Proceedings of the 2005 Winter Simulation Conference
M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, eds.

SIMULATION OF RISK AND RETURN PROFILES FOR PORTFOLIOS OF CDO TRANCHES

William J. Morokoff

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

ABSTRACT

Investments in Collateralized Debt Obligations (CDOs) often offer attractive yields relative to other similar debt instruments (corporate bonds, etc.). However, the risk profiles of CDO investments, and in particular portfolios of these investments, can be substantially different from straight credit portfolios due to complex correlation dependence across CDOs. Simulation is generally required to capture the intricate interaction of default and correlation risk that determines the risk and return profile of a portfolio of CDO investments. This paper considers some of the issues that must be addressed in determining the risk profiles with simulation and presents results on a simple example.

1 INTRODUCTION

Collateralized Debt Obligation (CDO) tranches are complex financial instruments that share many similarities to corporate bonds. Like bonds, tranches pay a periodic coupon that may be at risk if there is serious deterioration of the credit quality of the backing entity. For CDOs, the backing entity is generally a collateral pool of debt instruments such as bonds, loans, credit default swaps or asset backed securities (e.g. mortgage pools). Serious credit deterioration in a CDO means that many of the individual names in the pool default or suffer downgrades and loss in value. There are several variations on how the collateral pools can be arranged: cash CDOs require actual ownership by the deal of the underlying bonds or loans; synthetic CDOs simply reference a list of names that issue bonds or loans; and CDO squared deals have collateral pools consisting of investments in other CDOs (either as cash investments or synthetically referenced).

CDO tranches are partial claims on the performance of the underlying CDO collateral pool. They are generally ranked by seniority, this the most senior tranche being paid first before the other obligations. It is therefore the least risky and, as such, earns the smallest promised coupon. The collective set of tranches is known as the capital structure of the deal. As one moves down the capital structure from senior to mezzanine to equity tranches, the promised returns increase as does the risk of not being paid. The rules by which the tranches are paid are laid out in a cash flow waterfall that describes the structure of the deal. The waterfall is generally straightforward for synthetic deals, but can be extremely complex for cash deals due to complicated credit enhancement rules intended to protect the more senior tranches from collateral deterioration.

Quantitative modeling of credit valuation, default risk and portfolio risk is a well established field. Extended discussions of the modeling issues can be found in Arvantis and Gregory (2001), Bluhm et al (2002), Duffie and Singleton (2003) and Gordy (2003). For a discussion specific to CDOs see Goodman and Fabozzi (2002). Papers discussing quantitative methods for CDOs include Hull and White (2004) and Morokoff (2003).

2 MODELING PERFORMANCE OF CREDIT PORTFOLIOS

Consider a loan portfolio to be evaluated for potential losses over the next year. In order to evaluate the value at risk due to credit changes it is necessary to i) value each loan today; and ii) value each loan at the horizon as a function of the credit state at horizon, accounting for any coupon payments prior to horizon and any recovery on loans that may default before the horizon.

Valuation of a loan today requires a risk neutral default probability term structure that describes the current credit state of the borrower. The risk neutral default probabilities are adjustments to the actual, physical default probabilities (measured based on historical performance of similar credits). The adjustments, which lead to higher default probabilities, are required to account for the additional premium required by investors to take on non-diversifiable default risk. The default probability term structure is used to determine the probability of receiving each coupon and principal payment until the maturity of the loan.

The valuation of a loan at the horizon also requires a risk neutral default probability term structure that describes the credit state at horizon of the borrower looking forward from horizon. Future cash flows after horizon are weighted by probabilities of being received and discounted back to the horizon to establish the value at horizon.

It is important to note that the value either today or at horizon of a loan depends only on the credit state of a single borrower. The time evolution of this credit state can generally be described by a process for a single state variable. For example, in a structural model, the asset value of the borrowing firm would be the relevant variable. (Note that loan valuation may also depend on the stochastic evolution of interest rate terms structures. This aspect is not addressed here in order to focus on the credit risk).

To determine a probability distribution of the value of a loan at horizon, it is necessary to provide a probability distribution for the transition from the credit state today to the credit state at horizon. This credit state transition process establishes how the probability of default evolves over time. Note that the transition probability density for the credit state should be in the physical (also known as real world) measure – that is, it should be consistent with actual default and credit migration probabilities, not the risk neutral measure. We are interested in computing the actual risk associated with real probabilities of credit migration and default over the horizon period, although the valuation of the loans themselves must be done in the risk neutral space. It is in fact precisely the difference between the physical and risk neutral default probabilities that leads to a positive excess return on the loan as measured by the difference in the expected value at horizon and the value today.

In order to calculate the risk of a portfolio of loans, it is necessary to model the correlated change in the credit state across various borrowers. Unless the correlated credit state model has a particularly simple structure (e.g. it can be determined by a single scalar random variable), it is necessary to use Monte Carlo simulation to sample the correlated credit states. Once the credit states at horizon have been sampled, the value of each loan at horizon can be computed, and thus also the value of the portfolio at horizon. Given the value of the portfolio today, the change in value of the portfolio can be computed from which quantities such as unexpected loss (i.e. volatility of the portfolio returns) and credit value at risk can be computed.

If all the loans in the portfolio have approximately the same individual expected return, then the portfolio will also have this same expected return. The benefits of portfolio diversification can then be readily seen as the portfolio unexpected loss will decrease as more names are added (assuming they are not perfectly correlated). The less correlated the changes in credit states of the borrowers are, the lower the risk in the portfolio. A standard measure that balances portfolio return with portfolio unexpected loss is the Sharpe ratio which is defined as

$$S = \frac{\mu_\pi - r}{\sigma_\pi}.$$

Here μ_π, r, and σ_π are the portfolio expected return, the risk free return and the portfolio unexpected loss to horizon. Portfolio managers often look to maximize the Sharpe ratio of their portfolios either by increasing returns or by reducing unexpected loss through diversification.

3 CDO TRANCHE RISK AND RETURN

Similar to computing the credit risk associated with a loan, one can imagine computing the value of a CDO tranche today, the distribution of value of the tranche at a future horizon (say 1 year) conditional on the credit state at horizon, and thereby determine the return distribution to horizon for the tranche. Quantities such as unexpected loss and Sharpe ratio can be computed for the tranche, and by accounting for correlated tranche dependence, the portfolio return distribution for portfolios of tranches can be modeled.

One significant drawback to this approach is that the valuation of a CDO tranche today can be a fairly substantial calculation, depending on the type of CDO and the credit migration and correlation models being applied. Unlike a loan, the value of a tranche is not determined by the relatively simple discounting of future coupons weighted by a risk neutral probability of default (even more complicated loans with prepayment options and other features are still relatively straightforward to model). The value of a tranche depends on the credit migration and default probabilities of the entire collateral pool, which may contain several hundred correlated credits. Thus the value of the tranche depends on the credit state of hundreds of names. Moreover, for cash CDOs, the complex rules governing the cash flows to the tranche, which are credit state dependent, must be taken into account. There are significant other complexities not described here that indicate that Monte Carlo simulation must be used to value tranches if more sophisticated credit models are applied.

Valuation of a tranche at horizon is even more complicated, as it requires accounting for events prior to horizon (e.g. default losses on the collateral pool and cash flows to the tranche) as well as evaluating the expected future cash flows to the tranche after horizon. These future cash flows are conditional on the credit state at horizon, which depends on the correlated credit states of all the names in the collateral pool. A practical implementation of this approach requires careful calibration of functions that approximate tranche value to avoid the need for an extremely slow simulation within a simulation approach.

For the purposes of this paper, we sidestep many of these complications by considering a simplified example of a CDO that nonetheless reveals a number of interesting properties when analyzed with simulation. The CDO considered here is similar to a simple pass-through synthetic CDO with a five year maturity, but with "zero coupon" tranches. This means that each tranche receives no interest or other cash flows for five years, at which time it receives the par amount minus any losses due to default that would be assigned to the tranche. Since the maximum amount the tranche can receive is par at the five year horizon, the purchase price must be at a discount to par, with the size of the discount corresponding to the degree of risk of incurring losses to the tranche principal.

At the top of the capital structure there is a super-senior tranche corresponding to 30 – 100% of the collateral pool; this means that the tranche will be responsible for covering all losses in excess of 30%. The senior tranche covers losses from 15 – 30% of the original pool. The remaining tranches cover 10 – 15%, 7 – 10%, 3 – 7% and 0 – 3% of the losses. For example, the 0 – 3% tranche will receive full principal at five years only if no defaults occur, and it will receive nothing it the losses exceed 3%.

The underlying collateral reference pool consists of 100 identical names with asset correlation 0.3, a one year physical default probability of 1%, and a five year physical default probability of 6%. The loss that occurs upon default is assumed to be 50%. For simplicity, the risk free discount rate used to compute the present value of the principal payment at maturity is assume to be zero.

To further simplify matters, the default, correlation, credit migration and valuation models used for this example are chosen to facilitate ease of implementation and transparency. The methodology used to value the tranche at the analysis data and at the one year horizon is similar to the approach described in Hull and White (2004) and related to the Copula method discussed in Li (2002). The simulation methodology is two step implementation of the Gaussian dynamics multi-step method described in Morokoff (2003).

The simulation proceeds as follows. First compute the value of each tranche today (i.e. the analysis date) by computing the probability over the five year maturity of various loss levels and taking an expected value of the corresponding tranche value relative to this probability distribution. Note that the appropriate probability to use is the risk neutral measure. The probabilities of losses are computed based on a single step Gaussian copula using the asset correlation parameter and the risk neutral five year default probability.

The second step is to simulate one time step out to the one year horizon under the physical default probability measure. Thus each name in the collateral portfolio will default with the true one year default probability (set here to 1%). For names that do default, a loss of 0.5% is assigned. For names that do not default in the first year, a new four year default probability, corresponding to the period from the one year horizon to maturity, is computed based on the credit state sampled for that name as part of this one step simulation. For the model employed here, this is represented as a sample for a standard Normal distribution. The larger the sample, the better the credit state and the lower the associated four year default probability. Once the correlated forward four year default probabilities have been sampled for all the non-defaulted names in the portfolio, they can be converted to the risk neutral default probabilities, and the valuation method for the tranche described above can be used to determine the value for each tranche at the horizon conditional on the credit state of all the names in the collateral reference pool.

Because the value of the tranche at the analysis date and the value distribution at horizon are computed from the model, a consistent model requires that the expected change in value of the tranche must be positive (i.e., there is a positive return). This is the compensation for the risk. To achieve a valuation model consistent with the simulation, it is necessary to compute the risk neutral five year default probability based on a two step model. This effectively converts from physical to risk neutral based the assumption of a Brownian motion process that may default (i.e. be below a default barrier) at either the horizon or at the five year maturity. The default barriers are determined by the physical default probabilities to horizon and maturity, which are functions of the drift and volatility of the physical process. The risk neutral default probabilities are then computed as the probability of being below these barriers under a measure whereby the drift of the process has been shifted to the risk free rate. Instead of specifying the drift and volatility of the physical process, we derive these from the market risk premium, treated here as a parameter with value 0.4, and the correlation of each name in the collateral portfolio with the overall market, which is determined from the asset correlation value. Under this model, the five year risk neutral default probability is computed as about 13.9%. It is interesting to note that the expected loss on the collateral pool for the five year period under the physical measure is 3%, while under the risk neutral measure it is around 7%. Thus the physical and risk neutral measures provide a very different view of the performance of the 3-7% and 7 – 10% tranches.

Table 1 shows the summary of the tranche prices and performance measures. As expected, the 0 – 3% 'equity' tranche is discounted the most heavily reflecting the substantial probability of partial or complete loss. It prices to only 19.2 on a par of 100. Because the maximum price a tranche can have in this model is par, the equity tranche also has the most potential upside, and therefore the largest potential return to horizon. If a scenario were to arise in which all of the names in the collateral pool migrated over the horizon period to a perfect credit state (i.e., zero default probability), then the price at horizon would be par, and the return to horizon would be the maximum of 421%.

Table 1: Tranche Price and Performance Measures

Tranche	Price	Expected Return	Standard Deviation	Sharpe Ratio
0 – 3%	19.2	29.5%	80.0%	0.369
3 – 7%	50	15.9%	39.5%	0.401
7 – 10%	69.4	9.6%	23.5%	0.406
10 – 15%	82.3	5.6%	14.0%	0.403
15 – 30%	95.3	1.6%	4.3%	0.378
Portfolio	93.0	1.3%	3.1%	0.409

The portfolio row of Table 1 refers to the collateral portfolio, considered here as a portfolio of zero coupon bonds. The price of the portfolio reflects the 7% expected loss under the risk neutral measure.

As we go up the capital structure from the equity tranche to the senior 15 – 30% tranche, we see the price increase, and the expected return and risk (standard deviation) decrease. Most interesting, however, is that the Sharpe ratio of the tranches and the portfolio all remain about the same value near 0.4. This indicates that, at least with regard to the Sharpe ratio measure, all the tranches are being consistently compensated for their risk. It should be noted again that the prices and returns here are all model based and consistent. It is likely that for real CDOs there may be substantial differences in Sharpe ratio across the tranches if market prices are used.

Figure 1 shows the probability density function of returns for the portfolio and the 15 – 30% tranche as computed in the simulation. It is interesting to note that the portfolio has slightly higher maximum potential upside and a significantly smaller extreme potential downside. The expected return for the tranche is slightly higher. The volatility of the tranche is also higher; however, this stems from the fat tail, as the center of the distribution is somewhat tighter for the tranche than for the portfolio.

Figure 2 shows the probability density function for the returns on the mezzanine and equity tranches. Here we see that the equity returns range from -100% (complete loss of value) to upwards of 250%, with the distribution being almost flat over much of this interval. For the equity tranche there is in fact a substantial probability of a total loss which puts a delta function mass at -100%, not illustrated in the graph here. The 3 – 7 % and 7 – 10% tranche densities become progressively tighter and more peaked around their means.

It is interesting to note that while all the tranches have similar Sharpe ratios, they have very different risk profiles, in particular with regard to tail risk. The Sharpe ratio is not necessarily the best measure of performance for these tranches.

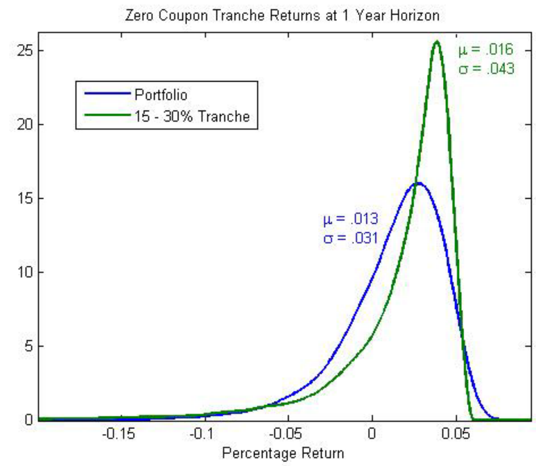

Figure 1: Return on Senior Tranche and Total Portfolio

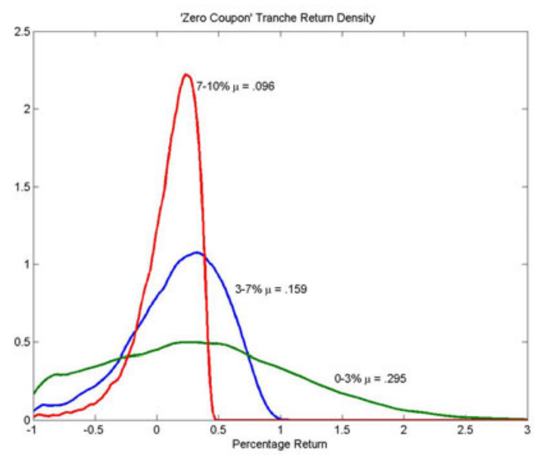

Figure 2: Return on Mezzanine and Equity Tranches

4 PORTFOLIO OF TRANCHES

In addition to the stand-alone risk/return profiles of individual tranches, it is of great interest to consider how portfolios of correlated tranches perform. It is becoming increasing common to also hedge a portfolio of tranches with a position in another tranche, a position in the entire underlying collateral portfolio (particular in the case of tranche written on an established credit index portfolio), or with positions in single names.

For the current example, we consider hedging a long position in the equity tranche with a short position in the portfolio. The hedge ratio is taken to be 11. This means that for every dollar invested in the equity tranche, 11 dollars of the portfolio should be sold. This ratio was chosen so that the portfolio consisting of the long equity position and short portfolio position would have approximately the same expected return as the 3 – 7% tranche. In this way, we can compare the risk profiles for two investment options that have the same expected return.

Figure 3 compares the return distributions for 3 – 7% tranche and the hedged equity portfolio. It is clear that the risk is distributed very differently although the expected return is the same. For the 3 – 7% tranche there is a small but significant probability of losing everything corresponding to the chance that 14% of the names default before horizon (under the physical measure). In the hedged portfolio, under this extreme scenario the equity tranche is also wiped out; however, the portfolio suffers substantial losses too, and with the 11 times leverage and short position, this scenario in fact leads to a large positive return. Thus we observe that the hedged portfolio has zero probability of having a -100% return (defined based on change in value of the amount invested in the equity tranche). The hedged portfolio also has a much large potential positive return than the straight tranche position. Thus the hedged portfolio has much less extreme downside and much better extreme upside. This is paid for, however, by a riskier center of the distribution which covers both a broader range of value and is also somewhat peaked over negative returns. Another interesting feature of the hedged portfolio is the density return fingers that peak around -50% and -40% returns. Closer inspections of the scenarios that lead to these returns show that they correspond to having four defaults prior to horizon and three defaults prior to horizon respectively.

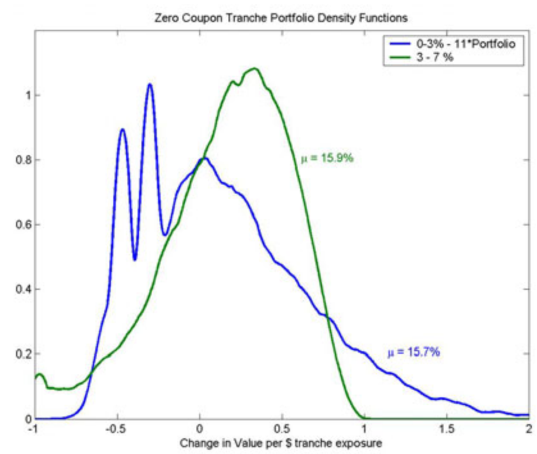

Figure 3: Return on Mezzanine and Hedged Equity Tranches

ACKNOWLEDGMENTS

The work in this paper was built upon a great deal of research carried out by a number of people in the research group at Moody's KMV. In particular, special thanks are due to Yim Lee, Saiyid Islam, Michele Freed, David Cao and Frank Li.

REFERENCES

Arvanitis, A and J Gregory. 2001. *Credit: The complete guide to pricing, hedging and risk management*. Risk Publications.

Bluhm, C., L. Overbeck, and C. Wagner. 2002. *An introduction to credit risk modeling.* Chapman & Hall/CRC.

Duffie, D., and K. J. Singleton. 2003. *Credit risk: Pricing, measurement and management.* Princeton University Press.

Goodman, L. and F. Fabozzi. 2002. *Collateralized debt obligations: Structure and analysis*. John Wiley & Sons.

Gordy, M. 2003. *Credit risk modelling: The cutting-edge collection – technical papers published in Risk 1999-2003*. Risk Books.

Hull, J. and A. White. 2004. Valuation of a CDO and an n^{th}-to-default CDS without Monte Carlo simulation. *Journal of Derivatives* 12 (2) (Winter 2004), 8-23.

Li, D. 2002. On default correlation: a copula function approach*. Journal of Fixed Income* 9 (4), 43-54.

Morokoff, W.J. 2003. Simulation methods for risk analysis of collateralized debt obligations [online]. Available via <http://www.wintersim.org/prog03.htm#RA> [accessed August 9, 2005].

AUTHOR BIOGRAPHY

WILLIAM MOROKOFF is a senior 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 Ph.D. from the Courant Institute at New York University. He can be contacted at <morokoff@mkmv.com>.