New Chapter 26: Credit Risk

Credit risk is the risk that a counterparty will fail to meet a contractual payment obligation. Such risk exists any time one party promises to make a future payment to another. Credit risk arises with loans, corporate bonds, and derivative contracts, and market-making activities generally leave dealers exposed to credit risk. We will refer to the failure to make a promised payment as *default*. The term **credit event** is often used in contracts to refer to occurrences that suggest that a default is likely. Examples of credit events are declaration of bankruptcy, failure to make a bond payment, repudiation of an obligation, or a credit downgrade.[51]

In this chapter we present a framework for analyzing credit risk and we

[51]A company that fails to make a promised payment may seek court protection—or a quick resolution of its dilemma—by declaring bankruptcy. Specific bankruptcy rules vary by country. We will use the terms "default" and "bankruptcy" interchangably and without regard to their precise legal meaning.

see how credit risk affects the prices of claims. We will discuss concepts and terminology that are essential for understanding credit risk. Finally, we look at the instruments used to modify, hedge, and trade credit risk, such as credit default swaps.

26.1 DEFAULT CONCEPTS AND TERMINOLOGY

In this section we introduce basic concepts and terminology related to default in the context of pricing a zero-coupon bond. Suppose that a firm with asset value A_0 issues a zero-coupon bond maturing at time T, with a promised payment of $\bar{B}$. Let B_T denote the market value of the bond at time T. At time T, there are two possible outcomes:

- $A_T > \bar{B}$. Since assets are worth more than the repayment owed bondholders, shareholders will repay bondholders in full, so $B_T = \bar{B}$.
- $A_T < \bar{B}$. Shareholders will walk away from the firm, surrendering it to bondholders. The value of the bonds at time T is then $B_T = A_T$.

Let $g^*(A_T; A_0)$ denote the risk-neutral probability density for the time T asset value, conditional upon the time 0 asset value, A_0. Then we can write the initial debt value, B_0, as

$$B_0 = e^{-rT}\left[\int_0^{\bar{B}} A_T g^*(A_T; A_0) dA_T + \bar{B}\int_{\bar{B}}^{\infty} g^*(A_T; A_0) dA_T\right] \qquad (26.109)$$

The first term in the brackets on the right-hand side is the partial expectation of the asset value. In the second term, $\int_{\bar{B}}^{\infty} g^* dA_T$ is the risk-neutral probability that the firm fully repays its debt, i.e., is solvent. Thus, we can rewrite the value of the bonds as

$$B_0 = e^{-rT}\left\{\bar{B} \times [1 - \text{Prob}^*(\text{Default})] + E^*\left(A_T|\text{Default}\right) \times \text{Prob}^*(\text{Default})\right\}$$

where E^* and Prob* are computed with respect to the risk-neutral measure. Since $B_T = A_T$ in default, we can also write this as

$$B_0 = e^{-rT}\left\{\bar{B} \times [1 - \text{Prob}^*(\text{Default})] + E^*\left(B_T|\text{Default}\right) \times \text{Prob}^*(\text{Default})\right\} \tag{26.110}$$

If we set the probability of default equal to zero, equation (26.110) yields the standard formula for the value of a default-free bond, $B_0 = e^{-rT}\bar{B}$. Equation (26.110) illustrates that default introduces two new elements: the default probability (Prob[Default]), and the payoff conditional on default ($E^*\left[B_T|\text{Default}\right]$).

The payoff conditional on default can be expressed in different ways. The **recovery rate** is the amount the debt-holders receive as a fraction of what they are owed. Thus, in the case where a firm issues a single zero-coupon bond, the expected recovery rate is

$$\text{Expected recovery rate} = \frac{E\left(B_T|\text{Default}\right)}{\bar{B}} \tag{26.111}$$

The **loss given default** is the difference between what the bond holders are owed and what they receive, as a fraction of the promised payment:

$$\text{Expected loss given default} = 1 - \text{Expected recovery rate} \tag{26.112}$$

Conventionally any such measure is expressed as a percentage.

Finally, we can express the **credit spread**—the difference between the yield to maturity on a defaultable bond and an otherwise equivalent default-free bond—in terms of the default probability and the loss given default. In equation (26.110), divide both sides by $\bar{B}$ and take the natural logarithm of both sides. Recall that the annual yield to maturity on the bond, ρ, is

$$\rho = \frac{1}{T} \ln \left(\frac{\bar{B}}{B_0} \right)$$

After some rearrangement, we obtain the following expression from equation (26.110):

$$\rho - r = \frac{1}{T} \ln \left[\frac{1}{1 - \text{Prob(Default)} \times \text{Expected loss given default}} \right] \tag{26.113}$$

The left-hand side in this expression is the credit spread. The expression in square brackets is greater than or equal to one (and hence the logarithm is greater than or equal to zero) because both the probability of default and the expected loss given default are less than one. As we would expect, if either the probability of default or the expected loss given default is zero, the bond yield equals the risk-free rate.

26.2 THE MERTON DEFAULT MODEL

In Section 16.1 we analyzed corporate securities as options. We saw that owning zero-coupon debt subject to default is the same thing as owning a default-free bond and writing a put option on the assets of the firm. This is an example of a *structural* approach to modeling bankruptcy: We create an explicit model for the evolution of the firm's assets, coupled with a rule governing default.

If we assume that the assets of the firm are lognormally distributed, then we can use the lognormal probability calculations of Chapter 18 to compute either the risk-neutral or actual probability that the firm will go bankrupt. This approach to bankruptcy modeling has come to be called the *Merton model* since Merton (1974) used continuous-time methods to provide a model of the credit spread. The Merton default model has in recent years been the basis for credit risk analyses provided by Moody's KMV.

Default at Maturity

Assume that the assets of the firm, A, follow the process

$$\frac{dA}{A} = (\alpha - \delta)dt + \sigma dZ \tag{26.114}$$

where α is the expected return on the firm assets and δ is the cash payout made to claim holders on the firm. Suppose the firm has issued a single zero-coupon bond that matures at time T and makes no payouts. If the promised payment on the bond is $\bar{B}$, then the probability of bankruptcy at time T, conditional on the value of assets at time t, is

$$\text{Prob}(A_T < \bar{B}; A_t) = N\left[-\frac{\ln(A_t/\bar{B}) + (\alpha - \delta - \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}\right] \tag{26.115}$$

This is essentially the Black-Scholes $N(-d_2)$ term, except that it contains α rather than the risk-free rate, r.

When assets are lognormally distributed, the **distance to default** is another way to express the probability of default. Default occurs at time T if $A_T < \bar{B}$. The expected log asset value at time T is

$$E[\ln(A_T)] = \ln(A_t) + (\alpha - \delta - 0.5\sigma^2)(T-t)$$

The distance between the expected asset value and the bankruptcy level $\bar{B}$, normalized by the standard deviation, is

$$\frac{E[\ln(A_T)] - \bar{B}}{\sigma\sqrt{T-t}} = \frac{\ln(A_t) + (\alpha - \delta - 0.5\sigma^2)(T-t) - \ln(\bar{B})}{\sigma\sqrt{T-t}}$$

Except for having α in place of r, this is the Black-Scholes d_2: the difference between current asset value the bankruptcy level, normalized by the standard deviation. This measures the size (in standard deviations) of the random shock required to induce bankruptcy. The default probability is $N(-\text{distance to default})$.

The expected recovery rate, conditional on default, is

$$E(A_T|A_T < \bar{B}) = A_t e^{(\alpha-\delta)(T-t)} \frac{N\left[-\frac{\ln(A_t/\bar{B})+(\alpha-\delta+\frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}\right]}{N\left[-\frac{\ln(A_t/\bar{B})+(\alpha-\delta-\frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}\right]} \tag{26.116}$$

This is the same as equation (18.28).

It is important to notice that the calculations in equations (26.115) and (26.116) are performed under the true probability measure (also sometimes called the *physical measure*). Thus, these equations provide estimates of the empirically observed default probability and recovery rate, but we cannot use them in pricing calculations. In order to compute the theoretical credit spread, for example, we must replace the actual asset drift, α, with the risk-free rate. This gives us the risk-neutral counterparts of equations (26.115) and (26.116), which we can use for pricing:

$$\text{Prob}^*(A_T < \bar{B}; A_t) = N\left[-\frac{\ln(A_t/\bar{B}) + (r-\delta-\frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}\right] \tag{26.117}$$

and

$$E^*(A_T|A_T < \bar{B}) = A_t e^{(r-\delta)(T-t)} \frac{N\left[-\frac{\ln(A_t/\bar{B})+(r-\delta+\frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}\right]}{N\left[-\frac{\ln(A_t/\bar{B})+(r-\delta-\frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}\right]} \tag{26.118}$$

The following example illustrates these calculations.

Example 26.1 Suppose that $\overline{B} = \$100$, $A_0 = \$90$, $\alpha = 10\%$, $r = 6\%$, $\sigma = 25\%$, $\delta = 0$ (the firm makes no payouts), and $T = 5$ years. As we saw in Example 16.1, which used the same assumptions, the theoretical debt value in this case is

\$62.928, which implies a yield of 9.2635%.

Using equations (26.115) and (26.117), the true and risk-neutral default probabilities are 33.49% and 47.26%. Thus, over a five-year horizon, we would expect to observe a default one-third of the time. Under the risk-neutral measure, however, defaults occur almost half the time. The greater risk-neutral default probability is due to the assets growing more slowly under the risk-neutral measure.

Using equations (26.116) and (26.118), the expected asset value conditional on default is \$71.867 under the true measure, and \$68.144 under the risk-neutral measure. Expected recovery rates are therefore

$$E(\text{Recovery rate}) = \frac{71.867}{100} = 0.71867$$

under the true measure, and

$$E^*(\text{Recovery rate}) = \frac{68.144}{100} = 0.68144$$

under the risk-neutral measure. Note that the risk-neutral expected loss given default is

$$E^*(\text{Loss given default}) = 1 - 0.68144 = 0.31866$$

Using the risk-neutral default probability and loss given default, we can compute the theoretical debt yield. From equation (26.113), the credit spread is

$$\frac{1}{5}\ln\left[\frac{1}{1-0.4726\times 0.31866}\right]=0.032635$$

This implies a debt yield to maturity of $0.060+0.032635 = 0.092635$, which is the same answer as that using the Black-Scholes formula to compute the theoretical debt value. ■

In practice, it is common to use data to estimate the parameters necessary to perform default calculations. As the preceding example shows, historical data on *defaults* provides different information than historical data on *prices*. The historical default frequency and recovery rate for example, correspond to equations (26.115) and (26.116), which are calculated under the true measure. If we examine credit spreads, contrast, we can infer an expected default frequency and recovery rate, however, these correspond to equations (26.117) and (26.118), which are calculated under the risk-neutral measure. Notice, however, that we would infer the same asset volatility from both sets of calculations.

Related Models

Suppose that the value of assets can jump to zero according to a Poisson process. Specifically, suppose that over an interval dt, the probability of a jump to zero is λdt, and that the occurrence of this jump is independent of the market and of other defaults. We saw in Section 25.4 that when a stock can independently jump to zero, the value of a European call is obtained by replacing the risk-free rate, r, with $r + \lambda$. As before, equity is a call option on the assets. If the firm makes no payouts, the value, B_t, of a single issue of zero coupon debt maturing at time T is

$$B_t = A_t - \text{BSCall}(A_t, \bar{B}, \sigma, r + \lambda, T - t, 0) \tag{26.119}$$

The possibility that assets can jump to zero will raise the bond yield.[52]

There is a special case where the effect of the jump probability on the bond yield is particularly easy to interpret. When the bond is default-free except for the possibility of a jump, then the bond yield is $r + \lambda$: The yield increases one-for-one with the default probability.

Example 26.2 Suppose that a firm has a single issue of zero coupon debt promising to pay \$10 in 5 years, and that $A_0 = \$90$, $\sigma = 30\%$, $r = 0.06$, and $\delta = 0$. From equation (26.119), when $\lambda = 0$, the value of the debt is

$$B_t = \$90 - \text{BSCall}(90, 10, 0.30, 0.06, 5, 0) = \$7.048$$

[52]By differentiating the expression for yield to maturity ($\ln[\bar{B}/B_t]/T$) with respect to λ, it is possible to show that the increase in yield from a 0.01 increase in λ is

$$\text{Yield increase} = \frac{\text{BSCallRrho}(A_t, \bar{B}, \sigma, r + \lambda, T - t, \delta)}{\text{bond price} \times T} \tag{26.120}$$

where "BSCallRho" is the formula for the rho (interest rate sensitivity) of a call.

The yield on debt is $\ln(10/7.048)/5 = 0.06$. This bond is priced as if it were default-free.

When $\lambda = 0.02$, the price of the bond is

$$B_t = \$90 - \text{BSCall}(90, 10, 0.30, 0.06 + 0.02, 5, 0) = \$6.703$$

The yield is then $\ln(10/6.703)/5 = 0.08$: The yield increases by the default probability.

When the bond can default apart from jumps to zero, then the increase in the bond yield is less than λ. For example, when $\bar{B} = \$100$, the bond yield is 9.263% without jumps, and 10.564% with a 2% jump to zero. ■

The models we have discussed are relatively simple: there are no coupon payments and bankruptcy occurs only at maturity. In practice, firms typically have complicated capital structures and the time of bankruptcy is uncertain. With barrier option pricing formulas, binomial valuation, Monte Carlo, or other numerical methods, it is possible to create more realistic models.

A variant of the Merton model is the Black and Cox (1976) model, in which case bankruptcy occurs if assets fall to a predetermined level, $\underline{A}$, prior to maturity. This assumption is meant to mimic a debt covenant that triggers default if the firm's finanical condition worsens sufficiently. In the Black-Cox model, equity is effectively a call option that knocks out if $A_t \leq \underline{A}$. A practical difficulty with applying this model is that the decision to default is not mechanical, but results from business judgment, which can be difficult to model.

26.3 BOND RATINGS AND DEFAULT EXPERIENCE

Bond ratings provide a measure of the credit risk for specific bonds. Such ratings, which are provided by third parties, attempt to measure the likelihood that a company will default on a bond.[53] In the U.S., the Securities and Exchange Commission (SEC) identifies specific credit-rating firms as "Nationally Recognized Statistical Rating Organizations" (NRSROs). The history and significance of this designation was explained by the Chairman of the SEC in congressional testimony:[54]

> The Commission originally used the term "Nationally Recognized Statistical Rating Organization", or NRSRO, with respect to credit rating agencies in 1975 solely to differentiate between grades of debt securities held by broker-dealers as capital to meet Commission capital requirements. Since that time, ratings by NRSROs have become benchmarks in federal and state legislation, domestic and foreign financial regulations and privately negotiated financial contracts.

[53] Companies pay ratings agencies to have their bonds rated. Some have criticized this practice, arguing that payment for ratings creates a conflict of interest.

[54] "Testimony Concerning The State of the Securities Industry", by William H. Donaldson Chairman, U.S. SEC, U.S. Senate Committee on Banking, Housing, and Urban Affairs, March 9, 2005. As of mid-2005, five firms were recognized as NRSROs: Moodys Investor Services, Standard and Poors, Fitch Ratings, Dominion Bond Rating Service, and A.M. Best.

Moodys rates bonds using the designations Aaa, Aa, A, Baa, Ba, B, Caa, Ca, and C. Within each ratings category, bonds may be further rated as a 1, 2, or 3, with 1 denoting the highest quality within a category. Standard and Poors and Fitch have a similar rating system, using the designations AAA, AA, A, BBB, BB, B, CCC, CC, C.

The market distinguishes between "investment grade" (a rating of Baa/BBB or above) and "below-investment grade" or "speculative grade" (a rating below Baa/BBB) bond. Some investors are permitted to hold only investment grade bonds, and some contracts have triggers based upon whether a company's bond rating is investment grade. For example, prior to Enron's bankruptcy, some of the company's deals contained clauses requiring that Enron make payments if Enron lost its investment grade status. Enron's financial difficulties were worsened when its rating fell below investment grade.

Using Ratings to Assess Bankruptcy Probability

A company that goes bankrupt will typically have had ratings downgrades prior to bankruptcy. By looking at the frequency with which bonds experience a ratings change, it is possible to estimate the ultimate bankruptcy probability. A change in ratings is called a **ratings transition.**

Table 26.1 is a *ratings transition matrix*, reporting the probability that a firm in a given ratings category will switch to another ratings category over the

Moody's average 1-year credit ratings transition matrix, 1970 to 2004. "WR" stands for "withdrawn rating."

Rating	Count	Aaa	Aa	A	Baa	Ba	B	Caa-C	Default	WR
Aaa	3,179	89.48	7.05	0.75	0.00	0.03	0.00	0.00	0.00	2.69
Aa	11,310	1.07	88.41	7.35	0.25	0.07	0.01	0	0	2.83
A	22,981	0.05	2.32	88.97	4.85	0.46	0.12	0.01	0.02	3.19
Baa	18,368	0.05	0.23	5.03	84.5	4.6	0.74	0.15	0.16	4.54
Ba	12,702	0.01	0.04	0.46	5.28	78.88	6.48	0.5	1.16	7.19
B	10,794	0.01	0.03	0.12	0.4	6.18	77.45	2.93	6.03	6.85
Caa-C	2,091	0	0	0	0.52	1.57	4	62.68	23.12	8.11

Source: Hamilton *et al.* (2005).

course of a year.[55] Firms rated Aaa, Aa, or A, however, all have about an 89% chance or retaining their rating over a one-year horizon, and almost no chance of suffering a default over that time. They do, however, have some chance of experiencing a downgrade, after which bankruptcy becomes likelier: The default probability increases as the rating decreases.

Given certain assumptions, we can use a short-term ratings transition matrix to compute the ultimate probability that a firm with a given rating will go bankrupt. Specifically, suppose we believe that a ratings transition matrix is constant over time and that the probability of moving from one rating to another in a given year does not depend on the rating in a previous year. Then we can use the matrix to compute the probability that a firm that is A-rated (for example), will move to any other rating, and the subsequent probability that it will move from one of the new ratings to a different rating, and so forth.

The following simple example will illustrate how to interpret and use a rat-

[55] A rating can be withdrawn if the rated obligation has matured or if the ratings agency deems there is insufficient information for a rating.

Ratings transition matrix. The entry in the i^{th} row and j^{th} column is the probability that a firm will, over one year, move from type i to type j.

		To:		
		Good	Bad	Ugly
	Good	0.90	0.07	0.03
From:	Bad	0.15	0.75	0.10
	Ugly	0.06	0.14	0.80

ings transition matrix. Suppose that securities can be in one of three categories: Good, Bad, and Ugly. The matrix in Table 26.2 displays the probability that a firm with a rating in the left hand column will, over the course of a year, move to a rating in the top row. For example, there is a 90% probability that a firm rated Good will still be good one year later. There is a 3% chance that the Good firm will become Ugly.

Notice that each row sums to one. This means that after one year each firm must be in one of the three categories. By contrast, in Table 26.1, there is a "Withdrawn Ratings" category, indicating that a firm has for some reason dropped out of the Moodys rating universe.

We can use the transition matrix to compute the probability that a firm rated Good will still be Good two years from now. To perform this calculation, recognize that there are three different paths by which a firm that is now Good can still be Good in two years.

- There is a 90% chance that a Good firm remains Good over one year. There is therefore an 81% (0.9×0.9) probability that the firm will be Good for both years.

- There is a 7% probability that the firm will be Bad next year, in which case there is a 15% chance the firm will be come Good the subsequent year. There is therefore a 1.05% (0.07×0.15) probability that the firm will go from Good to Bad and then back to Good.

- Finally, there is a 3% probability that the firm will become Ugly, and then a 6% probability that it will become Good again. There is therefore a 0.18% (0.03×0.06) probability that the firm become Ugly and then Good.

The total probablity that a Good firm will still be Good in two years is therefore 82.23%:

$$(0.90 \times 0.90) + (0.07 \times 0.15) + (0.03 \times 0.06) = 0.8223$$

In order to perform this calculation, it may at first seem necessary to enumerate the possible transitions. There is a trick, however. Notice that the calculation entails multiplying each element of the the first row of the transition matrix by the corresponding element of the first column, and then summing the results. It turns out that if we wish to know all possible transitions over a two year period, we can construct a new matrix from the one-year transition matrix, where the element in the i^{th} row and j^{th} column of the new matrix is created by multiplying the i^{th} row of the original matrix by the j^{th} column of the original matrix and summing the results. The result is in Table 26.3

In order to compute the ratings distribution after three years, we can duplicate the procedure, only taking the two year matrix and multiplying it by the

Ratings transition probabilities after two years.

		To:		
		Good	Bad	Ugly
	Good	0.8223	0.1197	0.0580
From:	Bad	0.2535	0.5870	0.1595
	Ugly	0.1230	0.2212	0.6558

one year matrix. This is the same as multiplying the one-year matrix by itself twice.

In general, let $p(i, t; j, t+s)$ denote the probability that, over an s-year horizon, a firm will move from the rating in row i to that in column j. The entries in Table 26.1 give us $p(i, t; j, t+1)$. Suppose there are n ratings. Over 2 years, the probability of moving from rating i to rating j is

$$p(i, t; j, t+2) = \sum_{k=1}^{n} p(i, t; k, t+1) \times p(k, t+1; j, t+2)$$

From the 2-year transitions we can go to 3 years, and then 4, and so on. Given the $s-1$-year transition probabilities, the s-year transition probability is

$$p(i, t; j, t+s) = \sum_{k=1}^{n} p(i, t; k, t+s-1) \times p(k, t+s-1; j, t+s) \qquad (26.121)$$

Thus, a transition probability matrix can be used to tell us the probability that a firm will go bankrupt after a given period of time.

The long-term experience of bonds with a given rating is reported in Table 26.4. Note that if a bond has a Aaa rating, even after 20 years there is only a 1.54% chance it will have gone bankrupt. However, if a bond is below-investment grade, there is a 47% chance of bankrutpcy over a 20-year horizon.

Cumulative default rates for rated bonds, 1970-2004.

	Years Since Rating						
Rating	1	2	3	5	10	15	20
Aaa	0	0	0	0.12	0.63	1.22	1.54
Aa	0	0	0.03	0.2	0.61	1.38	2.44
A	0.02	0.08	0.22	0.5	1.48	2.74	4.87
Baa	0.19	0.54	0.98	2.08	4.89	8.73	12.05
Ba	1.22	3.34	5.79	10.72	20.11	29.67	37.07
B	5.81	12.93	19.51	30.48	48.64	57.72	59.11
Caa-C	22.43	35.96	46.71	59.72	76.77	78.53	78.53
Investment grade	0.07	0.21	0.41	0.92	2.31	4.18	6.31
Below-investment grade	4.85	9.84	14.43	21.91	33.61	42.13	47.75
All Rated	1.56	3.15	4.6	6.94	10.53	13.51	16.13

Source: Hamilton *et al.* (2005).

Recovery rates (per $100 of par value) for different kinds of bonds, 1982-2003.

Priority	Mean
Senior Secured	$57.40
Senior Unsecured	$44.90
Senior Subordinated	$39.10
Subordinated	$32.00
Junior Subordinated	$28.90
All	$42.20

Source: Hamilton *et al.* (2005).

Recovery Rates

In addition to default probabilities, there is historical information about recovery rates. Table 26.5 displays historical average recovery rates for different kinds of bonds. As you would expect, bonds designated as more senior have higher recovery rates. When we modeled debt with different priorities in Section 16.1, we assumed that junior debt was not paid at all if senior debt was not completely repaid. This rule for assigning payments is called **absolute priority**. If the bankruptcy process respects absolute priority, we expect to more senior bonds with higher recovery rates.

There is considerable variation in recovery rates across firms. Hamilton *et al.* (2005) report, for example, that for senior subordinated bonds in 2004, the mean recovery rate was 44.4%, but realized recovery rates ranged between 8% and 90%, with a standard deviation of 25%.

Reduced Form Bankruptcy Models

The existence of data on corporate bond ratings, ratings changes, and defaults, suggests that one could construct statistical pricing models designed to match the behavior of bond prices. Such models are called *reduced form models*.[56] In order to price bonds we require risk-neutral probabilities, so we cannot directly use historical data.

To understand how reduced form models work, consider the simplest version of such a model. Suppose a T-year bond promises to pay $\bar{B}$ at maturity and there is a zero recovery rate in the event of default. The risk-free rate is r and constant over time. If default follows a Poisson process with the risk-neutral intensity λ, then the bond price depends only on time and on the occurrence of the jump. From Section 21.2, the partial differential equation for pricing the

[56] The reduced form approach was first used in Jarrow and Turnbull (1995). See Duffie and Singleton (2003) for a survey.

bond is

$$\frac{\partial B(t)}{\partial t} - \lambda B(t) = rB(t)$$

With the boundary condition that $B(T) = \$\bar{B}$, the bond price is[57]

$$B(t) = \bar{B}e^{-(r+\lambda)(T-t)} \tag{26.122}$$

Given our strong assumptions that recovery rate is zero and interest rates are non-stochastic, it would seem a simple matter to price this bond by observing r and inferring λ from data on defaults. The problem, however, is that equation (26.122) presumes that λ is a risk-neutral jump probability. Thus, we can infer λ from bond prices, but not from historical default data.

To understand the issue, suppose that bond defaults are idiosyncratic. In this case an investor can diversify default risk. We then expect λ in equation (26.122) to equal the historically observed λ.

If defaults are correlated, however, then even a portfolio containing numerous bonds will encounter systematic losses from correlated defaults. Defaults occur when firms perform poorly, i.e., when equity returns are low. Investors require a positive risk premium to hold such bonds, and therefore the risk-neutral default probability in equation (26.122) will exceed the historical default probability.

[57]Note that this the same bond price solution we obtain in equation (26.119) when $B < A$, and $\sigma = 0$.

A more general approach than that in equation (26.122) uses ratings transitions. Equation (26.122) does not take into account that default becomes more likely as ratings decline. With risk-neutral ratings transitions, it is possible to price bonds taking into account the various paths by which default can occur. Jarrow *et al.* (1997) show how to use observed bond prices and historical ratings transitions to infer risk-neutral ratings transition probabilities.

26.4 CREDIT INSTRUMENTS

The buyer of a corporate bond acquires both interest rate risk and the default risk of a specific firm. A particular investor may wish to hold a different combination of these risks. There are ways to alter the mix: investors can buy Treasury bonds instead of corporate bonds, thereby minimizing default risk, or use interest rate derivatives to reduce or increase exposure to interest rate risk. More recently, it has become possible to use *credit derivatives* to trade the credit risk of a specific firm. In this section, we explain these instruments.

Collateralized Debt Obligations

A **collateralized debt obligation** (CDO) is a financial structure that repackages the cash flows from a set of assets. You create a CDO by pooling the returns from a set of assets and issuing financial claims to this pool. The CDO claims reapportion the returns on the asset pool. Often, the CDO claims are *tranched*, meaning that the different CDO claims have differing priorities with respect to the cash flows generated by the collateral. With a CDO, it is possible to take a group of risky bonds, for example, and create new claims, some of which are less risky than the original bonds, and others which are riskier.

Given this general description, there are many different ways a CDO can be structured. First, the asset pool can be a fixed set of assets, in which case the CDO is *static*. If instead a manager buys and sells assets, the CDO is *managed*. Second, the CDO claims can directly receive the cash flows generated by the pool assets; this is a *cash flow* CDO. Alternatively, the CDO claims can receive payments based on cash flows and the gain or loss from asset sales; this is a *market value* CDO.

There are at least two reasons for creating CDOs. First, financial institutions will sometimes want to securitize assets, effectively removing them from the institution's balance sheet by selling them to other investors. A CDO can be used to accomplish this, in which case it is a *balance sheet CDO*. Second, a CDO can be created in response to institutional frictions. For example, some investors are permitted only to hold investment grade bonds. As we will see below, CDOs can potentially be used to create investment grade bonds from a pool of non-investment grade bonds. This is called an *arbitrage CDO*.

CDOs are relatively complex financial structures. The box on p. 511 provides an example of the legal arguments that sometimes result.

New Box:

Excerpted from "Russian Doll", The Economist*, Sept 23, 2004, p. 87.*

Considering its giant size, the decade-old credit-derivatives market is remarkably free of public squabbles. ... [However] a battle in public would be of great interest to outsiders who would like to understand the market better. Such a fight is looming.

... Some CDOs, known as "Russian dolls", contain investments in other CDOs, making their monitoring extremely complicated. To make matters more difficult still, some CDOs are actively managed—i.e., the composition of the underlying portfolio can be changed by the asset manager.

In December 2000 Barclays, a big British bank, launched a Russian-doll CDO called Corvus (the Latin for "crow"), with a face value of $950m in tranches rated from AAA to B. Corvus performed badly, particularly after September 11th 2001. Fitch began to downgrade it in December 2002, and in early 2003 took the unprecedented step of publishing details of the Corvus portfolio on its website. It included some unexpected stuff-for instance, exposure to airline leases and loans for prefabricated-housing in America. It also had exposure to other CDOs, all constructed by Barclays, some of which were never sold externally. By September 2003 the top three tranches of the CDO had fallen below a BBB rating (to "junk" status) (see Table 26.6). The other four tranches were rated CCC or lower. A similar Barclays CDO, called Nerva, fared even worse.

Investors were not amused: AAA ratings are not meant to deteriorate so rapidly.

...

If the case [a suit by HSH Nordbank, a German bank] does come to court, others will get to see how these complex financial instruments are put together and managed. One question for the court is whether Barclays correctly managed the potential conflict of interest when selecting credit risks from its own exposures for the CDO portfolio. Investors were not given details of each credit; Barclays argues that such disclosure was not required by the terms of the CDO. "The allegations in the suit are without merit," says a spokesman for the bank.

Note: The suit was settled out of court in February 2005.

Rating of tranches of the Corvus portfolio **[COMP: this accompanies the new box]**

Tranche	Amount (Millions of US $)	Dec 2000	Dec 2002	Sep 2003
A1	550	AAA	A+	BB
A2	200	AAA	A+	BB
B	65	AA	BBB	BB
C	60	A	CCC+	CCC
D	40	BBB	CC	C
E	25	BB	C	C
F	10	B	C	C

A CDO with Independent Defaults

Here is a simple example to illustrate how CDOs work. Suppose that there are three risky, speculative-grade bonds that each promise to pay \$100 in one year. Defaults are independent and occur only at maturity of the bond. Each bond has a 10% risk-neutral probability of default, a 40% recovery rate, and the risk-free rate is 6%. The price of each bond is

$$[(1 - 0.1) \times \$100 + 0.1 \times \$40] = \$88.526$$

The yield on each bond is $\ln(100/88.526) = 0.1219$.

Now suppose that there are investors wishing to invest in bonds, but some investors are happy to hold a speculative grade bond, while others seek safer bonds. We can create a CDO to rearrange the cash flows from a pool of bonds, in order to accomodate the different kinds of investors. The structure of the CDO is illustrated in Figure 26.4.1.1. The total promised payoff on the three bonds is \$300; the CDO apportions this payoff among three tranches of unequal size. The senior tranche (\$140) receives first claim to the bond payments, the mezzanine tranche (\$90) has the next claim, and the subordinated tranche (\$70) receives whatever is left. For the bond which is i^{th} in line, with a promised payment of $\bar{B}_i$, the payoff is

$$\text{Bond } i \text{ maturity payoff} = \min\left[\max\left(A_T - \sum_{j=1}^{i=1-1} \bar{B}_j, 0\right), \bar{B}_i\right] \qquad (26.123)$$

where A_T is the maturity value of the asset pool.

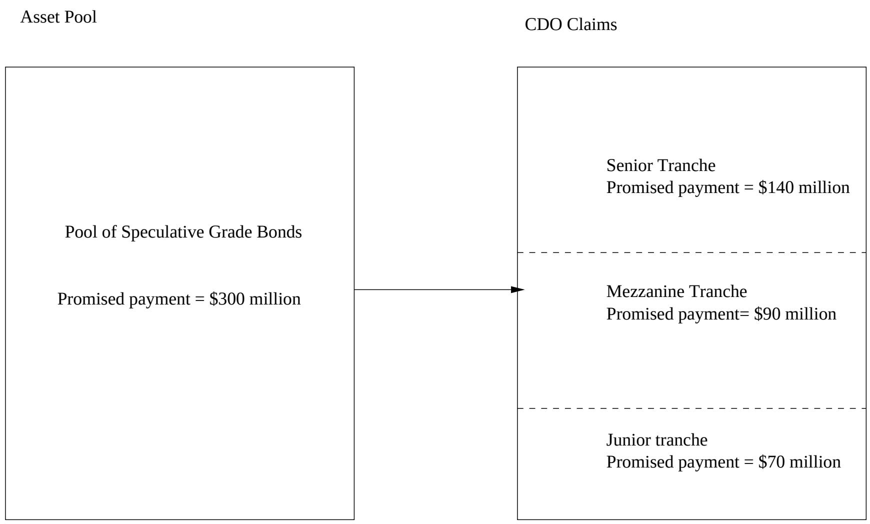

Structure of a CDO.

To understand the pricing of the CDO claims, recognize that there are four possible outcomes: no defaults ($0.90^3 = 72.9\%$ probability), one bond defaults ($3 \times 0.90^2 \times 0.10 = 24.3\%$ probability), two bonds default ($3 \times 0.90 \times 0.10^2 = 2.71\%$ probability) and three bonds default ($0.10^3 = 0.1\%$ probability). To compute the price of a CDO tranche, we can compute the expected payoff of the tranche using the risk-neutral default probabilities.

The CDO pricing is illustrated in Table 26.7. Note that the senior tranche is almost risk-free. The only time the senior tranche is not fully paid is in the unlikely (0.1% probability) event that all three bonds default. In that case, the senior tranche receives $120, a recovery rate of 85.7%. Since it is almost paid in full, investors will pay $131.828 fo the senior tranche, which is a yield of 6.02%.

The mezzanine tranche is fully paid if there is one default, but is not fully

Pricing of CDO in Figure 26.4.1.1 assuming that bond defaults are uncorrelated. Promised payoffs to the bonds are $140 (senior), $90 (mezzanine) and $70 (subordinated).

Number of Defaults	Probability	Total Payoff	Senior Bond Payoff	Mezzanine Bond Payoff	Subordinated Bond Payoff
0	0.729	300	140	90	70
1	0.243	240	140	90	10
2	0.027	180	140	40	0
3	0.001	120	120	0	0
		Price	131.828	83.403	50.347
		Yield	0.0601	0.07613	0.3296
		Default probability	0.0010	0.0280	0.2710
		Average recovery rate	0.8571	0.4286	0.1281

paid if there are two or three defaults. The yield is 7.61% and the average recovery rate is $40/90 \times 0.027/0.028 = 0.4285$. Finally, the subordinated tranche receives less than full payment if there are any defaults. Consequently, it is priced to yield 32.96%. Note that if you add up the prices of the three tranches, you $265.58. As you would expect, this is the same as the price of the three bonds put into the asset pool.

This example might remind you of the discussion of tranched debt in Chapter 16, especially Table 16.1. The idea is exactly the same, except that instead of valuing claims on corporate assets (the bonds in Chapter 16), we are valuing claims on corporate bonds.

A CDO With Correlated Defaults

In the preceding example we assumed that the bonds were uncorrelated. As you might have guessed, this is an important assumption. Table 26.8 shows how the CDO tranches are priced if the bonds are perfectly correlated, i.e., if at

Pricing of CDO in Figure 26.4.1.1 assuming that bond defaults are perfectly correlated. Promised payoffs to the bonds are \$140 (senior), \$90 (mezzanine) and \$70 (subordinated).

Number of Defaults	Probability	Total Payoff	Senior	Mezzanine Bond Payoff Bond Payoff	Subordinated
0	0.9	300	140	90	70
1	0	240	140	90	10
2	0	180	140	40	0
3	0.1	120	120	0	0
		Price	129.963	76.283	59.331
		Yield	0.074	0.165	0.165
		Default probability	0.1	0.1	0.1
		Average recovery rate	0.857143	0	0

maturity, either all firms default or none do. A comparison of Tables 26.7 and 26.8 shows the importance of default correlation in the pricing of CDOs.

Given the structure of the CDO and the assumptions about the recovery rate, with perfect correlation of defaults, the mezzanine and subordinated tranches have the same yield. The senior tranche becomes riskier because the probability of three defaults—the only circumstance in which the senior tranche is not fully paid—is greater with perfect correlation.

Nth to Default Baskets

The previous examples showed that it is possible to pool risky bonds and create riskier and less risky claims. A particular variant of this strategy is the *nth to default basket*.

Consider a CDO that contains equal quantities of N bonds. Over the life of the CDO there can be anywhere between 0 and N defaults. It is possible to create tranches where particular bond holders bear the consequences of a

Pricing of nth to default bonds. Assumes the bonds owned as assets have uncorrelated defaults.

		Probability	Payoffs		Expected		
Default	Prob	N or more	Default	No default	payoff	Price	Yield
First	0.243	0.271	40	100	83.74	78.863	0.237
Second	0.027	0.028	40	100	98.32	92.594	0.077
Third	0.001	0.001	40	100	99.94	94.120	0.061

Pricing of nth to default bonds. Assumes the bonds owned as assets have perfectly correlated defaults.

		Probability	Payoffs		Expected		
Default	Prob	N or more	Default	No default	payoff	Price	Yield
First	0.100	0.100	40	100	94.000	88.526	0.122
Second	0.100	0.100	40	100	94.000	88.526	0.122
Third	0.100	0.100	40	100	94.000	88.526	0.122

particular default.

The owner of the first-to-default tranche bears the most risk: If any of the bonds in the asset pool default, the first-to-default owner bears the loss from this default. The owner of the last-to-default geenerally bears the least risk, since all bonds must default in order for this claim to bear a loss.

Table 26.9 shows the pricing that results from this structure, assuming that the defaults are uncorrelated and occur only at time T. By comparing Table 26.9 with Table 26.7, you can observe a similarity between the subordinated tranche and the first-to-default on the one hand, and on the other, the senior tranche and the third to default.

Figure 26.3 illustrates the pricing of Nth to default bonds for a 10-bond pool assuming default correlations of 0, 0.25, and 1.[58] Again, we assume defaults

[58]How do we generate a default correlation of 0.25 among ten bonds? We can generate

occur only at time T. The results for a default correlation of 1 is like that in Table 26.10: When one bond defaults, all default, so all claims have the same yield. The graphs for correlations of 0 and 0.25 show significant differences for yields for low numbers of defaults. The first-to-default tranche, for example, has a yield of 16.02% when defaults are uncorrelated and 13.51% when the default correlation is 0.25. This again illustrates that Nth to default baskets are a way to speculate on default correlations.

Credit Default Swaps and Related Structures

Credit derivatives, which have been trading since the early 1990s, are contracts that permit the trading and hedging of credit risk. Credit derivatives permit institutions to hedge credit risk, in much the same way that, for example, gold futures permit the hedging of gold price risk. Table 26.11 shows that the outstanding notional principal covered by credit default swaps, an important kind of credit derivative, has grown significantly over the last few years. In ad-

ten independent normally distributed random variables and create correlated normal random variables, z_i, using the Cholesky decomposition (equation 19.17). Following this procedure, each *pairwise* correlation is 0.25. Assuming a 10% unconditional default probability, a firm defaults if $z_i < -1.2816$. Note that once we use this method, we can also let z_i determine *when in time* the default occurs; we simply interpret a smaller z_i to mean that default occurs earlier. It is then necessary to specify a distribution for default times. For example, suppose default occurs with a constant Poisson intensity. Then the time to default is exponentially distributed, so that

$$\Pr(\text{default} < t) = 1 - e^{-ht}$$

Given this assumption, the default time is then $\tau = -\ln[N(z_i)]/h$. See Duffie and Yurday (2004a) and Watkinson and Roosevelt (2004) for a discussion of pricing credit tranches with default correlations modeled in this fashion.

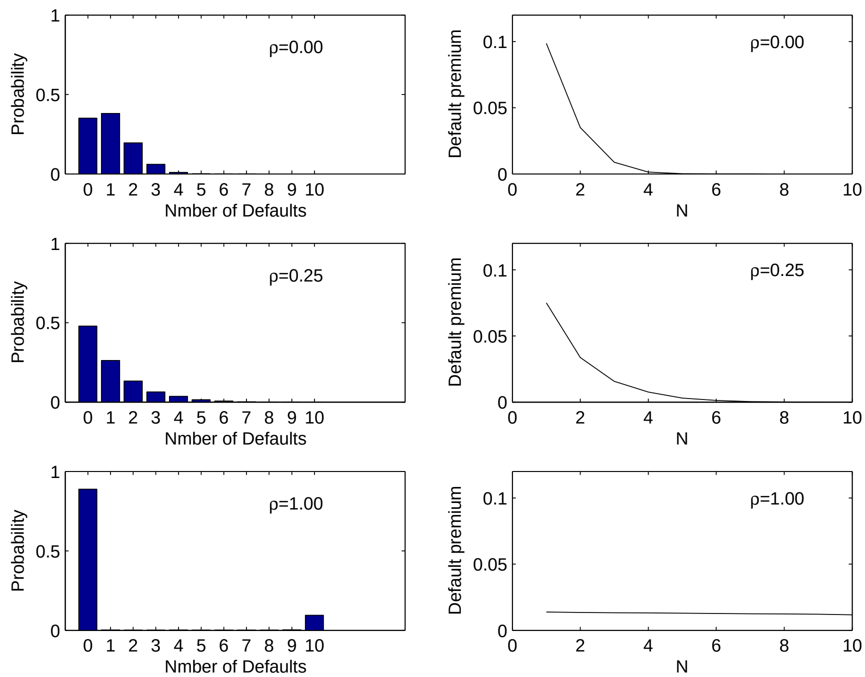

Pricing of an Nth to default basket containing equal quantities of 10 bonds, each with

Credit-default swaps outstanding, 2001–2004, billions of US dollars. Source: ISDA.

Year	06/01	12/01	06/02	12/02	06/03	12/03	06/04	12/04
Amount	631.50	918.87	1,563.48	2,191.57	2,687.91	3,779.40	5,441.86	8,422.26

dition to general growth, there has been a great deal of innovation in the credit derivatives market in the last few years.

Single Name Credit Default Swaps

A single name credit default swap (CDS) protects the owner of a bond issued by a specific company (the "single name") against a credit event sustained by the issuer of the bond. The buyer of the swap is the *protection buyer*. The market-maker or investor providing the credit insurance is the *swap writer* or *protection seller.*

If a credit event occurs, the protection buyer receives

$$\text{Payoff} = \text{Bond par value} - \text{Bond market value} \tag{26.124}$$

The protection buyer pays an insurance premium (to the seller), which is amortized. The premium payments stop once default occurs.

Figure 26.4 illustrates the cash flows and parties involved in a credit default swap on XYZ. Note in particular that there is no connection between XYZ and the swap writer. The default swap must specify a specific XYZ debt issue, called the *reference asset*. The reference asset is important because different bonds from the same issuer can have different prices after a default.

If there is an actual default, the default swap could settle either financially or physically. In a financial settlement, the swap writer would pay the bondholder the value of the loss on the bond. The bondholder would continue to hold the defaulted bond. In a physical settlement, the swap writer would buy the defaulted bond at the price it would have in the absence of default.

Financial settlement and physical settlement are economically equivalent in theory. However, the market for a defaulted corporate bond may not be liquid, and it may be difficult to determine a fair price upon which to base financial settlement. To avoid this problem, default swaps often call for, or at least permit, physical settlement.

For many firms, it is possible to trade CDSs for a variety of different expi-

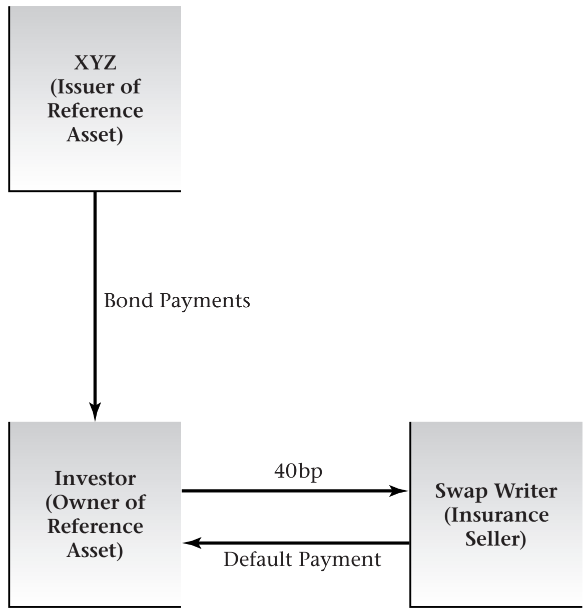

Depiction of the cash flows in a credit default swap. The investor owns the reference asset, which was issued by XYZ. The investor buys a credit default swap and pays 40 basis points per year in exchange for the swap writer's payment in the event of a default by XYZ.

ration. The set of prices with different maturities generates a *credit spread curve*, where, for example, you may observe that credit spreads are small at short horizons and substantially larger over five years. With this array of contracts it is possible to make sophisticated bets. For example, you could buy protection with a four year horizon and sell protection with a five year horizon; this is a bet that default will not occur in the fifth year. There are also swaps and options on credit spreads and CDSs.

Finally, recognize that there is credit risk in the default swap. The default swap buyer faces the possibility that the swap writer will go bankrupt at the

Comparison of a written CDS on XYZ's bond, and a CDO containing only an XYZ bond.

	Position		
Event	CDO Buyer	Written CDS	Difference (CDO − CDS)
No Default	$\$100 + c$	ρ	$\$100 + c - \rho$
Default	B_1	$\rho - [(\$100 + c) - B_1]$	$\$100 + c - \rho$

same time as a default occurs on the reference asset. The default swap market does permit investors to trade exposure to particular credits, but there is still counterparty default risk.

CDSs vs. CDOs

It may have already occurred to you that a CDO and a CDS are similar. Suppose that XYZ issues a bond promising to pay in one period $\$100+c$, where c is the coupon. Suppose that a CDS guarantees the payment of both principal and the coupon, and that the protection buyer pays the CDS premium, ρ, under any circumstances. We can see the similarity between a CDO and a CDS by comparing the cash flows for the following two positions:

- the payoffs to a *protection writer* for a CDS on XYZ's bond.
- the payoff to the buyer of a single tranche CDO containing the XYZ bond (this is just the XYZ bond).

Table 26.12 compares the payoffs on the two positions.

Both the written CDS and the CDO bear the costs of a default. The difference between the two positions—in the last column of Table 26.12—is the payoff to a default-free bond. The difference between the CDO and written CDS, therefore, is that the CDO is **funded**, meaning that the investor pays for it fully at the outset, while the CDS is **unfunded**, meaning that the investor pays nothing at the outset, and can have payment obligations in the future. We can think of the CDO as analogous to a prepaid forward on the bond, while the written CDS is a forward contract on the bond.

Pricing a default swap

The preceding comparison of a CDO and CDS shows that hedging the bond's credit risk converts the bond into the economic equivalent of a Treasury bond. This simple argument shows that the default swap premium should equal the difference between the yield on the reference asset and the yield on an otherwise equivalent default-free bond. In other words, $r = c - \rho$, or

$$\rho = c - r \tag{26.125}$$

In other words, the default swap premium equals the credit spread.

The idea that the default swap premium should equal the credit spread seems straightforward. Pricing is more complicated than this, however. First, consider a protection seller who wants to hedge, a situation described by Rooney (1998). The protection seller will buy synthetic protection by short-selling the reference asset and buying the equivalent default-free bond. The costs of short-selling, which can reflect scarcity, will be reflected in the price of the default swap.

Second, what exactly is an "otherwise equivalent default-free bond"? It seems natural to use government bonds as a benchmark, but government bonds can be unique in certain respects. Prices of government bonds may include a liquidity premium and sometimes reflect special tax attributes (for example, in the United States, federal government bonds are exempt from state taxes). Houweling and Vorst (2001) estimate a credit swap pricing model and find that, empirically, credit swap premiums are more related to the interest rate swap curve than to the government bond yield curve. This linkage between credit swaps and interest rate swaps suggests that the yield on an "equivalent default-free bond" is not the government bond curve, and in fact may not be directly observable. Rather the equivalent default-free yield may be inferred from the market for default swaps as the rate on the reference asset less the default swap premium.

Credit-linked notes Banks can issue securities known as credit-linked notes to hedge the credit risk of loans. A **credit-linked note** is a bond issued by one company with payments that depend upon the credit status (for example, bankrupt or not) of a different company. To see how a credit-linked note works, suppose that bank ABC lends money to company XYZ. At the time of the loan, ABC creates a trust that in turn issues notes bought by investors. These notes are the credit-linked notes. The funds raised by the issuance of these notes are invested in bonds with a low probability of default (such as government bonds), which are held in the trust. If XYZ remains solvent, ABC is obligated to pay the notes in full. If XYZ goes bankrupt, the note-holders receive the XYZ loans and become creditors of XYZ. ABC takes possession of the securities in the trust. Thus, the credit-linked note is in effect a bond issued by ABC, which ABC does not need to repay in full if XYZ goes bankrupt. This structure eliminates a third-party insurance provider.

Because of the trust, the credit-linked notes can be paid in full even if ABC defaults. Thus, even though they are issued by ABC, the interest rate on the notes is determined by the credit risk of XYZ. An arrangement like this was used by Citigroup when it made loans to Enron in 2000 and 2001.[59] When Enron went bankrupt in late 2001, Citigroup avoided losses on over $1 billion in loans because it had hedged the loans by issuing credit-linked notes.

Total Rate of Return Swaps

Total rate of return swaps, discussed in Chapter 8, can also be used to hedge credit risk. The difference between a total rate of return swap and the credit products we are discussing is that a total return swap for a bond makes payments based on changes in market value due to interest rate changes as well as due to credit changes. A CDS only makes payments if there is a credit event.

[59]See "How Citigroup Hedged Bets on Enron" by Daniel Altman, *New York Times*, p. C1, Feb. 8, 2002.

CDS Indices

A CDS index is an average of the premiums on a set of CDSs. Thus, a CDS index provides a way to track the overall market for credit. In this section we will discuss the kinds of structures and products that are traded.

To a first approximation, it is possible to replicate a CDS index by holding a pool of CDSs.[60] As with a single CDS, one party is a protection seller, receiving premium payments, and the other is a protection buyer, making the payments but receiving a payment from the seller if there is a credit event.

A CDS index can also be described as a *synthetic CDO*. The returns on a CDS index reflect the performance of the bonds in the index, just as the returns on a CDO reflect the performance of the bonds in the asset pool. There are numerous ways in which a CDS index product can be structured and traded:

- a CDS index can be funded or unfunded. The unfunded version is like a credit default swap, except that the underlying asset is a basket of firms. The funded index could be is a note linked to the index with a collateral arrangement in case the note buyer had to make a payment. Essentially, a synthetic CDO is an unfunded CDO.
- the claims can be tranched in various ways: for example, simple priority or

[60] The replication may not be exact since a default index can define a credit event differently than a single name CDS. For example, DJ TRAC-X, discussed below, specifies bankruptcy and failure to pay as credit events. A single name CDS can have a broader definition.

Nth to default.

- the underlying assets can represent different countries, currencies, maturities, or industries.

Credit indices have had a relatively brief history.[61] There are a number of potential challenges in trading credit risk. Historically, the corporate bond market has been less liquid than the stock market. Trading volume for bonds is lower than for stocks, and bid-ask spreads have been higher. A large firm can have dozens of different bond issues outstanding.

Single name credit default swaps first traded in the mid 1990s. In 2001 Morgan Stanley introduced TRACERS, a basket of corporate bonds that investors could trade as a unit, much like an exchange traded fund. There were both funded and non-funded TRACERS products. In 2003, Morgan Stanley and JP Morgan together created TRAC-X as a successor to TRACERS. Later in 2003, competing banks introduced iBoxx as an alternative credit index. In 2004, the TRAC-X and iBoxx products merged into Dow Jones CDX. These indices could be traded in both funded and non-funded products.

Firms can drop out of an index due to bankruptcy or illiquidity, but otherwise the make-up of a given CDX offering is set for the life of the offering. A new CDX is currently offered every six months with five and ten year maturities, with the most recent offering being the most heavily traded. For North Amer-

[61] See Duffie and Yurday (2004b) for an interesting account of the development of credit index products.

ica, there are versions of the index reflecting investment grade companies, high volatility companies, and high yield companies. The CDX North America index is an equally-weighted basket of 125 CDSs representing investment grade companies. There are also credit indices for North America, Europe, Asia, Japan, Australia, and emerging markets.

Figure 26.5 shows the CDS premium for the five-year on-the-run index. The series was constructed by Morgan Stanley from premiums on TRAC-X and CDX, and for years before TRAC-X, from data on credit defalt swaps from firms in the index. Credit default swap premia were four times larger in 2002 than at the beginning of 2005. This is due to the fact that a number of large companies defaulted during 2002, including Worldcom, Global Crossing, UAL Corp, Adelphia, and KMart, among others. The CDS premiums for these companies would have been high prior to their defaults. In interpreting Figure 26.5, keep in mind that after defaulting, a company drop out of the index, leaving behind a safer set of firms.

Tranched CDX is available, and has a payoff structure illustrated in Figure 26.6. There are both funded and unfunded products based on this structure. As you would expect from our earlier discussion of CDOs and default correlation, the pricing of the various tranches is senstive to correlation. In fact, tranche pricing is quoted using implied correlation in much the same way equity option premiums are quoted using implied volatility.

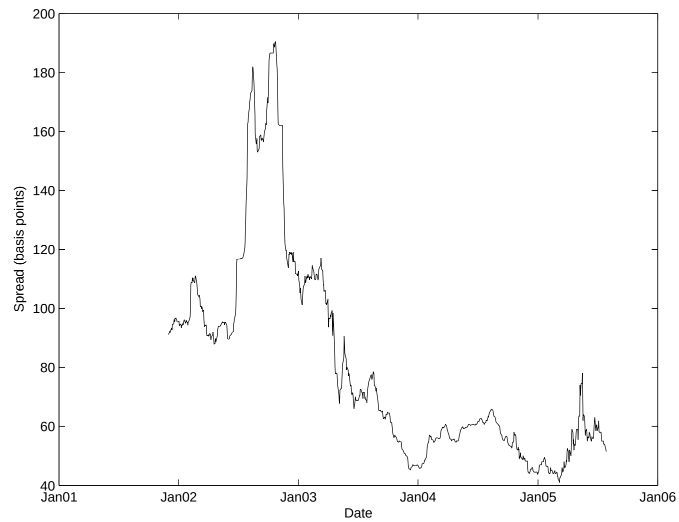

Premium for on-the-run five year basket of credit default swaps. Source: Morgan Stanley.

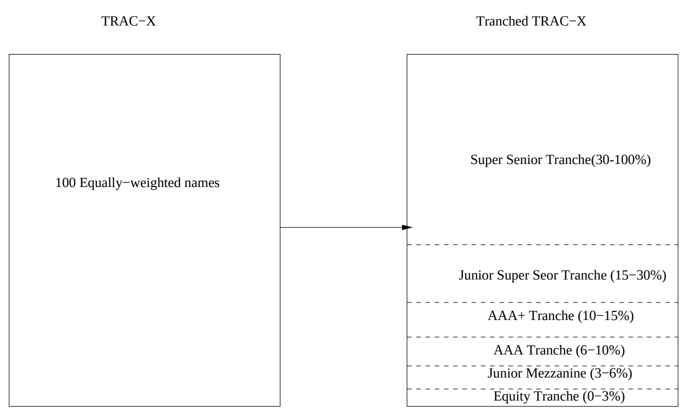

The structure of Tranched CDX. Source: Morgan Stanley.

*

Chapter Summary

A party to a contract may fail to make a required future payment. This possibility, which is called default, gives rise to credit risk. Credit risk is an important consideration in valuing corporate bonds, where two key inputs are the probability of default and the expected payoff to the bond if the firm does default. The Merton model uses option pricing to value debt subject to default. Credit agencies assign debt ratings to firms; these can be used to assess the probability of bankruptcy in the future.

There are various financial vehicles that permit the trading of credit risk. Collateralized debt obligations (CDOs) are claims to an asset pool. The claims are often structured (tranched) so as to create new claims, some of which are more and some of which are less sensitive to credit risk than the pool as a whole. The value of these claims depends importantly on the default correlation of the assets in the pool.

Credit default swaps pay to the protection buyer the loss on a corporate bond when there is a default. In exchange, the protection buyer makes a periodic premium payment to the seller. Baskets of credit default swaps are also called synthetic CDOs, since the holder of such a basket bears the default experience of the firms represented in the basket, as with a CDO. A synthetic CDO can be funded (the investor pays fully at the outset) or non-funded (there is no initial payment).

The average premium on a basket of default swaps is effectively a credit index. These indices, such as Dow Jones CDX, serve as the basis for a variety of investment vehicles and investment strategies.

FURTHER READING

Both the actual traded credit contracts and pricing theory continue to evolve. Books with a practitioner perspective include Goodman and Fabozzi (2002) and Tavakoli (2001). Frameworks for analyzing credit risk are discussed in Credit Suisse Financial Products (1997), J. P. Morgan (1997), and whitepapers on the Moody's KMV website, `www.moodyskmv.com`. A debate between advocates and critics of the KMV approach is in the February 2002 issue of *Risk* (Kealhofer and Kurbat 2002, Sobehart and Keenan 2002).

Books with a more academic and theoretical perspective include Cossin and Pirotte (2001), Duffie and Singleton (2003), Meissner (2005), and Schönbucher (2003). Duffie and Yurday (2004ba) provide a blended discussion of the history of some of the products, practical issues, and pricing models. Papers examining Merton-style models include Jones *et al.* (1984), Kim *et al.* (1993), Leland (1994), Leland and Toft (1996) and Longstaff and Schwartz (1995). Empirical studies of Merton-style models include Anderson and Sundaresan (2000) and Eom *et al.* (2004).

Data on bankruptcies is available from `www.bankrtupcydata.com`. Simple summary data, such as the largest bankruptcies sorted by year, is available for free.

PROBLEMS

[More to come...]