COLUMBIA
BUSINESS
SCHOOL

Probability Models of Credit Risk

In discussing financial risk, it is useful to distinguish between *market risk* and *credit risk*. Market risk refers to the possibility of losses due to changes in the prices of financial assets, whereas credit risk refers to the possible failure of a party to make a contractual payment. Shares of stocks are subject only to market risk because they do not carry a promise of payment. In contrast, the risk in a bank loan lies primarily in the possibility that the borrower may not be able to make scheduled payments.

The distinction between market and credit risk is not always very precise. A corporate bond, for example, carries both types of risk because its value is sensitive both to interest rates and to the creditworthiness of the issuer. Moreover, the two types of risk can interact because changes in the market value of a firm's assets may affect its ability to pay its debt.

The measurement of credit risk has grown substantially more complicated in recent years. In the past, banks made loans and generally held the loans on their books. Their success relied on their ability to gauge the creditworthiness of clients. But it is increasingly common for banks to sell loans, to securitize loans, or to enter into credit swaps, all of which are means of transferring credit risk. Fund managers are also taking advantage of new ways of transacting in credit risk through, for example, a burgeoning market for credit derivatives.

The growing complexity of credit risk has fueled the development of sophisticated methods for measuring credit risk at a *portfolio* level, rather than just at the level of an individual bond or loan. These methods are largely statistical and build on probabilistic models of creditworthiness and asset values. Some of the best known methods developed in recent years are the following:

- *CreditMetrics*, originally developed by J.P. Morgan
- *CreditRisk+*, developed by Credit Suisse Financial Products
- *CreditPortfolioView*, developed by McKinsey & Co.
- *Credit Monitor*, developed by KMV.[2]

A comprehensive discussion of these methods is far beyond the scope of this note.[3] We give only a rudimentary introduction to J.P. Morgan's CreditMetrics approach by discussing some examples from *Introduction to CreditMetrics* published by J.P. Morgan. The complete document can be downloaded from www.jpmorgan.com.

Credit Ratings

An essential feature of the CreditMetrics methodology is that it uses *credit ratings* of the type published by Moody's Investor Services and Standard & Poors. A credit rating is a measure of the creditworthiness of an issuer (which could be a company, a municipality, or a country, for example) or of an individual bond or note. Rating systems typically use letter grades, with AAA (for S&P) or Aaa (for Moody's) representing the highest ratings, and grades like BBB or Baa the minimum to qualify as investment grade; lower ratings, like CCC indicate what are sometimes called speculative issues, or simply junk bonds.

We will follow the CreditMetrics documentation in considering an example with eight ratings: AAA, AA, A, BBB, BB, B, CCC, and Default. A default rating generally indicates that the obligor has failed to make a scheduled payment, but it is important to stress that a bond in default is by no means worthless since the issuer may be able to make payments in the future. Indeed, creditors often end up receiving anywhere between 10% and 90% of a bond's face value in case of default.

Ratings change over time. A recession or a change in technology may adversely affect a company's profitability and thus its capacity to service its debt; this may result in a ratings downgrade. In the case of sovereign debt, a change in government economic policy may result in either a better or worse rating. Even universities will see their credit ratings change depending on how they manage their endowment and their operating expenses.[4]

For an investor holding a bond, a downgrade in the bond's rating can result in a financial loss, even if the bond's issuer has continued to make all

[2] The "K" in KMV is Stephen Kealhofer, a former member of the Columbia Business School faculty. Visit www.kmv.com to learn more about this company.

[3] For a highly readable introduction, see *Credit Risk Measurement* by Anthony Saunders, Wiley, 1999.

[4] Since 1998, Columbia has had the highest ratings, AAA and Aaa, from Standard & Poors and Moody's.

scheduled payments. If, for example, a AA bond is downgraded to A, it will be perceived as riskier and thus generally experience a drop in market value. Thus, ratings transitions influence the market value of bonds, even if no default occurs.

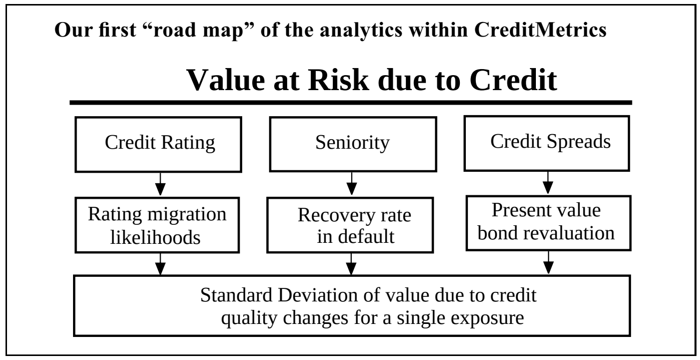

Figure 1: Schematic view of methodology, from *Introduction to CreditMetrics*, p.23.

CreditMetrics (and similar methodologies) attempt to measure the risk resulting from ratings transition. Figure 1, from the CreditMetrics document gives a schematic representation of the approach. An essential element of this approach is a *ratings transition matrix* giving the probabilities of rating changes over a period of, e.g., one year. Table 1 gives an example. Each row corresponds to an initial rating and each column corresponds to a year-end rating. To find the probability that a bond rated A today will be rated BBB one year from today, we read across the A row until we get to the BBB column and find the probability is 5.52%. Each row sums to 100% because a bond has to end the year in one of the eight column categories.[5] Notice that the largest probabilities are on the diagonal, indicating that most ratings do not change in the course of a year. Notice also that some transitions probabilities are 0 (at least to several decimal places). This indicates, for example, that the chance of a AAA-rated bond defaulting within one year is negligible. Figure 2 further illustrates the interpretation of each row of the transition matrix.

The second essential component of the CreditMetrics approach is a model of how a ratings transition affects the market value of a bond.[6] This idea is best illustrated through an example. As in Chapter 1 of the CreditMetrics document, consider a bond with an initial rating of BBB. Suppose the bond has a face value of $100 and pays an annual coupon of 6%.[7] Under some

[5]This is a simplifying assumption. In reality, a bond may have its rating withdrawn or the bond may be paid up, in which cases the bond would leave the ratings system and not end up in any of the eight columns in the table.

[6]The same approach can be used for other credit-risky assets; we discuss the case of a bond to be concrete.

[7]This means that the bond pays $6.00 each year and $100 at maturity.

One-year transition matrix (%)

Initial Rating	Rating at year-end (%)							
	AAA	AA	A	BBB	BB	B	CCC	Default
AAA	90.81	8.33	0.68	0.06	0.12	0	0	0
AA	0.70	90.65	7.79	0.64	0.06	0.14	0.02	0
A	0.09	2.27	91.05	5.52	0.74	0.26	0.01	0.06
BBB	0.02	0.33	5.95	86.93	5.30	1.17	0.12	0.18
BB	0.03	0.14	0.67	7.73	80.53	8.84	1.00	1.06
B	0	0.11	0.24	0.43	6.48	83.46	4.07	5.20
CCC	0.22	0	0.22	1.30	2.38	11.24	64.86	19.79

Source: Standard & Poor's CreditWeek (15 April 96)

Table 1: Rating transition matrix, from *Introduction to CreditMetrics*, p.20.

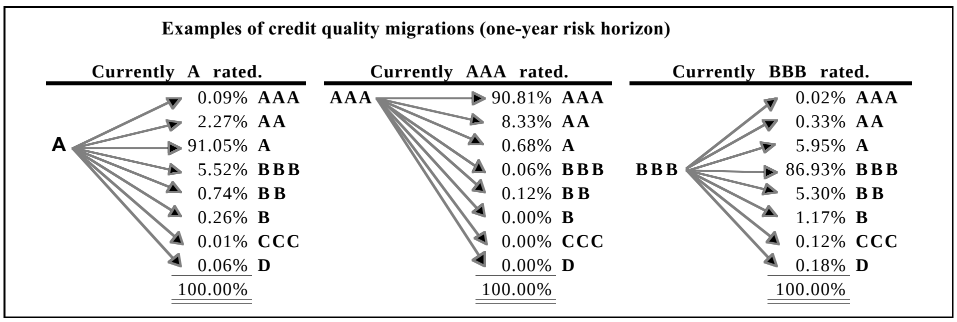

Figure 2: Ratings transitions, from *Introduction to CreditMetrics*, p.24.

pricing assumptions, Table 2 shows what the bond will be worth in one year as a function of its credit rating in one year. If, for example, it stays at BBB, it will be worth \$107.55; if it is upgraded to A, it will be worth \$108.66. According to the table, in case of default the value will be \$51.13; think of this as the amount bondholders can expect to recover in case of bankruptcy—roughly fifty cents on the dollar.

In practice, not all bonds with the same credit rating will have the same value (even if they have the same coupon and maturity). The values in the Table 2 may be thought of as *average* values within each category. The CreditMetrics document explains how to include information about the standard deviation of bond values within each category; however, for simplicity, we will ignore this issue and pretend that the possible bond values are exactly as given in the table.

Risk Measurement

We can now discuss how CreditMetrics combines a ratings transition matrix with a table of possible bond values to measure risk. Two measures of risk are discussed in the CreditMetrics document: standard deviation and a percentile or value-at-risk measure. We consider both as well.

Possible one-year forward values for a BBB bond plus coupon

Year-end rating	Value ($)
AAA	109.37
AA	109.19
A	108.66
BBB	107.55
BB	102.02
B	98.10
CCC	83.64
Default	51.13

Table 2: Possible values of a bond initially rated BBB as a function of year-end rating, from *Introduction to CreditMetrics*, p.10.

Although it is ultimately more important to consider risk at a portfolio level, we begin by considering a single bond. We consider the same BBB bond whose possible values one year from today are given in Table 2. Because we focus on just this one bond, only the BBB row of the ratings transition matrix is relevant. We can therefore pull this row out of the matrix to list the possible values of the bond together with their probabilities, as in Table 3.

Distribution of value of a BBB par bond in one year

Year-end rating	Value ($)	Probability (%)
AAA	109.37	0.02
AA	109.19	0.33
A	108.66	5.95
BBB	107.55	86.93
BB	102.02	5.30
B	98.10	1.17
CCC	83.64	0.12
Default	51.13	0.18

Table 3: Possible values and their probabilities for a bond initially rated BBB, from *Introduction to CreditMetrics*, p.11.

From this table, we can immediately determine a percentile or value-at-risk measure. Suppose we want to find a dollar amount such that the probability that the bond will be worth that amount or less is at most 5%. From the table we find that there is a 0.18% chance that the bond will be worth 51.13; there is a 0.30% (= 0.18%+0.12%) chance that the bond will be worth 83.64 or less; there is a 1.47% (= 0.18% + 0.12% + 1.17%) chance that the bond will be worth 98.10 or less; and there is a 6.77% that the bond will be worth 102.02 or less. So, we can say that the probability the bond will be worth 98.10 or less is below the 5% limit, but we cannot say that about 102.02. To hit 5% exactly, we may choose to interpolate:

$$\left(\frac{6.77 - 5}{6.77 - 1.47}\right) 98.10 + \left(\frac{5 - 1.47}{6.77 - 1.47}\right) 102.02 = 100.71.$$

This formula weights the two values by their proximity to 5%.

To calculate the standard deviation of the bond's value, we first need to find the expected value. We do this by weighting each possible value by its probability and then summing. This gives (see Table 3)

$$\begin{aligned}
&109.37 \cdot 0.0002 + 109.19 \cdot 0.0033 + 108.66 \cdot 0.0595 + 107.55 \cdot .8693 \\
&+102.02 \cdot .0530 + 98.10 \cdot 0.0117 + 83.64 \cdot 0.0012 + 51.13 \cdot 0.0018 \\
&= 107.09.
\end{aligned}$$

Thus, the expected value of the bond one year from today is \$107.09.

To proceed with the calculation of the standard deviation, we recall that the standard deviation is the square root of the variance, and that the variance of a random variable X is given by the formula

$$\mathsf{Var}[X] = \mathsf{E}[X^2] - (\mathsf{E}[X])^2.$$

We have already calculated the expected value, so the next step is to calculated the expected *squared* value. We do this by taking the probability-weighted average of the possible squared bond values:

$$(109.37)^2 \cdot 0.0002 + (109.19)^2 \cdot 0.0033 + (108.66)^2 \cdot 0.0595 + (107.55)^2 \cdot .8693 + (102.02)^2 \cdot .0530 + (98.10)^2 \cdot 0.0117 + (83.64)^2 \cdot 0.0012 + (51.13)^2 \cdot 0.0018 = 11476.77.$$

We now find that the variance is $11476.77 - (107.09)^2 = 8.95$ and the standard deviation is $\sqrt{8.95} = 2.99$, or \$2.99.

In the CreditMetrics document, this value is calculated using the alternative formula

$$\mathsf{Var}[X] = \mathsf{E}[(X - 107.09)^2].$$

To use this formula, we take the probability-weighted average of the squared differences between the bond values and the expected value of 107.09:

$$(109.37 - 107.09)^2 \cdot 0.0002 + (109.19 - 107.09)^2 \cdot 0.0033 + (108.66 - 107.09)^2 \cdot 0.0595 + (107.55 - 107.09)^2 \cdot .8693 + (102.02 - 107.09)^2 \cdot .0530 + (98.10 - 107.09)^2 \cdot 0.0117 + (83.64 - 107.09)^2 \cdot 0.0012 + (51.13 - 107.09)^2 \cdot 0.0018 = 8.95.$$

Once again, taking the square root produces the standard deviation of \$2.99. The intermediate steps are detailed in Table 4. (The values obtained using alternative formulas may differ slightly because of rounding.)

Calculating volatility in value due to credit quality changes

Year-end rating	Probability of state (%)	New bond value plus coupon ($)	Probability weighted value ($)	Difference of value from mean ($)	Probability weighted difference squared
AAA	0.02	109.37	0.02	2.28	0.0010
AA	0.33	109.19	0.36	2.10	0.0146
A	5.95	108.66	6.47	1.57	0.1474
BBB	86.93	107.55	93.49	0.46	0.1853
BB	5.30	102.02	5.41	(5.06)	1.3592
B	1.17	98.10	1.15	(8.99)	0.9446
CCC	0.12	83.64	1.10	(23.45)	0.6598
Default	0.18	51.13	0.09	(55.96)	5.6358
		Mean =	$107.09	Variance =	8.9477
				Standard deviation =	$2.99

Table 4: Calculation of standard deviation, from *Introduction to CreditMetrics*, p.28.

To summarize, the expected value of the bond is \$107.09 and its standard deviation is \$2.99. We estimated that there is a 5% chance that the bond will be worth \$100.71 or less. Notice that the difference $107.09 - 100.11$ is 6.38 and $6.38/2.99 \approx 2.5$. This indicates that 100.71 is about 2.5 standard deviations below the mean. The fact that there is approximately a 5% probability that the value will be more than 2.5 standard deviations below the mean reflects the skew in the distribution of value; see Figure 3.

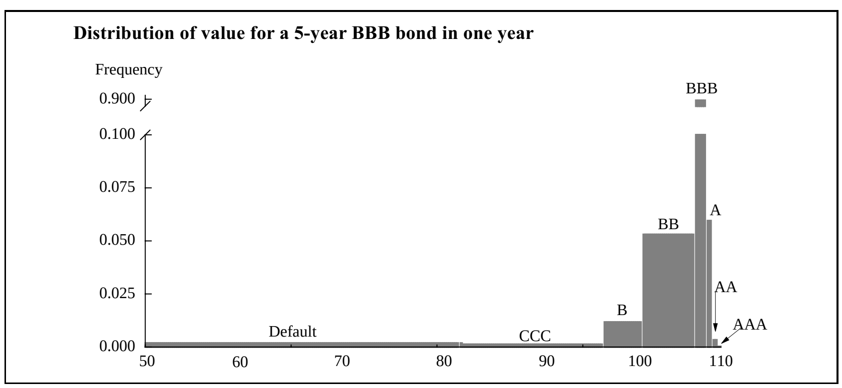

Figure 3: Distribution of possible values for one bond, from *Introduction to CreditMetrics*, p.11.

Portfolio Risk

Similar calculations can be used to determine the standard deviation or a percentile of the value distribution for a portfolio of bonds. The main additional consideration in moving from a single bond to two or more bonds is capturing the possibility that the ratings of the two bonds may not be independent. For example, a recession is likely to affect the creditworthiness of many companies; a downturn in a particular market sector may affect many companies in the same industry; geographical factors can also link the creditworthiness of issuers. An important feature of the CreditMetrics methodology is that attempts to measure the propensity of ratings to move together.

Consider a portfolio consisting of the same BBB bond as before and also a 5% coupon single-A bond. The possible values of the single-A bond are given in Table 5. (The probabilities that appear there are just the single-A row from the ratings transition matrix in Table 1.) The possible values of the portfolio are the possible values of the *sum* of the two individual bond values; because there are eight possible values for each bond, there are 64 possible values for the portfolio. These are given in Table 6.

Distribution of value of A bond in one year

Year-end rating	Value ($)	Probability (%)
AAA	106.59	0.09
AA	106.49	2.27
A	106.30	91.05
BBB	105.64	5.52
BB	103.15	0.74
B	101.39	0.26
CCC	88.71	0.01
Default	51.13	0.06

Table 5: Possible values and their probabilities for a bond initially rated A, from *Introduction to CreditMetrics*, p.12.

All possible 64 year-end values for a two-bond portfolio ($)

Obligor #1 (BBB)		Obligor #2 (single-A)							
		AAA	AA	A	BBB	BB	B	CCC	Default
		106.59	106.49	106.30	105.64	103.15	101.39	88.71	51.13
AAA	109.37	215.96	215.86	215.67	215.01	212.52	210.76	198.08	160.50
AA	109.19	215.78	215.68	215.49	214.83	212.34	210.58	197.90	160.32
A	108.66	215.25	215.15	214.96	214.30	211.81	210.05	197.37	159.79
BBB	107.55	214.14	214.04	213.85	213.19	210.70	208.94	196.26	158.68
BB	102.02	208.61	208.51	208.33	207.66	205.17	203.41	190.73	153.15
B	98.10	204.69	204.59	204.40	203.74	201.25	199.49	186.81	149.23
CCC	83.64	190.23	190.13	189.94	189.28	186.79	185.03	172.35	134.77
Default	51.13	157.72	157.62	157.43	156.77	154.28	152.52	139.84	102.26

Table 6: Possible values for two-bond portfolio, from *Introduction to CreditMetrics*, p.12.

What probabilities should we assign to these 64 possible values? If we assume the ratings of the two bonds move independently of each other, then the probability of any combination is the product of the probabilities of the individual transitions. For example,

$$
\begin{aligned}
& P(\text{BBB moves to BB and A stays at A}) \\
= \quad & P(\text{BBB moves to BB}) \times P(\text{A stays at A}) \\
= \quad & .0530 \times 0.9105 = .0483.
\end{aligned}
$$

The values of Table 7 are all computed this way. You can find 4.83% in row BB under column A.[8]

The CreditMetrics document provides a second table of joint probabilities for the two bonds in which their ratings transitions are no longer assumed independent. These are given in Table 8. In the terminology of CreditMetrics, these probabilities are calculated assuming that the *assets* of the two firms have a correlation of 0.30. (It would not make sense to refer to a correlation between ratings because ratings are categorical rather than numerical data.) We will not discuss how CreditMetrics calculates these joint probabilities—we will simply take them as given. Notice that we can recover the *marginal* transition probabilities of the two bonds by summing the rows and columns, but—in contrast to the independent case—there is no way to go from the marginal probabilities to the joint probabilities in this table.

Joint migration probabilities with zero correlation (%)

Obligor #1 (BBB)		Obligor #2 (single-A)							
		AAA	AA	A	BBB	BB	B	CCC	Default
		0.09	2.27	91.05	5.52	0.74	0.26	0.01	0.06
AAA	0.02	0.00	0.00	0.02	0.00	0.00	0.00	0.00	0.00
AA	0.33	0.00	0.01	0.30	0.02	0.00	0.00	0.00	0.00
A	5.95	0.01	0.14	5.42	0.33	0.04	0.02	0.00	0.00
BBB	86.93	0.08	1.97	79.15	4.80	0.64	0.23	0.01	0.05
BB	5.30	0.00	0.12	4.83	0.29	0.04	0.01	0.00	0.00
B	1.17	0.00	0.03	1.07	0.06	0.01	0.00	0.00	0.00
CCC	0.12	0.00	0.00	0.11	0.01	0.00	0.00	0.00	0.00
Default	0.18	0.00	0.00	0.16	0.01	0.00	0.00	0.00	0.00

Table 7: Joint distribution of ratings transitions for two bonds assuming independence, from *Introduction to CreditMetrics*, p.36.

Once we have determined the 64 possible portfolio values and the chosen a corresponding set of 64 probabilities, we can calculate the expected value, the standard deviation, and a percentile for the value of the portfolio. This is what CreditMetrics provides, except that it does it for portfolios containing thousands of bonds.

[8]This table should not be confused with the ratings transition matrix. Here, the sum of all 64 entries gives 100% because the table entries give probabilities of *pairs* of transitions, assuming one bond starts at BBB and one starts at A. In the ratings transition matrix, each row sums to 100% and each row corresponds to a different initial rating.

Joint migration probabilities with 0.30 asset correlation (%)

Obligor #1 (BBB)		**Obligor #2 (single-A)**							
		AAA	AA	A	BBB	BB	B	CCC	Default
		0.09	2.27	91.05	5.52	0.74	0.26	0.01	0.06
AAA	0.02	0.00	0.00	0.02	0.00	0.00	0.00	0.00	0.00
AA	0.33	0.00	0.04	0.29	0.00	0.00	0.00	0.00	0.00
A	5.95	0.02	0.39	5.44	0.08	0.01	0.00	0.00	0.00
BBB	86.93	0.07	1.81	79.69	4.55	0.57	0.19	0.01	0.04
BB	5.30	0.00	0.02	4.47	0.65	0.11	0.04	0.00	0.01
B	1.17	0.00	0.00	0.92	0.18	0.04	0.02	0.00	0.00
CCC	0.12	0.00	0.00	0.09	0.02	0.00	0.00	0.00	0.00
Default	0.18	0.00	0.00	0.13	0.04	0.01	0.00	0.00	0.00

Table 8: Joint distribution of ratings transitions for two bonds assuming correlation between the assets of the two firms, from *Introduction to CreditMetrics*, p.38.