Corporate Bond Portfolio Analysis

The CDO approach.

David P. Jacob, Mark Adelson, and Michiko Whetten

David P. Jacob
is a managing director at Nomura Securities International, Inc. in New York City.
djacob@us.nomura.com

Mark Adelson
is a director at Nomura Securities International, Inc., in New York.
madelson@us.nomura.com

Michiko Whetten
is an assistant vice president at Nomura Securities International, Inc., in New York.
mwhetten@us.nomura.com

The usual approach to portfolio management is to use measures of expected return and risk of assets, as well as the covariance of these returns, to create a portfolio with the highest possible expected return for a given level of risk. Very good portfolio managers will add value through their ability to choose undervalued credits.

The problem with using the standard deviation of returns as the measure of risk is that, as a summary measure, it fails to account for the whole distribution. As a result, portfolio managers may not discover their exposure to the risk of extreme events, and therefore take on risk without being compensated for it.

Exhibit 1 illustrates why it is important for a portfolio manager to focus on the whole distribution of losses, and not merely on a point estimate of expected losses. The graph shows the distribution of projected losses on two portfolios whose expected losses are identical. The standard deviations of both distributions are identical. Yet the distributions are noticeably different.

The likelihood of experiencing losses greater than 3% of market value is 30% for Portfolio A but just 24% for Portfolio B. All else equal, a portfolio manager who examines the loss distributions for both portfolios probably would prefer Portfolio B because it is less likely to suffer severe losses.

We propose applying the methodology used to construct and analyze collateralized debt obligations (CDOs) to measuring and managing the credit risk in corporate bond portfolios. A basic CDO invests in a

EXHIBIT 1
Distribution of Losses for Two Portfolios

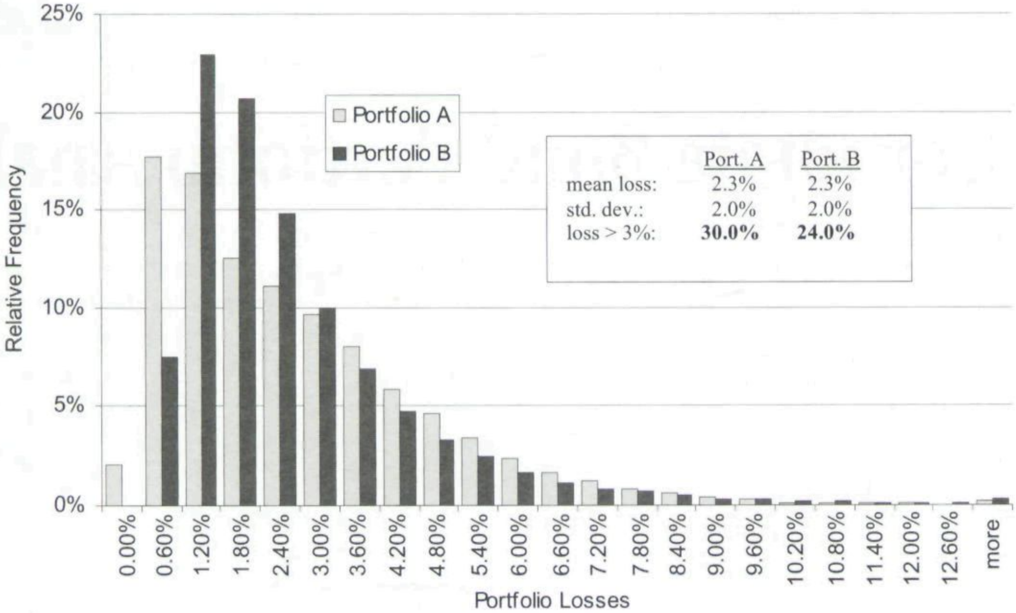

Source: Nomura Securities International.

portfolio of bonds or loans and issues its own bonds. In most cases, a CDO issues multiple classes of bonds with different levels of seniority. Each class protects the ones senior to it from losses on the underlying portfolio. Since different investors buy different tranches of a CDO, analyzing the whole distribution of potential losses in the portfolio, rather than just the expected (or average) losses, becomes essential in assessing the risk profile of each tranche.[1]

Market participants often analyze CDOs using computer simulation techniques that highlight the distribution of outcomes that the underlying portfolio might experience. With that knowledge in hand, the portfolio manager can adjust the portfolio's risk profile. The manager may use credit derivatives, such as credit default swaps (CDS), to increase or reduce exposure at different rating levels, industries, and maturities.

KEY INGREDIENTS OF THE SIMULATION

A simulation-based CDO analysis treats bonds as having mathematical properties that describe their behavior. More specifically, bond defaults are treated as if they are governed by the laws of probability. At first blush, the mathematical abstraction may seem fallacious to a traditional corporate bond portfolio manager. After all, in the real world, bonds default because of fundamental factors such as excessive leverage or a recessionary environment. For large groups of bonds viewed over extended periods, though, the mathematical abstraction produces a fair approximation of the real world. Remember, the CDO analysis described here should be viewed as a supplement to other tools, not a replacement for them.

The heart of the CDO analysis is a Monte Carlo simulation of a portfolio's credit performance over a specific time horizon (e.g., one year). The analysis requires assumptions about certain properties of the underlying bonds: 1) their probabilities of default, 2) recovery rates following defaults, and 3) correlation of default risk among different bonds.[2]

Probability of Default

The simulation process requires that we assign each bond in the portfolio a probability of default. For riskless securities, such as Treasury bonds, the probability of default would be zero. For all other bonds, the probability of default would be between zero and one.

One easy way to assign default probabilities is to use bond credit ratings from the rating agencies. For example, the historical default rate for Baa3-rated bonds arguably provides a reasonable estimate of the probability that such bonds will default in the future.

Moody's and S&P regularly publish default rates for corporate bonds. They show the average default rates for bonds in different rating categories over various time horizons.[3]

Bonds in the highest rating categories display zero defaults over time horizons of just one or two years. Therefore, for purposes of a simulation analysis, it is advisable to use default frequencies measured over longer time horizons and to convert those frequencies into equivalent one-year frequencies. For example, a five-year default rate can easily be translated into an equivalent one-year default rate using the formula:

$$\textit{Equivalent One–Year Default Rate} = 1 - \sqrt[5]{1 - (\textit{Five–Year Default Rate})}$$

Thus, the 43.55% five-year default rate for bonds rated B3 corresponds to an equivalent one-year default rate of 10.80%.

A portfolio manager who has a specific view about whether corporate bond defaults are likely to increase or decline can use higher or lower default probabilities to suit that view. For example, the historical default rates of a specific year could be selected as the basis for assigning default probabilities.

Beyond historical default rates, a portfolio manager has other potential sources for assigning default probabilities to specific bonds. For example, the portfolio manager can use equity prices, credit spreads, or other market data to estimate default probabilities. There are several third-party products, such as CreditEdge™ from Moody's KMV, that provide estimates of default probability based on such approaches.

Recovery Rates

As with default rates, historical data provide a reasonable basis for assumed recovery rates in a simulation. The simplest approach is to use a single recovery rate assumption for all bonds in a portfolio. Because most corporate bonds represent senior unsecured debt or senior subordinated debt, a reasonable starting point is a recovery assumption of 40%. A simulation model becomes only slightly more complex by using different assumed recovery rates for bonds at different levels of seniority (or for bonds at different rating levels).

In the most elaborate simulations, recovery rates can be tied to (i.e., correlated with) default rates, so that recoveries decline when there are many defaults (e.g., during a recession). There is some evidence of an inverse correlation between default rates and recovery rates.[4]

Correlation

Correlation is the trickiest part of a CDO analysis. There are many mathematical approaches to tackling correlation. Some are quite complicated and require a portfolio manager to assign a correlation factor to each separate pairing of bonds in the portfolio (i.e., construct a whole correlation matrix). A somewhat simpler approach is to assign correlation factors based on the specific industry classifications of the bonds' issuers. Fitch uses this approach in its CDO rating analysis.

An even simpler approach, and the one used by S&P, is to assume a single correlation factor between bonds from issuers in the same industry and a somewhat lower factor for the correlation between bonds of issuers in different industries. For example, S&P generally assumes 30% correlation between issuers in the same industry and 0% correlation between issuers in different industries.

The simplest approach is to use a single correlation factor to describe all correlations within a portfolio. This is the approach many trading desks use to manage their risk from trading credit derivatives.[5]

For simulating the credit risk dimension of a corporate bond portfolio, it is reasonable to balance the desire for detailed realism with the benefits of simplicity. We recommend using a framework with two correlation factors: one defining the correlation among bonds of issuers in the same industry, and one defining the correlation among bonds of issuers in different industries.[6]

A fair starting point is 15% interindustry correlation and 35% intraindustry correlation. This provides a somewhat more conservative correlation assumption than the ones Moody's and S&P use in their CDO rating analyses.

SIMULATION MODEL

Once we have the three key ingredients—default probabilities, recovery rates, and correlations—we need one more item before we can really perform a CDO analysis: a simulation model.

The simulation model for a CDO analysis takes the three key assumptions and processes them to generate a distribution of simulated outcomes for the subject portfolio. The simulation process consists of numerous iterations of the main calculations, often 10,000 or more. In each iteration, the computer generates random numbers and uses them, together with the key assumptions, to calculate simulated losses on the subject portfolio. Then, the simulation aggregates the results for all iterations and reports the distribution of outcomes.

The calculation process in each iteration can be either simple or complex. Simplicity offers the advantages of high speed and low cost. Greater complexity offers the potential advantage of greater realism. We generally favor a balance between the two extremes. Many financial professionals can create their own simple simulations using either Microsoft Excel by itself or in combination with specialized simulation add-ins such as @Risk from Palisade Corporation. Complex simulation models may require programming work by specialists, which can have the effect of distancing a portfolio manager from the simulation model.

In Adelson and Whetten [2004b] we describe a simple simulation model. Here is a quick summary of how it works.

The simulation uses the notion of *default time*, which is usually denoted by the Greek letter τ. In essence, the key idea is that all risky issuers *eventually* default, at some time between now and the end of the world. The converse of an issuer's default probability is its survival probability. For example, if an issuer's default probability is 5% annually, its survival probability for one year is 95%. Its survival probability for two years would be 90.25%, which is 95% squared. The issuer's survival probability for τ years would be 95% raised to the power of τ (i.e., 0.95^{τ}).

To simulate the default time for bonds of a particular issuer, generate a uniform random variable, u, and solve the formula for τ:

$$0.95^{\tau} = u$$

If τ is longer than the maturity of a given bond from the related issuer (or longer than the time horizon of the simulation), the bond does not default. If τ is shorter than the bond's maturity, a default occurs.

For a good model, it is not enough to simply simulate the default times of individual bonds in a portfolio. It is necessary further to address the correlation of default risk among the bonds. To do this, we use techniques to link the default times of different issuers. For our purposes, we use a model that allows us to specify individual correlation coefficients between each pair of issuers.

A somewhat simpler way is to use a single-factor approach with just one correlation coefficient (ρ) that describes the correlation among all pairs of bonds in the portfolio.

For example, a single-factor Gaussian copula approach uses four steps to produce correlated default times (τ) for different issuers in a portfolio:[7]

1. First, for a portfolio of n bonds from different issuers, generate $n + 1$ *independent*, normally distributed random variables. Designate one of them y and the rest of them ε_1 through ε_n.
2. Second, combine the correlation coefficient (ρ) with the variables from the previous step to generate n *correlated* and normally distributed random variables (designated x_1 through x_n) by using the formula:

$$x_i = \sqrt{\rho y} + \sqrt{1-\rho}\,\varepsilon_i$$

3. Third, use the cumulative distribution function of the standard normal distribution to convert each of the x_i variables into a corresponding uniformly distributed random variable in the range from 0 to 1, designated u_i.
4. Fourth, use the u_i variables and the default time formula above to calculate τ_i values, where τ_i is the default time of the i-th issuer of the portfolio.

Each iteration of the simulation uses simulated default times to tell whether bonds of each issuer default. That is, the simulation uses τ_i to tell whether bonds of the i-th issuer default. For each bond that defaults, the portfolio loses the principal amount of the bond less the recovery. Summing the losses for all defaulted bonds gives the total portfolio loss for that iteration. Repeating the process thousands of times provides a whole distribution of portfolio losses that a portfolio manager can examine and analyze.

RUNNING THE SIMULATION

Suppose a hypothetical portfolio manager manages a portfolio including corporate bonds from around 200 issuers. Suppose that the average rating of the issuers is around Baa3/BBB- and that the portfolio is slightly overweighted in certain industries. Exhibit 2 shows the industry diversification of the portfolio, and Exhibit 3 shows

EXHIBIT 2
Portfolio Industry Diversification

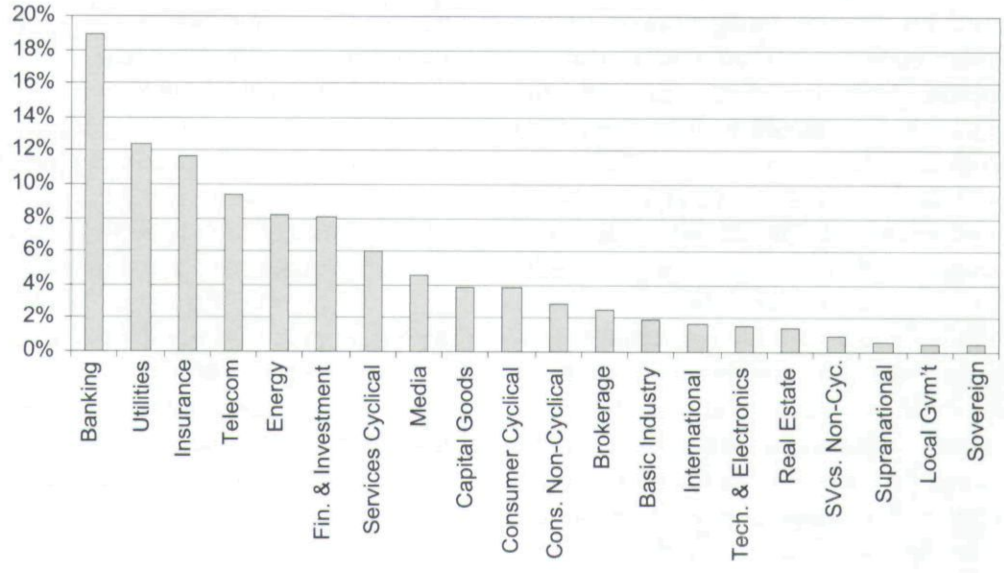

Source: Nomura Securities International.

EXHIBIT 3
Portfolio Credit Quality Distribution

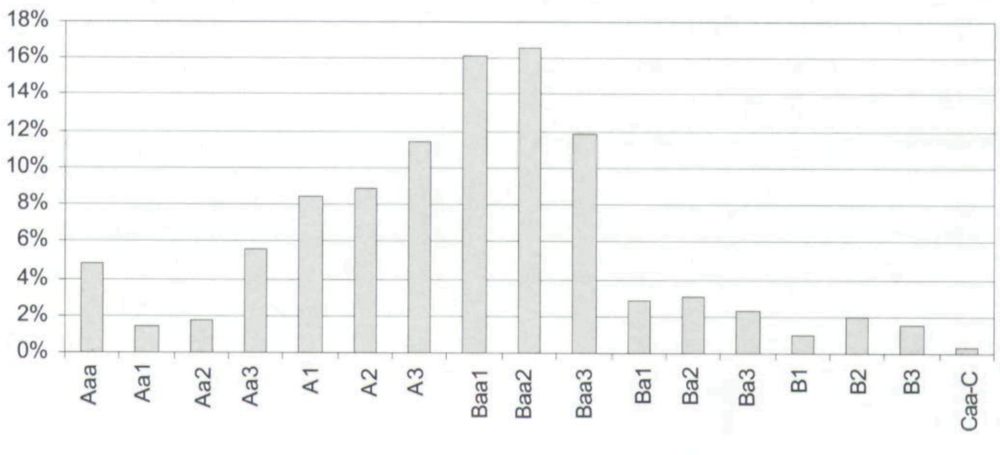

Source: Nomura Securities International.

the rating distribution.

The portfolio manager could run simulations on the portfolio to graph the distribution of simulated losses, expressed as a percentage of market value. Suppose that default probabilities are based on Moody's historical default rates, and interindustry correlation of 15% and intraindustry correlation of 35% are assumed. Suppose further that the manager assumes recovery rates that vary according to rating category—that is, that bonds with higher ratings experience higher recovery rates following default than bonds of lower ratings.

INTERPRETING THE RESULTS

Suppose the portfolio manager simulates the performance of the portfolio over a period of one year and gets the loss distributions shown in Exhibit 4. Note that the loss distribution for the portfolio has a tail extending out to the right. The loss distribution in Exhibit 4 shows there is roughly a 2.1% chance of losing more than 3% of the portfolio's value in one year. The portfolio manager must decide whether that likelihood is within acceptable bounds. If not, there are various strategies for fine-tuning the portfolio to alter its exposure to high-loss scenarios.

In Exhibit 4 there is a slight bump in the loss distribution at the 1.95% mark. That bump reflects the impact of the portfolio's lumpiness—it has bonds from only around 200 issuers. Without performing simulations, the portfolio manager probably could not discern that losses in the 1.95% bucket are somewhat more likely than losses in the 1.80% or 1.65% buckets.[8]

To explore the implications further, suppose the portfolio manager performs simulations over a longer time horizon, such as five years. The results are shown in Exhibit 5. Here, the simulated frequency of suffering losses of more than 3% increases to 30%.

The simulation results shown in Exhibits 4 and 5 treat the portfolio as entirely static. While that might be a reasonable basis for a one-year time horizon, it may be unrealistic for an actively managed portfolio over a longer interval; the portfolio manager might trade out of some positions that would otherwise have subsequently defaulted. On the other hand, for portfolios of illiquid securities, such as private placements, it may be not be possible to trade out of deteriorating credits. In that case, treating the portfolio as static over longer time horizons is entirely reasonable (and prudent).

In addition, the results shown in Exhibits 4 and 5 show losses on a gross basis. That is, they do not show the offsetting effect of the incremental spread that the portfolio earns from investing in risky securities. Indeed, as shown in Exhibit 3, nearly 58% of the portfolio is rated below single-A, and slightly over 13% is rated double-B or lower.

ACTING ON THE RESULTS

If the portfolio manager feels the portfolio has too much exposure to the risk of severe losses (i.e., > 3%), there are several options. One possibility is to sell some of the portfolio's weaker credits. Or, the portfolio manager might use credit derivatives to adjust the distribution of losses. Techniques for tuning performance may include any of the following:

- Buying protection through a credit default swap index (e.g., DJ CDX.NA.IG or DJ CDX.NA.HY) to reduce *macro* risk.

EXHIBIT 4

Simulated Portfolio Loss Distribution—Base Case, One-Year Time Horizon

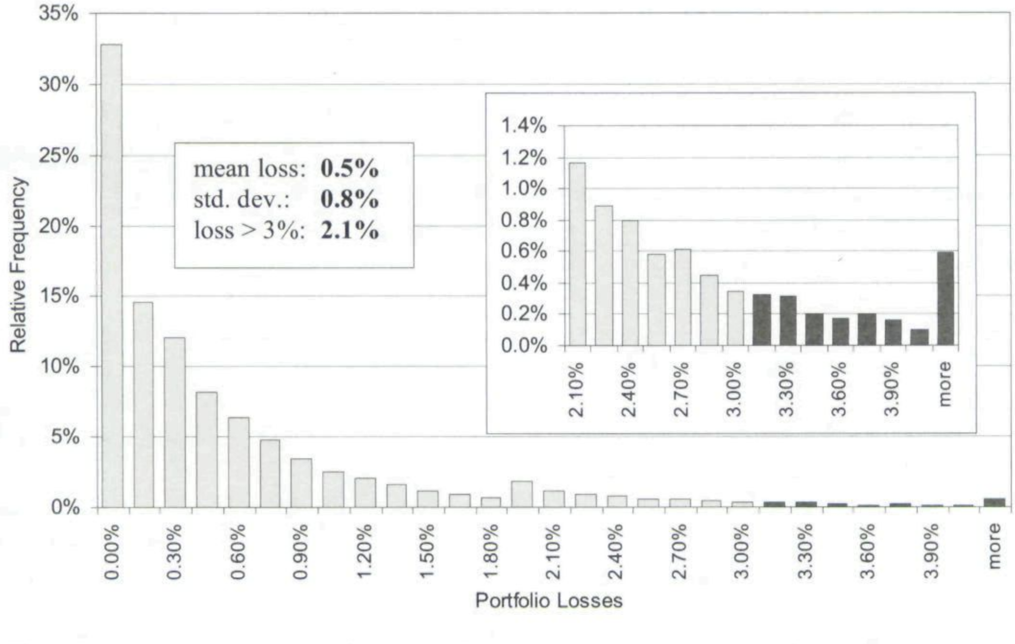

Source: Nomura Securities International.

EXHIBIT 5
Simulated Portfolio Loss Distribution—Five-Year Time Horizon

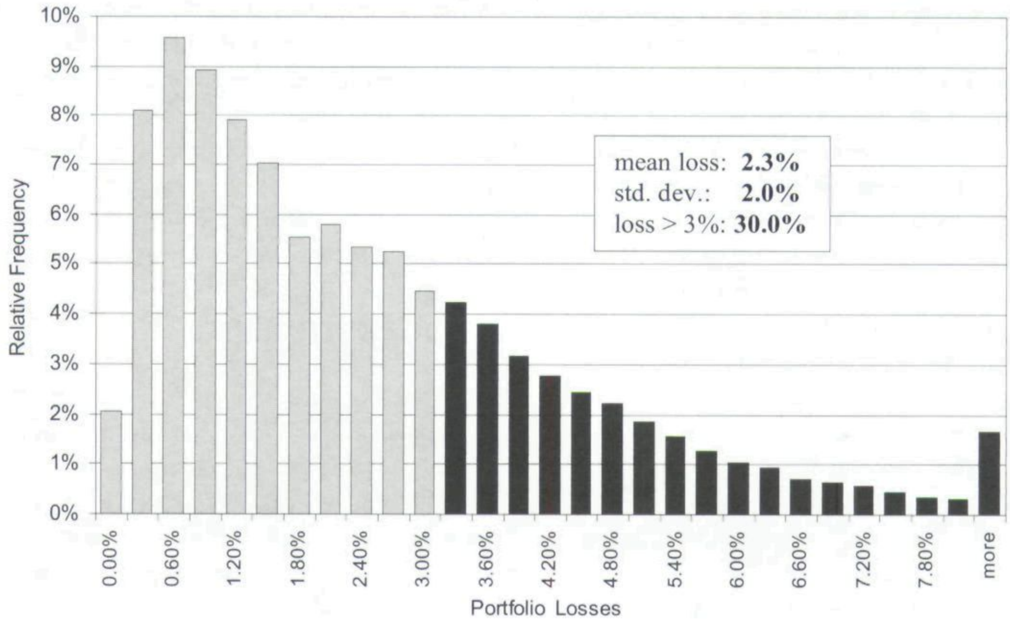

Source: Nomura Securities International.

- Buying protection through a CDS subsector index (e.g., DJ CDX.NA.ENRG or DJ CDX. NA.INDU) to reduce *sector-specific* risk.
- Buying single-name CDS to reduce *name-specific* risk.

The portfolio manager could buy protection for a five-year time horizon using standard contracts or purchase protection for a shorter period with the intention of adjusting the hedge periodically. Another option for the portfolio manager would be to rebalance the portfolio toward stronger credits by *selling* protection on issuers that are not already included in the portfolio.

Suppose the portfolio manager uses CDS to buy protection on the single-B-rated issuers in the portfolio. Exhibit 6 shows the distribution of simulated losses on the portfolio after incorporating the protection on the single-B-rated issuers. A comparison of Exhibit 6 with Exhibit 4 indicates that the tail of the distribution is *much* thinner. In fact, by hedging the single-B exposures, the simulated frequency of suffering losses of more than 3% declines to just 0.4%, compared to 2.1% in the base case (Exhibit 4). In addition, the mean simulated loss on the portfolio drops 23 basis points with the hedge.

Suppose that the cost of hedging the single-B credits is about 330 basis points on the single-B portion of the portfolio (on a running basis). The single-B credits represent roughly 4.5% of the whole portfolio. Therefore, on an overall basis, the cost of hedging the single-B credits translates into 15 basis points on the total portfolio. Thus, the benefits of the hedge—a 23 basis point reduction in mean losses and a thinner tail on the whole distribution of expected losses—appear to fully justify the cost of 15 basis points.

Alternatively, the portfolio manager might consider hedging the double-B credits in the portfolio. Exhibit 7 shows the distribution of simulated losses on the portfolio after incorporating the protection on the double-B-rated issuers. A comparison of Exhibit 7 with the base case in Exhibit 4 shows that the tail of the distribution is somewhat thinner; the simulated frequency with which losses exceed 3% is 1.4%, compared to 2.1% in the base case. In addition, the mean simulated loss on the portfolio drops 11 basis points with the hedge, compared to the base case.

Suppose that the cost of hedging the double-B credits is about 170 basis points on the double-B portion of the portfolio (on a running basis). The single-B credits represent roughly 8.2% of the whole portfolio. Therefore, on an overall basis, the cost of hedging the double-B credits translates into 14 basis points on the total portfolio. Thus,

EXHIBIT 6

Simulated Portfolio Loss Distribution—Hedging Single-B Exposures: One-Year Time Horizon

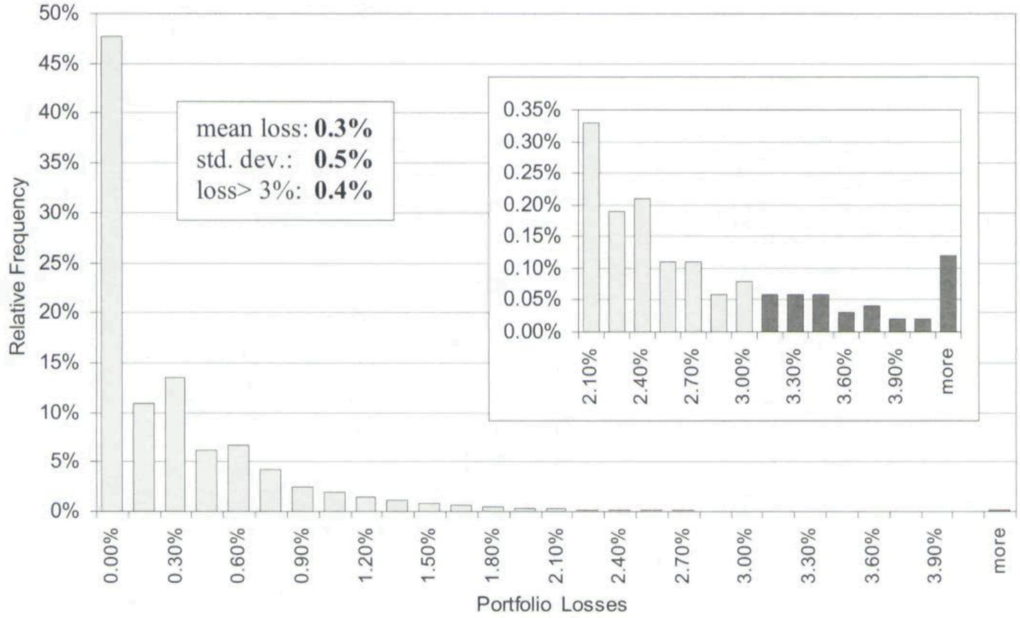

Source: Nomura Securities International.

the benefits of the double-B hedge—11 basis point reduction in mean losses and a slightly thinner tail—may not justify the 14 basis point cost of the hedge. Hedging the single-B portion of the portfolio would be a better strategy.

REALITY CHECK

A key aspect of any simulation-based analysis is sensitivity testing. A cautious portfolio manager generally would consider how the simulation results would be affected by changing the initial assumptions of the model. Doing so is an essential step toward establishing confidence in the model.

For example, the portfolio manager might experiment by testing higher levels of correlation among issuers. Exhibit 8 shows the effect of assuming interindustry correlation of 25% and intraindustry correlation of 45%:

In Exhibit 8, there is a 3.2% frequency of losses exceeding 3% of the portfolio's value, rather higher than the 2.1% frequency of the base case assumptions in Exhibit 4. Or, the portfolio manager might test a simulation that increases all default probabilities by one percentage point. This result is shown in Exhibit 9.

Increasing all the default probabilities has a very pronounced impact. Now losses exceed 3% up to roughly 7.9% of the time.

By experimenting with different assumptions, the portfolio manager learns the model's sensitivities while simultaneously gaining insight into the risk dimension of the portfolio.

CONCLUSION

The CDO analysis is a supplement to a portfolio manager's other tools. It does not replace either fundamental bond analysis or other types of portfolio optimization (efficient frontier) analyses. Nor does it replace experience, judgment, and insight as key ingredients for success in managing a corporate bond portfolio.

The CDO analysis enables a corporate bond portfolio manager to explore and analyze the whole distribution of potential future performance. The analysis ascribes mathematical properties to bonds and then uses computer simulations to illuminate the implications.

The key benefit of the CDO analysis is that it provides a whole distribution of future results, rather than just the most likely scenario. In addition, it offers the benefits of internal consistency, repeatability, ease of use, and results that can provide the basis for concrete action.

It is a useful and powerful instrument for the financial toolkit.

EXHIBIT 7
Simulated Portfolio Loss Distribution—Hedging Double-B Exposures: One-Year Time Horizon

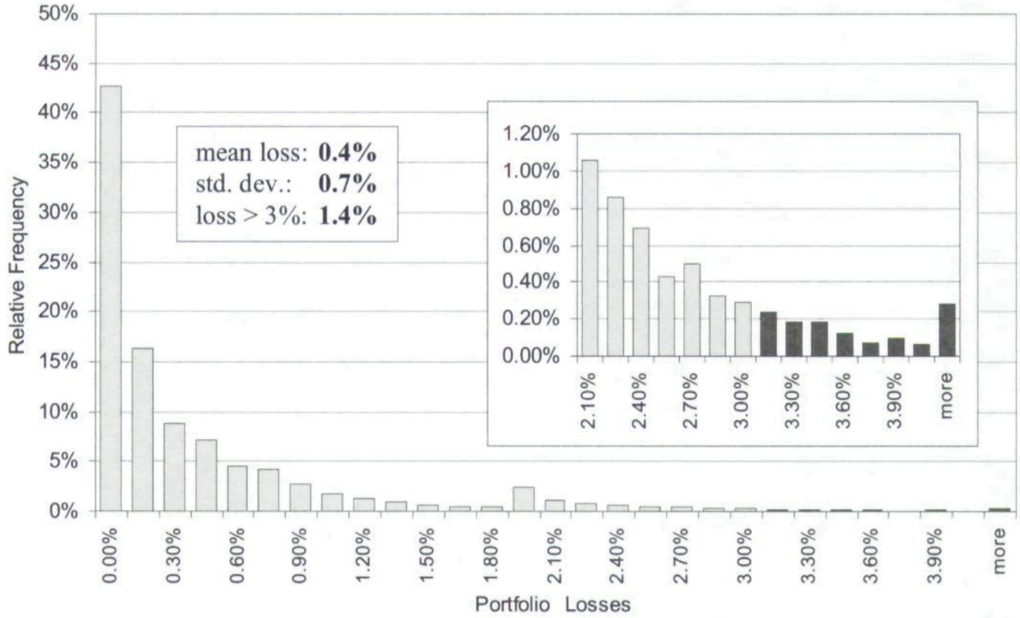

Source: Nomura Securities International.

EXHIBIT 8
Simulated Portfolio Loss Distribution—Higher Correlations: One-Year Time Horizon

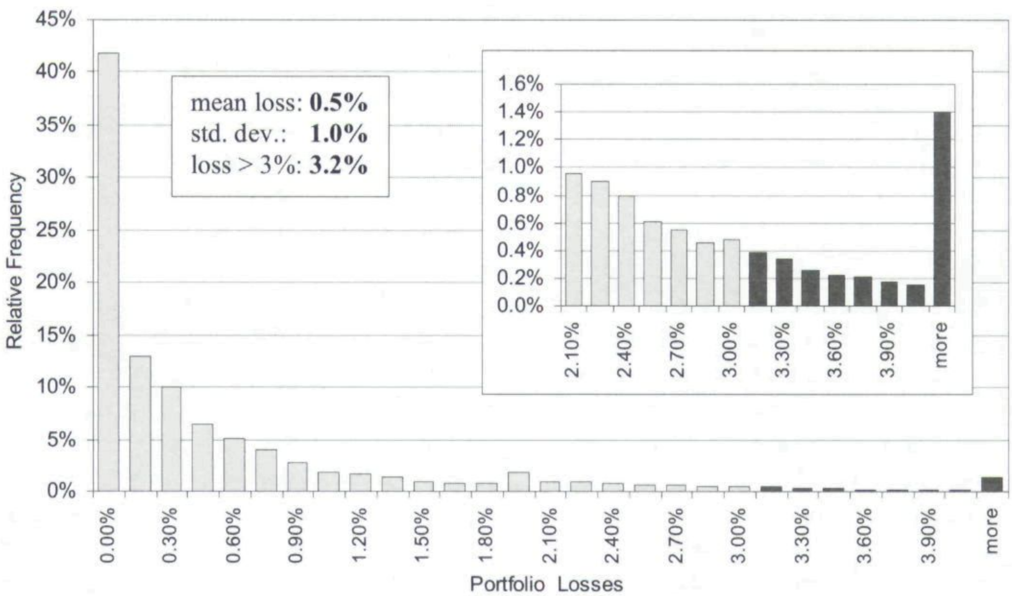

Source: Nomura Securities International.

EXHIBIT 9

Simulated Portfolio Loss Distribution—Higher Default Probabilities: One-Year Time Horizon

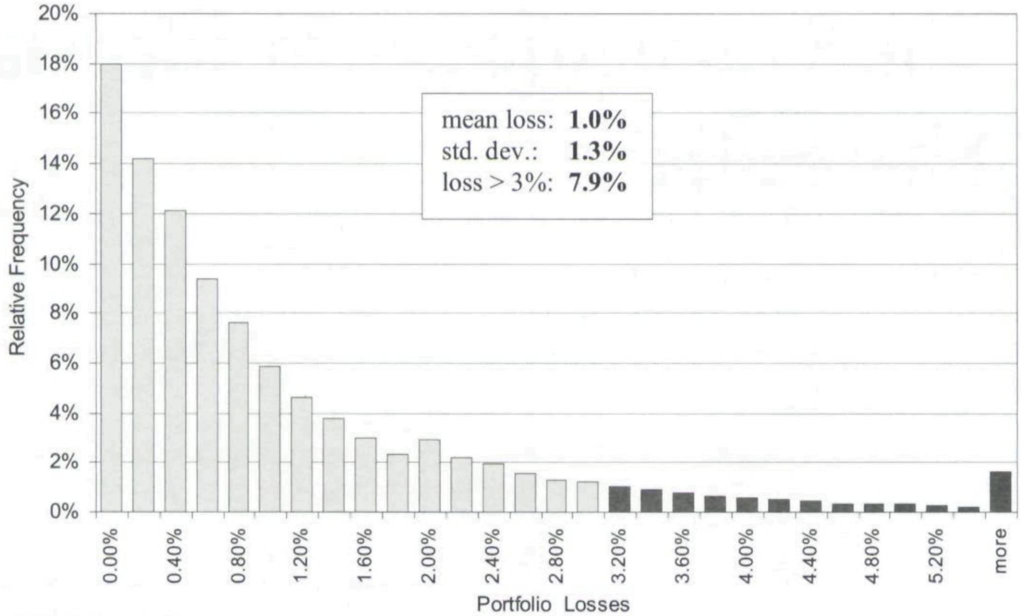

Source: Nomura Securities International.

ENDNOTES

[1]For a basic introduction to CDOs, see Adelson and Whetten [2004a].

[2]A Monte Carlo simulation solves a problem by generating random numbers as the inputs to a modeled process and then observing the distribution of results over many trials. The technique is most useful for obtaining numerical solutions to problems that are too complicated to solve analytically. See Weisstein [1999].

[3]The rating agencies periodically publish historical default rates and recovery. See, for example, Cantor et al. [2005].

[4]See, for example, Aurora, Schneck, and Vazza [2005].

[5]See "Global Cash Flow and Synthetic CDO Criteria" [2002].

[6]See Adelson and Whetten [2005b].

[7]Li [2000] originally proposed application of the copula approach in credit risk modeling. See also Schönbucher [2003].

[8]The 1.95% mark represents how often simulated losses are *greater than* 1.80% and *less than or equal* to 1.95%.

REFERENCES

Adelson, M., and M. Whetten. "CDO in Plain English." Nomura Fixed Income Research, September 2004a.

——. "Correlation Primer." Nomura Fixed Income Research, August 2004b.

Aurora, D., R. Schneck, and D. Vazza. "Annual Global Corporate Default Study: Corporate Defaults Poised to Rise in 2005." Standard & Poor's special report, January 2005.

Cantor, R., D. Hamilton, S. Ou, and P. Varma. "Default and Recovery Rates of Corporate Bond Issuers, 1920-2004." Moody's special report, January 2005.

"Global Cash Flow and Synthetic CDO Criteria." Standard & Poor's special report, March 2002.

Li, D.X. "On Default Correlation: A Copula Function Approach." Working paper, Risk Metrics Group, April 2000.

Schönbucher, P. *Credit Derivatives Pricing Models: Models, Pricing, and Implementation.* New York: John Wiley & Sons, 2003.

Weisstein, E. *CRC Concise Encyclopedia of Mathematics.* London: Chapman & Hall/CRC, 1999.

To order reprints of this article, please contact Dewey Palmieri at dpalmieri@iijournals.com or 212-224-3675.