

Credit and Counterparty Risk

To understand credit risk and how to measure it, we need both a set of analytical tools and an understanding of such financial institutions as banks and rating agencies. In this chapter

- We define credit risk and its elements, such as the likelihood that a company goes bankrupt, or the amount the investor loses if it happens.
- A great deal of effort goes into assessing the credit risk posed by borrowers. This is, in fact, one of the oldest activities of banks. We will look at different ways this is done, including both time-sanctioned, relatively non-quantitative techniques and more recently developed modeling approaches.
- As with market risk, we sometimes want to summarize credit risk in one number, such as credit Value at Risk, so we will also look at quantitative approaches to measuring credit risk.

This is the first, covering basic concepts, in a sequence of chapters on credit risk. One way credit risk is expressed is through the spread, or the difference between credit-risky and risk-free interest rates. Since the market generally demands to be compensated for credit risk, credit-risky securities are priced differently from securities that promise the same cash flows without credit risk. They are discounted by a credit risk premium that varies with the perceived credit risk and market participants' desire to bear or avoid credit risk. When measured in terms of the interest rate paid on a debt security, this premium is part of the *credit spread*. In the next chapter, we will extend Section 4.2's discussion of interest-rate analytics and risk measurement to credit-risky securities.

In practice, very few investors or financial institutions have exposure to only one credit-risky security. The present chapter also sets up the concepts we need to study portfolios of credit-risky securities in Chapter 8. Finally, we

apply the techniques of portfolio credit risk measurement to the valuation and risk measurement of structured credit products in Chapter 9.

6.1 DEFINING CREDIT RISK

Let's begin by defining some terms. *Credit* is an economic obligation to an "outsider," an entity that doesn't own equity in the firm. *Credit risk* is the risk of economic loss from *default* or changes in ratings or other *credit events*.

Credit-risky securities include:

Corporate debt securities are the only type that can default in the narrowest sense of the word. The most common members of this group are fixed and floating rate bonds, and bank loans.

Sovereign debt is denominated either in the local currency of the sovereign entity or in foreign currency. It may be issued by the central government or by a state-owned or state-controlled enterprise. State or provincial and local governments also issue debt. In the United States, such issues are called *municipal bonds*. We discussed credit risk issues around sovereign debt in Chapter 2 and have more to say in the context of financial crises in Chapter 14.

Credit derivatives are contracts whose payoffs are functions of the payoffs on credit-risky securities. The most important and widespread are credit default swaps (CDS), which we introduced in Chapter 2 and discuss in more detail in Chapter 7.

Structured credit products are bonds backed by pools of mortgage, student, and credit card loans to individuals, by commercial mortgages and other business loans, and by other types of collateral. They are often not defaultable in the narrow sense that the issuer can file for bankruptcy. They are, however, credit risky in the sense that, when enough loans in the collateral pool default, at least some of the liabilities issued against the collateral must be written down, that is, the creditor takes a loss.

All of these types have in common that their interest rates include a credit spread; interest rates on these securities are higher than credit risk-free securities with the same promised future cash flows.

6.2 CREDIT-RISKY SECURITIES

6.2.1 The Economic Balance Sheet of the Firm

We start with an *economic balance sheet* for the firm:

Assets	Liabilities
Value of the firm (A_t)	Equity (E_t)
	Debt (D_t)

While this looks familiar, it differs from the *accounting* balance sheet in that asset values are not entered at book or accounting values, but at market values, or at some other value, such as an option delta equivalent, that is more closely related to the market and credit risk generated by the asset. In Chapter 12, we use this concept to create more accurate measures of the firm's indebtedness. Here, we will use an economic balance sheet to more accurately value the firm's equity, that is, the part of the value of the assets belonging to the owners of the firm once the debt has been deducted.

The assets of the firm, equal to A_t at current market prices, produce cash flows, and, hopefully, profits.¹ These assets are financed by:

Debt obligations are contractually bound to pay fixed amounts of money. Occasionally, debt may be repaid in the form of securities, as in the case of pay-in-kind (PIK) bonds, discussed just below. A debtor or issuer of debt securities is called the *obligor*.

Equity is the *capital* invested by the firm's owners. Once the creditors—the owners of the firm's debt securities—are paid any interest and principal they are owed in full, the firm's owners can keep any remaining cash flow, either as a *dividend* they permanently extract from the firm, or to be added to their equity capital and reinvested in the firm. Equity capital absorbs any losses fully until it is exhausted. Only then does debt take a loss.

The ratio of equity to assets $\frac{E_t}{A_t}$ is called the *equity ratio*. The ratio of assets to equity $\frac{A_t}{E_t}$ is (often) called the *leverage ratio*. We discuss leverage in more detail in Chapter 12.

¹The asset value of the firm is greater than its *enterprise value* by the value of its cash and cash equivalent assets.

6.2.2 Capital Structure

So far, we've presented a simplified version of the firm's balance sheet that distinguishes only between equity and debt. Equity receives the *last* free cash flow, and suffers the *first* loss. Within the capital structure of a firm, however, different securities have different rights. These rights determine their *seniority* or *priority* within the capital structure.

Debt seniority refers to the order in which obligations to creditors are repaid. *Senior debt* is paid first, while *subordinated* or *junior debt* is repaid only if and when the senior debt is paid.

Many corporate debt securities combine characteristics of equity and debt. The issues raised have become particularly important for financial firms in the context of regulatory policy:

Preferred stock or “pref shares” are similar to bonds in that they pay a fixed dividend or coupon if and only if all other debt obligations of the firm are satisfied. However, they often do not have a fixed maturity date (*perpetual preferred stock*), and failure of the pref shares to pay their dividend does not usually constitute an event of default. Rather, if the shares are *cumulative*, the dividends cumulate, and must be paid out later before the common shares receive a dividend. The pref shares may then also receive voting rights, which they do not generally have.

Pref shares have a priority between that of equity and bonds; in the event of default, preferred stock bears losses before the bonds but after common equity is wiped out. In spite of the name, they behave for the most part like highly subordinated bonds with a distant maturity date.

Convertible bonds are bonds that can be converted into common shares. They therefore have some characteristics of options, such as option time value and risk sensitivity to the implied volatility of the issuer's stock price. But they retain, as well, characteristics of corporate bonds, such as risk sensitivity to interest rates and credit risk.

Conventional, or “plain-vanilla” convertible bonds, act broadly speaking like bonds with equity options attached. Some varieties of convertible bonds, however, act almost entirely like options or like bonds:

Mandatory convertible bonds are bonds that must be converted into equity at a future date, and pay a fixed coupon during that period. They generally have terms to maturity of about three years, so the present value of the coupon payment is not

generally a large part of the value. Rather, their values are close to that of a portfolio consisting of a long out-of-the-money call option and a short out-of-the-money put option.

Convertible preferred shares are pref shares that are convertible into common stock.

Payment in Kind (PIK) bonds do not pay interest in cash, but rather in the form of additional par value of bonds. The amount of bonds issued thus rises over time. They are typically issued by borrowers who do not have sufficient cash flow to meet cash interest obligations and have historically typically been issued by purchasers of entire companies with borrowed money. In these *leveraged buyouts*, which we discussed in Chapter 1, the purchasing firm may have a high debt ratio and seeks *ab initio*, rather than conditionally on encountering financial difficulties, to defer cash interest payments as long as possible.

Such bonds are usually much riskier than bonds paying regular cash interest, since they increase the firm's indebtedness over time, and have higher nominal interest rates. If PIK bonds are not the most junior in the capital structure, they also have an adverse effect on the credit quality of any bonds to which they are senior. The PIK feature may also be an option, to be exercised by the company if it has concerns about its future cash flow (*PIK toggle notes*).

6.2.3 Security, Collateral, and Priority

One of the major problems in engaging in credit transactions is to ensure performance by the obligor. Over the many centuries during which credit transactions have been developed, a number of mechanisms have been developed for providing greater assurance to lenders. These mechanisms benefit both parties, since, by reducing the expected value of credit losses, they reduce the cost of providing credit.

Credit obligation may be *unsecured* or *secured*. These are treated differently in the event of default. Unsecured obligations have a general claim on the firm's assets in bankruptcy. Secured obligations have a claim on specific assets, called *collateral*, in bankruptcy. A claim on collateral is called a *lien* on the property. A lien permits the creditor to seize specific collateral and sell it, but the proceeds from the sale must be used only to discharge the specific debt collateralized by the property. Any proceeds left over are returned to the owner. They cannot be used by the creditor even to repay other debt owed to him by the owner of the collateral.

Most liens are on real estate, but bonds, subsidiaries of large firms, a specific factory, or even personal property can also serve as collateral. A

pawnbroker, for example, has a lien on the goods pledged or pawned; if the debt is not discharged within the 30 days typically set as the term for a pawnshop loan, the pawnbroker can sell the goods to discharge the debt.

In collateralized lending, a *haircut* ensures that the full value of the collateral is not lent. As the market value of the collateral fluctuates, the haircut may be increased or lowered. The increment or reduction in haircut is called *variation margin*, in contrast to the *initial haircut* or *margin*.

Secured loans can be with or without *recourse*. The proceeds from the sale of the property collateralizing a loan may not suffice to cover the entire loan. In a non-recourse or *limited liability* loan, the lender has no further claim against the borrower, even if he has not been made whole by the sale of the collateral. If the loan is with recourse, the lender can continue to pursue a claim against the borrower for the rest of what is owed.

Secured debt is said to have *priority* over unsecured debt in the event of bankruptcy. Even within these classes, there may be priority distinctions. Secured claims may have a first, second, or even third lien on the collateral. The claims of a second-lien obligation cannot be met until those of the first lien are fully satisfied. Similarly, unsecured debt may be senior or junior. A *debenture*, for example, is a junior claim on the assets that are left after all the secured claims and all the senior unsecured claims have been satisfied.

6.2.4 Credit Derivatives

Credit risk can be assumed not only in the form of the *cash securities*—bonds, notes, and other forms of corporate debt—but also in the form of derivatives contracts written on an underlying cash security. The most common type of credit derivative are CDS, in which one party makes fixed payments each period to the other party unless a specified firm goes bankrupt. In case of bankruptcy, the other party to the contract will pay the value of the underlying firm's bonds. We describe CDS in more detail in the next chapter.

Derivatives can be written not only on the obligations of individual firms, but also on portfolios. We describe some common types of portfolio credit derivatives in our discussions of structured credit in Chapter 9, and of model risk in Chapter 11.

6.3 TRANSACTION COST PROBLEMS IN CREDIT CONTRACTS

Credit contracts have a number of problems that can be subsumed under the concepts of *transaction costs* and *frictions*. Credit contracts are rife with conflicts of interest between the contracting parties. Many of these conflicts arise from information problems that are inherent in credit transactions and

are costly to overcome. Acquiring information about a borrower's condition is costly, and harmonizing the actions of market participants involves negotiation, also costly. In understanding credit risk, it is helpful to be familiar with concepts from economics that help in identifying and analyzing these conflicts:

Asymmetric information describes a situation in which one party has different information than another. In credit contracts, the borrower generally has more information than the lender about the project the loan proceeds have been applied to, and thus about his ability to repay. Information disparities can be mitigated through *monitoring* by the lender and reporting by the borrower, but only incompletely and at some cost.

Principal-agent problems arise because it is costly to align incentives when a principal employs an agent, and the latter has better information about the task at hand. A common example is investment management; the manager, though employed as the investor's agent, may maximize his own fee and trading income rather than the investor's returns. Another is *delegated monitoring*, which arises for depositors and other creditors of banks. Bank managers are charged with monitoring, on behalf of the bank's depositors and other creditors, how bank loan proceeds are being used. Apart from banking, as we see in Chapter 9, principal-agent problems are particularly difficult to address in structured credit products, since managers of the underlying loans may also own the securities.

Risk shifting can occur when there is an asymmetry between the risks and rewards of market participants who have different positions in the capital structure of the firm or different contracts with a firm's managers. The classic example is the conflict between equity investors and lenders. Increasing risk to the firm's assets can benefit the equity investor, since their potential loss is limited to their equity investment, while their potential return is unlimited. Debt holders have no benefit from increased risk, since their return is fixed, only the increased risk of loss. Increasing risk therefore shifts risk from equity to bondholders.

The problem of risk shifting, however, also occurs in the context of regulation of financial intermediaries, and in particular, the problem of "too-big-to-fail." If at least some positions in the capital structure, such as the senior unsecured debt, will be protected in the event of a failure of the firm, then increasing the risk to the firm's assets may shift risk to the public rather than to bondholders. Bondholders then will not need to be compensated for the increased risk. We discuss these and related issues in more detail in Chapter 15.

Moral hazard is an old term, originating in economic analysis of the insurance business. The problem it describes is that buying insurance reduces the incentives of the insured to avoid the insurable event. Moral hazard describes a situation in which (a) the insured party has some ability to mitigate the risk of occurrence of the event against which he is insured, (b) the insurer cannot monitor the action the insured does or doesn't take to avert the event, and (c) mitigating the risk is costly for the insured. For example, a person insuring a residence against fire might not buy costly smoke detectors that could reduce the risk of fire damage, or a person with medical insurance might take less care of himself, or use more medical services, than someone without health insurance.

In finance, it arises in the propensity of financial firms that can expect to be rescued from bankruptcy by public policy actions to take greater risks than otherwise. Moral hazard arises in many financial contracts aside from insurance. It can occur whenever one party to a contract has incentives to act in a way that does harm or diminishes economic benefit to another party in a way the other party can't easily see or find out about.

Adverse selection or the "lemons problem" also occurs when transaction parties possess asymmetric information. This issue arises in all trading of financial assets. The fact that a seller is bringing an asset to market is a bit of evidence that he knows something negative about the asset that is not yet incorporated into its price. It may or may not be true that the seller has negative information, but even a low probability that such is the case must lower the price the buyer is prepared to pay. In models of liquidity, this phenomenon is one of the fundamental explanations of the bid-ask spreads (see Chapter 12).

An example is the sale by a bank of loans or securities into the secondary markets or into a structured credit product. A bank is expected to be well-informed and exercise care about the credit quality of its assets, especially that of loans it has originated. The incentive to do so may be diminished if it plans to sell the assets shortly after originating them. We discuss this phenomenon and its effects in Chapter 9.

Externalities are benefits or costs that the actions of one actor cause another without the ability to exact compensation via a market transaction. A common example arises in short-term lending markets, where the excessive risk-taking of one or a small number of borrowers can raise borrowing costs for prudent borrowers, as potential lenders are uncertain about the risk-taking of each

individual firm. Lenders have less information than borrowers about how much risk each is taking, and asymmetric information generates the externality.

Collective action problems or *coordination failures* are situations in which all would benefit from all taking a course of action that is not to the individual's benefit if he alone takes it. The classic example studied in game theory is the Prisoner's Dilemma. Another important example is the "tragedy of the commons," in which a limited resource is depleted by overuse.

Examples of coordination failure in finance occur when creditors of a particular class cannot come to agreement with one another and are made worse off in corporate bankruptcy as a result. A typical way in which this occurs is a restructuring plan in which junior creditors receive equity in the reorganized company. This class of creditors will be worse off if the company is liquidated and the assets sold. But any single holdout among them will be better off if all the rest accept the reorganization plan and they keep their credit claim against the company. The collective action problem is to induce them all to agree to accept the less-valuable equity claims.

Another example from financial crises are bank runs (see Chapters 12 and 14), which can be modeled as resulting from a collective action problem among bank depositors.

These information problems do not just increase the costs of credit transactions. To the extent they involve externalities, they may also introduce problems of public policy that are difficult and costly to mitigate.

Conflicts of interest may be resolvable through pricing. Conflicts between creditors with claims of different priorities, for example, can be resolved *ex ante* through the pricing of different securities.

6.4 DEFAULT AND RECOVERY: ANALYTIC CONCEPTS

In this section, we introduce some basic analytical concepts around credit events. We use these concepts extensively in credit risk modeling.

6.4.1 Default

Default is failure to pay on a financial obligation. Default events include *distressed exchanges*, in which the creditor receives securities with lower value or an amount of cash less than par in exchange for the original debt. An

alternative definition of default is based on the firm's balance sheet: Default occurs when the value of the assets is smaller than that of the debt, that is, the equity is reduced to zero or a negative quantity. *Impairment* is a somewhat weaker accounting concept, stated from the standpoint of the lender; a credit can be impaired without default, in which case it is permissible to write down its value and reduce reported earnings by that amount.

Example 6.1 (Distressed Exchange) CIT Group Inc. is a specialty finance company that lent primarily to small businesses, and fell into financial distress during the subprime crisis. On August 17, 2009, it obtained agreement from bondholders via a tender offer to repurchase debt that was to mature that day, at a price of 87.5 cents on the dollar. The company stated that, had it been obliged to redeem the bonds at par, it would have been forced to file for bankruptcy protection. CIT Group was downgraded to selective default status as a result.

Bankruptcy is a legal procedure in which a person or firm “seeks relief” from its creditors to either reorganize and restructure its balance sheet and operations (Chapter 11 in the United States), or liquidate and go out of business in an orderly way (Chapter 7). During the first half of the nineteenth century, limited liability of corporations in the United Kingdom and the United States became generally recognized, paving the way for public trading in their securities. Creditors of limited liability corporations and partnerships do not have recourse to property of shareholders or partners apart from the latter's invested capital.

In practice, firms generally file for bankruptcy protection well before their equity is reduced to zero. During bankruptcy, the creditors are prevented from suing the bankrupt debtor to collect what is owed them, and the obligor is allowed to continue business. At the end of the bankruptcy process, the debt is extinguished or *discharged*. There are a very few exceptions. For example, many student loans cannot be discharged through personal bankruptcy.

6.4.2 Probability of Default

In formal models, the *probability of default* is defined over a given time horizon τ , for instance, one year. Each credit has a random default time t^* . The probability of default π is the probability of the event $t^* \leq \tau$. Later in this chapter, we discuss various models and empirical approaches to estimating π .

Three points of time need to be distinguished in thinking about default modeling, default timing, and default probabilities. Incorporating these time

dimensions into default analytics is a potential source of confusion:

1. The time t from which we are viewing default: The point of view is usually “now,” that is, $t = 0$, but in some contexts we need to think about default probabilities viewed from a future date. The “point of view” or “perspective” time is important because it determines the amount of information we have. In the language of economics, the perspective time determines the information set; in the language of finance, it determines the filtration.
2. The time interval over which default probabilities are measured: If the perspective time is $t = 0$, this interval begins at the present time and ends at some future date T , with $\tau = T - 0 = T$ the length of the time interval. But it may also be a future time interval, with a beginning time T_1 and ending time T_2 , so $\tau = T_1 - T_2$. The probability of default will depend on the length of the time horizon as well as on the perspective time.
3. The time t^* at which default occurs. In modeling, this is a random variable, rather than a parameter that we choose.

6.4.3 Credit Exposure

The *exposure at default* is the amount of money the lender can potentially lose in a default. This may be a straightforward amount, such as the par or market value of a bond, or a more difficult amount to ascertain, such as the net present value (NPV) of an interest-rate swap contract.

For derivatives, exposure depends in part on whether the contract is linear. As noted in Chapter 4, linear derivatives such as futures have zero NPV at initiation, while nonlinear derivatives such as options have a positive or negative NPV at (nearly) all times.

6.4.4 Loss Given Default

If a default occurs, the creditor does not, in general, lose the entire amount of his exposure. The firm will likely still have assets that have some value. The firm may be unwound, and the assets sold off, or the firm may be reorganized, so that its assets continue to operate. Either way, there is likely to be some *recovery* for the investor that is greater than zero, but less than 100 percent of the exposure. The *loss given default* (LGD) is the amount the creditor loses in the event of a default. The two sum to the exposure:

$$\text{exposure} = \text{recovery} + \text{LGD}$$

The *recovery amount* is the part of the money owed that the creditors receive in the event of bankruptcy. It depends on debt seniority, the value

of the assets at the time of default, and general business conditions. Typical recovery rates for senior secured bank debt are in excess of 75 percent, while for junior unsecured bonds, it can be much closer to zero. In credit modeling, recovery is often conventionally assumed to be 40 percent. As we see in a moment, it can also be assumed to be a random variable, or linked to the *ex ante* default probability.

Recovery is usually expressed as a *recovery rate* R , a decimal value on $[0, 1]$:

$$R = \frac{\text{recovery}}{\text{exposure}}$$

We have

$$R = 1 - \frac{\text{LGD}}{\text{exposure}}$$

LGD and recovery are in principle *random* quantities. They are not known for certain in advance of default. This raises two very important issues. First, the uncertainty about LGD makes it more difficult to estimate credit risk. Second, because it is random, the LGD may be correlated with the default probability, adding an additional layer of modeling difficulty. There is a large body of recovery research, primarily conducted by the rating agencies, focusing on the distribution of losses if default occurs. In many applications, however, the expected LGD is treated as a *known* parameter.

The recovery rate can be defined as a percent of the current value of an equivalent risk-free bond (*recovery of Treasury*), of the market value (*recovery of market*), or of the face value (*recovery of face*) of the obligation. In modeling, assumptions about recovery are made for modeling tractability as well as to stay close to empirical behavior of recovery rates.

6.4.5 Expected Loss

The *expected loss* (EL) is the expected value of the credit loss. From a balance sheet point of view, it is the portion of the loss for which the creditor should be provisioning, that is, treating as an expense item in the income statement and accumulating as a reserve against loss on the liability side of the balance sheet.

If the only possible credit event is default, that is, we are disregarding the potential for changes in ratings (referred to as *credit migration*), then the expected loss is equal to

$$\text{EL} = \pi \times (1 - R) \times \text{exposure} = \pi \times \text{LGD}$$

If credit migration as well as default are possible, the expected loss is equal to the probability-weighted sum of the changes in value that occur under the various migration scenarios.

LGD and recovery are *conditional* expectations. The LGD has therefore been “divided” by the default probability:

$$E[\text{loss}|\text{default}] = \text{LGD} = \frac{\text{EL}}{P[\text{default}]} = \frac{\text{EL}}{\pi}$$

Thus, the LGD can be large, even if the expected loss is small.

Example 6.2 (Recovery and Loss Given Default) Suppose our exposure is \$1,000,000, and we know with certainty that the LGD is \$400,000. The recovery is then \$600,000, and we have $R = 0.60$.

Next, suppose the default probability is 1 percent. The expected loss is then $0.01 \times 400,000 = \$4,000$.

Why would an investor hold a security that has an expected loss? Because the investor believes that the credit spread more than compensates for the expected loss. Suppose an investor compares a defaultable one-year bond that pays an annual coupon of $r + z$ percent with a risk-free bond that pays a coupon of r . The coupon spread z is the compensation for default risk. If z is large enough, the expected future value of the defaultable bond will be greater than the expected future value of a riskless bond.

For simplicity, suppose the credit-risky bond can only default in exactly one year, just before it is scheduled to pay the coupon, if it defaults at all. The probability of default is π , and if it occurs, the recovery value is a decimal fraction R of the par value. There are two possible payoffs on the credit risky bond:

1. With probability $1 - \pi$, the investor receives $1 + r + z$.
2. With probability π , the investor receives R .

The future value of the risk-free bond is $1 + r$ with certainty. Therefore, if the expected value of the risky bond

$$(1 - \pi)(1 + r + z) + \pi R > 1 + r$$

the investor *may* find the credit-risky bond preferable. The expected loss is $\pi(1 - R)$. In the event of default, the unexpected loss is $(1 - \pi)(1 - R)$.

6.4.6 Credit Risk and Market Risk

A common source of confusion in discussions of fixed-income risks is the distinction between credit risk and market risk. There is no universally accepted distinction between the two. *Market risk* is the risk of economic loss from change in market prices, *including fluctuations in market prices of credit-risky securities*.

An example of a pure credit event is this: A previously AAA-rated company downgraded to AA, but no change in AAA spreads or in risk-free rates. An example of a pure market event is a widening spread between AAA and risk-free rates, or a rise in risk-free rates.

There is some ambiguity in the distinction between credit and market risk. Spreads may change even in the absence of a default or migration, because the likelihood of the event, as perceived by the market has changed, or because of a change in risk premiums. But migration generally results in a change in market value, and may oblige the lender to write the credit up or down, that is, record a higher or lower value in its balance sheet. A change in market perception of a firm's credit quality, even if it does not result in migration, may cause a change in spreads. In the credit risk context, this is called *mark-to-market* risk, as opposed to default risk.

6.5 ASSESSING CREDITWORTHINESS

Lenders and investors in credit-risky securities need a way to assess the creditworthiness of borrowers. This takes place in a variety of ways. All are “quantitative,” at least in the sense of relying on balance sheet and other business data of the firm as well as data on the state of the economy or the firm's industry. Some approaches are based on more formal mathematical or statistical modeling.

6.5.1 Credit Ratings and Rating Migration

A *credit rating* is an alphanumeric grade that summarizes the creditworthiness of a security or a corporate entity. Credit ratings are generally assigned by *credit rating agencies* that specialize in credit assessment. The most prominent in the United States are Standard and Poor's (S&P), Moody's, Fitch Ratings, and Duff and Phelps. Along with a handful of others, they have been granted special recognition by the Securities and Exchange Commission (SEC). Their ratings are used as part of the bank and securities markets' regulatory system, though the 2010 Dodd-Frank Act mandates a reduction of the regulatory role of ratings in the United States (see

TABLE 6.1 The S&P and Moody's Long-Term Rating Systems

Investment grade		Speculative grade	
S&P	Moody's	S&P	Moody's
AAA	Aaa	BB+	Ba1
AA+	Aa1	BB	Ba2
AA	Aa2	BB-	Ba3
AA-	Aa3	B+	B1
A+	A1	B	B2
A	A2	B-	B3
A-	A3	CCC+	Caa1
BBB+	Baa1	CCC	Caa2
BBB	Baa2	CCC-	Caa3
BBB-	Baa3	CC	Ca
		C	C

These ratings apply to debt with an original term to maturity of one year or longer. There is a comparable system applied to short-term securities such as commercial paper.

Chapter 14). For example, a high rating from one of these *Nationally Recognized Statistical Rating Organization* (NRSRO) can reduce the compliance burden in issuing a security or the regulatory capital requirement for owning it. Table 6.1 lists the rating categories of the two largest NRSROs.

Rating agencies also assess the probability of default for companies based on their letter ratings. These probabilities can be compared with the annual rates at which firms with different ratings actually default, plotted in Figure 6.1. Rating agencies assess not only the probability of default, but also of rating migration, or change in letter rating, which occurs when one or more of the rating revises the rating of a firm (or a government) or its debt securities. These probability estimates are summarized in *transition matrices*, which show the estimated likelihood of a company with any starting rating ending a period, say, one year, with a different rating or in default.

Typically, the diagonal elements in a transition matrix, which show the probability of finishing the year with an unchanged rating, are the largest elements (see Table 6.2). Also typically, the probability of ending in default rises monotonically as the letter rating quality falls. Finally, note that there is no transition *from* default to another rating; default is a terminal state.

The ratings business provides a good example of conflicts of interest in credit markets. Since the advent of photocopy technology, ratings agencies have generally been compensated for ratings by bond issuers rather than by sale of ratings data to investors. In this so-called *issuer-pays model*, a

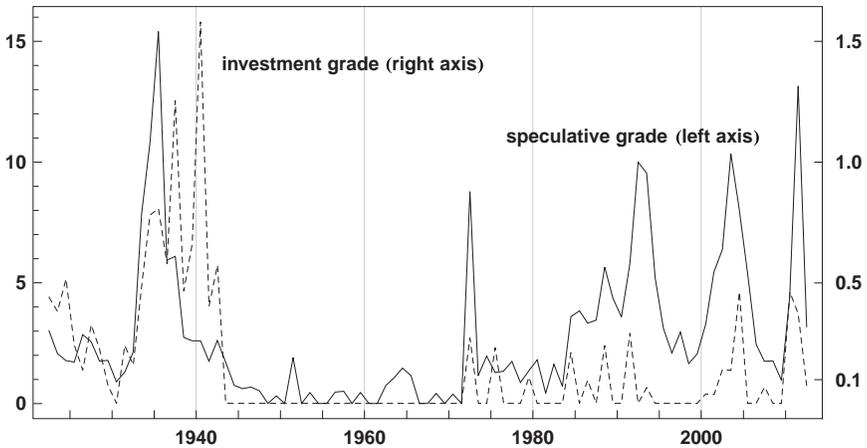


FIGURE 6.1 Default Rates 1920–2010
Issuer-weighted default rates, that is, the fraction of rated issuers defaulting each year, in percent. For recent years, Moody’s also reports *volume-weighted* default rates, that is, rates based on the notional value outstanding of each issuer’s debt. Solid line plots speculative defaults and dashed line plots investment grade defaults. *Source:* Moody’s Investor Service (2011), Exhibit 31.

potential conflict arises between investors in rated bonds, who expect ratings to be based on objective methodologies, and issuers of bonds. This conflict has been identified by some observers as having played a material role in causing the subprime crisis. Structured product ratings are said to have been more permissive than appropriate because of rating agencies’ competition for issuers’ business. We discuss this problem further in Chapter 15.

TABLE 6.2 One-Year Ratings Transition Matrix

	AAA	AA	A	BBB	BB	B	CCC/C	Default
AAA	91.42	7.92	0.51	0.09	0.06	0.00	0.00	0.00
AA	0.61	90.68	7.91	0.61	0.05	0.11	0.02	0.01
A	0.05	1.99	91.43	5.86	0.43	0.16	0.03	0.04
BBB	0.02	0.17	4.08	89.94	4.55	0.79	0.18	0.27
BB	0.04	0.05	0.27	5.79	83.61	8.06	0.99	1.20
B	0.00	0.06	0.22	0.35	6.21	82.49	4.76	5.91
CCC/C	0.00	0.00	0.32	0.48	1.45	12.63	54.71	30.41

One-year ratings transitions, based on ratings history of 11,605 global S&P-rated companies 1981–2005, adjusted for withdrawn ratings. *Source:* Standard & Poor’s (2006), p. 17.

6.5.2 Internal Ratings

Ratings provide third-party assessments of the credit quality of a borrower. Many firms also carry out their own assessments of credit quality. Such assessments are used in making decisions about whether to buy a debt security or extend credit, and are in some ways similar to the work of equity analysts. Credit analysis includes detailed attention to the legal documentation accompanying debt contracts.

In larger firms, credit analysis may be formalized into *internal ratings*. In such firms, internal ratings can play a role in setting regulatory capital, as we see in Chapter 15. Internal credit assessments may use quantitative techniques such as *credit scoring*, in which a numerical ranking is a function of balance-sheet or other firm data.

6.5.3 Credit Risk Models

A widely used approach to estimating credit risk is via formal credit risk models. There are two broad approaches to assessing the credit risk of a single company or other issuer of debt.

Reduced-form models are, in a sense, not risk models at all. Rather than producing, as a model output, an estimate of the default probability or LGD, they take these quantities as inputs. The models are most often used to simulate default times as one step in portfolio credit risk measurement. We study how to estimate such models in the next chapter, and apply them in Chapters 8 and 9.

Structural models derive measures of credit risk from fundamental data, particularly elements of the firm's balance sheet such as the volumes of assets and debt.

A third type of model, *factor models*, falls somewhere in between; they are "informed" by company, industry, and economy-wide fundamentals, but in a highly schematized way that lends itself to portfolio risk modeling.

Credit risk models may also be distinguished by whether they take into account credit migration, in which case the model is said to operate in *migration mode*, or only default, in which case the model is said to operate in *default mode*.

6.6 COUNTERPARTY RISK

Counterparty risk is a special form of credit risk. It arises whenever *two* conditions both need to be fulfilled in order for a market participant to profit from an investment: The investment must be profitable, and another party must fulfil a contractual obligation to him.

Every credit transaction is, of course, a contractual relationship between two counterparties. But the term “counterparty risk” is not typically applied to situations in which credit risk is “one-way,” that is, in which only one of the counterparties can incur a loss due to default or impairment. Only the lender bears the credit risk of a loan or bond, not the borrower. Counterparty risk is “two-way,” and it is often not clear who is the borrower and who the lender. It typically arises in two contexts:

OTC derivatives trading. Every derivatives contract has two sides. Depending on how asset prices evolve, either party may end up owing money to the other. Just as in their impact on leverage, which we discuss in Chapter 12, there is a difference in this regard between swaps and other linear derivatives on the one hand, and options on the other. Swaps can have a positive net present value for either counterparty, so both counterparties are potentially exposed to counterparty risk. Only the purchaser of an option is exposed to counterparty risk, since his obligation to the option-selling counterparty is limited to payment of the premium.

The issue of counterparty risk arises particularly for OTC derivatives, since these are purely bilateral contracts between two private counterparties. Exchange-traded futures and options have far less counterparty risk, since the exchanges interpose a *clearinghouse* as a counterparty to every trade. The clearinghouses historically have been well capitalized, since each clearing member is obliged to put up *clearing margin*, and clearinghouse failure is quite rare. No individual customer is exposed to another individual counterparty. Rather, each is matched up against the clearinghouse.

Title VII of the Dodd-Frank Act of 2011 (see Chapter 15) requires exchange clearing of OTC derivatives, but also permits many exemptions. Rules specifying what types of OTC derivatives and what types of counterparties will be subject to “mandatory clearing” have not yet been adopted.

Brokerage relationships. The financing of positions has been increasingly closely related to intermediation and trading services, particularly with the increase in hedge fund trading and concomitant expansion of the prime brokerage business (see Section 12.2). Brokerage clients often assume significant counterparty risk to their brokers. Prior to the subprime crisis, the presumption had been that the broker, but not the client, had credit exposure. The crisis experience, in which clients experienced losses on exposures to failed brokerages such as Lehman, but broker losses on client exposures remained comparatively rare, changed this presumption.

6.6.1 Netting and Clearinghouses

It is very natural in derivatives trading for contracts to proliferate. Consider a single forward commodities contract with a single maturity date. A dealer might buy and sell various quantities of the same contract in the course of a trading day. Two dealers might find they have traded with one another at various prices, obliging each one to deliver a large amount of the commodity to the other in exchange for a large amount of money. *Netting* their contracts is an obvious efficiency gain for both dealers. It is easy to carry out *bilateral netting* by having one of the dealers owe the other only a net amount of the commodity, against a net amount of money determined by the prices at which the contracts were struck. For example, if traders A and B made three trades with each other, each obliging one to deliver 100,000 bushels of wheat to the other, with trader A having gone long at \$100 and \$110, and short at \$105, either of his long contracts can be netted against the short. If the \$100 contract is netted out, trader A will receive \$5 from B, and only remain long the \$110 contract. If the \$110 contract is netted out, trader A will pay \$5 to B, and only remain long the \$100 contract.

Multilateral netting, or netting among many counterparties, is far more complicated, but brings the same efficiencies. Netting, as well as settlement services, are among the reasons clearinghouses were introduced in U.S. futures markets in the latter half of the nineteenth century. Clearinghouses and clearing associations were also important features of U.S. banking in the nineteenth and early twentieth centuries.

6.6.2 Measuring Counterparty Risk for Derivatives Positions

The counterparty risk of a derivatives position is closely related to, but distinct from, its market risk. Market risk is the risk that the market factors underlying the position will move against the trader, so that the NPV of the position is negative. Counterparty risk is the risk that the NPV is positive, but the counterparty fails to perform, so that no gain is reaped. Counterparty risk is thus a conditional risk. In order to realize a counterparty loss, the NPV of the contract must be positive. The amount at risk from a counterparty default will be quite close to the amount at risk from market moves. Adjustments may be made to the fair NPV to account for counterparty risk. Such an adjustment is called a *counterparty valuation adjustment* (CVA).

A key mechanism in protecting against counterparty risk in derivatives trading is *margin*, a form of collateralization in which one or both counterparties set aside a cash amount to cover losses in the event of counterparty default. We will return to the subject in Chapter 12's discussion of

leverage in securities markets. Initial margin, as noted above, is the amount of cash collateral set aside when a trade is initiated. Initial margin tends to be relatively small on most derivatives transactions, though, as we see in Chapter 14, margins generally rose rapidly during the subprime crisis. The bulk of margin on derivatives, once enough time has elapsed since initiation, is therefore the NPV of swap contracts. Generally, as prices and NPVs fluctuate, changes in NPV are remitted daily as variation margin, that is, as increments to collateral rather than as settlement, by that day's losing counterparty to the gaining counterparty. This is a general practice, however, rather than an ironclad rule.

In some cases, initial margin on swaps can be relatively large. Examples are CDS, CDS indexes, and structured-credit CDS such as the ABX and CMBX—credit derivatives that provide exposure to subprime mortgage and commercial real estate loans—and synthetic CDO tranches, when these products are trading in *points upfront*, or “points up.” In these cases, the counterparty selling protection receives an initial payment equal to the points up. The points upfront serve the function of initial margin. We further discuss point upfront and other CDS refinements in Chapter 7.

The treatment of the counterparty exposure and the scope for netting is often governed by legal agreements called *master agreements* or “ISDAs” (after the International Swap Dealer Association that drafted the templates for these contracts). As is the case for margin on cash securities, the net exposure at any moment is determined by the effect of the netting agreements and the collateral exchanged up to that point in time.

Derivatives dealers have historically had a privileged role in this process, much like that of brokers vis-à-vis their clients. They may be in a position to demand initial margin from their customers on both the long and short protection sides of a swap. The dealer, but not the client, might reasonably consider making CVAs to OTC derivatives' position values. This presumption proved faulty during the Lehman bankruptcy. In the wake of the Lehman bankruptcy, many of its swap counterparties suffered losses because Lehman held margin. Any collateral Lehman held became part of the bankruptcy estate, and its counterparties became general unsecured creditors of Lehman. In some cases, counterparties of Lehman such as hedge funds suffered losses on derivatives contracts in which Lehman held margin from them although the NPV of the contract was positive for the hedge fund. The counterparty would then become an unsecured general creditor of the Lehman bankruptcy estate both for the margin it had paid to Lehman and the claim arising from the derivatives contract. The details of any netting depended in part on the netting and cross-margining agreements in place and with which entities of the bankrupt broker these transactions had been carried out.

6.6.3 Double Default Risk

Double default risk arises from the possibility that the counterparty of a credit derivative such as a CDS that protects you against the default of a third party will default at the same time as that third party. This type of risk is both a counterparty risk and a form of credit correlation risk.

An important example of double default risk was prominent during the subprime crisis. It involved the counterparty exposures of American International Group (AIG), an internationally active insurance company with large property and casualty as well as life insurance businesses. AIG was a well-capitalized firm with an AAA rating, and was therefore an attractive counterparty for firms, especially banks and broker-dealers, looking to buy protection on a wide range of credit exposures via CDS. Because of AIG's strong capital base, the financial intermediaries were able to hedge their credit exposures while incurring only minimal apparent counterparty risk.

AIG's protection-selling activity was housed mainly within a subsidiary called AIG Financial Products (AIGFP), and one of its focuses was the most highly-rated tranches of structured credit products, particularly mortgage-backed securities. It was implicitly long a very large quantity of senior CDO tranches. Its counterparties had been long these exposures, and had now substituted counterparty risk for their initial long exposures. If AIGFP proved unable to meet its obligations under the CDS contracts at the same time that the bonds they owned suffered a material impairment, they would lose their hedges. This scenario became far likelier during the subprime crisis.

Another, similar, example of double default risk is the "wrap" business carried on by *monoline insurers*. In this business, an insurance company with a strong capital base and an AAA rating guarantees payment on a debt security issued by a company, another customer, or a municipality in exchange for an annual fee. The guarantee, called the "wrap," raises the creditworthiness of the bond sufficiently for rating agencies to grant it a higher rating. A number of monoline insurers suffered large losses during the subprime crisis, as the bonds they had guaranteed suffered credit downgrades or impairments, increasing the risk they would not have sufficient capital to make good on the guarantees they had written.

6.6.4 Custodial Risk

Securities in the contemporary financial system almost invariably exist in electronic, rather than paper form. *Bearer securities* have virtually disappeared. Even so, securities have to be "kept somewhere," dividends and interest have to be collected, and the securities have to be available for delivery if they are sold, lent, or transferred to another account. These services are called *custodial services*.

In margin lending and prime brokerage relationships, customers of large intermediaries such as banks and broker-dealers typically keep much of their securities in *margin accounts*, where the intermediary providing financing also has custody of the securities. These securities are collateral against the financing, and are said to be in “street name.” In the event the borrower defaults on the financing, securities in street name can be sold immediately in order to protect the lender against credit losses.

Many retail brokerage accounts are also margin accounts, and are subject to much the same risks. Securities in cash or nonmargin accounts are in customer name. If the broker defaults, the retail customer’s ownership of the security is not called into question. Securities in margin accounts, that is, in street name, are subject to custodial risk.

Many brokers perform custodial as well as credit intermediation services for their clients. The customers keep securities in custody with the broker that may at times be pledged as collateral and at other times be “in the box.” The broker can *rehypothecate* pledged securities to fund its own operations. That is, the broker can itself employ, as collateral to borrow money, the very securities it holds as collateral against money it has lent to customers. The customer securities can also be lent to generate fee and interest income. The rules governing rehypothecation differ in important ways internationally.

These arrangements pose a counterparty risk to the broker’s customers that was not widely noted until it was realized during the Lehman bankruptcy in September 2008. Customers of the firm’s U.K. subsidiary were particularly badly situated, as even their unpledged assets were typically not segregated, but might be subject to rehypothecation. If the customer’s securities are rehypothecated by the broker, the customer becomes a creditor of the broker. If the broker files for bankruptcy protection, the customer might not receive the securities back, but instead be treated as an unsecured lender of the broker-dealer. The amount of the unsecured exposure is then equal to the amount arrived at by netting the customer’s margin across all exposures, including equity in margin loans and the positive NPV in his derivatives contracts.

6.6.5 Mitigation of Counterparty Risk

The ideal approach to mitigating counterparty risk would be to accurately measure exposures, maintain assessments of the credit condition of counterparty, maintain a diverse set of counterparties, and promptly limit exposure to weaker counterparties. For large firms with many counterparties, this can be a complex undertaking requiring considerable staff and systems investment. Exposure to specific counterparties can be reduced by reducing the volume of OTC contracts with them or by increasing the amount of

collateral held against them. One difficulty in managing counterparty risk in CDS trades is that typically CDS trades are not netted or cancelled. Rather, an offsetting trade with another counterparty is initiated, a process called *novation*.

One tool for limiting counterparty risk is therefore CDS *compression*, or reducing the volume of redundant contracts with the same reference entity. A firm might put on long and short protection positions on the same name as it varies the size of its position. Compression reduces the set of CDS to a single net long or short position, thus eliminating a certain amount of nominal exposure and counterparty risk. There are many practical difficulties in carrying out compression trades, since in general the contracts will not be identical as to counterparty, premium, and maturity.

Individual firms have less control over systemic counterparty risk, that is, the risk that all counterparties become weaker in a systemic risk event, other than to keep overall counterparty exposure low.

6.7 THE MERTON MODEL

In the rest of this and the next chapter, we will focus on *single-obligor credit risk models*, that is, models of a single issuer of debt obligations. In Chapter 8, we will extend our discussion to *portfolio credit risk models*, which treat the credit risk of portfolios containing obligations of several obligors. The specific additional problem that portfolio credit models deal with is that of the correlation between credit events of different obligors.

In structural credit risk models, the evolution of the firm's balance sheet drives credit risk. The approach is sometimes called the *Merton model*. It applies the Black-Scholes option pricing model to value credit-risky corporate debt.

The setup for a simple variety of the Merton structural model combines a set of assumptions we need so that we can apply the Black-Scholes option pricing model, with a set of additional assumptions that tailor the model to our credit-risk valuation context. We'll set out the model assumptions, and then use the model to derive the firm's default probability:

- The value of the firm's assets A_t is assumed to follow a geometric Brownian motion:

$$dA_t = \mu A_t dt + \sigma_A A_t dW_t$$

Two of the parameters, the market value of the assets A_t and the expected return μ , are related to one another. In equilibrium, if r is the

riskless continuously compounded interest rate for the same maturity as the firm's debt, the market's assessment of the asset value will be such that, given investors' risk appetites and the distribution of returns they expect, the risk premium $\mu - r$ on the assets is a sufficient, but not too generous, reward.

- The balance sheet of the firm is simple:

$$A_t = E_t + D_t$$

The debt consists entirely of one issue, a zero-coupon bond with a nominal payment of D , maturing at time T . The notation D , with no subscript, is a constant referring to the par value of the debt. The notation D_t , with a time subscript, refers to the value of the debt at that point in time. In reality, most firms with tradeable debt have different types of issues, with different maturities and different degrees of seniority. In this model, the firm can default only on the maturity date of the bond. Similarly, we also assume the entire equity consists of common shares.

- Limited liability holds, so if the equity is wiped out, the debtholders have no recourse to any other assets.
- Contracts are strictly enforced, so the equity owners cannot extract any value from the firm until the debtholders are paid in full. In reality, when a firm is expected to reorganize rather than liquidate in bankruptcy, there is usually a negotiation around how the remaining value of the firm is distributed to debtholders and equity owners, and all classes may have to lose a little in order to maximize the value with which they all emerge from bankruptcy.
- There is trading in the assets of the firm, not just in its equity and debt securities, and it is possible to establish both long and short positions. This rules out intangible assets such as goodwill.
- The remaining assumptions are required to “enforce” limited liability: The firm can default only on the maturity date of the bond. There are no cash flows prior to the maturity of the debt; in particular, there are no dividends.

Default takes place if and only if, at time T , the firm's assets are less than its debt repayment obligation:

$$A_T < D$$

The probability of default over the next T years is therefore the probability that the Brownian motion A_t hits the level D within the interval $(0, T_0)$. The quantity $A_T - D$ is called the *distance to default*. In this setup, we can view

both the debt and equity securities as European options on the value of the firm's assets, maturing at the same time T as the firm's zero-coupon debt. We can value the options using the Black-Scholes formulas of Appendix A.3. The model will then help us obtain estimates of the probability of default and other default and recovery parameters.

However, in contrast to the way in which option pricing models are usually applied in finance, we are interested, in the credit risk context, in *both* risk-neutral and “physical” or true quantities, so we will have to be careful to identify clearly which one we are talking about, and we have to use the correct formulas.

We are also making some unrealistic assumptions about what parameters we know. In particular, it is unrealistic to imagine we will know at every point in time what exactly is the current market value of the firm's assets. It is even less likely we will know their volatility.

With these assumptions, we can use option-pricing theory to compute the equity and debt values. These in turn will lead us to the default probability:

Equity value of the firm. For expositional purposes, we treat the current value of the firm's assets A_t and the volatility of the assets σ_A as known quantities. The equity can then be treated as a call option on the assets of the firm A_t with an exercise price equal to the face value of the debt D . If, at the maturity date T of the bond, the asset value A_T exceeds the nominal value of the debt D , the firm will pay the debt. If, in contrast, we have $A_T < D$, the owners of the firm will be wiped out, and the assets will not suffice to pay the debt timely and in full.

The equity value at maturity is therefore

$$E_T = \max(A_T - D, 0)$$

Denoting $\tau = T - t$, and the Black-Scholes value of a τ -year European call struck at D by $v(A_t, D, \tau, \sigma_T, r, 0)$ —remember that we assume no dividends—can value the firm's equity as a call on its assets, struck at the value of the debt:

$$E_t = v(A_t, D, \tau, \sigma, r, 0)$$

Market value of the debt. We can also apply option theory from the point of view of the lenders. We can treat the debt of the firm as a portfolio consisting of a risk-free bond with par value D plus a short position in a put option on the assets of the firm A_t with exercise

price D . In other words, if the bondholders received, as a gift, a put option on the firm's assets struck at the par value of the debt, with, crucially, no counterparty risk, they would be indifferent between that portfolio and a risk-free bond.

The present value of the risk-free bond is $De^{-r\tau}$. We can state the future value of the debt as

$$D_T = D - \max(D - A_T, 0)$$

Denoting the Black-Scholes value of a European put $w(A_t, D, \tau, \sigma_A, r, 0)$, we have the current value of the bond, as adjusted for risk by the market:

$$D_t = e^{-r\tau} D - w(A_t, D, \tau, \sigma_A, r, 0)$$

Firm balance sheet. The firm's balance sheet now also expresses put-call parity:

$$\begin{aligned} A_t &= E_t + D_t \\ &= v(A_t, D, \tau, \sigma_A, r, 0) + e^{-r\tau} D - w(A_t, D, \tau, \sigma_A, r, 0) \end{aligned}$$

This means that the firm's assets are equal in value to a portfolio consisting of a risk-free discount bond in the nominal amount of the firm's debt, plus a long call and a short put, each struck at the nominal value of the debt.

Leverage. As a simple consequence of the balance sheet quantities, we can also compute balance sheet ratios. The leverage of the firm, expressed as the equity ratio at market prices, is

$$1 - \frac{e^{-r\tau} D - w(A_t, D, \tau, \sigma_A, r, 0)}{A_t}$$

We could specify the leverage ratio and the principal amount of the debt, and deduce the level of assets. In fact, we can specify any two of the three quantities—assets, debt principal, or leverage—and solve for the remaining quantity.

Default probabilities. The probability of default is identical to the probability of exercise of the put and call options we have been describing. However, we now have to distinguish between the true, actuarial or “physical” probability of exercise, and the risk-neutral probability.

The actuarial probability of default can be computed from the stochastic process followed by the firm's assets, provided we know or have a plausible estimate of the return on assets μ . The firm's assets are then lognormally distributed with parameters μ and σ_A . The probability of default is equal to

$$P[A_T < D] = \Phi \left[-\frac{\log\left(\frac{A_t}{D}\right) + \left(\mu - \frac{1}{2}\sigma_A^2\right)\tau}{\sigma_A\sqrt{\tau}} \right]$$

The default probability estimated this way has in common with the rating agencies' estimates that it is an estimate of the true probability. It is different from agencies' estimates in that it is not a "through the cycle" rating, but rather a short-term estimate over the term of the debt.

In the Black-Scholes model, the risk-neutral probability of exercise is given by

$$\frac{\partial}{\partial D} v(A_t, D, \tau, \sigma, r, 0) = \Phi \left[-\frac{\log\left(\frac{A_t}{D}\right) + \left(r - \frac{1}{2}\sigma_A^2\right)\tau}{\sigma_A\sqrt{\tau}} \right]$$

Since this is a non-dividend paying stock (an inessential assumption), we can switch from true to risk-neutral by changing the asset return to the risk-free rate.

Credit spread. We now have enough information to compute the yield to maturity and the credit spread of the firm's debt. The yield to maturity y_t solves

$$D_t e^{y_t \tau} = D$$

Substituting the current market value of the debt, we have

$$[e^{-r\tau} D - w(A_t, D, \tau, \sigma_A, r, 0)] e^{y_t \tau} = D$$

so after taking logarithms, we have

$$y_t = \frac{1}{\tau} \log[(1 - e^{-r\tau})D + w(A_t, D, \tau, \sigma_A, r, 0)]$$

The credit spread is

$$y_t - r = \frac{1}{\tau} \log[(1 - e^{-r\tau})D + w(A_t, D, \tau, \sigma_A, r, 0)] - r$$

Loss given default. The LGD in the Merton model is a *random* quantity. It depends on how far short of the par value of the debt the firm's assets fall on the maturity date, if the firm defaults at all.

The default loss is equal to $\max(D - A_T, 0)$, which, of course, is the value *at expiry* of the “virtual” put option we are using to price the debt. However, the actuarial expected default loss is *not* equal to the current market value of the put option. The put option value is the risk-neutral value of protection against default in the Black-Scholes world of perfect and costless hedging via trading in the firm's assets. In computing its expected value, we use the risk-neutral, not the physical probability distribution that takes the growth rate of the assets μ into account. The value of the put option is greater than the actuarial expected loss, because there is compensation to the put writer for taking on the risk as well as the expected cost of providing the default protection.

To get the actuarial expected loss, we need to compute the expected value of $\max(D - A_T, 0)$, conditional on $A_T < D$

$$\text{expected loss} = E[A_T | A_T < D] \mathbf{P}[A_T < D] - DP[A_T < D]$$

Fortunately, we don't need to compute this; rather, we can use the Black-Scholes formula, but with μ in place of r and taking the future rather than the present value of the payoff:

$$\begin{aligned} E[\text{LGD}] &= e^{r\tau} w(A_t, D, \tau, \sigma_A, \mu, 0) \\ &= D\Phi\left[-\frac{\log\left(\frac{A_t}{D}\right) + (\mu - \frac{1}{2}\sigma_A^2)\tau}{\sigma_A\sqrt{\tau}}\right] \\ &\quad - e^{r\tau} A_t \Phi\left[-\frac{\log\left(\frac{A_t}{D}\right) + (\mu + \frac{1}{2}\sigma_A^2)\tau}{\sigma_A\sqrt{\tau}}\right] \end{aligned}$$

Dividing this by the actuarial default probability conditions the expected loss on the occurrence of default and gives us the expected LGD: that is, the amount the debtholder can expect to lose in the event of default.

$$\text{expected LGD} = \frac{e^{r\tau} w(A_t, D, \tau, \sigma_A, \mu, 0)}{\mathbf{P}[A_T < D]}$$

We can use this expression to compute the recovery rate on the debt, or to be precise, the *expected* recovery rate. The LGD is a conditional quantity, that is, the loss *if* default occurs. So to find it, we divide the expected value of the loss by the probability of its occurrence to get the conditional expected value of default loss.

We can then state the expected value of the recovery rate as

$$R = 1 - \frac{1}{D} \frac{e^{r\tau} w(A_t, D, \tau, \sigma_A, \mu, 0)}{\mathbf{P}[A_T < D]}$$

Figure 6.2 illustrates the idea behind the Merton model using the parameters of Example 6.3.

Example 6.3 (Merton Model) We apply the model to a firm that has an asset value of \$140. We'll assume the firm's sole debt issue is a bond, with one 6 percent coupon left, to be paid in one year along with the principal at the maturity of the bond. This is effectively a zero-coupon bond with a par value of 106, but looking at the debt this way conveniently centers current debt prices near 100 percent of par.

The parameters for the example are:

Firm's current asset value	A_t	140
Nominal value of debt	D	100
Coupon on debt	c	0.06
Asset volatility	σ_A	0.25
Debt maturity in years	$\tau = T - t$	1
Risk-free rate	r	0.05
Return on firm's assets	a	0.10

The market value of the debt is equivalent to a portfolio consisting of the present value of the debt, discounted at the risk-free rate, plus a short put on the assets. The put has a value of

$$w(A_t, D, \tau, \sigma, r, 0) = w(140, 106, 1, 0.25, 0.05, 0) = 1.3088$$

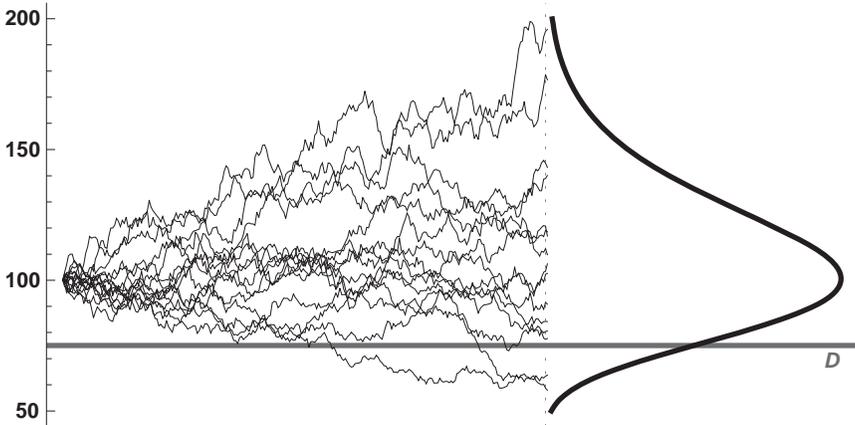


FIGURE 6.2 Merton Model

To the left, the graph displays 15 daily-frequency sample paths of the geometric Brownian motion process.

$$dA_t = \mu A_t dt + \sigma_A A_t dW_t$$

with parameters as stated in the example. To the right, the graph displays the probability density of the firm's assets on the maturity date of the debt. The grid line represents the debt's par value D .

The bonds must be priced so as to incorporate the cost of insuring the bondholders against a default, so the current market value of the debt is:

$$De^{-r\tau} - w(A_t, D, \tau, \sigma, r, 0) = 106 \times 0.9512 - 1.3088 = 99.5215$$

We can now calculate the fair market value of the firm's equity:

$$A_t - [De^{-r\tau} - p(A_t, D, \tau, \sigma, r, 0)] = 140 - 99.5215 = 40.4785$$

which put-call parity tells us is also equal to the value today of a European call on the future residual value of the firm's assets over and above the par value of the debt. The economic balance sheet of the firm is:

Assets	Liabilities
Value of the firm $A_t = 140$	Equity $E_t = 40.4785$
	Debt $D_t = 99.5215$

So the leverage (defined as the reciprocal of the equity ratio) is 3.4586.
The physical probability of default is

$$\Phi \left[-\frac{\log\left(\frac{140}{106}\right) + (0.05 - \frac{1}{2}0.25^2) \cdot 1}{0.25 \cdot \sqrt{1}} \right] = 0.0826$$

or 8.3 percent, while the risk-neutral probability is somewhat higher at

$$\Phi \left[-\frac{\log\left(\frac{140}{106}\right) + (0.10 - \frac{1}{2}0.25^2) \cdot 1}{0.25 \cdot \sqrt{1}} \right] = 0.1175$$

or 11.8 percent.

The expected value of default losses is represented by the future value of the “actuarial default put,”

$$\begin{aligned} e^{r\tau} w(A_t, D, \tau, \sigma, \mu, 0) &= e^{0.05} w(140, 106, 1, 0.25, 0.05, 0) \\ &= 1.05127 \times 0.8199 \end{aligned}$$

This value, of course, is lower than the risk-neutral default put value.

To arrive at the LGD and recovery rate, we convert the loss to a conditional expected loss in the event of default

$$\frac{e^{r\tau} w(A_t, D, \tau, \sigma, \mu, 0)}{\mathbf{P}[A_T < D]} = \frac{1.05127 \times 0.8575}{0.0826} = 10.4349$$

The conditional expected recovery rate is

$$1 - \frac{1}{106} 10.4349 = 0.9016$$

We can now summarize our results:

Discount factor	0.9512
Exposure	106.0000
Present value of exposure	100.8303
Put value	1.3088
Current market value of debt	99.5215
Current market price of debt	0.9952
Current market value of equity	40.4785
Leverage: Equity ratio	0.2891
Leverage: Reciprocal of equity ratio	3.4586
Actuarial default probability	0.0826
Risk-neutral default probability	0.1175
Fair yield to maturity	0.0631
Fair credit spread	0.0131
Actuarial expected loss	0.8619
Risk-neutral expected loss	1.3759
Expected LGD	10.4349
Expected recovery	95.5651
Expected recovery rate	0.9016

One drawback of the Merton model becomes clearer from this example. Leverage of nearly $3\frac{1}{2}$ times is quite high, at least for a nonfinancial company. Yet the default probability is low, at 8.3 percent, and the recovery rate is very high at 90 percent. Companies with such high leverage would typically have ratings that imply higher default probabilities, and recovery would be expected to be considerably lower.

We've laid out the basic structure of the Merton model, and along the way we have flagged some of its features that lack realism. The model has been adapted in commercial applications, particularly Moody's *KMV* and RiskMetrics' *CreditGrades*. These applications also attempt to address the two main respects in which the basic Merton model differs "too much" from reality. First, the capital structure of a typical firm is much more complex, particularly in its debt structure, than we have assumed. Second, in contrast to other applications of Black-Scholes, the underlying asset value A_t and the volatility σ_A , whether implied or historical, are not directly observable.

6.8 CREDIT FACTOR MODELS

Factor models can be seen as a type of structural model, since they try to relate the risk of credit loss to fundamental economic quantities. In contrast to other structural models, however, the fundamental factors have their

impact directly on asset returns, rather than working through the elements of the firm's balance sheet.

A simple but widely used type is the *single-factor model*. The model is designed to represent the main motivating idea of the Merton model—a random asset value, below which the firm defaults—while lending itself well to portfolio analytics.

The horizon of the model is fixed at a future date $T = t + \tau$. The logarithmic asset return is

$$a_T = \log\left(\frac{A_T - A_t}{A_t}\right)$$

so the event that $A_T < D$ is identical to the event that

$$a_T < \log\left(\frac{D - A_t}{A_t}\right) = -\log\left(\frac{E_t}{A_t}\right) = -\log(\text{equity ratio})$$

or, in terms of financial ratios, that the asset return is negative and greater in absolute value than the initial equity ratio.² The horizon is constant, so to keep the notation from cluttering, we suppress the time subscript from here on.

The firm's asset return is represented as a function of two random variables: the return on a "market factor" m that captures the correlation between default and the general state of the economy, and a shock ϵ_t capturing idiosyncratic risk. However, the fundamental factor is not explicitly modeled: It is *latent*, meaning that its impact is modeled indirectly via the model parameters. We can write the model as:

$$a_T = \beta m + \sqrt{1 - \beta^2} \epsilon$$

We assume that m and ϵ are standard normal variates, and are not correlated with one another:

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

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

$$\text{Cov}[m, \epsilon] = 0$$

²We have expressed the model in terms of the asset return. Some presentations in the literature are done in terms of the level of asset value. Because of the way the model is set up, this doesn't matter. As long as the statistical behavior of the latent and idiosyncratic factors are as we have specified, the results are the same.

Under these assumptions, a is a standard normal variate. Since both the market factor and the idiosyncratic shocks are assumed to have unit variance (since they're standard normals), the beta of the firm's asset return to the market factor is equal to β :

$$\begin{aligned} \mathbf{E}[a_T] &= 0 \\ \text{Var}[a_T] &= \beta^2 + 1 - \beta^2 = 1 \end{aligned}$$

The role of σ_A in the Merton model is taken in the single-factor model by $\sqrt{1 - \beta^2}$. Figure 6.3 illustrates the role of the latent factor, the market index, and of the correlation, in driving asset returns.

The β in this model is related, but not identical, to a firm's equity beta. In this model, β captures comovement with an unobservable index of market conditions, rather than with an observable stock index. Also, β here relates the firm's asset, rather than equity, return to the market index. Still, they are related, because cyclical firms—that is, firms that do well when the economy overall is doing well—have high equity betas and will also have a higher β in the credit single-factor model. Defensive stocks will have low equity betas and low β .

Typically, in applications, we have an estimate of the default probability π , derived externally to the model, either market-implied or fundamentals/ratings-based. We can then use this probability to “calibrate” the model, that is, to determine values for the parameters. Rather than an output, as in the Merton model, the default probability is an input in the single-factor model. The default probability calibrates the default threshold asset value. Since under our assumptions a_T is a standard normal variate, the following holds

$$\pi = \mathbf{P}[a_T \leq k] \Leftrightarrow k = \Phi^{-1}(\pi)$$

where $\Phi^{-1}(\cdot)$ is the quantile function of the standard normal distribution and $\Phi^{-1}(\pi)$ is the π -th quantile of a_T (see Appendix A.2). The firm defaults if $a_T \leq k$, the logarithmic distance to the default asset value, measured in standard deviations. The asset return could bump into the default threshold via an infinite number of combinations of market factor realizations and idiosyncratic shocks, but the probability is π .

Example 6.4 (Single-Factor Model) For what value of β do systematic and idiosyncratic risk contribute equally to total risk of a credit, as measured by

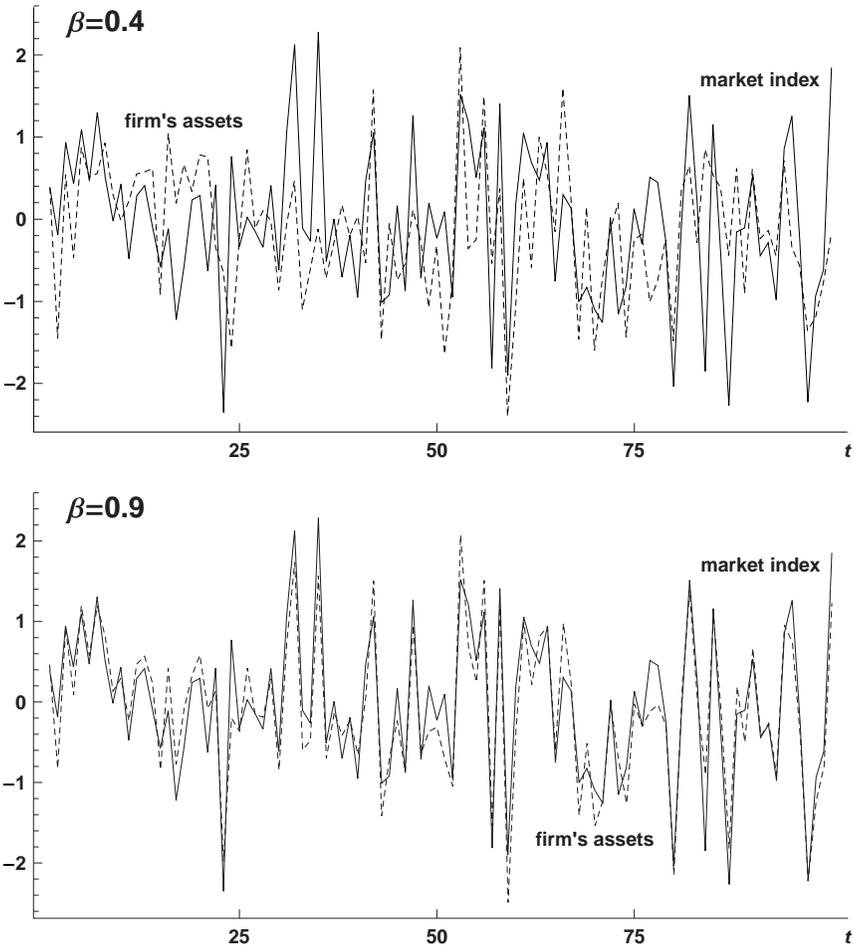


FIGURE 6.3 Asset and Market Index Returns in the Single-Factor Model
 Each panel shows a sequence of 100 simulations from the single-factor model for the indicated value of beta. The solid line plots the returns on the market index (a sequence of $N(0, 1)$ realizations). The dashed line plots the corresponding returns on a firm's assets with the specified β to the market, generated as $\beta m + \sqrt{1 - \beta^2} \epsilon$ using a second independent set of $N(0, 1)$ realizations.

its asset return variance? It is the β for which

$$\beta^2 = 1 - \beta^2 = \frac{1}{2} \quad \Rightarrow \quad \beta = \frac{1}{\sqrt{2}}$$

In sum, the single-factor model has two parameters, β and k . In applications, β and k may be estimated using the firm's balance-sheet and stock price data, or using the firm's rating and rating agency transition matrices. The mapping from default probability to default threshold is just that from standard normal cumulative probabilities to the associated z -values, for example:

π	k
0.05	-1.65
0.01	-2.33
0.001	-3.09

The parameters can be set independently, since a firm with a high default probability may be cyclical or defensive.

One of the main motivations of the single-factor model is to model portfolio credit risk, and we spend more time with it in Chapter 8.

6.9 CREDIT RISK MEASURES

The models we have sketched help us estimate corporate default probabilities and recovery rates. The next step is to use these models to estimate a risk statistic. The most common credit risk statistics, unexpected loss and credit VaR, are closely related, in that they incorporate the notion of potential loss at a future time horizon with a stated probability.

But the concepts are not quite identical, and, as we will see, they have different implications for how banks and investors should determine the appropriate amount of capital to set aside as a buffer against losses. These issues arise because

- The typical time horizon for measuring credit risk is much longer, on the order of one year, than for measuring market risk, where time horizons are almost always between one day and one month. An immediate consequence is that expected credit returns, the credit "drift," cannot be assumed to be immaterial, as the drift is for market risk. This in turn creates additional issues, involving the treatment of promised coupon

payments and the cost of funding positions, which do not arise in the same form in the market risk context.

- A second set of issues is driven by the extreme skewness of credit return distributions. For most unleveraged individual credit-risky securities and credit portfolios, the overwhelming likelihood is that returns will be relatively small in magnitude, driven by interest payments made on time or by ratings migrations. But on the rare occasions of defaults or clusters of defaults, returns are large and negative. This contrasts with market risk in most cases, although similar market risk skewness can be seen in option portfolios.

Because of the skewness of credit portfolios, the confidence level for credit VaR measures tend to be somewhat higher than for market risk; 95 percent confidence levels appear less frequently, 99.9 percent more so.

As an example of the skewness of credit returns, consider the distribution of future bond value in our Merton model Example 6.3, illustrated in Figure 6.4. Most of the probability mass is located at a single point, $D = 106$. The rest is distributed smoothly below D .

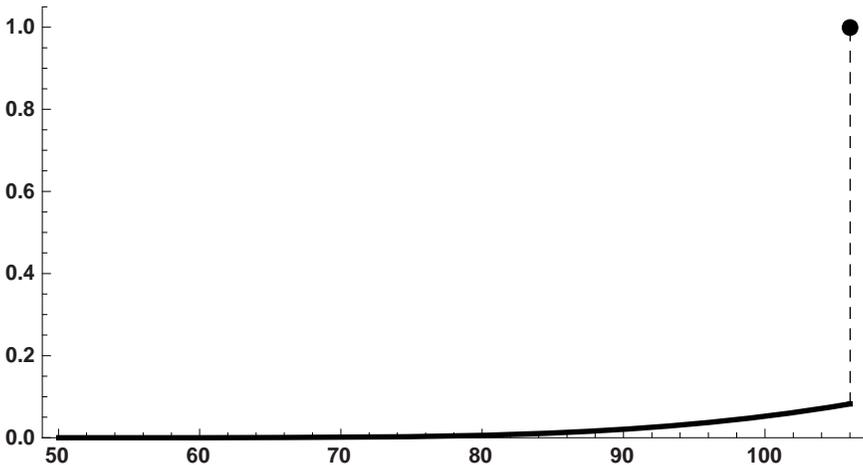


FIGURE 6.4 Distribution of Bond Value in the Merton Model
Cumulative probability distribution of bond value in the Merton model. The parameters are as set out in the text.

6.9.1 Expected and Unexpected Loss

Credit losses can be decomposed into three components: expected loss, unexpected loss, and the loss “in the tail,” that is, beyond the unexpected.

Unexpected loss (UL) is a quantile of the credit loss in excess of the expected loss. It is sometimes defined as the standard deviation, and sometimes as the 99th or 99.9th percentile of the loss in excess of the expected loss. The standard definition of *credit Value-at-Risk* is cast in terms of UL: It is the worst case loss on a portfolio with a specific confidence level over a specific holding period, minus the expected loss.

This is quite different from the standard definition of VaR for market risk. The market risk VaR is defined in terms of P&L. It therefore compares a future value with a current value. The credit risk VaR is defined in terms of differences from EL. It therefore compares two future values.

To make this concept clearer, let’s continue the Merton model example. The results are illustrated in Figure 6.5, the probability density function of the bond’s future value.

Example 6.5 (Credit VaR in the Merton Model) Figure 6.5 displays a portion of the density corresponding to the cumulative distribution of Figure 6.4, to the left of the default threshold. The graph shows how outcomes for the future value of the debt are decomposed into expected and unexpected

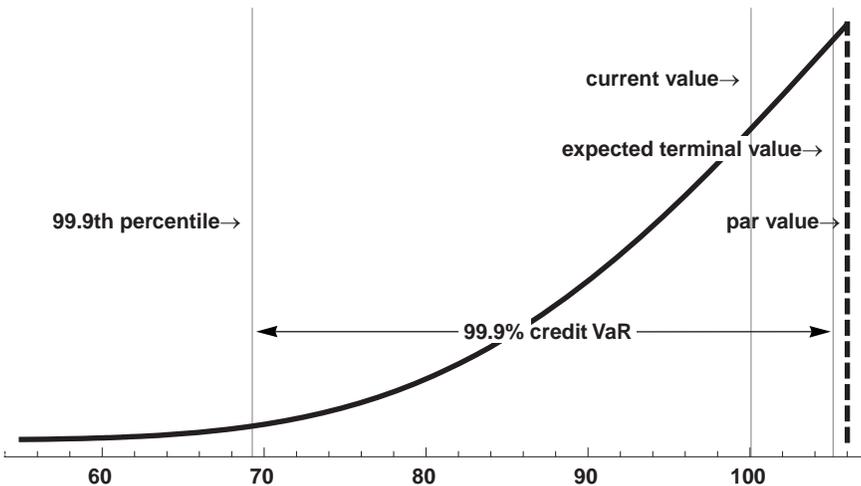


FIGURE 6.5 Credit VaR in the Merton Model
Probability density function of bond future value in the Merton model. The parameters are as set out in the text.

loss, and how the unexpected loss is decomposed into the credit VaR and losses beyond the credit VaR. The expected loss is the small gap between the par value and the expected future value of the bond, taking account of the default probability and recovery.

Actuarial expected loss	0.8619
Expected future value of debt	105.1381
0.001 quantile of bond value	69.2577
Credit VaR at the 99.9 confidence level	35.8804

6.9.2 Jump-to-Default Risk

Jump-to-default risk is an estimate of the loss that would be realized if a position, were to default instantly. It is based on the market value of the position or that of the underlying credit if the trade is expressed through derivatives. The jump-to-default value of a position consisting of x units of a bond with a value p is xpR .

Jump-to-default risk can be computed without using a probability of default. It is a form of stress testing, in that it realizes a worst-case scenario (see Section 13.3). It is a valuable tool for assessing the size and risk of individual positions. However, it can be misleading for portfolios. If there are long and short positions in a portfolio, the jump-to-default value of the portfolio is a net value that is likely to be small, even if the portfolio has significant risk. If the portfolio is large and consists entirely of long positions, the jump-to-default value will be misleadingly large, since it does not take diversification into account, that is, the fact that the probability of all the positions defaulting is lower than the typical probability of a single position defaulting. We discuss more precise ways of measuring credit risk for portfolios in the next chapter.

FURTHER READING

Duffie and Singleton (2003), Schönbucher (2003), and Lando (2004) are textbook introductions to credit risk models.

Although it has a focus on distressed rather than general corporate debt investing, Moyer (2005) is a good source on legal and quantitative credit analysis, and on capital structure issues. Baskin and Miranti (1997) discuss the historical development of limited liability. A classic paper on credit scoring is Altman (1968). Default rate histories are provided by Moody's Investors Service (2011) and Standard and Poor's (2010). See Duffie and Singleton (1999) on definitions of recovery in default, and Altman and Kishore (1996) on empirical recovery research. Frye (2000, 2001) discuss

the relationship between recovery and default, and its impact on the use of credit risk models.

Williamson (1985), Chapters 1 and 2, gives an overview of information cost problems in contracts. Gertler (1988), p. 568ff, and Mehran and Stulz (2007) provides applications to financial intermediation. Jensen and Meckling (1976) is a seminal text explaining the typical debt-equity capital structure of a firm as a response to the problem of agency costs. The term “moral hazard” was introduced into economics in Arrow (1963). The classic paper on adverse selection is Akerlof (1970).

Segoviano Basurto and Singh (2008) provide an overview of counterparty credit risk. Williams (1986) discusses the development of clearinghouses for futures markets. Timberlake (1984) and Kroszner (2000) discuss bank clearinghouses. The Lehman episode is described in Aragon and Strahan (2009). See also the Further Reading sections of Chapters 12 and 14.

Basel Committee on Banking Supervision (1999) and Task Force of the Market Operations Committee of the European System of Central Banks (2007) are surveys of credit risk models. O’Kane and Schloegl (2001) and Crouhy, Mark, and Galai (2000a) provide surveys focusing on commercial applications. Gupton, Finger, and Bhatia (2007) focus on the CreditMetrics methodology while providing a great deal of credit modeling background.

The original presentation of the Merton model is Merton (1974). Finger (2002) describes the *CreditGrades* variant on the Merton model. Credit factor models are presented in Finger (2001) and Gordy (2003). See Kupiec (2002) on issues in the definition of credit VaR.