

SIMULATION METHODS FOR RISK ANALYSIS OF COLLATERALIZED DEBT OBLIGATIONS

William J. Morokoff

New Product Research
Moody's KMV
1620 Montgomery Street
San Francisco, CA 94111, U.S.A.

ABSTRACT

Collateralized Debt Obligations (CDOs) are sophisticated financial products that offer a range of investments, known as tranches, at varying risk levels backed by a collateral pool typically consisting of corporate debt (bonds, loans, default swaps, etc.). The analysis of the risk-return properties of CDO tranches is complicated by the highly non-linear and time dependent relationship between the cash flows to the tranche and the underlying collateral performance. This paper describes a multiple time step simulation approach that tracks cash flows over the life of a CDO deal to determine the risk characteristics of CDO tranches.

1 INTRODUCTION

The term Collateralized Debt Obligation (CDO) covers a broad range of structured finance products. They may be supported by a variety of underlying collateral, from bonds, loans, credit default swaps, asset-backed securities, and sovereign debt, to more exotic securities such as equity default swaps or CDO tranches from other deals. The CDO structures, which describe the size and number of tranches and the rules for how to distribute the collateral proceeds to the tranches, also vary widely. The structures may be simple pass-throughs, whereby interest payments are made in order of tranche seniority. However, there may also be complicated rules to redirect cash flows to the senior tranches based on the quality and performance of the underlying collateral. Each deal has a unique structure determined by market conditions, collateral properties and investor demand, among other factors, at the time of issuance. Many features of a structure are intended to insure that the most senior tranche will be rated at the highest credit quality level by the debt rating agencies (usually Moody's, S&P, and Fitch). See Goodman and Fabozzi (2002) for a detailed discussion of the CDO market; Duffie and Singleton (2003) includes a chapter on credit modeling

methods applied to CDOs. A further discussion of credit modeling can be found in Arvanitis and Gregory (2001).

From a legal perspective, a CDO deal is generally set up as a Special Purpose Entity (SPE) that functions as an independent company, often incorporated in Bermuda. The capital structure of this company is very simple: the assets owned by the SPE are the collateral (e.g., 100 corporate bonds), while the liabilities are the tranches issued by the SPE. Investors purchase the tranches, and the SPE uses the proceeds from the sale of the tranches to purchase the collateral assets. Periodically (typically quarterly or semi-annually), the interest and principal cash flows generated by the collateral assets over the period are collected together into accounts that are then used to make interest and principal payments to the tranches. The set of rules for how the funds are distributed at a given payment date is known as the cash flow waterfall for the CDO.

In a typical waterfall, taxes and management fees are paid first, followed by the interest due to the senior tranche. This senior tranche generally accounts for the largest percentage of invested principal (75% - 90%) but gets paid the smallest coupon. This is consistent with the senior tranche being the least risky due to its position of getting paid first. The tranches are paid in order of seniority with coupons and risk increasing as payments move down the structure. The most junior tranche at the bottom of the waterfall is known as the equity tranche. The equity of a CDO deal usually does not receive a predetermined interest payment on its initial investment; instead the equity receives all the remaining collateral interest payments that were not required to make the interest payments on the more senior tranches. In deals with substantial "excess spread" (the amount of interest generated by the collateral portfolio beyond what is due to the CDO tranches), the equity tranche performs well. As defaults occur in the collateral portfolio, the amount of excess spread decreases and the equity tranche suffers the first losses. Many deals have collateral quality triggers that divert all cash flows away from the equity if the collateral quality deteriorates too much.

The collateral manager of a CDO deal is responsible for managing the collateral assets as their credit quality changes. This involves buying and selling assets, as well as reinvesting funds that have been recovered from defaulting or maturing names. The manager's skill and strategy around managing the collateral portfolio can have a large effect on the performance of the CDO deal.

From a risk management perspective, the most important factor affecting the performance of a CDO deal is the total loss in collateral portfolio value over the life of the deal due to correlated defaults among the collateral. Each tranche can withstand a characteristic level of loss on the collateral pool before it does not receive its promised interest and principal payments. Performance is also greatly affected by the timing of the defaults, particularly for the equity tranche. Other risk factors are interest rates, maturing and prepayment rates of collateral, and recovery rates on defaulted collateral. There is also price risk, i.e. the change in value of the collateral due to changes in credit quality or interest rates, and the related issue of reinvestment risk.

2 MODELING DEFAULT PROBABILITIES

A key component of accurately capturing the risks associated with correlated collateral defaults is the determination of current default probabilities over the life of the deal. There are many methods currently in use, and frequently components of various approaches are used together. Broadly, there are four distinct approaches. The first is qualitative analyst review of a company's financials, management, business plan, etc. to determine credit worthiness. This approach is used internally by banks making decisions to lend as well as by rating agencies, which provide the market with an independent qualitative review. The second approach is statistical, in which multi-dimensional regressions are carried out on large data sets of company financial information to determine default indicators and associated probabilities of default. This is particularly effective when market data on prices of the corporate debt or equity are not available, as is often the case with private firms; see Falkenstein (2000) for a discussion of the implementation of this method. A third method uses an approach first put forth by Merton (1974) that derives from the idea that the equity of a firm is actually a kind of call option on the underlying asset value of the firm, and that the firm is in default when that asset value drops below liabilities owed by the firm. Merton's model takes as input the observable stock price for publicly traded firms to back out the unobservable firm asset value, which when combined with firm liability information, leads to default probabilities for public companies. Finally, default probabilities may be inferred from prices of corporate debt (bonds, loans, default swaps) since price is strongly influenced by the market's perception of default probability. However, accurate pricing information

is only available on a relatively small number of names. Also, there are numerous factors other than default probability that determine the price of debt, and these factors, such as liquidity, may be difficult to quantify. Prices are often used to calibrate parameters for stochastic models for the evolution of default intensities.

The models used at Moody's KMV are based on the largest database of corporate defaults in the world, which is important for accurate parameter and model estimation. For publicly traded companies, MKMV uses the Vasicek-Kealhofer model, described in Kealhofer (2003), which is an extension of the Merton approach that incorporates extensive asset volatility modeling and historical default data, to produce the Expect Default Frequency™, or EDF™, credit measure. Most of the largest 100 financial institutions in the world use the EDF credit measure to monitor the credit quality of their loans and investments. The Vasicek-Kealhofer model is an example of a structural model, named so because there are explicit economic drivers that are quantified in the model and that determine the default probability of the company. These drivers are firm asset value, firm asset volatility and firm liability structure.

There are two important further benefits of using this structural model. First, based on the model, MKMV has produced weekly time series of asset returns and asset volatilities for over 28,000 publicly traded companies world-wide. From these time series, a factor model for asset returns has been derived that provides the most accurate means available for describing the correlated behavior of corporate asset values. When combined with default probabilities, the correlation model of asset returns provides an estimate of joint default probabilities necessary to describe the correlated default behavior of a portfolio of CDO collateral.

The second important consequence of the structural model is the concept of the Distance to Default (DD). This is essentially the number of standard deviations that the asset value is above the default level (a function of the liability structure); it is a scalar measure that captures the key relationship among the three structural drivers (Distance to Default also depends on the time horizon of interest; for the purposes here we will consider one year DD). It has been shown empirically to be a stable measure of credit quality over time and economic conditions (credit cycles) as well as over geographic regions. MKMV has an extensive database of time series of Distance to Default for public companies. From this database, empirical transition probability distributions have been determined that provide the probability of a firm with DD_0 at time T_0 migrating to DD_1 at time T_1 . These empirical transition densities provide a much more realistic description of credit quality migration than the standard approach of assuming geometric Brownian motion for the asset value process. Although this model for asset value is reasonable, in order to adequately capture changes in DD (and the corresponding EDF), it would be necessary to also model a correlated process for

how firms change their liability structure. From the data it is clear this is not a continuous stochastic process; while a jump process would be a possible model, the calibration of the parameters would be difficult because the decision to add liabilities is driven by the firm's management style and the economic opportunities available. The empirical DD distributions capture these effects as well as changes in asset volatility without requiring explicit modeling.

The factor model for asset returns and the empirical DD distributions, both derived from the structural model for public firms, can also be used with the statistical models for private firms or any other default probability model, as long as an estimate for the R-squared for the firm in question's asset return regressed against the factors (i.e., the percentage of variance of the asset return explained by the factors) is available, together with industry and country information about the firm.

3 MODELING DEFAULT TIMES

This section describes the methodology most commonly employed today for simulating correlated defaults. It is known as the default time or copula approach and is described by Li (2000) and Schmidt & Ward (2002).

For an exposure in a CDO collateral pool, the default probability to maturity (either of the CDO deal or the exposure, whichever is sooner) gives the probability of that name defaulting at some point during the life of a CDO deal. The timing of the default, however, can also play an crucial role in determining the performance of the deal. Default timing is determined from a default probability "term structure" which may be represented as a vector of cumulative default probabilities

$$(CEDF_1, CEDF_2, \dots, CEDF_N)$$

specified at times

$$(T_1, T_2, \dots, T_N).$$

The quantity $CEDF_i$ is interpreted to mean the probability of default in the interval $(0, T_i)$. Thus the $CEDF$ are increasing. This may be generalized to a time continuous default probability function $CEDF(t)$; however, default probabilities are usually reported at discrete times, and a continuous function is obtained from interpolation.

The default time/copula method of randomly sampling default times works as follows. The first step is to randomly sample a uniform $(0,1)$ variate u . Assuming that T_N is the maturity, if $u > CEDF_N$ then the exposure does not default. If $CEDF_{i-1} < u \leq CEDF_i$ then the exposure defaults in period i . This procedure is closely related to

sampling a stopping time for a random process crossing a default boundary.

A key feature of this approach is the process for determining correlated default times. This requires sampling a set of correlated uniform variates (u_1, \dots, u_M) , where M is the number of exposures in the portfolio. This is done by specifying a copula function $C(u_1, \dots, u_M)$, which is a probability distribution function defined on the M -dimensional unit cube. The copula function is often related to the asset return distribution function at time T_N , $F(R_1, \dots, R_M)$, by the formula

$$C(u_1, \dots, u_M) = F(F_1^{-1}(u_1), \dots, F_M^{-1}(u_M))$$

where $F_j^{-1}(\cdot)$ is the inverse of the marginal probability distribution for the j^{th} exposure. However, any copula function may be used for this purpose. The most commonly used are Gaussian and T-copulas, although a variety of other methods, including Archimedean copulas, have been considered.

For the Gaussian copula, the sampling procedure is particularly simple. Based on the correlation matrix for the asset returns, a correlated sample of standard Normal variates $(\varepsilon_1, \dots, \varepsilon_M)$ is sampled, either from a Cholesky decomposition of the correlation matrix or from a factor model decomposition. The uniform variates are then obtained from the formula

$$u_j = \Phi(\varepsilon_j).$$

Here Φ is the one dimensional standard cumulative Normal distribution function.

If the factor model underlying the correlation structure has more than a few dimensions, it is necessary to use Monte Carlo simulation to sample correlated defaults and default times that are then used to evaluate expectation integrals such as the probability of having more than k defaults or the expected value of the cash flows to a tranche. Under more restrictive assumptions on the correlation structure, semi-analytical solutions can be derived. For example, the latent variable approach, proposed by Vasicek (1987) for credit portfolio risk problems, has been extended to CDOs by Gregory and Laurent (2003). The idea is that there exists a low dimensional underlying latent variable x such that conditional on x the default probabilities and times for the exposures are independent. The law of conditional expectations then allows the portfolio properties of interest to be expressed as an expectation over x of the portfolio properties of an independent portfolio. Often x is taken to be one dimensional, so the problem reduces to a one dimensional quadrature.

4 MULTI-STEP SIMULATION

An alternative to the default time approach based on simulating the firm asset value as a stochastic random variable has been described by Hull and White (2001), Arvanitis and Gregory (2001) and Finger (2000). In this section we describe an implementation of this approach and describe a new multi-step approach based on the empirically derived Distance to Default distributions.

While the default time approach captures the marginal default probabilities of each individual exposure correctly over the life of the simulation, substantial error may be introduced into the correlated default structure, depending on how the correlation structure and the underlying stochastic default process are viewed. Time series of asset, equity or debt price returns are usually based on daily or weekly time intervals. Given the relatively high default probability of most assets over time horizons of five years or longer, using a correlation structure based on weekly returns as a proxy for multi-year horizon correlations can lead to skewed results. In particular, the single step approach may not adequately capture the absorbing nature of the default state (i.e., the stochastic process has an absorbing boundary). Thus it is better to consider a simulation based on a sequence of shorter time steps that one single step to maturity.

It is possible to model the credit migration of a single asset as a continuous time stochastic process, such as geometric Brownian motion or an Ornstein-Uhlenbeck process, with an absorbing boundary implied by the cumulative default probability function $CEDF(t)$. In this formulation a free boundary problem PDE can be derived as described by Avellaneda and Zhu (2001). However, the existence of $CEDF(t)$ as a time continuous function usually arises from imposed model or interpolation assumptions; there is generally not enough market data or financial information available to imply forward default probabilities over short time windows. Thus the continuous approach does not add accuracy relative to a discrete approach as long as the correlated behavior of asset over the time step is consistent with the correlation modeling. In any case, unless a low-dimensional latent variable approach is applied, computation of the properties of a portfolio of many exposures will require a Monte Carlo simulation based on discrete time steps.

For analyzing a single CDO deal, it is most convenient to use simulation time steps based on the CDO payment dates. For one simulation step, the names defaulting during that period are identified, recoveries on defaulted names are determined, interest cash flows from non-defaulted collateral are aggregated, scheduled and unscheduled principal payments from the collateral are collected, etc. The resulting pools of interest and principal cash flows are then passed to the cash flow waterfall engine to be distributed to the CDO tranches. If desired, the

exact default time of an exposure can be sampled using the default time methodology described above within one simulation period; in practice, however, the default on a particular exposure will occur on a coupon date, not at a random time. The key question for the simulation is thus whether the default occurs in a given period.

There are numerous approaches that can lead to multi-step simulations for correlated defaults depending on how the default process is modeled. We focus here on two methods related to structural models for which correlated default behavior is derived from the underlying firm asset value correlations. Both methods take as input the cumulative default function $CEDT_j(t)$ specified at discrete times (T_1, \dots, T_N) for each obligor in the collateral portfolio, indexed by j . In addition, the firm asset value correlation matrix for all obligors must be specified.

The first approach assumes that the asset value process for each obligor follows correlated geometric Brownian motion. The associated asset value (log) return process therefore follows a standard Brownian motion process. An obligor j defaults during a period $(T_{i-1}, T_i]$ if the asset return R_j^i at time T_i is less than some threshold level α_j^i , while $R_j^k > \alpha_j^k$ for all $k < i$ (i.e., there was no previous default). In a continuous time formulation, the function $\alpha_j(t)$ is the default boundary such that the default time is the stopping time of the Brownian motion process associated with crossing the boundary. Obviously the default thresholds must be related to the default probability. Specifically the relationship is

$$1 - P(R_j^1 > \alpha_j^1, \dots, R_j^i > \alpha_j^i) = CEDF_j(T_i).$$

As this equation suggests, the determination of the default thresholds requires a non trivial calculation as it relates to inverting an i -variate cumulative Normal distribution (in the continuous case, the default boundary is the solution to a free boundary PDE). One approach that gets around the need to invert a multi-dimensional distribution is to determine the distribution of R_j^{i-1} , conditional on no defaults up to time T_{i-1} . Assuming we know this distribution and using the fact that

$$R_j^i = R_j^{i-1} + \phi_j^i$$

where ϕ_j^i is an increment independent of R_j^{i-1} (since the return process is Brownian motion) with a Normal distribution, we can obtain by convolution the distribution of R_j^i , conditional on no defaults up to T_{i-1} , from the conditional

distribution for R_j^{i-1} and φ_j^i . We can then solve for the default threshold α_j^i from the equation

$$P(R_j^i \leq \alpha_j^i \mid \text{no defaults up to } T_{i-1}) \cdot (1 - CEDF(T_{i-1})) \\ = CEDF(T_i) - CEDF(T_{i-1}).$$

Once α_j^i has been determined, the distribution of R_j^i conditional on no defaults up to T_i can be determined by truncating the distribution of R_j^i conditional on no defaults up to time T_{i-1} . By repeated application of this procedure, the entire set of default thresholds can be determined. The main computational cost is associated with the convolution. This can be handled easily with the fast Fourier transform algorithm, which is effective since the conditional distribution is always convolved with a Normal distribution.

Once the default thresholds are determined, the simulation proceeds by sampling correlated Brownian motion paths for the asset returns at the specified times. Default occurs for a given obligor during the first period for which its return falls below the associated threshold. For names that don't default, conditional default probabilities at each time step can be used as input in valuation algorithms to provide consistent, correlated mark-to-model pricing for the collateral.

As mentioned above, the assumption of geometric Brownian motion for the asset value process often does not adequately capture how a firm's credit quality changes over time because it does not take into account the associated changes in liability structure. It is known that as firms do well (e.g. as the asset value of the firm increases), they tend to take on more debt, thereby keeping their credit quality more stable over time. For example, a Baa rated firm will tend to maintain that rating by borrowing more when opportunities arise. It would be unusual for such a firm to grow without adding leverage to become a Aaa rated. However, this tends to be the consequence of the geometric Brownian motion model: over longer time horizons, firms that do not default undergo systematic improvement in their credit quality.

To capture the effects of changes to both asset value and liability structure on credit quality in long horizon multi-step simulations, at MKMV we have developed a multi-step simulation based on the Distance to Default transition densities. We now consider the implementation of this second, empirically-based method.

A key point to consider when working with historically observed data is the need to bucket the data in order to build a suitable sample size. For example, the first step in determining the probability of transitioning from a DD value of 3 over a one year horizon to a DD value of 4 is to identify all names in the historical sample that have at some time point a DD value of 3. However, since DD is determined as a con-

tinuous variable, it is unlikely that any of the sample will have a DD value of exactly 3. Thus it is necessary to repose the question as to the probability of transition from a bucket, or interval, containing the DD value 3 to a DD value less than 4. The distribution of arrival DD 's after one year does not necessarily have to be bucketed – a parametric distribution for the cumulative transition probability distribution can be selected and the actual data used to estimate the distribution's parameters. However, for use in a multi-step simulation, it is convenient to work with the transition probabilities from one bucket to another bucket in the form of a transition matrix. The multi-step simulation is then carried out as a discrete Markov chain by repeated application of the transition matrix to an initial state vector. The size of the transition matrix, which is determined by the size of the DD buckets, is chosen to balance the desire for high resolution in DD space with the need to minimize the statistical errors arising from small sample sizes. Ultimately this is a question of the size of the original data set. The MKMV simulation is based on 9 years of monthly data on over 12000 firms.

There are a number of important observations to be made about the DD transition matrix. First, the default state, conveniently labeled as $DD = 0$, is an absorbing state. The total probability of transitioning to this default state over a given time period is the forward EDF. This EDF is different for each firm; however, the transition matrix was determined by pooling data on many firms. Thus the transition matrix must be viewed as firm aggregate behavior. In order to capture the firm-specific behavior dictated by the input EDF term structure for each firm, it is necessary to make a firm-specific calibration of the transition matrix. The calibration consists of satisfying the constraint that over a given time period, the probability of transitioning from a non-default state to the default state must be the unconditional (or more precisely, conditional only on data specified at T_0) forward default probability:

$$FWD\ EDF(T_{i-1}, T_i) = \frac{CEDF(T_i) - CEDF(T_{i-1})}{1 - CEDF(T_{i-1})}.$$

There are numerous ways this constraint could be enforced. One simple approach is to rescale all the original, firm aggregate transition probabilities to default by a single factor such that their sum, weighted by the unconditional probabilities of being in each non-default state at time T_{i-1} , matches the forward EDF. Once the transition probabilities are adjusted by this scaling, the unconditional probabilities for each state at time T_i can be determined, thereby allowing the calibration for the next time step. This is equivalent to the convolution and truncation steps employed for the geometric Brownian motion model.

A second consideration for the transition matrix is whether the underlying data supports the model of a

Markov process. Not surprisingly, the firm-aggregate transition matrices for time horizons of 6 months, 1 year, 2 years, 5 years, etc., derived from the data do not fit perfectly in a Markov framework. In other words, the one year matrix is not exactly the convolution of the 6 month matrix with itself; nor is the five year transition matrix exactly the five-fold convolution of the one year transition matrix. The agreement of these transition matrices is however sufficient, particularly given the complexity of the underlying factors which drive credit migration of firms as well as the firm-aggregate nature of the transitions themselves, to warrant the approximation by a single, Markov transition matrix, which is determined by optimally fitting, in a least-squares sense, one matrix (and its convolutions) to the empirical transition matrices. This avoids the exceptionally difficult task of specifying and calibrating a non-Markov process for the credit migration.

Once the transition matrix is specified for each obligor at each time step, the simulation proceeds by sampling from $F_i(DD|DD_{i-1})$, the probability distribution of DD states at time T_i determined from the appropriate probability distribution (as given by the transition matrix) conditional on the DD state at time T_{i-1} . By interpolation from the cumulative probabilities for the discrete transition matrix DD states, $F_i(DD|DD_{i-1})$ can be assumed to be a continuous, non-decreasing function with inverse $F_i^{-1}(u)$ defined on the unit interval $[0,1]$. For values of u in the interval $[0, P(DD_{i-1} \rightarrow 0)]$ (i.e., between 0 and the conditional probability of defaulting), it follows that $F_i^{-1}(u) = 0$. We introduce correlations among obligors by assuming multi-variate Brownian motion for the asset return process and sampling the correlated asset return increments according to the specified asset return correlation matrix. The cumulative Normal distribution function is then used to map the sampled asset return increments to the unit interval; this value is then used as the argument for $F_i^{-1}(u)$. More precisely, the DD sample for obligor j at time i is given by

$$DD = F_i^{-1}(\Phi(\varepsilon_j))$$

where the ε_j are the normalized, correlated Normal samples of asset returns.

If the random asset return sample falls below the default threshold (determined by the DD state at the previous time step and the original EDF term structure), the default state of $DD = 0$ is sampled. In this case, a random recovery may be drawn from an appropriate distribution of recovery rates. If the obligor does not default, the sampled DD state at T_i can be used to determine a conditional EDF

term structure looking forward that can be used to discount future cash flows according to their credit risk in order to obtain a price for the exposure at time T_i . (Note that a discussion of the modeling of a stochastic interest rate process, important for determine both price and cash flow characteristics of debt instruments, has not been included here).

5 COMPARISON OF MODELS

In this section we use simulation experiments to compare the two model choices discussed above. The first concerns the differences between the single step default time model and the multi-step simulation, while the second concerns the difference between modeling the asset value process as geometric Brownian motion and modeling credit migration through the Distance to Default empirical transition distributions.

For the default time/multi-step comparison, we consider a portfolio of 120 high yield bonds of maturity greater than five years issued by 120 correlated firms. A histogram of the cumulative five year EDFs are shown in Figure 1. This shows the portfolio to have a substantial component of distressed names, with a portfolio mean 5 year EDF of around 18.7% (corresponding to annualized default rate of 4%). The expected number of defaults, computed as the average EDF times the number of bonds, is 22.4. However, based on the multi-step simulation there is a 5.4% chance of having more than 41 defaults. Figure 2 plots the cumulative probability distribution for the total number of defaults over the five year period for each method. This shows that the default time model overestimates the probabilities of the extreme events (very few defaults, say under 10, or very many defaults, say over 40) relative to the multi-step model. For example, the default time model puts the probability of having more than 60 defaults at around ten times greater than the multi-step model (2.6% versus 0.26%). This difference may be put into the CDO context if we assume all exposures are of equal size

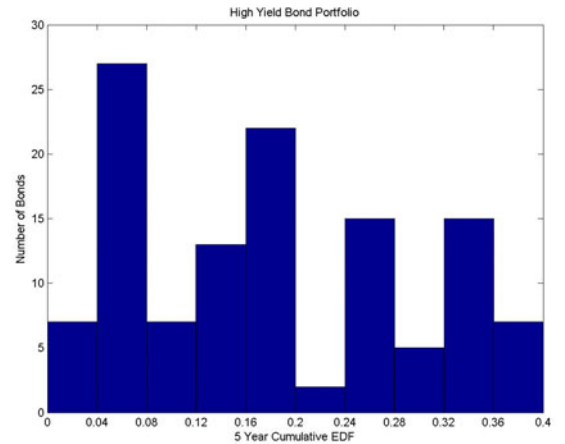


Figure 1: Distribution of 5 Year EDFs for Bond Portfolio

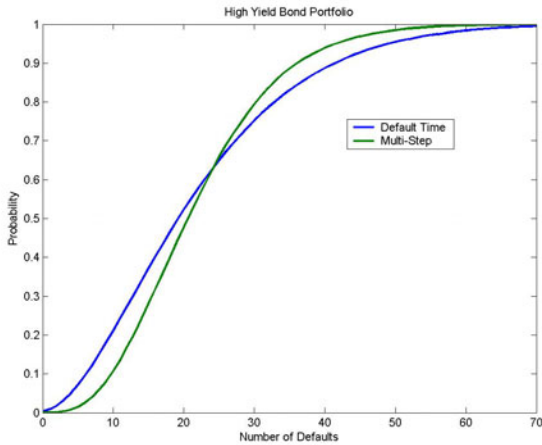


Figure 2: Comparison of Default Time and Multi-Step Models for Generating Correlated Defaults

and a recovery of 50%; in this case 60 defaults corresponds to a loss of 25%. For a senior tranche with 25% subordination, this gives us sense of the default probability of the tranches over the 5 year horizon (not accounting for structural effects of the cash flow waterfall). Under the default time model, the tranche would be considered a moderate investment grade asset, while under the multi-step model this would be a triple A investment.

For the second question of comparing the Brownian motion model with the empirical distribution model, we consider a two horizon case and compare the relative role of each time step’s asset return draw in determining default. For each model the total default probability is the same. However, depending on the credit migration over the first step (determined by the first asset return draw), the range of asset return draws require for default in the second time period can be substantially different. This is illustrated in Figure 3, which plots regions in the two dimen-

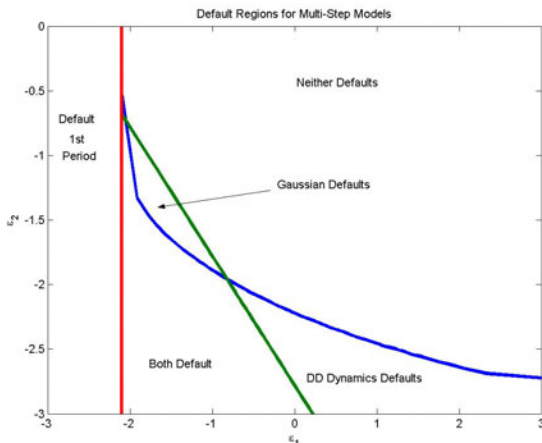


Figure 3: Comparison of Gaussian Process and Empirical Distance to Default Process for Two Steps of a Multi-Step Simulation

sional space (ϵ_1, ϵ_2) , where ϵ_i is the scaled asset return for period i ($i = 1, 2$) normalized to be a standard Normal variate. In this example the time interval for each period are equal and the cumulative default probabilities are 1.8% for the first period and 3.57% for the first and second periods together.

This plot shows that under the Gaussian process, the conditional default probability following a positive asset return on the first step is much lower than for the empirical distribution. By the same token, a negative asset return on the first step under the Brownian motion model is much more likely to lead to default in the second step. This model behavior differs from the Distance to Default empirical distribution behavior because it fails to capture the firms response to good asset returns (adding more debt) and bad asset returns (taking measures to avoid default).

Another striking difference between the models can be seen by considering the case for which the firm behaves according to expectation in the first period. This corresponds to $\epsilon_1 = 0$. The example has been chosen such that the default probability in each period is 1.8%. Yet for the Gaussian model, if the firm behaves at its expected level for the first period, the conditional default probability for the second period drops to 0.27%. For the DD Dynamics model, the second period default probability conditional on the expected asset return is around 1.4%. Note that the Gaussian process default boundary is linear because there is a fixed threshold such that if the two period cumulative asset return (the sum of the two one period asset returns) falls below it, the firm is in default.

6 CONCLUSIONS

The complexity of CDOs and the default behavior of the underlying collateral demand sophisticated simulations to capture the behavior accurately. True multi-step simulations have been shown here to yield significantly different results from single step default time approximations. In addition, a substantially more realistic credit migration behavior can be captured by using empirical distributions in place of the standard mathematical modeling approach.

ACKNOWLEDGMENTS

The work in this paper was built upon a great deal of research carried out by a number of people in the San Francisco research group at Moody’s KMV. In particular, special thanks are due to Yim Lee and Jeff Bohn.

REFERENCES

Arvanitis, A and J Gregory. 2001. *Credit: The Complete Guide to Pricing, Hedging and Risk Management*. Risk Publications.

- Avellaneda, M and J Zhu. 2001. Distance to default. *Risk* (December).
- Duffie, D and K Singleton. 2003. *Credit Risk: Pricing, Management and Measurement (Princeton Series in Finance)*. Princeton University Press.
- Falkenstein, E. 2000. RiskCalc™ for private companies: Moody's default model [online]. Available via <www.Moodyskmv.com/research/riskmodeling.html>[accessed August 5, 2003].
- Finger, C. 2000. A comparison of stochastic default rate models. *RiskMetrics Journal* 1 (November), 49-73.
- Goodman, L. and F. Fabozzi. 2002. *Collateralized Debt Obligations: Structure and Analysis*. John Wiley & Sons.
- Gregory, J and Laurent, J. 2003. I will survive. *Risk*. (June), 103-107.
- Hull, J. and A. White. 2001. Valuing credit default swaps II: modeling default correlations. *Journal of Derivatives* 8 (3), 12-22.
- Kealhofer, S. 2003. Quantifying credit risk I: default prediction. *Financial Analyst Journal* (January/February), 30-44.
- Li, D. 2002. On default correlation: a copula function approach. *Journal of Fixed Income* 9 (4), 43-54.
- Merton, R. 1974. On the pricing of corporate debt: the risk structure of interest rates. *Journal of Finance* 29 (2).
- Schmidt, W. and I. Ward. 2002. Pricing default baskets. *Risk* (January), 111-114.
- Vasicek, O. 1987. Probability of loss on loan portfolio [online]. Available via <www.Moodyskmv.com/research/portfoliotheory.html> [accessed August 5, 2003].

AUTHOR BIOGRAPHY

WILLIAM MOROKOFF is a director at Moody's KMV in the San Francisco research department where he heads the New Product Research Group. His research focuses on Monte Carlo simulation applied to risk management and derivatives pricing in finance. He holds a Ph.D. from the Courant Institute at New York University. He can be contacted at <morokoff@mkmv.com>.