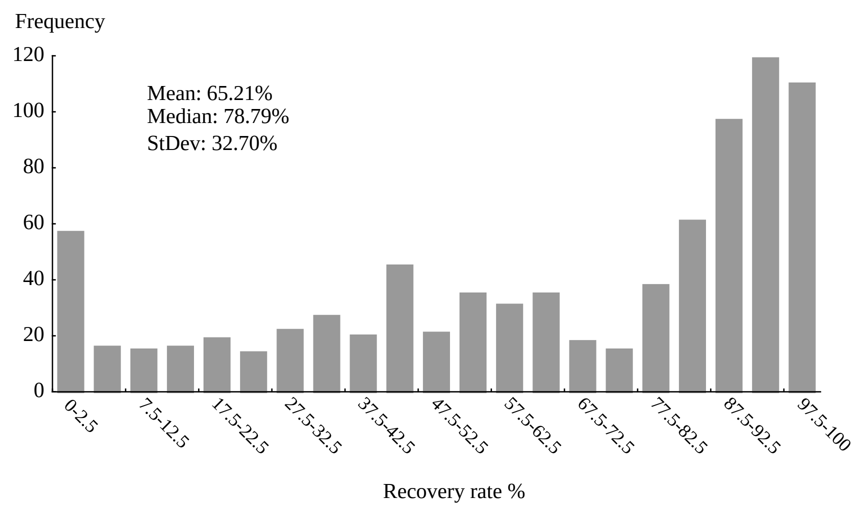


*Source: Asarnow & Edwards [95]*

A legitimate concern is that all of the studies referenced here are either exclusively based on, or primarily driven by, U.S. bankruptcy experience. Since bankruptcy law and practice differs from jurisdiction to jurisdiction (and even across time within a jurisdiction), it is not clear that these historical estimates of recovery rate will be directly applicable internationally.

[2] For this reason, our software implementation of CreditMetrics, CreditManager, will allow recovery rate estimates to be overwritten on the individual exposure level

## 7.2 Distribution of recovery rate

Recovery rates are best characterized, not by the predictability of their mean, but by their consistently wide uncertainty. Loss rates are bounded between 0% and 100% of the amount exposed. If we did not know anything about recovery rate, that is, if we thought that all possible recovery rates were equally likely, then we would model them as a flat (i.e., uniform) distribution between the interval 0 to 1. Uniform distributions have a mean of 0.5 and a standard deviation of 0.29 ($\sigma = \sqrt{1/12}$). The standard deviations of 25.45% for senior unsecured bonds and 32.7% for bank facilities are on either side of this and so represent relatively high uncertainties.

We can capture this wide uncertainty and the general shape of the recovery rate distribution – while staying within the bounds of 0% to 100% – by utilizing a *beta distribution*. Beta distributions are flexible as to their shape and can be fully specified by stating the desired mean and standard deviation. *Chart 7.2* illustrates beta distributions for different seniority classes using some of statistics reported in *Table 7.1*.

*Chart 7.2*
**Example beta distributions for seniority classes**

.

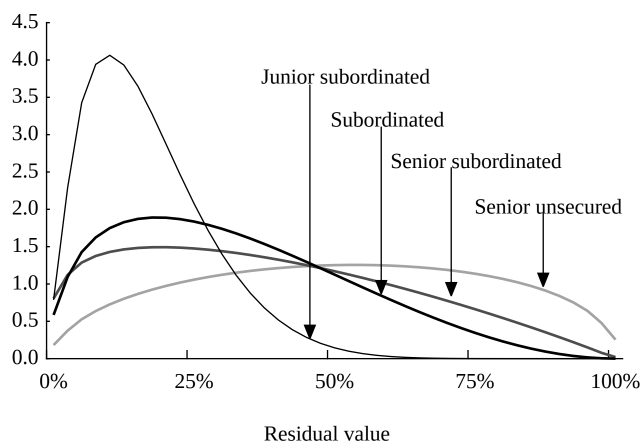


This full representation of the distribution is unnecessary for the analytic engine of CreditMetrics. It is used later in our simulation framework, where the shape of the full distribution is required.

# Chapter 8. Credit quality correlations

Central to our view of credit risk estimation is a diligent treatment of the portfolio effect of credit. Whereas market risks can be diversified with a relatively small portfolio or hedged using liquid instruments, credit risks are more problematic. For credit portfolios, simply having many obligors' names represented within a portfolio does not assure good diversification (i.e., they may all be large banks within one country). When diversification is possible, it typically achieved by much larger numbers of exposures than for market portfolios.

The problem of constructing a Markowitz-type portfolio aggregation of credit risk has only recently been widely examined. We know of two academic papers which address the problems of estimating correlations within a credit portfolio: Gollinger & Morgan [93] use time series of default likelihoods (Zeta-Scores™ published by Zeta Services, Inc.) to correlate across 42 constructed indices of industry default likelihoods, and Stevenson & Fadil [95] correlate the default experience, as listed in Dun & Bradstreet's *Business Failure Record*, across 33 industry groups. Both of these studies note the practical difficulties of estimating default correlations.

Our portfolio treatment of credit risk was greatly influenced by various engagements with KMV, which has studied models of credit correlations for a number of years.

The structure of this chapter is as follows:

- First, we discuss evidence from default histories which supports our assertion that credit correlations actually exist.
- Second, we investigate the possibility of modeling joint rating changes directly using historical rating change data.
- Third, we discuss the estimation of credit correlations through the observation of bond spread histories.
- Fourth, we present a model which connects rating changes and defaults to movements in an obligor's asset value. This allows us to model joint rating changes across multiple obligors without relying on historical rating change or bond spread data.
- Last, we discuss methods to estimate the parameters of the asset value model, and present a dataset for this purpose.

### 8.1 Finding evidence of default correlation

In this section, before moving on to modeling correlations in credit rating changes, we examine several histories of rating changes and defaults in order to establish that such correlations in fact exist. One might claim that each firm is in many ways unique and its changes in credit quality often are driven by events and circumstances specific to that firm; this would argue for little correlation between different firms' rating changes and defaults. Thus, it would be desirable for us to first find evidence of defaults across a large body of companies.

We can do this by examining the default statistics reported by the major rating agencies over many years. Since the studies we consider are based on a very large number of observations, if defaults were uncorrelated, then we would expect to observe default rates which are very stable from year to year. On the other hand, if defaults were perfectly correlated, then we would observe some years where every firm in the study defaults and others where no firms default. That our observations lie somewhere between these two extremes (that is, we observe default rates which fluctuate, but not to the extent that they would under perfect correlation) is evidence that some correlation exists. We make this observation more precise below.

We will use the formula below to compute average default correlation $\rho$ from the data; for a full derivation, see *Appendix F*.

$$[8.1] \qquad \rho = \frac{N\left(\dfrac{\sigma^2}{\mu - \mu^2}\right) - 1}{N - 1} \approx \frac{\sigma^2}{\mu - \mu^2}$$

where the approximation is for large values of $N$, the number of names covered by the data. In the formula, $\mu$ denotes the average default rate over the years in the study and $\sigma$ is the standard deviation of the default rates observed from year to year.

Both Moody's and S&P publish default rate statistics which could be used to make this type of statistical inference of average default correlations. In *Table 8.1*, we use data from Tables 3 and 6 from Moody's most recent default study (see Carty & Lieberman [96a]).

We can infer from these figures that the number of firm-years supporting the default rate, $\mu$, is in the thousands for all credit rating categories. Thus, our approximation formula for $\rho$ is appropriate. However, there are only 25 yearly observations supporting the calculation of $\sigma$ (and it is reported with significant rounding), so the confidence levels around the resulting inferred correlation will be high.

*Table 8.1*
**Inferred default correlations with confidence levels**

| Credit rating category | Default rate $\mu$ | Standard deviation defaults $\sigma$ | Implied default correlation $\rho$ | Lower confidence $Pr\{\rho < X\} = 2.5\%$ | Upper confidence $Pr\{\rho > X\} = 2.5\%$ |
|---|---|---|---|---|---|
| Aa | 0.03% | 0.1% | 0.33% | 0.05% | 1.45% |
| A | 0.01% | 0.1% | 1.00% | 0.15% | 4.35% |
| baa | 0.13% | 0.3% | 0.69% | 0.29% | 1.83% |
| ba | 1.42% | 1.4% | 1.40% | 0.79% | 2.91% |
| B | 7.62% | 4.8% | 3.27% | 1.95% | 6.47% |

*Source: Moody's 1970-1995 1-year default rates and volatilities (Carty & Lieberman [96a])*

There are at least four caveats to this approach:

- the standard deviations of default rates, $\sigma$, are calculated over a very limited number of observations which lead to wide confidence levels;

- the underlying periodic default rates for investment grade categories are not normally distributed; thus the confidence levels for the investment grades will be wider than those calculated;

- the average default rate, $\mu$, is assumed to be constant across all firms within the credit rating category and constant across time; and

- the approach is sensitive to the proportion of recession versus growth years which – in the 25-year sample – may not be representative of the future.

The inferred default correlations are all positive and – using the confidence interval technique discussed above – are all statistically greater than zero to at least the 97.5% level. This is a fairly objective indication that default events have statistically significant correlations which cannot be ignored in a risk assessment model such as CreditMetrics.

In fact, our needs go beyond estimations of default correlations; we must estimate the joint likelihood of any possible combination of credit quality outcomes. Thus, if the credit rating system recognizes eight states (i.e., *AAA, AA, …, CCC* plus *Default*), then between two obligors there are 8•8 or 64 possible joint states whose likelihoods must to be estimated.

## 8.2 Direct estimation of joint credit moves

Perhaps the most direct way to estimate joint rating change likelihoods is to examine credit ratings time series across many firms which are synchronized in time with each other. We have done this with a sample of 1,234 firms who have senior unsecured S&P credit ratings reported quarterly for as much as the last 40 quarters. We note that this data set does not include much of the default experience that S&P reports in their more comprehensive studies and stress that we have assembled this data set only to illustrate the principle that joint credit quality migration likelihoods can be estimated directly. With this method, it is possible to avoid having to specify a correlation estimate and an accompanying descriptive model.

Since we are interested in tabulating all possible pairwise combinations between firms, there are over 1.13 million pairwise combinations within our particular sample. In general, if a rating series data set offers $N$ observations in a tabulated transition matrix then it will offer on the order of $N^2$ observations of joint migration. For a rating system with seven non-default categories, there will be 28 unique joint likelihood tales. In *Table 8.2* we show one of these 28 tabulated results for the case where one firm starts the period as a BBB and another firm starts the period with a single-A rating.

*Table 8.2*
**Historically tabulated joint credit quality co-movements**

| Firm starting in BBB | Firm starting in A | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | AAA | AA | A | BBB | BB | B | CCC | Default |
| AAA | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| AA | 0 | 15 | 1,105 | 54 | 4 | 0 | 0 | 0 |
| A | 0 | 978 | 44,523 | 2,812 | 414 | 224 | 0 | 0 |
| BBB | 0 | 12,436 | 621,477 | 40,584 | 5,075 | 2,507 | 0 | 0 |
| BB | 0 | 839 | 41,760 | 2,921 | 321 | 193 | 0 | 0 |
| B | 0 | 175 | 7,081 | 532 | 76 | 48 | 0 | 0 |
| CCC | 0 | 55 | 2,230 | 127 | 18 | 15 | 0 | 0 |
| Default | 0 | 29 | 981 | 67 | 7 | 0 | 0 | 0 |

This yields our non-parametric estimate of joint credit quality probabilities to be as in *Table 8.3*:

*Table 8.3*
**Historically tabulated joint credit quality co-movement (%)**

| Firm starting in BBB | Firm starting in A | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | AAA | AA | A | BBB | BB | B | CCC | Default |
| AAA | - | - | - | - | - | - | - | - |
| AA | - | 0.00 | 0.14 | 0.01 | 0.00 | - | - | - |
| A | - | 0.12 | 5.64 | 0.36 | 0.05 | 0.03 | - | - |
| BBB | - | 1.57 | 78.70 | 5.14 | 0.64 | 0.32 | - | - |
| BB | - | 0.11 | 5.29 | 0.37 | 0.04 | 0.02 | - | - |
| B | - | 0.02 | 0.90 | 0.07 | 0.01 | 0.01 | - | - |
| CCC | - | 0.01 | 0.28 | 0.02 | 0.00 | 0.00 | - | - |
| Default | - | 0.00 | 0.12 | 0.01 | 0.00 | - | - | - |

We emphasize again that this illustration is only to demonstrate a technique for estimating joint credit quality migration likelihoods directly. Unfortunately, our own access to the rating agency's data sets is inadequate to fully estimate a production quality study.

This method of estimation has the advantage that it does not make assumptions as to the underlying process, the joint distribution shape, or rely on distilling the data down to a single parameter – the correlation. However, it carries the limitation of treating all firms with a given credit rating as identical. So two banks would be deemed to have the same relationship as a bank and an oil refiner. In the following sections, we discuss a method of estimating credit quality correlations which are sensitive to the characteristics of individual firms.

## 8.3 Estimating credit quality correlations through bond spreads

A second way to estimate credit quality correlations using historical data would be to examine price histories of corporate bonds. Because it is intuitive that movements in bond prices reflect changes in credit quality, it is reasonable to believe that correlations of bond price moves might allow for estimations of correlations of credit quality moves. Such an approach has two requirements: adequate data on bond price histories and a model relating bond prices to credit events.

Where bond price histories are available, it is possible to estimate some type of credit correlation by first extracting credit spreads from the bond prices, and then estimating the correlation in the movements of these spreads. It is important to note that such a correlation only describes how spreads tend to move together. To arrive at the parameters we require for CreditMetrics (that is, likelihoods of joint credit quality movements), it is necessary to adopt a model which links spread movements to credit events.

Models of risky bonds typically have three state variables: the first is the risk free interest rate, the second is the credit spread, and the third indicates whether the bond has defaulted. A typical approach (see for example Duffee [95] or Nielsen and Ronn [94]) is to assume that the risk free rate and credit spread evolve independently[1] and that defaults are linked to the credit spread through some pricing model. This pricing model allows us to infer the probability of the issuer defaulting from the observed bond spread[2]. An extension of this type of model to two or more bonds would allow for the inference of default correlations from the correlation in bond spread moves.

While an approach of this type is attractive because it is elegant and consistent with other models of risky assets, its biggest drawback is practical. Bond spread data is notoriously scarce, particularly for low credit quality issues, making the estimation of bond spread correlations impossible in practice.

## 8.4 Asset value model

In this section, we present the approach which we introduced in *Chapter 3* and which we will use in practice to model joint probabilities of upgrades, downgrades, and defaults (all of which will be referred to generically as credit rating changes). We are motivated to pursue such an approach by the fact that practical matters (such as the lack of data on joint defaults) make it difficult to estimate such probabilities directly. Our approach here then will be indirect. It involves two steps:

1. Propose an underlying process which drives credit rating changes. This will establish a connection between the events which we ultimately want to describe (rating changes), but which are not readily observable, and a process which we understand and can observe.
2. Estimate the parameters for the process above. If we have been successful in the first part, this should be easier than estimating the joint rating change probabilities directly.

In this section, we propose that a firm's asset value be the process which drives its credit rating changes and defaults. This model is essentially the option theoretic model of Merton [74], which is discussed further in Kealhofer [95]. We describe the model which links changes in asset values to credit rating changes and explain how we parameterize

[1] The evolution of these quantities is generally modeled by diffusion processes with some drift and volatility.

[2] This is similar to the inference of implied volatilities from observed option premiums.

the asset value model. We postpone the discussion of parameter estimation to *Section 8.5*.

It is evident that the value of a company's assets determines its ability to pay its debt holders. We may suppose then that there is a specific level such that if the company's assets fall below this level in the next year, it will be unable to meet its payment obligations and will default. Were we only treating value changes due to default, this would be a sufficient model. However, since we wish to treat portfolio value changes resulting from changes in credit rating as well, we need a slightly more complex framework.

Extending the intuition above, we assume there is a series of levels for asset value that will determine a company's credit rating at the end of the period in question. For example, consider a hypothetical company that is BB rated and whose assets are currently worth \$100 million. Then the assumption is that there are asset levels such that we can construct a mapping from asset value in one year's time to rating in one year's time, as in *Chart 8.1*. Essentially, the assumption is that the asset value in one year determines the credit rating (or default) of the company at that time. The asset values in the chart which correspond to changes in rating will be referred to as asset value thresholds. We reiterate that we are not yet claiming to know what these thresholds are, only that this relationship exists.

*Chart 8.1*
**Credit rating migration driven by underlying BB firm asset value**

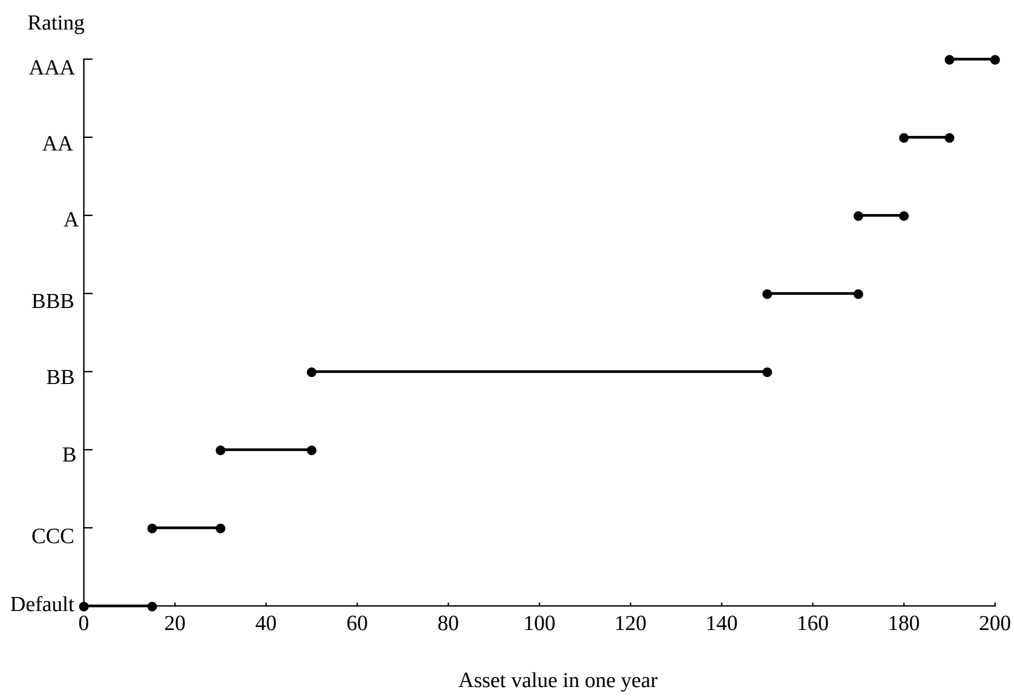


Assuming we know the asset thresholds for a company, we only need to model the company's change in asset value in order to describe its credit rating evolution. To do this, we assert that the percent changes in asset value (that is, asset "returns," which we will denote by $R$) are normally distributed, and parameterized by a mean $\mu$ and standard deviation (or volatility) $\sigma$. Note that this volatility is not the volatility of value of a credit instrument (which is an output of CreditMetrics) but simply the volatility of asset returns for a given name. For ease of exposition, we will assume $\mu=0$[3].

Given this parameterization of the asset value process, we may now establish a connection between the asset thresholds in the chart above and the transition probabilities for our company. Continuing with our example of the BB rated obligor, we read from the transition matrix that the obligor's one-year transition probabilities are as in the second column of *Table 8.4*.

On the other hand, from the discussion of asset thresholds above, we know that there exist asset return thresholds $Z_{Def}$, $Z_{CCC}$, $Z_{BBB}$, etc., such that if $R< Z_{Def}$, then the obligor goes into default; if $Z_{Def}<R<Z_{CCC}$, then the obligor is downgraded to CCC; and so on. So for example, if $Z_{Def}$ were equal to -70%, this would mean that a 70% (or greater) decrease in the asset value of the obligor would lead to the obligor's default.

Since we have assumed that $R$ is normally distributed, we can compute the probability that each of these events occur:

[8.2]
$$Pr\{Default\} = Pr\{R < Z_{Def}\} = \Phi(Z_{Def}/\sigma) ,$$
$$Pr\{CCC\} = Pr\{Z_{De(f<}\ )R < Z_{CCC}\} = \Phi(Z_{CCC}/\sigma) - \Phi(Z_{Def}/\sigma) ,$$

and so on. ($\Phi$ denotes the cumulative distribution for the standard normal distribution.) These probabilities are listed in the third column of *Table 8.4*.

*Table 8.4.*
**One year transition probabilities for a BB rated obligor**

| Rating | Probability from the transition matrix (%)( | Probability according to the asset value model |
|---|---|---|
| AAA | 0.03 | $1-\Phi(Z_{AA}/\sigma)$ |
| AA | 0.14 | $\Phi(Z_{AA}/\sigma)-\Phi(Z_{A}/\sigma)$ |
| A | 0.67 | $\Phi(Z_{A}/\sigma)-\Phi(Z_{BBB}/\sigma)$ |
| BBB | 7.73 | $\Phi(Z_{BBB}/\sigma)-\Phi(Z_{BB}/\sigma)$ |
| BB | 80.53 | $\Phi(Z_{BB}/\sigma)-\Phi(Z_{B}/\sigma)$ |
| B | 8.84 | $\Phi(Z_{B}/\sigma)-\Phi(Z_{XXX}/\sigma)$ |
| CCC | 1.00 | $\Phi(Z_{XXX}/\sigma)-\Phi(Z_{Def}/\sigma)$ |
| Default | 1.06 | $\Phi(Z_{Def}/\sigma)$ |

The connection between asset returns and credit rating may be represented schematically as in *Chart 8.2*, where we present the return thresholds superimposed on the distribution of asset returns. The integral between adjacent thresholds corresponds to the probability that the obligor assumes the credit rating corresponding to this region.

[3] This likely will not be the case in practice, but for our purposes here, the value of $\mu$ will not influence the result. It is in fact true that $\sigma$ does not influence the final result either – and the reader may choose to ignore $\sigma$ in the expressions to follow – but we retain it for illustrative purposes.

*Chart 8.2*
**Distribution of asset returns with rating change thresholds**

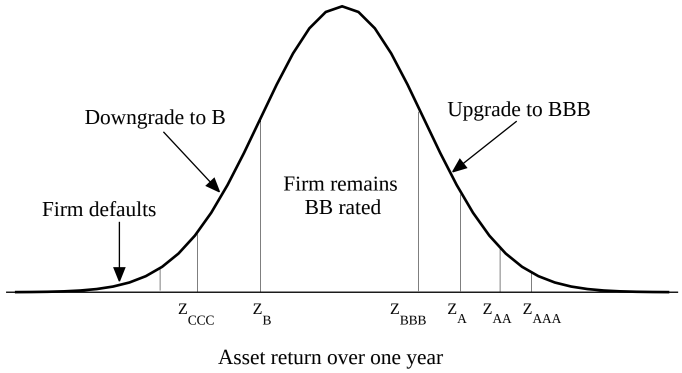


Now in order to complete the connection, we simply observe that the probabilities in the two columns of the Table 1 must be equal. So considering the default probability, we see that $\Phi(Z_{Def}/\sigma)$ must equal 1.06%, which lets us solve for $Z_{Def}$:

$$Z_{Def} = \Phi^{-1}(1.06\%) \cdot \sigma = -2.30\sigma, \quad [8.3]$$

where $\Phi^{-1}(p)$ gives the level below which a standard normal distributed random variable falls with probability $p$. Using this value, we may consider the *CCC* probability to solve for $Z_{CCC}$, then the *B* probability to solve for $Z_B$, and so on, obtaining the values in *Table 8.5*. Note there is no threshold $Z_{AAA}$, since any return over 3.43σ implies an upgrade to AAA.[4]

*Table 8.5*
**Threshold values for asset return for a BBB rated obligor**

| Threshold | Value |
|---|---|
| $Z_{AA}$ | 3.43σ |
| $Z_A$ | 2.93σ |
| $Z_{BBB}$ | 2.39σ |
| $Z_{BB}$ | 1.37σ |
| $Z_B$ | –1.23σ |
| $Z_{CCC}$ | –2.04σ |
| $Z_{Def}$ | –2.30σ |

Now consider a second obligor, A rated. Denote this obligor's asset return by R′, the standard deviation of asset returns for this obligor by σ′, and its asset return thresholds

[4] We comment that to this point, we have not added anything to our model. For one obligor, we only need the transition probabilities to describe the evolution of credit rating changes, and the asset value process is not necessary. The benefit of the asset value process is only in the consideration of multiple obligors.

by $Z'_{Def}$, $Z'_{CCC}$, and so on. The transition probabilities and asset return thresholds are listed in *Table 8.6*.

*Table 8.6*
**Transition probabilities and asset return thresholds for A rating**

| Rating | Probability | Threshold | Value |
|---|---|---|---|
| AAA | 0.09% | | |
| AA | 2.27% | $Z'_{AA}$ | $3.12\sigma'$ |
| A | 91.05% | $Z'_{A}$ | $1.98\sigma'$ |
| BBB | 5.52% | $Z'_{BBB}$ | $-1.51\sigma'$ |
| BB | 0.74% | $Z'_{BB}$ | $-2.30\sigma'$ |
| B | 0.26% | $Z'_{B}$ | $-2.72\sigma'$ |
| CCC | 0.01% | $Z'_{CCC}$ | $-3.19\sigma'$ |
| Default | 0.06% | $Z'_{Def}$ | $-3.24\sigma'$ |

At this point, we have described the motion of each obligor individually according to its asset value processes. To describe the evolution of the two credit ratings jointly, we assume that the two asset returns are correlated and normally distributed,[5] and it only remains to specify the correlation ρ between the two asset returns. We then have the covariance matrix for the bivariate normal distribution:

$$[8.4] \qquad \Sigma = \begin{pmatrix} \sigma^2 \rho\sigma\sigma' \\ \rho\sigma\sigma'\sigma'^2 \end{pmatrix}$$

This done, we know how the asset values of the two obligors move together, and can then use the thresholds to see how the two credit ratings move together.

To be specific, say we wish to compute the probability that both obligors remain in their current credit rating. This is the probability that the asset return for the BB rated obligor falls between $Z_B$ and $Z_{BB}$ while at the same time the asset return for the A rated obligor falls between $Z'_{BBB}$ and $Z'_A$. If the two asset returns are independent (i.e., $\rho=0$), then this joint probability is just the product of 80.53% (the probability that the BB rated obligor remains BB rated) and 91.05% (the probability that the A rated obligor remains A rated). If ρ is not zero, then we compute:

$$[8.5] \qquad Pr\{Z_B < R < Z_{BB}, Z'_{BBB} < R' < Z'_A\} = \int_{Z_B}^{Z_{BB}} \int_{Z'_{BBB}}^{Z'_A} f(r, r'; \Sigma)(dr')dr$$

where $f(r,r';\Sigma)$ is the density function for the bivariate normal distribution with covariance matrix $\Sigma$[6]. We may use the same procedure to calculate the probabilities of each of

[5] Technically, we assume that the two asset returns are bivariate normally distributed. We remark, however, that it is not necessary to use the normal distribution. Any multivariate distribution (including those incorporating fat tails or skewness effects) where the joint movements of asset values can be characterized fully by one correlation parameter would be applicable.

[6] The variables $r$ and $r'$ in Eq. [8.5] represent the values that the two asset returns may take on within the specified intervals.

the 64 possible joint rating moves for the two obligors. As an example, suppose that ρ=20%. We would then obtain the probabilities in *Table 8.7*.

*Table 8.7*
**Joint rating change probabilities for BB and A rated obligors (%)**

| Rating of first company | Rating of second company | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | AAA | AA | A | BBB | BB | B | CCC | Def | Total |
| AAA | 0.00 | 0.00 | 0.03 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.03 |
| AA | 0.00 | 0.01 | 0.13 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.14 |
| A | 0.00 | 0.04 | 0.61 | 0.01 | 0.00 | 0.00 | 0.00 | 0.00 | 0.67 |
| BBB | 0.02 | 0.35 | 7.10 | 0.20 | 0.02 | 0.01 | 0.00 | 0.00 | 7.69 |
| BB | 0.07 | 1.79 | 73.65 | 4.24 | 0.56 | 0.18 | 0.01 | 0.04 | 80.53 |
| B | 0.00 | 0.08 | 7.80 | 0.79 | 0.13 | 0.05 | 0.00 | 0.01 | 8.87 |
| CCC | 0.00 | 0.01 | 0.85 | 0.11 | 0.02 | 0.01 | 0.00 | 0.00 | 1.00 |
| Def | 0.00 | 0.01 | 0.90 | 0.13 | 0.02 | 0.01 | 0.00 | 0.00 | 1.07 |
| Total | 0.09 | 2.29 | 91.06 | 5.48 | 0.75 | 0.26 | 0.01 | 0.06 | 100 |

This table is sufficient to compute the standard deviation of value change for a portfolio containing only issues of these two obligors. Note that the totals for each obligor are just that obligor's transition probabilities. To compute the standard deviation for a larger portfolio, it is only necessary to repeat this analysis for each pair of obligors in the portfolio.[7]

The effect of the correlation merits further comment. Consider the worst case event for a portfolio containing these two obligors – that both obligors default. If the asset returns are independent, then the joint default probability is the product of the individual default probabilities, or 0.0006%. On the other hand, if the asset returns are perfectly correlated (ρ=1), then any time the A rated obligor defaults, so too does the BB rated obligor. Thus, the probability that they both default is just the probability that the A rated obligor defaults, or 0.06%, 100 times greater than in the uncorrelated case.

In *Chart 8.3*, we illustrate the effect of asset return correlation on the joint default probability for our two obligors.

[7] Note that if all pairs of obligors have the same correlation, then the maximum number of matrices like *Table 8.7* which would be needed is 28, regardless of the size of the portfolio. Notice that *Table 8.7* depends only on the ratings of the two obligors and on the correlation between them, and not on the particular obligors themselves. Thus, since there are only seven possible ratings for each obligor, there are only 28 possibilities for the ratings of each pair of obligors, and 28 possible matrices.

*Chart 8.3*
**Probability of joint defaults as a function of asset return correlation**

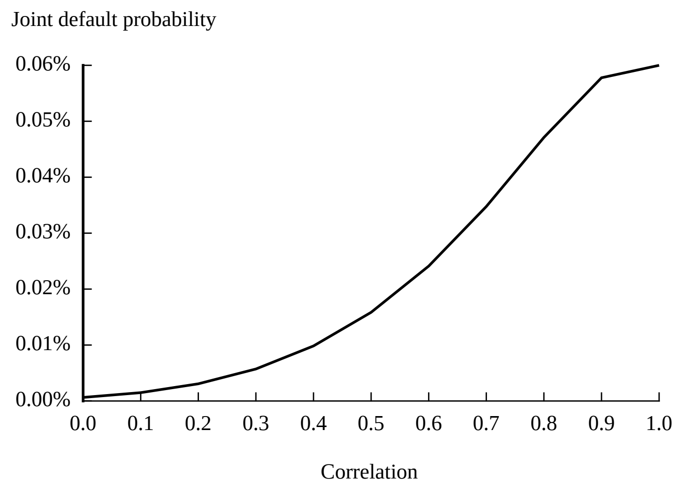


We have pointed out before that for pairs of obligors, it is only necessary to specify joint probabilities of rating changes and defaults, and that actual default correlations are not used in any calculations. However, many people are accustomed to thinking in terms of default correlations, and so we touch briefly on them here. For an asset correlation $\rho_A$, we have shown that it is possible to compute $p_{12}$, the probability that obligors 1 and 2 both default. The default correlation between these two obligors can then be written as

$$[8.6] \qquad \rho_D = \frac{p_{12} - p_1 p_2}{\sqrt{p_1(1-p_1)p_2(1-p_2)}},$$

where $p_1$ and $p_2$ are the probabilities that obligor 1 and obligor 2 default, respectively.

The translation from asset to default correlation lowers the correlation significantly. Asset correlations in the range of 40% to 60% will typically translate into default correlations of 2% to 4%. We see then that even the very small default correlation estimates in *Section 8.1* require that asset value moves exhibit relatively high correlations.

*Chart 8.4* shows how the default correlation is a function of the two obligor's default probabilities. An asset correlation of 30% was assumed and default probabilities range from 1bp to above 10%. The high "mound" towards the back indicates that junk bond defaults will be far more correlated with each other than will investment grade defaults.

*Chart 8.4*
**Translation of equity correlation to default correlation**

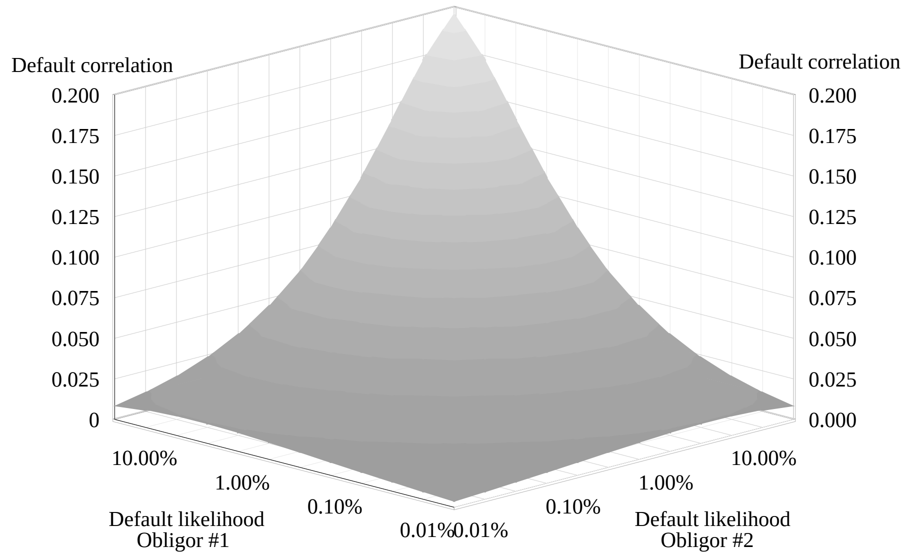


Before moving on to estimation of parameters, we make one important observation: Equation [8.5] above does not depend on either of the volatilities $\sigma$ or $\sigma'$. This may seem counterintuitive, that in a risk model we are ignoring asset volatility, but essentially all of the volatility we need to model is captured by the transition probabilities for each obligor. As an example, consider two obligors which have the same rating (and therefore the same transition probabilities), but where the asset volatility for one obligor is ten times greater than the other. We know that the credit risk is the same to either obligor. One obligor does have a more volatile asset process, but this just means that its asset return thresholds are greater than those of the other firm. In the end, the only parameters which affect the risk of the portfolio are the transition probabilities for each obligor and the correlations between asset returns.

The consequence of this is that we may consider *standardized* asset returns, that is, asset returns adjusted to have mean zero and standard deviation one. The only parameter to estimate then is the correlation between asset returns, which is the focus of the next section.

One last comment is that it is a simple matter to adjust for different time horizons. For example, to perform this analysis for a six-month time horizon, the only change is that we use the six-month transition probabilities to calibrate the asset return thresholds.

## 8.5 Estimating asset correlations

The user can pursue different alternatives to estimate firm asset correlations. The simplest is just to use some fixed value across all obligor pairs in the portfolio. This precludes the user having to estimate a large number (4,950 for a 100-obligor portfolio) of individual correlations, while still providing reasonable portfolio risk measures. However, the ability to detail risk due to overconcentration in a particular industry, for example, is lost. A typical average asset correlation across a portfolio may be in the range of 20% to 35%.[8]

For more specific correlations, there are independent data providers that can provide models which are independent of – but can be consistently used in – CreditMetrics. Below, we present our own interpretation of this type of underlying firm asset correlation estimation.

One fundamental – and typically very observable – source of firm-specific correlation information are equity returns. Here, we use the correlation between equity returns as a proxy for the correlation of asset returns. While this method has the drawback of overlooking the differences between equity and asset correlations, it is more accurate than using a fixed correlation, and is based on much more readily available data than credit spreads or actual joint rating changes.

In the best of all possible worlds, we could produce correlations for any pair of obligors which a user might request. However, the scarcity of data for many obligors, as well as the impossibility of storing a correlation matrix of the size that would be necessary, make this approach untenable. Therefore, we resort to a methodology which relies on correlations within a set of indices and a mapping scheme to build the obligor-by-obligor correlations from the index correlations.

Thus, to produce individual obligor correlations, there are two steps:

- First, we utilize industry indices in particular countries to construct a matrix of correlations between these industries. The result is that we obtain the correlation, for example, of the German chemical industry with the United States insurance industry. For reasons which will become clear below, we also report the volatility for each of these indices.[9]

- Next, we map individual obligors by industry participation. For example, a company might be mapped as 80% Germany and 20% United States, and 70% chemicals and 30% finance, resulting in 56% participation in the German chemicals industry, 24% in German finance, 14% in American chemicals, and 6% in American finance. Using these weights and the country-industry correlations from above, we obtain the correlations between obligors.

In *Section 8.5.1*, we discuss the data we provide and the methodology which goes into its construction. In the following subsection, we present an example to describe the methods by which the user specifies the weightings for individual obligors and arrives at individual obligor correlations. The last subsection is a generalization of this example.

---

[8] Based on conversations with Patrick H. McAllister in 1994 when he was an Economist at the Board of Governors of the Federal Reserve System. Part of his research inferred average asset correlations of corporate & industrial loan portfolios within mid-sized US banks to be in the range 20%-to-25%. Our own research suggests that it is easier to construct higher correlation portfolios versus lower correlation portfolios, hence a 20%-to-35% range.

[9] Recall from *Section 8.4* that volatilities do not figure into the model for joint rating changes. We will see that the volatilities of the indices are necessary, however, for mapping individual obligors to the indices.

### 8.5.1 Data

As mentioned above, we provide the user a matrix of correlations between industries in various countries. In this section, we discuss the data and the methods by which we construct this matrix.

In *Table 8.8*, we list the countries for which we provide data, along with the family of industry specific indices we use for each country. For each country, the broad country index used is the MSCI index. For countries where no index family appears, insufficient industry index data was available and we utilize only the data for the broad country index.

*Table 8.8*
**Countries and respective index families**

| Country | Index family | Country | Index family |
|---|---|---|---|
| Australia | ASX | Mexico | Mexican SE |
| Austria | | New Zealand | |
| Belgium | | Norway | Oslo SE |
| Canada | Toronto SE | Philippines | Philippine SE |
| Finland | Helsinki SE | Poland | |
| France | SBF | Portugal | |
| Germany | CDAX | Singapore | All-Singapore |
| Greece | Athens SE | South Africa | |
| Hong Kong | Hang Seng | Spain | |
| Indonesia | | Sweden | Stockholm SE |
| Italy | Milan SE | Switzerland | SPI |
| Japan | Topix | Thailand | SET |
| Korea | Korea SE | United Kingdom | FT-SE-A |
| Malaysia | KLSE | United States | S&P |

In *Table 8.9*, we list the industries for which we provide indices in one or more of the countries. We choose these industry groups by beginning with the major groups used by Standard & Poor for the United States, and then eliminating groups which appear redundant. For instance, we find that the correlation between the Health Care and Pharmaceuticals indices is over 98%, and so consolidate these two groups into one, reasoning that the two indices essentially explain the same movements in the market.

*Table 8.9*
**Industry groupings with codes**

| Grouping | Code | Grouping | Code |
|---|---|---|---|
| General country index | GNRL | Insurance | INSU |
| Automobiles | AUTO | Machinery | MACH |
| Banking & finance | BFIN | Manufacturing | MANU |
| Broadcasting & media | BMED | Metals Mining | MMIN |
| Chemicals | CHEM | Oil & gas – refining & marketing | OGAS |
| Construction & building materials | CSTR | Paper & forest products | PAPR |
| Electronics | ELCS | Publishing | PUBL |
| Energy | ENRG | Technology | TECH |
| Entertainment | ENMT | Telecommunications | TCOM |
| Food | FOOD | Textiles | TXTL |
| Health care & pharmaceuticals | HCAR | Transportation | TRAN |
| Hotels | HOTE | Utilities | UTIL |

Because the industry coverage in each country is not uniform, we also provide data on MSCI worldwide industry indices. In a case such French chemicals, where there is no country-industry index, the user may then choose to proxy the French chemical index with a combination of the MSCI France index and the MSCI worldwide chemicals index. Finally, realizing that it may at times be more feasible to describe a company by a regional index rather than a set of country indices, we provide data on six MSCI regional indices. In the end, we select the indices for which at least three years of data are available, leaving us with 152 country-industry indices, 28 country indices, 19 worldwide industry indices, and 6 regional indices. The available country-industry pairs are presented in *Table 8.10*. For the specific index titles used in each case, refer to *Appendix I*.

*Table 8.10*
**Country-industry index availability**

| Country | GNRL | AUTO | BFIN | BMED | CHEM | CSTR | ELCS | ENRG | ENMT | FOOD | HCAR | HOTE | INSU | MACH | MANU | MMIN | OGAS | PAPR | PUBL | TECH | TCOM | TXTL | TRAN | UTIL | Total |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Australia | X | | X | X | X | X | | X | | X | | | X | | | | | X | | | | | X | | 10 |
| Austria | X | | | | | | | | | | | | | | | | | | | | | | | | 1 |
| Belgium | X | | | | | | | | | | | | | | | | | | | | | | | | 1 |
| Canada | X | X | X | X | X | X | X | X | | X | X | X | X | | | X | | | X | | | | X | | 15 |
| Finland | X | | X | | | | | | | | | | X | | | X | | X | | | | | | | 5 |
| France | X | X | X | | | X | | X | | X | | | | | | | | | | | | | | | 6 |
| Germany | X | X | X | | X | X | | | | | | | X | X | | | | X | | | | X | X | X | 11 |
| Greece | X | | X | | | | | | | | | | X | | | | | | | | | | | | 3 |
| Hong Kong | X | | X | | | | | | | | | | | | | | | | | | | | | X | 3 |
| Indonesia | X | | | | | | | | | | | | | | | | | | | | | | | | 1 |
| Italy | X | | X | | X | | | | | X | | | | | | X | | X | | | | | | | 6 |
| Japan | X | | X | X | X | X | X | X | | X | X | | X | X | | X | X | X | | | | X | X | | 16 |
| Korea | X | | X | | X | X | | | | X | | | X | X | | X | | X | | | | X | X | | 11 |
| Malaysia | X | | X | | | | | | | | | | | | | X | | | | | | | | | 3 |
| Mexico | X | | | | | X | | | | | | | | | | X | | | | | | | X | | 4 |
| New Zealand | X | | | | | | | | | | | | | | | | | | | | | | | | 1 |
| Norway | X | | X | | | | | | | | | | X | | | | | | | | | | | | 3 |
| Philippines | X | | | | | | | | | | | | | | | X | X | | | | | | | | 3 |
| Poland | X | | | | | | | | | | | | | | | | | | | | | | | | 1 |
| Portugal | X | | | | | | | | | | | | | | | | | | | | | | | | 1 |
| Singapore | X | | X | | | | | | | | | X | | | | | | | | | | | | | 3 |
| South Africa | X | | X | | | | | | | | | | | | | X | | | | | | | | | 3 |
| Spain | X | | | | | | | | | | | | | | | | | | | | | | | | 1 |
| Sweden | X | | X | | X | X | | | | | | | | | | | | X | | | | | | | 6 |
| Switzerland | X | | X | | X | X | X | | | | | | | | | | | | | | | | | | 5 |
| Thailand | X | | X | | X | X | X | X | | X | X | X | X | X | | X | | X | X | X | | X | X | | 17 |
| United Kingdom | X | | X | X | X | X | X | X | | X | X | X | X | | | X | X | X | | | X | X | X | | 17 |
| United States | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | 24 |
| MSCI worldwide | | X | X | X | X | X | X | X | X | X | X | X | X | X | | X | | X | | | X | X | X | X | 19 |
| Total | 28 | 5 | 20 | 6 | 12 | 13 | 7 | 8 | 2 | 10 | 6 | 6 | 12 | 6 | 1 | 13 | 4 | 11 | 3 | 2 | 3 | 7 | 10 | 4 | 199 |

For each of the indices, we consider the last 190 weekly returns, and compute the mean and standard deviation of each return series. Thus, if we denote the $t^{th}$ week's return on the $k^{th}$ index by $R_t^{(k)}$, we compute the average weekly return on this index by

[8.7] $$\bar{R}^{(k)} = \frac{1}{T}\sum_{t=1}^{T} R_t^{(k)},$$

where $T$ is 190 in our case, and the weekly standard deviation of return by

[8.8] $$\sigma_k = \sqrt{\frac{1}{T-1}\sum_{t=1}^{T}\left(R_t^{(k)} - \bar{R}^{(k)}\right)^2}.$$

As mentioned above, we provide the user with the standard deviations (volatilities), and discuss their use in the next section. In addition, for all pairs of indices, we compute the covariance of weekly returns by

[8.9] $$\mathrm{COV}(k, l) = \frac{1}{T-1}\sum_{t=1}^{T}\left(R_t^{(k)} - \bar{R}^{(k)}\right)\left(R_t^{(l)} - \bar{R}^{(l)}\right),$$

and the correlation of weekly returns by

[8.10] $$\rho_{k,l} = \frac{\mathrm{COV}(k, l)}{\sigma_k \sigma_l}.$$

We provide these correlations to the user.

Note that our computations of volatilities and correlations differ from the standard volatility computations in RiskMetrics in that we weight all of the returns in each time series equally. The motivation for this is that we are interested in computing correlations which are valid over the longer horizons for which CreditMetrics will be used. The statistics here tend to be more stable over time, and reflect longer term trends, whereas the statistics in RiskMetrics vary more from day to day, and capture shorter term behavior.

Note also that the correlations we compute are based on historical weekly returns. It is therefore an assumption of the model that the weekly correlations which we provide are accurate reflections of the quarterly or yearly asset moves which drive the CreditMetrics model.

### *8.5.2 Obligor correlations – example*

Now that we have described how to calculate correlations between country-industry pairs, it only remains to illustrate how to apply these to obtain correlations between individual obligors. The steps of this computation are as follows:

1. Assign weights to each obligor according to its participation in countries and industries, and specify how much of the obligor's equity movements are not explained by the relevant indices.
2. Express the standardized returns for each obligor as a weighted sum of the returns on the indices and a company-specific component.
3. Use the weights along with the index correlations to compute the correlations between obligors.

By specifying the amount of an obligor's equity price movements are not explained by the relevant indices, we are describing this obligor's firm-specific, or idiosyncratic, risk. Generally, prices for companies with large market capitalization will track the indices closely, and the idiosyncratic portion of the risk to these companies is small; on the other hand, prices for companies with less market capitalization will move more independently of the indices, and the idiosyncratic risk will be greater.

We will explain each of the steps above through an example.

Suppose we wish to compute the correlation between two obligors, ABC and XYZ. Assume that we decide that ABC participates only in the United States chemicals industry, and that its equity returns are explained 90% by returns on the United States chemicals index and 10% by company-specific movements. We assume that these company-specific movements are independent of the movements of the indices, and also independent of the company-specific movements for all other companies. Assume that XYZ participates 75% in German insurance and 25% in German banking and finance and that 20% of the movements in XYZ's equity are company-specific.

To apply these weights and describe the standardized returns for the individual obligors, we need the volatilities and correlations of the relevant indices. We present these in the *Table 8.11*. The volatilities listed are for weekly returns.

*Table 8.11*
**Volatilities and correlations for country-industry pairs**

| Index | Volatility | Correlations | | |
|---|---|---|---|---|
| | | U.S. Chemicals | Germany Insurance | Germany Banking |
| U.S.: Chemicals | 2.03% | 1.00 | 0.16 | 0.08 |
| Germany: Insurance | 2.09% | 0.16 | 1.00 | 0.34 |
| Germany: Banking | 1.25% | 0.08 | 0.34 | 1.00 |

For the firm ABC, the volatility explained by the U.S. chemicals index is 90% of the firm's total volatility. The remainder is explained by ABC's firm specific movements. Thus, we consider two independent standard normal random variables, $r^{(USCm)}$ and $\hat{r}^{(ABC)}$, which represent the standardized returns of the U.S. chemical index and ABC's firm specific standardized returns, respectively. We then write ABC's standardized returns as

$$r^{(ABC)} = w_1 r^{(USCm)} + w_2 \hat{r}^{(ABC)} . \tag{8.11}$$

We know that 90% of ABC's volatility is explained by the index, and thus we know that $w_1 = 0.9$. We also know that the total volatility must be one (since the returns are standardized), and thus $w_2 = \sqrt{1 - w_1^2} = 0.44$.

For XYZ, we proceed in a similar vein. We first figure the volatility of the index movements for XYZ, that is, the volatility of an index formed by 75% German insurance and 25% German banking, by

[8.12]

$$\hat{\sigma} = \sqrt{0.75^2 \cdot \sigma_{DeIn}^2 + 0.25^2 \cdot \sigma_{DeBa}^2 + 2 \cdot 0.75 \cdot 0.25 \cdot \rho(DeIn, DeBa) \cdot \sigma_{DeIn} \cdot \sigma_{DeBa}} = 0.017.$$

We then scale the weights so that the total volatility of the index portion of XYZ's standardized returns is 80%. Thus, the weight on the German insurance index is

[8.13] $$0.8 \cdot \frac{0.75 \cdot \sigma_{DeIn}}{\hat{\sigma}} = 0.74,$$

and the weight on the German banking index is

[8.14] $$0.8 \cdot \frac{0.25 \cdot \sigma_{DeBa}}{\hat{\sigma}} = 0.15.$$

Finally, in order that the total standardized return of XYZ have variance one, we know that the weight on the idiosyncratic return must be $\sqrt{1 - 0.8^2} = 0.6$.

At this point, we have what we will refer to as each firm's *standard weights*, that is, the weightings on standardized index returns which allow us to describe standardized firm returns. Recall that for our example we describe the returns for ABC and XYZ by:

[8.15] $$r^{(ABC)} = 0.90 r^{(USCm)} + 0.44 \hat{r}^{(ABC)},$$

and

[8.16] $$r^{(XYZ)} = 0.74 r^{(DeIn)} + 0.15 r^{(DeBa)} + 0.6 \hat{r}^{(XYZ)},$$

where $\hat{r}^{(ABC)}$ and $\hat{r}^{(XYZ)}$ are the idiosyncratic returns for the two firms. Since the idiosyncratic returns are independent of all the other returns, we may compute the correlation between ABC and XYZ by:

[8.17] $$\rho(ABC, XYZ) = 0.90 \cdot 0.74 \cdot \rho(USCm, DeIn) + 0.90 \cdot 0.15 \cdot \rho(USCm, DeBa) = 0.11$$

The above illustrates the method for computing correlations between pairs of obligors, and suggests a more general framework. In the next subsection, we present the same methods, but generalized to handle obligors with participations in more industries and countries.

Note that the index volatilities do not actually enter into the correlation calculations, but do play a role when we convert industry participations to standard weights. This allows us to account for cases like our example, where industry participation is split 75% and 25%, or 3 to 1, but since the industry with 75% participation (insurance) is more volatile than the other industry (banking), the standard weight on insurance is actually more than three times greater than the standard weight on banking.

### *8.5.3 Obligor correlations – generalization*

To complete our treatment of obligor correlations, we provide generalizations of the methods above for computing standard weights and for calculating correlations from these weights.

First, to compute standard weights, consider a firm with industry participations of $\hat{w}_1$, $w_2$, and $w_3$, where the indices account for $\alpha$ of the movements of the firm's equity. We compute the firm's standard weights in the following steps:

Compute the volatility of the weighted index for the firm, that is,

$$[8.18] \quad \hat{\sigma} = \sqrt{w_1^2\sigma_1^2 + w_2^2\sigma_2^2 + w_3^2\sigma_3^2 + 2w_1w_2\rho_{1,2}\sigma_1\sigma_2 + 2w_2w_3\rho_{2,3}\sigma_2\sigma_3 + 2w_1w_3\rho_{1,3}\sigma_1\sigma_3}$$

Scale the weights on each index such that the indices represent only $\alpha$ of the volatility of the firm's standardized returns. The scaling is as below:

$$[8.19] \quad w_1 = \alpha \cdot \frac{w_1\sigma_1}{\hat{\sigma}},\ w_2 = \alpha \cdot \frac{w_2\sigma_2}{\hat{\sigma}},\ \text{and}\ w_3 = \alpha \cdot \frac{w_3\sigma_3}{\hat{\sigma}}.$$

Compute the weight on the idiosyncratic returns by taking $\sqrt{1-\alpha^2}$ .

The generalization to the case of four or more indices should be clear.

Now suppose we have *n* different firms with standard weightings on *m* indices, and we wish to compute the equity correlations between these firms. Let the correlation matrix for the indices be denoted by *C*. Since the weightings are on both the indices and the idiosyncratic components, we need to create a correlation matrix, $\bar{C}$, which covers both of these. This matrix will be $m+n$ by $m+n$, and constructed as below:

$$
\begin{array}{rc}
 & \begin{array}{cc} \overbrace{\hspace{5em}}^{m \text{ columns}} & \overbrace{\hspace{5em}}^{n \text{ columns}} \end{array} \\
\begin{array}{r} m \text{ rows} \left\{ \vphantom{\begin{matrix}0\\ \vdots\\ \vdots\\ \vdots\\ 0\end{matrix}} \right. \\ n \text{ rows} \left\{ \vphantom{\begin{matrix}0\\ \vdots\\ \vdots\\ 0\end{matrix}} \right. \end{array} &
\left(\begin{array}{c|c}
C & \begin{matrix} 0 & \ldots & \ldots & 0 \\ \vdots & \ldots & \ldots & \vdots \\ \vdots & \ldots & \ldots & \vdots \\ \vdots & \ldots & \ldots & \vdots \\ 0 & \ldots & \ldots & 0 \end{matrix} \\
\hline
\begin{matrix} 0 & \ldots\ldots & 0 \\ \vdots & \ldots\ldots & \vdots \\ \vdots & \ldots\ldots & \vdots \\ 0 & \ldots\ldots & 0 \end{matrix} & \begin{matrix} 1 & 0 & \ldots & 1 \\ 0 & \ddots & \ldots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & \ldots & \ldots & 1 \end{matrix}
\end{array}\right)
\end{array}
$$

Thus, the upper left of $\bar{C}$ is the *m* by *m* matrix *C*, representing the correlations between indices; the lower right is the *n* by *n* identity matrix, reflecting that each firm's idiosyncratic component has correlation one with itself and is independent of the other firms' idiosyncratic components; and the remainder consists of only zeros, reflecting that there is no correlation between the idiosyncratic components and the indices. For the example in the previous subsection (where $m = 3$ and $n = 2$), we would have

$$[8.20] \qquad \bar{C} = \begin{bmatrix} 1 & 0.16 & 0.08 & 0 & 0 \\ 0.16 & 1 & 0.34 & 0 & 0 \\ 0.08 & 0.34 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

We then create a *m*+*n* by *n* weight matrix *W*, where each column represents a different firm, and each row represents weights on indices and idiosyncratic components. Thus, in the $k^{\text{th}}$ column of *W*, the first *m* entries will give the first firm's weights on the indices, the *m*+*n*+*k* entry will give the firm's idiosyncratic weight, and the remaining entries will be zero. For our example, the matrix $W$ would be given by

$$[8.21] \qquad W = \begin{bmatrix} 0.90 & 0 \\ 0 & 0.74 \\ 0 & 0.15 \\ 0.44 & 0 \\ 0 & 0.60 \end{bmatrix}.$$

The *n* by *n* matrix giving the correlations between all of the firms is then given by $W' \cdot \bar{C} \cdot W$.

# *Part III*

# *Applications*