

THE ESSENTIALS OF RISK MANAGEMENT

Third Edition (EoRM3)

ONLINE CONTENT FOR CHAPTER 11

Excerpt Ch. 11 [Updated] and Appendix 11-1 of EoRM2

From pp. 392–400 of *The Essentials of Risk Management, Second Edition* (EoRM2) [Updated]

The Actuarial and Reduced-Form Approaches to Measuring Credit Risk

Two other approaches to estimating portfolio credit risk have been proposed. These are the actuarial approach, based on statistical models used by the insurance industry, and the reduced-form approach.

The structural model of default (e.g., the Merton 1974 framework) is based on information extracted from equity prices and the fundamental analysis of the balance sheet of the firm; it models the way in which firms default when their asset value falls below a certain boundary. By contrast, the actuarial model and the reduced-form models treat the firm's road to bankruptcy, including recovery, as factors external to the modeling process—that is, they make assumptions about the number and timing of default events rather than attempting to derive them internally. The actuarial approach uses historical data on defaults to calibrate default probabilities while the reduced-form approach uses the information from the market credit spreads in bonds or CDSs.

The actuarial approach

CreditRisk+, released in late 1997 by the investment bank Credit Suisse Financial Products (CSFP), is a purely actuarial model, based on mortality models developed by insurance companies. The probabilities of default that the model employs are based on historical statistical data on default experience. Unlike the Moody's Analytic approach, there is no attempt to relate default to a firm's capital structure or balance sheet.

CreditRisk+ makes a number of assumptions:

- For a loan, the probability of default in a given period (e.g., one month), is the same as in any other period of the same length, say, another month.
- For many obligors, the probability of default by any particular obligor is small, and the number of defaults that occur in any given period is independent of the number of defaults that occur in any other period.

Under these assumptions, and based on empirical observation, the probability distribution for the number of defaults during a given period of time (e.g., one year) is well represented by a certain shape

of statistical distribution known as a Poisson distribution. We expect the mean default rate to change over time depending on the business cycle. This suggests that the distribution can be used to represent the default process only if, as CreditRisk+ suggests, we make the additional assumption that the mean default rate is itself changing, following a certain distribution.

In CreditRisk+, obligors are divided into bands, or subportfolios, and all obligors in a band are characterized by approximately the same average exposure, net of the recovery adjustments. If we know the distribution of defaults in each band, then we can find the distribution of defaults over all bands for the whole portfolio. CreditRisk+ derives a closed-form solution for the loss distribution of the loan portfolio.

CreditRisk+ has the advantage that it is relatively easy to implement. First, closed-form expressions—that is, explicit formulas—can be derived for the probability of portfolio bond or loan losses, and this makes CreditRisk+ very attractive from a computational point of view. In addition, marginal risk contributions by obligor can be computed. Second, CreditRisk+ focuses on default, and therefore it requires relatively few estimates and inputs. For each instrument, only the probability of default and the loss given default statistics are required.

One disadvantage of CreditRisk+ is that it ignores migration risk: the exposure for each obligor is fixed and is not sensitive to possible future changes in the credit quality of the issuer. This is the major difference between the approaches of CreditRisk+ and CreditMetrics.

Default correlations between obligors only arise through dependence on a common set of risk factors—for example, macroeconomic indexes for the areas of capital markets, consumption, employment, foreign investments, income, and prices. These risk factors are incorporated into the specification of the default rates, allowing the default rate itself to have a probability distribution. When the default rate volatilities are set to zero, the default events become independent.

Even in its most general form, where the probability of default depends upon several stochastic risk factors, the credit exposures in the portfolio under examination are taken to be constant and are not related to changes in these factors. In reality, credit exposure is often linked quite closely to risk factors such as interest rates and the probability of default. For example, in the case of a loan commitment, a corporate borrower has the option of drawing on its credit line—and is more likely to exercise this option when its credit standing is deteriorating. Finally, like CreditMetrics and Moody's Analytics, CreditRisk+ is not able to cope satisfactorily with nonlinear products such as options and foreign currency swaps.

The reduced-form approach

Structural models imply that default events are predictable, since the evolution of the value of the firm follows a continuous process without jumps.¹ However, default may happen at any time—a phenomenon that is known as “jump to default.” For example, at the time of the millennial stock boom and bust, certain large firms such as Enron and Parmalat surprised investors by suddenly defaulting due to accounting frauds.² A firm may be solvent from the perspective of its long-term economics but simply run out of cash and become unable to pay salaries and make other critical payments; if the firm

1. This is true for the standard Merton model. However, more sophisticated models now incorporate random default barriers, and models have been developed under which the value of the firm follows a process with jumps.

2. However, Moody's Analytics has shown that both companies' EDFs had gradually deteriorated over the two or three years preceding the defaults.

is unable to borrow additional money from the bank, it is likely to be pushed into default. In addition to these issues, it has proved difficult to reproduce the credit spreads observed in the real world using structural models.

Reduced-form models, also called intensity-based models, were introduced to address some of these shortcomings.³ One of the fundamental advantages of the reduced-form approach is that it produces explicit formulas for the value of risky debt and credit derivative products, and provides a bridge between default probabilities and the credit spreads seen in the financial markets.

These advantages arise from the nature of reduced-form models. Unlike structural model approaches, such as Moody's Analytics, reduced-form models don't attempt to predict default by looking at its underlying causes; they are essentially statistical (like CreditRisk+) and are calibrated using credit spreads that are observable in the world's financial markets. The reduced-form approach is therefore less intuitive than a structural model from an economic point of view, and it does not necessarily require balance sheet information. Instead, the data used to feed the models are principally credit instrument prices derived from markets such as the corporate bond, corporate loan, and credit derivatives markets (in contrast to the equity price data from stock markets employed by the Moody's Analytics approach).

Although this distinction is important in principle, it is not absolute. Reduced-form models depend principally on credit market spreads, but they may also use other input factors—including equity prices and balance sheet information—to better disentangle the estimation of default probabilities from loss-given-default.⁴

In theory, by looking at the price of credit-risky securities over time and subtracting the price of similar securities that do not incur credit risk (such as U.S. government bonds), the “price of credit” can be made transparent. Unfortunately, there are many complications and other real-world problems. The reduced-form modelers' task is to overcome these problems and derive the term structure of risk-adjusted implied default probabilities from the term structure of credit spreads apparent in the market—and then to find the most valid way to reapply this information to a particular bank loan or portfolio in the pursuit of better risk analysis.

Another more challenging problem is that the world's credit markets are very imperfect as sources of data. Unfortunately, credit risk is not the only determinant of prices for credit-risky securities—various other risk factors and market inefficiencies interfere with the credit price signals. In particular, although the corporate bond markets are large, the market for each individual bond tends to be illiquid and much less transparent than share prices in an equity market, not least because many bond transactions are conducted over-the-counter rather than on a formal exchange. The heterogeneous nature of bonds as financial instruments is also tricky: many are structured with embedded options (e.g., convertible bonds, callable bonds), and bond prices may be affected by various regulations and taxes in local markets. Jarrow's generalization of the reduced-form model proposes a general formulation for the impact of liquidity on bond market prices, as discussed in Appendix 11-1 of EoRM2.

3. This approach models default as an exogenous event whose instantaneous rate of occurrence, the hazard rate, is the key parameter to calibrate.

4. In the model introduced by Jarrow and Turnbull, the loss given default and the determinants of default can be empirically estimated. Duffie and Singleton (1999) describe an extension of the model. However, their extension suffers from the fact that the loss given default and the determinants of default cannot be separately identified. This is known as the “indetermination problem”—that is, the calibration of the model produces one number, the expected loss by unit of exposure ($PD \times LGD$). The estimation of PD requires making an assumption about the LGD.

These are empirical challenges, but there are also more fundamental analytical challenges. For example, how do the various credit-risk factors interact over a period of time to produce the credit spread visible in bond market data? The relative contribution of default probability and loss given default (or recovery rate) is not at all clear in the market data, yet distinguishing between the effects of these factors in the historical market data is important if the modeling results are to be applied to predict the future default and loss given default rates of loan portfolios.⁵ Appendix 11-1 offers a more technical introduction to the basic ideas underlying reduced-form models, as pioneered by Jarrow and Turnbull in 1995.⁶

Kamakura-KRIS (Kamakura Risk Information Services) credit portfolio model⁷

Kamakura's credit portfolio model (KRIS) is based on Jarrow's generalization of the reduced-form model sketched in Appendix A11-1 of EoRM2. The approach proposed by Kamakura is quite general and can be applied to the construction of default models for all types of borrowers: retail, small business, large unlisted, municipal, and sovereign. Jarrow's model can be fitted to bond prices alone, to bond prices and equity prices simultaneously, to credit derivatives prices, and to historical data on defaults.

The fitting process involves hazard rate modeling, known as logistic regression⁸, an extension of credit scoring technology that has been used for retail and small business default probability estimation for many years. The logistic regression formula, which predicts the default hazard rate (i.e., the probability of default $P(t)$ for a given time period t , provided the firm has survived until that time) is a function of explanatory variables that take the following form:

$$P(t) = 1 / [1 + \exp(-\alpha - \sum_i \beta_i X_i)]$$

where the explanatory variables, X_i , $i=1, \dots, n$, include accounting financial ratios such as return on assets and leverage ratio; macroeconomic factors such as the unemployment rate; multiple inputs from the stock price, such as the monthly excess return over a stock market index, equity volatility, and their history for each company; multiple inputs from the comprehensive stock market index for the relevant country; the company's size relative to the total market capitalization of the relevant country; industry variables; seasonality; the depth of trading activity in a given company's stock, and so on. KRIS uses historical default data from the KRIS proprietary default database to calibrate the model. Jarrow's specification of the reduced-form model, which incorporates a liquidity premium, allows the endogenous estimation of the LGD.⁹

KRIS provides monthly default probabilities for a full 10-year horizon: $P(1)$ is the probability of default for the first month; $P(2)$ is the probability of default for the second month, conditional on

5. The problem is important because the default probability and the loss given default associated with an instrument vary over time and also depend on the nature of the credit portfolio. At the top of their business cycle, for example, airlines tend not to default, and if they do, any collateral to a bank loan can be sold for a top price. Five years later, the story on both counts will be very different.

6. R. A. Jarrow and S. M. Turnbull, "Pricing Derivatives on Financial Securities Subject to Credit Risk," *Journal of Finance* 50(1), March 1995, pp. 53–85.

7. SAS acquired Kamakura Corporation in June 2022.

8. Kamakura's initial default probability versions used logistic regression consistently, but a wide variety of techniques are now employed to take advantage of Stein's Paradox and its recent Bayesian extensions.


9. Jarrow's initial work on the liquidity premium has been extended by Hilscher, Jarrow, and van Deventer; see J. Hilscher, R. A. Jarrow, and D. R. van Deventer, "The Valuation of Corporate Bonds," working paper, January 2022.

surviving the first month; and so on. $P(120)$ is the probability of default for the last month of the forecasting period, conditional on surviving the first 119 previous months.¹⁰

Empirical studies show that the reduced-form approach has been better able to predict U.S. corporate defaults¹¹ than other approaches. For example, adding Moody's Analytic EDF or "distance-to-default" makes only a marginal contribution to the predictive accuracy of the model.¹²

Figure 11-9 shows the annualized KRIS term structure of default for Toronto Dominion Bank on September 19, 2022, during a period of heightened uncertainty for banks from both an interest rate risk and credit perspective. The upper part shows the annualized monthly probability of default over a 10-year horizon. For example, the annualized three-month default rate for Toronto Dominion on this date is 0.19 percent. The lower part of Figure 11-9 shows the same default risk for Toronto Dominion, this time displayed versus the annualized term structure for Toronto Dominion as the GFC unfolded as of November 28, 2008.¹³

FIGURE 11-9 Term Structure of KRIS Default Probabilities for Toronto Dominion on September 19, 2022¹⁴



Source: Kamakura.

10. These default probabilities can be combined to produce a term structure of cumulative default probabilities (Appendix A11-1 of EoRM2).

11. The KRIS version 7.0 models are benchmarked on data from 24 countries, and the same performance results apply internationally as well.

12. J. Y. Campbell, J. Hilscher, and J. Szilagyi, "In Search of Distress Risk," *The Journal of Finance*, 63(6), December 2008, pp. 2899–2939; J. Y. Campbell, J. Hilscher, and J. Szilagyi, "Predicting Financial Distress and the Performance of Distressed Stocks," *Journal of Investment Management* 9(2), 2011, pp. 1–21; and S. Bharath and T. Shumway, "Forecasting Default with the KMV-Merton Model," *Review of Financial Studies* 21(3), May 2008, pp. 1339–1369.

13. These default probabilities are benchmarked versus traded bond prices, both in-sample and out-of-sample, on a daily basis.

14. The first line, JC7, corresponds to the version of the KRIS model recommended by Kamakura.

Hybrid Structural Models

While the structural model approach is theoretically appealing, the predicted default probabilities and credit spreads calculated from the Merton model and some of its later extensions are too low compared to those observed empirically.¹⁵ The Moody's Analytics proprietary model represented one attempt to circumvent this limitation. More recently, researchers have proposed hybrid structural models that combine the structural model approach with additional accounting and credit information.

The underlying reason for this effort is that default is a complex process and cannot be described simply in terms of an asset value crossing a particular default point. Instead, we must try to take into account the firm's behavior as it approaches the default point. For example, firms that are solvent according to the Merton model can still default on their obligations as a result of severe liquidity problems. As the credit quality of a firm deteriorates, its capacity to borrow and to refinance its debt can determine whether or not it actually defaults.

The firm's borrowing capacity is the result of the borrower's ability to generate revenues for servicing its debt and the value of any assets it can employ as collateral. Both are clearly related to accounting information, such as the borrower's profitability, liquidity, and capital structure, as well as information about the business environment and the borrower's competitiveness (see Chapter 10 of EoRM3).

For this reason, proponents of a hybrid approach are attempting to bring together various accounting and market variables to describe the value of the firm's assets and its borrowing capacity. In one such recent approach, for example, the variables are the firm's market equity, stock volatility, stock return, book value of total assets, current liabilities, long-term debt, and net income.¹⁶

15. See D. Galai, A. Raviv, and Z. Wiener, "Liquidation Triggers and the Valuation of Equity and Debt," *Journal of Banking and Finance*, 2007, for extensions of Merton's economic model of the firm, taking into account various legal structures regarding bankruptcy laws.

16. J. Sobehart and S. Keenan, "Hybrid Probability of Default Models: A Practical Approach to Modeling Default Risk," working paper, Citigroup Global Markets, October 2003.

From pp. 407–410 of *The Essentials of Risk Management*, Second Edition (EoRM2)

Appendix 11-1

The Basic Idea of the Reduced-Form Model

Consider a defaultable zero-coupon debt instrument issued by a corporation with a promised payment at maturity, say, one year, of \$100.

Notations:

RR: recovery rate as a percentage of the promised payment = $1 - \text{LGD}$ (loss given default)

PD: one-year risk-neutral probability of default¹

i: one-year risk-free interest rate

y: one-year risk-adjusted yield

P: zero-coupon bond price

There are two approaches to value this defaultable zero-coupon bond:

- The risk-neutral valuation where expected cash flows (using risk-neutral probabilities) are discounted at the risk-free rate:

$$P = [100 (1 - \text{PD}) + 100 * \text{PD} * \text{RR}] / (1 + i) \quad (\text{A11-1})$$

1. This cumulative probability of default, PD, over the time horizon of one year for the purpose of this illustration, is derived from the term structure of “hazard rates”—that is, the probabilities of default over small time intervals. Denote by $\lambda(t)$ the hazard rate (or intensity), then the probability of default in the time interval $[t, t + \Delta t]$ is, by definition, $P(t) = \lambda(t)\Delta t$, having not defaulted before time t .

Then, the probability of survival—that is, the absence of default in the time interval $[t, t + \Delta t]$ —is: $1 - \lambda(t)\Delta t$.

Now, consider the time interval $[0, T]$ and assume that there are n time intervals $[t_{k-1}, t_k]$ of length Δt , with $k = 1$ to n , with a constant hazard rate $\lambda(k)$ in each time interval $[t_{k-1}, t_k]$. Then, the survival probability over the period $[0, T]$ —that is, the probability of no default between 0 and T —is:

$$S(T) = [1 - \lambda(1)\Delta t] [1 - \lambda(2)\Delta t] \dots [1 - \lambda(n)\Delta t] = [1 - P(1)] [1 - P(2)] \dots [1 - P(n)]$$

And the probability of default between time T and $T + \Delta t$, as seen at time 0, is:

$$q(T) \Delta t = \lambda(T) \Delta t \times S(T)$$

If we assume a constant hazard rate, λ , over the period $[0, T]$, then the risk-neutral survival probability, $S(T)$, and the risk-neutral default probability, $\text{PD}(T)$, over the time interval $[0, T]$ are respectively:

$$S(T) = \exp(-\lambda T) \text{ and } \text{PD}(T) = 1 - \exp(-\lambda T)$$

This formula tells us that survival probabilities have the same structure as discount factors, with the default intensity, or hazard rate, playing the role of interest rates. Thus, hazard rates can be viewed as credit spreads.

Assuming for simplicity that the hazard rate and interest rates are constant, then it can be shown that the fair value of the CDS spread is:

$$S = \lambda (1 - \text{RR})$$

- The risk-adjusted valuation where promised payments are discounted at a risk-adjusted rate of return, y :

$$P = 100 / (1 + y) \quad (\text{A11-2})$$

From these two expressions for the price of this defaultable zero-coupon bond, it follows:

$$1 + i = [1 - \text{PD} \times \text{LGD}] (1 + y) \quad (\text{A11-3})$$

Numerical example:

$$P = 95, i = 5\%, \text{ and } \text{RR} = 50\%$$

then, from (A11-2) $y = 5.26\%$ and from (A-11-3) it follows that $\text{PD} = 0.5\%$

The continuous time equivalent of (A11-3) is:

$$y = i + \text{PD} \times \text{LGD} \quad (\text{A11-4})$$

This one-period example can be generalized to a multiperiod setting. In Chapter 6 of EoRM2 we showed how to derive forward interest rates given a term structure of zero-coupon interest rates. Applying the reasoning that led to (A11-3) to forward rates allows to derive the forward default probabilities (hazard rates). In a two-period setting:

$P(1)$ denotes the probability of default during the first period, say one year.

$P(2)$ denotes the probability of default during the second period, say year 2, conditional on no default occurring during the first year. Then, the two-year cumulative probability of default, $\text{PD}(2)$, is: $1 - \text{the survival probability over the first two years}$ —that is:

$$\text{PD}(2) = 1 - (1 - P(1))(1 - P(2))$$

Numerical example:

$$\text{Given } P(1) = 0.2\% \text{ and assuming that } P(2) = 0.3\% \text{ then } \text{PD}(2) = 0.5\%$$

Simple reduced-form models are confronted with the identification problem expressed in formula A11-4. The credit spread is the expected loss—that is, the product of the probability of default and the loss given default. Jarrow and Turnbull (1995) assume a constant LGD and an independent default intensity process with a Poisson distribution that determines the time to default. Default probabilities are nonrandom time-dependent functions. The default and interest rates are not correlated.

Jarrow has generalized the Jarrow-Turnbull model in several directions.² First, default probabilities are assumed to be stochastic with an explicit dependence on stochastic interest rates and an arbitrary number of risk factors. Some of the risk factors are firm specific, while the others are macro-factors driving default correlations, such as commodity prices and real estate prices.

Jarrow also incorporates a liquidity factor that affects the prices of bonds, but not equities. This liquidity factor can be random and different for each issuer, and it can be a function of the same macro-risk factors that determine the default intensity.

2. Jarrow R., "Default Parameter Estimation Using Market Prices," *Financial Analysts Journal* 57(5), 2001.

Jarrow's generalization of the Jarrow-Turnbull model has two key features: it examines both debt and equity prices with the assumption that, in the event of default, the equity value is zero; and it incorporates liquidity risk and illiquid spreads. These features allow the model to resolve the indetermination problem mentioned earlier by permitting the calculation of an implied recovery parameter expressed as a fraction of the market value of the risky debt, an instant prior to bankruptcy—that is, rather than as a fraction of the principal of the bond, as assumed earlier in our illustration. This model is the analytical framework implemented in the Kamakura-KRIS credit portfolio model.

The term structure of hazard rates can be fitted to bond prices (as well as bond prices and equity prices, simultaneously), to credit derivatives prices (CDS spreads), and to historical data on defaults. The reduced-form model can be used to improve portfolio hedging.³

3. H. Doshi, J. Ericsson, K. Jacobs and S. M. Turnbull, "Pricing Credit Default Swaps with Observable Covariate," *Review of Financial Studies*, forthcoming.