

Portfolio Credit Risk

In this chapter, we extend the study of credit risk to portfolios containing several credit-risky securities. We begin by introducing the most important additional concept we need in this context, default correlation, and then discuss approaches to measuring portfolio credit risk.

A portfolio of credit-risky securities may contain bonds, commercial paper, off-balance-sheet exposures such as guarantees, as well as positions in credit derivatives such as credit default swaps (CDS). A typical portfolio may contain many different obligors, but may also contain exposures to different parts of one obligor's capital structure, such as preferred shares and senior debt. All of these distinctions can be of great importance in accurately measuring portfolio credit risk, even if the models we present here abstract from many of them.

In this chapter, we focus on two approaches to measuring portfolio credit risk. The first employs the factor model developed in Chapter 6, the key feature of which is latent factors with normally distributed returns. Conditional on the values taken on by that set of factors, defaults are independent. There is a single future time horizon for the analysis. We will specialize the model even further to include only default events, and not credit migration, and only a single factor. In the CreditMetrics approach, this model is used to compute the distribution of credit migrations as well as default. One could therefore label the approach described in this chapter as “default-mode CreditMetrics.” An advantage of this model is that factors can be related to real-world phenomena, such as equity prices, providing an empirical anchor for the model. The model is also tractable.

The second approach we take in this chapter uses simulation, together with the intensity models developed in Chapters 6 and 7, to measure credit portfolio risk. We sketch the theoretical basis for the approach, which employs a mathematical structure called a copula, and offer a simple example. Chapter 9 provides a full-blown application of the approach to valuation and risk measurement of securitizations.

8.1 DEFAULT CORRELATION

In modeling a single credit-risky position, the elements of risk and return that we can take into consideration are

- The probability of default
- The loss given default (LGD), the complement of the value of recovery in the event of default
- The probability and severity of rating migration (nondefault credit deterioration)
- Spread risk, the risk of changes in market spreads for a given rating
- For distressed debt, the possibility of restructuring the firm's debt, either by negotiation among the owners of the firm and of its liabilities, or through the bankruptcy process. Restructuring opens the possibility of losses to owners of particular classes of debt as a result of a negotiated settlement or a judicial ruling

To understand credit portfolio risk, we introduce the additional concept of *default correlation*, which drives the likelihood of having multiple defaults in a portfolio of debt issued by several obligors. To focus on the issue of default correlation, we'll take default probabilities and recovery rates as given and ignore the other sources of return just listed.

8.1.1 Defining Default Correlation

The simplest framework for understanding default correlation is to think of

- Two firms (or countries, if we have positions in sovereign debt)
- With probabilities of default (or restructuring) π_1 and π_2
- Over some time horizon τ
- And a joint default probability—the probability that both default over τ —equal to π_{12}

This can be thought of as the distribution of the product of two Bernoulli-distributed random variables x_i , with four possible outcomes. We must, as in the single-firm case, be careful to define the Bernoulli trials as default or solvency over a specific time interval τ . In a portfolio credit model, that time interval is the same for all the credits in the book.

We have a new parameter π_{12} in addition to the single-name default probabilities. And it is a genuinely new parameter, a primitive: It is what it

is, and isn't computed from π_1 and π_2 , *unless* we specify it by positing that defaults are independent.

Since the value 1 corresponds to the occurrence of default, the product of the two Bernoulli variables equals 0 for three of the outcomes—those included in the event that at most one firm defaults—and 1 for the joint default event:

| Outcome | x_1 | x_2 | x_1x_2 | Probability |
|----------------------|-------|-------|----------|--------------------------------|
| No default | 0 | 0 | 0 | $1 - \pi_1 - \pi_2 + \pi_{12}$ |
| Firm 1 only defaults | 1 | 0 | 0 | $\pi_1 - \pi_{12}$ |
| Firm 2 only defaults | 0 | 1 | 0 | $\pi_2 - \pi_{12}$ |
| Both firms default | 1 | 1 | 1 | π_{12} |

These are proper outcomes; they are distinct, and their probabilities add up to 1. The probability of the event that at least one firm defaults can be found as either 1 minus the probability of the first outcome, or the sum of the probabilities of the last three outcomes.

$$P[\text{Firm 1 or Firm 2 or both default}] = \pi_1 + \pi_2 - \pi_{12}$$

We can compute the moments of the Bernoulli variates:

- The means of the two Bernoulli-distributed default processes are

$$E[x_i] = \pi_i, \quad i = 1, 2$$

- The expected value of the product—representing joint default—is $E[x_1x_2] = \pi_{12}$.
- The variances are

$$E[x_i]^2 - (E[x_i])^2 = \pi_i(1 - \pi_i) \quad i = 1, 2$$

- The covariance is

$$E[x_1x_2] - E[x_1]E[x_2] = \pi_{12} - \pi_1\pi_2$$

- The default correlation, finally, is

$$\rho_{12} = \frac{\pi_{12} - \pi_1\pi_2}{\sqrt{\pi_1(1 - \pi_1)}\sqrt{\pi_2(1 - \pi_2)}} \quad (8.1)$$

We can treat the default correlation, rather than joint default probability, as the primitive parameter and use it to find the joint default probability:

$$\pi_{12} = \rho_{12}\sqrt{\pi_1(1-\pi_1)}\sqrt{\pi_2(1-\pi_2)} + \pi_1\pi_2$$

The joint default probability if the two default events are independent is $\pi_{12} = \pi_1\pi_2$, and the default correlation is $\rho_{12} = 0$. If $\rho_{12} \neq 0$, there is a linear relationship between the probability of joint default and the default correlation: The larger the “excess” of π_{12} over the joint default probability under independence, $\pi_1\pi_2$, the higher the correlation. Once we specify or estimate the π_i , we can nail down the joint default probability either directly or by specifying the default correlation. Most models, including those set out in this chapter, specify a default correlation rather than a joint default probability.

Example 8.1 (Default Correlation) Consider a pair of credits, one BBB+ and the other BBB-rated, with $\pi_1 = 0.0025$ and $\pi_2 = 0.0125$. If the defaults are uncorrelated, then $\pi_{12} = 0.000031$, less than a third of a basis point. If, however, the default correlation is 5 percent, then $\pi_{12} = 0.000309$, nearly 10 times as great, and at 3 basis points, no longer negligible.

In a portfolio containing more than two credits, we have more than one joint default probability and default correlation. And, in contrast to the two-credit portfolio, we cannot specify the full distribution of defaults based just on the default probabilities and the pairwise correlations or joint default probabilities. To specify all the possible outcomes in a three-credit portfolio, we need the three single-default probabilities, the three two-default probabilities, and the no-default and three-default probabilities, a total of eight. But we have only seven conditions: the three single-default probabilities, three pairwise correlations, and the constraint that all the probabilities add up to unity. It’s the latter constraint that ties out the probabilities when there are only two credits. With a number of credits $n > 2$, we have 2^n different events, but only $n + 1 + \frac{n(n-1)}{2}$ conditions:

| n | 2^n | $n + 1 + \frac{n(n-1)}{2}$ |
|-----|-------|----------------------------|
| 2 | 4 | 4 |
| 3 | 8 | 7 |
| 4 | 16 | 11 |
| 10 | 1,024 | 56 |

We can’t therefore build an entire credit portfolio model solely on default correlations. But doing so is a pragmatic alternative to estimating or

stipulating, say, the 1,024 probabilities required to fully specify the distribution of a portfolio of 10 credits.

Even if all the requisite parameters could be identified, the number would be quite large, since we would have to define a potentially large number of pairwise correlations. If there are N credits in the portfolio, we need to define N default probabilities and N recovery rates. In addition, we require $N(N - 1)$ pairwise correlations. In modeling credit risk, we often set all of the pairwise correlations equal to a single parameter. But that parameter must then be non-negative, in order to avoid correlation matrices that are not positive-definite and results that make no sense: Not all the firms' event of default can be negatively correlated with one another.

Example 8.2 Consider a portfolio containing five positions:

1. A five-year senior secured bond issued by Ford Motor Company
2. A five-year subordinate unsecured bond issued by Ford Motor Company
3. Long protection in a five-year CDS on Ford Motor Credit Company
4. A five-year senior bond issued by General Motors Company
5. A 10-year syndicated term loan to Starwood Resorts

If we set a horizon for measuring credit risk of $\tau = 1$ year, we need to have four default probabilities and 12 pairwise default correlations, since there are only four distinct corporate entities represented in the portfolio. However, since the two Ford Motor Company bonds are at two different places in the capital structure, they will have two different recovery rates.

This example has omitted certain types of positions that will certainly often occur in real-world portfolios. Some of their features don't fit well into the portfolio credit risk framework we are developing:

- Guarantees, revolving credit agreements, and other contingent liabilities behave much like credit options.
- CDS basis trades are not essentially market- or credit-risk-oriented, although both market and credit risk play a very important role in their profitability. Rather, they may be driven by "technical factors," that is, transitory disruptions in the typical positioning of various market participants.

A dramatic example, which we discuss in Chapter 13, occurred during the subprime crisis. The CDS basis widened sharply as a result of the dire lack of funding liquidity.

- Convertible bonds are both market- and credit-risk oriented. Equity and equity vega risk can be as important in convertible bond portfolios as credit risk.

8.1.2 The Order of Magnitude of Default Correlation

For most companies that issue debt, most of the time, default is a relatively rare event. This has two important implications:

1. Default correlation is hard to measure or estimate using historical default data. Most studies have arrived at one-year correlations on the order of 0.05. However, estimated correlations vary widely for different time periods, industry groups, and domiciles, and are often negative.
2. Default correlations are small in magnitude.

In other contexts, for example, thinking about whether a regression result indicates that a particular explanatory value is important, we get used to thinking of, say, 0.05 as a “small” or insignificant correlation and 0.5 as a large or significant one. The situation is different for default correlations because probabilities of default tend to be small—on the order of 1 percent—for all but the handful of CCC and below firms. The probability of any particular pair of credits defaulting is therefore also small, so an “optically” small correlation can have a large impact, as we saw in Example 8.1.

8.2 CREDIT PORTFOLIO RISK MEASUREMENT

To measure credit portfolio risk, we need to model default, default correlation, and loss given default. In more elaborate models, we can also include ratings migration. We restrict ourselves here to default mode. But in practice, and in such commercial models as Moody’s KMV and CreditMetrics, models operate in migration mode; that is, credit migrations as well as default can occur.

8.2.1 Granularity and Portfolio Credit Value-at-Risk

Portfolio credit VaR is defined similarly to the VaR of a single credit. It is a quantile of the credit loss, minus the expected loss of the portfolio.

Default correlation has a tremendous impact on portfolio risk. But it affects the volatility and extreme quantiles of loss rather than the expected loss. If default correlation in a portfolio of credits is equal to 1, then the portfolio behaves as if it consisted of just one credit. No credit diversification is achieved. If default correlation is equal to 0, then the number of defaults in

the portfolio is a binomially distributed random variable. Significant credit diversification may be achieved.

To see how this works, let's look at diversified and undiversified portfolios, at the two extremes of default correlation, 0 and 1. Imagine a portfolio of n credits, each with a default probability of π percent and a recovery rate of zero percent. Let the total value of the portfolio be \$1,000,000,000. We will set n to different values, thus dividing the portfolio into larger or smaller individual positions. If $n = 50$, say, each position has a value of \$20,000,000. Next, assume each credit is in the same place in the capital structure and that the recovery rate is zero; in the event of default, the position is wiped out. We'll assume each position is an obligation of a different obligor; if two positions were debts of the same obligor, they would be equivalent to one large position. We can either ignore the time value of money, which won't play a role in the example, or think of all of these quantities as future values.

Now we'll set the default correlation to either 0 or 1.

- If the default correlation is equal to 1, then either the entire portfolio defaults, with a probability of π , or none of the portfolio defaults. In other words, with a default correlation of 1, regardless of the value of n , the portfolio behaves as though $n = 1$.

We can therefore continue the analysis by assuming all of the portfolio is invested in one credit. The expected loss is equal to $\pi \times 1,000,000,000$. But with only one credit, there are only the two all-or-nothing outcomes. The credit loss is equal to 0 with probability $1 - \pi$. The default correlation doesn't matter.

The extreme loss given default is equal to \$1,000,000,000, since we've assumed recovery is zero. If π is greater than the confidence level of the credit VaR, then the VaR is equal to the entire \$1,000,000,000, less the expected loss. If π is less than the confidence level, then the VaR is less than zero, because we always subtract the expected from the extreme loss. If, for example, the default probability is $\pi = 0.02$, the credit VaR at a confidence level of 95 percent is negative (i.e., a gain), since there is a 98 percent probability that the credit loss in the portfolio will be zero. Subtracting from that the expected loss of $\pi \times 1,000,000,000 = 20,000,000$ gives us a VaR of $-\$20,000,000$. The credit VaR in the case of a single credit with binary risk is well-defined and can be computed, but not terribly informative.

- If the default correlation is equal to 0, the number of defaults is binomially distributed with parameters n and π . We then have many intermediate outcomes between the all-or-nothing extremes.

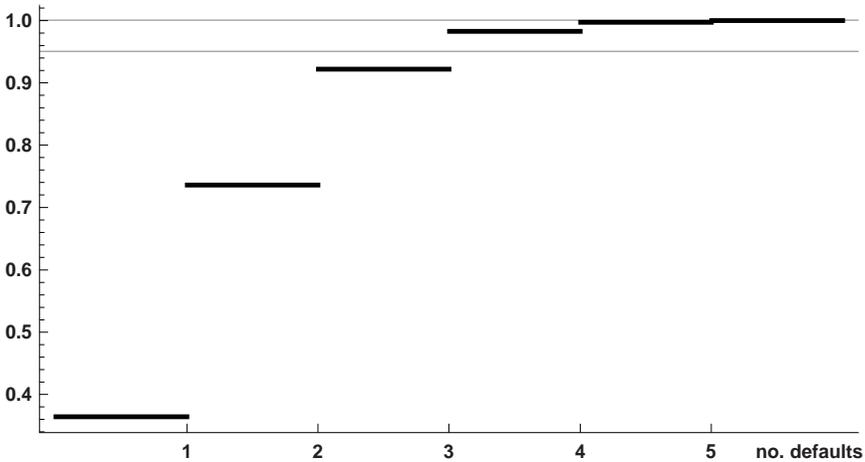


FIGURE 8.1 Distribution of Defaults in an Uncorrelated Credit Portfolio
Cumulative probability distribution function of the number of defaults in a portfolio of 50 independent credits with a default probability of 2 percent.

Suppose there are 50 credits in the portfolio, so each position has a future value, if it doesn't default, of \$20,000,000. The expected loss is the same as with one credit: $\pi \times 1,000,000,000$. But now the extreme outcomes are less extreme. Suppose again that $\pi = 0.02$. The number of defaults is then binomially distributed with parameters 50 and 0.02. The 95th percentile of the number of defaults is 3, as seen in Figure 8.1; the probability of two defaults or less is 0.92 and the probability of three defaults or less is 0.98. With three defaults, the credit loss is \$60,000,000. Subtracting the expected loss of \$20,000,000, which is the same as for the single-credit portfolio, we get a credit VaR of \$40,000,000.

As we continue to increase the number of positions and decrease their size, keeping the total value of the portfolio constant, we decrease the variance of portfolio values. For $n = 1,000$, the 95th percentile of defaults is 28, and the 95th percentile of credit loss is \$28,000,000, so the credit VaR is \$8,000,000.

We summarize the results for $n = 1, 50, 1,000$, for default probabilities $\pi = 0.005, 0.02, 0.05$, and at confidence levels of 95 and 99 percent in Table 8.1 and in Figure 8.2.

TABLE 8.1 Credit VaR of an Uncorrelated Credit Portfolio

| | $\pi = 0.005$ | $\pi = 0.02$ | $\pi = 0.05$ |
|------------------------------------|---------------|--------------|--------------|
| Expected loss | 5,000,000 | 20,000,000 | 50,000,000 |
| <i>n = 1</i> | | | |
| <i>95 percent confidence level</i> | | | |
| Number of defaults | 0 | 0 | 0 |
| Proportion of defaults | 0.000 | 0.000 | 0.000 |
| Credit Value-at-Risk | -5,000,000 | -20,000,000 | -50,000,000 |
| <i>99 percent confidence level</i> | | | |
| Number of defaults | 0 | 1 | 1 |
| Proportion of defaults | 0.000 | 1.000 | 1.000 |
| Credit Value-at-Risk | -5,000,000 | 980,000,000 | 950,000,000 |
| <i>n = 50</i> | | | |
| <i>95 percent confidence level</i> | | | |
| Number of defaults | 1 | 3 | 5 |
| Proportion of defaults | 0.020 | 0.060 | 0.100 |
| Credit Value-at-Risk | 15,000,000 | 40,000,000 | 50,000,000 |
| <i>99 percent confidence level</i> | | | |
| Number of defaults | 2 | 4 | 7 |
| Proportion of defaults | 0.040 | 0.080 | 0.140 |
| Credit Value-at-Risk | 35,000,000 | 60,000,000 | 90,000,000 |
| <i>n = 1000</i> | | | |
| <i>95 percent confidence level</i> | | | |
| Number of defaults | 9 | 28 | 62 |
| Proportion of defaults | 0.009 | 0.028 | 0.062 |
| Credit Value-at-Risk | 4,000,000 | 8,000,000 | 12,000,000 |
| <i>99 percent confidence level</i> | | | |
| Number of defaults | 11 | 31 | 67 |
| Proportion of defaults | 0.011 | 0.031 | 0.067 |
| Credit Value-at-Risk | 6,000,000 | 11,000,000 | 17,000,000 |

What is happening as the portfolio becomes more *granular*, that is, contains more independent credits, each of which is a smaller fraction of the portfolio? The credit VaR is, naturally, higher for a higher probability of default, given the portfolio size. But it decreases as the credit portfolio becomes more granular for a given default probability. The convergence is more drastic with a high default probability. But that has an important converse: It is harder to reduce VaR by making the portfolio more granular, if the default probability is low.

Eventually, for a credit portfolio containing a very large number of independent small positions, the probability converges to 100 percent that

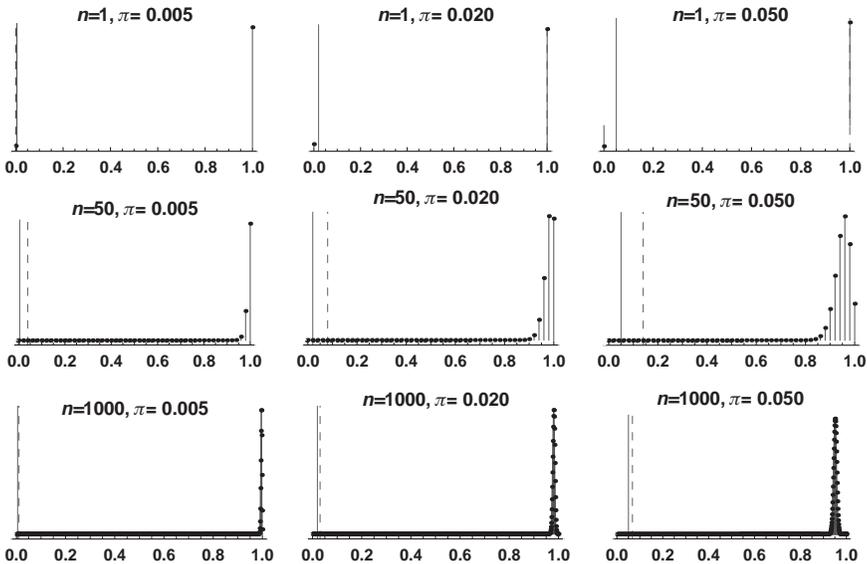


FIGURE 8.2 Distribution of Losses in an Uncorrelated Credit Portfolio

The graph displays the probability density of losses for each combination of a number of equally sized credits and default probabilities. The initial future value of the portfolio is \$1,000,000,000. The values on the x -axis can be interpreted as the fraction of credit losses or as the dollar value of loss in billions. The dashed grid line marks the 99th percentile of loss. The solid grid line marks the expected loss and is the same in each panel.

the credit loss will equal the expected loss. While the single-credit portfolio experiences no loss with probability $1 - \pi$ and a total loss with probability π , the granular portfolio experiences a loss of 100π percent “almost certainly.” The portfolio then has zero volatility of credit loss, and the credit VaR is zero.

In the rest of this chapter, we show how models of portfolio credit risk take default correlation into account, focusing on two models in particular:

1. The single-factor model, since it is a structural model, emphasizes the correlation between the fundamental driver of default of different firms. Default correlation in that model depends on how closely firms are tied to the broader economy.
2. Intensity models emphasize the timing of defaults. Default correlation depends on how many firms default within a given timeframe. It is driven by how default simulation exercises are set up and parameterized.

8.3 DEFAULT DISTRIBUTIONS AND CREDIT VAR WITH THE SINGLE-FACTOR MODEL

In the example of the last section, we set default correlation only to the extreme values of 0 and 1, and did not take account of idiosyncratic credit risk. In the rest of this chapter, we permit default correlation to take values anywhere on $(0, 1)$. We took a first look at the single-factor model, along with other structural credit models, in Chapter 6. The single-factor model enables us to vary default correlation through the credit's beta to the market factor and lets idiosyncratic risk play a role.

8.3.1 Conditional Default Distributions

To use the single-factor model to measure portfolio credit risk, we start by imagining a number of firms $i = 1, 2, \dots$, each with its own correlation β_i to the market factor, its own standard deviation of idiosyncratic risk $\sqrt{1 - \beta_i^2}$, and its own idiosyncratic shock ϵ_i . Firm i 's return on assets is

$$a_i = \beta_i m + \sqrt{1 - \beta_i^2} \epsilon_i \quad i = 1, 2, \dots$$

As in Chapter 6, we assume that m and ϵ_i are standard normal variates, and are not correlated with one another. We now in addition assume the ϵ_i are not correlated with one another:

$$m \sim N(0, 1)$$

$$\epsilon_i \sim N(0, 1) \quad i = 1, 2, \dots$$

$$\text{Cov}[m, \epsilon_i] = 0 \quad i = 1, 2, \dots$$

$$\text{Cov}[\epsilon_i, \epsilon_j] = 0 \quad i, j = 1, 2, \dots$$

Under these assumptions, each a_i is a standard normal variate. Since both the market factor and the idiosyncratic shocks are assumed to have unit variance, the beta of each credit i to the market factor is equal to β_i . The correlation between the asset returns of any pair of firms i and j is $\beta_i \beta_j$:

$$E[a_i] = 0 \quad i = 1, 2, \dots$$

$$\text{Var}[a_i] = \beta_i^2 + 1 - \beta_i^2 = 1 \quad i = 1, 2, \dots$$

$$\begin{aligned} \text{Cov}[a_i, a_j] &= E \left[\left(\beta_i m + \sqrt{1 - \beta_i^2} \epsilon_i \right) \left(\beta_j m + \sqrt{1 - \beta_j^2} \epsilon_j \right) \right] \\ &= \beta_i \beta_j \quad i, j = 1, 2, \dots \end{aligned}$$

Just as in the single-credit version of the model, firm i defaults if $a_i \leq k_i$, the logarithmic distance to the default asset value, measured in standard deviations.

Example 8.3 (Correlation and Beta in Credit Single-Factor Model) Suppose firm 1 is “cyclical” and has $\beta_1 = 0.5$, while firm 2 is “defensive” and has $\beta_2 = 0.1$. The asset return correlation of the two firms is then $\beta_1\beta_2 = 0.5 \times 0.1 = 0.05$.

The single-factor model has a feature that makes it an especially handy way to estimate portfolio credit risk: conditional independence, the property that once a particular value of the market factor is realized, the asset returns—and hence default risks—are independent of one another. Conditional independence is a result of the model assumption that the firms’ returns are correlated only via their relationship to the market factor.

To see this, let m take on a particular value \bar{m} . The distance to default—the asset return—increases or decreases, and now has only one random driver ϵ_i , the idiosyncratic shock:

$$a_i - \beta_i \bar{m} = \sqrt{1 - \beta_i^2} \epsilon_i \quad i = 1, 2, \dots$$

The mean of the default distribution shifts for any $\beta_i > 0$ when the market factor takes on a specific value. The variance of the default distribution is reduced from 1 to $\sqrt{1 - \beta_i^2}$, even though the default threshold k_i has not changed. The change in the distribution that results from conditioning is illustrated in Figure 8.3.

Example 8.4 (Default Probability and Default Threshold) Suppose a firm has $\beta_i = 0.4$ and $k_i = -2.33$, it is a middling credit, but cyclical (relatively high β_i). Its unconditional probability of default is $\Phi(-2.33) = 0.01$. If we enter a modest economic downturn, with $\bar{m} = -1.0$, the conditional asset return distribution is $N(-0.4, \sqrt{1 - 0.40^2})$ or $N(-0.4, 0.9165)$, and the conditional default probability is found by computing the probability that this distribution takes on the value -2.33 . That probability is 1.78 percent.

If we were in a stable economy with $m = 0$, we would need a shock of -2.33 standard deviations for the firm to die. But with the firm’s return already 0.4 in the hole because of an economy-wide recession, it takes only a 1.93 standard deviation *additional* shock to kill it.

Now suppose we have a more severe economic downturn, with $\bar{m} = -2.33$. The firm’s conditional asset return distribution is $N(-0.932, 0.9165)$

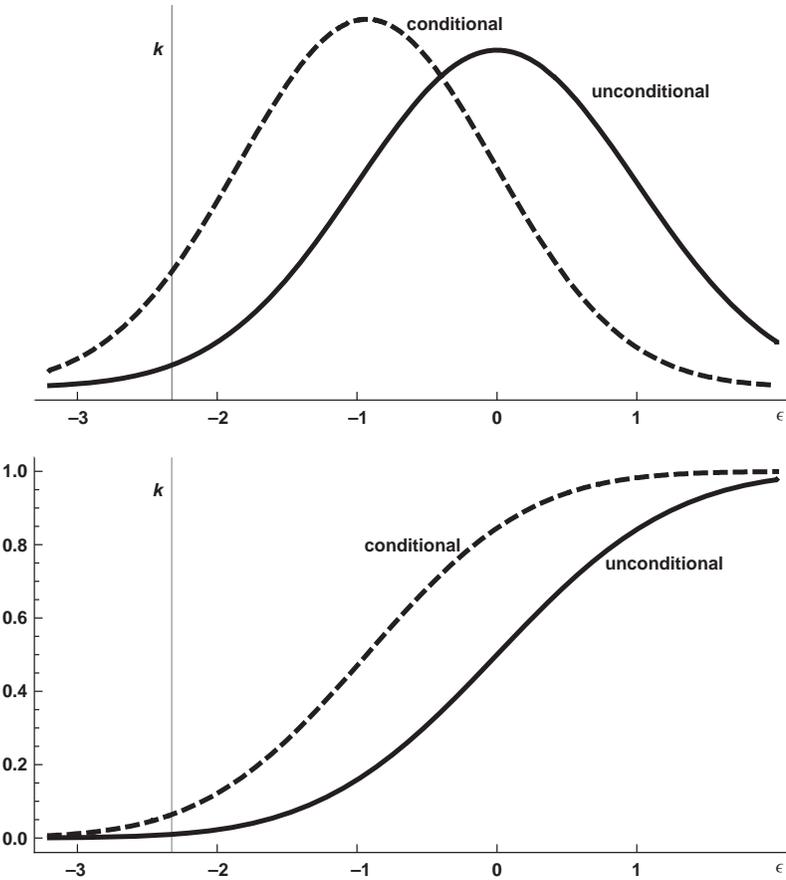


FIGURE 8.3 Default Probabilities in the Single-Factor Model
 The graph assumes $\beta_i = 0.4$, $k_i = -2.33 (\Leftrightarrow \pi_i = 0.01)$, and $\bar{m} = -2.33$. The unconditional default distribution is a standard normal distribution, while the conditional default distribution is $N(\beta_i \bar{m}, \sqrt{1 - \beta_i^2}) = N(-0.4, 0.9165)$.

Upper panel: Unconditional and conditional probability density of default. Note that the mean as well as the volatility of the conditional distribution are lower.

Lower panel: Unconditional and conditional cumulative default distribution function.

and the conditional default probability is 6.4 percent. A 0.93 standard deviation shock ($\epsilon_i \leq -0.93$) will now trigger default.

To summarize, specifying a realization $m = \bar{m}$ does three things:

1. The conditional probability of default is greater or smaller than the unconditional probability of default, unless either $\bar{m} = 0$ or $\beta_i = 0$, that

is, either the market factor shock happens to be zero, or the firm's returns are independent of the state of the economy.

There is also no longer an infinite number of combinations of market and idiosyncratic shocks that would trigger a firm i default. Given \bar{m} , a realization of ϵ_i less than or equal to

$$k_i - \beta_i \bar{m} \quad i = 1, 2, \dots$$

triggers default. This expression is linear and downward sloping in \bar{m} : As we let \bar{m} vary from high (strong economy) to low (weak economy) values, a smaller (less negative) idiosyncratic shock will suffice to trigger default.

2. The conditional variance of the default distribution is $1 - \beta_i^2$, so the conditional variance is reduced from the unconditional variance of 1.
3. It makes the asset returns of different firms *independent*. The ϵ_i are independent, so the conditional returns $\sqrt{1 - \beta_i^2} \epsilon_i$ and $\sqrt{1 - \beta_j^2} \epsilon_j$ and thus the default outcomes for two different firms i and j are independent.

Putting this all together, while the unconditional default distribution is a standard normal, the conditional distribution can be represented as a normal with a mean of $-\beta_i \bar{m}$ and a standard deviation of $\sqrt{1 - \beta_i^2}$.

The *conditional cumulative default probability function* can now be represented as a function of m :

$$p(m) = \Phi \left(\frac{k_i - \beta_i m}{\sqrt{1 - \beta_i^2}} \right) \quad i = 1, 2, \dots$$

It is plotted in the lower panel of Figure 8.4 for different correlations. This function is the standard normal distribution function of a random variable that has been standardized in a specific way. The mean, or “number of standard deviations,” is set to the new distance to default, given the realization of the market factor, while the standard deviation itself is set to its value $\sqrt{1 - \beta_i^2}$ under conditional independence. The intuition is that, for a given value of the market factor, the probability of default depends on how many standard deviations below its mean of 0 is the realization of ϵ_i . The density function corresponding to the cumulative default function is plotted in the upper panel of Figure 8.4.

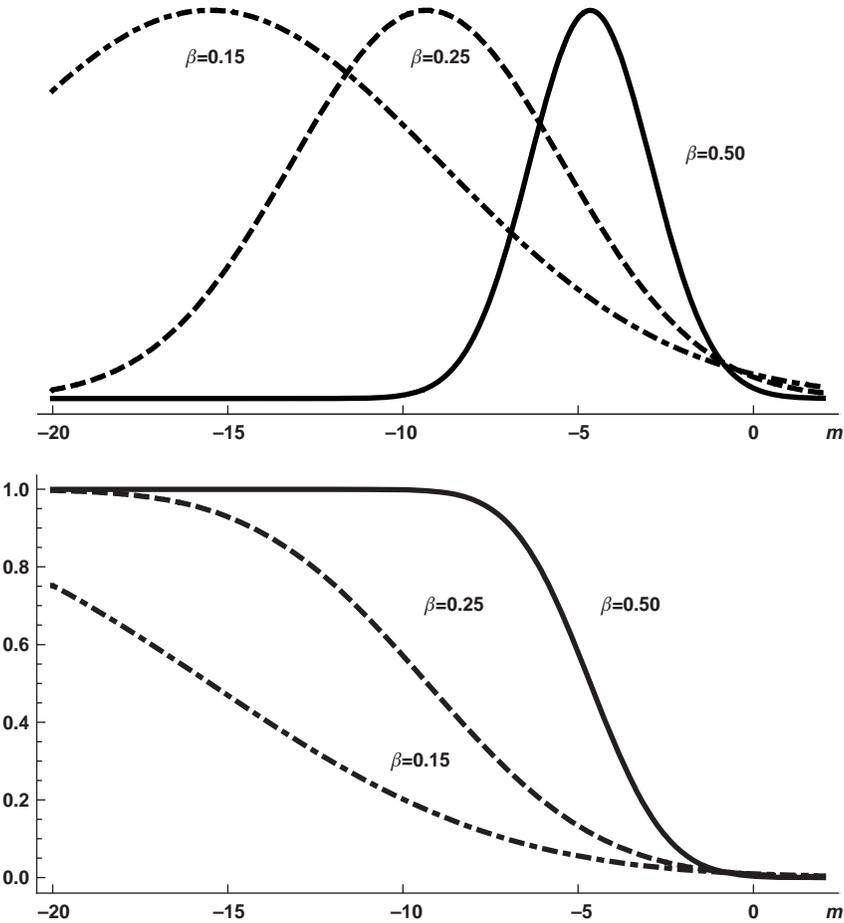


FIGURE 8.4 Single-Factor Default Probability Distribution
 Probability of default of a single obligor, conditional on the realization of m (x -axis). The default probability is set to 1 percent ($k = -2.33$), and the default correlation is set to different values as specified by the plot labels.
Upper panel: Conditional default density function, that is, the density function corresponding to $p(m)$.
Lower panel: Conditional cumulative distribution function of default $p(m)$.

8.3.2 Asset and Default Correlation

We began earlier to discuss the difference between the asset return and the default correlation. Let’s look for a moment at the relationship between the two.

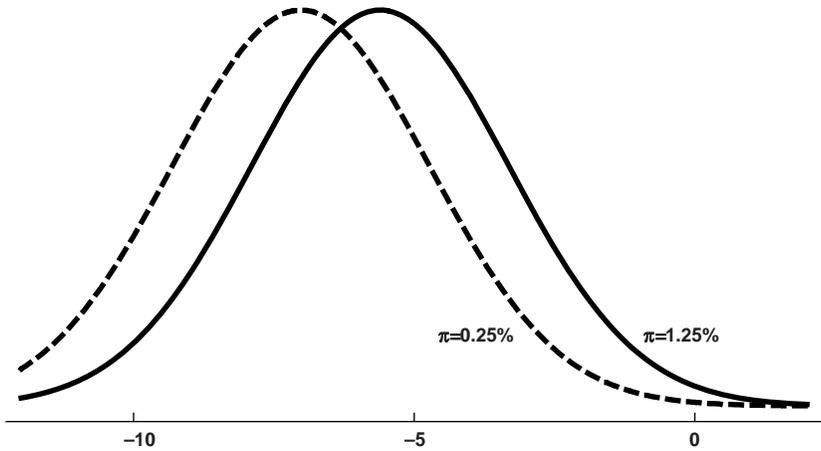


FIGURE 8.5 Conditional Default Density Function in the Single-Factor Model
 Both plots take $\beta = 0.40$. For a given correlation, the probability of default changes the location of the default distribution, but not its variance.

In the single-factor model, the cumulative return distribution of any pair of credits i and j is a bivariate standard normal with a correlation coefficient equal to $\beta_i\beta_j$:

$$\begin{pmatrix} a_i \\ a_j \end{pmatrix} \sim N \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \beta_i\beta_j \\ \beta_i\beta_j & 1 \end{pmatrix} \right]$$

Its cumulative distribution function is $\Phi_{(a_i)}$. The probability of a joint default is then equal to the probability that the realized value is in the region $\{-\infty \leq a_i \leq k_i, -\infty \leq a_j \leq k_j\}$:

$$\Phi \begin{pmatrix} k_i \\ k_j \end{pmatrix} = \mathbf{P}[-\infty \leq a_i \leq k_i, -\infty \leq a_j \leq k_j]$$

To get the default correlation for this model, we substitute $\pi_{ij} = \Phi \begin{pmatrix} k_i \\ k_j \end{pmatrix}$ into Equation (8.1), the expression for the linear correlation:

$$\rho_{ij} = \frac{\Phi \begin{pmatrix} k_i \\ k_j \end{pmatrix} - \pi_i\pi_j}{\sqrt{\pi_i(1 - \pi_i)}\sqrt{\pi_j(1 - \pi_j)}}$$

From here on, let's assume that the parameters are the same for all firms; that is, $\beta_i = \beta$, $k_i = k$, and $\pi_i = \pi$, $i = 1, 2, \dots$. The pairwise asset return correlation for any two firms is then β^2 . The probability of a joint default for any two firms for this model is

$$\Phi\left(\frac{k}{k}\right) = \mathbf{P}[-\infty \leq a \leq k, -\infty \leq a \leq k]$$

and the default correlation between any pair of firms is

$$\rho = \frac{\Phi\left(\frac{k}{k}\right) - \pi^2}{\pi(1 - \pi)}$$

Example 8.5 (Default Correlation and Beta) What β corresponds to a “typical” low investment-grade default probability of 0.01 and a default correlation of 0.05? We need to use a numerical procedure to find the parameter β that solves

$$\rho = 0.05 = \frac{\Phi\left(\frac{k}{k}\right) - \pi^2}{\pi(1 - \pi)}$$

With $\pi = 0.01$, the results are $\beta = 0.561$, the asset correlation $\beta^2 = 0.315$, and a joint default probability of 0.0006, or 6 basis points. Similarly, starting with $\beta = 0.50$ ($\beta^2 = 0.25$), we find a joint default probability of 4.3 basis points and a default correlation of 0.034.

8.3.3 Credit VaR Using the Single-Factor Model

In this section, we show how to use the single-factor model to estimate the credit VaR of a “granular,” homogeneous portfolio. Let n represent the number of firms in the portfolio, and assume n is a large number. We will assume the loss given default is \$1 for each of the n firms. Each credit is only a small fraction of the portfolio and idiosyncratic risk is de minimis.

Conditional Default Probability and Loss Level Recall that, for a given realization of the market factor, the asset returns of the various credits are independent standard normals. That, in turn, means that we can apply the law of large numbers to the portfolio. For each level of the market factor, the *loss level* $x(m)$, that is, the fraction of the portfolio that defaults, converges

to the *conditional probability* that a single credit defaults, given for any credit by

$$p(m) = \Phi\left(\frac{k - \beta m}{\sqrt{1 - \beta^2}}\right) \quad (8.2)$$

So we have

$$\lim_{N \rightarrow \infty} x(m) = p(m) \quad \forall m \in \mathbb{R}$$

The intuition is that, if we know the realization of the market factor return, we know the level of losses realized. This in turn means that, given the model's two parameters, the default probability and correlation, portfolio returns are driven by the market factor.

Unconditional Default Probability and Loss Level We are ultimately interested in the unconditional, not the conditional, distribution of credit losses. The unconditional probability of a particular loss level is equal to the probability that the the market factor return that leads to that loss level is realized. The procedure for finding the unconditional distribution is thus:

1. Treat the loss level as a random variable X with realizations x . We don't simulate x , but rather work through the model analytically for each value of x between 0 (no loss) and 1 (total loss).
2. For each level of loss x , find the realization of the market factor at which, for a single credit, default has a probability equal to the stated loss level. The loss level and the market factor return are related by

$$x(m) = p(m) = \Phi\left(\frac{k - \beta m}{\sqrt{1 - \beta^2}}\right)$$

So we can solve for \bar{m} , the market factor return corresponding to a given loss level \bar{x} :

$$\Phi^{-1}(\bar{x}) = \frac{k - \beta \bar{m}}{\sqrt{1 - \beta^2}}$$

or

$$\bar{m} = \frac{k - \sqrt{1 - \beta^2} \Phi^{-1}(\bar{x})}{\beta}$$

3. The probability of the loss level is equal to the probability of this market factor return. But by assumption, the market factor is a standard normal:

$$P[X \leq \bar{x}] = \Phi(\bar{m}) = \Phi\left(\frac{k - \sqrt{1 - \beta^2}\Phi^{-1}(\bar{x})}{\beta}\right)$$

4. Repeat this procedure for each loss level to obtain the probability distribution of X .

Another way of describing this procedure is: Set a loss level/conditional default probability x and solve the conditional cumulative default probability function, Equation (8.2), for \bar{m} such that:

$$\bar{m} = \frac{k - \sqrt{1 - \beta^2}\Phi^{-1}(x)}{\beta}$$

The loss distribution function is thus

$$P[X \leq x] = \Phi\left(\frac{k - \sqrt{1 - \beta^2}\Phi^{-1}(x)}{\beta}\right)$$

Example 8.6 (Loss Level and Market Level) A loss of 0.01 or worse occurs when—converges to the event that—the argument of $p(m)$ is at or below the value such that $p(m) = 0.01$.

$$p(\bar{m}) = 0.01 = \Phi\left(\frac{k - \beta\bar{m}}{\sqrt{1 - \beta^2}}\right)$$

The value \bar{m} at which this occurs is found by solving

$$\Phi^{-1}(0.01) \approx -2.33 = p^{-1}(\bar{m}) = \frac{k - \beta\bar{m}}{\sqrt{1 - \beta^2}}$$

for \bar{m} . This is nothing more than solving for the \bar{m} that gives you a specific quantile of the standard normal distribution.

With a default probability $\pi = 0.01$ and correlation $\beta^2 = 0.50^2 = 0.25$, the solution is $\bar{m} = -0.6233$. The probability that the market factor ends up at -0.6233 or less is $\Phi(-0.6233) = 0.2665$.

As simple as the model is, we have several parameters to work with:

- The probability of default π sets the unconditional expected value of defaults in the portfolio.
- The correlation to the market β^2 determines how spread out the defaults are over the range of the market factor. When the correlation is high, then, for any probability of default, defaults mount rapidly as business conditions deteriorate. When the correlation is low, it takes an extremely bad economic scenario to push the probability of default high.

To understand the impact of the correlation parameter, start with the extreme cases:

- $\beta \rightarrow 1$ (perfect correlation). Recall that we have constructed a portfolio with no idiosyncratic risk. If the correlation to the market factor is close to unity, there are two possible outcomes. Either $m \leq k$, in which case nearly all the credits default, and the loss rate is equal to 1, or $m > k$, in which case almost none default, and the loss rate is equal to 0.
- $\beta \rightarrow 0$ (zero correlation). If there is no statistical relationship to the market factor, so idiosyncratic risk is nil, then the loss rate will very likely be very close to the default probability p .

In less extreme cases, a higher correlation leads to a higher probability of either very few or very many defaults, and a lower probability of intermediate outcomes. This can be seen in Figure 8.6 in the cumulative loss distribution and loss density functions, which converge to an L-shape. The loss density converges to a ray over the default probability as the correlation goes to zero, that is, the volatility goes to zero. Figure 8.7 compares loss densities for a given correlation and different default probabilities.

8.4 USING SIMULATION AND COPULAS TO ESTIMATE PORTFOLIO CREDIT RISK

The big problem with which portfolio credit risk models grapple is the likelihood of joint default. The single-factor model introduces a latent factor that drives joint default. Factor models make sense, because they link the probability of joint default with our intuition that default is driven by the state of the economy and perhaps an industry sector, that is common to all or to a group of credits, as well as to a company's unique situation.

An alternative approach is more agnostic about the fundamental forces driving defaults. It relies on simulations, and ties the simulations together

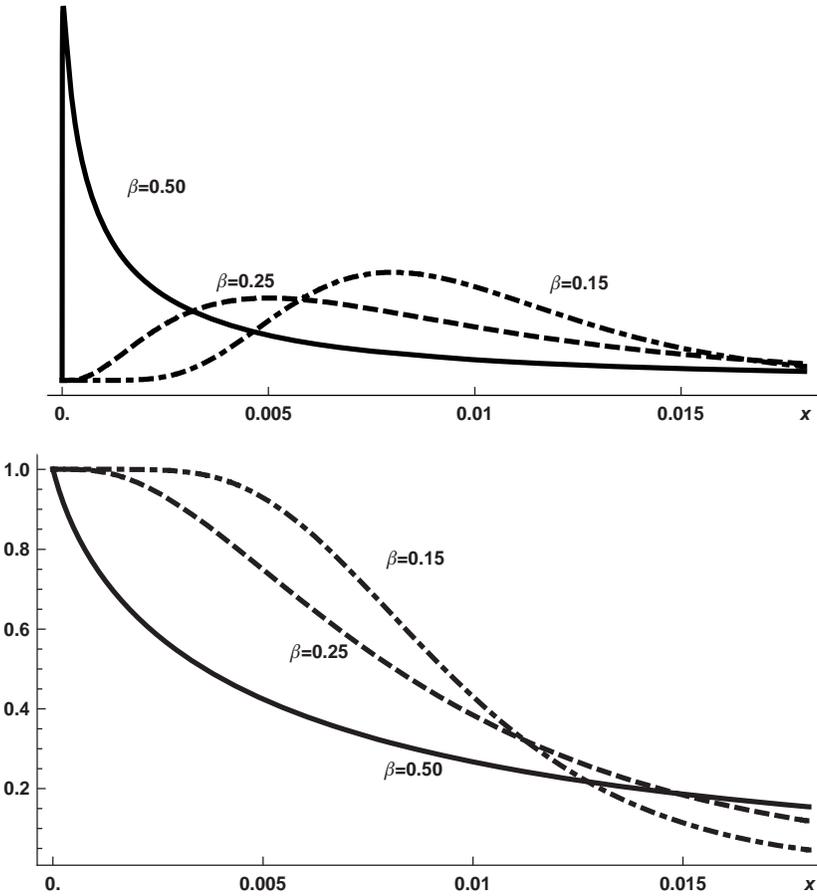


FIGURE 8.6 Distribution of Losses in the Single-Factor Model
 Homogeneous and completely diversified portfolio; default probability
 1.0 percent.

Upper panel: Density function of the loss rate. Note that the mean as well as the volatility of the conditional distribution are lower.

Lower panel: Cumulative distribution function of the loss rate.

using a “light” modeling structure. This approach uses a particular mathematical trick, the *copula*, to correlate defaults.

We assume in this section that we have a default time distribution for each of the credits in the portfolio, either by risk-neutral estimation from bond or CDS prices, or using ratings, or from a structural model. The rest

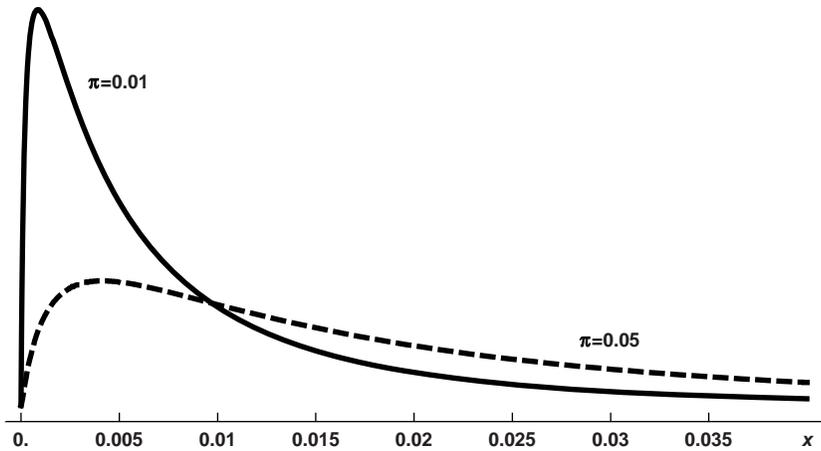


FIGURE 8.7 Density Function of Portfolio Losses in the Single-Factor Model
 $\beta = 0.40$.

of this chapter thus builds on Chapter 7, in which we defined the concept of a default time distribution and showed how to estimate it from market data. We will also assume that credit spreads, recovery rates, and risk-free interest rates are deterministic, and focus on default risk.

8.4.1 Simulating Single-Credit Risk

To explain how this approach works, we first describe how to estimate the default risk of a single credit via simulation. Simulation is not really necessary to estimate single-credit risk, and we are going to describe a needlessly complicated way to do so. We're taking this trouble in order to build up to portfolio credit risk estimation.

To see why simulation is not required, imagine we have a portfolio consisting of a bond issued by a single-B credit with an estimated one-year default probability of 0.05. The portfolio has a probability of 0.95 of being worth a known future market value and a probability of 0.05 of being worth only its recovery value in one year. If we assume that yield and credit curves are flat (so that we do not roll up a curve as we move the maturity date closer), the portfolio credit VaR at a confidence level of 0.95 or more is equal to the recovery value less the expected loss.

Let's begin by describing a simple technique for simulating default distributions for a single credit. We use the cumulative default time distribution, which we defined and derived from credit spread data in Chapter 7.

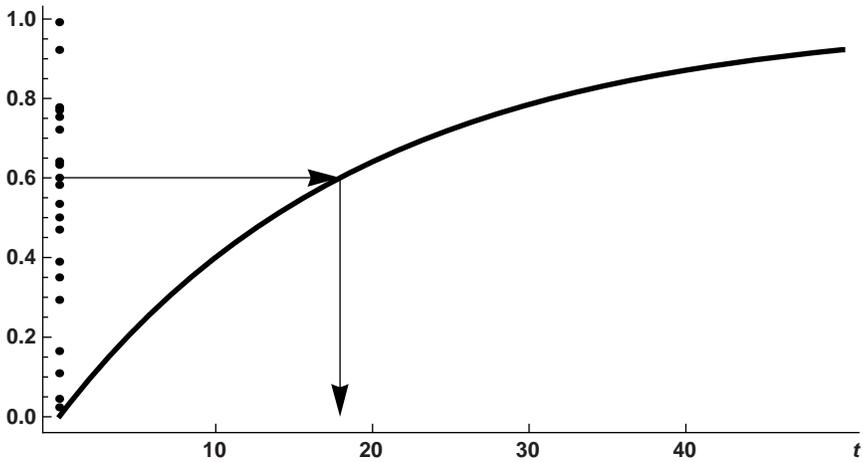


FIGURE 8.8 Estimated Single-Credit Default Risk by Simulation

The graph shows the cumulative default time distribution for a credit with a one-year default probability of 0.05. The hazard rate is 0.0513. The points represent 20 simulated values of the uniform distribution.

Figure 8.8 illustrates the procedure using a default time distribution similar to that illustrated in Figure 7.3. The points parallel to the y -axis represent simulated values of the uniform distribution. For each of the uniform simulations, we can find the default time with a probability equal to the uniform variate. For example, the arrows trace one of the simulation threads, with a value of about 0.6. The probability of a default within the next 18 years or so is about 0.6. So that simulation thread leads to the simulated result of default in about 18 years.

We can repeat this process for a large number of simulated uniform variates. We will find that very close to 5 percent of our simulations lead to default times of one year or less; we can get arbitrarily close to about 5 percent defaulting within one year by taking enough simulations. This does not add anything to what we already knew, since the one-year 5 percent default probability was our starting point.

Now we make the simulation procedure even more complicated by transforming the cumulative default time distribution function. Up to now, we've worked with a mapping from a future date or elapsed time to a probability. We now transform it into a mapping from a future date to a standard normal z value.

This procedure will seem to add needless complexity, but it is a crucial part of applying copula techniques to measuring portfolio credit risk. It enables us to use the joint normal distribution to simulate credit returns for

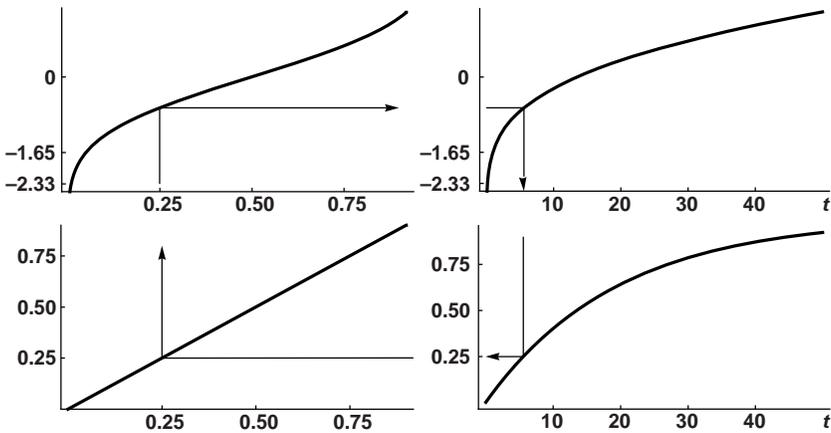


FIGURE 8.9 Shifting from Uniform to Normal Distribution Simulations
Starting in the lower right panel, the graph traces how to change one thread of a uniform simulation to a normal simulation. The lower right panel shows the default time distribution for a single-B credit with a one-year default probability of 5 percent.

portfolios containing more than one credit. The procedure is illustrated in Figure 8.9 for our single-B credit with a one-year default probability of 0.05.

We can move from simulations of the uniform distribution to simulations of the univariate normal distribution using the *transformation principle*, explained in Appendix A.5:

- We start with an arbitrary default time, say, 5.61 years. The probability of a default over the next 5.61 years is 0.25, as illustrated in the lower right panel of Figure 8.9.
- The second step is to find, in the upper left panel, the corresponding standard normal distribution value, -0.675 . That is, $\Phi(-0.675) = 0.25$.
- The final step closes the loop by mapping the normal distribution value to the default time with a probability of 0.25. This step is illustrated in the two right panels of Figure 8.9.
- The upper right panel displays the result, a function mapping from default times to standard normal values.

8.4.2 Simulating Joint Defaults with a Copula

Next, we simulate joint defaults. To do so, we need a multivariate default time distribution. But the default time distributions that we have are univariate, covering one issuer at a time. We don't have a statistical theory that tells us how these distributions might be associated with one another.

We could model default times as multivariate normal, a distribution with which we are very familiar. The problem with that is that the marginal distributions would then also be normal, and as different as can be from the default time distributions we have been using. We would, for example, have default times that are potentially negative. The problem would not be solved by using an alternative to the normal, since the alternative distribution, too, would have marginal distributions that bear no resemblance to our default time distributions.

To summarize, the problem is this: We have a set of univariate default time distributions that we believe in, or are at least willing to stipulate are realistic. But we do not know how to connect them with one another to derive the behavior of portfolios of credits. On the other hand, we are familiar with a few families of multivariate distributions, but none of them have marginal distributions we feel are appropriate for the analysis of default risk.

A now-standard solution to problems such as this is the use of copulas. The great benefit of a copula for our purposes is that it permits us to separate the issue of the default-time distribution of a single credit from the issue of the dependence of default times for a portfolio of credits, that is, their propensity or lack of propensity to default at the same time.

We can then combine the default-time distributions we believe in with a distribution that makes talking about joint events easier, namely, the multivariate normal distribution. One of the reasons the copula approach has become so popular is that it exploits the familiarity of multivariate normal distributions and the ease of simulation using them, without having to accept a multivariate normal model of defaults.

We will spend a little time sketching the theory that justifies this approach, and then illustrate by continuing the example of the previous section. Mathematically, a copula has these properties:

- It is a function $c : \{[0, 1]^n\} \mapsto [0, 1]$. That is, it maps from the Cartesian product of $[0, 1]$, repeated n times, to $[0, 1]$.
- It is therefore an n -dimensional distribution function; that is, it takes as its argument n uniformly $[0, 1]$ distributed random variables and returns a probability.
- The marginal distributions are all uniformly distributed on $[0, 1]$.

Formally, our problem is this: Suppose we have a portfolio with securities of n issuers. We have estimated or specified single-issuer default-time distributions $F_1(t_1), \dots, F_n(t_n)$. We do not know the joint distribution $F(t_1, \dots, t_n)$. We can, however, somewhat arbitrarily specify a copula

function $c(F(t_1, \dots, t_n))$, stipulating that

$$c(F_1(t_1), \dots, F_n(t_n)) = F(t_1, \dots, t_n)$$

Since $c(F_1(t_1), \dots, F_n(t_n))$ is a copula, its marginal distributions are the single-issuer default-time distributions $F_1(t_1), \dots, F_n(t_n)$. For any multivariate distribution, we can always find a copula.

Let's reduce the level of generality and consider a portfolio consisting of two credits, one single-B and one CCC-rated, each with a known/stipulated hazard rate λ_B or λ_{CCC} . The default-time distribution functions are

$$\begin{aligned} F(t_B) &= 1 - e^{-\lambda_B t_B} & t_B &\in [0, \infty) \\ F(t_{CCC}) &= 1 - e^{-\lambda_{CCC} t_{CCC}} & t_{CCC} &\in [0, \infty) \end{aligned}$$

and we can use them to define corresponding uniform-[0, 1] random variates

$$\begin{aligned} u_B &= F(t_B) \\ u_{CCC} &= F(t_{CCC}) \end{aligned}$$

as well as corresponding quantile functions

$$\begin{aligned} t_B &= F^{-1}(u_B) = -\frac{1}{\lambda_B} \log(1 - u_B) & u_B &\in [0, 1] \\ t_{CCC} &= F^{-1}(u_{CCC}) = -\frac{1}{\lambda_{CCC}} \log(1 - u_{CCC}) & u_{CCC} &\in [0, 1] \end{aligned}$$

The transformation principle (see Appendix A.5) tells us that u_B and u_{CCC} , which are ranges of distribution functions, are uniform-[0, 1]. In the last section, we saw how to move back and forth between distribution and quantile functions. We can do this in a multivariate context, too. We do not know the joint default-time distribution function $F(t_B, t_{CCC})$. But by virtue of being a distribution function,

$$\begin{aligned} F(t_B, t_{CCC}) &= \mathbf{P}[\tilde{t}_B \leq t_B \wedge \tilde{t}_{CCC} \leq t_{CCC}] \\ &= \mathbf{P}[F^{-1}(\tilde{u}_B) \leq t_B \wedge F^{-1}(\tilde{u}_{CCC}) \leq t_{CCC}] \\ &= \mathbf{P}[\tilde{u}_B \leq F(t_B) \wedge \tilde{u}_{CCC} \leq F(t_{CCC})] \\ &= \mathbf{P}[\tilde{u}_B \leq u_B \wedge \tilde{u}_{CCC} \leq u_{CCC}] \\ &= c(u_B, u_{CCC}) \end{aligned}$$

The tildes identify the symbols representing the random times and their probabilities. The first line follows from the definition of a distribution function, while the last line follows from the copula theorem, known as *Sklar's theorem*, which tells us *some* copula must exist.

So far, all we have done is define a type of mathematical object called a copula and seen how it can be related to the “known” (or at least stipulated) single-issuer default-time distributions $F_1(t_1), \dots, F_n(t_n)$ and the “unknown” joint distribution $F(t_1, \dots, t_n)$. How do we compute a credit VaR? There are four steps:

1. Specify the copula function that we’ll use.
2. Simulate the default times.
3. Apply the default times to the portfolio to get the market values and P&Ls in each scenario.
4. Add results to get portfolio distribution statistics.

So next we need to actually specify a copula. The most common type or family of copulas is the *normal copula*. The user provides the $F_1(t_1), \dots, F_n(t_n)$ and an estimate of a multivariate normal correlation matrix Σ .

In our bivariate example, the normal copula is

$$\begin{aligned} c(u_B, u_{CCC}) &= \Phi(\Phi^{-1}[F(t_B)], \Phi^{-1}[F(t_{CCC})]; 0, \Sigma) \\ &= \Phi(\Phi^{-1}(u_B), \Phi^{-1}(u_{CCC}); 0, \Sigma) \end{aligned}$$

with

$$\Sigma = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$$

Once we have chosen this copula function and have set the parameter ρ , we can use it to simulate joint defaults. Figure 8.10 illustrates the procedure for our two-credit portfolio.

The starting point is a simulation of the joint standard normal distribution $N(0, \Sigma)$, as seen in the lower left panel of Figure 8.10. Each of these pairs can be mapped to the standard normal quantile function to get the pair

$$(\Phi^{-1}(z_B^{(i)}), \Phi^{-1}(z_{CCC}^{(i)})) \quad i = 1, \dots, I$$

with I the number of simulations. Each of the latter is a pair in $[0, 1]^2$.

Next, map the first element of each of these pairs to the single-B default time that has the probability $\Phi^{-1}(z_B^{(j)})$ and the second element of each of these pairs to the CCC default time that has the probability $\Phi^{-1}(z_{CCC}^{(j)})$.

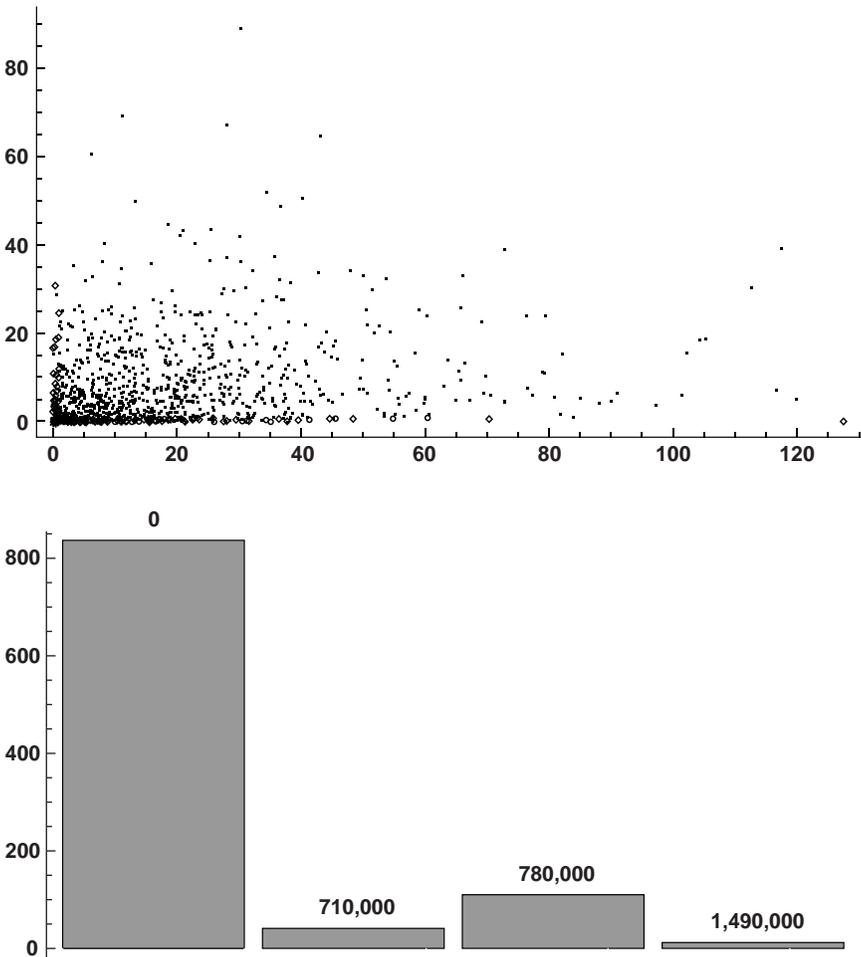


FIGURE 8.10 Distribution of Losses in the Single-Factor Model
 Estimate of credit VaR for a portfolio consisting of two credit-risky securities based on 1,000 simulations.

Upper panel: Simulated default times. The simulation trials are partitioned and marked as follows:

- * trials leading to losses in the 0.01 quantile
- trials leading to losses in the 0.05 quantile, but smaller than the 0.01 quantile
- ◇ trials leading to losses in the 1-s.d. quantile, but smaller than the 0.05 quantile
- trials leading to losses smaller than the 1-s.d. quantile

Lower panel: Histogram of loss levels. Each bar is labelled by the portfolio loss realized in the number of simulation trials indicated on the y-axis.

This step is illustrated by the arrows drawn from the lower left panel to the upper left and lower right panels.

Each of these pairs is then plotted to the upper-right panel of Figure 8.10. We now have our correlated joint default-time simulation.

Two features of this procedure seem arbitrary. First, we chose a normal copula; there are alternatives. Second, how do we estimate or otherwise assign the correlation parameter ρ , or in a more general context, the correlation matrix? One answer is provided by the prices of credit derivatives written on credit indexes, which we will study in the next chapter.

We next give a detailed example of the procedure for a portfolio consisting of just two speculative-grade credits, each with a current notional and market value of \$1,000,000. As we saw above, for a two-credit portfolio, if we have the default probabilities of each credit and their default correlation, we can determine the entire credit distribution of the portfolio, so we are carrying out this example for illustrative purposes. We imagine the credits to have single-B and CCC ratings and assume:

| | p | λ | Notional | Coupon | Spread |
|----------|------|-----------|-----------|--------|--------|
| CCC | 0.10 | 0.1054 | 1,000,000 | 0.18 | 0.13 |
| Single-B | 0.05 | 0.0513 | 1,000,000 | 0.11 | 0.06 |

We assume a recovery rate of 40 percent. The horizon of the credit VaR is one year.

There are four possible outcomes over the next year: no default, only the single-B loan defaults, only the CCC loan defaults, and both default. To keep things simple, we ignore the distinction between expected and unexpected credit loss, assuming, in effect, that the lender does not set aside any provision for credit losses. If there is a default, we assume the coupon is not paid. The credit losses then consist of forgone principal and coupon, mitigated by a recovery amount paid one year hence. The losses for each of the four scenarios are:

| | Default time realization | Terminal value | Loss |
|------------------|--|----------------|-----------|
| No default | $(\tau_{B,i} > 1, \tau_{CCC,i} > 1)$ | 2,290,000 | 0 |
| Single-B default | $(\tau_{B,i} \leq 1, \tau_{CCC,i})$ | 1,580,000 | 710,000 |
| CCC default | $(\tau_{B,i} > 1, \tau_{CCC,i} \leq 1)$ | 1,510,000 | 780,000 |
| Both default | $(\tau_{B,i} \leq 1, \tau_{CCC,i} \leq 1)$ | 800,000 | 1,490,000 |

To estimate the VaR, we first simulate correlated default times using the normal copula. We apply a correlation of 0.25. Following the procedure outlined above, the results are

- We first generate 1,000 realizations of the bivariate standard normal distribution using a correlation coefficient $\rho = 0.25$, giving us 1,000 pairs of real numbers.
- Each of these 2,000 real numbers is mapped to its standard univariate normal quantile, giving us 1,000 pairs of numbers in $(0, 1)$.
- The first element of each pair is mapped to the single-B default time with that probability. The second element of each pair is mapped to the CCC default time with that probability. We now have 1,000 pairs of simulated default times $(\tau_{B,i}, \tau_{CCC,i})$. These are illustrated in the upper right panel of Figure 8.10.
- Each default time is either greater than, less than, or equal to the one-year horizon of the credit VaR. We can accordingly assign a terminal value to each loan in each simulation, sum across the two loans, and subtract the sum from the no-default future value to get the loss. There are four distinct possible values for each simulation trial. In the upper panel of Figure 8.11, each trial is marked by its pair of default times.
- Finally, we tally up the number of simulation trials resulting in each loss level. This is displayed as a histogram in the lower panel of Figure 8.10.

The credit VaR estimates (in dollars) are:

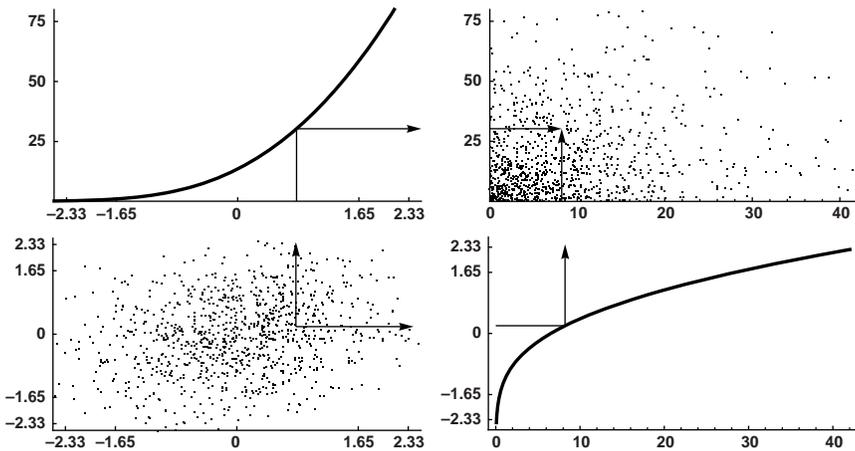


FIGURE 8.11 Simulating Multiple Defaults

Starting in the lower right panel, the graph traces how to change one thread of a uniform simulation to a normal simulation. The lower right panel shows the default time distribution for a single-B credit with a one-year default probability of 5 percent.

| Confidence level | VaR |
|-------------------------|-----------|
| 99 percent | 1,490,000 |
| 95 percent | 780,000 |
| 1 s.d. (84.135 percent) | 710,000 |

Copulas are a very attractive modeling technique, since they permit the model to generate quite detailed results—the entire probability distribution of portfolio credit outcomes—with a very light theoretical apparatus and requiring the estimation of only one additional parameter, the correlation, beyond those used in single-credit modeling. However, the copula approach also has a number of pitfalls. Most important among these is that the choice of copula is arbitrary and that we simply do not know enough to reliably estimate the copula correlation. It is difficult enough to estimate default correlations, and the copula correlation is only related to, not identical to it. Yet once a parameter value is assigned, the temptation to rely on the wide range of model results that can then be generated is enormous. Reliance on a poorly understood and hard-to-estimate parameter in a simplified model is dangerous. This particular example was important in the subprime crisis. We explore model risk further in Chapter 11.

Copula techniques are widely used in the valuation and risk management of credit portfolios. The most frequent application, however, is in modeling portfolio credit products, such as securitizations and credit index products. We describe these in more detail in Chapter 9. Variants of the model let default intensity vary over time stochastically, or correlate the intensities of different firms with one another.

The models presented in this chapter focus on default, but ratings migration, of course, is also an important driver of credit risk. The Gaussian single-factor approach can be applied in migration mode as well as default mode. In addition to a default probability and a corresponding default threshold, the model requires a set of migration or transition probabilities, for example those contained in the transition matrices described in Chapter 6.

FURTHER READING

Lucas (1995) provides a definition of default correlation and an overview of its role in credit models. See also Hull and White (2001).

Credit Suisse First Boston (2004) and Lehman Brothers (2003) are introductions by practitioners. Zhou (2001) presents an approach to modeling

correlated defaults based on the Merton firm value, rather than the factor-model approach.

The application of the single-factor model to credit portfolios is laid out in Finger (1999) and Vasicek (1991). Accessible introduction to copula theory are Frees and Valdez (1998) and in Klugman, Panjer, and Willmot (2008). The application to credit portfolio models and the equivalence to Gaussian CreditMetrics is presented in Li (2000).

The correlated intensities approach to modeling credit portfolio risk, as well as other alternatives to the Gaussian single-factor approach presented here, are described in Schönbucher (2003), Chapter 10, and Lando (2004), Chapter 5.