



Introduction to CreditMetrics™

The benchmark for understanding credit risk

New York
April 2, 1997

- 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.*

J.P. Morgan

Co-sponsors:

Bank of America

Bank of Montreal

BZW

Deutsche Morgan Grenfell

KMV Corporation

Swiss Bank Corporation

Union Bank of Switzerland

Table of Contents

1. Introduction to CreditMetrics	3
2. The case for a portfolio approach to credit risk	9
3. The challenges of estimating portfolio credit risk	12
4. An overview of CreditMetrics methodology	14
5. Practical applications	30

Introduction to CreditMetrics

Copyright © 1997 J.P. Morgan & Co. Incorporated. All rights reserved.
J.P. Morgan Securities, Inc., member SIPC. J.P. Morgan is the marketing name for J.P. Morgan & Co. Incorporated and its subsidiaries worldwide.

CreditMetrics™, CreditManager™, FourFifteen™, and RiskMetrics™ are trademarks of J.P. Morgan in the United States and in other countries. They are written with the symbol ™ on their first occurrence in the publication, and as CreditMetrics, CreditManager, FourFifteen or RiskMetrics thereafter.

Windows NT and Windows 95 are trademarks of the Microsoft Corporation

J.P. Morgan Securities Inc.
Global Credit Derivatives
Blythe Masters
(1-212) 648-1432
masters_blythe@jpmorgan.com

1. Introduction to CreditMetrics

1.1 What is CreditMetrics?

CreditMetrics is the first readily available portfolio model for evaluating credit risk. The CreditMetrics approach enables a company to consolidate credit risk across its entire organization, and provides a statement of value-at-risk (VaR) due to credit caused by upgrades, downgrades, and defaults.

CreditMetrics will be useful to all companies worldwide that carry credit risk in the course of their business. It provides a methodology to quantify credit risk across a broad range of instruments, including traditional loans, commitments and letters of credit; fixed income instruments; commercial contracts such as trade credits and receivables; and market-driven instruments such as swaps, forwards and other derivatives.

CreditMetrics has three components:

- **a methodology** for assessing portfolio value-at-risk (VaR) due to changes in obligor credit quality. This is fully explained in the CreditMetrics Technical Document that is freely available from the Internet (<http://www.jpmorgan.com>) or from J.P. Morgan and the co-sponsors in printed form
- **a data set** that is also freely available from the internet
- **a software package (CreditManager™)** that implements the methodology of CreditMetrics, that can be purchased from J.P. Morgan and the co-sponsors

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 own activities.

1.2 What is the aim of CreditMetrics?

J.P. Morgan developed CreditMetrics to:

- **Create a benchmark for credit risk measurement:** Without a common point of reference for credit risks, it is difficult to compare different sources and measures of risk. CreditMetrics provides the yardstick for measuring these risks, making them comparable.
- **Promote credit risk transparency and better risk management tools, leading to improved market liquidity:** CreditMetrics provides a more precise and systematic understanding of risks. With that knowledge comes the ability and greater willingness to manage risks more actively. Transparent methodology and better risk management tools will promote more liquid markets that are essential to effective risk management.
- **Encourage a regulatory capital framework that more closely reflects economic risk:** Many regulated institutions today are subject to capital requirements that do not reflect economic risk. A portfolio VaR model for credit can be used as a risk-based capital allocation tool, equally appropriate for both internal and regulatory capital allocation.
- **Complement other elements of credit risk management decisions:** CreditMetrics is intended to complement other, more traditional forms of credit analysis that are integral to an institution's credit underwriting and monitoring standards. In particular, there is no substitute for sound credit analysis, and CreditMetrics methodology does not obviate this need. CreditMetrics

is neither a credit rating tool nor a pricing model, nor does it provide, in its current version, a portfolio optimization methodology. However, CreditMetrics is designed to provide measures of portfolio risk which take account of the relationship between each asset and the existing portfolio, thereby making credit risk management decisions more systematic.

In the interest of establishing a benchmark in a field with as little standardization and precise data as credit risk measurement, J.P. Morgan has 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 J.P. Morgan is pleased to have received their input and support in the development of CreditMetrics. With this sponsorship we all hope to send one clear and consistent message to the marketplace in an area with little clarity to date.

1.3 Why release CreditMetrics now?

Credit risk may be the key risk management challenge of the late 1990s. In recent years, a booming global economy and a healthy credit cycle have created a business environment in which a growing number of institutions are taking on more, and increasingly complex, forms of credit risk. For example:

- As credit spreads continue to narrow, banks in competitive lending markets are retaining more credit risks: U.S. primary loan syndication activity reached \$900 billion in 1996, yet secondary loan trading volumes were only \$40 billion. Elsewhere, such statistics are even more skewed.
- The proliferation of complex financial instruments has created uncertain and market-sensitive counterparty exposures that are significantly more challenging to manage than traditional instruments such as bonds.
- As investors find fewer opportunities in interest rate and currency markets, they're moving towards yield enhancement through extending and trading credit. In Europe, for example, the post-EMU world will likely see participating government bond markets become credit markets.
- An increasingly varied array of institutions is intermediating and extending credit. Global credit markets have experienced a significant inflow of funds from mutual funds, pension plans, hedge funds and other non-bank institutional investors. Similarly, corporates, insurance companies and their reinsurers are taking on increasing credit exposures through commercial contracts, insurance and derivatives activities.
- Recently, as in the 1980s, the high yield and emerging market sectors have grown significantly, and asset securitization globally is spreading rapidly.

As credit exposures have multiplied and become more complex, the need for more sophisticated risk management techniques for credit risk has also increased. CreditMetrics provides the methodology, data, and software to meet this need.

A risk measurement system is of limited use if it is not accompanied by tools to take action to manage that risk. Today, however, innovative credit risk management tools such as credit derivatives are evolving rapidly. Credit derivatives have made possible more active trading of credit risks without interfering with other business objectives, such as relationship management. At the same time, market liquidity for secondary loans, and high yield and emerging market credits is also improving. CreditMetrics provides the framework necessary to evaluate credit derivative and other credit transactions, whether for hedging or investment. By combining better credit risk trading tools with a portfolio approach to evaluating credit risk, CreditMetrics makes more active credit portfolio management possible.

1.4 The importance of the portfolio approach

CreditMetrics takes a *portfolio* approach to credit risk analysis. This has two aspects: First, credit risks to each obligor across the whole portfolio are restated on an equivalent basis and aggregated in order to be treated consistently, regardless of the underlying asset class. Second, correlations of credit quality moves across obligors are taken into account. Consequently, portfolio effects – the benefits of diversification and costs of concentrations – can be properly quantified.

Concentration risk has been the specific cause of many occurrences of financial distress (e.g., agricultural loans in the U.S. mid-west; oil loans in Texas; the Latin American loan crisis, and so on). It is only in the context of a portfolio model that concentration risk can be evaluated on anything other than an intuitive level.

A portfolio approach allows risk managers to:

- Quantify and control concentration risk that arises from increased exposure to one obligor or groups of correlated obligors, and which can be mitigated only through diversification or hedging.
- Consider concentrations along almost any dimension such as industry, rating category, country, or type of instrument.
- Interpret portfolio credit risk in terms comparable to market value-at-risk calculated using models such as RiskMetrics™ – a benchmark for estimation of market risk.
- Evaluate investment decisions, credit extension, and risk mitigating actions more precisely based on systematic quantitative analysis.
- Set consistent risk-based credit limits, rather than intuitive, but arbitrary, limits based on exposure amounts.
- Make rational risk-based capital allocations.

1.5 The importance of the mark-to-market framework

CreditMetrics considers credit risk in a mark-to-market framework. The risk measures provided by CreditMetrics include not just expected losses but value-at-risk (VaR). That is, the uncertainty or volatility of value – due to changes in obligor credit quality, both across the entire portfolio and for marginal transactions. Credit VaR arises from changes in value due to credit events – that is, changes in obligor credit quality or “migrations”. These credit events include not only defaults but also upgrades and downgrades.

Upgrades and downgrades cause market pricing reactions that result in immediate gains or losses in a mark-to-market accounting regime. Book value accounting ignores these. Regardless of the accounting framework, however, a risk management tool is a better indicator of actual risk if it recognizes changes in value as they occur rather than on a deferred basis. If we fail to recognize the impact of portfolio value due to credit events other than outright defaults, we miss a significant component of risk.

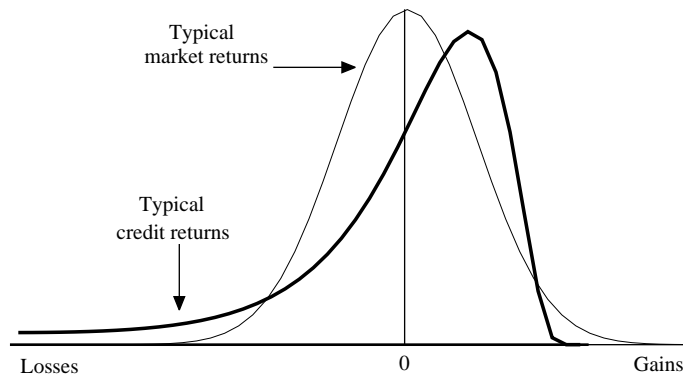
1.6 Is VaR due to credit comparable with VaR in market portfolios?

Market risk is significantly different in nature from credit risk. Typically, market value distributions are relatively symmetrical and well approximated by bell-shaped (mathematically speaking, normal) distributions. In credit portfolios, however, value changes will be relatively small upon minor

up(down)grades, but can be substantial given default. As illustrated in Chart 1, this remote probability of large losses produces skewed return distributions with heavy downside tails that differ significantly from the more normally distributed returns typically addressed by market value-at-risk models.

Chart 1

Comparison of distribution of market returns and credit returns



Nonetheless, this difference in risk profile does not preclude assessing risk on a comparable basis; indeed, it is only by doing so that one can move toward the goal of a fully integrated credit and market risk management system. RiskMetrics and other market VaR models look to a horizon and estimate value-at-risk across a distribution of estimated market outcomes. Likewise, CreditMetrics looks to a horizon and constructs a distribution of value given different estimated credit outcomes. Therefore, although the relevant time horizon is naturally longer for credit risk, CreditMetrics computes credit risk on a comparable basis with market risk.

1.7 The challenge of modeling credit portfolio risk

As we describe more fully in section 3, modeling portfolio risk in credit portfolios is neither analytically nor practically easy, presenting at least two significant challenges. The first problem relates to the long, fat tails observed in credit portfolio distributions, illustrated in Chart 1. Because of this feature, to understand the risks of credit portfolios completely requires that the nature of these tails be explored. To do this demands much more information than is yielded by simple summary statistics such as the mean (expected value) and standard deviation (volatility of value). In fact, to examine the nature of the tails in credit risk portfolios requires moving beyond simple, closed-form analytical equations and deriving the entire shape of the portfolio distribution through simulation, a computationally onerous exercise.

The second problem is empirical. Unlike market portfolios where the data necessary to compute correlations are readily available, correlations in credit portfolios cannot easily be directly observed. Consequently, credit quality correlations must either be derived indirectly from other sources, such as equity prices, or tabulated at a relatively high level of aggregation (e.g., treating all A-rated obligors identically).

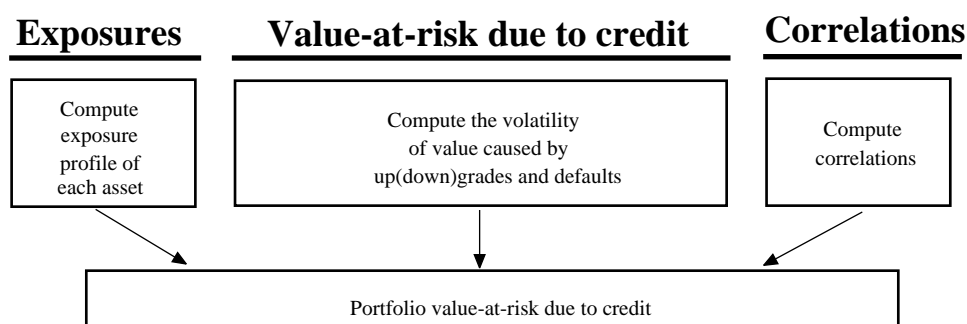
Recognizing these, and other, difficulties associated with computing credit portfolio risk, CreditMetrics seeks to provide the user with a flexible framework within which to express any approach to data generation. The data set provided represents certain possible solutions to each problem, but the CreditMetrics framework neither necessitates, nor advocates, one particular approach over another. Its contribution is thus to encourage a systematic but nonetheless individually tailored approach to risk estimation.

1.8 CreditMetrics methodology

Illustrated in the analytical “road map” presented in chart 2, the CreditMetrics methodology assesses individual and portfolio value-at-risk due to credit in three steps: First, it establishes the exposure profile of each obligor in a portfolio. Second, it computes the volatility in value of each instrument caused by possible upgrades, downgrades, and defaults. Third, taking into account correlations between each of these events, it combines the volatility of the individual instruments to give an aggregate portfolio volatility.

Chart 2

Simplified “road map” of the analytics with CreditMetrics



We consider each step in turn:

1.8.1 Exposure profiles

CreditMetrics easily incorporates the exposures of conventional instruments such as floating rate bonds or drawn loans. It also provides a framework in which to consider less straightforward exposure profiles such as undrawn or non-interest-bearing instruments, including loan commitments, letters of credit and commercial credit arrangements, such as trade credits or receivables. Exposures deriving from undrawn instruments, such as commitments, are captured on the basis of assumptions as to likely changes in drawn amounts upon default or up(down)grade. CreditMetrics also incorporates exposures of market driven instruments, such as swaps and fixed rate bonds, stating all on an equivalent basis to other credit instruments.

1.8.2 Volatility of each exposure from up(down)grades and defaults

Likelihoods are attributed to each possible credit event, including upgrades and downgrades, not just defaults. The probability that an obligor will migrate over a given time horizon to any other rating is derived from a “transition matrix”. Each migration results in an estimated change in value (derived from credit spread data and, in default, recovery rates). Each value outcome is weighted by its likelihood, to create a distribution of value across each credit state, from which each asset’s expected value and volatility (standard deviation) of value are computed.

1.8.3 Correlations

Finally, individual value distributions for each exposure are combined to yield a portfolio result. To compute the volatility of portfolio value from the volatility of individual asset values requires estimates of correlation in credit quality changes. Since credit quality correlations cannot easily be directly observed from historical data many different approach to estimating correlations, including a simple constant correlation, can be used within CreditMetrics.

1.8.4 Advanced modeling features

CreditMetrics also incorporates certain more advanced modeling features to make credit risk estimates more precise:

- **Risk due to recovery rate uncertainty:** The model takes account of volatility of recovery rates, which are notoriously uncertain. The CreditMetrics data set provides volatility estimates for recovery rates based on historical data.
- **Exposure uncertainty in market-driven instruments:** In market-driven instruments, revaluation of positions at the risk horizon is complicated by the interaction between credit and market risk, which are coupled because of an inherent optionality. This optionality derives from the fact that a credit loss can occur only if the position is in-the-money at the point at which the counterparty undergoes a change in credit quality. To evaluate this, CreditMetrics considers expected exposures at the risk horizon, which are, in turn, derived from market rates and volatilities. The necessary expected exposure calculations for market-driven instruments are not performed by CreditManager, but can be performed in and imported from J.P. Morgan's FourFifteen¹, or any other source.
- **Simulation engine:** In addition to providing quantitative solutions to portfolio expected value and volatility of value (solutions to closed-form mathematical expressions), CreditMetrics includes a simulation engine that enables the user to estimate the entire distribution of a credit portfolio.

1.9 What is the output of CreditMetrics?

CreditMetrics computes several different measures of value-at-risk due to credit:

- **Standard deviation (volatility):** This is a measure of symmetrical dispersion about the mean (average) portfolio value. If credit risks were normally distributed, the mean and standard deviation would be sufficient to specify fully the distribution. As shown in Chart 1, however, credit risk is clearly not symmetrical. Thus, the standard deviation measure cannot capture the fact that, for instance, the maximum upside might be only one standard deviation above the average, while meaningful occurrences of loss can be many standard deviations below the average. Consequently, there is a need for more information about the distribution to understand risk in credit portfolios.
- **Percentile levels:** These reflect the likelihood that the portfolio value will fall below a specified level, e.g., that the likelihood of its falling below the 1st percentile level is 1 percent. To calculate a percentile level the full distribution of portfolio values must be specified. This, requires a potentially lengthy simulation, which is computationally complex.

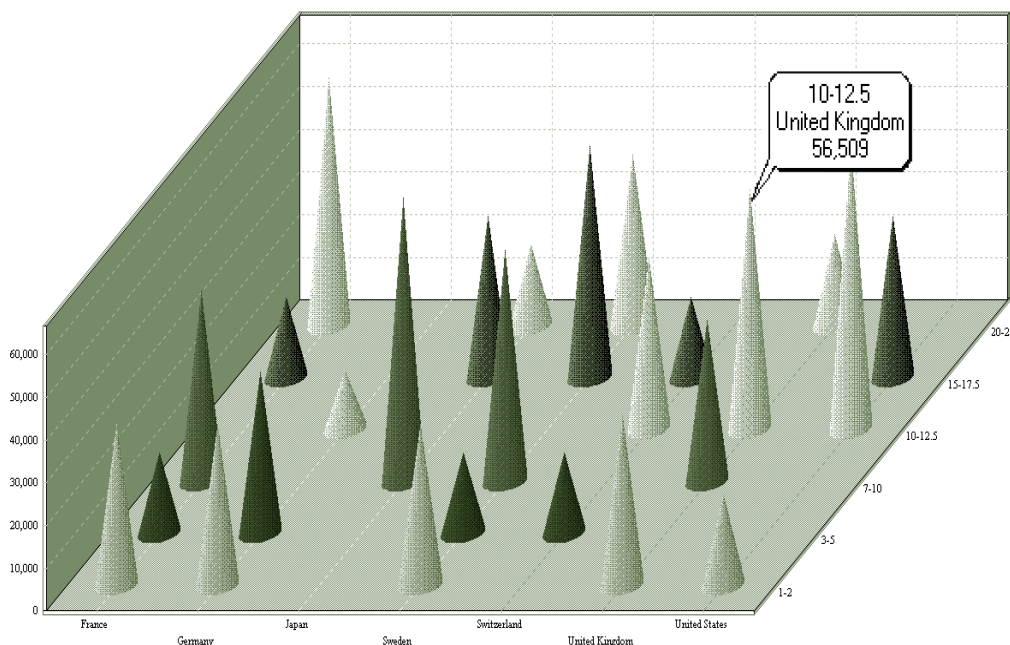
1.10 CreditManager – the software implementation of CreditMetrics

Using these risk measures, CreditMetrics enables managers to generate a value-at-risk report to quantify the amount of credit risk across their firm. These reports can be expressed in the aggregate, or broken down by country, industry, maturity, rating, product type, or any other category of credit exposure. Managers are then able to identify different pockets, or concentrations of risk within a single portfolio, or across an entire firm. Moreover, reports can be generated using mean or standard deviation as a risk measure or any chosen percentile level, as well as valuing risk on a marginal or relative basis. Chart 3 is a value-at-risk report showing credit risk broken down by country and maturity using mean exposure.

¹ FourFifteen™ is the software implementation of J.P. Morgan's market VaR methodology, RiskMetrics.

In this report one can see that there are four pockets of credit risk to the United Kingdom, with a total of \$56,509 having a maturity of 10-12.5 years.

Chart 3
Value-at-risk due to credit



To implement CreditMetrics methodology, a software package, CreditManager, has been developed to provide a flexible analysis and reporting tool in a Microsoft Windows NT or Windows 95 environment. This desktop software provides both sophisticated graphics and statistics to represent the computations and analysis outlined in this and other CreditMetrics documents. Chart 3 is a direct output of CreditManager as are many of the charts and tables of this document. With CreditManager users can assess the different credit risks of their firm on an absolute, relative or marginal basis, across the various portfolios, categories, sectors, or product lines using the risk statistics that they feel most appropriate. CreditManager can be used with the CreditMetrics data set, and/or alternative data provided by the user.

2. The case for a portfolio approach to credit risk

2.1. The credit market environment

As discussed in section 1.3, credit risk is emerging as a key challenge for risk management in the late 1990s. Globally, buoyed by the momentum of a healthy business cycle and a relatively benign credit environment, more institutions are taking on an increasing amount of credit risk.

Meanwhile, many institutions are experiencing constraints on regulatory and/or economic capital, and certainly, following recent high-profile risk management mishaps, there has been intensified focus on risk monitoring and controls. Consequently, as credit exposures have multiplied, the need for more-sophisticated risk management techniques for credit risk has also increased. Common sense dictates that if businesses are demanding better performance in terms of return on economic capital, management must have a solid grasp of all forms of risks being taken to achieve this.

Of course, more active credit risk management could be achieved by more rigorous enforcement of traditional credit processes such as stringent underwriting standards, limit enforcement, and counterparty monitoring. Increasingly, however, risk managers are also seeking to quantify and integrate overall credit risk within a benchmark value-at-risk statement capturing exposure to both market and credit risks.

2.2 The need for a portfolio approach

Why has a quantitative portfolio approach to credit risk management become so important?

2.2.1 Concentration risk

The primary reason is to adequately address and quantify *concentration* risk. Concentration risk refers to additional portfolio risk resulting from increased exposure to one obligor or groups of correlated obligors (perhaps in a particular industry or location). Concentration risk can be mitigated only by diversification or transactions that hedge the specific risk of the concentrated exposure. Such a model creates a framework to consider and stress-test concentrations along almost any dimension (by industry sector, rating category, country or instrument type).

2.2.2 Risk-based limit setting

Traditionally, portfolio managers have relied on a qualitative feeling for the concentration risk in their credit portfolios. Intuitive – but arbitrary – exposure-based credit limits have been the primary defense against unacceptable concentrations of credit risk. Fixed exposure limits, however, do not recognize the relationship between risk and return.

A more quantitative approach, such as the one presented here, would make credit lines a function of marginal portfolio volatility (that is, an *output* of the portfolio management model rather than an *input* to it).

2.2.3 Rational investment decisions and risk-mitigating actions

Another important reason to take a portfolio view of credit risk is to more rationally and accountably address credit extension decisions and risk-mitigating actions.

For example, rightly or wrongly, financial markets are currently indicating a widespread perception of diminished risk due to credit, as illustrated by the historically tight level of credit spreads. In this environment, the bank lending marketplace has become increasingly competitive. As a result, good customer relationships have often become synonymous with heavily concentrated exposures as corporate borrowers command smaller bank groups and larger commitments from relationship banks. Yet, banks are often caught in a paradoxical trap whereby those customers with whom they have developed the most valued relationships are precisely the customers to whom they have the least capacity to take incremental risk. Bank portfolio managers have begun to harbor suspicions that they may be vulnerable to a possible turn for the worse in global credit cycles, and that current levels of spread income may not justify the concentration of risks being accumulated.

Such concerns cannot easily be evaluated nor systematically reflected in pricing and credit extension decisions in the absence of a portfolio model. In a portfolio context, the decision to take on ever higher exposure to an obligor will meet with ever higher risk – risk that grows geometrically with the concentration on that obligor. If relationship demands the extension of credit to a customer to whom the portfolio is overexposed, a portfolio model allows the portfolio manager to quantify (in units of undercompensated risk) exactly the extent of envisaged investment in relationship development. Consequently the risk-return trade-off of concentrated lending activity can be better managed.

Conversely, the portfolio manager can rationally take increased exposure to under-concentrated names. Indeed, such names may be *individually* risky yet offer a relatively small marginal contribution to overall *portfolio* risk due to diversification benefits.

2.2.4 Risk-based economic and regulatory capital allocation

Finally, by capturing portfolio effects (diversification benefits and concentration risks), recognizing that risk accelerates with declining credit quality, and treating credit risk consistently, regardless of asset class, a portfolio credit risk model can provide the foundation for rational risk-based capital allocation.

Such a model is equally appropriate for economic and regulatory capital purposes, but would differ fundamentally from the capital measures currently mandated for bank regulation by the Bank for International Settlements (BIS). For a portfolio of non-trading positions, the BIS risk-based capital accord of 1988 requires capital that is a simple summation of the capital required on each of the portfolio's individual transactions. In turn, each transaction's capital requirement depends on: (i) broad categorization of the obligor, (ii) the transaction's exposure type (e.g., drawn loans versus undrawn commitments), and, (iii) for off-balance-sheet exposures, whether the transaction's maturity is more or less than one year. The weaknesses of this structure – such as its one-size-fits-all risk weighting for all corporate loans regardless of credit rating, and its inability to distinguish between diversified and undiversified portfolios – are increasingly apparent to regulators and market participants alike. Particular concern has been paid to the uneconomic incentives created by the regulatory regime and the inability of regulatory capital adequacy ratios to accurately portray actual bank risk levels. In response to these concerns, bank regulators are increasingly looking for insights in internal credit risk models that generate expected losses and a probability distribution of unexpected losses.²

2.2.5 Responding to market innovation

There are also other, more practical, reasons why a quantitative approach to credit risk is important in response to continuing innovations in financial markets:

- Financial products have become more complex. The growth of derivatives activity has created uncertain and dynamic counterparty exposures that are significantly more challenging to manage than the static exposures of traditional instruments such as bonds or loans. End users and providers of these instruments need to understand such credit risk and its interaction with market risk.
- There has been a proliferation of credit enhancement mechanisms that make it increasingly necessary to assess credit risk at the portfolio as well as the individual asset level. These include: third-party guarantees, credit derivatives, posted collateral, margin arrangements, and netting.
- Improved liquidity in secondary cash markets and the emergence of credit derivatives have made possible more active trading of credit risk based on rational pricing. Prudence requires that institutions thoroughly review existing risks before hedging or trading them.
- Innovative credit instruments explicitly derive value from correlation risk or credit events such as upgrades, downgrades or default. Such risks are best understood in the context of a portfolio model that also explicitly accounts for credit quality migrations.

² See Remarks by Alan Greenspan, Board of Governors of the Federal Reserve System, before the 32nd Annual Conference on Bank Structure and Competition, Federal Reserve Board of Chicago, May 2, 1996.

This section addressed the factors making a portfolio approach to credit risk both necessary and timely. The following section discusses why estimating portfolio credit risk is a much harder problem than estimating portfolio market risk.

3. The challenges of estimating portfolio credit risk

3.1 Unexpected versus expected losses

The *expected* loss calculation is, in one sense, the most straightforward aspect of portfolio theory. That is, the ability to estimate credit quality and the expected size of losses given changes in credit quality, allows the risk manager to price, and reserve for, expected loss (probability of loss \times expected size of loss = expected loss). If there were no further uncertainty relating to possible credit losses, that would be the extent of the risk management problem: Predictable credit losses year after year would be no more than a budgeted expense.

Risk, however, entails not just an estimated possibility of loss but also the *uncertainty* of loss. It turns out that if it is difficult enough to estimate even the *expected* portfolio values, it is harder still to predict *uncertainties* around these values. Even a preliminary excursion into credit analysis reveals that not only are credit-related losses uncertain, but the distribution of outcomes is heavily skewed. It is not uncommon for meaningful probabilities of loss in a credit portfolio to occur many standard deviations distant from the mean. This reveals the inadequacy of an analysis that goes only so far as to characterize expected portfolio values without addressing the uncertainty of those values (VaR).

The essence of prudent credit portfolio management is the establishment of a portfolio balance with adequate diversification. This mitigates the consequences of the portfolio's volatility of value (sometimes termed *unexpected* losses) to a level where an institution can survive such losses given its reserves and capital.

3.2 Characteristics of credit risk distributions

Modern portfolio theory has taken enormous strides in its application to equity and other market price risks. Fundamental differences between credit risks and equity price risks, however, make equity portfolio theory problematic when applied to credit portfolios.

The most immediate problem is that equity returns are relatively symmetrical and are well approximated by normal distributions, while credit returns are highly skewed and fat-tailed (as illustrated in Chart 1). Because of this asymmetry in credit returns, modeling the full distribution of portfolio values requires a great deal of information beyond simple summary statistics such as the mean and standard deviation. Without a full specification of the portfolio value distribution, it is not possible to compute the percentile levels necessary to describe risk in credit portfolios.

By considering every possible combination of credit states across every obligor in the portfolio, the full distribution of a credit portfolio can be constructed mechanistically, but, this is computationally complex for portfolios of more than a few obligors. Consequently, the portfolio distribution can be estimated only by a process of simulation. Simulation reduces the computational burden by sampling outcomes randomly across all possibilities. Once the portfolio distribution has been approximated in this way, it is possible to compute percentile levels and summary statistics that describe the shape of the distribution.

To understand intuitively why credit returns are so different in nature from market returns, consider the risk/return profile of debt investment. This is effectively a skewed bet in which the lender runs a small risk of incurring a large loss (default), balanced by a much larger probability of earning a (relatively) small excess return (net interest earnings), given no default.

Considering an entire portfolio rather than a single obligor has the effect of smoothing the distribution and capturing diversification effects. Nevertheless, the limitation of upside opportunity, combined with the remote possibility of severe losses, still causes the asymmetry and fat, long tails in typical credit portfolio distributions, where these risks are not easily diversified away.

3.3 The importance of diversification

Intuitively, two loans held to maturity will have a default correlation much lower than their corresponding equity price correlation (due to the low likelihood that two extremely remote events will occur simultaneously). For the layperson, a natural conclusion to draw might be that the benefits of diversification in a credit portfolio are not significant precisely because default correlations are so low. But this is not a correct conclusion. The implication of very low default correlations is that the systematic risk in a credit portfolio is small relative to the nonsystematic or individual contribution to risk of each asset. Nonsystematic risk is hedgable or diversifiable risk. The greater the component of nonsystematic risk, the greater the benefits of diversification, and vice versa.

The problem can be viewed another way. Indices provide great hedges of risk in equity portfolios because most equity portfolios are sufficiently diversified to resemble the market. However, because a portfolio of debt of those same names is unlikely to be sufficiently diversified to resemble the market, this same type of index hedge will not work in debt portfolios. The portfolio management consequences of a full characterization of credit risks are thus not insignificant: It takes many more names to fully diversify a credit portfolio than an equity portfolio,³ but when those diversification benefits are achieved, they are considerable. An inadequately diversified portfolio, on the other hand, can result in significantly lower return on risk ratios than would seem intuitively obvious.

3.4 The importance of liquidity and active risk management

Credit exposure has sometimes been modeled as analogous to a portfolio of short, deep out-of-the-money put options on firm assets.⁴ The analogy is intuitively sound given the limited upside and remote, but large, downside profile of credit risk. This “short option” analogy allows us to draw some insights on the consequences of illiquidity in credit portfolios. In equity portfolios, it has been argued, independence of daily returns allows time to diversify risk. If credit portfolios are similar to a portfolio of out-of-the-money puts, however, it can be argued that as the market declines (credit quality deteriorates) the “delta” equivalent⁵ of that portfolio increases and the portfolio becomes more leveraged (riskier). Consequently, any persistent serial correlation in credit returns, as suggested by the historical tendency of one downgrade to be followed by another, can cause poor performance to increase volatility and create accelerating portfolio riskiness. An ability to rebalance the portfolio in response to credit deterioration

³ In *Challenges of Managing Credit Portfolios*, Ron Levin, (J.P. Morgan Securities Inc., 1997) it is argued that compared with an equity portfolio of 30 names, 350 names are required to achieve an equivalent degree of diversification in a debt portfolio.

⁴ This framework was first proposed by Robert Merton (*On the Pricing of Corporate Debt: The Risk Structure of Interest Rates*, *The Journal of Finance*, Vol. 29, 1974), and is often referred to as the option theoretic valuation of debt. It builds on Black-Scholes option pricing theory by stating that debt can be valued as if it were an option on the value of the underlying assets of the firm.

⁵ An option’s “delta” measures the sensitivity of the value of that option to changes in the underlying market – in this case, to changes in credit quality.

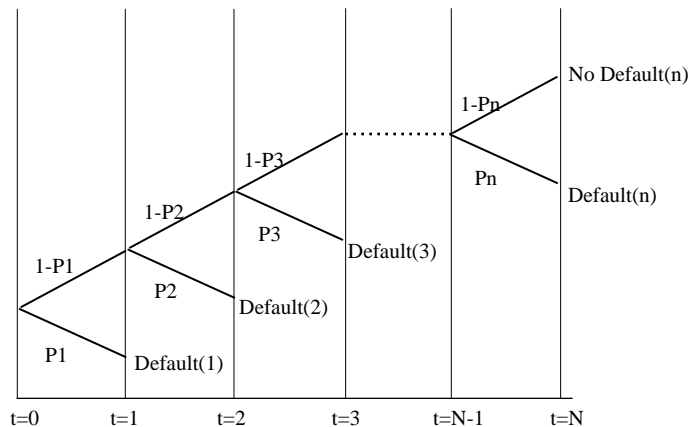
is the only effective way to materially offset this effect. The consequences of illiquidity and/or absence of active risk management in credit portfolios are therefore more severe than in market risk portfolios.

4. An overview of CreditMetrics methodology

4.1 How CreditMetrics relates to other theoretical approaches

There has never been a consensus view on the best way to quantify credit risk, and there are at least two potentially competing frameworks in existence. One popular approach considers only two states of the world: default and no default. The model constructs a binomial tree of default versus no default outcomes until maturity (see Chart 4), seeking to capture each instrument's entire credit risk profile based only on default probability and recovery estimates.

Chart 4
A binomial model of credit risk




An alternative approach, often referred to as the RAROC approach, holds that risk is the observed volatility of corporate bond values within each credit rating category, maturity band, and industry grouping. Implementations vary, but the idea is to track a benchmark corporate bond (or index) which has observable pricing. The resulting estimate of volatility of value is then used to proxy the volatility of the exposure (or portfolio) under analysis.

In some sense, these two approaches reside on opposite extremes, one focusing exclusively on default events prior to maturity, the other exclusively on day-to-day credit spread volatility. Both frameworks have deficiencies. The binomial model causes investment horizon and correlation problems in the calculation of portfolio volatility across exposures of different maturity. In other words, this approach completely ignores the risks associated with changes in value that are recognized only in a mark-to-market framework.

On the other hand, focusing exclusively on observable spread volatility is likely to lead to inefficiencies in estimating the impact of infrequent but important realizations of (up)downgrade or default. Observing a benchmark bond over, for example, the last year will yield one of two qualitative results: no realized migration, resulting in relatively little volatility, or a realized migration, resulting in relatively large volatility. Over many trials, this bias is reduced, but the risk of such estimation error persists due to the infrequency of meaningful credit quality migrations.

CreditMetrics resides between these two extremes. The model estimates portfolio value-at-risk at the risk horizon due to credit events that include upgrades and downgrades, rather than just defaults. Thus it adopts a mark-to-market framework. However, it uses long term estimates of migration likelihood, rather than observations within some recent sample period, so avoiding the problem of biased estimates. Consider Chart 5. Bonds within each credit rating category have volatility of value due to day-to-day credit spread fluctuations. The RAROC approach measures these fluctuations, but will sometimes realize a potentially large move due to a migration. The CreditMetrics approach is probabilistic. It assumes that all migrations might have been realized, weighting each by a migration likelihood.

Chart 5
Construction of volatility across credit quality categories



4.2 Choice of risk horizon

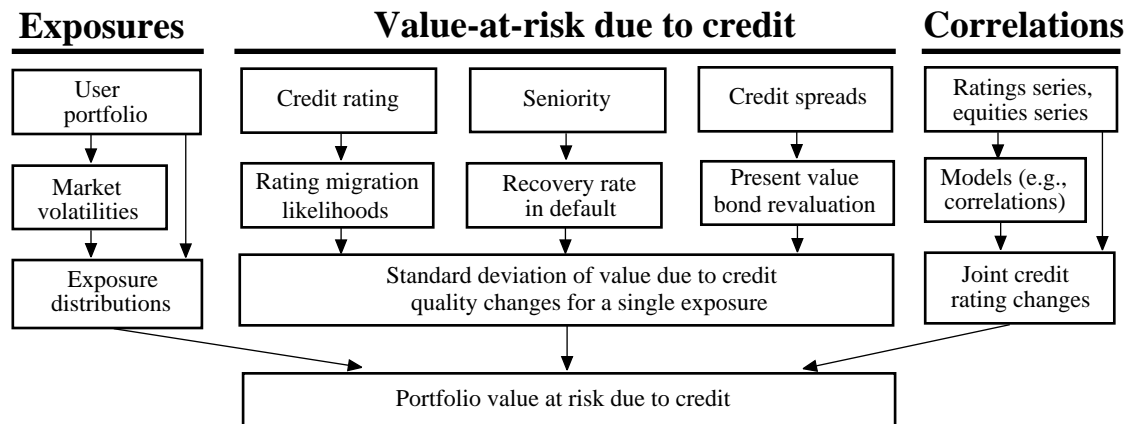
CreditManager, in its initial version, adopts a one-year risk horizon. This is largely because, by convention, much of the academic and credit agency data is stated on an annual basis. However, there is nothing about the CreditMetrics framework that requires a one-year horizon. Indeed, it is difficult to argue that any one particular risk horizon is best, since there is no explicit theory to guide on this point. In a sense, the use of a one-year horizon is merely a convenient convention as is the use of annualized interest rates.

The choice of time horizon implies the degree of activity in portfolio management, so that if risk-mitigating actions are taken more frequently, one-year projections of risk will not be realized. Almost any risk measurement system, however, is better at stating relative rather than absolute risk. Since relative risk measurements are likely to drive decisions, the choice of risk horizon will probably not make an appreciable difference. So long as the risk horizon is not shorter than it would take to perform risk mitigation actions, it is likely to lead to the same qualitative decisions.

4.3 A road map of CreditMetrics methodology

Each step in the CreditMetrics methodology is illustrated in the roadmap presented in Chart 6, which provides greater detail than the simplified roadmap of Chart 2.

Chart 6
Detailed road map of the analytics with CreditMetrics



In the following three sections, each of the major steps in this road map are described in more detail.

The sections describe:

- Step 1: Calculation of the different exposure profiles and dynamics for each exposure type on a comparable basis (*section 4.4*).
- Step 2: Calculation of the volatility of value due to credit quality migration for each individual exposure and the data required to accomplish each step (*section 4.5*).
- Step 3: Calculation of the volatility of value due to credit quality migration across the entire portfolio and different approaches to estimation of correlations of credit quality migrations required for this calculation, (*section 4.6*).

4.4 Estimating credit exposure amounts (step 1)

There are many different types of instruments which embody credit risk. Some have simply a fixed exposure amount, but others create exposures which can vary. Variable exposures can either change in a way that is directly related to up(down)grades – as is the case with commitments to lend – or they can change due to some non credit-related market rate move – as is the case with swaps or forwards. The instrument types we consider are: receivables, bonds, loans, commitments to lend, financial letters of credit and market-driven instruments such as swaps and forwards. There is no reason that the list must end here since the CreditMetrics framework is flexible.

It is important to distinguish between two features of the model: (i) the likelihood of credit quality changes, and (ii) the change in value – assessed against an exposure amount – in the event of each credit quality change. The likelihoods of both default and other credit quality changes are discussed in *section 4.5*. Here, we discuss only how to estimate the amount of credit exposure.

4.4.1. Receivables

Many commercial and industrial firms have credit exposure to their customers through non-interest-bearing *receivables*, or *trade credit*. The exposure on a receivable is treated as its full face amount. It will commonly be the case that a receivable will have a maturity which is shorter than the risk horizon (e.g., one year or less). This simplifies matters in that there is then no need to revalue the exposure upon

up(down)grade, but only upon default. For receivables with a maturity greater than one year revaluation upon up(down)grade is necessary, but the credit risk is – in concept – no different than the risk in a comparable bond issued to the customer, and so it can be revalued accordingly.

4.4.2 Bonds and loans

The exposure on a floating rate bond or loan will always be very close to par. For fixed rate instruments – especially those with longer maturities – there can be more or less exposure since movements in rates can take the current value away from par. The user can elect to treat fixed rate bonds or loans as market-driven instruments, (discussed in section 4.4.5 below) or to ignore this uncertainty in exposure amounts and treat the exposure as the face, or par, amount.

For bonds and loans, the value at the risk horizon is the present value of the remaining cash flows. These cash flows consist of the remaining coupon payments and the principal payment at maturity. To discount the cash flows, one can use the discount rates derived from the forward zero curve for each specific rating category, which will depend on the market credit spread for that category. This forward curve is calculated as of the end of the risk horizon.

4.4.3 Commitments

A loan commitment is essentially a loan (equal to the current amount drawn) and an option to increase the amount of the loan, up to the limit of the commitment. The borrower pays interest on the drawn amount, and a fee on the undrawn amount in return for the option to draw down further.

Historically, the amount drawn under commitments has been closely related to obligor credit quality. If an obligor deteriorates, it is likely to draw down additional funds. If it improves, it is unlikely to need the additional borrowings. This aspect of commitment behavior is closely akin to the “negative convexity” of certain market instruments, although in this case the underlying risk is credit quality rather than market prices. The lender is essentially short an option which is more likely to be exercised, the weaker the credit quality of the borrower.

Consequently, three factors influence the revaluation of commitments in future credit rating states:

- the amount currently drawn
- expected changes in the amount drawn *that are due to credit rating changes*
- the spreads and fees needed to revalue both the drawn and undrawn portions

CreditMetrics captures this aspect of commitment behavior by assuming that the amount drawn under commitments increases upon downgrade and default, and, conversely, declines upon upgrade. Published studies provide estimates of how much is likely to become drawn in the case of default and upon up(down)grade. While the figures from these studies suggest lines are typically *not* fully drawn in default, but that the amount drawn upon default is greater the weaker the credit rating at the time of default, each institution is free to input any other assumptions.

In practice, revaluation of loans, or loan commitments, in different states can be complicated by the presence of repricing grids and covenants specific to each facility. These will tend to reduce volatility, either by allowing the lender to be released from its commitment to lend as the obligor declines in credit quality, or by causing the loan fee or coupon to change to reflect changes in credit rating. In such cases, the user can override the prompted suggestions.

4.4.4 Financial letters of credit

Financial or stand-by letters of credit typically stand as a guarantee against default by the obligor and so they clearly will be fully drawn in default. Thus, for risk assessment purposes, the full nominal amount is considered “exposed”, whether currently drawn or not. Consequently, financial letters of credit – whether or not any portion is actually drawn – are treated exactly like loans.

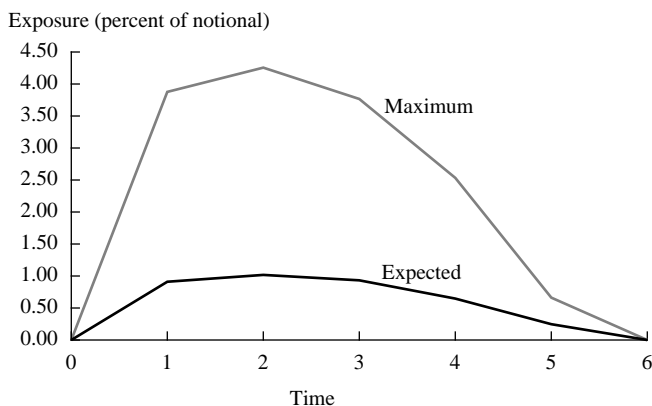
Other types of letters of credit may be either secured by a specific asset or triggered only by some contingent event. The unique “two-name” or “dual contingency” features of these types of letters of credit are not addressed within the current specification of CreditMetrics.

4.4.5 Market-driven instruments

For instruments whose credit exposure depends on the moves of underlying market rates, such as swaps, forwards and, to a lesser extent, fixed rate bonds⁶, revaluation at future rating states is more difficult. The complexity of these instruments comes from the fact that credit and market risk components are intimately coupled. This stems from the fact that losses are incurred on the transaction only if the counterparty owes money on a mark-to-market basis when defaulting or declining in credit quality. Whether or not the mark-to-market of the swap is positive at that time will, of course, depend on market rates.

Not only is exposure uncertain due to market rate moves, but it also changes over the life of the instrument. Chart 7 illustrates this changing profile of expected and peak exposure over the life of an interest rate swap. CreditMetrics uses only expected exposures. Peak exposures cannot all occur simultaneously and so there is no clear way to aggregate them across a portfolio.

Chart 7
Profile of market-driven exposures over time



In default, the future value of market driven instruments is estimated using the expected exposure of the instrument at the risk horizon. Expected exposure depends on both current market rates and their

⁶ On a percentage basis, the uncertainty or standard deviation of exposure amounts for bonds will be far less than that of swaps and other off-balance-sheet instruments. To illustrate, consider a swap, with \$100 million notional, expected exposure of \$10 million and standard deviation of exposure (arising from interest rate volatility) of 10 percent of notional — that is, \$10 million. Expressed as a percentage of expected exposure, this standard deviation is very large — 100 percent. Compare this with a fixed rate bond, with \$100 million face amount, and the same coupon and interest rate sensitivity. Expected exposure is \$100 million, and standard deviation of exposure \$10 million, i.e., only 10 percent of expected exposure. Because the percentage uncertainty of expected exposure is smaller, it is less important to address exposure uncertainty in bond portfolios than swap portfolios.

volatilities. In non-default states, the revaluation consists of two parts: the present value of future cashflows, and the amount that might be lost should the counterparty default at some future time. First, the present value of all future cashflows at the risk horizon is computed as if these were risk-free. Second, an amount equal to the expected amount of losses due to credit events over the remaining life of the instrument is subtracted. This expected loss depends on the average market-driven exposure over the remaining life (estimated in a similar fashion to expected exposure), the probability that the counterparty will default over the same time (which probability is determined by the credit rating at the risk horizon), and the recovery rate given default. As mentioned earlier, expected and average exposure calculations for market driven instruments are not performed by CreditManager, but can be performed in and imported from FourFifteen, or any other source.

This methodology, which is applicable to single instruments or to any groups of instruments whose exposures can be netted, is explained in greater detail in the Technical Document and also in the Risk-Metrics Monitor, Peter Zangari, *On measuring credit exposure* (first Quarter 1997).

4.5 Calculating volatility of value due to credit quality changes (step 2)

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

- *Step A:* The senior unsecured credit rating of the issuer determines the chance of either defaulting or migrating to any possible credit quality state at the risk horizon (*section 4.5.1*).
- *Step B:* Revaluation at the risk horizon can take two forms: (i) the seniority of the exposure determines its recovery rate in the case of default, and (ii) the forward zero curve for each credit rating category determines the revaluation upon up(down)grade (*section 4.5.2*).
- *Step C:* The likelihoods from Step A and the values from Step B are then combined to calculate volatility of value due to credit quality changes (*section 4.5.3*).

In this section, each is considered in turn.

4.5.1 Step A: Estimating credit quality migrations (transition matrices)

CreditMetrics recognizes risk not only due to default but also to changes in value caused by up(down)grades. Thus, it is important to estimate not only the likelihood of default but also the chance of migrating to any possible credit quality state at the risk horizon. The likelihood of any migration is conditioned on the senior unsecured credit rating of the obligor.

Table 1 below shows a transition matrix with data sourced from Standard & Poor's. To read this table, find today's rating on the left and follow along that row to the column representing the rating at the risk horizon. For instance, the likelihood that a BBB will upgrade to a single-A in one year is 5.95%.

Table 1
One-year transition matrix (%)

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 (15 April 96)

Transition matrices published by rating agencies have been calculated by observing the historical pattern of rating change and default. In addition to providing transition matrices from these sources, the CreditMetrics data set includes "smoothed" transition matrices adjusted to satisfy certain desirable long-term characteristics which are often violated when using historically tabulated data because of unavoidably small sample sizes. These adjustments are discussed in detail in the CreditMetrics Technical Document.

Since CreditMetrics does not attempt to provide, nor advocate, any particular credit-rating or scoring methodology, the user has the option to use any alternative transition matrix, including a proprietary methodology. The Technical Document discusses how transition matrices can be fitted for different rating systems.

4.5.2 Step B: Estimating changes in value upon credit quality migration

Given the likelihood of a credit migrating to any other credit rating at the risk horizon, each instrument must be revalued in each different state of credit quality. There are two different types of revaluation:

1. Revaluation in default, which is driven by the seniority standing of the issue or instrument.
2. Revaluation upon up(down)grade, which is driven by credit spread changes.

4.5.2.1 Revaluation in default: Recovery rates

In default, likely recovery rates depend on the seniority class of the debt. CreditMetrics provides suggested recovery rates based upon several historical studies of this dependence. Table 2 summarizes the recovery rates as reported by one of the available studies (relating to rated public bonds). Also provided are separate recovery rates for bank loans based upon studies that have essentially treated bank facilities as a seniority class of its own – with this being senior to all public bond seniority classes

Table 2

Recovery rates by seniority class (% of face value, i.e., “par”)

Seniority Class	Mean (%)	Standard Deviation (%)
Senior Secured	53.80	26.86
Senior Unsecured	51.13	25.45
Senior Subordinated	38.52	23.81
Subordinated	32.74	20.18
Junior Subordinated	17.09	10.90

Source: Carty & Lieberman [96a] –Moody’s Investors Service

Notice that this study reports a wide uncertainty to these mean (average) recovery rate estimates. This contributes to the overall risk of the position, and is addressed in more detail below.

The available studies rely heavily on U.S. bankruptcy experience. Since bankruptcy practice differs by jurisdiction, the user may wish to adjust recovery rates in different jurisdictions to take account of differing bankruptcy practice. Also, major differences will apply to secured versus unsecured debt and there are no studies which isolate the effects of pledged security. Consequently, this will become a practical problem requiring attention on a case-by-case basis.

The user can either use the recovery rate information provided in the CreditMetrics data set, or override this with any other alternative.

4.5.2.2 Revaluation upon up(down)grade: Credit spreads


To obtain the values at the risk horizon corresponding to rating up(down)grades, requires a straightforward present value revaluation. This involves the following:

1. Obtaining the forward zero curves for each rating category. These forward curves are stated as of the risk horizon and go to the maturity of the bond. The CreditMetrics data set will provide current market yield and credit spread data by currency, rating category, industry and product as available, although the user may specify more detailed credit spreads curves.
2. Using these zero curves, revaluing the bond's remaining cash flows at the risk horizon for each rating category; a simple present-valuation exercise.

4.5.3 Step C: Compute distribution of bond value

To summarize, the likelihood of all possible credit events and the instrument’s value, given each credit event, have now been established. Given this information, the volatility of value due to credit quality changes for this one exposure, on a stand-alone basis, can be calculated. This process of constructing the distribution of values for a single bond is illustrated in Chart 8 below.

Chart 8
Constructing the distribution of value for a BBB bond



$$\sigma_T = \sqrt{\sum_{i=1}^s p_i (\mu_i^2 + \sigma_i^2) - \mu_T^2} \quad \text{where } \mu_T = \sum_{i=1}^s p_i \mu_i$$


The computation of mean and standard deviation from the likelihoods and values obtained from Steps 1 and 2, respectively, is illustrated in Table 3 below.

Table 3
Calculating volatility in value due to credit quality changes for a single exposure

Year-end rating	Probability of state (%)	New bond value plus coupon (\$)	Probability weighted value (\$)	Difference of value from mean (\$)	Probability weighted difference squared
AAA	0.02	109.37	0.02	2.28	0.0010
AA	0.33	109.19	0.36	2.10	0.0146
A	5.95	108.66	6.47	1.57	0.1474
BBB	86.93	107.55	93.49	0.46	0.1853
BB	5.30	102.02	5.41	(5.06)	1.3592
B	1.17	98.10	1.15	(8.99)	0.9446
CCC	0.12	83.64	1.10	(23.45)	0.6598
Default	0.18	51.13	0.09	(55.96)	5.6358
		Mean =	\$107.09	Variance =	8.9477
				Standard deviation =	\$2.99

This information can be used to specify the full distribution of possible values at the one year horizon, illustrated in Chart 9. CreditMetrics uses this type of distribution – and more complex ones representing portfolios – to estimate percentile levels directly.

Chart 9
Distribution of value for a five-year BBB bond in one year



4.5.4 Uncertainty of recovery rates

In the calculations of Table 3 and the resulting Chart 9, it was assumed for simplicity that there is no uncertainty in value for recovery values in default or up(down)grade. This approach is simplistic because a bond can actually take on a distribution of values within each state. In other words, there is uncertainty in the value assumptions made. In particular, recovery rates in the event of default are notoriously uncertain, as illustrated by the very large historical standard deviations in Table 2.

CreditMetrics methodology takes account of the uncertainty in recovery rates. Essentially this requires the addition of a term for the standard deviation of the recovery rate to the standard deviation formula. A more precise description of the mathematics required to achieve this is given in the Technical Document.

To illustrate, incorporating an uncertainty of 25.45 percent — the recovery rate uncertainty of senior, unsecured bonds taken from Table 2 — increases the standard deviation of value of the bond in this example from \$2.99 to \$3.18 (an increase of 6.3 percent). The magnitude of this adjustment highlights the importance of addressing recovery rate uncertainty. It is noteworthy that allowing for uncertainty in recovery rates makes the results less sensitive to errors in *expected* recovery rates.

Just as there is an uncertainty in value in the default state, there is uncertainty in value in other rating states, caused by volatility in credit spreads. For now, CreditMetrics methodology assumes credit spread volatility is zero since data limitations have prevented any analysis of what portion of it is systematic versus diversifiable (again, a function of correlations, but in this case credit spread correlations). Given sufficient data to resolve this issue, future versions of CreditMetrics will allow for credit spread volatility.

4.6 Estimating credit quality correlations and calculating portfolio risk (step 3)

For a portfolio with one bond and an 8 states rating system, there are 8 values that could be observed at the end of the year. For a portfolio of two bonds, the number of possible values increases to 64 (that is 8²). Generalizing, if there are *N* bonds, the portfolio can be in any one of the 8^{*N*} possible states of credit quality. The value of the portfolio in each of those 8^{*N*} states is simply obtained by adding the values of the individual instrument values in the corresponding states. However, to obtain the portfolio value distribution requires the estimation of the probability of observing these values. This would be

simple if the rating outcomes of different obligors were independent of each other. In this case, the “joint likelihood” – the likelihood that an obligor will be in any rating state at the end of the risk horizon, given the state of the other assets – is simply a product of the individual likelihoods.

Unfortunately, both empirical evidence and intuition suggest this is incorrect. Rating outcomes on different obligors are not independent of each other because they are affected in part by the same economic factors. Estimation of joint likelihoods therefore requires some measure of the interdependence or correlation between rating outcomes, in addition to the individual likelihoods provided by transition matrices. In a portfolio of N names, with eight states of credit quality, this involves estimating not simply default correlations, but the correlations between credit quality migrations to 8^N possible outcomes.

Empirically, correlation data are the most complex and potentially controversial element in credit portfolio modeling. This is because data from which this information can directly be observed are typically sparse and of poor quality, but on the other hand, models that infer this information from other, more readily-observable data are forced to employ ambitious assumptions in converting from the available to the desired data set. Some alternative approaches are discussed in the next section.

4.6.1 Alternative approaches to estimating credit quality correlations

There are a number of alternative approaches to correlation estimation, including:

- **actual rating and default correlations**⁷: derived from rating agency data, these provide an objective measure of actual experience, but suffer from sparse sample sets, requiring as a practical matter that all obligors with a given credit rating be treated as identical
- **bond spread correlations**: these provide perhaps the most objective measure of actual correlation in bond values and credit quality but are plagued by data quality problems, particularly for low credit quality issuers, making this approach all but impossible in practice
- **uniform constant correlation**: benefiting from simplicity and speed of calculation, a uniform correlation assumption, while imprecise, suffices to highlight significant overconcentrations to individual obligors. This simplified approach is certainly preferable to ignoring correlations altogether. Nonetheless it limits the analysis of concentrations along other dimensions, such as by industry or country
- **equity price correlations**: equity prices provide forward-looking, efficient market information and offer the advantage of uninterrupted time series and a broad universe of names. However, they require much more heavy processing before they will yield information about likely credit quality correlations

CreditMetrics neither advocates nor necessitates the use of any one methodology over others. The model allows the user to input any desired correlation assumptions. For users desiring a higher degree of accuracy than a constant correlation assumption, the CreditMetrics data set provides correlations aggregated at the industry and country level, derived from equity price data, discussed in the following section.


⁷ In the Technical Document we use analyses of actual default and rating experience to provide empirical evidence of positive correlations. This revealed positive, statistically significant, default and rating change correlations, supporting the argument that correlations cannot be ignored in a risk assessment model.

4.6.2 Inferring credit quality correlations from equity prices

To infer information about likely credit quality correlations from equity price data first requires a model linking firm asset value to changes in firm credit quality. This model is described more fully in the Technical Document, but it is touched on briefly here. Illustrated in Chart 10, underlying (and volatile) firm value can be thought of as being randomly distributed according to some distribution. Assuming liabilities to be constant, then if the value of assets should happen to decline below the value of outstanding liabilities, then the firm may be unable to meet its obligations and thus default.⁸

Chart 10

Model of firm value and default




It follows that the volatility of asset values should directly predict the chance of default by any firm. However, since CreditMetrics does not propose a credit rating methodology, the approach does not suggest that default likelihoods must be estimated based on the volatility of underlying firm value. However, extending the analysis, it follows that asset volatilities will also drive the joint default probability between any pair of firms. A positive correlation between asset returns should directly imply some positive correlation in default expectations. The actual transformation of asset correlations to default correlations involves a downscaling that reflects the low probability that two unlikely events will occur simultaneously.

Default correlations are not sufficient in the framework, which requires consideration of correlations in all changes in credit quality. Fortunately, this model is easily extended to include rating changes by generalizing the default threshold to include thresholds for all credit quality categories, as illustrated in Chart 11. Ultimately, this framework creates a link between the underlying firm value and the firm's credit rating, allowing the joint probabilities for two firms to be built from a knowledge of the correlation between the two firm's asset values.

⁸ This framework is the same option theoretic valuation of debt, originally proposed by Robert Merton (1974), referenced in footnote 4, and later discussed in Stephen Kealhofer, "Managing default risk in portfolios of derivatives," *Derivative Credit Risk: Advances in Measurement and Management*, (Renaissance Risk Publications, 1995).

Chart 11
Model of firm value and migration

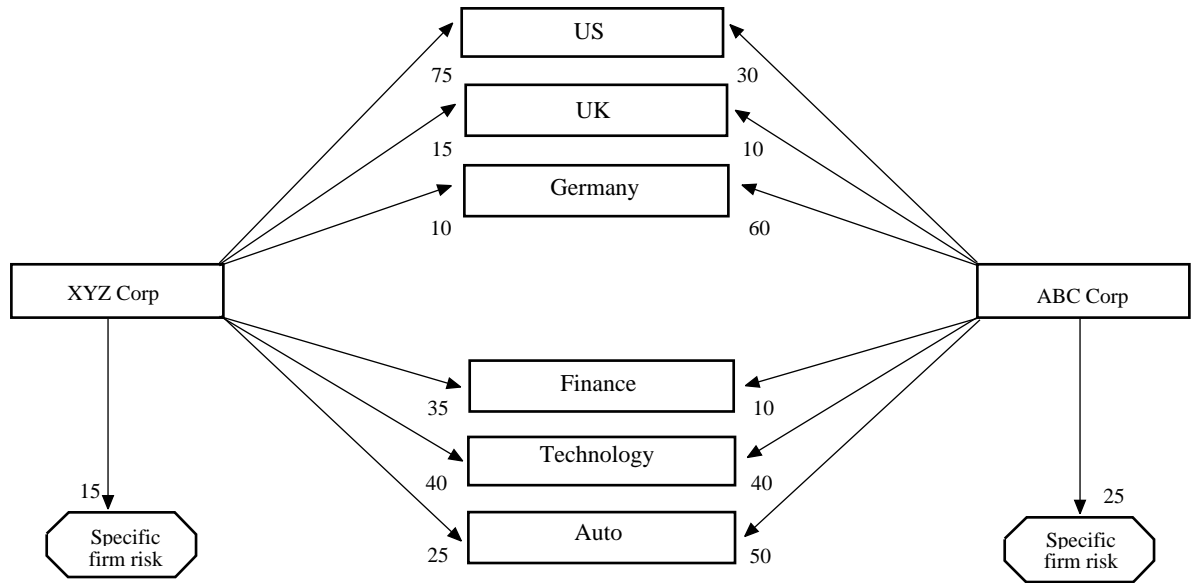


The final link, then, in estimating credit quality correlations is to estimate firm asset correlations, of which, fortunately, equity prices are a reasonable approximation.

For the tens of thousands of potential obligors, the model described would ideally require calculating an enormous matrix of default correlations for each possible combination of obligors. To simplify considerably, CreditMetrics provides correlation data calculated using an approach which maps each obligor to the industries and countries most likely to determine its performance. The portfolio manager must attribute these industry and country weightings to each obligor.⁹ Each firm will also have an element of firm-specific risk, which cannot be explained by its sensitivity to a particular sector or country. Firm-specific risk can be specified, but is prompted by CreditManager as a function of obligor market capitalization (the smaller the obligor, the more idiosyncratic its likely behavior). Having mapped each obligor to an industry/country/specific risk combination, firms can be related to one another via their common sensitivity to industry/country sectors, as illustrated in Chart 12. This essentially reduces the size of each axis of the required correlation matrix from the number of names in the portfolio to the number of countries and industries in the analysis. This mapping also allows users to calculate correlations on firms that are not publicly traded, or that have illiquid equity issues. In fact, any firm can be correlated against any other as long as they participate in the industries and/or countries covered.

⁹ This could be based on intuition or perhaps on an analysis of revenues or sales by industry and country. We expect that industry and country allocations will be provided by data-providers, if not in the first release, then in subsequent versions of the CreditMetrics data set.

Chart 12
A framework for consolidating correlations



4.7 Obtaining a distribution of values for a portfolio of many bonds


Given a possible eight states of credit quality, a portfolio of only five bonds could take on 32,768 different possible values at year end. As mentioned earlier, the number of possible outcomes multiplies quickly as the number of obligors in the portfolio increases.¹⁰ A plot of the distribution of values therefore starts looking more like a smooth curve than the collection of a few discrete points shown on Chart 13. In practice, it is computationally infeasible to cycle through all the possible portfolio states to obtain the value distribution for a portfolio of many bonds. Instead, it is possible to concentrate on a reduced set of these portfolio values, selected randomly through simulation so as not to introduce any bias in the selection process. The larger the random sample, the more closely it will approach the smooth distribution, as illustrated in Chart 13.

Simulation therefore provides a way of obtaining the value distribution in a manageable time. Even then, simulations are computationally complex and take a long time to run. CreditManager automatically tracks several statistics to describe the shape of the portfolio distribution during a simulation. These include not only portfolio mean and standard deviation, but the first four moments¹¹ of the distribution, as well as percentile levels, discussed in more detail in Section 4.8.

¹⁰This number will also vary depending on how many possible states of credit quality there are. The examples given have been using the standard rating agency "eight state" grading for illustration, but any other number of states could be used within CreditMetrics.

¹¹The first four moments of any distribution are: mean, variance, skewness (a measure of distribution asymmetry), and kurtosis (a measure of distribution peakedness or flatness). Credit portfolios typically have negative skew (i.e., downside biases) and a downside tail that is "fat" (i.e., kurtosis greater than 3.0 - also termed leptokurtosis).

Chart 13
Simulating the distribution of credit returns



4.8 Risk measures (model output)

Both measures of risk discussed in section 1.9 can now be computed at the portfolio level. These are the portfolio standard deviation and percentile levels.

4.8.1 Standard deviation

The standard deviation measures the symmetrical dispersion around the average portfolio value. If credit returns were normally distributed, then the standard deviation measure could be easily interpreted. A standard statistical table would show that 1.64 standard deviations would cover all but 5% of possible downside moves. However, since credit returns are lopsided, such interpretations are not possible.

Table 3 provides a specific numerical example for a BBB-rated bond. In this example, 1.64 standard deviations is equal to \$4.90 (The standard deviation of \$2.99 multiplied by 1.64). However, \$4.90 below the mean portfolio value of \$107.09 is \$102.19. The first column of the table reveals that it is not 5 percent, but 6.77 percent likely that the bond value will fall below this level. To compute the 6.77 simply requires calculating the cumulative probability of values below 102.19, i.e., $0.18 + 0.12 + 1.17 + 5.30 = 6.77$ percent. This illustrates that for a credit portfolio, using a given number of standard deviations of portfolio value will tend to understate the risk or provide less "coverage" of downside events than for a normally-distributed (market) portfolio.

In spite of this, there are two advantages to being able to calculate a standard deviation analytically. First, it can be done quite a bit faster – especially for marginal risk analysis – than the alternative calculation which is simulation. Second, increases in standard deviation are *diagnostic* of increases in the downside tail. Thus, even if the standard deviation cannot give particular percentile levels, it will never the less give a directional signal as to relative changes in risk.

4.8.2 Percentile levels

The interpretation of the percentile levels is much simpler than the standard deviation; the likelihood that the portfolio value falls below the 5th percentile level is 5 percent. The 5th percentile level could thus provide a probabilistic lower bound on the year-end portfolio value. No particular percentile level is "best"; the figure should simply reflect a portfolio manager's preference.

To calculate a percentile level requires full specification of the distribution of portfolio values. While this can only be achieved through a computationally complex simulation, the calculation of the percentile level itself is simple. Consequently, the user can potentially view a series of percentile levels simultaneously.

Once again in the example of Table 3, where the number of possible outcomes is small (8 credit quality states), simulation is not necessary. Instead, the entire set of possible values and corresponding likelihoods can be specified. The calculation proceeds by starting from the bottom of the first column labeled “Probability of state”, and keeping a running total of the likelihoods while moving up the table. The value at which this running total first becomes equal to or greater than 5 percent is the 5th percentile level. Using these numbers, the 5th percentile level is crossed at the BB-rating threshold, which has a revaluation of \$102.02. The amount at risk due to credit is equal to the difference between this and the average value, or \$5.07 ($= \$107.09 - \102.02). (One can compare this result with the \$4.90 arrived at using standard deviation).

Computational requirements introduce a trade-off between using the standard deviation and using percentile levels. The percentile level is intuitively appealing because it shows a precise downside limit that is not possible using a standard deviation. On the other hand, it is computationally simpler to calculate the standard deviation.

Users should evaluate this trade-off carefully and use the risk measure that best fits their purpose. It is worth noting, however, that the tails of a credit portfolio distribution are not likely to be subject to constant or frequent change. Consequently, the use of a standard deviation measure to monitor changing portfolio riskiness on an interim basis is more useful once the portfolio manager has established a sense of the tails of the credit portfolio distribution using percentile level analysis.

4.8.3 Marginal risk statistics

Decisions to buy, sell or hold an exposure are likely to be made within the context of an existing portfolio. The relevant calculation is then not the stand-alone risk of that exposure but the marginal increase to the portfolio risk that would be created by adding an exposure to it. Consequently, CreditMetrics also provides marginal risk statistics.

Marginal risk refers to the difference between the total portfolio risk before the marginal transaction versus after. If the new transaction adds to an already over-concentrated portion of the portfolio then this marginal risk is likely to be large. If the new transaction is diversifying (or in the extreme is actually a hedge position), then the marginal risk may be quite small or even negative.

The importance of calculating the marginal risk is that it captures the specific characteristics of a particular portfolio. It would not be unusual for a given transaction to be considered risky in one institution's portfolio but of considerably lower risk in another institution's portfolio.

Mechanically, a marginal risk statistic can be calculated using either standard deviations or percentile levels. The point is the same: to show the change in total portfolio risk upon the addition of a new transaction.

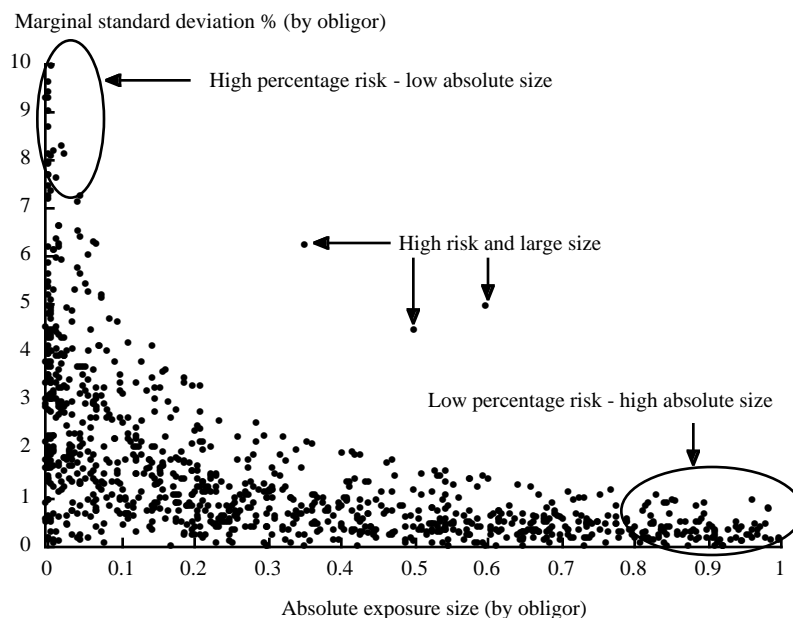
5. Practical applications

Finally, having completed an overview of CreditMetrics methodology and key data requirements, this section illustrates three potentially powerful applications of CreditMetrics in a practical risk management setting.

5.1 Prioritizing risk-reducing actions

The primary purpose of any risk management system is to direct and prioritize actions. To illustrate this discussion, Chart 14 shows risk versus size of exposures within a typical credit portfolio.

Chart 14
Exposure size versus risk



When considering risk-mitigating actions, there are various features of risk worth targeting, including perhaps obligors having the largest:

- *absolute size* (lower right hand corner of the chart)
- *percentage level* of risk (upper left corner of the chart)
- *absolute amount of risk* (upper right hand corner of the chart)

Although each approach is valid, the last is most appealing since it sets as priority those obligors that are *both* relatively high percentage risk and relatively large exposure. In practice, such outliers may be the result of fallen-angels, whose now excessive exposures were appropriate when originated, or simply relationship-driven concentrations. Clearly, the portfolio manager can readily identify these in the CreditMetrics framework.

Thus far, the analysis has made no reference to returns, and whether they adequately compensate risk. A consequence of portfolio analysis is that it highlights where an asset may contribute differently to the risk of distinct portfolios, and yet yield the same returns in either case. Consequently, it is easy to imag-

ine a situation in which two managers identify two assets of the same maturity, yield and credit rating, but – because of the composition of the two portfolios – the risk of *both* portfolios is reduced by swapping these assets. This is accomplished without a loss in return: the financial equivalent of a *free lunch*. The importance of identifying the nature of the contribution of each asset to portfolio risk is obvious. The risk of credit assets is largely due to concentrations particular to the portfolio. Thus, opportunities may exist to restructure the portfolio to reduce risk, with no change to profitability.

5.2 Risk-based exposure limits

Traditionally, credit risk limits have been based on intuitive – but arbitrary – exposure amounts. This approach is unsatisfactory because resulting decisions are not risk-driven. Consequently, the next step beyond using risk statistics for prioritization is to use them for limit-setting.

In the previous section, it was argued that it is best to address exposures with the highest level of absolute risk first, since these have the greatest impact on the total portfolio risk. Similarly, it is most sensible to set credit limits according to the absolute contribution to portfolio risk. This would correspond to a limit resembling the curve defined by the boundary of non-outlying scatter points in Chart 14. Such a limit would prevent the addition to the portfolio of any exposure which increased portfolio risk by more than a given amount.

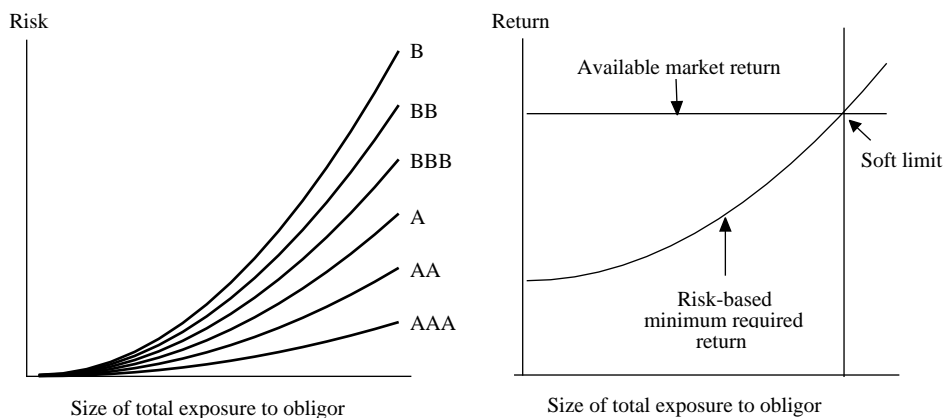
Limiting absolute risk is consistent with the natural tendencies of portfolio managers. Intuitively, exposures posing a greater chance of decreases in value due to credit should be smaller and vice versa. Thus, setting limits based on absolute risk takes the qualitative intuition that currently drives decisions and makes it quantitative.

It is worth mentioning that such risk limits are not meant to replace existing limits to individual names. Limits based on the notion that there is a maximum amount of desired exposure to a given counterparty, regardless of this counterparty's credit standing, are certainly appropriate. Such limits may be thought of as conditional, in that they reflect the maximum acceptable amount of loss conditioned on a counterparty's defaulting, regardless of the probability that the counterparty actually defaults. The limits proposed in this section can supplement, but not replace, these conditional limits. Sound business judgement will never be replaced by any quantitative tool, rather the tool should be used as an aid to expert opinion.

As a portfolio manager considers limit setting, he can use several statistics: first, marginal or stand alone risk statistics, and second, the marginal standard deviation or percentile levels. The case for marginal risk statistics is clear: these allow the user to examine an exposure in the context of the portfolio, capturing the effects of concentration and diversification. As discussed earlier, the choice of standard deviation or percentile levels involves a trade-off between computational complexity and the greater precision of percentile levels.

Again, this approach has thus far made no reference to returns. Chart 15 illustrates how marginal risk statistics can be used to make credit limits sensitive to the trade-off between risk and return.

Chart 15
Risk-based exposure limits



The left hand graphic reflects how marginal contribution to portfolio risk increases geometrically with exposure size of an individual obligor, more noticeably so for weaker credits. Consequently, as illustrated in the right hand graphic, to maintain a constant balance between risk and return, proportionately more return is required with each increment of exposure to an individual obligor.

The marginal risk of the obligor is shown by the curve labeled “risk-based minimum required return.” This represents the set of constant return-on-risk opportunities. The horizontal straight line represents the available market return for the obligor. Areas to the right of the intersection between the two lines represent situations in which the portfolio manager might decrease exposure, because doing so would more than proportionately reduce risk relative to the loss in return. Areas to the left of the intersection represent situations in which the portfolio manager might increase exposure, because doing so would more than proportionately increase return relative to the increase in risk.

The vertical line marking the point of intersection is thus effectively the “soft” risk-based limit for the obligor, given available market pricing, the portfolio targeted risk-return balance and the relationship between the obligor and the existing portfolio. The appeal of a limit that is an output of a model that is sensitive to both risk and return is unmistakable.

5.3 Risk-based capital allocation

Finally, this section examines the application of credit risk measures to assessment of economic capital that a firm puts at risk by holding a credit portfolio. In this context, we are no longer trying to compare different exposures and decide which contribute most to the riskiness of the portfolio. Rather, we are seeking to understand the risk of the entire portfolio with regard to what this implies about the stability of an organization.

Essentially, in this framework, risk is measured in terms of its threat to shareholder capital. The idea is that if a firm’s liabilities are constant, then it is taking risk by holding assets that are volatile. Such risk-taking capacity is not unlimited and must be allocated as a scarce resource. For example, if a manager found that there was a ten percent chance of a decline in portfolio value occurring in the next year severe enough to cause organization-wide insolvency, then he would likely seek to decrease the risk of the asset portfolio. For a portfolio with a more reasonable level of risk, the manager cannot add new exposures indiscriminately, since eventually the portfolio risk will surpass the “comfort level.” Thus, each additional exposure utilizes a scarce resource, which is commonly thought of as risk-taking capacity, or alternatively, as economic capital.

Consequently, as an indicator of economic capital, a percentile level seems quite appropriate. Using, for example, the first percentile level, economic capital could be defined as the level of losses on the portfolio that, with 99 percent certainty, will not be exceeded in the next year.

Such an approach to economic capital allocation contrasts starkly with the framework currently mandated for bank regulation under the BIS accord described in section 2.2.4. The approach described here benefits from the comparison in at least three respects:

- sensitivity to obligor credit quality
- sensitivity to portfolio concentrations
- uniform treatment of value-at-risk due to credit, irrespective of the underlying instrument.

5.4 Conclusion

In conclusion, CreditMetrics is a significant innovation for risk managers seeking to apply recent advances in portfolio theory and value-at-risk methodology to credit risk. Recognizing the need for an industry benchmark for credit risk measurement, we have set out to provide as good a methodology for capturing credit risk as the practical constraints of available data quality will currently allow. We have focused consistently on the goal of implementing this methodology in a practical environment and provided data and software to facilitate this end.

Once implemented, the methodology has several powerful applications, with important implications for the way in which institutions think about pricing, trading and carrying credit risks. The applications include prioritizing and evaluating investment and risk-mitigating transactions, setting rational, risk-based limits, and ultimately, the maximization of shareholder value based on risk-based capital allocation. The implications include the longer-run liquidity of credit markets, the emergence of a mark-to-market approach to credit positions, and the potential for closer alignment of regulatory and economic capital.





CreditMetrics™ Products

Introduction to CreditMetrics™: An abbreviated document which broadly describes the CreditMetrics™ methodology for measuring portfolio credit risk.

CreditMetrics™ – Technical Document: A manual describing the CreditMetrics™ methodology for estimating credit risks. It fully specifies how we construct the volatility of value due to credit quality changes for both stand-alone exposures and portfolios of exposures. It also discloses our approach to estimating credit exposures by instrument type and a method of estimating correlations of credit quality co-movements.

CreditMetrics™ Monitor: A semiannual publication which will discuss broad credit risk management issues, statistical questions as well as new software implementations and enhancements.

CreditMetrics™ data sets: A set of historical statistics and results of academic and industry studies which will be updated periodically.

All the above can be downloaded from the Internet at <http://www.jpmorgan.com/RiskManagement/CreditMetrics>

CreditManager™ PC Program: A desktop software tool that implements the methodology of CreditMetrics and produces value-at-risk reports and other analysis of credit risk such as those outlined in the CreditMetrics documents. CreditManager can be purchased from J.P. Morgan and any of the the co-sponsors.

Trouble accessing the Internet? If you encounter any difficulties in either accessing the J.P. Morgan home page on <http://www.jpmorgan.com> or downloading the CreditMetrics™ data files, you can call (1-800) JPM-INET in the United States.

Worldwide CreditMetrics™ Contacts

For more information about CreditMetrics™, please contact the author or any co-sponsors listed below:

J.P. Morgan

Americas	(1-212) 648-3461 cmx_amer@jpmorgan.com
Europe	(44-171) 325-8007 cmx_euro@jpmorgan.com
Asia Pacific	(852) 2973-5646 cmx_asia@jpmorgan.com

Co-sponsors

Bank of America	Janet M. Tavakoli (1-312) 828-4732 Philip Basil (44-171) 634-4482 Walter Bloomenthal (1-312) 828-1668
Bank of Montreal	Barry Campbell (1-416) 867-4809 Loretta Hennessey (1-212) 605-1541
BZW	Jo Ousterhout (1-212) 412-6893 Michael Dyson (44-171) 696-3045 Loren Boston (852) 2903-2000
Deutsche Morgan Grenfell	Hugo Bänziger (44-171) 545-2562
KMV Corporation	David Nordby (1-415) 756-3337 edfs@kmv.com
Swiss Bank Corporation	Robert Gumerlock (41-1) 239-5739 robert.gumerlock@swissbank.com Linda Bammann (1-212) 335-1085 linda.bammann@swissbank.com
Union Bank of Switzerland	Hei Wai Chan (1-212) 821-5547 nycnh@ny.ubs.com

CreditMetrics™ is based on, but differs significantly from, the credit risk management systems developed by J.P. Morgan for its own use. J.P. Morgan does not warrant any results obtained from the use of the CreditMetrics™ data, methodology, documentation or any information derived from the data (collectively the "Data") and does not guarantee its sequence, timeliness, accuracy, completeness or continued availability. The Data is calculated on the basis of historical observations and should not be relied upon to predict future credit upgrades, downgrades, defaults or market movements. Examples are for illustrative purposes only; actual risks will vary depending on specific circumstances.

Additional information is available upon request. Information herein is believed to be reliable, but J.P. Morgan does not warrant its completeness or accuracy. Opinions and estimates constitute our judgment and are subject to change without notice. Past performance is not indicative of future results. This material is not intended as an offer or solicitation for the purchase or sale of any financial instrument. J.P. Morgan may hold a position or act as market maker in the financial instruments of any issuer discussed herein or act as advisor or lender to such issuer. Morgan Guaranty Trust Company is a member of FDIC and SFA. Copyright 1997 J.P. Morgan & Co. Incorporated. Clients should contact analysts at and execute transactions through a J.P. Morgan entity in their home jurisdiction unless governing law permits otherwise.