CHAPTER 1

The Framework: Definitions and Concepts

Commercial credit is the creation of modern times and belongs in its highest perfection only to the most enlightened and best governed nations. Credit is the vital air of the system of modern commerce. It has done more, a thousand times more, to enrich nations than all of the mines of the world.

—Daniel Webster, 1934 (excerpt from speech in the U.S. Senate)

Theories of the known, which are described by different physical ideas, may be equivalent in all their predictions and are hence scientifically indistinguishable. However, they are not psychologically identical when trying to move from that base into the unknown. For different views suggest different kinds of modifications which might be made and hence are not equivalent in the hypotheses one generates from them in one's attempt to understand what is not yet understood.

—Richard Feynman, 1965

Objectives

After reading this chapter, you should understand the following:

- Definition of credit.
- Evolution of credit markets.
- The importance of a portfolio perspective of credit.
- Conceptual building blocks of credit portfolio models.
- Conceptually how credit models are used in practice.
- The impact of bank regulation on portfolio management.
- Why we advocate active credit portfolio management (ACPM).

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

WHAT IS CREDIT?

Credit is one of the oldest innovations in commercial practice. Historically, credit has been defined in terms of the borrowing and lending of money. Credit transactions differ from other investments in the nature of the contract they represent. Contracts where fixed payments are determined up front over a finite time horizon differentiate a credit instrument from an equity instrument. Unlike credit instruments, equity instruments tend to have no specific time horizon in their structure and reflect a claim to a share of an entity's future profits, no matter how large these profits become. While some equity instruments pay dividends, these payments are not guaranteed, and most equity is defined by not having any predetermined fixed payments.

In contrast, traditional credit instruments facilitate transactions in which one party borrows from another with specified repayment terms over a specific horizon. These instruments include fixed-coupon bonds and floating-rate loans (the coupon payments are determined by adding a spread to an underlying benchmark rate such as the U.S. Treasury rate or LIBOR[1]). Corporations are well-known issuers of these types of debt instruments; however, they are not the only borrowers. The past several decades have seen an explosion of consumer credit (particularly in the United States) in the form of home mortgages, credit card balances, and consumer loans. Other borrowers (also called *obligors*) include governments (usually termed *sovereigns*) and supranational organizations such as the World Bank. The credit risk of these instruments depends on the ability of the sovereign, corporation, or consumer to generate sufficient future cash flow (through operations or asset sales) to meet the interest and principal payments of the outstanding debt.

As financial engineering technology has advanced, the definition of credit has expanded to cover a wider variety of exposures through various derivative contracts whose risk and payoffs are dependent on the credit risk of some other instrument or entity. The key characteristic of these instruments is that, here again, the risk tends to lie in a predetermined payment stream over the life of the security or contract. Credit default swaps (CDS) exemplify this trend which aims to isolate the credit risk of a particular firm, the *reference obligor*, by linking a derivative's value to the solvency of the reference obligor, only. These contracts require the *protection buyer* to pay a regular fee (or spread) to the *protection seller*. In the event the reference

[1] The London Interbank Offer Rate (LIBOR) is the rate at which large banks are willing to lend to each other. The interest rate swap market provides an indication of how LIBOR is expected to change over time.

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

obligor defaults (per the specification of the CDS contract), the protection seller is required to make the protection buyer whole per the terms of the contract. Conceptually, the contract represents an insurance policy between the buyer (the insured) and the seller (the insurance provider). Extending the metaphor, the regular fee represents an insurance premium and the payout in the event of default represents an insurance claim under the policy. While a myriad of contract types now trade in the market, fundamentally they all represent a view on the credit risk of the underlying reference obligor.

While a CDS refers to a single name, derivative contracts on indexes of many named obligors can also be purchased as contracts on a specific basket of assets. These instruments expand the ability of credit portfolio managers to manage a large number of exposures without always resorting to hedging on a name-by-name basis or selling assets outright.

A related set of securities requiring financial engineering are broadly defined as *structured credit*. Popular forms of structured credit (also known as *securitization*) include collateralized debt obligations (CDOs) and asset-backed securities (ABS). In recent times, the credit crisis has made discussion of CDOs and ABSs more common in the media. Many commentators have called for drastic measures to curtail the use of structured credit. While abuse of these instruments can increase risk in institutions and markets, structured financial products can also be used responsibly to reduce risk in the financial system. Some regional banks, for example, have successfully hedged the concentration risk in their portfolios that results from most of their loans being originated in a single geography. They do this by selling some of their portfolio risk via structured credit. Other investors have purchased this risk and integrated it into their own portfolios as diversifying investments, creating lower volatility portfolios with improved return per unit of risk profiles. All market participants benefit from this kind of trading. Of course, these instruments can be abused when combined with excessive leverage or when market participants attempt to speculate using structures they do not fully understand.

But even the *simplest* of financial instruments such as equity can be inappropriate for particular investors in certain situations. The same is true of structured credit. We try to be careful to distinguish the *purpose* from the *characteristics* of particular instruments.

Conceptually, the basic structure of these instruments is straightforward: A number of securities or derivative contracts called *collateral* are placed in a structure called a *special purpose vehicle* (SPV) or *special purpose company* (SPC), creating a corporate vehicle to direct the cash flows from the collateral. In its simplest form, the purpose of the SPV is to borrow cash, typically through debt issuance, and to use this cash to purchase the collateral: some type of credit-sensitive obligation. The collateral may be provided by a financial institution, such as a bank that issues mortgages, or purchased in the secondary market, such as the case of corporate bonds.

Why could not a financial institution just issue the bonds directly rather than through an SPV? The purpose of an SPV is typically to create *bankruptcy remoteness* for the issuance of the debt. This means that the ownership of collateral is legally transferred from, say, the bank that made the loans, to the SPV. The objective is to ensure that if the bank goes into default, the collateral will not be considered part of the assets of the bank. Said another way, the SPV structure ensures that the collateral will be used only for the benefit of the holders of the structured securities issued by the SPV, regardless of where it was originated.

The SPV uses the cash flow from the collateral to pay back the debt as the collateral generates payment income through, for example, amortization and interest payments. The cash flow from the collateral is paid out to holders of each class of the liability structure (called a *tranche*) of the SPV per a set of rules called a *cash flow waterfall*. The tranching of debt creates a priority of payments (or of loss positions) such that more junior tranches (i.e., those lower in the capital structure) absorb losses first, followed by the next most senior, and so on. The motivation behind these structures is the desire to change the return/risk profile of the collateral into a set of securities or tranches with different return/risk profiles, with lower tranches exposed to more risk and higher tranches enjoying greater protection from collateral losses. In many structures there are also rules that specify that all cash be directed to more senior tranches if the performance of the collateral begins to deteriorate, providing still further protection for the higher tranches. It should be obvious that the analysis of many types of structured instruments is therefore quite similar to the analysis of a portfolio of assets in any financial institution but with the added complication of waterfalls and other structural provisions.

The names of these structures, such as CDO or ABS, reflect this collateralized nature of these instruments. Each specific structure name refers to the nature of the collateral:

- CLO: Collateralized loan obligation.
- CBO: Collateralized bond obligation.
- CDO-squared: CDO of tranches issued by other CDOs.
- RMBS: Residential mortgage–backed security.
- CMBS: Commercial mortgage–backed security.

Even without the added complexity of a securitization, credit instruments can be fairly complicated on a stand-alone basis. For example, many corporate bonds incorporate an attached call option designed to give the issuer the opportunity to pay back the debt earlier, should market conditions favor doing so. The call option identifies a price at which the issuer (i.e., obligor or borrower) can purchase back the debt. In an environment of falling interest rates or improving credit quality for the borrower, this option opens the door for the borrower to take advantage of better terms as they become available. For example, a fixed-rate bond will rise in price as interest rates fall. At some point, the issuer of a callable bond will find it advantageous to purchase back the debt so they can reissue at the lower rate. The call option provides this opportunity. As another example, many bank loans are structured with triggers and other features that change the payoff of the loan conditional on various metrics related to the borrower's performance. Such loan covenants may increase the loan's coupon rate if the financial performance of the firm, based on a predefined metric such as a leverage ratio (e.g., total debt/total equity or total debt/total assets), deteriorates.

Sometimes a credit exposure does not even reflect actual cash being loaned right away. Instead of a straight term loan, a bank may extend a commitment to lend with a variety of conditions as to the terms of borrowing. We typically refer to loans where cash is actually disbursed as *funded* and commitments to lend as *unfunded*. Note, however, that a contractual commitment to lend exposes the bank to risk even if funds have not actually been transferred to the obligor.[2] As this brief discussion highlights, credit exposures like these can be decomposed into a risk-free debt instrument and a collection of other (e.g., default, prepayment, interest rate, etc.) options. In fact, most credit instruments represent a portfolio of options.

Credit exposure also arises in the context of more traditional derivative transactions such as equity options and interest rate swaps. When such a derivative is *in-the-money*,[3] the market risk (i.e., risk arising from changes in quantities driving the value of the derivative) must be separated out from the credit risk. This implicit credit risk may become significant when systemic events impact the entire market. The recent financial crisis has highlighted how the solvency of large counterparties to derivatives transactions can have widespread impact on the financial system overall. The most recent global credit crisis is not, however, the only example in modern times of increased

[2] A common oversight of some banks is to ignore their unfunded commitments since the commitments are made to potential borrowers at times when these borrowers are financially healthy. The problem that can arise is that these obligors tend to borrow at times when they face difficulties.

[3] The counterparty who is *out-of-the-money* owes payments to the counterparty who is *in-the-money*, such as would be the case, for example, for the holder of the fixed-rate leg of a floating-fixed-rate interest swap when fixed rates were above floating rates.

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

counterparty-default risk. The latter part of 1998 also saw a substantial increase in the likelihood of counterparty default. After Russia defaulted on its domestic currency debt and LTCM (a large hedge fund) came to the brink of insolvency, many investment banks appeared to face unprecedented difficulty. In this situation, the risk of a counterparty not paying became significant. Counterparty credit risk always exists, and even if a derivative counterparty does not default, the value of an in-the-money derivative may be adversely affected by the difficulties faced by the counterparty. A firm or counterparty does not have to default in order to result in a loss of value for a particular credit-risky instrument. Counterparty credit risk has become a much more important topic as the volume of derivatives has mushroomed and market participants have become more cognizant of this risk.

The salient feature of all these different types of credit exposure is the shape of the distribution of losses. Credit exposures are typically characterized by skewed, fat-tailed return distributions. That is, the lender or originator of an exposure has a high probability of receiving its principal back plus a small profit over the life of the exposure and a low probability of losing a significant portion of the exposure. An example of a credit loss distribution can be seen in Figure 1.1.

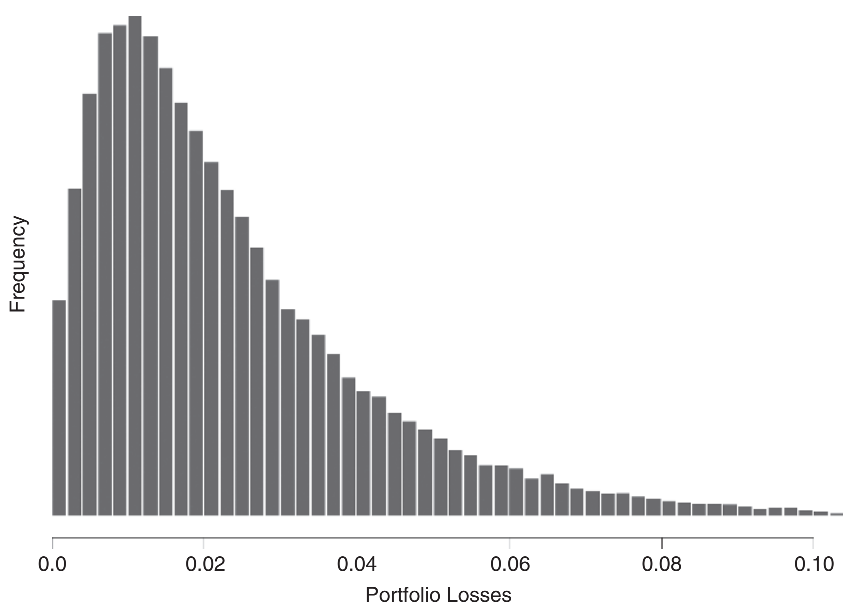

FIGURE 1.1 Simulated Loss Distribution

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

Said another way, many borrowers have a high chance of repayment but if they do fail, they tend to fail severely. The correlation among these types of exposures tends to be quite low compared with, say, the correlation of equity exposures. However, ironically, the diversification of these exposures tends to take a larger number of names than is the case with equity or other instruments with less skewed payoffs. This low correlation coupled with the chance of losing a substantial amount on any one exposure makes these securities particularly well suited for management in the context of a large, *well-diversified* portfolio. If a bank's portfolio contains only small bits of each exposure, the occasional extreme loss for any one exposure will tend not to affect the portfolio's overall performance. Thus, diversification buys stability in the portfolio's loss profile. Importantly, unlike the case of other instruments, even a well-diversified portfolio will typically exhibit significant skewness *that cannot be diversified away*. We return to this conclusion a number of times throughout this book.

EVOLUTION OF CREDIT MARKETS

While the idea of debt extends back into ancient societies, the more modern notion of credit really began in preindustrial Europe in the context of commercial payments. Credit was typically extended by way of deferred payment for goods sold or advance payment for future delivery of goods purchased (see Kohn 2001 for more details on the history of banks and credit). Over time these debts began to be treated as fungible and would be assigned to other merchants, and eventually systems of settlement evolved. Deposit banking developed in response to the need for assignment of third-party debt among strangers. Since the bank became the counterparty for multiple transactions, it could net a large number of payments without resorting to final cash settlement.

This set of circumstances enabled preindustrial banks to offer a solution to the endemic problem of liquidity risk faced by merchants, namely a short-term lack of cash preventing completion of a particular transaction. Since depositors in the bank found it convenient to leave their money with the banker so that settlement of transactions could be done without having to lug around actual coins, the bank now had a store of deposits to use as the basis of an overdraft loan. The bankers discovered that they could extend credit beyond the quantity of actual coins or gold on deposit since most depositors did not demand all of their deposits most of the time. Here we find the beginnings of leverage in financial institutions. Since the banker knew his clients well, the bank could use its knowledge of the capacity of a potential borrower (who is also likely a depositor) to repay a loan and allow this individual to periodically overdraw his account. Eventually, these short-duration, relatively small overdraft loans were supplemented and then overtaken by longer, larger commercial loans. (Again refer to Kohn 2001 for more details on the evolution of banks and credit markets.)

From these humble beginnings, credit evolved along a myriad of dimensions. Credit could be extended not only in the form of loans, but also in the form of bonds traded in a global capital market. Computers replaced written ledgers and money become *tokenized*—represented as digitized bits stored in a hard drive. However, the characteristics of credit remain the same. Yet along with this technological progress developed a capacity for higher volume lending. As a result, a number of difficulties appeared as the institutions and markets developed for the origination and management of credit and this evolution progressed.

The first difficulty the financial world encountered was that of *bank runs*. Since the process of lending depends on depositors not demanding their money in cash all at once, the reputation of the bank, and depositor's confidence in its solvency, is critical. If a large enough number of depositors perceive a bank to be unsound and demand their cash all at once (creating a run on the bank), that bank may fail even if the perception is false. The creation of a lender of last resort such as a central bank and the provision of deposit insurance from the government are institutional responses to this bank failure risk due to runs on banks.

The second difficulty developed from the challenge in managing potentially large losses on the bank's loan book. In these cases, the trouble arises when a sizable portion of a bank's portfolio of loans simultaneously cannot be repaid as promised. In this case the bank, in a sense, becomes the victim of its own success. Typically, a bank develops expertise in originating loans within a particular geography and sector. For example, large Japanese banks in the 1980s became very good at lending money to large Japanese trading companies. While economic times were good, this concentration of loans in one geography and one sector did not seem to pose a problem. However, such concentrations obviously create significant correlation in the payoffs of the loans in the portfolio. When Japan's economic bubble burst and the 1990s uncovered the disastrous impact of holding a concentrated portfolio, large Japanese banks watched the loans in their portfolios deteriorate together. This problem is by no means unique to Japan. It is hard to find any country with a functioning banking system that has not seen this kind of bank crisis at some point in its financial history. Origination expertise in a particular area leads to concentrations that create problems when that sector or geography becomes troubled.

The third difficulty concerns the inefficiencies in the market for corporate credit. The corporate bond market developed in parallel to the expansion in the origination of bank loans. In preindustrial Europe, some merchants traded bills of exchange with each other. Over time, a dealer market emerged for corporate debt. The problem with this market was a lack of standardization and in turn a lack of transparency in pricing. These inefficiencies resulted in a lack of liquidity, making it difficult to trade in and out of positions and to thus manage a portfolio of corporate debt.

These challenges notwithstanding, debt markets continued to mature, albeit at a leisurely pace. However, the 1990s ushered in a new era for corporate credit markets in which several trends converged to create an environment where credit could be priced and managed in a relatively efficient manner. The first trend involved the successful implementation of objective, quantitative analytic tools to facilitate rigorous evaluation of credit exposures. This environment arose from the marriage of modern finance and powerful computer technology. However, the ability to analyze the risk of a credit portfolio was only the first step; a portfolio manager also needed to have the ability to act on this analysis and trade at a reasonable cost. This second step, which has only become fully implemented in the past decade due to the availability of cheap telecommunications, has created a trend that facilitates inexpensive trading in credit-risky instruments. While corporate bonds have always been traded, a market in secondary trading of corporate loans has also developed.

The third step in this evolution was the ability to complete the cycle of analysis and trading and to thus diversify *portfolio* holdings. Modern financial theory emphasizes the power of portfolio diversification. A variety of financial institutions ranging from banks to pension funds now manage their portfolios using measures of diversification. This third trend has set the stage for a dramatic increase in the number of market participants trading credit for reasons other than just exiting a distressed position (although corporate distress will always motivate a significant number of trades).

In recent years, some of the most sophisticated banks have used portfolio analysis technology to devise transfer pricing mechanisms allowing them to separate the management of the bank's credit portfolio from the creation of valuable service businesses. (We discuss this organizational change in more detail later in this chapter.) Clearly, today the motivation for trading credit goes beyond avoiding a default and ranges from perceived market inefficiencies to portfolio rebalancing designed to improve the return/risk profile of an institution's entire credit portfolio.

Another important trend has been the change in the regulatory environment as financial regulators come to grips with the importance of measuring and managing credit risk. The first global bank accord, known as Basel I, defined a simple notion of how much capital a bank should hold, given the credit risk of its loan book. Currently, a more complex accord known as Basel II is being debated. While regulators have now acknowledged the feasibility and importance of estimating quantitative measures of credit such as probability of default (PD) and loss given default (LGD), the most advanced banks have already been running systems that not only evaluate PDs and LGDs, but also incorporate the correlations among exposures in their portfolios. Some regulators have made efforts to incorporate a portfolio view into bank regulations, but the progress has been slow. The benefit of this new regulatory focus on credit is that it motivates many financial institutions to invest in the systems that enable them to do better credit portfolio management. Regulators have also improved market transparency. In the United States, regulatory pressure resulted in the creation of the Trade Reporting and Compliance Engine (TRACE) data initiative requiring bond dealers to post their transaction prices for corporate bonds.

A fifth trend is the sudden appearance of a deep and liquid market in corporate credit derivatives. At the time of this writing, the CDS market exceeds USD$60 trillion in notional value. The availability of credit indexes such as the iTraxx and CDX makes it much easier to hedge portfolio exposure. Synthetic CDOs have become common transactions in the world of credit management. These instruments create a mechanism for more efficient management and transfer of risk exposure. A portfolio manager can now isolate the credit risk components of price from other types of risks impacting the value of a bond or a loan (e.g., market risk and liquidity risk). In this way, portfolio decisions are no longer held hostage by the inability to trade a particular risk by itself. Furthermore, research can now begin to sort out the relationships among credit risk and some of these other kinds of risks. The draining of liquidity in the structured credit market in 2007 and 2008, particularly for collateralized loan obligations, has set the market back somewhat as the ability to hedge with structured credit has diminished. More recently, questions have arisen regarding potential misuse of leverage in constructing portfolios of CDS contracts, and more investment and transparency is needed in the infrastructure of settling CDS trades. These challenges have made all market participants more focused on how to better develop this important tool for managing credit risk.

Though still evolving, the markets for corporate credit risk, whether they involve bank loans, bonds, or credit derivatives, are becoming more liquid and more transparent. This does not imply that they are anywhere near fully mature. The development of these markets has not been smooth,

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

as exemplified by the recent credit crisis resulting in the dramatic reduction in issuance volume in many sectors of the market for CDOs in late 2007 and the overall difficulties across most credit markets in 2008. That said, the CDS market remains the primary place to trade corporate credit risk and it appears to be here to stay despite recent drops in volume. This market is generally much more liquid than other markets involving credit. While these markets still have much room for improvement (particularly outside of the United States), we have the benefit today of tools and the understanding to manage a portfolio of corporate credit exposures actively in a way that substantially decreases the risk of extreme losses. Tools and methods are also being developed for analyzing portfolios of ABS and retail exposures, though the quantitative literature on these types of exposures lags in many cases that of the corporate literature.

The challenge lies in choosing the right models and systems to support this active corporate portfolio management effort. Even more important is modifying the way that risk is managed within a financial organization in a manner that motivates employees to make decisions that result in efficient allocation of the bank's economic capital. In our judgment, proper organizational incentives informed by useful portfolio insight will lead to less risky, more valuable banks.

DEFINING RISK

Throughout our discussions in this book, we define risk as the possible change in value of a security or asset over a particular time horizon. Change in value is not the only way to define risk. Some practitioners have focused on risk as defined only in terms of the probability of default (i.e., firms with low PDs (high ratings) are safe and those with high PDs are risky). The trouble with this definition is that a portfolio can store up "time bombs," in effect, that are not readily appreciated until it is too late when many firms in the same industry or geography default at the same time. Since the probability of default of one loan is the same regardless of the concentrations in a portfolio, the potential for large losses on a portfolio can change dramatically with portfolio correlation. Furthermore, the tracking of credit migration or changes in value prior to maturity becomes essential to capturing the true risk of a portfolio through time. Otherwise, the portfolio manager may be surprised by a cluster of sudden defaults. In this context, an approach that considers both the underlying risk and the change in the values of securities within a portfolio is a superior focus for risk assessment than just the risk of defaults within a portfolio.

Other authors have argued that risk should be defined only in terms of a decline in value. However, our experience suggests that a focus on only downside variance (sometimes called semivariance) ignores important information about the future. For example, Internet firms in the late 1990s experienced a few years of skyrocketing growth in value. Their later fall was even faster than their rise. Focusing just on downside variance in those cases would have led to a severe underestimation of overall risk.

In the case of *credit risk*, this change in value derives from the changing probability that the obligor will fulfill its obligation to pay interest and ultimately repay principal. This is fundamentally different along a number of dimensions than *market risk*, which encompasses changes in a security's value as driven by variables such as interest rates, equity prices, commodity prices, and foreign exchange. Financial practitioners have settled on models and systems in the field of market risk much more quickly than in the field of credit risk. The availability of data and liquid markets in instruments such as interest-rate swaps and other derivatives has made it easier to introduce quantitative hedging and portfolio management techniques in the field of market risk for equity and other instruments, while the absence of data and the more complicated statistical relationships made it more difficult historically to do the same for credit risk. That said, recent advances in both fields have produced a convergence of models and systems. Increasingly, we are encountering demands to integrate credit and market risk.

We touch briefly in this book on the state of this integration. Our primary interest lies in understanding how interest rates and credit spreads are related. The portfolio factor model structure we introduce in Chapter 8 can be modified to handle both credit-risky securities and market-risky securities. The challenge lies in defining the function that transforms *factor realizations* (i.e., economy- or sector-wide shifts in the drivers of default) into a security value. The increasingly heterogeneous (in terms of asset classes) nature of most financial institutions' portfolios makes it even more important to build models with the flexibility to handle a variety of instruments. As part of our exploration of reduced-form models (Chapter 6), we also discuss the similarities between market-risk models and some of the reduced-form models used for credit risk.

As previously noted, in credit risk modeling, we attribute much of the change in value of credit-risky securities to changes in the likelihood that the obligor will pay its coupons and repay principal. Some models, such as structural models, rely on specific economic reasoning to describe why an obligor defaults—namely that the market value of the borrower's assets falls to a point at which it no longer covers the total amount of its

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

obligations. Other more statistically focused models such as reduced-form models do not rely on a specific causal economic relationship, but rather focus on default as an unpredictable event that can be captured in a coherent mathematical model that is consistent with financial theory. Even so, reduced-form models tend to focus on processes that drive credit quality. They can also be extended to include processes that drive the state of market liquidity.

What can substantially muddy this modeling challenge is the possibility of a liquidity-based default or liquidity-based change in security value. In a circumstance in which market liquidity has dried up, a firm with sufficient market value may still default because it cannot roll over its short-term debt as it comes due. The claims represented in the issued loans and bonds of a particular obligor may still relate to that obligor's valuable assets, but the absence of liquidity in the market prevents a portfolio manager or credit trader from finding new financing or selling positions in its portfolio to cover existing claims. These liquidity-driven difficulties may result from different processes than the ones driving changes in credit quality (although the credit problems and liquidity difficulties are often related). From a model perspective, we attempt to separate (when possible) the effect of credit factors from the effect of liquidity factors on estimates of relevant metrics that characterize risk and value.

A WORD ABOUT REGULATION

Given the importance of banks to most national economies, governments have an interest in ensuring the prudent management of these institutions. Such efforts to reduce systemic financial risk often focus on instituting regulations. At the international level, the Bank of International Settlements (BIS) has taken on the task of coordinating proposals for bank regulations internationally. These proposals may or may not be implemented in each domicile; however, the ideas spark discussion throughout the regulatory community. As mentioned earlier in this chapter, in the late 1980s, BIS published a global banking accord designed to eliminate the advantage Japanese banks seemed to hold in gaining access to cheap funding. Basel I, as the accord is now called, outlined for banks the appropriate levels of capital they should hold for given classes of risk. It did this in broad terms with the goal of creating a common language of regulatory capital risk rather than of outlining detailed risk management strategies.

However, in hindsight, while it was an important step forward, the blunt nature of Basel I created opportunities for regulatory arbitrage in which a bank could take advantage of situations in which the rules unintentionally led institutions *away* from economically profitable transactions.

In recent years, the BIS has struggled to finalize the next generation of regulation, Basel II. Basel II is intended to create more sensible guidelines within which banks can develop systems for credit risk assessment and economic capital allocation. The promise of Basel II lies in aligning the regulatory guidelines with the way in which decisions are actually made at financial institutions. Unfortunately, the tendency of government entities to create broad-ranging proposals that attempt to satisfy many different interest groups has resulted in regulations that some market participants feel fall short along certain dimensions. One positive result of the Basel II efforts is the impact it has had on the way in which senior bank managers think about and now focus on the notion of quantitative credit risk modeling and capital allocation. As a consequence, risk management efforts within banks now receive better funding to build and implement systems that not only facilitate regulatory compliance, but that can also be used to implement economic capital systems, which in turn result in more efficient and, importantly, less risky banks.

To our knowledge, most regulators still do not publicly promote the idea of active portfolio management.[4] Their efforts focus more on establishing rules that reduce systemic risk in the financial markets. However, the regulatory perspective with respect to quantitative risk management has become far more sophisticated than it was at the time Basel I was introduced. In fact, some of the leading researchers on credit risk now reside within central banks and other regulatory bodies. As a result we expect that over time, newly formed organizations such as the International Association of Credit Portfolio Managers (IACPM) will assist banks in the process of coordinating with regulators to improve the dialogue around implementing new systems and new organizational structures.

WHAT ARE CREDIT MODELS GOOD FOR?

One of the authors recalls an experience a number of years ago teaching a group of old and wizened loan originators at a bank implementing quantitative tools for credit risk management. In the middle of the training session, one frustrated participant complained that we "rocket scientists" were

[4] In fact, Basel II has relatively less to say about portfolio correlation in general, compared to PD and LGD estimation.

destroying relationship banking. He went on to proclaim that a computer model could never match his capability for assessing a company's credit quality. While his track record was never verified to our knowledge, we are aware of several studies[5] at banks that show that *on balance*, subjective credit risk assessment alone is decidedly inferior to quantitative-based approaches (in a later chapter we explore in more detail how to evaluate models). Further, the credit officer's first statement in this anecdote about the destruction of relationship banking seemed to imply a simplified view of how models should be used.

While some computer scientists still assert "true" artificial intelligence is possible in the near term, typical businesspeople do not expect that a model or computer will fully replace a human being in the credit assessment process in the foreseeable future or that this would even be a good thing. In fact, relationship banking is alive and well and relies primarily on the strength of human intuition. Rather than destroying relationships, quantitative models change the way a bank can be organized and, more importantly, change the way credit analysts and relationship managers can do their jobs. Well-implemented systems improve the development of relationship banking and increase the efficiency and accuracy of credit analysts. Good models can provide a means to reduce some of the more tedious aspects of credit analysis and focus the analyst on the obligors, data, and processes that need attention as the bank manages its risk.

With quantitative models at the foundation of a bank's credit assessment process, qualitative assessment can be overlaid when appropriate. It is crucial that when such systems are developed and implemented, they facilitate ongoing rigorous assessment of how well models are performing and what the models' limitations are, regardless of whether they are quantitative, qualitative, or a mixture.

Qualitative assessment becomes more important when evaluating borrowers where market observable information is lacking. Even in these circumstances where data is scarce, a quantitative model can assist in directing the conversation to meaningful characterizations of what drives a borrower's risk. In many ways, these models become a *lingua franca* for risk discussions throughout the bank and transaction discussions outside the bank. We find that the most successful institutions benchmark (on a regular basis) internal models to ensure that the language of risk maintains the same meaning from transaction to transaction. Models are best used in environments in which

[5] Such studies are often internal and thus are not often published externally; one standard older reference on a nonfinance topic is Dawes (1979).

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

the organization maintains a balance of healthy skepticism—reviewing the models underlying this language of risk and reconsidering model assumptions regularly—and healthy enthusiasm for the efficiency and insights that quantitative approaches to credit risk management can bring to their credit processes and internal communication about risk. If implemented correctly, this language of risk can be used to transform a financial institution's business, moving it from origination of single exposures to *active credit portfolio management*.

ACTIVE CREDIT PORTFOLIO MANAGEMENT (ACPM)

Throughout this chapter we refer back to the importance of managing a portfolio and improving its diversification. In the equity market, symmetric return distributions coupled with the diversified nature of what is available in the market often means that active management does not pay high dividends. In fact, most active equity managers *under*perform their risk-adjusted benchmark.

Credit is different. Credit markets do not originate well-diversified portfolios, and the asymmetric nature of credit return distributions makes avoiding a deteriorating credit material to overall portfolio return. Moreover, the lack of good benchmarks makes it difficult to offer index funds that do not suffer from substantial idiosyncratic risk. These characteristics of credit markets create an opportunity to earn outsize returns given a particular level of risk on a portfolio of credits that is actively managed. This starts with models and systems that discriminate good from bad obligors. Further returns can be earned by refining the correlation estimates—a difficult but achievable proposition to some degree for certain segments of the credit markets.

Another important reason that active management is beneficial in the world of credit has to do with the heterogeneous nature of liquidity across credit-risky instruments. While robust liquidity models are still being researched, good credit models can move a manager a step closer to identifying profitable trades and reduce the uncertainty with respect to the question of liquidity. In many circumstances, these models provide an interpretive framework to discern the different factors driving value and focus analysis. Developing these strategies in the context of portfolio trades helps reduce the idiosyncratic impact of inexplicable behaviors of particular securities. A portfolio perspective complemented with quantitative systems sets the stage for generating high Sharpe (return per unit of volatility) and Vasicek (return per unit of tail-risk contribution or return per unit of capital) ratios for a credit portfolio that is actively managed.

LIQUIDITY

Defining liquidity can be difficult. In general, we think of liquidity as a measure of the depth of a market and the ease with which a trade can be made. For some, liquidity is the label researchers place on the things that economists or financial modelers cannot explain (i.e., the residual in their analyses). With the development of a variety of markets pricing risk associated with the same names (e.g., equity, bonds, loans, CDSs), we have begun to catch glimpses of pricing differences that are a function of differences in liquidity. We do not yet have the full framework to sort out these differences. In the meantime, we are left with cruder methods, such as matrices of liquidity premia that reflect geography, industry, and size of the obligor.

While we currently do not have fully developed models of liquidity, we do understand the following:

1. Many theoretical credit models underestimate credit spreads, in part because they do not account for a liquidity premium.
2. Large transactions or trades tend to be heavily impacted by lack of market liquidity.
3. While available approaches are still evolving, some measure of liquidity (even if *ad hoc*) should be incorporated into mark-to-market and transfer pricing frameworks.
4. With the availability of CLOs and bespoke synthetic CDOs, we can develop an estimate of the cost of hedging through these vehicles that can assist us in finding an indirect estimate of the illiquidity premium. The difficulty lies in disentangling the credit risk premium from the illiquidity premium.

The topic of liquidity will continue to be a focus of research as more financial institutions build up their portfolio management capabilities. The dramatic changes in liquidity seen throughout the credit markets since late 2007 should provide important new data on liquidity premia.

While still in its infancy, the development of ACPM groups within financial institutions and the increasingly common discussions of the importance of tracking a portfolio's mark-to-market value suggest that some banks will start to look more like trading houses than classical commercial lending institutions. This shift will continue to blur the difference among different types of financial firms. Hedge funds, large corporations, insurance companies, asset managers, and investment banks are joined by commercial banks as financial institutions discover the value in separating the *management of their credit portfolios* from the *development of franchise businesses*. In some cases (e.g., hedge funds), the only business of a firm may be managing a portfolio, while in other cases (e.g., large financial conglomerates) the ACPM business is just one of many. Though still developing slowly, this convergence bodes well for the global capital market's ability to originate, distribute, and manage credit risk without creating dangerous concentrations in any one location.

An implication of this shift in managing a bank's portfolio separately from developing its franchise businesses (which includes the business of loan origination) is that a bank moves from an originate-and-hold strategy to an originate-to-distribute (also called "originate-and-distribute") strategy. This means that loans may be sold or hedged right after origination and not necessarily held to maturity. Said differently, the bank now manages its portfolio or credit risk based on portfolio concerns rather than assuming it will hold each originated loan to maturity.

Some critics have pointed to the originate-and-distribute model of commercial banks as a key cause of credit market difficulties such as the recent subprime crisis. In a world where the portfolio managers (whether they are CDO collateral managers or ACPM portfolio managers) do not rigorously evaluate the securities for which they have responsibility or where outright fraud is perpetrated by borrowers, an originate-and-distribute model can result in agency problems in which market participants do not pay sufficient attention to (or have transparency into) what kinds of borrowers are creating credit exposure. The problem that can arise when this happens on a large scale is that dramatic market corrections that occur in a systemic manner across the economy can have undesirable external impact in other parts of the financial markets. The only environment in which the originate-and-distribute model can function is one in which there is ample transparency with respect to instruments and assets and in which the incentives and structure of the lending process makes fraud difficult and its penalties severe.

While this is a tall order, unfortunately, the alternative of returning to the originate-and-hold model leaves the economy open to a greater risk of widespread systemic problems as commercial banks end up holding concentrated portfolios that cannot withstand cyclical economic downturns. The numerous bank crises seen throughout history illustrate this risk. Each system has its strengths and weaknesses. Our view, however, is that the originate-and-distribute model has much more to offer to counterbalance the possibilities of widespread market difficulties. Recent events will more than likely increase market transparency and set the stage for much more robust

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

institutional response to liquidity crises. In our view, while the originate-and-distribute model must still evolve to provide more closely aligned incentives for market participants, turning back to the former model of originate-and-hold will not do much to improve the resilience of financial markets.

The tools we describe in this book and the framework we suggest for their application within banks and other financial institutions provide a means to achieving ACPM by coordinating a set of models and systems with organizational change to improve dramatically the growth opportunities at a bank. The mechanism lies in aligning incentives at the nonportfolio business and relationship manager levels with the overall objective of a bank's management to build new and growing channels of cash flow. In the process of making these system and organizational changes, the bank will manage its credit portfolio such that the likelihood of extreme loss can be significantly reduced (though some systematic risk will always remain).

The ACPM function becomes critical to making the most of the models and systems available. The necessity of holding concentrated portfolios to leverage internal bank local expertise disappears. Discussions about business strategy and new transactions become much more meaningful as a quantitative framework provides context for framing and testing assertions. By coupling this with a performance evaluation system tied into this framework, the bank's management can credibly justify higher valuation in the equity market and lower spreads in the debt market. This objective of higher share valuation becomes the ultimate motivation for moving the bank to an active portfolio management mind-set and investing in the models and internal processes to make this happen.

FRAMEWORK AT 30,000 FEET

At a conceptual level, the models discussed in this book provide insight into the return and risk trade-off among exposures in a credit portfolio. The stand-alone risk of a particular exposure tends to be the easiest to understand and act upon. Most analysts look to their wins in terms of which names they labeled correctly as high or low risk. The industry tends to remember the analyst who identified a deteriorating credit well before this deterioration was reflected in that obligor's loan, bond, or CDS price. A financial professional who identified problems at WorldCom or Ford before they became newspaper headlines will emphasize this in describing his abilities as a credit analyst.

The problem with these isolated examples is the lack of focus on the overall performance of the portfolio. If an analyst is consistently negative about all obligors, he will successfully identify the handful of big names that default. But that analyst has not necessarily helped that financial institution: Anyone can recommend avoiding all prospective borrowers. While it would be unusual for an analyst to deem all prospects poor risks, many qualitative assessments can tend toward the negative. Some analysts will use this negative bias to highlight the borrowers that do default. A similar difficulty will arise from an always optimistic analyst. Stand-alone risk assessment should clearly distinguish the strong from the weak borrowers. The ability to make this distinction should be regularly benchmarked and tested regardless of whether the assessment is done by a model, an analyst, or both.

What we and others have discovered over the past 20 years is that by itself, qualitative, stand-alone risk assessment typically does not (on average) lead to better-performing portfolios. Analysts who can regularly separate winners from losers in high volume are few and far between. More importantly, from a bank's perspective, the risk of any particular exposure is less interesting than is the performance of the portfolio as a whole. Thus, single exposures should be evaluated in the context of a portfolio, which requires characterization of credit exposure correlations. Stand-alone risk is only one piece of the portfolio puzzle. Moreover, we need measures that place each exposure on the same scale. The framework we emphasize in this book enables an analyst to calculate a portfolio value distribution. This distribution reflects the likelihood of different value outcomes over a specified time horizon.

The probability of a loss exceeding the point on the distribution associated with a target threshold can be interpreted as the probability that the portfolio will become insolvent—in other words, that the capital will be exhausted in the remote event of an extreme loss beyond the threshold. (In the next section we describe how this threshold may be set.) Each exposure's contribution to aspects of the portfolio loss distribution reflects a consistent, portfolio-referent measure of risk. We will refer to these portfolio-based risk measures as *risk contribution*.

Later we will be more specific about how we calculate risk contribution. At this stage, risk contribution can be interpreted as the post-diversification contribution of an exposure to a portfolio's overall risk—that is, the exposure's contribution to overall portfolio risk after accounting for how much the exposure adds to or subtracts from the portfolio's overall diversification. Risk contribution reflects the marginal risk of an exposure. An important point is that this marginal risk or risk contribution measure will be specific to a unique portfolio. What looks like a good addition at the margin in a Japanese bank portfolio may be a terrible loan to make for a U.S. bank portfolio.

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

Given the importance of understanding the concept of risk contribution, consider two examples where a portfolio risk contribution measure will motivate different conclusions than a stand-alone risk measure.

Japanese banks in the 1990s tended to hold portfolios heavily concentrated in large, Japanese companies. Some of these companies such as Toyota were quite safe on a stand-alone basis. However, the risk contribution of a safe, large company such as Toyota to a Japanese bank portfolio at this time would likely have been larger than the risk contribution of a moderately risky mid-size European company. The stand-alone measure for Toyota may imply it is a good addition to the portfolio, while the portfolio-referent risk measure may suggest that a riskier, non-Japanese company is a better choice. The portfolio perspective accounts not just for an exposure's stand-alone risk, but also the correlation and concentration of that exposure in the context of a given portfolio. Typically the correlation across geographies is lower than the correlation across industries within any particular geography.

Consider a similar example in the United States: U.S. banks in the 1970s tended to hold portfolios of high-quality, large U.S. corporate borrowers. Some of these banks even characterized themselves as diversifying across industries, but did not validate their assertion of effective portfolio diversification in any objectively quantitative way. We now know that the risk contribution of one more large corporate borrower in the context of these U.S. bank portfolios was typically higher than the risk contribution of a small to mid-size company even if the smaller company was from an industry already heavily represented in the portfolio. (Typically the correlation across company size groups within the same industry is lower than the correlation across industries for a given company size group.) A portfolio perspective requires uncovering the underlying factors that drive correlation across the portfolio. (Sometimes this risk is labeled *systemic* or *systematic*.)

With this conceptual understanding of why the risk side of the equation should focus on the portfolio, we now introduce the notion of *return*.

At its core, the motivation behind putting money at risk is to earn some kind of return. We can measure both the cash payments and the change in value of a credit-risky security when calculating return. In the context of the framework we develop in this book, the value of a security or asset is a function of the size and likelihood of cash payments we expect to receive as a consequence of holding that security or asset. In order to place all securities in a portfolio on the same measurement scale, we specify a time horizon. The change in value of a particular security over that time horizon requires us to know the value at the starting time of analysis (often referred to as the *as-of date*) and the time horizon date in the future over which the risk analysis is performed. These valuation exercises require a model to convert

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

the characteristics of a particular security and its concomitant risk into a currency value.[6] Many default probability models are variants of valuation models.

Returning to our discussion: While the total return for a particular exposure is a useful first measure, we need to make two adjustments before we can draw any strong conclusions about a particular exposure. First, we must adjust the return for the time value of money. Conceptually, this means subtracting a measure of the risk-free return—that is, return earned from investing in a risk-free security. (In practice, we lean toward subtracting a measure of the cost of funds for the bank as that is the cost of securing the funds to put at risk.) Second, we must adjust for the amount we expect to lose. (Because there is credit risk associated with credit-risky securities and because there is an upper bound on the payoffs, the risk of a loss is always positive. This expected loss is the cost of running a credit business.) The result is a measure of return over a particular time horizon of analysis that is the premium earned for taking credit risk. The credit risk premium is what we expect to earn above and beyond the cost of funds and the exposure's expected loss.

Now we have the conceptual pieces for building a high-performance (lower risk/higher return) credit portfolio. We will always be faced with some constraints limiting the type of credit exposures that can be placed in the portfolio (e.g., limits on position sizes, availability of borrowers in some sectors, etc.). Subject to these constraints, we can compare the credit risk premium to the risk contribution for each existing exposure as well as each possible new exposure to determine which exposures to hold and which to sell out of the portfolio or buy protection on.

A final criterion for identifying a useful return or risk measure is that the measure can be coherently aggregated to characterize the health of the overall portfolio as well as subportfolios. For example, stand-alone risk cannot be coherently aggregated—in other words, a portfolio's stand-alone risk is not a simple weighted average of each of its exposure's stand-alone risks. Risk contribution, however, can be aggregated based on each exposure's weight in the portfolio. Expected return can also be aggregated based on the portfolio exposures' weights. In general, return measures are easier to aggregate than risk measures since return measures typically represent mean quantities, while risk measures typically represent higher moments of distributions (e.g, the ninety-ninth percentile). In this book, we focus on

[6] In credit we often convert the currency value into a spread, which is another way to represent the same value subject to several conditions. We discuss the notion of spread in detail several times throughout this book.

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

measures that meet this criterion. We introduce several different approaches for modeling each component of this framework. We also describe how to interpret and implement these measures in ways that will materially improve the performance of a financial institution actively managing its credit portfolio.

BUILDING BLOCKS OF PORTFOLIO RISK

Understanding the portfolio framework requires definitions of the key components used for credit portfolio analysis:

- Probability of default (PD): The probability that an obligor will not meet a stated obligation. In the case of a loan, the obligor is the borrower and the obligation is to pay a regular coupon and repay the principal at maturity. A PD will have a time horizon attached to it.
- Loss given default (LGD): The amount lost when an obligor fails to meet a stated obligation. Many times the focus is on recovery, or 1-LGD.
- Time horizon of analysis (H): Meaningful credit portfolio analysis requires the specification of a time horizon over which the analytics are calculated. Later we will be more specific with respect to the criteria for specifying H. Most analyses begin with the assumption that H is one year. Note that we often denote time with the letter T. In this book, we distinguish time to maturity as T from time horizon of analysis, which is H.
- Default correlation: The co-movement into default of two obligors.
- Value correlation: The co-movement in the value of the credit-risky securities within a portfolio.

With these definitions, we can sketch out the framework for evaluating a credit portfolio. Initially, we will determine expected loss, which is a primary cost of building a credit portfolio.

- Expected loss (EL): PD times LGD. This quantity is typically calculated over the time horizon, H. In this definition, we assume that the exposure at default (EAD) is par. This definition can be modified for other instrument types.
- Economic capital: The amount of (value) cushion available to absorb extreme losses—that is, absorb losses after using other sources of reserves such as provisions based on EL and earnings from exposures. The economic capital amount is calculated based on a target probability associated with an estimated portfolio loss distribution (estimated

for the time horizon, H). That is, the economic capital corresponds to the present value (of amounts at time H) of the loss level at which the probability of exceeding that loss level equals the target probability.

As we have alluded, it may be tempting to interpret EL as a measure of risk; however, it is better thought of as a measure of the cost of building credit portfolios. Then when we discuss capital as a cushion for unexpected losses, we have a clean separation of costs and capital.[7] The occasional surprise loss (or losses) becomes the focus of portfolio risk assessment. The following are two preferred measures of portfolio risk:

1. Unexpected loss (UL): A measure of the volatility or standard deviation for a portfolio loss distribution.
2. Tail risk (TR): A measure of the likelihood of extreme losses in the portfolio (this is similar to the concept of value-at-risk or VaR; we will also introduce the concept of conditional VaR or CVaR, which is sometimes referred to as *expected shortfall*). Tail risk corresponds to the area of the portfolio loss distribution from which we typically calculate economic capital.

Figure 1.2 shows a graphical depiction of a portfolio value distribution with an indication of the UL and TR. Note that this figure displays the value distribution. We often analyze a portfolio loss distribution, which is a linear transformation of the value distribution. Simply explained, to convert a value distribution to a loss distribution, we identify a loss point (i.e., the point at which we start counting losses), and subtract that point from each point in the value distribution. A typical candidate loss point is the risk-free value of the portfolio at the horizon date. We discuss these calculations in more detail in Chapter 8.

[7] It is important to highlight that in this discussion and throughout this book we focus on *economic capital*, which is reflected in the market capitalization of a financial institution. Economic capital typically differs from *book capital*, which is an accounting concept. Book capital is not really the value cushion available to absorb extreme losses; book capital reflects the accumulation of accounting entries that have a backward-looking bias. Whether the financial institution possesses the resources to absorb loss depends entirely on the current ability of the financial institution to make use of its equity's market value. Another type of capital results from regulations. This is known as *regulatory capital*. Calculation of regulatory capital results from an attempt on the part of regulators to determine a minimum cushion that will coincide with economic capital.

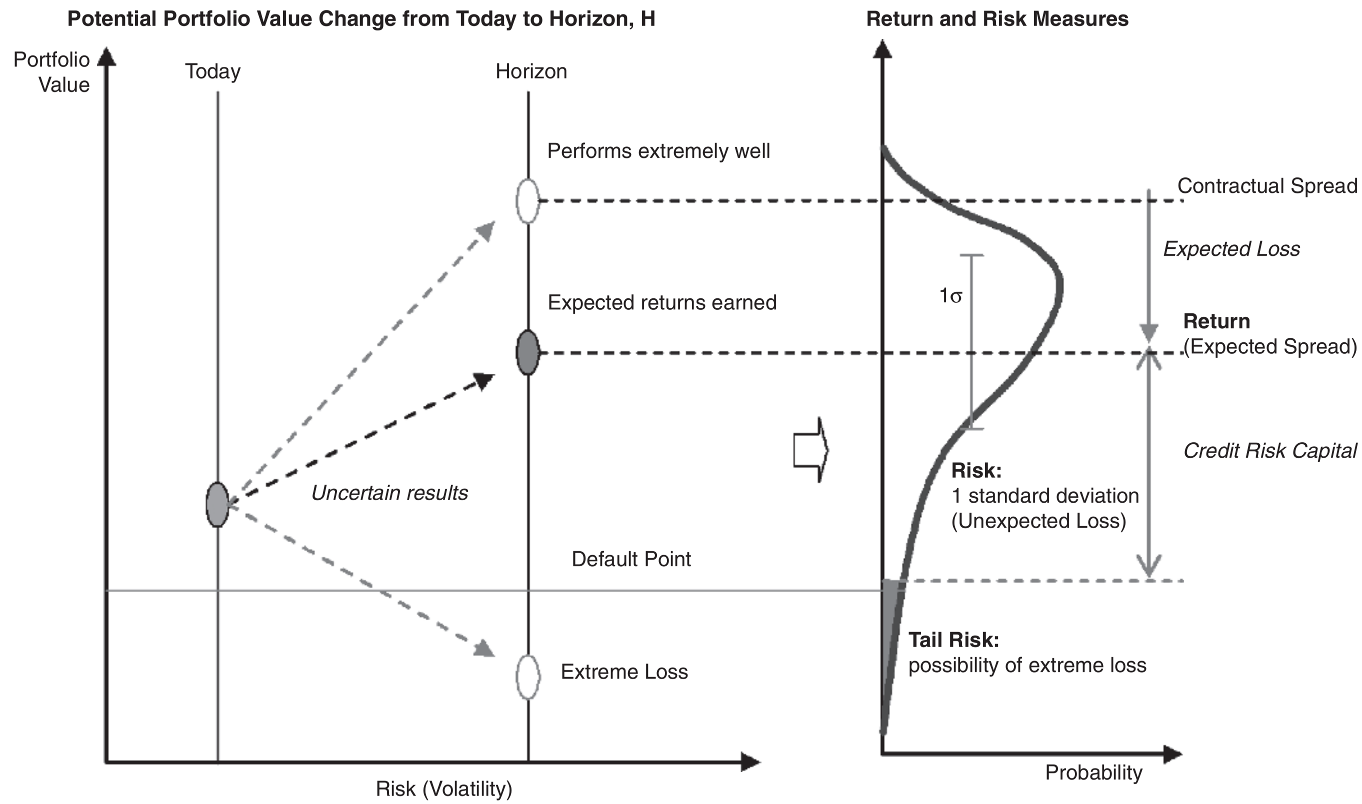

FIGURE 1.2 Portfolio Value Distribution with Unexpected Loss and Tail Risk

Both of these measures are essential for characterizing the risk of a credit portfolio, but they are necessarily summary statistics. In fact, the entire portfolio loss distribution contains important information about the risk of the portfolio; however, a financial institution with a credit portfolio needs to develop one or two analytics that can be communicated throughout the firm. Since tracking the entire loss distribution over time can be difficult in a conceptual sense (practically speaking, we can certainly calculate the portfolio loss distribution at different points in time; the difficulty arises in understanding the implications of changes in the distribution), focusing on measures such as UL and TR provides a current assessment of the portfolio risk as well as some historical context (e.g., the UL is higher or lower than before, which provides some sense of how the portfolio risk has changed).

Since we have only discussed UL and TR abstractly thus far, let us consider how the measures are interpreted in practice. Unexpected loss tells us something about the variation we expect to see in the size of most losses experienced by the portfolio. Since a large portion of a bank's earnings derive from the financial portfolio, this variation will directly impact the bank's earnings volatility. In this way, UL provides guidance as to how the composition of an existing portfolio will impact a bank's earnings' volatility. However, a portfolio that experiences little volatility, but every once in a while is hit with a loss so large so as to put the entire bank at risk, is not a portfolio a bank should be interested in holding. We turn to TR for a characterization of this extreme loss risk.

In recent years, the interpretation of TR has arisen from the ideas underpinning *value-at-risk* modeling, which is often called VaR. To better understand VaR in a credit context, we need to take a short digression to review the motivation behind a financial firm's target capital structure.

The owners of a financial institution's equity make use of debt to improve the return on their equity. Banks, for example, typically have a depositor base to provide fairly low-cost debt. But even if depositors are not the primary source of a bank's funding, the degree of leverage will typically affect a bank's credit quality, which in turn will determine how much a bank will pay (in terms of ongoing interest expense) for its debt. While this book tends to emphasize concepts in the context of banking institutions, the principles are equally relevant to any institution building a capital structure to support a credit portfolio. In theory, different levels of leverage in the bank's capital structure will directly impact the cost of funding for that bank. At some level of leverage, the cost of funding becomes so high that the bank cannot profitably employ those funds. Furthermore, at some level of high leverage, the credit quality of the bank will be so low so as to reduce or eliminate many profitable areas of business such as entering into derivative transactions, serving as custodian for assets, or generally providing services that require a strong balance sheet and a strong reputation.

Thus we have two countervailing motivations for deciding on the degree of leverage in a bank's capital structure. On the one hand, equity holders want to use as much leverage as possible to improve their return. On the other hand, these same equity holders realize that their ability to run a profitable business depends on maintaining a suitable level of credit quality, which constrains the desired degree of leverage. More fundamentally, the event of bankruptcy by the bank would cause the equity holders to lose all of their investment.[8] For these reasons, we can assume that at some point a bank can have too much leverage.

Though we do not have a coherent theory of optimal bank capital structure, we assume that a bank will desire to maintain a strong investment-grade level of credit quality to profitably construct a portfolio and build service businesses. While not rigorously verified empirically, our experience suggests that in many bank managers' views, this level does not need to be the highest (Aaa in the parlance of rating agencies) for all banks. That is, the cost of obtaining and maintaining a Aaa rating may exceed its benefit in some cases. Casual observation of market spreads for debt issued by different financial institutions suggests that a strong A borrower often pays about the same spread as a Aa or Aaa borrower in many settings. In a similar vein, Aaa, Aa, and strong A borrowers appear to have the same kind of access to banking-related service businesses. The extra capital needed to achieve a Aaa rating may materially impact the bank's overall return on equity while not necessarily changing its funding cost structure or range of business opportunities.[9] Figure 1.3 provides more intuition with respect to the relationship of economic capital and rating on the bank's debt.

A related issue that has taken center stage in the recent credit crisis of 2007 and 2008 is the benefit of having diversified sources of funding. In the months and years before summer 2007, a financial institution could often fund itself entirely in the wholesale finance market, borrowing directly from other financial institutions or investors. Many of the institutions offering this type of funding have either stopped doing so or been acquired. Surviving

[8] Some debate continues as to whether holding bank equity in a portfolio of investments changes the incentives so that the risk of bankruptcy of any one bank is overshadowed by the possibility of better returns across a number of bank equity investments.

[9] Note that certain businesses require that counterparties be rated Aa or Aaa. For example, some investors require Aaa-rated securities for their investment portfolios. Banks interested in providing guarantees on instruments that target such investors may require a Aaa rating to be able to do so.

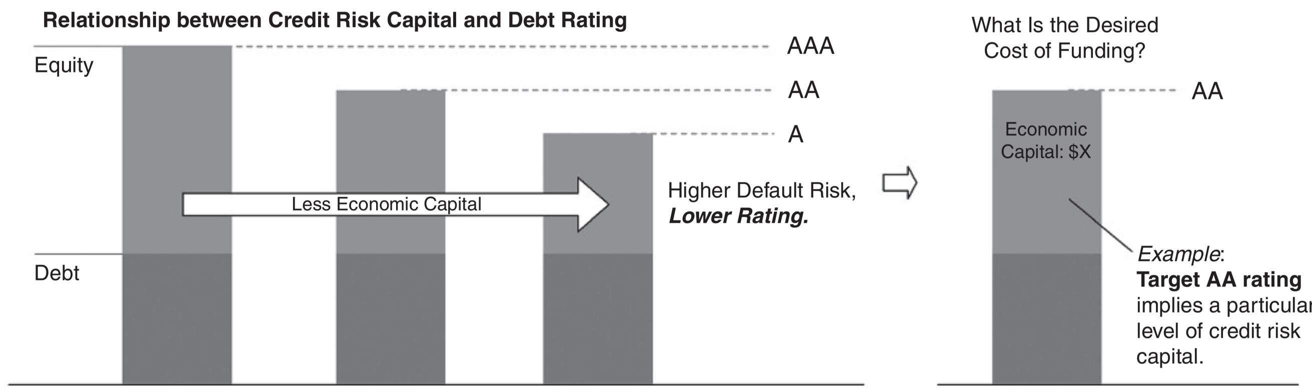

FIGURE 1.3 Relationship between Credit Risk and Capital and Debt Trading

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

financial institutions maintained a diversified mix of funding sources anchored with deposit-based funding. While we are not aware of any systematic research at this time, we suspect that in some cases, a higher rating such as Aa or even Aaa will make a difference going forward in a bank's ability to access funds at reasonable (i.e., profitable) cost. The key point is a bank's management should think carefully about its overall objectives and contingencies when considering what its target probability of default should be.

Why have we taken this digression in our discussion of VaR and TR to discuss bank capital structure? If we accept the assumption that a bank targets a capital structure that results in a strong (but not necessarily the strongest possible) investment-grade credit quality, we can interpret bank equity as the cushion available to absorb portfolio losses. In other words, equity value can be considered a reflection of the bank's available economic capital. The size of this equity or economic capital cushion makes the credit quality of the debt dependent on the likelihood that the portfolio will suffer a loss so large as to exhaust all the available equity value. This likelihood is a direct function of the bank portfolio's composition—that is, which securities or assets are held and in what quantities. If we convert this likelihood into a specific target probability, we can use our portfolio loss distribution to help determine which loss threshold corresponds to a specific target probability of exhausting all available economic capital. The same is true of other financial institutions.

Assuming that a particular probability of exhausting all capital is the same as the probability of a financial institution defaulting on its outstanding debt, we can convert a given target probability into a simplified "rating" for easier interpretation. For example, using a market-based measure of default probability, a one-year target probability of 0.15 percent is roughly associated with an A rating.[10] If a bank's management concludes that an A rating is sufficient to run a profitable business given its portfolio composition, then it will adjust its leverage (i.e., issue more or pay down debt) such that the probability of exhausting its equity is 0.15 percent (or 0.02 percent or another level, depending on the bank's view of which PD maps into which risk category). In this way, the portfolio loss distribution leads to a VaR interpretation of the risk for the debt issued by the bank. Please refer to Figure 1.4 for a graphical characterization of this relationship. Note that

[10] We base this statement on our research using MKMV's Expected Default Frequency (EDF) measure of the probability of default. The typical one-year EDF for the past 10 to 15 years for an A-rated borrower has been about 0.15 percent. However, this depends on the measure chosen. For example, the historical average issuer weighted historical one-year default rate for a Moody's A-rated instrument was about 0.02 percent between 1986 and 2006.

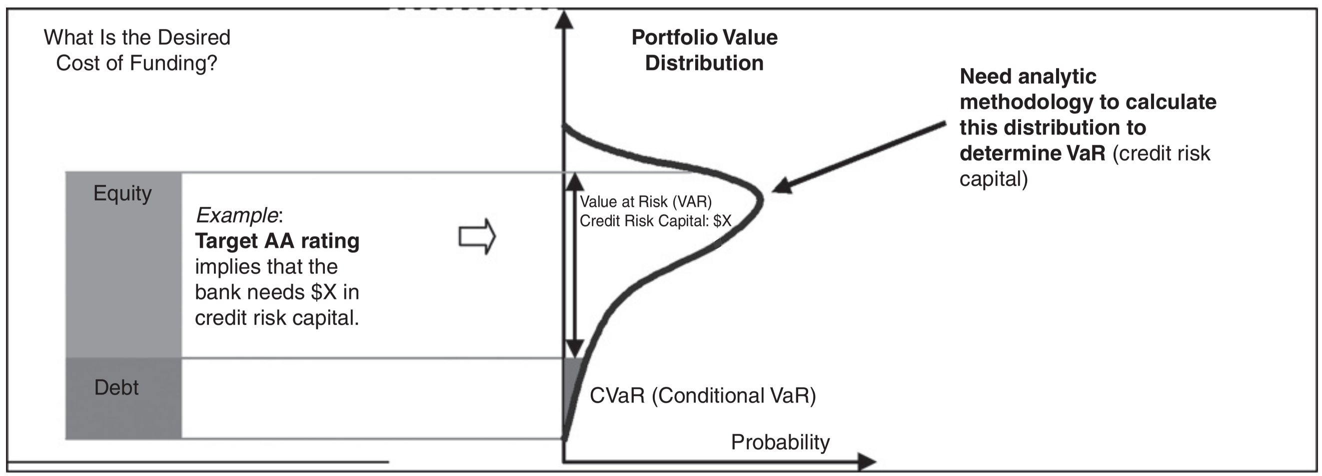

FIGURE 1.4 Relationship of Simple Bank Capital Structure to the Value at Risk (VaR) Characterized by a Target Probability for Portfolio Value That Indicates Bank Insolvency (Equity Value of Zero)

the target probability matches the area under the curve that is shaded in at the extreme edge of the value distribution. Portfolio value realizations that are below the beginning of the debt in this simple example will correspond to the equity losing all its value. As we have stated, the time horizon of analysis is typically one year.

Tail risk is what we measure to determine how a particular exposure contributes to this extreme portfolio loss event at a particular target probability. Thus far, we have not been too rigorous about defining TR. In Chapter 8 we distinguish TR calculated from conventional VaR (i.e., the contribution to the likelihood of a portfolio loss exceeding a particular threshold, which is associated with a specific target probability) from *conditional VaR* or CVaR (sometimes called *expected shortfall*—the contribution to the amount expected to be lost conditional on the portfolio loss being beyond the threshold). Tail risk calculated based on VaR does not distinguish between losses that may be considerably beyond the target-probability threshold. Tail risk calculated based on CVaR, by contrast, tells us whether the loss will likely be just a bit beyond the threshold or whether it will substantially exceed the threshold. We explore the details of this difference in Chapter 8. In recent years, various researchers have raised concerns regarding the coherency and usefulness of VaR-based measures. As a result, a number of practitioners are now inclined to focus on CVaR.

Once a target probability is chosen, the amount of capital required can be calculated by examining the quantile of the loss distribution corresponding to this probability. (Note that this approach assumes that the target quantity is a PD. In some environments, targets are described in terms of expected loss, in which case a more complicated calculation is required.) Given the preceding discussion, we find it convenient to consider bank equity value as the measure of capital that the bank needs to ensure a particular target probability of capital exhaustion given the composition of the bank's portfolio.[11]

Once we have determined an aggregate capital amount for a given portfolio, we can next look to allocate that capital to each exposure in the portfolio. For example:

- If we are focused on UL or volatility as our measure of risk, we may choose to base capital allocation on a given exposure's contribution to

[11] Note that as we move out into the far tail, caution is warranted since the drivers of extreme losses may be different than in our general model assumptions. Stress-testing should be considered along with correlation models for extreme losses.

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

UL, which is sometimes called *risk contribution* (RC). This approach will favor exposures that do not vary in value much over time, such as loans to large, high-quality corporate borrowers.

- Alternatively, we may choose to allocate capital based on a given exposure's contribution to portfolio TR, which is sometimes called *tail-risk contribution* (TRC). This approach will favor exposures that do not have correlation with the portfolio or do not represent a large concentration.

  Ultimately, the choice of objective function is a management decision, so it is important that the management of a financial institution all share a comfortable understanding of the underlying meaning of each of these measures.

USING PDs IN PRACTICE

In the previous section, we provided a first introduction to the language of credit portfolio risk. Now consider an example of how this language can be used in practice. When a financial firm shifts to a quantitative orientation of risk assessment, PDs become the foundation for evaluating and monitoring obligors. Probabilities of default move considerably through the credit cycle: From trough to peak of the cycle, the typical market-based PD of a public firm in a sector or geography may change by as much as a factor of five or six. Probabilities of default also vary considerably across different classes of obligors.[12]

This change over time and change across obligors at any given point in time make it important to estimate PDs accurately. The first step in doing so is to rank current and prospective obligors in terms of their PDs. In most banks, PDs are converted into internal ratings. This bucketing of PDs creates more stability in the estimates and facilitates better communication with nonquantitative employees at the bank. (For those interested in bucketing PDs, we discuss a number of approaches to doing so in Chapter 4.) However, it also discards a good deal of information on the credit quality of the borrower.[13]

[12] For example, MKMV's EDF values (one commonly used measure of PDs) range from 1 basis point (0.01 percent) to 3,500 basis points (35 percent).

[13] One of the authors was once speaking with a banker who complained that PDs moved much more frequently than ratings. The author responded, "You don't have to look at them every day." By bucketing PDs, banks effectively do this.

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

EVOLVING APPRECIATION FOR CREDIT MEASURES IN BANKS

As a bank develops more experience with PDs, the tendency is toward converting them into financial values such as bond prices or CDS spreads. This value orientation is the primary focus of credit hedge funds and other teams that trade credit.* Note that the first step of simply ranking by PD provides a rough way of determining which credit exposures are more likely to get into trouble. Converting a PD into a value offers an extra dimension that facilitates comparison to existing market prices since it describes in financial terms the *cost* of the risk. Eventually, the most sophisticated users of these analytics will incorporate both the physical PDs and their corresponding valuations into a portfolio model by introducing exposure weights and correlations.

*For valuation applications, the physical PD must be converted into a *risk-neutral* PD and an LGD needs to be estimated. We discuss valuation, risk-neutral measures, and so on in more detail in a later section.

Probabilities of default are an integral, if not the most important, part of a portfolio model. Using PDs, an analyst can not only make a relative statement about obligor risk (e.g., this small manufacturing company is a better risk than that medium-size business services firm), but he can also quantify the differences (e.g., this small manufacturing company is one-third as likely to default as that medium-size business services firm). Layering in the portfolio view facilitates analysis in the context of diversification (e.g., while the medium-size business services firm is much riskier than the small manufacturing company on a stand-alone basis, in the context of this European bank portfolio both firms contribute about the same level of risk).

The final piece of the puzzle concerns return: How much will a bank earn from each credit exposure? Valuation helps sort out this question. We can use PDs in the context of a valuation model to determine how much an exposure is worth today and how much we expect it to be worth at a particular date in the future. This expected gain or loss can be added to the expected cash flow stream to determine the overall expected return from holding a particular credit security. We can also incorporate these models into simulation engines so that we determine values conditional on the realization of various factors that reflect a particular state of the economy. Probabilities of default are the threads that run throughout the fabric of all these risk models.

In evaluating the quality of various PD models, we typically make a distinction between a model's power and its calibration. The power of a PD model tells us how well it distinguishes good from bad borrowers. In Chapter 7, we describe evaluation tools known as *power curves* that statistically characterize a model's power. A powerful PD model is not necessarily well calibrated. A well-calibrated model produces probability estimates that agree well with the observed default rates for obligors. If the model is used only for ranking the risk of obligors, poor calibration will not make a large difference as long as the model properly distinguishes the obligors. However, if the objective is to allocate economic capital or to value a security or exposure, both power and calibration matter. A poorly calibrated model will produce values that will not correspond to market prices. Moreover, portfolio analyses may be skewed if the valuations do not conform at all to existing market prices.

VALUE, PRICE, AND SPREAD

In the course of our model discussions, we will return often to the theme of valuation. The objective of most credit analyses in a portfolio context is to determine how the value of individual exposures and the overall portfolio will change over time. An important concept in this context is the transformation of actual or physical or actuarial PDs, which measure the actual probability of default events into risk-neutral or risk-adjusted PDs, which measure the probability of default implied by market prices after accounting for uncertainty. When assessing risk, we estimate the actual, *physical* PDs. When we estimate a credit instrument's value, we also use PDs; however, we first convert these actual PDs into *risk-neutral* PDs. While the physical PD represents the estimate of the expected rate of default for a particular entity, this conversion to risk-neutral PDs for the purpose of valuation reflects the inherent risk aversion of investors. This adjustment reflects the fact that risk-averse investors require an extra premium beyond payment for expected loss to compensate them for the risk associated with purchasing a security that may lose value in excess of everyone's expectations. In other words, investors require additional compensation for accepting a gamble versus a sure thing with the same expected payout. The change of measure from physical to risk-neutral PDs can be accomplished in a variety of ways, but the key objective is to adjust the probability for the (extra) risk that needs to be implied in order to adjust for risk aversion. We will be careful to distinguish actual from risk-neutral PDs.[14]

As we stated before, the value of a credit instrument derives from the likelihood of receiving cash payments over time in the future. This value may be reflected in a market price, in quotes provided by market participants, or it may be the output of a model. Unlike accounting measures, valuation-based measures can fluctuate substantially. For example, if in general the market is demanding a higher premium for credit risk due to pessimism about the economy, *all* of the loans in a bank's portfolio may be affected.

In several of our discussions, we will talk about *value*, *price*, and s*pread*. These terms have multiple definitions and tend to be used in ambiguous and sometimes confusing ways in the finance literature. In this section, we define these terms as we will use them in this book.

When we discuss *value*, we refer to a measure that reflects the model framework's assessment of the present value of the future risk-adjusted cash flows expected to be generated by the asset or security under analysis. *Price*, by contrast, is the amount at which an asset or security is bought or sold in a market. The asset or security's value does not necessarily equal the price at which one could buy or sell the asset or security at any given point in time. Instead, value provides an indication, per a model, of the price at which the asset or security *should* be bought or sold, but which may not reflect the actual current market price due to various market conditions. A useful model will generate values to which prices converge over some reasonable time horizon (usually less than a year). In some cases, the asset is not traded (e.g., the market value of an entire firm[15] or the bank loans of a small firm) and in some cases the security's price is driven by factors outside a model (e.g., market liquidity effects). These are two examples of many in which value may differ from price. Of course, the model generating a particular value may be wrong and never converge to any observable price. (We recommend rejecting models to which relevant prices never converge.)

While most people tend to think in terms of a price or value in currency terms (which for credit is typically quoted on a scale of 100—thus, a price that is 98 percent of par value would be quoted as "98"), we typically

[14] If you are not familiar with this concept, please refer to one of the financial texts referenced in this book. Hull (2005) and Neftci (2000) are good choices.

[15] In many instances throughout this book, we refer to "securities" and "assets" in a portfolio, which we use to mean a financial instrument. We also refer at times to the "asset value" of a firm, which reflects the value of the firm's entire enterprise. An asset in a portfolio is in most cases different than the assets of a firm (unless the portfolio is composed of owning many firms in their entirety). Readers should be aware of this distinction when they encounter the term *asset*.

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

convert price or value into a spread relative to a risk-free benchmark. The word *spread* is perhaps the most overloaded (and sometimes abused) word in finance. For example, the difference between a bid and an offer on a security is called a spread; however, we do not focus for the most part on this type of spread and when we do, we refer to this type of spread as the *bid-offer spread*. In contrast, most of the time when we refer simply to a *spread*, we generally mean the extra premium with respect to a reference benchmark (e.g., a risk-free rate) that represents a suitable conversion from the benchmark to the security's value or price.

To calculate a spread, we must first specify a risk-free benchmark. While most of the finance literature adopts the U.S. Treasury curve or possibly another reserve-currency sovereign such as the UK gilt curve, we maintain that credit models are best fit to spreads relative to a corporate-risk-free rate. This corporate-risk-free rate is the rate at which a risk-free corporate borrower could borrow. (While this type of borrower does not exist in practice, we can consider near risk-free borrowers as a benchmark.[16]) The swap curve is a good first approximation to this curve.

Once we have determined the appropriate risk-free benchmark, r_f, we can convert a price or value into a spread. We do this conversion in the context of a particular model for the price or value of debt, D. We can sometimes solve for spread analytically (e.g., in the case of zero-coupon bonds); but for most debt securities—particularly ones with coupon, C—we must find the spread, s, such that the functional relationship holds: $D = f(r_f, s, T, C, \text{etc.})$

In the case of callable or putable bonds, we incorporate the optionality into the calculation to estimate what is termed an *option-adjusted spread* (OAS). The OAS calculation generally requires some kind of *lattice* or *tree* construction depicting the possible paths and path probabilities for the debt value. If the debt security does not have optionality attached to it, fitting the spread in the context of a lattice construction results in what is termed a *Z-spread*.

In the context of the models in this book, we are careful to distinguish the following spreads:

- *Total spread (TS)*. We usually call the spread to the U.S. Treasury (or some similar reserve-currency sovereign curve) the total spread. This spread includes compensation for a variety of risks extending beyond just credit risk. In a portfolio modeling context, we define TS relative

[16] At MKMV we coined the term *0-EDF rate* to describe this benchmark. The 0-EDF curve seems to be 10 to 20 basis points less than the swap curve (see Bohn 2000b). We can think of this curve as 0-PD or rates for a corporate borrower that has a PD of 0.

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

to a 0-EDF curve (see footnote 16) or swap curve.[17] If we define TS in this way, the difference between the benchmark and the U.S. Treasury curve can be considered compensation for less liquidity in the corporate debt market.

- *Expected spread (ES).* As we discuss in Chapter 8, we subtract expected loss (EL) from TS to arrive at ES. This spread reflects the premium for holding credit exposures. This spread may comprise both a credit component and a liquidity component specific to the firm under evaluation.
- *Credit spread.* Part of the challenge in credit modeling is isolating that part of the spread compensating credit risk only. Strictly speaking, we define the credit spread as that portion of the total spread related just to credit risk. It is composed of the EL and the premium earned for holding credit risk.
- *Liquidity spread.* Unfortunately, we do not have good models of liquidity, so this spread tends to be the residual of the ES that cannot be explained by the credit model. Conceptually, there is likely to be some premium earned for exposure to liquidity risk. In practice, we may inadvertently characterize portions of the TS or ES as liquidity spread when in fact our model is misspecified.
- *Zero-coupon spread.* When discussing term structures of spreads, we normally characterize the zero-coupon spread as that which would be earned on a zero-coupon debt security at different maturity dates.
- *Par spread.* An alternative way to characterize the term structure of spreads is in terms of the spread earned on a floating-rate bond such that the price of the bond is equal to par—that is, equal to 100. We call this characterization the par spread.

We will be explicit about which type of spread is calculated when discussing model calculations. In most modeling discussions, we will refer to the spread associated with a debt security. Spreads tend to put all securities on a similar measuring scale for comparison. While not a perfect measure even in a credit-modeling context (it is sometimes better to think about a term structure of spreads), most practitioners focus on spreads and they tend to be a good way to refer to the output of credit valuation models. We sometimes distinguish spreads that come from models (i.e., *value-based* spreads) and spreads that come from market prices (i.e., *price-based* spreads). In a portfolio context, we tend to focus on ES to determine the relevant marginal

[17] The reference to a *swap* here refers to swapping a fixed-rate obligation for a floating-rate obligation (or vice versa). Swap curves define the rates at which an approximately AA obligor can borrow. While some credit risk is reflected in these rates, they are generally good proxies for the corporate risk-free or 0-PD rates.

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

return we expect to earn relative to the risk contributed to the portfolio. Expected spread is a common component of metrics used to evaluate a credit portfolio's performance.

DEFINING DEFAULT

Central to the various discussions in this book is the concept of an event of default. A simple definition of default is the nonpayment of interest or principal on a credit obligation. We have found, however, that sometimes default can be usefully defined in different ways. Another possible definition of default is bankruptcy. For example, a firm may behave differently in a bankruptcy resulting in full liquidation of its assets than if it just defaulted on one or two outstanding debt issues. This may, in fact, be mandated by law. The difference in behavior may materially impact the recovery obtained in the event of default. Thus, in such a context it may be useful to examine the difference between probability of bankruptcy and probability of default.

Situations can also arise in which a firm appears to have defaulted in an economic sense, but not in practical terms. Consider a bank that becomes insolvent, but whose country's government continues to inject funds to keep the institution running. (Recall Japan in the 1990s.) This bank is economically in default while still making good on its obligations. Firms may also restructure their debt in such a way as to adversely affect the value of their debt without actually defaulting. This restructuring may be interpreted as a technical default despite the fact that interest and principal continue to be paid.

Further complicating definitional issues is that conventions may differ from country to country regarding when a firm is considered to be in default. For some countries, the convention may be to interpret any missed payment as an immediate default while other countries may wait until the payment is 90 or 120 days late.

It is important to understand clearly the definition of default when interpreting model output and evaluating model performance. Definitions may differ depending on the circumstances. As long as the definition is consistent with the procedures taken to estimate and use the model in question, the definition of default may vary without difficulty.

PORTFOLIO PERFORMANCE METRICS

We have touched on a measure of the extra premium or expected spread (ES) earned to hold a risky security in a portfolio. We have also characterized contribution to portfolio volatility, or risk contribution (RC), and contribution to portfolio tail risk (TRC). The RC and TRC can be used to allocate economic capital, usually defined as a capitalization rate (CR), to a particular exposure. Finally, we characterized the portfolio's stand-alone risk or unexpected loss (UL) and a particular portfolio's target probability–based economic capital (C). The question remains: How do we use these measures in practice? While we delve into these measures in more detail toward the end of this book, we introduce them briefly here to motivate the material that follows.

Performance cannot be managed if it is not measured and benchmarked. In the equity portfolio management business, performance is measured in terms of return per unit of volatility, known as the Sharpe ratio (named after the Nobel laureate, William Sharpe). In the case of equity portfolios, the Sharpe ratio is compared to an index that matches a portfolio manager's style of equity investing and the manager is evaluated based on whether he over- or underperformed relative to the index. While many retail investors still tend to look only at historical return, sophisticated investors evaluate equity portfolios on the basis of their return per unit of volatility relative to an appropriate benchmark.

In the world of credit, notions of return per unit of risk and benchmarks are still quite new. Nonetheless, the Sharpe ratio measure has started gaining popularity for credit portfolios. To port the Sharpe ratio to the credit context, analysts make a few modifications. Return is measured in terms of ES: not just the risk-free rate but also expected loss is subtracted from total return. The denominator for the portfolio is UL.[18] (At the exposure level, UL is replaced by RC.)

Once the Sharpe ratio is calculated, though, a more fundamental difficulty arises in that there is a lack of an appropriate benchmark. While many credit indexes do exist, they tend to suffer from the endemic problem that the market does not originate well-diversified portfolios. Unlike in the equity world, where indexes can be constructed that more or less diversify away idiosyncratic risk, in the credit world, a borrower such as Ford may account for several percentage points of the market portfolio in exposure value terms. Thus the benchmark is overshadowed by the idiosyncratic risk of its largest constituents. Traded CDS indexes are rapidly developing into benchmarks which may not suffer from the diversification problems of current bond indexes. That said, the issues' changing liquidity and their relative

[18] Note that the calculation of the Sharpe ratio for a credit portfolio maintains the conceptual underpinning of the traditional Sharpe ratio: return divided by risk. However, in our context, we replace the traditional measures of excess return and the standard deviation of returns with measures appropriate to credit risk.

newness make them not always the most reliable choice at present, though they will likely eventually develop into benchmarks in the future.

Where does this leave us?

First, we can begin with the portfolio's current Sharpe ratio as a reasonable target for motivating each transaction within the bank. In this way, each new transaction should earn at least as much return per unit of volatility risk as the current portfolio.

An alternative benchmark is the Sharpe ratio for the "best" possible[19] portfolio that the firm can construct given its constraints. All financial institutions have constraints governing what and how much they can own. This benchmark pushes the institution toward building the best possible portfolio available to it. Each transaction's Sharpe ratio (calculated in this case as ES/RC) is compared to this best-case portfolio's Sharpe ratio (calculated as ES/UL). This approach to decision making works for motivating value-adding transactions.

For single transactions and subportfolios, the performance can be measured against the larger portfolio or the best-case portfolio. But how can we benchmark the larger portfolios? The lack of good market credit indexes requires construction of something like a well-diversified market portfolio. Though we will provide some ideas and guidance based on what is best practice today, this field of benchmarking still leaves much to be desired.

An alternative measure of performance to the Sharpe ratio is something we call the Vasicek ratio (named after its developer, Oldrich Vasicek, a pioneer in the field of modern quantitative finance). A Vasicek ratio is analogous to a Sharpe ratio in the sense that it is a measure of return per unit of risk. But in this case, we replace the volatility-based denominator with a capital number reflecting a calculation of the portfolio tail risk. In the early 1990s, a similar measure was developed at Banker's Trust, called *risk-adjusted*

[19] We have purposely refrained from using the word *optimal* as it is still difficult to construct an optimal credit portfolio. We discuss this further in Chapter 8. At this point, we should clarify the difference between *ex ante* and *ex post* measures of performance. In this discussion, we have tended to focus on *ex ante* measures, which are reflected in expected performance over a given time horizon. We recommend using *ex ante* measures to drive efficient decision making. *Ex post* measures describe what actually happened over a given time horizon. Here we must focus on realized return and realized volatility. Several practical difficulties arise when calculating realized quantities that we discuss in more detail later in the book. The point to remember here is that while portfolio construction can only be done based on *ex ante* measures, employee compensation should be paid on realized (*ex post*) performance relative to a benchmark. This motivates the firm to ensure that the credit decision-making tools it uses are as accurate and realistic as possible.

return on capital (RAROC). The credit modeling firm KMV introduced a variant of this measure in the mid-1990s called *return on risk-adjusted capital* (RORAC). (The authors have even come across *risk-adjusted return on risk-adjusted capital* (RARORAC).) The "rocks" and "racks" measures tend to include the risk-free return component of asset value growth, and many of these metrics mix in a variety of expenses and capital calculations that can sometimes distort decision making. In the end, most of these adjustments tend to diminish the usefulness of the measure.

The Vasicek ratio is much more straightforward. It is defined as ES/C, so the estimation effort stays focused on properly calculating the capital, or C. We will have much to say about components of this calculation throughout this book. The conceptual point to take away from this discussion is that the Vasicek ratio provides some sense of the return earned on capital that has been put at risk to support a particular transaction or portfolio. A credit portfolio manager still faces the difficulty of choosing the proper benchmark and, similar to results using the Sharpe ratio, the best choices for a benchmark would appear to be the entire portfolio's Vasicek ratio or the best-case portfolio's Vasicek ratio.

A further benefit of the Vasicek ratio is the ease with which this concept of return on economic capital can be extended to nonportfolio businesses within the firm. In Chapter 2 we will discuss how a transfer pricing mechanism within the firm can centralize the management of credit risk in an active credit portfolio management (ACPM) function. This differentiation enables the firm to measure separately capital allocated to the portfolio and capital allocated to each of the nonportfolio businesses. The operating cash flow earned by these nonportfolio businesses becomes a measure of return. One can then calculate a business-based Vasicek ratio, assuming the capital allocation can be properly estimated. In this way, a common notion of return to capital informs performance evaluation and strategy discussions throughout the firm.

A related metric to the Vasicek ratio and one that can also serve as a common measure for evaluating portfolio and nonportfolio businesses is *shareholder value added* (SVA), or economic profit. Shareholder value added is calculated by first calculating a net income number adjusted downward for expected loss (similar in concept to the calculation of ES). Next, the cost of capital allocated against a portfolio or business (defined as allocated capital times the cost of that capital) is also subtracted from the net income. The residual is the economic profit or SVA of that activity.

While the Vasicek ratio provides some sense of the percentage return earned on capital allocated, SVA provides a sense of the absolute level of contribution to shareholder value. A business may have a high Vasicek ratio but be so small as to not materially increase shareholder value. In contrast,

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

a lower Vasicek ratio business may have a high SVA given its size. In our view, though, the Vasicek ratio provides a sense of the ongoing opportunity associated with a particular business, as a high growth rate will eventually result in a high SVA (all else equal). Both the Vasicek ratio and SVA are useful tools for benchmarking performance.

DATA AND DATA SYSTEMS

Many of the modeling choices we recommend follow from our attempts to accommodate real-world data limitations. One of the primary factors behind the gap between many models introduced in the financial literature and those used in practice is the difficulty researchers have in finding sufficient data to estimate some academic models. Over time the models driving risk analytics have become more widely accepted and in some cases, somewhat commoditized. However, the manner in which the models are implemented in systems and the manner in which the organization changes to make use of the models continue to be debated. We will illustrate in several different contexts how we think data should influence this debate.

In our judgment, it makes sense to have a bias toward market observable information. While private information about management plans or new products and other qualitative analysis can be useful supplements, we have been firmly convinced over years of work with dozens of clients using a wide variety of analytic approaches that quantitative, observable data (as long as the historical period is sufficiently long) provides the best starting point and foundation for modeling in almost every case. This does not mean that qualitative analysis has no value. Rather, we find that it is hard to achieve consistent and repeatable performance across a large portfolio using subjective measures alone. Given this bias, we tend to choose and recommend models that rely on market data.

In circumstances where market data is not available, such as when modeling private small and medium enterprises (SMEs), our experience suggests large pools of data from several different sources constitute the best data foundation from which to start. Some banks erroneously believe that their own portfolio is large enough so that they have sufficient internal data to build robust credit models. In the process of collecting data for corporate (as well as retail) credit modeling, we have found that multiple banks' data provided substantial improvement in out-of-sample performance of default probability models. We discuss an example of this in Chapter 7. In the absence of market data, we have found that the best approach relies on collecting a diverse and large sample of data that usually requires collaboration among several institutions and/or government entities.

P1: a/b P2: c/d QC: e/f T1: g
c01 JWBT038-Bohn February 16, 2009 13:16 Printer: Yet to come

Significant investment in data collection, data cleaning, and data systems will be a common implication of our recommendations. No matter how robust a model appears or how confident an analyst is of a particular approach, bad data will quickly render the model and system useless. Since credit events happen so infrequently, the modeling exercises in this context tend to be focused on outliers. Any system whose calibration is heavily impacted by outliers becomes highly sensitive to data quality. We will demonstrate in several contexts the importance of this point and how to practically address it.

With this introduction as background, we now turn to the question of how an ACPM function is implemented in practice. This chapter and Chapter 2 provide the overall context for the model and system discussions that follow in subsequent chapters.

REVIEW QUESTIONS

1. What distinguishes credit risk from market risk?
2. What distinguishes credit risk from liquidity risk?
3. What is an asset-value driven default?
4. What is a liquidity-driven default?
5. Explain why the nature of credit instrument return distributions make them difficult to hold in isolation (i.e., not in portfolios).
6. What is the purpose of holding capital?
7. How does economic capital differ from book capital and regulatory capital?
8. Why are credit portfolios more amenable to active management than are equity portfolios?
9. Generally speaking, over which dimension should a credit portfolio be diversified and why?
10. Explain the difference between unexpected loss (UL) and tail risk (TR).
11. Describe the motivation behind the creation of a CDO.