

Spread Risk and Default Intensity Models

This chapter discusses credit spreads, the difference between risk-free and default-risky interest rates, and estimates of default probabilities based on credit spreads. Credit spreads are the compensation the market offers for bearing default risk. They are not pure expressions of default risk, though. Apart from the probability of default over the life of the security, credit spreads also contain compensation for risk. The spread must induce investors to put up not only with the uncertainty of credit returns, but also liquidity risk, the extremeness of loss in the event of default, for the uncertainty of the timing and extent of recovery payments, and in many cases also for legal risks: Insolvency and default are messy.

Most of this chapter is devoted to understanding the relationship between credit spreads and default probabilities. We provide a detail example of how to estimate a risk neutral default curve from a set of credit spreads. The final section discusses spread risk and spread volatility.

7.1 CREDIT SPREADS

Just as risk-free rates can be represented in a number of ways—spot rates, forward rates, and discount factors—credit spreads can be represented in a number of equivalent ways. Some are used only in analytical contexts, while others serve as units for quoting prices. All of them attempt to decompose bond interest into the part of the interest rate that is compensation for credit and liquidity risk and the part that is compensation for the time value of money:

Yield spread is the difference between the yield to maturity of a credit-risky bond and that of a benchmark government bond with the

same or approximately the same maturity. The yield spread is used more often in price quotes than in fixed-income analysis.

i-spread. The benchmark government bond, or a freshly initiated plain vanilla interest-rate swap, almost never has the same maturity as a particular credit-risky bond. Sometimes the maturities can be quite different. The *i*- (or interpolated) spread is the difference between the yield of the credit-risky bond and the linearly interpolated yield between the two benchmark government bonds or swap rates with maturities flanking that of the credit-risky bond. Like yield spread, it is used mainly for quoting purposes.

z-spread. The *z*- (or zero-coupon) spread builds on the zero-coupon Libor curve, which we discussed in Section 4.2. It is generally defined as the spread that must be added to the Libor spot curve to arrive at the market price of the bond, but may also be measured relative to a government bond curve; it is good practice to specify the risk-free curve being used. Occasionally the *z*-spread is defined using the forward curve.

If the price of a τ -year credit-risky bond with a coupon of c and a payment frequency of h (measured as a fraction of a year) is $p_{\tau,b}(c)$, the *z*-spread is the constant z that satisfies

$$p_{\tau,b}(c) = ch \sum_{i=1}^{\frac{\tau}{h}} e^{-(r_{ib}+z)ih} + e^{-(r_{\tau}+z)\tau}$$

ignoring refinements due to day count conventions.

Asset-swap spread is the spread or quoted margin on the floating leg of an asset swap on a bond.

Credit default swap spread is the market premium, expressed in basis points, of a CDS on similar bonds of the same issuer.

Option-adjusted spread (OAS) is a version of the *z*-spread that takes account of options embedded in the bonds. If the bond contains no options, OAS is identical to the *z*-spread.

Discount margin is a spread concept applied to floating rate notes. It is the fixed spread over the current (one- or three-month) Libor rate that prices the bond precisely. The discount margin is thus the floating-rate note analogue of the yield spread for fixed-rate bonds. It is sometimes called the *quoted margin*.

Example 7.1 (Credit Spread Concepts) Let's illustrate and compare some of these definitions of credit spread using the example of a U.S. dollar-denominated bullet bond issued by Citigroup in 2003, the $4\frac{7}{8}$ percent

fixed-rate bond maturing May 7, 2015. As of October 16, 2009, this (approximately) $5\frac{200}{360}$ -year bond had a semiannual pay frequency, no embedded options, and at the time of writing was rated Baa1 by Moody's and A- by S&P. These analytics are provided by Bloomberg's YAS screen.

Its yield was 6.36, and with the nearest-maturity on-the-run Treasury note trading at a yield of 2.35 percent, the yield spread was 401 bps.

The i -spread to the swap curve can be calculated from the five- and six-year swap rates, 2.7385 and 3.0021 percent, respectively. The interpolated $5\frac{200}{360}$ -year swap rate is 2.8849 percent, so the i -spread is 347.5 bps.

The z -spread, finally, is computed as the parallel shift to the fitted swap spot curve required to arrive at a discount curve consistent with the observed price, and is equal to 351.8 bps.

To see exactly how the z -spread is computed, let's look at a more stylized example, with a round-number time to maturity and pay frequency, and no accrued interest.

Example 7.2 (Computing the z -Spread) We compute the z -spread for a five-year bullet bond with semiannual fixed-rate coupon payments of 7 percent per annum, and trading at a dollar price of 95.00. To compute the z -spread, we need a swap zero-coupon curve, and to keep things simple, we assume the swap curve is flat at 3.5 percent per annum. The spot rate is then equal to a constant 3.470 percent for all maturities, as we saw in Section 4.2.

The yield to maturity of this bond is 8.075 percent, so the i -spread to swaps is $8.075 - 3.50 = 4.575$ percent. The z -spread is the constant z that satisfies

$$0.95 = \frac{0.07}{2} \sum_{i=1}^{5.2} e^{-(0.03470+z)\frac{i}{2}} + e^{-(0.03470+z)5}$$

This equation can be solved numerically to obtain $z = 460.5$ bps.

7.1.1 Spread Mark-to-Market

In Chapter 4.2, we studied the concept of DV01, the mark-to-market gain on a bond for a one basis point change in interest rates. There is an analogous concept for credit spreads, the "spread01," sometimes called DVCS, which measures the change in the value of a credit-risky bond for a one basis point change in spread.

For a credit-risky bond, we can measure the change in market value corresponding to a one basis point change in the z -spread. We can compute the spread01 the same way as the DV01: Increase and decrease the z -spread

by 0.5 basis points, reprice the bond for each of these shocks, and compute the difference.

Example 7.3 (Computing the Spread01) Continuing the earlier example, we start by finding the bond values for a 0.5-bps move up and down in the z -spread. The bond prices are expressed per \$100 of par value:

$$\frac{0.07}{2} \sum_{i=1}^{5.2} e^{-(0.03470+0.04605-0.00005)\frac{1}{2}} + e^{-(0.03470+0.04605-0.00005)5} = 0.950203$$

$$\frac{0.07}{2} \sum_{i=1}^{5.2} e^{-(0.03470+0.04605+0.00005)\frac{1}{2}} + e^{-(0.03470+0.04605+0.00005)5} = 0.949797$$

The difference is $95.0203 - 94.9797 = 0.040682$ dollars per basis point per \$100 of par value. This would typically be expressed as \$406.82 per \$1,000,000 of par value. The procedure is illustrated in Figure 7.1.

The spread01 of a fixed-rate bond depends on the initial level of the spread, which in turn is determined by the level and shape of the swap curve, the coupon, and other design features of the bond. The “typical” spread01 for a five-year bond (or CDS) is about \$400 per \$1,000,000 of

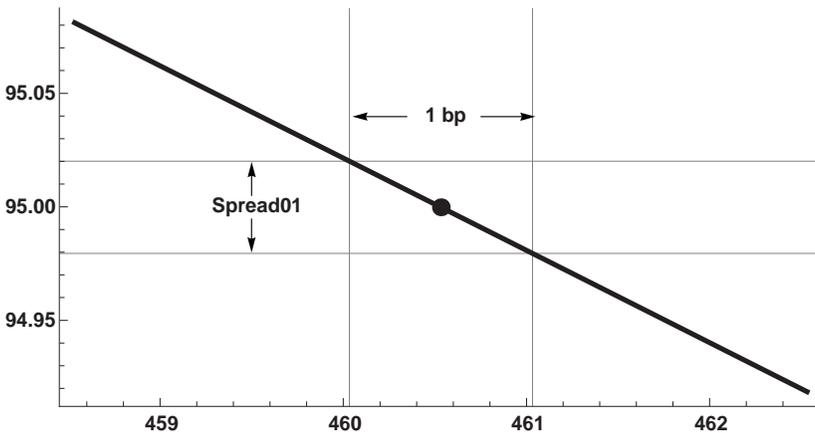


FIGURE 7.1 Computing Spread01 for a Fixed-Rate Bond
 The graph shows how spread01 is computed in Example 7.3 by shocking the z -spread up and down by 0.5 bps. The plot displays the value of the bond for a range of z -spreads. The point represents the initial bond price and corresponding z -spread. The vertical grid lines represent the 1 bps spread shock. The horizontal distance between the points on the plot where the vertical grid lines cross is equal to the spread01 per \$100 par value.

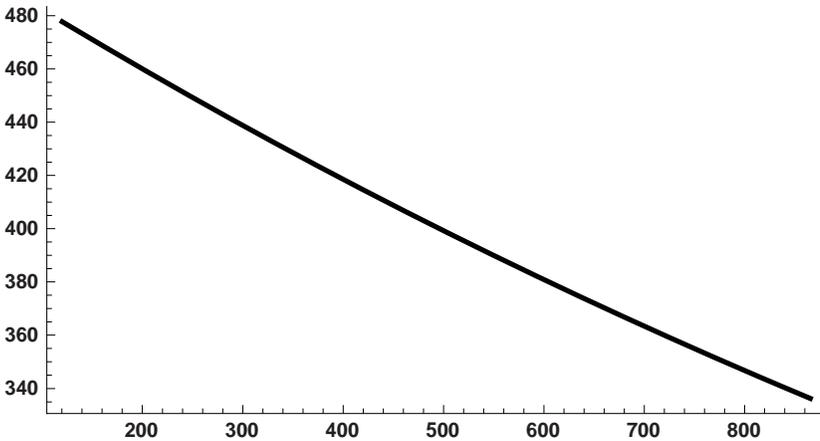


FIGURE 7.2 Spread01 a Declining Function of Spread Level

The graph shows how spread01, measured in dollars per \$1,000,000 of bond par value, varies with the spread level. The bond is a five-year bond making semiannual fixed-rate payments at an annual rate of 7 percent. The graph is constructed by permitting the price of the bond to vary between 80 and 110 dollars, and computing the z -spread and spread01 at each bond price, holding the swap curve at a constant flat 3.5 percent annual rate.

bond par value (or notional underlying amount). At very low or high spread levels, however, as seen in Figure 7.2, the spread01 can fall well above or below \$400.

The intuition is that, as the spread increases and the bond price decreases, the discount factor applied to cash flows that are further in the future declines. The spread-price relationship exhibits convexity; any increase or decrease in spread has a smaller impact on the bond's value when spreads are higher and discount factor is lower. The extent to which the impact of a spread *change* is attenuated by the high *level* of the spread depends primarily on the bond maturity and the level and shape of the swap or risk-free curve.

Just as there is a duration measure for interest rates that gives the proportional impact of a change in rates on bond value, the *spread duration* gives the proportional impact of a spread change on the price of a credit-risky bond. Like duration, spread duration is defined as the ratio of the spread01 to the bond price.

7.2 DEFAULT CURVE ANALYTICS

Reduced-form or *intensity models* of credit risk focus on the analytics of default timing. These models are generally focused on practical applications

such as pricing derivatives using arbitrage arguments and the prices of other securities, and lead to simulation-friendly pricing and risk measurement techniques. In this section, we lay out the basics of these analytics. Reduced form models typically operate in default mode, disregarding ratings migration and the possibility of restructuring the firm's balance sheet.

Reduced-form models, like the single-factor model of credit risk, rely on estimates of default probability that come from "somewhere else." The default probabilities can be derived from internal or rating agency ratings, or from structural credit models. But reduced form models are most often based on market prices or spreads. These *risk-neutral estimates of default probabilities* can be extracted from the prices of credit-risky bonds or loans, or from credit derivatives such as credit default swaps (CDS). In the next section, we show how to use the default curve analytics to extract default probabilities from credit spread data. In Chapter 8, we use the resulting default probability estimates as an input to models of credit portfolio risk.

Default risk for a single company can be represented as a *Bernoulli trial*. Over some fixed time horizon $\tau = T_2 - T_1$, there are just two outcomes for the firm: Default occurs with probability π , and the firm remains solvent with probability $1 - \pi$. If we assign the values 1 and 0 to the default and solvency outcomes over the time interval $(T_1, T_2]$, we define a random variable that follows a *Bernoulli distribution*. The time interval $(T_1, T_2]$ is important: The Bernoulli trial doesn't ask "does the firm ever default?," but rather, "does the firm default over the next year?"

The mean and variance of a Bernoulli-distributed variate are easy to compute. The expected value of default on $(T_1, T_2]$ is equal to the default probability π , and the variance of default is $\pi(1 - \pi)$.

The Bernoulli trial can be repeated during successive time intervals $(T_2, T_3]$, $(T_3, T_4]$, \dots . We can set each time interval to have the same length τ , and stipulate that the probability of default occurring during each of these time intervals is a constant value π . If the firm defaults during any of these time intervals, it remains defaulted forever, and the sequence of trials comes to an end. But so long as the firm remains solvent, we can imagine the firm surviving "indefinitely," but not "forever."

This model implies that the Bernoulli trials are *conditionally independent*, that is, that the event of default over each future interval $(T_j, T_{j+1}]$ is independent of the event of default over any earlier ($j > i$) interval $(T_i, T_{i+1}]$. This notion of independence is a potential source of confusion. It means that, from the current perspective, if you are told that the firm will survive up to time T_j , but have no idea when thereafter the firm will default, you "restart the clock" from the perspective of time T_j . You have no more or less information bearing on the survival of the firm over $(T_j, T_j + \tau]$ than you did at an earlier time T_i about survival over $(T_i, T_i + \tau]$. This property is also

called *memorylessness*, and is similar to the martingale property we noted for geometrical Brownian motion in Chapter 2.

In this model, the probability of default over some longer interval can be computed from the binomial distribution. For example, if τ is set equal to one year, the probability of survival over the next decade is equal $(1 - \pi)^{10}$, the probability of getting a sequence of 10 zeros in 10 independent Bernoulli trials.

It is inconvenient, though, to use a discrete distribution such as the binomial to model default over time, since the computation of probabilities can get tedious. An alternative is to model the random time at which a default occurs as the first arrival time—the time at which the modeled event occurs—of a *Poisson process*. In a Poisson process, the number of events in any time interval is *Poisson-distributed*. The time to the next arrival of a Poisson-distributed event is described by the *exponential distribution*. So our approach is equivalent to modeling the time to default as an exponentially distributed random variate. This leads to the a simple algebra describing default-time distributions, illustrated in Figure 7.3.

In describing the algebra of default time distributions, we set $t = 0$ as “now,” the point in time from which we are considering different time horizons.

7.2.1 The Hazard Rate

The *hazard rate*, also called the *default intensity*, denoted λ , is the parameter driving default. It has a time dimension, which we will assume is annual.¹ For each future time, the probability of a default over the tiny time interval dt is then

$$\lambda dt$$

and the probability that no default occurs over the time interval dt is

$$1 - \lambda dt$$

In this section, we assume that the hazard rate is a constant, in order to focus on defining default concepts. In the next section, where we explore how to derive risk-neutral default probabilities from market data, we’ll relax this assumption and let the hazard rate vary for different time horizons.

¹In life insurance, the equivalent concept applied to the likelihood of death rather than default is called the *force of mortality*.

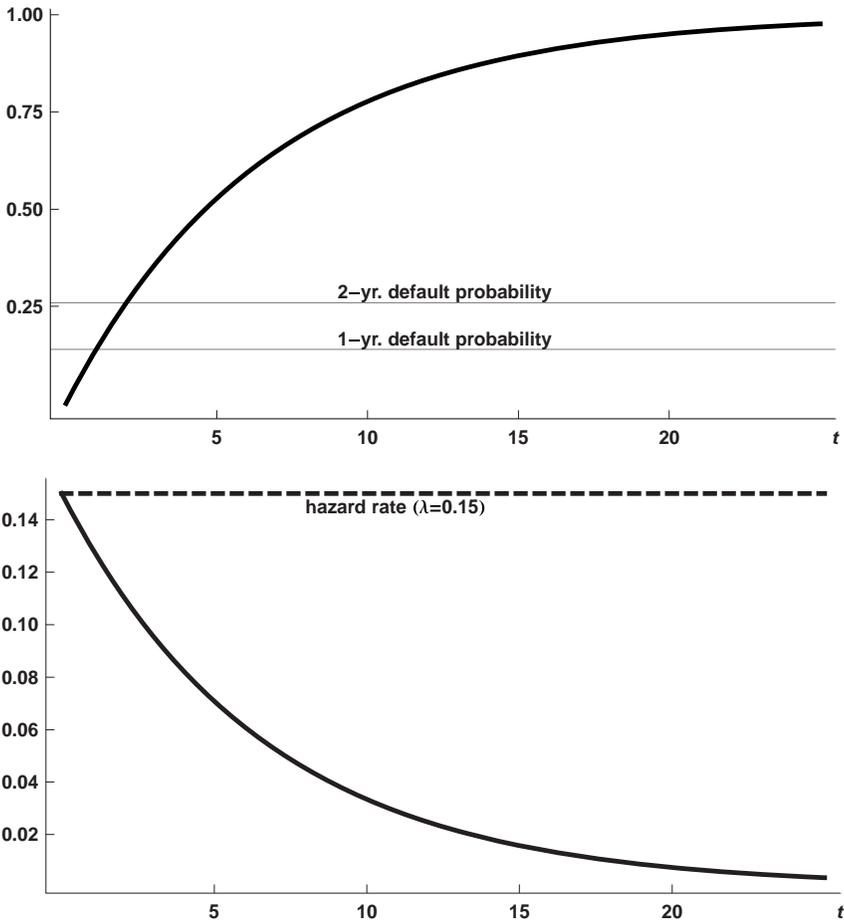


FIGURE 7.3 Intensity Model of Default Timing

The graphs are plotted from the perspective of time 0 and assume a value $\lambda = 0.15$, as in Example 7.4.

Upper panel: Cumulative default time distribution $1 - e^{-\lambda t}$. The ordinate of each point on the plot represents the probability of a default between the time 0 and the time t represented by the abscissa.

Lower panel: Hazard rate λ and marginal default probability $\lambda e^{-\lambda t}$. The ordinate of each point on the plot represents the annual rate at which the probability of a default between time 0 and the time t is changing. The marginal default probability is decreasing, indicating that the one-year probability of default is falling over time.

7.2.2 Default Time Distribution Function

The *default time distribution function* or *cumulative default time distribution* $F(t)$ is the probability of default sometime between now and time t :

$$\mathbf{P}[t^* < t] \equiv F(t) = 1 - e^{-\lambda t}$$

The survival and default probabilities must sum to exactly 1 at every instant t , so the probability of *no* default sometime between now and time t , called the *survival time distribution*, is

$$\mathbf{P}[t^* \geq t] = 1 - \mathbf{P}[t^* < t] = 1 - F(t) = e^{-\lambda t}$$

The survival probability converges to 0 and the default probability converges to 1 as t grows very large: in the intensity model, even a “bullet-proof” AAA-rated company will default eventually. This remains true even when we let the hazard rate vary over time.

7.2.3 Default Time Density Function

The *default time density function* or *marginal default probability* is the derivative of the default time distribution w.r.t. t :

$$\frac{\partial}{\partial t} \mathbf{P}[t^* < t] = F'(t) = \lambda e^{-\lambda t}$$

This is always a positive number, since default risk “accumulates”; that is, the probability of default increases for longer horizons. If λ is small, it will increase at a very slow pace. The survival probability, in contrast, is declining over time:

$$\frac{\partial}{\partial t} \mathbf{P}[t^* \geq t] = -F'(t) = -\lambda e^{-\lambda t} < 0$$

With a *constant* hazard rate, the marginal default probability is positive but declining, as seen in the lower panel of Figure 7.3. This means that, although the firm is likelier to default the further out in time we look, the rate at which default probability accumulates is declining. This is not necessarily true when the hazard rate can change over time. The default time density is still always positive, but if the hazard rate is rising fast enough with the time horizon, the cumulative default probability may increase at an increasing rather than at a decreasing rate.

7.2.4 Conditional Default Probability

So far, we have computed the probability of default over some time horizon $(0, t)$. If instead we ask, what is the probability of default over some horizon $(t, t + \tau)$ given that there has been no default prior to time t , we are asking about a *conditional default probability*. By the definition of conditional probability, it can be expressed as

$$\mathbf{P}[t^* < t + \tau | t^* > t] = \frac{\mathbf{P}[t^* > t \cap t^* < t + \tau]}{\mathbf{P}[t^* > t]}$$

that is, as the ratio of the probability of the joint event of survival up to time t and default over some horizon $(t, t + \tau)$, to the probability of survival up to time t .

That joint event of survival up to time t and default over $(t, t + \tau)$ is simply the event of defaulting during the discrete interval between two future dates t and $t + \tau$. In the constant hazard rate model, the probability $\mathbf{P}[t^* > t \cap t^* < t + \tau]$ of surviving to time t and then defaulting between t and $t + \tau$ is

$$\begin{aligned} \mathbf{P}[t^* > t \cap t^* < t + \tau] &= F(t + \tau) - F(t) \\ &= 1 - e^{-\lambda(t+\tau)} - (1 - e^{-\lambda t}) \\ &= e^{-\lambda t} (1 - e^{-\lambda \tau}) \\ &= [1 - F(t)]F(\tau) \\ &= \mathbf{P}[t^* > t] \mathbf{P}[t^* < t + \tau | t^* > t] \end{aligned}$$

We also see that

$$F(\tau) = \mathbf{P}[t^* < t + \tau | t^* > t]$$

which is equal to the unconditional τ -year default probability. We can interpret it as the probability of default over τ years, if we started the clock at zero at time t . This useful result is a further consequence of the memorylessness of the default process.

If the hazard rate is constant over a very short interval $(t, t + \tau)$, then the probability the security will default over the interval, given that it has

not yet defaulted up until time t , is

$$\lim_{\substack{\tau \rightarrow 0 \\ \tau > 0}} \mathbf{P}[t < t^* < t + \tau | t^* > t] = \frac{F'(t)\tau}{1 - F(t)} = \lambda\tau$$

The hazard rate can therefore now be interpreted as the instantaneous conditional default probability.

Example 7.4 (Hazard Rate and Default Probability) Suppose $\lambda = 0.15$. The unconditional one-year default probability is $1 - e^{-\lambda} = 0.1393$, and the survival probability is $e^{-\lambda} = 0.8607$. This would correspond to a low speculative-grade credit.

The unconditional two-year default probability is $1 - e^{-2\lambda} = 0.2592$. In the upper panel of Figure 7.3, horizontal grid lines mark the one- and two-year default probabilities. The difference between the two- and one-year default probabilities—the probability of the joint event of survival through the first year and default in the second—is 0.11989. The conditional one-year default probability, given survival through the first year, is the difference between the two probabilities (0.11989), divided by the one-year survival probability 0.8607:

$$\frac{0.11989}{0.8607} = 0.1393$$

which is equal, in this constant hazard rate example, to the unconditional one-year default probability.

7.3 RISK-NEUTRAL ESTIMATES OF DEFAULT PROBABILITIES

Our goal in this section is to see how default probabilities can be extracted from market prices with the help of the default algebra laid out in the previous section. As noted, these probabilities are risk neutral, that is, they include compensation for both the loss given default and bearing the risk of default and its associated uncertainties. The default intensity model gives us a handy way of representing spreads. We denote the spread over the risk-free rate on a defaultable bond with a maturity of T by z_T . The constant risk-neutral hazard rate at time T is λ_T^* . If we line up the defaultable securities by maturity, we can define a *spread curve*, that is, a function that relates the credit spread to the maturity of the bond.

7.3.1 Basic Analytics of Risk-Neutral Default Rates

There are two main types of securities that lend themselves to estimating default probabilities, bonds and credit default swaps (CDS). We start by describing the estimation process using the simplest possible security, a credit-risky zero-coupon corporate bond.

Let's first summarize the notation of this section:

p_τ	Current price of a default-free τ -year zero-coupon bond
p_τ^{corp}	Current price of a defaultable τ -year zero-coupon bond,
r_τ	Continuously compounded discount rate on the default free bond
z_τ	Continuously compounded spread on the defaultable bond
R	Recovery rate
λ_τ^*	τ -year risk neutral hazard rate
$1 - e^{-\lambda_\tau^*}$	Annualized risk neutral default probability

We assume that there are both defaultable and default-free zero-coupon bonds with the same maturity dates. The issuer's credit risk is then expressed by the discount or price concession at which it has to issue bonds, compared to the that on government bonds, rather than the coupon it has to pay to get the bonds sold. We'll assume there is only one issue of defaultable bonds, so that we don't have to pay attention to seniority, that is, the place of the bonds in the capital structure.

We'll denote the price of the defaultable discount bond maturing in τ years by p_τ^{corp} , measured as a decimal. The default-free bond is denoted p_τ . The continuously compounded discount rate on the default-free bond is the spot rate r_τ of Chapter 4, defined by

$$p_\tau = e^{-r_\tau \tau}$$

A corporate bond bears default risk, so it must be cheaper than a risk-free bond with the same future cash flows on the same dates, in this case \$1 per bond in τ years:

$$p_\tau \geq p_\tau^{\text{corp}}$$

The continuously compounded τ -year spread on a zero coupon corporate is defined as the difference between the rates on the corporate and default-free bonds and satisfies:

$$p_\tau^{\text{corp}} = e^{-(r_\tau + z_\tau)\tau} = p_\tau e^{-z_\tau \tau}$$

Since $p_\tau^{\text{corp}} \leq p_\tau$, we have $z_\tau \geq 0$.

The credit spread has the same time dimensions as the spot rate r_τ . It is the constant exponential rate at which, if there is no default, the price difference between a risky and risk-free bond shrinks to zero over the next τ years.

To compute hazard rates, we need to make some assumptions about default and recovery:

- The issuer can default any time over the next τ years.
- In the event of default, the creditors will receive a deterministic and known recovery payment, but only at the maturity date, regardless of when default occurs. Recovery is a known fraction R of the par amount of the bond (recovery of face).

We'll put all of this together to estimate λ_τ^* , the risk-neutral constant hazard rate over the next τ years. The risk-neutral τ -year default probability is thus $1 - e^{-\lambda_\tau^* \tau}$. Later on, we will introduce the possibility of a time-varying hazard rate and learn how to estimate a term structure from bond or CDS data in which the spreads and default probabilities may vary with the time horizon. The time dimensions of λ_τ^* are the same as those of the spot rate and the spread. It is the conditional default probability over $(0, T)$, that is, the constant annualized probability that the firm defaults over a tiny time interval $t + \Delta t$, given that it has not already defaulted by time t , with $0 < t < T$.

The risk-neutral (and physical) hazard rates have an exponential form. The probability of defaulting over the next instant is a constant, and the probability of defaulting over a discrete time interval is an exponential function of the length of the time interval.

For the moment, let's simplify the setup even more, and let the recovery rate $R = 0$. An investor in a defaultable bond receives either \$1 or zero in τ years. The expected value of the two payoffs is

$$e^{-\lambda_\tau^* \tau} \cdot 1 + (1 - e^{-\lambda_\tau^* \tau}) \cdot 0$$

The expected present value of the two payoffs is

$$e^{-r_\tau \tau} [e^{-\lambda_\tau^* \tau} \cdot 1 + (1 - e^{-\lambda_\tau^* \tau}) \cdot 0]$$

Discounting at the risk-free rate is appropriate because we want to estimate λ_τ^* , the risk-neutral hazard rate. To the extent that the credit-risky bond price and z_τ reflect a risk premium as well as an estimate of the true default probability, the risk premium will be embedded in λ_τ^* , so we don't have to discount by a risky rate.

The risk-neutral hazard rate sets the expected present value of the two payoffs equal to the price of the defaultable bond. In other words, if market prices have adjusted to eliminate the potential for arbitrage, we can solve Equation (7.1) for λ_τ^* :

$$e^{-(r_\tau+z_\tau)\tau} = e^{-r_\tau\tau} [e^{-\lambda_\tau^*\tau} \cdot 1 + (1 - e^{-\lambda_\tau^*\tau}) \cdot 0] \quad (7.1)$$

to get our first simple rule of thumb: If recovery is zero, then

$$\lambda_\tau^* = z_\tau$$

that is, the hazard rate is equal to the spread. Since for small values of x we can use the approximation $e^x \approx 1 + x$, we also can say that the spread $z_\tau \approx 1 - e^{-\lambda_\tau^*}$, the default probability.

Example 7.5 Suppose a company's securities have a five-year spread of 300 bps over the Libor curve. Then the risk-neutral annual hazard rate over the next five years is 3 percent, and the annualized default probability is approximately 3 percent. The exact annualized default probability is 2.96 percent, and the five-year default probability is 13.9 percent.

Now let the recovery rate R be a positive number on $(0, 1)$. The owner of the bond will receive one of two payments at the maturity date. Either the issuer does not default, and the creditor receives par (\$1), or there is a default, and the creditor receives R . Setting the expected present value of these payments equal to the bond price, we have

$$e^{-(r_\tau+z_\tau)\tau} = e^{-r_\tau\tau} [e^{-\lambda_\tau^*\tau} + (1 - e^{-\lambda_\tau^*\tau})R]$$

or

$$e^{-z_\tau\tau} = e^{-\lambda_\tau^*\tau} + (1 - e^{-\lambda_\tau^*\tau})R = 1 - (1 - e^{-\lambda_\tau^*\tau})(1 - R)$$

giving us our next rule of thumb: The additional credit-risk discount on the defaultable bond, divided by the LGD, is equal to the τ -year default probability:

$$1 - e^{-\lambda_\tau^*\tau} = \frac{1 - e^{-z_\tau\tau}}{1 - R}$$

We can get one more simple rule of thumb by taking logs in Equation (7.1):

$$-(r_\tau + z_\tau)\tau = -r_\tau\tau + \log[e^{-\lambda_\tau^*\tau} + (1 - e^{-\lambda_\tau^*\tau})R]$$

or

$$z_\tau \tau = -\log[e^{-\lambda_\tau^* \tau} + (1 - e^{-\lambda_\tau^* \tau})R]$$

This expression can be solved numerically for λ_τ^* , or we can use the approximations $e^x \approx 1 + x$ and $\log(1 + x) \approx x$, so $e^{-\lambda_\tau^* \tau} + (1 - e^{-\lambda_\tau^* \tau})R \approx 1 - \lambda_\tau^* \tau + \lambda_\tau^* \tau R = 1 - \lambda_\tau^* \tau(1 - R)$. Therefore,

$$\log[1 - \lambda_\tau^* \tau(1 - R)] \approx -\lambda_\tau^* \tau(1 - R)$$

Putting these results together, we have

$$z_\tau \tau \approx \lambda_\tau^* \tau(1 - R) \Rightarrow \lambda_\tau^* \approx \frac{z_\tau}{1 - R}$$

The spread is approximately equal to the default probability times the LGD. The approximation works well when spreads or risk-neutral default probabilities are not too large.

Example 7.6 Continuing the example of a company with a five-year spread of 300 bps, with a recovery rate $R = 0.40$, we have a hazard rate of

$$\lambda_\tau^* \approx \frac{0.0300}{1 - 0.4} = 0.05$$

or 5 percent.

So far, we have defined spot hazard rates, which are implied by prices of risky and riskless bonds over different time intervals. But just as we can define spot and forward risk-free rates, we can define spot and forward hazard rates. A forward hazard rate from time T_1 to T_2 is the constant hazard rate over that interval. If $T_1 = 0$, it is identical to the spot hazard rate over $(0, T_2)$.

7.3.2 Time Scaling of Default Probabilities

We typically don't start our analysis with an estimate of the hazard rate. Rather, we start with an estimate of the probability of default π over a given time horizon, based on either the probability of default provided by a rating agency—the rightmost column of the transition matrix illustrated in Table 6.2—or a model, or on a market credit spread.

These estimates of π have a specific time horizon. The default probabilities provided by rating agencies for corporate debt typically have a horizon

of one year. Default probabilities based on credit spreads have a time horizon equal to the time to maturity of the security from which they are derived. The time horizon of the estimated default probability may not match the time horizon we are interested in. For example, we may have a default probability based on a one-year transition matrix, but need a five-year default probability in the context of a longer-term risk analysis.

We can always convert a default probability from one time horizon to another by applying the algebra of hazard rates. But we can also use a default probability with one time horizon directly to estimate default probabilities with longer or shorter time horizons. Suppose, for example, we have an estimate of the one-year default probability π_1 . From the definition of a constant hazard rate,

$$\pi_1 = 1 - e^{-\lambda}$$

we have

$$\lambda = \log(1 - \pi_1)$$

This gives us an identity

$$\pi_1 = 1 - e^{-\log(1-\pi_1)}$$

We can then approximate

$$\pi_t = 1 - (1 - \pi_1)^t$$

7.3.3 Credit Default Swaps

So far, we have derived one constant hazard rate using the prices of default-free and defaultable discount bonds. This is a good way to introduce the analytics of risk-neutral hazard rates, but a bit unrealistic, because corporations do not issue many zero-coupon bonds. Most corporate zero-coupon issues are commercial paper, which have a typical maturity under one year, and are issued by only a small number of highly rated “blue chip” companies. Commercial paper even has a distinct rating system.

In practice, hazard rates are usually estimated from the prices of CDS. These have a few advantages:

Standardization. In contrast to most developed-country central governments, private companies do not issue bonds with the same cash flow structure and the same seniority in the firm’s capital structure

at fixed calendar intervals. For many companies, however, CDS trading occurs regularly in standardized maturities of 1, 3, 5, 7, and 10 years, with the five-year point generally the most liquid.

Coverage. The universe of firms on which CDS are issued is large. Markit Partners, the largest collector and purveyor of CDS data, provides curves on about 2,000 corporate issuers globally, of which about 800 are domiciled in the United States.

Liquidity. When CDS on a company’s bonds exist, they generally trade more heavily and with a tighter bid-offer spread than bond issues. The liquidity of CDS with different maturities usually differs less than that of bonds of a given issuer.

Figure 7.4 displays a few examples of CDS credit curves.

Hazard rates are typically obtained from CDS curves via a bootstrapping procedure. We’ll see how it works using a detailed example. We first need more detail on how CDS contracts work. We also need to extend our discussion of the default probability function to include the possibility of time-varying hazard rates. CDS contracts with different terms to maturity can have quite different prices or spreads.

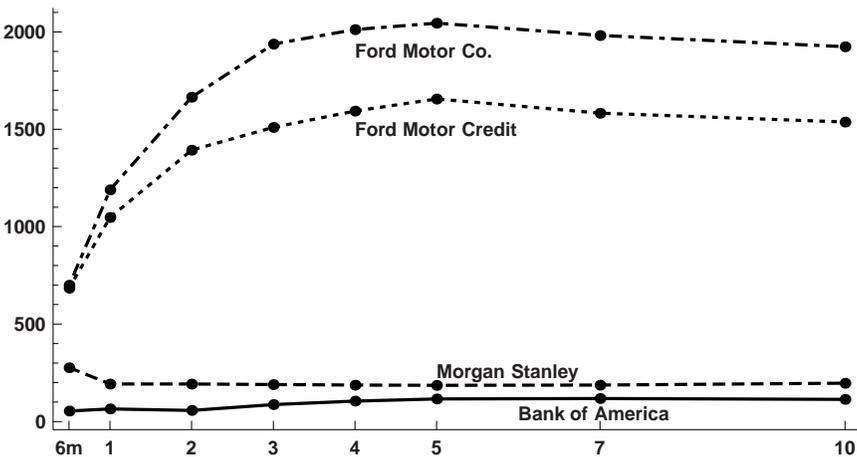


FIGURE 7.4 CDS Curves

CDS on senior unsecured debt as a function of tenor, expressed as an annualized CDS premium in basis points, July 1, 2008.

Source: Bloomberg Financial L.P.

To start, recall that in our simplified example above, the hazard rate was found by solving for the default probability that set the expected present value of the credit spread payments equal to the expected present value of the default loss. Similarly, to find the default probability function using CDS, we set the expected present value of the spread payments by the protection buyer equal to the expected present value of the protection seller's payments in the event of default.

The CDS contract is written on a specific *reference entity*, typically a firm or a government. The contract defines an event of default for the reference entity. In the event of default, the contract obliges the protection seller to pay the protection buyer the par amount of a deliverable bond of the reference entity; the protection buyer delivers the bond. The CDS contract specifies which of the reference entity's bonds are "deliverable," that is, are covered by the CDS.

In our discussion, we will focus on single-name corporate CDS, which create exposure to bankruptcy events of a single issuer of bonds such as a company or a sovereign entity. Most, but not all, of what we will say about how CDS work also applies to other types, such as CDS on credit indexes.

CDS are traded in spread terms. That is, when two traders make a deal, the price is expressed in terms of the spread premium the counterparty buying protection is to pay to the counterparty selling protection. CDS may trade "on spread" or "on basis." When the spread premium would otherwise be high, CDS trade points upfront, that is, the protection buyer pays the seller a market-determined percentage of the notional at the time of the trade, and the spread premium is set to 100 or 500 bps, called "100 running." Prior to the so-called "Big Bang" reform of CDS trading conventions that took place on March 13, 2009, only CDS on issuers with wide spreads traded "points up." The running spread was, in those cases, typically 500 bps. The reformed convention has all CDS trading points up, but with some paying 100 and others 500 bps running.

A CDS is a swap, and as such

- Generally, no principal or other cash flows change hands at the initiation of the contract. However, when CDS trade points upfront, a percent of the principal is paid by the protection buyer. This has an impact primarily on the counterparty credit risk of the contract rather than on its pricing, since there is always a spread premium, with no points up front paid, that is equivalent economically to any given market-adjusted number of points upfront plus a running spread.

There are generally exchanges of collateral when a CDS contract is created.

- Under the terms of a CDS, there are agreed future cash flows. The protection buyer undertakes to make spread payments, called the *fee leg*, each quarter until the maturity date of the contract, unless and until there is a default event pertaining to the underlying name on which the CDS is written. The protection seller makes a payment, called the *contingent leg*, only if there is a default. It is equal to the estimated loss given default, that is, the notional less the recovery on the underlying bond.²
- The pricing of the CDS, that is, the market-adjusted spread premium, is set so that the expected net present value of the CDS contract is zero. In other words, on the initiation date, the expected present value of the fee leg is equal to that of the contingent leg. If market prices change, the net present value becomes positive for one counterparty and negative for the other; that is, there is a mark-to-market gain and loss.

The CDS contract specifies whether the contract protects the senior or the subordinated debt of the underlying name. For companies that have issued both senior and subordinated debt, there may be CDS contracts of both kinds.

Often, risk-neutral hazard rates are calculated using the conventional assumption about recovery rates that $R = 0.40$. An estimate based on fundamental credit analysis of the specific firm can also be used. In some cases, a risk-neutral estimate is available based on the price of a *recovery swap* on the credit. A recovery swap is a contract in which, in the event of a default, one counterparty will pay the actual recovery as determined by the settlement procedure on the corresponding CDS, while the other counterparty will pay a fixed amount determined at initiation of the contract. Subject to counterparty risk, the counterparty promising that fixed amount is thus able to substitute a fixed recovery rate for an uncertain one. When those recovery swap prices can be observed, the fixed rate can be used as a risk-neutral recovery rate in building default probability distributions.

²There is a procedure for cash settlement of the protection seller's contingent obligations, standard since April 2009 as part of the "Big Bang," in which the recovery amount is determined by an auction mechanism. The seller may instead pay the buyer the notional underlying amount, while the buyer delivers a bond, from a list of acceptable or "deliverable" bonds issued by the underlying name rather than make a cash payment. In that case, it is up to the seller to gain the recovery value either through the bankruptcy process or in the marketplace.

7.3.4 Building Default Probability Curves

Next, let's extend our earlier analysis of hazard rates and default probability distributions to accommodate hazard rates that vary over time. We will add a time argument to our notation to indicate the time horizon to which it pertains. The conditional default probability at time t , the probability that the company will default over the next instant, given that it has survived up until time t , is denoted $\lambda(t)$, $t \in [0, \infty)$.

The default time distribution function is now expressed in terms of an integral in hazard rates. The probability of default over the interval $[0, t)$ is

$$\pi_t = 1 - e^{\int_0^t \lambda(s) ds} \quad (7.2)$$

If the hazard rate is constant, $\lambda(t) = \lambda$, $t \in [0, \infty)$, then Equation (7.2) reduces to our earlier expression $\pi_t = 1 - e^{\lambda t}$. In practice, we will be estimating and using hazard rates that are not constant, but also don't vary each instant. Rather, since we generally have the standard CDS maturities of 1, 3, 5, 7, and 10 years available, we will extract 5 piecewise constant hazard rates from the data:

$$\lambda(t) = \begin{Bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \\ \lambda_4 \\ \lambda_5 \end{Bmatrix} \quad \text{for} \quad \begin{Bmatrix} 0 < t \leq 1 \\ 1 < t \leq 3 \\ 3 < t \leq 5 \\ 5 < t \leq 7 \\ 7 < t \end{Bmatrix}$$

The integral from which default probabilities are calculated via Equation (7.2) is then

$$\int_0^t \lambda(s) ds = \begin{Bmatrix} \lambda_1 t \\ \lambda_1 + (t-1)\lambda_2 \\ \lambda_1 + 2\lambda_2 + (t-3)\lambda_3 \\ \lambda_1 + 2\lambda_2 + 2\lambda_3 + (t-5)\lambda_4 \\ \lambda_1 + 2\lambda_2 + 2\lambda_3 + 3\lambda_4 + (t-7)\lambda_5 \end{Bmatrix} \quad \text{for} \quad \begin{Bmatrix} 0 < t \leq 1 \\ 1 < t \leq 3 \\ 3 < t \leq 5 \\ 5 < t \leq 7 \\ 7 < t \end{Bmatrix}$$

Now, let's look at the expected present value of each CDS leg. Denote by s_τ the spread premium on a τ -year CDS on a particular company. The protection buyer will pay the spread in quarterly installments if and only if the credit is still alive on the payment date. The probability of survival up

to date t is π_t , so we can express this expected present value, in dollars per dollar of underlying notional, as

$$\frac{1}{4 \times 10^3} s_\tau \sum_{u=1}^{4\tau} p_{0.25u} (1 - \pi_{0.25u})$$

where p_t is the price of a risk-free zero-coupon bond maturing at time t . We will use a discount curve based on interest-rate swaps. The summation index u takes on integer values, but since we are adding up the present values of quarterly cash flows, we divide u by 4 to get back to time measured in years.

There is one more wrinkle in the fee leg. In the event of default, the protection buyer must pay the portion of the spread premium that accrued between the time of the last quarterly payment and the default date. This payment isn't included in the summation above. The amount and timing is uncertain, but the convention is to approximate it as half the quarterly premium, payable on the first payment date following default. The implicit assumption is that the default, if it occurs at all, occurs midway through the quarter. The probability of having to make this payment on date t is equal to $\pi_{t-0.25} - \pi_t$, the probability of default during the interval $(t - \frac{1}{4}, t]$. This probability is equal to the probability of surviving to time $(t - \frac{1}{4}]$ minus the smaller probability of surviving to time t .

Taking this so-called fee accrual term into account, the expected present value of the fee leg becomes

$$\frac{1}{4 \times 10^4} s_\tau \sum_{u=1}^{4\tau} p_{0.25u} \left[(1 - \pi_{0.25u}) + \frac{1}{2} (\pi_{0.25(u-1)} - \pi_{0.25u}) \right]$$

Next, we calculate the expected present value of the contingent leg. If a default occurs during the quarter ending at time t , the present value of the contingent payment is $(1 - R)p_t$ per dollar of notional. We assume that the contingent payment is made on the quarterly cash flow date following the default. The expected present value of this payment is obtained by multiplying this present value by the probability of default during the quarter:

$$(1 - R)p_t (\pi_{t-0.25} - \pi_t)$$

The expected present value of the contingent leg is therefore equal to the sum of these expected present values over the life of the CDS contract:

$$(1 - R) \sum_{u=1}^{4\tau} p_{0.25u} (\pi_{0.25(u-1)} - \pi_{0.25u})$$

The fair market CDS spread is the number s_τ that equalizes these two payment streams, that is, solves

$$\begin{aligned} & \frac{1}{4 \times 10^4} \sum_{u=1}^{4\tau} p_{0.25u} \left[(1 - \pi_{0.25u}) + \frac{1}{2} (\pi_{0.25(u-1)} - \pi_{0.25u}) \right] \\ &= (1 - R) \sum_{u=1}^{4\tau} p_{0.25u} (\pi_{0.25(u-1)} - \pi_{0.25u}) \end{aligned} \quad (7.3)$$

Now we're ready to estimate the default probability distribution. To solve Equation (7.3), the market must "have in its mind" an estimate of the default curve, that is the π_t . Of course, it doesn't: The s_τ are found by supply and demand. But once we observe the spreads set by the market, we can infer the π_t by backing them out of Equation (7.3) via a bootstrapping procedure, which we now describe.

The data we require are swap curve interest data, so that we can estimate a swap discount curve, and a set of CDS spreads s_τ on the same name and with the same seniority, but with different terms to maturity. We learned in Chapter 4 how to generate a swap curve from observation on money-market and swap rates, so we will assume that we can substitute specific numbers for all the discount factors p_t .

Let's start by finding the default curve for a company for which we have only a single CDS spread, for a term, say, of five years. This will result in a single hazard rate estimate. We need default probabilities for the quarterly dates $t = 0.25, 0.50, \dots, 5$. They are a function of the as-yet unknown hazard rate λ : $\pi_t = e^{-\lambda t}$, $t > 0$. Substituting this, the five-year CDS spread, the recovery rate and the discount factors into the CDS valuation function (7.3) gives us

$$\begin{aligned} & \frac{1}{4 \times 10^4} s_\tau \sum_{u=1}^{4\tau} p_{0.25u} \left[e^{-\lambda \frac{u}{4}} + \frac{1}{2} \left(e^{-\lambda \frac{u-1}{4}} - e^{-\lambda \frac{u}{4}} \right) \right] \\ &= (1 - R) \sum_{u=1}^{4\tau} p_{0.25u} \left(e^{-\lambda \frac{u-1}{4}} - e^{-\lambda \frac{u}{4}} \right) \end{aligned}$$

with $\tau = 5$. This is an equation in one unknown variable that can be solved numerically for λ .

Example 7.7 We compute a constant hazard rate for Merrill Lynch as of October 1, 2008, using the closing five-year CDS spread of 445 bps. We assume a recovery rate $R = 0.40$. To simplify matters, we also assume a flat

swap curve, with a continuously compounded spot rate of 4.5 percent for all maturities, so the discount factor for a cash flow t years in the future is $e^{0.045t}$. As long as this constant swap rate is reasonably close to the actual swap rate prevailing on October 1, 2008, this has only a small effect on the numerical results.

With $\tau = 5$, $s_\tau = 445$, $R = 0.40$, we have

$$\begin{aligned} & \frac{445}{4 \times 10^4} \sum_{u=1}^{4.5} e^{0.045 \frac{u}{4}} \left[e^{-\lambda \frac{u}{4}} + \frac{1}{2} \left(e^{-\lambda \frac{u-1}{4}} - e^{-\lambda \frac{u}{4}} \right) \right] \\ &= 0.60 \sum_{u=1}^{4.5} e^{0.045 \frac{u}{4}} \left(e^{-\lambda \frac{u-1}{4}} - e^{-\lambda \frac{u}{4}} \right) \end{aligned}$$

This equation can be solved numerically to obtain $\lambda = 0.0741688$.

The bootstrapping procedure is a bit more complicated, since it involves a sequence of steps. But each step is similar to the calculation we just carried out for a single CDS spread and a single hazard rate. The best way to explain it is with an example.

Example 7.8 We will compute the default probability curve for Merrill Lynch as of October 1, 2008. The closing CDS spreads on that date for each CDS maturity were

i	τ_i (yrs)	s_{τ_i} (bps/yr)	λ_i
1	1	576	0.09600
2	3	490	0.07303
3	5	445	0.05915
4	7	395	0.03571
5	10	355	0.03416

The table above also displays the estimated forward hazard rates, the extraction of which we now describe in detail. We continue to assume a recovery rate $R = 0.40$ and a flat swap curve, with the discount function $p_t = e^{0.045t}$.

At each step i , we need quarterly default probabilities over the interval $(0, \tau_i]$, $i = 1, \dots, 5$, some or all of which will still be unknown when

we carry out that step. We progressively “fill in” the integral in Equation (7.2) as the bootstrapping process moves out the curve. In the first step, we find

$$\pi_t = 1 - e^{-\lambda_1 t} \quad t \in (0, \tau_1]$$

We start by solving for the first hazard rate λ_1 . We need the discount factors for the quarterly dates $t = \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1$, and the CDS spread with the shortest maturity, τ_1 . We solve this equation in one unknown for λ_1 :

$$\begin{aligned} & \frac{1}{4 \times 10^3} s_{\tau_1} \sum_{u=1}^{4\tau_1} p_{0.25u} \left[e^{-\lambda_1 \frac{u}{4}} + \frac{1}{2} \left(e^{-\lambda_1 \frac{u-1}{4}} - e^{-\lambda_1 \frac{u}{4}} \right) \right] \\ &= (1 - R) \sum_{u=1}^{4\tau_1} p_{0.25u} \left(e^{-\lambda_1 \frac{u-1}{4}} - e^{-\lambda_1 \frac{u}{4}} \right) \end{aligned}$$

With $\tau_1 = 1$, $s_{\tau_1} = 576$, and $R = 0.40$, this becomes

$$\begin{aligned} & \frac{576}{4 \times 10^3} \sum_{u=1}^4 e^{0.045 \frac{u}{4}} \left[e^{-\lambda_1 \frac{u}{4}} + \frac{1}{2} \left(e^{-\lambda_1 \frac{u-1}{4}} - e^{-\lambda_1 \frac{u}{4}} \right) \right] \\ &= 0.60 \sum_{u=1}^4 e^{0.045 \frac{u}{4}} \left(e^{-\lambda_1 \frac{u-1}{4}} - e^{-\lambda_1 \frac{u}{4}} \right) \end{aligned}$$

which we can solve numerically for λ_1 , obtaining $\lambda_1 = 0.0960046$. Once the default probabilities are substituted back in, the fee and the contingent legs of the swap are found to each have a fair value of \$0.0534231 per dollar of notional principal protection.

In the next step, we extract λ_2 from the data, again by setting up an equation that we can solve numerically for λ_2 . We now need quarterly default probabilities and discount factors over the interval $(0, \tau_2] = (0, 3]$. For any t in this interval,

$$\pi_t = e^{-\int_0^t \lambda(s) ds} = \begin{cases} e^{-\lambda_1 t} \\ e^{-[\lambda_1 + (t-1)\lambda_2]} \end{cases} \quad \text{for} \quad \begin{cases} 0 < t \leq 1 \\ 1 < t \leq 3 \end{cases}$$

The default probabilities for $t \leq \tau_1 = 1$ are known, since they use only λ_1 .

Substitute these probabilities, as well as the discount factors, recovery rate, and the three-year CDS spread into the expression for CDS fair value to get:

$$\begin{aligned} & \frac{1}{4 \times 10^3} s_{\tau_2} \sum_{u=1}^{4\tau_1} p_{0.25u} \left[e^{-\lambda_1 \frac{u}{4}} + \frac{1}{2} \left(e^{-\lambda_1 \frac{u-1}{4}} - e^{-\lambda_1 \frac{u}{4}} \right) \right] \\ & + \frac{1}{4 \times 10^3} s_{\tau_2} e^{-\lambda_1 \tau_1} \sum_{u=4\tau_1+1}^{4\tau_2} p_{0.25u} \left[e^{-\lambda_2 (\frac{u}{4} - \tau_1)} + \frac{1}{2} \left(e^{\lambda_2 [\frac{u-1}{4} - \tau_1]} - e^{\lambda_2 (\frac{u}{4} - \tau_1)} \right) \right] \\ & = (1 - R) \sum_{u=1}^{4\tau_1} p_{0.25u} \left(e^{-\lambda_1 \frac{u-1}{4}} - e^{-\lambda_1 \frac{u}{4}} \right) \\ & + (1 - R) e^{-\lambda_1 \tau_1} \sum_{u=4\tau_1+1}^{4\tau_2} p_{0.25u} \left(e^{-\lambda_2 [\frac{u-1}{4} - \tau_1]} - e^{-\lambda_2 (\frac{u}{4} - \tau_1)} \right) \end{aligned}$$

and solve numerically for λ_2 .

Notice that the first term on each side of the above equation is a known number at this point in the bootstrapping process, since the default probabilities for horizons of one year or less are known. Once we substitute the known quantities into the above equation, we have

$$\begin{aligned} & \frac{490}{4 \times 10^3} \sum_{u=1}^4 e^{0.045 \frac{u}{4}} \left[e^{-0.0960046 \frac{u}{4}} + \frac{1}{2} \left(e^{-0.0960046 \frac{u-1}{4}} - e^{-0.0960046 \frac{u}{4}} \right) \right] \\ & + \frac{490}{4 \times 10^3} e^{-0.0960046} \sum_{u=5}^{4.3} e^{0.045 \frac{u}{4}} \left\{ e^{-\lambda_2 (\frac{u}{4} - 1)} + \frac{1}{2} \left[e^{\lambda_2 (\frac{u-1}{4} - 1)} - e^{\lambda_2 (\frac{u}{4} - 1)} \right] \right\} \\ & = 0.04545 \\ & + \frac{490}{4 \times 10^3} e^{-0.0960046} \sum_{u=5}^{4.3} e^{0.045 \frac{u}{4}} \left\{ e^{-\lambda_2 (\frac{u}{4} - 1)} + \frac{1}{2} \left[e^{\lambda_2 (\frac{u-1}{4} - 1)} - e^{\lambda_2 (\frac{u}{4} - 1)} \right] \right\} \\ & = 0.60 \sum_{u=1}^4 e^{0.045 \frac{u}{4}} \left(e^{-0.0960046 \frac{u-1}{4}} - e^{-0.0960046 \frac{u}{4}} \right) \\ & + 0.60 \sum_{u=5}^{4.3} e^{0.045 \frac{u}{4}} \left[e^{\lambda_2 (\frac{u-1}{4} - 1)} - e^{\lambda_2 (\frac{u}{4} - 1)} \right] \\ & = 0.05342 + 0.60 \sum_{u=5}^{4.3} e^{0.045 \frac{u}{4}} \left[e^{\lambda_2 (\frac{u-1}{4} - 1)} - e^{\lambda_2 (\frac{u}{4} - 1)} \right] \end{aligned}$$

which can be solved numerically to obtain $\lambda_2 = 0.0730279$.

Let's spell out one more step explicitly and extract λ_3 from the data. The quarterly default probabilities and discount factors we now need cover the interval $(0, \tau_3) = (0, 5)$. For any t in this interval,

$$\pi_t = e^{-\int_0^t \lambda(s) ds} = \begin{cases} e^{-\lambda_1 t} \\ e^{-[\lambda_1 + (t-1)\lambda_2]} \\ e^{-[\lambda_1 + 2\lambda_2 + (t-3)\lambda_3]} \end{cases} \quad \text{for} \quad \begin{cases} 0 < t \leq 1 \\ 1 < t \leq 3 \\ 3 < t \leq 5 \end{cases}$$

The default probabilities for $t \leq \tau_2 = 3$ are known, since they are functions of λ_1 and λ_2 alone, which are known after the second step.

Now we use the five-year CDS spread in the expression for CDS fair value to set up:

$$\begin{aligned} & \frac{s_{\tau_3}}{4 \times 10^3} \sum_{u=1}^{4\tau_1} p_{0.25u} \left[e^{-\lambda_1 \frac{u}{4}} + \frac{1}{2} \left(e^{-\lambda_1 \frac{u-1}{4}} - e^{-\lambda_1 \frac{u}{4}} \right) \right] \\ & + \frac{s_{\tau_3}}{4 \times 10^3} e^{-\lambda_1 \tau_1} \sum_{u=4\tau_1+1}^{4\tau_2} p_{0.25u} \left[e^{-\lambda_2 \frac{u}{4}} + \frac{1}{2} \left(e^{-\lambda_2 \frac{u-1}{4}} - e^{-\lambda_2 \frac{u}{4}} \right) \right] \\ & + \frac{s_{\tau_3}}{4 \times 10^3} e^{-[\lambda_1 \tau_1 + \lambda_2 (\tau_2 - \tau_1)]} \sum_{u=4\tau_2+1}^{4\tau_3} p_{0.25u} \left[e^{-\lambda_3 \frac{u}{4}} + \frac{1}{2} \left(e^{\lambda_3 \frac{u-1}{4}} - e^{\lambda_3 \frac{u}{4}} \right) \right] \\ & = (1 - R) \sum_{u=1}^{4\tau_1} p_{0.25u} \left(e^{-\lambda_1 \frac{u-1}{4}} - e^{-\lambda_1 \frac{u}{4}} \right) \\ & + (1 - R) e^{-\lambda_1 \tau_1} \sum_{u=4\tau_1+1}^{4\tau_2} p_{0.25u} \left(e^{-\lambda_2 \frac{u-1}{4}} - e^{-\lambda_2 \frac{u}{4}} \right) \\ & + (1 - R) e^{-[\lambda_1 \tau_1 + \lambda_2 (\tau_2 - \tau_1)]} \sum_{u=4\tau_2+1}^{4\tau_3} p_{0.25u} \left(e^{\lambda_3 \frac{u-1}{4}} - e^{\lambda_3 \frac{u}{4}} \right) \end{aligned}$$

Once again, at this point in the bootstrapping process, since the default probabilities for horizons of three years or less are known, the first two

terms on each side of the equals sign are known quantities. And once we have substituted them, we have

$$\begin{aligned}
 &0.10974 \\
 &+ \frac{445}{4 \times 10^3} e^{-[\lambda_1+2\lambda_2]} \sum_{u=4\tau_2+1}^{4\tau_3} p_{0.25u} \left\{ e^{-\lambda_3(\frac{u}{4}-1)} + \frac{1}{2} \left[e^{\lambda_3(\frac{u-1}{4}-1)} - e^{\lambda_3(\frac{u}{4}-1)} \right] \right\} \\
 &= 0.12083 + 0.60 \sum_{u=4\tau_2+1}^{4\tau_3} p_{0.25u} \left[e^{\lambda_3(\frac{u-1}{4}-1)} - e^{\lambda_3(\frac{u}{4}-1)} \right]
 \end{aligned}$$

which can be solved numerically to obtain $\lambda_3 = 0.05915$.

The induction process should now be clear. It is illustrated in Figure 7.4. With our run of five CDS maturities, we repeat the process twice more. The intermediate results are tabulated by step in the table below. Each row in the table displays the present expected value of either leg of the CDS after finding the contemporaneous hazard rate, and the values of the fee and contingent legs up until that step. Note that last period’s value of either leg becomes the next period’s value of contingent leg payments in previous periods:

i	Either leg $\rightarrow \tau_i$	Fee leg $\rightarrow \tau_{i-1}$	Contingent leg $\rightarrow \tau_{i-1}$
1	0.05342	0.00000	0.00000
2	0.12083	0.04545	0.05342
3	0.16453	0.10974	0.12083
4	0.18645	0.14605	0.16453
5	0.21224	0.16757	0.18645

The CDS in our example did not trade points up, in contrast to the standard convention since 2009. However, Equation (7.3) also provides an easy conversion between pure spread quotes and points up quotes on CDS.

To keep things simple, suppose that both the swap and the hazard rate curves are flat. The swap rate is a continuously compounded r for any term to maturity, and the hazard rate is λ for any horizon. Suppose further that the running spread is 500 bps. The fair market CDS spread will then be a constant s for any term τ . From Equation (7.3), the expected present value of all the payments by the protection buyer must equal the expected present

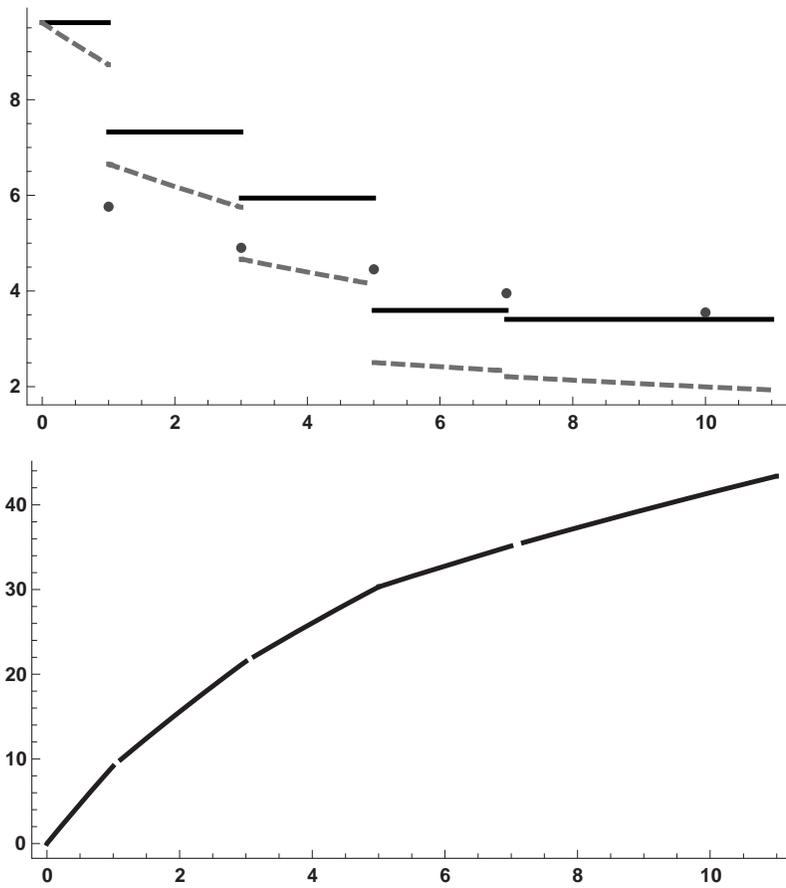


FIGURE 7.5 Estimation of Default Curves

Upper panel shows the CDS spreads from which the hazard rates are computed as dots, the estimated hazard rates as a step function (solid plot). The default density is shown as a dashed plot.

Lower panel shows the default distribution. Notice the discontinuities of slope as we move from one hazard rate to the next.

value of loss given default, so the points upfront and the constant hazard rate must satisfy

$$\frac{\text{points upfront}}{100} = (1 - R) \sum_{u=1}^{4\tau} e^{-r \frac{u}{4}} \left(e^{-\lambda \frac{u-1}{4}} - e^{-\lambda \frac{u}{4}} \right) - \frac{500}{4 \times 10^4} \sum_{u=1}^{4\tau} e^{-r \frac{u}{4}} \left[e^{-\lambda \frac{u}{4}} + \frac{1}{2} \left(e^{-\lambda \frac{u-1}{4}} - e^{-\lambda \frac{u}{4}} \right) \right] \quad (7.4)$$

Once we have a hazard rate, we can solve Equation (7.4) for the spread from the points upfront, and vice versa.

7.3.5 The Slope of Default Probability Curves

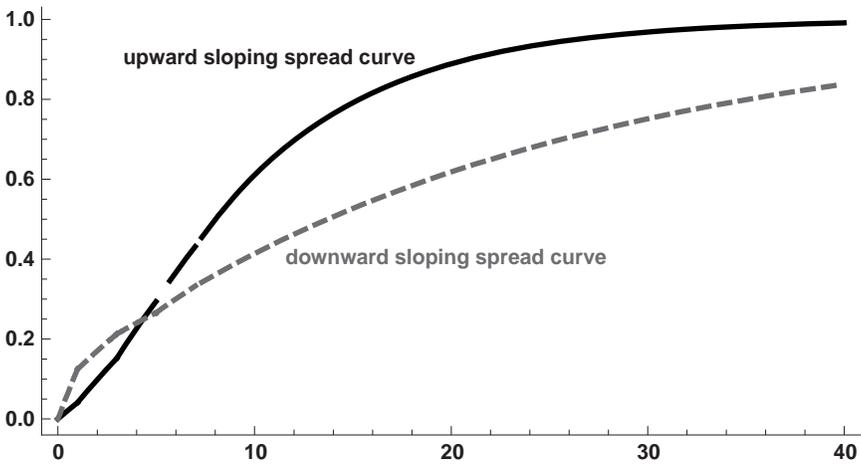
Spread curves, and thus hazard curves, may be upward- or downward-sloping. An upward-sloping spread curve leads to a default distribution that has a relatively flat slope for shorter horizons, but a steeper slope for more distant ones. The intuition is that the credit has a better risk-neutral chance of surviving the next few years, since its hazard rate and thus unconditional default probability has a relatively low starting point. But even so, its marginal default probability, that is, the conditional probability of defaulting in future years, will fall less quickly or even rise for some horizons.

A downward-sloping curve, in contrast, has a relatively steep slope at short horizons, but flattens out more quickly at longer horizons. The intuition here is that, if the firm survives the early, “dangerous” years, it has a good chance of surviving for a long time.

An example is shown in Figure 7.6. Both the upward- and downward-sloping spread curves have a five-year spread of 400 basis points. The downward-sloping curve corresponds to an unconditional default probability that is higher than that of the upward-sloping curve for short horizons, but significantly lower than that of the upward-sloping curve for longer horizons.

Spread curves are typically gently upward sloping. If the market believes that a firm has a stable, low default probability that is unlikely to change for the foreseeable future, the firm’s spread curve would be flat if it reflected default expectations only. However, spreads also reflect some compensation for risk. For longer horizons, there is a greater likelihood of an unforeseen and unforeseeable change in the firm’s situation and a rise in its default probability. The increased spread for longer horizons is in part a risk premium that compensates for this possibility.

Downward-sloping spread curves are unusual, a sign that the market views a credit as distressed, but became prevalent during the subprime crisis. Figure 7.7 displays an example typical for financial intermediaries, that of Morgan Stanley (ticker MS), one of the five large broker-dealers not associated with a large commercial bank within a bank holding company during the period preceding the crisis. (The other large broker-dealers were Bear Stearns, Lehman Brothers, Merrill Lynch, and Goldman Sachs.) Before the crisis, the MS spread curve was upward-sloping. The level of spreads was, in retrospect, remarkably low; the five-year CDS spread on Sep. 25, 2006 was a mere 21 basis points, suggesting the market considered a Morgan Stanley bankruptcy a highly unlikely event.



Graph displays cumulative default distributions computed from these CDS curves, in basis points:

Term	Upward	Downward
1	250	800
3	325	500
5	400	400
7	450	375
10	500	350

A constant swap rate of 4.5 percent and a recovery rate of 40 percent were used in extracting hazard rates.

FIGURE 7.6 Spread Curve Slope and Default Distribution

The bankruptcy of Lehman Brothers cast doubt on the ability of any of the remaining broker-dealers to survive, and also showed that it was entirely possible that senior creditors of these institutions would suffer severe credit losses. Morgan Stanley in particular among the remaining broker-dealers looked very vulnerable. Bear Stearns had already disappeared; Merrill Lynch appeared likely to be acquired by a large commercial bank, Bank of America; and Goldman Sachs had received some fresh capital and was considered less exposed to credit losses than its peers.

By September 25, 2008, the five-year CDS spread on MS senior unsecured debt had risen to 769 basis points. Its 6-month CDS spread was more than 500 basis points higher at 1,325 bps. At a recovery rate of 40 percent, this corresponded to about a 12 percent probability of bankruptcy over the

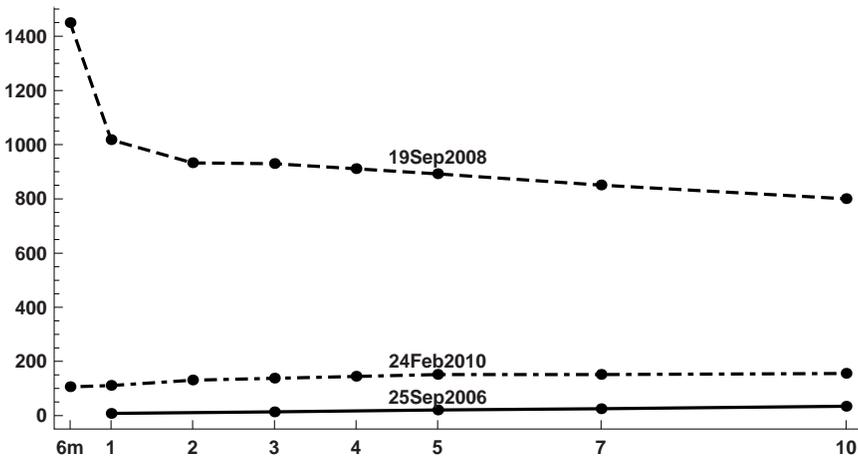


FIGURE 7.7 Morgan Stanley CDS Curves, select dates
 Morgan Stanley senior unsecured CDS spreads, basis points.
 Source: Bloomberg Financial L.P.

next half-year. The one-year spread was over 150 times larger than two years earlier. Short selling of MS common equity was also widely reported, even after the company announced on September 25 that the Federal Reserve Board had approved its application to become a bank holding company.

One year later, the level of spreads had declined significantly, though they remained much higher than before the crisis. On Feb. 24, 2010, the MS five-year senior unsecured CDS spread was 147 basis points, and the curve was gently upward-sloping again.

7.4 SPREAD RISK

Spread risk is the risk of loss from changes in the pricing of credit-risky securities. Although it only affects credit portfolios, it is closer in nature to market than to credit risk, since it is generated by changes in prices rather than changes in the credit state of the securities.

7.4.1 Mark-to-Market of a CDS

We can use the analytics of the previous section to compute the effect on the mark-to-market value of a CDS of a change in the market-clearing

premium. At initiation, the mark-to-market value of the CDS is zero; neither counterparty owes the other anything. If the spread increases, the premium paid by the fixed-leg counterparty increases. This causes a gain to existing fixed-leg payers, who in retrospect got into their positions cheap, and a loss to the contingent-leg parties, who are receiving less premium than if they had entered the position after the spread widening. This mark-to-market effect is the spread01 of the CDS.

To compute the mark-to-market, we carry out the same steps needed to compute the spread01 of a fixed-rate bond. In this case, however, rather than increasing and decreasing one spread number, the z -spread, by 0.5 bps, we carry out a parallel shift up and down of the entire CDS curve by 0.5 bps. This is similar to the procedure we carried out in computing DV01 for a default-free bond, in which we shifted the entire spot curve up or down by 0.5 bps.

For each shift of the CDS curve away from its initial level, we recompute the hazard rate curve, and with the shocked hazard rate curve we then recompute the value of the CDS. The difference between the two shocked values is the spread01 of the CDS.

7.4.2 Spread Volatility

Fluctuations in the prices of credit-risky bonds due to the market assessment of the value of default and credit transition risk, as opposed to changes in risk-free rates, are expressed in changes in credit spreads. Spread risk therefore encompasses both the market's expectations of credit risk events and the credit spread it requires in equilibrium to put up with credit risk. The most common way of measuring spread risk is via the *spread volatility* or "spread vol," the degree to which spreads fluctuate over time. Spread vol is the standard deviation—historical or expected—of changes in spread, generally measured in basis points per day.

Figure 7.8 illustrates the calculations with the spread volatility of five-year CDS on Citigroup senior U.S. dollar-denominated bonds. We return to the example of Citigroup debt in our discussions of asset price behavior during financial crises in Chapter 14. The enormous range of variation and potential for extreme spread volatility is clear from the top panel, which plots the spread levels in basis points. The center panel shows daily spread changes (also in bps). The largest changes occur in the late summer and early autumn of 2008, as the collapses of Fannie Mae and Freddie Mac, and then of Lehman, shook confidence in the solvency of large intermediaries, and Citigroup in particular. Many of the spread changes during this period are extreme outliers from the average—as measured by the root mean square—over the entire period from 2006 to 2010.

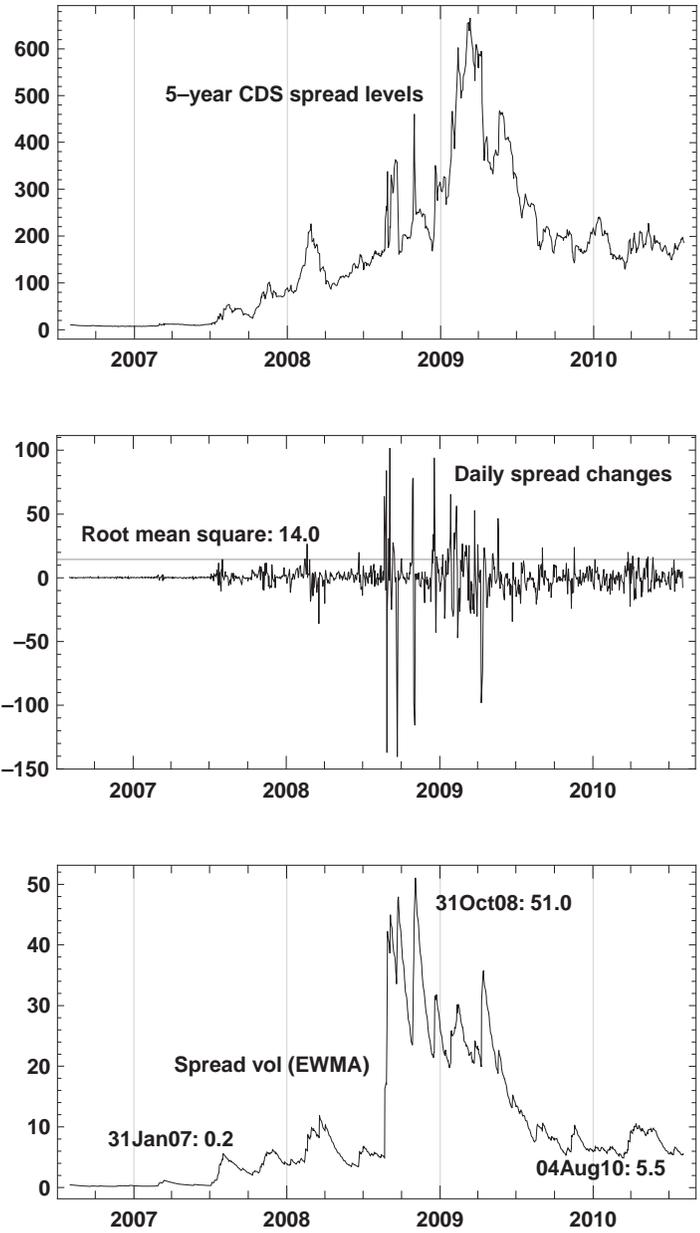


FIGURE 7.8 Measuring Spread Volatility: Citigroup Spreads 2006–2010
Citigroup 5-year CDS spreads, August 2, 2006, to September 2, 2010. All data expressed in bps.

Source: Bloomberg Financial L.P.

Upper panel: Spread levels.

Center panel: Daily spread changes.

Lower panel: Daily EWMA estimate of spread volatility at a daily rate.

The bottom panel plots a rolling daily spread volatility estimate, using the EWMA weighting scheme of Chapter 3. The calculations are carried out using the recursive form Equation (3.2), with the root mean square of the first 200 observations of spread changes as the starting point. The volatility is expressed in basis points per day. A spread volatility of, say, 10 bps, means that, if you believe spread changes are normally distributed, you would assign a probability of about 2 out of 3 to the event that tomorrow's spread level is within ± 10 bps of today's level. For the early part of the period, the spread volatility is close to zero, a mere quarter of a basis point, but spiked to over 50 bps in the fall of 2008.

FURTHER READING

Klugman, Panjer, and Willmot (2008) provides an accessible introduction to hazard rate models. Litterman and Iben (1991); Berd, Mashal, and Wang (2003); and O'Kane and Sen (2004) apply hazard rate models to extract default probabilities. Schönbucher (2003), Chapters 4–5 and 7, is a clear exposition of the algebra.

Duffie (1999), Hull and White (2000), and O'Kane and Turnbull (2003) provide overviews of CDS pricing. Houweling and Vorst (2005) is an empirical study that finds hazard rate models to be reasonably accurate.

See Markit Partners (2009) and Senior Supervisors Group (2009a) on the 2009 change in CDS conventions.