Valuation of Forward Starting CDOs*

Ken Jackson[†] Wanhe Zhang[‡]

February 10, 2007

Abstract

A forward starting CDO is a single tranche CDO with a specified premium starting at a specified future time. Pricing and hedging forward starting CDOs has become an active research topic. We present a method for pricing a forward starting CDO by converting it to an equivalent synthetic CDO. The value of the forward starting CDO can then be computed by the well developed methods for pricing the equivalent synthetic one. We illustrate our method using the industry-standard Gaussian-factor-copula model. Numerical results demonstrate the accuracy and efficiency of our method.

1 Introduction

A forward starting CDO is a forward contract obligating the holder to buy or sell protection on a specified CDO tranche for a specified periodic premium at a specified future time. For example, a forward starting CDO might obligate the holder to buy protection on a CDO tranche with attachment point a and detachment point b over a future period of $[T, T^*]$ for a predetermined spread s. Hence, the maturity of the forward contract is T, and the maturity

*This research was supported in parted by the Natural Sciences and Engineering Research Council (NSERC) of Canada.

[†]Department of Computer Science, University of Toronto, 10 King's College Rd, Toronto, ON, M5S 3G4, Canada; krj@cs.toronto.edu

[‡]Department of Computer Science, University of Toronto, 10 King's College Rd, Toronto, ON, M5S 3G4, Canada; zhangw@cs.toronto.edu

of the forward starting CDO is T^*. At time T, the contract turns into a single tranche CDO over $[T, T^*]$ with attachment point $(a + L_T)$ and detachment point $(b + L_T)$, where L_T is the pool losses before T.

Pricing and hedging of forward starting CDOs has become an active research area. The most common approach is Monte Carlo simulation. Such methods are flexible, but are computationally expensive. Therefore, more efficient analytical or semi-analytical approaches are being developed by researchers. Bennani [4], Schönbucher [14], and Sidenius, Piterbarg, and Andersen [15] propose a dynamic modeling approach to capture the evolution of the aggregate portfolio losses. In order to price forward starting CDOs, they first simulate the pool losses L_T. Conditional on the simulated path, they price the forward contract by specifying the dynamics of the aggregated losses over $[T, T^*]$. The approximation of L_T constrains the accuracy and efficiency of these methods. In addition, their models require a large amount of data to calibrate, so they are not applicable to bespoke CDOs now.

Another class of forward starting CDOs, in which the tranche attachment and detachment points remain the same as a and b at time T, is straightforward to price using methods for synthetic CDOs, as shown in Hull and White [7]. For this type of contract, Hull and White [6] introduce a relatively simple dynamic process. They model the dynamics of the representative company's cumulative default probability by a simple jump process in the form of a binomial tree. In the homogeneous pool, all underlying companies have the same cumulative default probability, while in the heterogenous pool, it is possible to find a representative company matching the CDS spreads of the underlying pool. Thus, the forward starting CDOs can be priced through the dynamics of the representative company. Walker [17] extracts the tranche loss distributions from the market quotes, then the pricing of forward starting CDOs becomes straightforward with known tranche loss distributions. For nonstandard tranches, he employs an interpolation and extrapolation process.

In this paper, we price the first type of forward starting CDOs (with attachment point $(a+L_T)$ and detachment point $(b+L_T)$ at time T) by converting it to an equivalent synthetic CDO and then pricing the equivalent synthetic one based on the market-standard Gaussian-factor-copula model. Our approach avoids the consideration of the pool losses before T and is applicable to both index tranches and bespoke CDOs. (While preparing this paper, we learned that Ben De Prisco and Alex Kreinin [13] developed a similar method to price forward starting CDOs and Leif Andersen applied a similar approach to numerically test the correlation of losses across time in forward starting CDOs in his recent paper [1], although the method is not explained in detail there.)

The rest of the paper is organized as follows. Section 2 describes the pricing equations for forward starting CDOs. Section 3 derives a method to convert forward starting CDOs to equivalent synthetic CDOs. Section 4 reviews the widely used Gaussian factor copula model. Section 5 introduces a valuation method for synthetic CDOs based on the conditional independence framework. Section 6 tests two numerical examples. Section 7 discusses the extension of our method. Section 8 concludes the paper.

2 Pricing equation

In a forward starting CDO, the protection seller absorbs the pool losses specified by the tranche structure. That is, if the pool losses over $[T, T^*]$ are less than the tranche attachment point a, the seller does not suffer any loss; otherwise, the seller absorbs the losses up to the tranche size $S = b - a$. In return for the protection, the buyer pays periodic premiums at specified times $T_1 < T_2 < \ldots < T_n = T^*$, where $T_i > T = T_0$, for $i = 1, \ldots, n$.

We consider a forward starting CDO containing K instruments with recovery-adjusted notional value N_k for name k in the original pool. Assume that the recovery rates are constant, and the interest rate process is deterministic. Let d_i denote the discount factors corresponding to T_i. Denote the original pool losses up to time T_i by L_i, then the effective pool losses over $[T, T_i]$ is $\hat{L}_i = L_i - L_T$. Therefore, the losses absorbed by the specified tranche are

$$\mathcal{L}_i = \min(S, (\hat{L}_i - a)^+), \quad \text{where } x^+ = \max(x, 0) \tag{1}$$

In general, valuation of a CDO tranche balances the expectation of the present values of the premium payments (premium leg) against the effective tranche losses (default leg), such that

$$\mathbb{E}\Big[\sum_{i=1}^{n} s(S-\mathcal{L}_i)(T_i-T_{i-1})d_i\Big] = \mathbb{E}\Big[\sum_{i=1}^{n}(\mathcal{L}_i-\mathcal{L}_{i-1})d_i\Big] \tag{2}$$

The fair spread s is therefore given by

$$s = \frac{\mathbb{E}\big[\sum_{i=1}^{n}(\mathcal{L}_i-\mathcal{L}_{i-1})d_i\big]}{\mathbb{E}\big[\sum_{i=1}^{n}(S-\mathcal{L}_i)(T_i-T_{i-1})d_i\big]} = \frac{\sum_{i=1}^{n}(\mathbb{E}\mathcal{L}_i-\mathbb{E}\mathcal{L}_{i-1})d_i}{\sum_{i=1}^{n}(S-\mathbb{E}\mathcal{L}_i)(T_i-T_{i-1})d_i}$$

Alternately, if the spread is set, the value of the forward starting CDO is the difference between the two legs:

$$V_{\text{fwd}} = \sum_{i=1}^{n} s(S-\mathbb{E}\mathcal{L}_i)(T_i-T_{i-1})d_i - \sum_{i=1}^{n}(\mathbb{E}\mathcal{L}_i-\mathbb{E}\mathcal{L}_{i-1})d_i$$

Therefore, the problem is reduced to the computation of the mean tranche losses, $\mathbb{E}\mathcal{L}_i$.

3 Forward starting CDOs to synthetic CDOs

From (1), we know that the expectation of the tranche losses $\mathbb{E}\mathcal{L}_i$ is determined by the distribution of the effective pool losses $\hat{L}_i$. If we denote the default time of name k by τ_k and define the indicator function $\mathbf{1}_{\{\tau_k \le t\}}$ by

$$\mathbf{1}_{\{\tau_k \le t\}} = \begin{cases} 1, & \tau_k \le t \\ 0, & \text{otherwise} \end{cases}$$

then we have

$$\hat{L}_i = L_i - L_T = \sum_{k=1}^{K} N_k \mathbf{1}_{\{\tau_k \le T_i\}} - \sum_{k=1}^{K} N_k \mathbf{1}_{\{\tau_k \le T\}} = \sum_{k=1}^{K} N_k \mathbf{1}_{\{T < \tau_k \le T_i\}} \tag{3}$$

The right most sum in (3) is the expression of the pool losses in a synthetic CDO starting at time T. Therefore, the pool losses in our forward starting CDOs are equivalent to the pool losses in this synthetic CDO. The distributions of the effective pool losses $\hat{L}_i$ is determined by whether the underlying names default in $[T, T_i]$, and they can be computed through the equivalent synthetic CDO with modified default probabilities. That is, instead of using the probability that name k defaults before T_i in the synthetic CDO, we use the probability that name k defaults during the period $[T, T_i]$ in the equivalent synthetic CDO.

Remark: According to the argument above, a synthetic CDO can be treated as a special case of a forward starting CDO with $T = 0$.

In the next section, we specify the default process for each name and the correlation structure of the default events needed to evaluate $\mathbb{E}\mathcal{L}_i$. This will allow us to price forward starting CDOs using the well-known methods for pricing the equivalent synthetic CDO.

4 Gaussian factor copula model

Due to their tractability, Gaussian factor copula models are widely used to specify a joint distribution for default times consistent with their marginal distribution. A one factor model was first introduced by Vasicek [16] to evaluate the loan loss distribution, and the Gaussian copula was first applied to multi-name credit derivatives by Li [11]. After that, the model was generalized by Andersen, Sidenius, and Basu [3], Andersen and Sidenius [2], Hull and White [5], and Laurent and Gregory [9], to name just a few. In this section, we review the one-factor Gaussian copula model to illustrate the conditional independence framework and introduce the conditional forward default probabilities.

4.1 One factor copula

Assume the risk-neutral (cumulative) default probabilities

$$\pi_k(t) = \mathbb{P}(\tau_k \leq t), \quad k = 1, 2, \ldots, K$$

are known[1]. In order to generate the dependence structure of default times, we introduce random variables U_k, such that

$$U_k = \beta_k X + \sigma_k \varepsilon_k, \quad \text{for } k = 1, 2, \ldots, K \tag{4}$$

where X is the systematic risk factor reflecting the health of the macroeconomic environment; ε_k are idiosyncratic risk factors, which are uncorrelated with each other and also uncorrelated with X; the constants β_k and σ_k, satisfying $\beta_k^2 + \sigma_k^2 = 1$, are assumed to be known[2]. The random variables X and ε_k follow zero-mean unit-variance distributions, so the correlation between U_i and U_j is $\beta_i \beta_j$.

The default times τ_k and the random variables U_k are connected by a percentile-to-percentile transformation, such that

$$\pi_k(t) = \mathbb{P}(\tau_k \leq t) = \mathbb{P}(U_k \leq u_k(t))$$

where each $u_k(t)$ can be viewed as a default barrier. Thus the dependence among default times is captured by the common factor X. If we assume X and ε_k follow standard normal distributions, each U_k also follows a standard normal distribution. Hence we have

$$u_k(t) = \Phi^{-1}(\pi_k(t)). \tag{5}$$

where Φ is the standard normal cumulative distribution function.

Conditional on a particular value x of X, the conditional risk-neutral default probabilities are defined as

$$\pi_k(t, x) \equiv \mathbb{P}(\tau_k \leq t \mid X = x) = \mathbb{P}(U_k \leq u_k(t) \mid X = x) \tag{6}$$

[1]Usually, the risk-neutral default probabilities are implied from the market price of defaultable bonds or credit default swaps. For more details, see [10].

[2]Generally, the correlation factors β_k are calculated from the correlation matrix by principal component analysis as proposed by Andersen, Sidenius, and Basu [3]. The correlation matrix is usually estimated from the historical correlations of asset returns.

Substituting (4) and (5) into (6), we have

$$\pi_k(t,x) = \mathbb{P}\big[\beta_k x + \sigma_k \varepsilon_k \le \Phi^{-1}(\pi_k(t))\big] = \Phi\left[\frac{\Phi^{-1}(\pi_k(t)) - \beta_k x}{\sigma_k}\right]$$

In this framework, the default events of the names are assumed to be conditionally independent. Thus, the problem of correlated names is reduced to the problem of independent names. The mean tranche losses $\mathbb{E}\mathcal{L}_i$ satisfies

$$\mathbb{E}\mathcal{L}_i = \int_{-\infty}^{\infty} \mathbb{E}_x[\mathcal{L}_i] d\Phi(x) \tag{7}$$

where $\mathbb{E}_x[\mathcal{L}_i] = \mathbb{E}_x[\min(S, (\hat{L}_i - a)^+)]$ is the expectation of $\mathcal{L}_i$ conditional on a specified value x of X; and $\hat{L}_i = \sum_{k=1}^{K} N_k \mathbf{1}_{\{u_k(T)<U_k\le u_k(T_i)\}}$, where $\mathbf{1}_{\{u_k(T)<U_k\le u_k(T_i)\}}$ are mutually independent, conditional on $X = x$. Therefore, if we have the conditional distributions of $\mathbf{1}_{\{u_k(T)<U_k\le u_k(T_i)\}}$, the conditional distributions of $\hat{L}_i$ can be computed easily, as can $\mathbb{E}_x[\mathcal{L}_i]$. To approximate the integral (7), we use a quadrature rule, such as the Gaussian-Legendre rule or the Gaussian-Hermite rule. Thus, the integral (7) reduces to

$$\mathbb{E}\mathcal{L}_i \approx \sum_{m=1}^{M} w_m \mathbb{E}_{x_m}[\min(S, (\hat{L}_i - a)^+)]$$

where the w_m and x_m are the quadrature weights and nodes, respectively. Therefore, the main challenge in CDO pricing lies in the evaluation of the distribution of $\hat{L}_i$, conditional on a given value x of X.

4.2 Conditional forward default probabilities

Conditional on a given x, to compute the distributions of $\hat{L}_i$, we need to specify the distributions of $\mathbf{1}_{\{T<\tau_k\le T_i\}}$, which are equal to the conditional distributions of $\mathbf{1}_{\{u_k(T)<U_k\le u_k(T_i)\}}$. To this end, we introduce conditional forward default probabilities

$$\hat{\pi}_k(t,x) = \pi_k(t,x) - \pi_k(T,x), \quad \text{for } t \ge T \tag{8}$$

so that the conditional distributions of $\mathbf{1}_{\{T<\tau_k\le T_i\}}$ satisfy

$$\mathbb{P}_x(\mathbf{1}_{\{T<\tau_k\le T_i\}}=1)=\hat{\pi}_k(T_i,x)$$
$$\mathbb{P}_x(\mathbf{1}_{\{T<\tau_k\le T_i\}}=0)=1-\hat{\pi}_k(T_i,x)$$

where $\mathbb{P}_x$ is the probability conditional on $X=x$. Armed with the conditional forward default probabilities, the conditional distribution of $\hat{L}_i$ for a forward starting CDO can be computed using the methods developed for synthetic CDOs.

5 Valuation methods for synthetic CDOs

Based on the conditionally independent framework, researchers have developed many methods to evaluate the conditional loss distribution for synthetic CDOs. There are generally two kinds of approaches: the first one computes the conditional loss distribution exactly by a recursive relationship or the convolution technique, e.g., Andersen, Sidenius, and Basu [3], Hull and White [5], Laurent and Gregory [9], Jackson, Kreinin, and Ma [8]; the second approach computes the conditional loss distribution approximately by, for example, the normal power or compound Poisson approximations, e.g., De Prisco, Iscoe, and Kreinin [12] and Jackson, Kreinin, and Ma [8]. Here we review one of the exact methods – JKM proposed by Jackson, Kreinin, and Ma [8] – and employ it to solve our numerical examples in the next section. Other methods for pricing synthetic CDOs are equally applicable.

A homogeneous pool has identical recovery-adjusted notional values, denoted by N_1, but different default probabilities and correlation factors. Hence, conditional on a specified common factor x, the pool losses satisfy

$$\hat{L}_i=\sum_{k=1}^{K}N_k\mathbf{1}_{\{T<\tau_k\le T_i\}}=N_1\sum_{k=1}^{K}\mathbf{1}_{\{T<\tau_k\le T_i\}}$$

Therefore, we can compute the conditional distribution of $\hat{L}_i$ through computing the conditional distribution of the number of defaults $\sum_{k=1}^{K} \mathbf{1}_{\{T<\tau_k\le T_i\}}$.

Suppose the conditional distribution of the number of defaults over a specified time horizon $[T, T_i]$ in a homogeneous pool with k names is already known. Denote it by $V_k = (p_{k,k}, p_{k,k-1}, \ldots, p_{k,0})^T$, where $p_{k,j} = \mathbb{P}_x(\sum_{l=1}^{k} \mathbf{1}_{\{T<\tau_l\le T_i\}} = j)$. The conditional distribution of the number of defaults in a homogeneous pool containing these first k names plus the $(k+1)$st name with conditional forward default probability $Q_{k+1} = \hat{\pi}_{k+1}(T_i, x)$ satisfies

$$V_{k+1} = \begin{pmatrix} p_{k+1,k+1} \\ p_{k+1,k} \\ \vdots \\ p_{k+1,1} \\ p_{k+1,0} \end{pmatrix} = \begin{pmatrix} V_k & 0 \\ 0 & V_k \end{pmatrix} \begin{pmatrix} Q_{k+1} \\ 1-Q_{k+1} \end{pmatrix}$$

Using this relationship, V_K can be computed after $K-1$ iterations with initial value $V_1 = (p_{1,1}, p_{1,0})^T = (Q_1, 1-Q_1)^T$. The method has been proved numerically stable by Jackson, Kreinin, and Ma [8].

An inhomogeneous pool, which has different recovery-adjusted notional values, different default probabilities, and different correlation factors, can be divided into I small homogeneous pools with notionals $N_1, N_2, \ldots, N_I$. The conditional loss distribution for the ith group can be computed using the method above. We denote it by $(p_{i,0}, \ldots, p_{i,d_i})$, where d_i is the maximum number of defaults in group i. Suppose the conditional loss distribution of the first i groups is available. Denote it by $(p_0^{(i)}, \ldots, p_{S_i}^{(i)})$, where $p_s^{(i)}$ is the probability that s units of the pool default out of the first i groups, for $s = 0, 1, \ldots, S_i = \sum_{j=1}^{i} d_j N_j$. The conditional loss distribution of the pool containing these first i groups plus the $(i+1)$st group satisfies

$$p_s^{(i+1)} = \sum_{\substack{l \in \{0,\ldots,S_i\} \\ (s-l)/N_{i+1} \in \{0,\ldots,d_{i+1}\}}} p_l^{(i)} \cdot p_{i+1,(s-l)/N_{i+1}}, \quad \text{for } s = 0, 1, \ldots, S_{i+1} = S_i + d_{i+1} N_{i+1}$$

To start the iteration, we need to initialize the conditional loss distribution of the first group $(p_0^{(1)}, p_1^{(i)}, \ldots, p_{d_1 N_1}^{(i)})$ by setting

$$p_s^{(1)} = \begin{cases} p_{1,s/N_1}, & s/N_1 \in \{0, 1, \ldots, d_1\} \\ 0, & \text{otherwise} \end{cases}$$

6 Numerical examples

We compare the results generated by the Monte Carlo method to those obtained by our analytical method. The numerical experiments are based on two forward starting CDOs: one is a homogeneous pool; the other is an inhomogeneous pool. The contracts are 5-year CDOs starting one year later with annual premium payments, i.e., $T = T_0 = 1$, $T_1 = 2$, ..., $T_5 = 6 = T^*$. The CDO tranche structures are described in Table 1. The continuously compounded interest rates are listed in Table 2. The recovery rate of the instruments in the pool is 40%. The risk-neutral cumulative default probabilities for two credit ratings are listed in Table 3. The pool structure of the inhomogeneous CDO is defined in Table 4, while the homogenous pool has the same structure except that the notional values are 30 for all names.

Tranche	Attachment	Detachment
Super-senior	12.1%	100%
Senior	6.1%	12.1%
Mezzanine	4%	6.1%
Junior	3%	4%
Equity	0%	3%

Table 1: CDO tranche structure

time	1Y	2Y	3Y	4Y	5Y	6Y
Rate	0.046	0.050	0.056	0.058	0.060	0.061

Table 2: Risk-free interest rate curve

We employ Latin hypercube sampling to accelerate the Monte Carlo simulation. Each experiment consists of 100,000 trials, and 100 runs (with different seeds) of each experiment

Credit rating	Time					
	1Y	2Y	3Y	4Y	5Y	6Y
Baa2	0.0007	0.0030	0.0068	0.0119	0.0182	0.0223
Baa3	0.0044	0.0102	0.0175	0.0266	0.0372	0.0485

Table 3: Risk-neutral cumulative default probabilities

Notional	Credit Rating	β_k	Quantity
10	Baa2	0.5	5
10	Baa3	0.5	2
10	Baa2	0.6	5
10	Baa3	0.6	5
10	Baa3	0.7	4
10	Baa3	0.8	4
20	Baa3	0.5	7
20	Baa2	0.6	10
20	Baa3	0.6	8
30	Baa2	0.5	15
30	Baa3	0.5	10
60	Baa2	0.4	10
60	Baa2	0.4	8
60	Baa3	0.5	7

Table 4: Inhomogeneous pool structure

are made. Base on the results of these 100 experiments, we calculate the mean and the 95% non-parametric confidence interval. Table 5 presents the risk premiums for these two forward starting CDOs. The results demonstrate that our method is accurate for the valuation of forward starting CDOs.

The performance of the two methods are compared in Matlab 7 on a Celeron 2.6GHZ PC with 256M RAM. For the homogeneous forward starting CDO, the running time of one Monte Carlo experiment with 100,000 trials is about 14 times that used by our method; for the inhomogeneous forward starting CDO, the Monte Carlo method uses about 6 times the CPU time used by our method.

Pool	Tranche	Monte Carlo	95% CI	Analytic
Homogeneous	Equity	1151.57	[1148.56, 1154.66]	1151.79
	Junior	380.61	[377.96, 383.35]	380.82
	Mezzanine	232.47	[230.45, 234.18]	232.57
	Senior	80.39	[79.52, 81.30]	80.40
	Super-Senior	1.24	[1.18, 1.29]	1.24
Inhomogeneous	Equity	1208.57	[1204.12, 1212.46]	1208.66
	Junior	406.37	[403.53, 409.47]	406.30
	Mezzanine	228.83	[227.06, 230.71]	228.76
	Senior	68.02	[66.92, 68.95]	67.95
	Super-Senior	0.77	[0.72, 0.81]	0.76

Table 5: Tranche premiums (bps)

7 Extensions of the method

Besides standard forward starting CDOs, our method works well for the exotic forward starting contracts with prematurity underlying assets. In the normal contract, we assume that all underlying assets mature after T^*; in the prematurity contract, we allow some instruments to mature before T^*.

Suppose name j's maturity t_j satisfies $T < t_j < T^*$. Before t_j, the contract is the same as the normal one. Therefore, conditional on $X = x$, the computation of $\hat{L}_i$'s distribution is the same. After t_j, we still have $\hat{L}_i = \sum_{k=1}^{K} N_k \mathbf{1}_{\{T<\tau_k \le T_i\}}$, but we need to modify the conditional distribution of $\mathbf{1}_{\{T<\tau_j \le T_i\}}$ to reflect the prematurity of name j. After maturity, name j will never default, so its default probability will never change. Therefore, the conditional distribution of $\mathbf{1}_{\{T<\tau_j \le T_i\}}$ for $t_j \le T_i$ satisfies

$$\mathbb{P}_x(\mathbf{1}_{\{T<\tau_k \le T_i\}} = 1) = \mathbb{P}_x(\mathbf{1}_{\{T<\tau_k \le t_j\}} = 1)$$
$$\mathbb{P}_x(\mathbf{1}_{\{T<\tau_k \le T_i\}} = 0) = \mathbb{P}_x(\mathbf{1}_{\{T<\tau_k \le t_j\}} = 0)$$

The modification can be realized by changing the conditional forward default probabilities of name j in (8) to

$$\hat{\pi}_j(t,x) = \begin{cases} \pi_j(t,x) - \pi_j(T,x), & t \le t_j \\ \pi_j(t_j,x) - \pi_j(T,x), & t > t_j \end{cases}$$

8 Conclusions

In this paper, we study a valuation method for forward starting CDOs based on the Gaussian factor copula model. The effective pool losses in forward starting CDOs are converted to the pool losses in the equivalent synthetic CDOs. Based on the conditional independence framework, computing the conditional distribution of the effective pool losses in forward starting CDOs is converted to computing the conditional pool loss distribution in the equivalent synthetic CDO. The latter problem is well studied by researchers. We apply one of the loss distribution evaluation methods in our numerical examples. The numerical results for both homogeneous and inhomogeneous forward starting CDOs demonstrate the accuracy and efficiency of our method. The method can also be applied to the prematurity problem by modifying the conditional forward default probabilities.

Our method cannot be applied directly to all other classes of credit derivatives. For example, it does not apply to forward starting basket default swaps (BDS), because, in a forward starting BDS, the distribution of the terminal default time depends on the pool losses before T. Our method works for forward starting CDOs, because, the pool losses before T do not influence the effective pool losses, and hence the tranche losses.

Future work includes calibrating the correlation factors from market quotes and pricing options on a CDO tranche.

Acknowledgments

The authors thank Alex Kreinin for proposing this interesting topic, for several very informative discussions that lead to the development of our method, and for his helpful comments on earlier drafts of the paper.

References

[1] L. Andersen. Portfolio losses in factor models: Term structures and intertemporal loss dependence. Working paper, September 2006.

[2] L. Andersen and J. Sidenius. Extensions to the Gaussian copula: Random recovery and random factor loadings. *Journal of Credit Risk*, 1:29–70, 2004.

[3] L. Andersen, J. Sidenius, and S. Basu. All your hedges in one basket. *Risk*, 16(11):67–72, 2003.

[4] N. Bennani. The forward loss model: A dynamic term structure approach for the pricing of portfolio credit derivatives. Working paper, 2005.

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

[6] J. Hull and A. White. Dynamic models of portfolio credit risk: A simplified approach. Working paper, December 2006.

[7] J. Hull and A. White. Forwards and European options on CDO tranches. Working paper, December 2006.

[8] K. Jackson, A. Kreinin, and X. Ma. Loss distribution evaluation for synthetic CDOs. Working paper, December 2005.

[9] J. Laurent and J. Gregory. Basket default swaps, CDOs and factor copulas. *Journal of Risk*, 7:103–122, 2005.

[10] D. Li. Constructing a credit curve. *Credit Risk*, pages 40–44, 1998.

[11] D. Li. On default correlation: A copula approach. *Journal of Fixed Income*, 9:43–54, 2000.

[12] B. De Prisco, I. Iscoe, and A. Kreinin. Loss in translation. *Risk*, June 2005.

[13] B. De Prisco and A. Kreinin. Valuation of forward-starting CDOs. Working paper, September 2006.

[14] P. Schönbucher. Portfolio losses and the term structure of loss transition rates: a new methodology for the pricing of portfolio credit derivatives. Working paper, September 2005.

[15] J. Sidenius, V. Piterbarg, and L. Andersen. A new framework for dynamic credit portfolio loss modeling. Working paper, 2005.

[16] O. Vasicek. Probability of loss distribution. Technical report, KMV Corporation, 1987.

[17] M. Walker. CDO models – towards the next generation: Incomplete markets and term structure. Working paper, 2006.