

## CALIBRATING CREDIT PORTFOLIO LOSS DISTRIBUTIONS

Menghui Cao  
William J. Morokoff

New Product Research  
Moody's KMV  
99 Church Street  
New York, NY 10007, U.S.A.

### ABSTRACT

Determination of credit portfolio loss distributions is essential for the valuation and risk management of multi-name credit derivatives such as CDOs. The default time model has recently become a market standard approach for capturing the default correlation, which is one of the main drivers for the portfolio loss. However, the default time model yields very different default dependency compared with a continuous-time credit migration model. To build a connection between them, we calibrate the correlation parameter of a single-factor Gaussian copula model to portfolio loss distribution determined from a multi-step credit migration simulation. The *deal correlation* is produced as a measure of the portfolio average correlation effect that links the two models. Procedures for obtaining the portfolio loss distributions in both models are described in the paper and numerical results are presented.

### 1 INTRODUCTION

A synthetic Collateralized Debt Obligation (CDO) is a synthetic securitization of a pool of credit risky assets in which investors agree to insure the credit portfolio against a specified range of losses in exchange for the periodic payment of an insurance premium. The different loss ranges are referred to as the tranches of the CDO. In a typical synthetic CDO transaction, the portfolio loss distribution directly determines the tranche cash flows and thus the tranche valuation and risk analysis. The main drivers for the portfolio loss distribution are default probability, default correlation and the loss given default (LGD). While there are many credit risk models that address these issues, we focus here on the structural model from which the default time model and the multi-step credit migration model can be derived.

The structural approach was first proposed by Merton (1974). It makes assumptions on the underlying firm asset value process that drives the credit migration and defaults. A multi-step simulation model (Hull and White 2001, Morokoff 2003) can capture the default correlation and credit mi-

gration in a practical manner, but requires a step-wise simulation that may often be time-consuming.

The default time approach was developed by Li (2000) for pricing CDOs. It can be understood as a single-step version of the structural model, requiring only one-step simulation to generate the default times based on the marginal default probabilities and copula function. While the Gaussian copula is commonly used or implicitly implied for modeling the joint defaults, some papers claimed that other copula functions, such as Student- $t$  copula (Mashal and Naldi 2001) and Clayton copula (Rogge and Schönbucher 2003), may better capture the tail dependency.

The single-factor Gaussian copula is a simple form of the default time model with a semi-analytical solution so that Monte Carlo simulation is not required for pricing (Gregory and Laurent 2003, Hull and White 2004). Specifically, the loss distribution can be expressed as a low dimensional integral that may be evaluated with quadrature methods. An earlier work on this was by Vasicek (1987) on loan loss distribution. The main idea is: conditional on the latent factor, marginal default and survival probabilities are independent. An explicit form of the number-of-default distribution, or loss distribution, can be computed and used to value basket default swaps or synthetic CDO tranches. The sensitivity measures can be produced in a similar way.

In recent years, the default time model, especially the single-factor Gaussian copula version, has become an industry standard mostly due to its tractability. However, the default time model yields very different default dependency from using multi-step simulation. As pointed out in Morokoff (2003), the default time model tends to overestimate the default correlation when using a weekly asset return correlation to determine the joint defaults in a longer time horizon. What correlation parameter to use in the default time model is thus a problem worth careful consideration.

### 2 CREDIT PORTFOLIO LOSS DISTRIBUTION

We consider here a portfolio of  $M$  exposures, with notional amount  $IN_j$  and loss given default  $LGD_j$ ,  $j=1, \dots, M$ . The

cumulative default probability of an exposure up to time  $t$  is given by the function  $F_j(t)$  (note that this function is non-decreasing by definition). The key to determining the portfolio loss distribution at time  $t$  is modeling how the defaults up to time  $t$  are correlated. The model must be calibrated so the expected number of defaults for each individual name is given by  $F_j(t)$ .

## 2.1 Default Time Model

This section describes the methodology most commonly employed today for describing correlated defaults. It is known as the default time or copula approach and is described by Li (2000) and Schmidt and 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 as some point during the life of a CDO deal. The timing of the default, however, can also play a crucial role in determining the performance of the deal. Default timing is described by the function  $F_i(t)$ .

The default time/copula method is most easily described as a simulation for randomly sampling default times. The first step is to randomly sample a uniform (0,1) variate  $u$ . Let  $T$  be the maturity of the CDO deal. If  $u > F_i(T)$  then the exposure  $i$  does not default. If  $F_i(t) < u \leq F_i(t + \Delta t)$  then the exposure defaults in period  $(t, t + \Delta t]$ . 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  $(u_1, \dots, u_M)$ . 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  $F_j^{-1}(\square)$  is the inverse of the marginal probability distribution for the  $j^{\text{th}}$  exposure. However, any copula function may be used for this purpose. The most commonly used are Gaussian copula and  $t$ -copula, although a variety of other methods, including Clayton and 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 set of correlated standard normal variates  $(\varepsilon_1, \dots, \varepsilon_M)$  is sampled, either from a Cholesky decomposition of the correlation matrix or from a factor model de-

composition. The correlated uniform variates are then obtained from the formula

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

where  $\Phi$  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  $V$  conditional on which the default probabilities and default times are independent. As an example, the asset returns associated with a single-factor model with a Gaussian copula can be represented as

$$\varepsilon_i = \sqrt{R_i^2} V + \sqrt{1 - R_i^2} z_i,$$

where the latent variable  $V$  is assumed to have a standard normal distribution and the idiosyncratic component of the return  $z_i$  are iid standard normal and independent of the latent variable. The R-squared parameter ( $R_i^2$ ) determines the degree of variance of the asset return attributable to the latent factor. Conditional on the latent factor, the asset returns are all independent. The law of conditional expectations then allows the portfolio properties of interest to be expressed as an expectation of the portfolio properties of an independent portfolio with respect to the latent factor.

## 2.2 Single-Factor Gaussian Copula Model

This is a special case of the default time model in that it assumes there is one latent factor that explains all the systematic risk and the joint default is modeled through a Gaussian copula function as described above. It has the advantage of using a semi-analytical calculation instead of simulation.

To compute the loss distribution under this model, we first standardize each exposure's loss given default as an integer multiple of a loss unit amount. This can be considered as a special case of the second approach in Hull and White (2004), but may be computed faster than that approach in general and the fast Fourier transform in Gregory and Laurent (2003).

The unit of standardized loss,  $s/l$ , can be chosen as a small percentage of total collateral notional, a fraction of the minimum loss given default in the collateral, or a com-

bination of these two. By assuming constant LGD and constant exposure at default, we can calculate the number of standardized losses for each underlying exposure by

$$l_i = \text{int} \left( \frac{IN_i LGD_i}{sl} \right),$$

where  $\text{int}(x)$  is the function that rounds  $x$  to the closest integer.  $N_L = \sum_{i=1}^M l_i$  gives the maximum number of standardized losses. Given common factor  $V$ , we can calculate the conditional default probability  $p_i(t|V)$  by

$$p_i(t|V) = \Phi \left( \frac{\Phi^{-1}(F_i(t)) - R_i V}{\sqrt{1 - R_i^2}} \right),$$

where  $R_i$  is the correlation of exposure  $i$  with the latent factor,  $\Phi(x)$  is the cumulative normal distribution function,  $q_i(t|V) = 1 - p_i(t|V)$  is the conditional survival probability. A common practice is to assume constant asset correlation  $\rho$  throughout the portfolio, for which we need to set all  $R_i^2$  equal to  $\rho$ .

Given the conditional default probabilities, we can calculate the conditional number-of-default distribution  $Q(N(t)=n|V)$  and conditional standardized loss distribution  $Q(L(t)=l|V)$  through the probability generating function. Specifically, they can be calculated as the coefficient of  $u^n$  and  $u^l$  respectively from the following expansions according to a simple iterative process,

$$\sum_{n=0}^M Q(N(t)=n|V) u^n = \prod_{i=1}^M (q_i(t|V) + p_i(t|V) u),$$

$$\sum_{l=0}^{N_L} Q(L(t)=l|V) u^l = \prod_{i=1}^M (q_i(t|V) + p_i(t|V) u^{l_i}).$$

The two expansions are similar, except that in the second equation  $l_i$  is used as the exponent to account for the exposure-specific losses.

The number-of-default distribution  $Q(N(t)=n)$  and standardized loss distribution  $Q(L(t)=l)$  can be obtained through numerical integration of the corresponding conditional distributions with respect to  $V$ . This may be done through Gaussian quadrature with Hermite polynomials,

$$Q(N(t)=n) = \sum_m w_m Q(N(t)=n|V_m),$$

$$Q(L(t)=l) = \sum_m w_m Q(L(t)=l|V_m),$$

where  $V_m$  are the chosen quadrature points and  $w_m$  are the associated weights.

### 2.3 Multi-Step Credit Migration Model

The multi-step credit migration approach used in this paper has been given in Morokoff (2003). This alternative to the default time approach is based on simulating the firm asset value as a stochastic process. Similar methods have been proposed by Hull and White (2001), Arvanitis and Gregory (2001) and Finger (2000). In this section we describe an implementation of a 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 than one single step to maturity. The effects of the stochastic process for credit migration is also not accounted for in the default time model since it focuses only on the function  $F_i(t)$ , which for the multi-step credit migration process acts only as a calibration parameter.

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  $F(t)$ . In this formulation a free boundary problem PDE can be derived as described by Avellaneda and Zhu (2001). However, the existence of  $F(t)$  as a continuous-time 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 periods. 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 un-

scheduled 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 a discrete time process that makes use of distributions derived from empirical data. The method takes as input the cumulative default function  $F_j(t)$  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 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 ( $DD$ ) transition densities. Distance-to-Default refers to the number of standard deviations a firm's asset value is from a default point derived from the firm's liabilities. We now consider the implementation of this 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 continuous variable, it is unlikely that any of the sample will have a  $DD$  value of exactly 3. Thus it is necessary to rephrase 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 tran-

sition 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 Moody's KMV model is calibrated from 9 years of monthly data on 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 default probability. This forward default probability 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  $F(t)$  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:

$$f(T_{i-1}, T_i) = \frac{F(T_i) - F(T_{i-1})}{1 - F(T_{i-1})}.$$

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 default probability. 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.

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. It 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 on  $DD$  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  $\varepsilon_j$  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 default probability 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 default probability 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).

### 3 RESULTS AND DISCUSSIONS

Using the procedures described above, we obtain the portfolio loss distributions from multi-step simulation and from single-factor Gaussian copula with a group of constant correlations. Then we calculate the sum-squared-errors between a pair of loss distributions, one from multi-step simulation and one from single-factor Gaussian copula, across the group of constant correlations. A minimization procedure returns a single correlation number that best calibrates the loss distribution from multi-step simulation, which we name the *deal correlation*. It is a single number that captures the average correlation effect of the collateral portfolio, which could be easily communicated and compared across different CDO deals. It is similar, in some sense, to the Moody's Diversity Score.

In general, we can infer a deal correlation either from a number-of-default distribution or from a portfolio loss distribution. For a homogeneous portfolio in which the exposure sizes and LGD values do not vary much throughout the portfolio, they will return similar results of deal correlation, otherwise the results may deviate a lot from each other. The deal correlation implied by a number-of-default distribution is more meaningful for an  $n^{\text{th}}$ -to-default swap or a 'number-of-default' CDO, while the deal correlation implied by the portfolio loss distribution is more relevant for a regular CDO tranche.

As the numerical example, we consider a collateral portfolio of 100 names. Each name represents a notional amount of 10MM with 60% LGD, all maturing after 5 years. The 1-year default probabilities vary from 3 bps to 8% with an average of 0.86%, R-squared vary from 10% to 65% with an average of 32%, pair-wise asset correlations vary from 6% to 62% with an average of 24%. We are interested in the total portfolio loss within a 5-year horizon. The multi-step simulation is carried out with a semi-annual frequency, i.e., 10 time steps for 5 years.

Figure 1 shows a comparison of the portfolio loss distribution from multi-step simulation and from single-factor Gaussian copula. As the correlation number for the latter increases, the portfolio loss distribution tends to have both a higher 'head' (low loss scenario) and heavier 'tail' (high loss scenario), indicating a higher default correlation.

The deal correlation for this portfolio is 16%, which is lower than the portfolio average asset correlation (24%) and portfolio average R-squared (32%). It's because the default time model tends to overestimate default correlation when using instantaneous asset return correlation to infer the long-term joint defaults. Figure 2 shows that the portfolio loss distribution from multi-step simulation matches almost perfectly with the loss distribution using single-factor Gaussian copula with 16% correlation.

As we can see from Figure 1 and Figure 2, different correlations lead to completely different loss distributions under the copula model, which, of course, incurs different tranche valuation and risk analysis results.

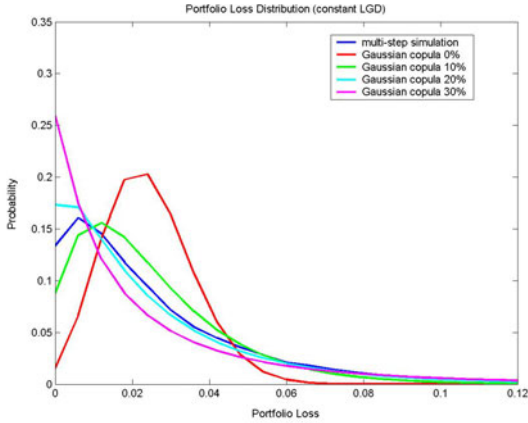


Figure 1: Portfolio Loss Distribution with Constant LGD Multi-Step Simulation vs. Single-Factor Gaussian Copula

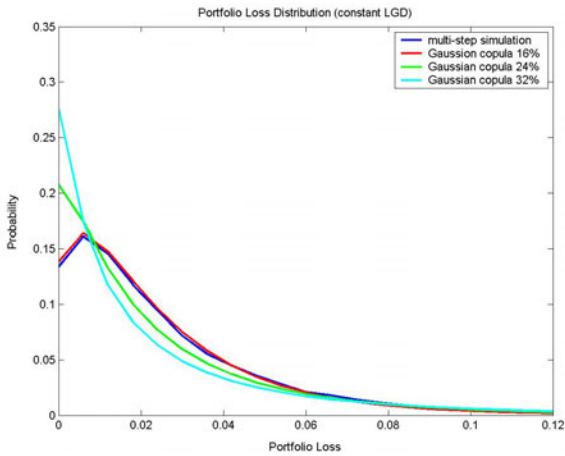


Figure 2: Portfolio Loss Distribution with Constant LGD (Deal Correlation=16%, Portfolio Average Asset Correlation=24%, Portfolio Average R-squared=32%)

Table 1 lists the sum-squared-errors (between loss distributions from multi-step simulation and Gaussian copula model) and tail loss probabilities, corresponding to different correlation numbers. Consider a 10-20% tranche for example (i.e., a tranche that absorbs the portfolio loss between 10% and 20%), 24% correlation implies a 3.17% probability of getting hit while 16% correlation only incurs a 1.67% chance.

Table 1: Sum-Squared-Errors and Tail Loss Probabilities (Portfolio Loss Distribution with Constant LGD)

$\rho$	0%	8%	16%	24%	32%
S.S.E. ( $10^{-4}$ )	537	84	1	67	245
P(L>10%)	0%	0.35%	1.67%	3.17%	4.48%

As a comparison, we study the portfolio loss distribution with stochastic LGD, or equivalently stochastic recovery rate. We assume the LGD follows a Beta distribution

with mean 0.6 and shape parameter 4, there is no correlation between default probability and LGD. The portfolio loss distribution in this case (Figure 3) is similar in shape to the loss distribution with constant LGD, while the former has a little higher ‘head’ and heavier ‘tail’ than the latter, since the randomness in LGD flattens the loss distribution on both ends.




Figure 3: Portfolio Loss Distribution with Stochastic LGD (Multi-Step Simulation vs. Single-Factor Gaussian Copula)

Sum-squared-errors between distributions and tail loss probabilities under stochastic LGD are given in Table 2.

Table 2: Sum-Squared-Errors and Tail Loss Probabilities (Portfolio Loss Distribution with Stochastic LGD)

$\rho$	0%	8%	16%	24%	32%
S.S.E. ( $10^{-4}$ )	456	79	1	69	247
P(L>10%)	0%	0.41%	1.70%	3.17%	4.46%

## 4 CONCLUSION

We describe in this paper the methods to obtain the credit portfolio loss distributions from multi-step credit migration model and single-factor Gaussian copula model. The loss distributions are compared to produce a deal correlation number, which measures the portfolio average correlation effect and can be compared across different deals.

## ACKNOWLEDGMENTS

The work in this paper is part of a research project undergone by a number of people in the New Product Research Group at Moody’s KMV. In particular, the authors would like to thank Yim Lee and Frank Li for their help and contributions to this work.

## 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: 125-129.
- Finger, C. 2000. A Comparison of Stochastic Default Rate Models. *RiskMetrics Journal* 1 (November): 49-73.
- Gregory, J., and J-P. Laurent. 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.
- Hull, J., and A. White. 2004. Valuation of a CDO and an  $n^{\text{th}}$ -to-default CDS without Monte Carlo Simulation. Working paper. University of Toronto.
- Li, D. 2000. On Default Correlation: A Copula Function Approach. *Journal of Fixed Income* 9(4): 43-54.
- Mashal, R., and M. Naldi. 2001. Pricing Multi-name Credit Derivatives: Heavy Tailed Hybrid Approach. Working paper. Columbia University and Lehman Brothers.
- Merton, R. 1974. On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. *Journal of Finance* 29: 449-470.
- Morokoff, W. 2003. Simulation Methods for Risk Analysis of Collateralized Debt Obligations. In *Proceedings of the 2003 Winter Simulation Conference*, ed. S. Chick, P. J. Sánchez, D. Ferrin and D. J. Morris, 335-342.
- Rogge, E., and P. Schönbucher. 2003. Modeling Dynamic Portfolio Credit Risk. Working paper. Imperial College, ABN AMRO Bank and ETH Zurich.
- Schmidt, W., and I. Ward. 2002. Pricing Default Baskets. *RISK* January: 111-114.
- Vasicek, O. 1987. Probability of Loss on Loan Portfolio. Moody's KMV. Available online via <[http://www.moodyskmv.com/research/whitepaper/Probability\\_of\\_Loss\\_on\\_Loan\\_Portfolio.pdf](http://www.moodyskmv.com/research/whitepaper/Probability_of_Loss_on_Loan_Portfolio.pdf)> [accessed July 31st, 2004].

## AUTHOR BIOGRAPHIES

**MENGHUI CAO** is a financial researcher at Moody's KMV in the New Product Research Group. His research interest is on Monte Carlo simulation and derivatives pricing, currently focusing on pricing, hedging and risk analysis of CDOs. He's finishing his PhD in Operations Research at Columbia University and can be contacted at <Menghui.Cao@mkmv.com>.

**WILLIAM MOROKOFF** is a director at Moody's KMV in the 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 PhD in Mathematics from the Courant Institute at New York University. He can be contacted at <William.Morokoff@mkmv.com>.