2 June 2004

Extensions to the Gaussian Copula: Random Recovery and Random Factor Loadings

by

Leif Andersen* **and Jakob Sidenius**†
Bank of America

Abstract: This paper presents two new models of portfolio default loss that extend the standard Gaussian copula model, yet preserve tractability and computational efficiency. In one extension, we randomize recovery rates, explicitly allowing for the empirically well-established effect of inverse correlation between recovery rates and default frequencies. In another extension, we build into the model random systematic factor loadings, effectively allowing default correlations to be higher in bear markets than in bull markets. In both extensions, special cases of the models are shown to be as tractable as the Gaussian copula model and to allow efficient calibration to market credit spreads. We demonstrate that the models—even in their simplest versions—can generate highly significant pricing effects such as fat tails and a correlation "skew" in synthetic CDO tranche prices. When properly calibrated, the skew effect of random recovery is quite minor, but the extension with random factor loadings can produce correlation skews similar to the steep skews observed in the market. We briefly discuss two alternative skew models, one based on the Marshall-Olkin copula, the other on a spread-dependent correlation specification for the Gaussian copula.

1 Introduction

The valuation of synthetic CDO tranches requires the modeling of joint default losses in portfolios of credit default swaps. A common technique for the description of co-dependence of defaults is to specify a copula that governs the joint distribution of default times. While many copulas have been proposed for this purpose (see e.g. Schonbucher (2001)), the market standard copula is the Gaussian copula Li (2000) and its convenient Student t extension Frey and McNeil (2003). In this model each underlying credit risk is associated with a Gaussian (or Student t) random default variable whose value determines the time of default. The map from default variables to default times can be chosen to reproduce observed default swap quotes, allowing an interpretation of the default variables as firm asset values; see Merton (1974). The random default variables are assumed correlated with a constant correlation matrix, the entries of which can conveniently be

*leif.andersen@bofasecurities.com
†jakob.sidenius@bankofamerica.com

thought of as correlations of firm asset returns. If the correlation matrix can be represented by a low number of systematic factors, portfolio loss distributions—and thereby CDO tranche prices—can be computed efficiently using either recursions or Fourier methods; see Andersen et al. (2003) or Gregory and Laurent (2003).

While convenient and intuitive, the Gaussian copula has a number of obvious shortcomings as a model of the real world. For instance, standard implementations of the model make the assumption that recovery rates on default are known firm-specific constants. In reality, however, forecasting recovery rates of companies in default is notoriously difficult. While a number of modeling frameworks exist (e.g. Gupton and Stein (2002)), the reality is that the certainty with which one can predict recovery at present is not significantly better than when guessing a draw from a uniform distribution. Among the few undisputed facts about recovery rates, however, is the observation that average recovery rates tend to be inversely related to default rates: in a bad year, not only are there many defaults, but recoveries are also low. The empirical literature supporting this observation is large and growing; relevant papers include Acharya et al. (2003), Altman et al. (2003), Frye (2000), Frye (2003), and Hu and Perraudin (2002), to name a few. Altman et al. (2003) reviews much of the empirical research in detail.

Models for stochastic recovery rates have been proposed by a number of authors, see Frye (2000), Jarrow (2001), Jokivuolle and Peura (2000), and Pykhtin (2003). The primary focus of most of the existing literature is on applications to risk management and the management of tail risk. In this paper, however, we shall tailor the modeling of recovery rates to the pricing of CDOs and other structured credit derivatives. This, in turn, imposes a number of constraints. First, precision in the fit to market observables now becomes of prime importance, necessitating the construction of explicit calibration routines. Second, to efficiently risk-manage and hedge CDO positions, tractability and numerical efficiency is critical. Third, given the nature of the underlying collateral (credit default swaps), recovery rates as a percentage of notional are bounded on [0,1] and should be modeled as such[1]. This paper introduces, and embeds into a Gaussian copula, a multi-dimensional factor-type recovery model that satisfies these constraints.

With the introduction of a quoted market on standardized tranches of credit indices, another imperfection of the Gaussian copula has recently become evident: the allocation of value across tranches implied by the model does not match the observable market. This effect is perhaps best discussed in the context of implied "base correlation", defined as the constant correlation (for all pairs of credits) required for a Gaussian copula to match the price of an all-equity CDO tranche with fixed upper detachment level[2]. Graphed against the upper detachment level, the implied base correlations observable in the market tend to form a steeply increasing skew, as illustrated in Figure 1 for 5-year tranches on the I-Boxx index.

[1] In Frye (2000) and Pykhtin (2003), for example, loss amounts are Gaussian and log-normal, respectively, and thus unbounded.

[2] That is, we consider loss tranches spanning the first x% of the total pool notional. Non-equity tranches spanning a percentage loss interval of, say, $[l,u],\ 0 < l < u$, can trivially be priced as differences of two equity trances with $x = u$ and $x = l$, respectively.

Figure 1: Base correlation skew in I-Boxx NA (May 2004)

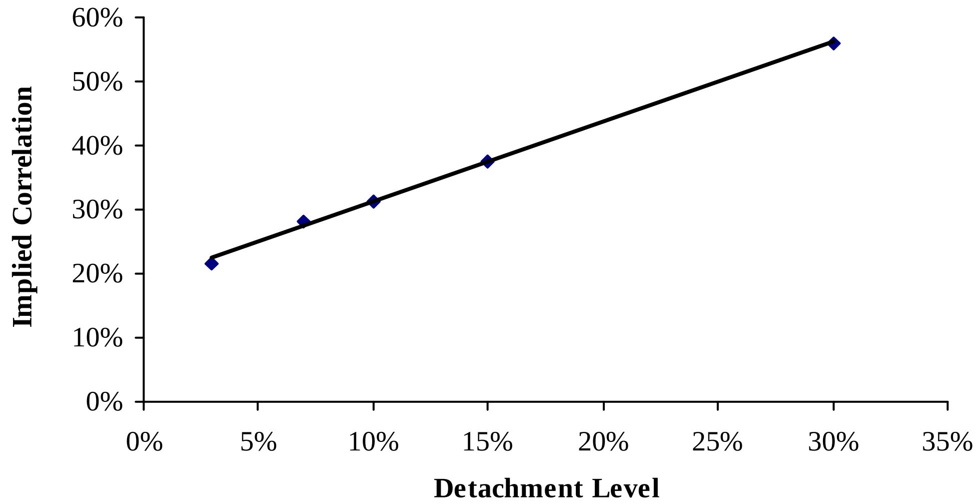

Equipped with interpolation and/or extrapolation rules, one can use graphs such as that on Figure 1 to price arbitrary tranches on standard indices. Apart from the fact that such schemes may be subject to arbitrage[3], this practice has obvious shortcomings. For instance, for customized portfolios with composition and/or maturities different from that of a traded index it is unclear how to use the base correlation information in a systematic way. It is also unclear how to use the market information for the pricing of non-vanilla structured instruments, such as CDOs on portfolios of CDOs ("CDO-squared"). To address these issues and to predict the evolution of the skew with changes in spreads as required for hedge computations, we need a fundamental model mechanism capable of producing the correlation skew.

To understand the demands on a model for the base correlation skew, consider that tranche prices are essentially integrals over the risk-neutral portfolio loss distribution, allowing us to back out the market-implied loss distributions from base correlation skew graphs such as that of Figure 1. As demonstrated schematically in Figure 2, the market loss distribution is more "kinked" than that of a Gaussian copula with empirical correlations, with a fat upper tail and a relatively low probability of generating small losses.

While it is easy to make a Gaussian copula fat-tailed by switching to a Student t copula (or by raising all correlations), it is difficult to do so without increasing too much the probability of generating zero losses[4]. Families of copulas exist that can produce market-similar loss distributions directly, see Section 5, but these tend to be unwieldy numerically

[3] The observed I-Boxx base skew is often (suspiciously) close to a straight line. This can lead to arbitrage for sufficiently high detachment points.

[4] One, somewhat contrived, way of doing this is to simply assume—contrary to empirical evidence—that pairwise asset correlations are suitable functions of credit spreads. We consider this case in Section 5.

Figure 2: Market implied portfolio loss distribution (schematic)

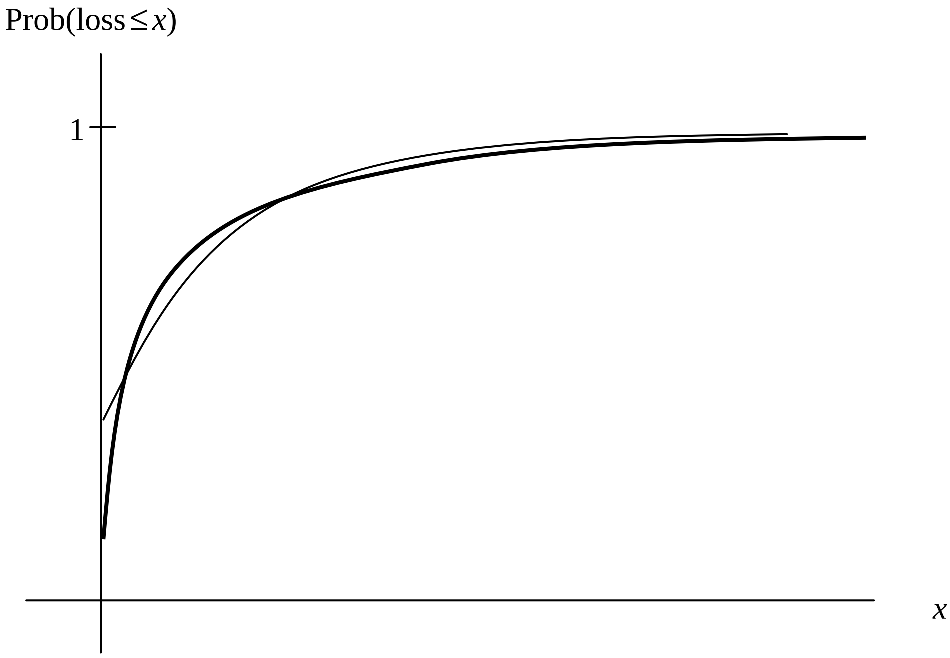

Notes: Schematic portfolio loss distribution in the presence of an upward-sloping base correlation skew (fat curve). The loss distribution for a Gaussian copula is given for reference (thin curve).

and difficult to parameterize and make operational. In this paper, we suggest a straightforward extension of the Gaussian copula where factor loadings are made functions of the systematic factors themselves. Interpreting systematic factors as the "state of the market", we can use this mechanism to mimic the well-known empirical effect that equity (and thereby asset) correlations are higher in a bear market than in a bull market. As we shall show, this mechanism will also induce a strong correlation skew.

The rest of the paper is organized as follows. First, in Section 2, we present a general framework which can accommodate a wide range of factor models. For reference, we introduce the standard Gaussian copula model and briefly review the efficient numerical technique of Andersen et al. (2003) available for this model. In Sections 3 and 4, we present our models for random recovery and random factor loadings, respectively. For both models, our treatment puts particular emphasis on financial engineering aspects such as model implementation and calibration. For intuition and for traditional portfolio risk management, we also give easy-to-use results for large-portfolio limits. In Section 5 we briefly consider some alternatives to the main developments: first, a model based on the Marshall-Olkin copula, and second, the imposition of specific structures on the correlation matrix for the Gaussian copula. Section 6 contains numerical examples and model comparison and we present our conclusions in Section 7. Appendix A collects a number of useful results for Gaussian integrals.

2 General framework

Consider a portfolio of N default-risky obligors. For a fixed time horizon $[0,T]$, we associate the obligors with known default probabilities (eg, inferred from default swap quotes) $p_i(T),\ i=1,\dots,N$, for default before time T. Note that we usually suppress the dependence of default probabilities on the time horizon T.

On default, the obligors generate (possibly random) loss amounts of $l_i,\ i=1,\dots,N$. Define the random default time of obligor i to be τ_i. The total portfolio loss experienced on $[0,T]$ is then

$$L=\sum_{i=1}^{N} l_i 1_{\tau_i\le T}, \tag{1}$$

with expected value

$$\mathsf{E}(L)=\sum_{i=1}^{N} p_i \mathsf{E}\left(l_i|\tau_i<T\right).$$

In our notation the *indicator function*, $1_{\tau_i\le T}$, in (1) is one if the i'th obligor defaults no later than time T and zero otherwise. We assume that each l_i is bounded, ie, $l_i\in[0,l_i^{\max}],\ l_i^{\max}\in\mathbb{R}_+$, such that $0\le L\le\sum l_i^{\max}=L^{\max}$. Often, it is convenient to introduce the notion of a *recovery rate*, R_i, whereby

$$l_i=l_i^{\max}(1-R_i) \tag{2}$$

for (possibly random) $R_1,\dots,R_N$ taking values on $[0,1]$.

To provide further structure to the loss distribution, we assume the existence of continuous random variables $X_1,\dots,X_N$ and fixed thresholds $c_1,\dots,c_N$, such that

$$1_{\tau_i\le T}\equiv 1_{X_i\le c_i},\ i=1,\dots,N.$$

Let the distribution function of X_i be F_i^X. Assuming that F_i^X is invertible, we have

$$c_i=(F_i^X)^{-1}(p_i),\ i=1,\dots,N.$$

Often, it is convenient to think of the X_i and c_i as proxies for firm asset and liquidation values, respectively, as in the structural model of Merton (1974).

Throughout, we make the assumption that there exists a d-dimensional vector of independent conditioning variables $Z=(Z_1,\dots,Z_d)$ such that all components of the augmented vector $(X_1,\dots,X_N,l_1,\dots,l_N)$ are independent when conditioned on Z. Without loss of generality we assume that the $Z_i,\ i=1,\dots,d$, all have zero mean and unit variance.

We shall refer to the vector Z as the *systematic factors* of our model, and write $p_i(Z) = \mathsf{E}(1_{X_i \le c_i}|Z) = \mathsf{Prob}(X_i \le c_i|Z)$ and $l_i(Z) = \mathsf{E}(l_i|Z)$. By conditional independence,

$$\mathsf{E}(L) = \mathsf{E}\left(\sum_{i=1}^{N} p_i(Z) l_i(Z)\right).$$

While Z may simply be an abstract set of latent variables, it is often useful to think of Z as representing the industry or economy data that characterize the systematic credit risk of the portfolio. Existence of a low-dimensional conditioning variable is often critical to the numerical efficiency of a model. In particular, the portfolio loss distribution can be efficiently computed with the techniques developed in Andersen et al. (2003) when the dimension of the conditioning variable is low. As shown in, eg, Andersen et al. (2003), knowledge of the portfolio loss distribution for a grid of horizon times suffices for the purpose of pricing CDO tranches.

For practical purposes it is crucial that the market quotes of default swaps be reproduced. The valuation of the default swap coupon leg is straightforward and standard given the default probabilities for all times to default swap maturity, but the valuation of the default leg needs to be revisited for the general model framework considered here.

Recall that the default leg of a credit default swap pays $(1-R_i)$ at the time of default if and only if default happens before swap maturity. If $P(0,t)$ denotes the risk-free discount factor for time t we have[5]

$$\begin{aligned} \mathsf{PV}_{i,\text{default}}(0) &= \mathsf{E}\left[\int_0^T (1-R_i)P(0,t)1_{\tau_i \in [t,t+dt]}\right] \\ &= \int_0^T P(0,t)\mathsf{E}\left[(1-R_i)\middle|\, \tau_i = t\right] p_i'(t)dt, \end{aligned} \tag{3}$$

where the prime denote differentiation. In the second equation we used the definition of conditional expectation to factorize the expectation.

From (3) we see that the valuation of a default swap requires (only) the expectation of recovery *conditioned on default*. Once this is given the default swap value can be computed as usual from default probabilities. As a consequence, the default probabilities may be inferred by standard methods such as, eg, “bootstrapping”, from default swap quotes.

For the special case where recovery rates are not random, ie, each R_i is just a number in $[0,1]$, the requisite recovery expectations require no calculation, but in models with random recovery it is natural to let the conditional recovery expectation be given exogenously. This not only makes calibration simple, it is also avoids the possibly nonsensical and certainly non-intuitive concept of “recovery without default”.

Remark. As just demonstrated, the valuation of default swaps under random recovery requires the knowledge of the expectation of recovery conditioned on default taking place

[5] Here we make the standard assumption of independence between risk-free discounting rates and default.

at a given time. Sometimes we need or are given the expectation conditioned instead on default taking place *before* a certain time. If

$$\begin{aligned} f_i(t) &:= \mathsf{E}\left[(1-R_i)\middle|\,\tau_i \in [t, t+dt]\right]; \\ F_i(t) &:= \mathsf{E}\left[(1-R_i)|\,\tau_i \le t\right], \end{aligned}$$

then we note the following general relationships[6]

$$F_i(T)p_i(T) = \int_0^T f_i(t)p_i'(t)dt \tag{4a}$$

$$f_i(t) = F_i(t) + F_i'(t)\frac{p_i(t)}{p_i'(t)}. \tag{4b}$$

We note that the special case of horizon-independent expectations

$$f_i(t) \equiv F_i(t) \equiv \mathrm{const}_i, \text{ for all } t,$$

is a solution to (4).

2.1 Gaussian copula model

In the general framework set out above the Gaussian copula model is described by letting the factors Z be Gaussian and by assuming the following relationships

$$\left.\begin{aligned} X_i &= a_i \cdot Z + \sqrt{1-||a_i||^2}\varepsilon_i \\ l_i &= l_i^{\max}(1-R_i) \end{aligned}\right\} i = 1,\ldots,N, \tag{5}$$

where $R_i \in [0,1]$, $i = 1,\ldots,N$, are non-random and ε_i, $i = 1,\ldots,N$, define an iid sequence of zero-mean, unit-variance Gaussian variables independent of Z. The *factor loadings* a_i are d-dimensional vectors of length less than one. We note that Andersen et al. (2003) give an efficient algorithm for computing the factor loadings by optimally approximating an arbitrary correlation structure with a low-dimensional factor structure. Clearly, in (5) the conditional independence of the X_i's and therefore of defaults is manifest.

In Andersen et al. (2003) quasi-analytical techniques were developed for studying this model starting from the observation that since defaults are independent when conditioned on factors, we can build the conditional distribution of losses over the time interval $[0,T]$ by a simple recursion. To do this we require an arbitrary *loss unit*, u, to be defined such that the loss amounts l_i (and thereby all portfolio losses) can be well-approximated by integer multiples of u, say $l_i = k_i u$. Now let L_n, $1 \le n < N$, be the loss (measured in loss units) over $[0,T]$ in the subportfolio consisting of the first n obligors in some arbitrary

[6] These follow from rewriting $F_i(T)p_i(T)$ as $\mathsf{E}[(1-R_i)1_{\tau_i \le T}]$ which is clearly identical to $\int_0^T \mathsf{E}\left[(1-R_i)1_{\tau_i \in [t,t+dt]}\right]$.

ordering. Evidently we have the following recursive relation between the conditional distributions of L_n and L_{n+1}:

$$\begin{aligned}\mathsf{Prob}(L_{n+1}=K|Z) &= p_{n+1}(Z)\mathsf{Prob}(L_n=K-k_{n+1}|Z)\\ &\quad +(1-p_{n+1}(Z))\mathsf{Prob}(L_n=K|Z)\end{aligned} \tag{6}$$

and we may use this to compute the loss distribution of the portfolio from the boundary case of the empty portfolio for which we have $\mathsf{Prob}(L_0=K|Z)=\delta_{K,0}$. For sensitivity computations it is useful to remark that clearly the portfolio loss distribution cannot depend on the ordering of the obligors.

In general the distribution of a sum of independent random variables is given by the convolution of the distributions of each random variable; the recursion method can be seen as a particular way of performing the (discrete) convolution of the single-obligor loss distributions to obtain the portfolio loss distribution. An alternative is the use of Fourier transformation techniques (see Gregory and Laurent (2003)) although this was found in Andersen et al. (2003) to be less computationally efficient for the Gaussian copula model.

As shown in Andersen et al. (2003) sensitivities of expectations over the portfolio loss distribution can be efficiently computed using (6). Let $X(L)$ be some function of the portfolio loss and consider its sensitivities to default probabilities[7], ie, $\partial\mathsf{E}(X)/\partial p_i$. These can be computed as

$$\begin{aligned}\frac{\partial\mathsf{E}(X)}{\partial p_i} &= \int\frac{\partial\mathsf{E}(X|Z)}{\partial p_i}d\Phi_d(Z)\\ &= \int\frac{dp_i(Z)}{dp_i}\frac{\partial\mathsf{E}(X|Z)}{\partial p_i(Z)}d\Phi_d(Z)\\ &= \int\frac{dp_i(Z)}{dc_i}\left(\frac{dp_i}{dc_i}\right)^{-1}\frac{\partial\mathsf{E}(X|Z)}{\partial p_i(Z)}d\Phi_d(Z),\end{aligned} \tag{7}$$

where Φ_d is the d-dimensional cumulative Gaussian distribution function with all correlations zero. Here the first two factors of the integrand are easy to compute analytically and the last factor follows from (6):

$$\begin{aligned}\frac{\partial\mathsf{E}(X|Z)}{\partial p_i(Z)} &= \sum_K X(K)\frac{\partial\mathsf{Prob}(L_N=K|Z)}{\partial p_i(Z)}\\ &= \sum_K X(K)\left[\mathsf{Prob}(L_{N-1}^{(i)}=K-k_i|Z)-\mathsf{Prob}(L_{N-1}^{(i)}=K|Z)\right].\end{aligned}$$

Here $L_{N-1}^{(i)}$ is the loss of the portfolio with the i'th obligor removed and we may obtain the distribution of $L_{N-1}^{(i)}$ by solving (6) recursively for given portfolio loss distribution.

[7] These are related to *hazard rate* sensitivities and *credit spread deltas* by simple Jacobian factors.

3 Random recovery

In this section we consider a special case of the general framework of the previous section in which the loss amounts—or equivalently, the recovery rates—associated with defaults are random. To specify the model let $a_i,\ b_i,\ i=1,\ldots,N$, be d-dimensional vectors with a_i non-negative and of less than unit length. Also, let $\varepsilon_i,\ i=1,...,N$, define a sequence of independent, zero-mean, unit variance random variables independent of Z. Finally, let $\xi_i,\ i=1,...,N$, be a sequence of independent zero-mean random variables with variances $\sigma^2_{\xi_i}$, independent of Z and the ε_i's.

Now consider the model

$$\left.\begin{array}{rcl} X_i &=& a_i\cdot Z+\sqrt{1-\|a_i\|^2}\varepsilon_i \\ l_i &=& l_i^{\max}\left(1-C_i(\mu_i+b_i\cdot Z+\xi_i)\right) \end{array}\right\}\ i=1,\ldots,N, \tag{8}$$

where the $C_i:\mathbb{R}\to[0,1],\ i=1,\ldots,N$, are arbitrary mapping functions, and the $\mu_i,\ i=1,\ldots,N$, are constants. Note that in this model the i'th recovery has a systematic term $(b_i\cdot Z)$ as well as an idiosyncratic one (ξ_i). The interdependence between default of the i'th obligor and recovery on the j'th obligor is controlled by $a_i\cdot b_j$; in particular, $a_i\cdot b_j=0$ implies independence.

We note that the conditional default probabilities are straightforward:

Proposition 1 If F_i^ε is the distribution functions of ε_i, then

$$p_i(Z)=F_i^\varepsilon\left(\frac{(F_i^X)^{-1}(p_i)-a_i\cdot Z}{\sqrt{1-\|a_i\|^2}}\right). \tag{9}$$

Proof: Simply note that

$$\begin{aligned} p_i(Z) &= \mathsf{Prob}(X_i\le c_i|Z)=\mathsf{Prob}\left(\varepsilon_i\le\frac{c_i-a_i\cdot Z}{\sqrt{1-\|a_i\|^2}}\middle| Z\right) \\ &= F_i^\varepsilon\left(\frac{c_i-a_i\cdot Z}{\sqrt{1-\|a_i\|^2}}\right)=F_i^\varepsilon\left(\frac{(F_i^X)^{-1}(p_i)-a_i\cdot Z}{\sqrt{1-\|a_i\|^2}}\right). \end{aligned}$$ ∎

Let us consider characterizing the correlation structure of the setup in (8). Given the non-linear dependence of recovery rates on the stochastic drivers, we wish to focus on rank correlation, rather than the usual linear (Pearson) correlation coefficient. Specifically, let $\tau(x,y)$ be the Kendall's Tau (rank correlation) of random variables x and y, such that, by definition, $\tau(x,y)=\mathsf{Prob}\left((x-\tilde{x})(y-\tilde{y})>0\right)-\mathsf{Prob}\left((x-\tilde{x})(y-\tilde{y})<0\right)$, where $(\tilde{x},\tilde{y})$ is an independent copy of (x,y). For the practically important class of elliptical distributions, the following characterization is possible.

Proposition 2 *Consider (8) with the recovery mapping functions C strictly increasing. Define* $Y_i = \mu_i + b_i \cdot Z + \xi_i$ *and* $\sigma_i = \sqrt{b_i \cdot b_i + \sigma^2_{\xi_i}}, i = 1, \ldots, N$. *If* $(X_1, \ldots, X_N, Y_1, \ldots, Y_N)$ *belong to the class of continuous elliptical distributions, then*

$$\tau(X_i, X_j) = 2\pi^{-1} \sin^{-1}(a_i \cdot a_j),\ i \neq j \tag{10a}$$
$$\tau(R_i, R_j) = \tau(Y_i, Y_j) = 2\pi^{-1} \sin^{-1}\left(b_i \cdot b_j/(\sigma_i \sigma_j)\right),\ i \neq j \tag{10b}$$
$$\tau(R_i, X_j) = \tau(Y_i, X_j) = 2\pi^{-1} \sin^{-1}(b_i \cdot a_j/\sigma_i). \tag{10c}$$

Proof: All three statements are derived the same way, so we just concentrate on (10c). First, notice that the normalization assumption on the X_i implies a linear correlation coefficient between Y_i and X_j of just $b_i \cdot a_j/\sigma_i$ $(i \neq j)$. The equality

$$\tau(Y_i, X_j) = 2\pi^{-1} \sin^{-1}(b_i \cdot a_j/\sigma_i)$$

is a consequence of a general result for elliptical distributions in Lindskog et al. (2003); the equation $\tau(R_i, X_j) = \tau(Y_i, X_j)$ follows from the fact that Kendall's Tau is invariant under strictly increasing transformations (see e.g. Embrechts et al. (2001)). ∎

The class of elliptical distributions is characterized in detail in Embrechts et al. (2001) and includes the Gaussian and Student's *t* distributions, among many others. We note that if recovery mapping functions are strictly decreasing, the sign on statement c) should simply be reversed.

3.1 Portfolio loss distribution

The computation of the portfolio loss distribution in this model is more involved than for the Gaussian copula model. This is because the conditional loss distributions of individual obligors are in general—due to the idiosyncratic randomness of recovery—not just two-point distributions. However, since losses from individual obligors are independent conditional on the factors, we can still compute the conditional portfolio loss distribution as the convolution product of all of the single-name loss distributions.

To do this we assume that a positive loss unit u is given such that we can approximate the distribution of each l_i by the discrete distribution of the integer valued k_i, $l_i = k_i u$. Since $l_i \in [0, l_i^{\max}]$ the support of the discrete distribution must lie on the non-negative integers up to some finite value $k_i^{\max}$ $(\approx l_i^{\max}/u$, but depending on the details of the chosen discretization). What we need in building the portfolio loss distribution is a discretization of the quantity $\ell_i := l_i \mathbf{1}_{\tau_i \leq T}$, which describes the distribution of losses due to default before time T. Since recovery and default are independent when we condition on Z, we have

$$\begin{aligned} \mathsf{Prob}(\ell_i \leq x|Z) &= \mathsf{Prob}(l_i \leq x|Z) \cdot \mathsf{Prob}(\tau_i \leq T|Z) + \mathsf{Prob}(\tau_i > T|Z) \\ &= 1 - p_i(Z)(1 - \mathsf{Prob}(l_i \leq x|Z)) \end{aligned}$$

and we can therefore get a discrete distribution for ℓ_i, denoted $P^{(i)}(\cdot|Z)$, from the discrete distribution for l_i. The support of $P^{(i)}(\cdot|Z)$ lies on the non-negative integers up to $k_i^{\max}$.

Now let L denote the total portfolio loss (measured in loss units) and let '$\wedge$' denote the discrete convolution product over the integers. Then the conditional portfolio loss distribution is given by

$$\mathsf{Prob}(L = k|Z) = \left(\bigwedge_{i=1}^{N} P^{(i)}(\cdot|Z) \right)(k),$$

where we note that the convolution product is well-defined since all factors have finite support. Clearly, the conditional portfolio loss distribution also has finite support. In the computation of the convolution product we can make use of standard Fourier transform techniques, but we may also do a recursive computation. Specifically, let $L_n,\ 1 \le n < N$, be the loss (measured in loss units) in the subportfolio consisting of n obligors. Then we have the following recursive relation between the conditional distributions of L_n and L_{n+1}:

$$\mathsf{Prob}(L_{n+1} = K|Z) \quad = \quad \sum_{k=0}^{k_{n+1}^{\max}} \mathsf{Prob}(L_n = K - k|Z) P^{(n+1)}(k). \tag{11}$$

We may use this recursion to compute the conditional loss distribution of the portfolio from the boundary case of the empty portfolio for which we have $\mathsf{Prob}(L_0 = K|Z) = \delta_{K,0}$.

Sensitivities may be computed from (11), but in general the direct recursive solution of (11) for the loss distribution of the portfolio with one obligor removed is numerically unstable. This problem is of a purely numerical nature and may be overcome, eg, by the use of Fourier transform or matrix inversion techniques. Note that the problem exists only for the general case where recovery has an idiosyncratic component. If there is no idiosyncratic component, ie, recovery is completely determined by the conditioning factors, then the conditional loss distributions of individual obligors are two-point distributions and (11) can be solved by a stable recursion.

3.2 Specific model: cumulative Gaussian recovery

We consider the specific model defined from (8) by letting Z, ε_i and ξ_i be Gaussian and by assuming that the mapping functions $C_i(\cdot)$ are all given by the standard cumulative Gaussian distribution: $C_i = \Phi$, for all i.

For later use, and to further characterize the model, we now list some useful results for losses and conditional defaults. First note that the X_i's are also standard Gaussian and that

$$p_i(Z) = \Phi\left(\frac{\Phi^{-1}(p_i) - a_i \cdot Z}{\sqrt{1 - \|a_i\|^2}} \right),\ i = 1, \ldots, N, \tag{12}$$

as a special case of Proposition 1.

Proposition 3 Define $\sigma_i = \sqrt{b_i \cdot b_i + \sigma_{\xi_i}^2}$ *and* $\rho_i = a_i \cdot b_i / \sigma_i$ *and let* $\varphi(x) = (2\pi)^{-\frac{1}{2}} \exp(-\frac{1}{2}x^2)$ *be the standard Gaussian density. Also, let* $\Phi_2(\cdot, \cdot; \rho)$ *be the bivariate cumulative Gaussian distribution function with correlation* ρ*. Then we have*

$$\mathsf{Prob}(R_i \le x) = \Phi\left(\frac{\Phi^{-1}(x) - \mu_i}{\sigma_i}\right); \tag{13a}$$

$$\mathsf{Prob}(R_i \in [x, x+dx]) = dx \cdot \varphi\left(\frac{\Phi^{-1}(x) - \mu_i}{\sigma_i}\right) \left(\varphi\left(\Phi^{-1}(x)\right)\sigma_i\right)^{-1}; \tag{13b}$$

$$\mathsf{E}(R_i) = \Phi\left(\frac{\mu_i}{\sqrt{1+\sigma_i^2}}\right); \tag{13c}$$

$$\mathsf{V}(R_i) = \Phi_2\left(\frac{\mu_i}{\sqrt{1+\sigma_i^2}}, \frac{\mu_i}{\sqrt{1+\sigma_i^2}}; \frac{\sigma_i^2}{1+\sigma_i^2}\right) - \Phi\left(\frac{\mu_i}{\sqrt{1+\sigma_i^2}}\right)^2; \tag{13d}$$

$$l_i(Z) = l_i^{\max}\left(1 - \Phi\left(\frac{\mu_i + b_i \cdot Z}{\sqrt{1+\sigma_{\xi_i}^2}}\right)\right); \tag{13e}$$

$$\mathsf{Prob}(R_i \le x | Z) = \Phi\left(\frac{\Phi^{-1}(x) - \mu_i - \beta_i \cdot Z}{\sigma_{\xi_i}}\right); \tag{13f}$$

$$\mathsf{E}(R_i | \tau_i \le T) = p_i^{-1}\Phi_2\left(\frac{\mu_i}{\sqrt{1+\sigma_i^2}}, c_i; \frac{-\rho_i\sigma_i}{\sqrt{1+\sigma_i^2}}\right). \tag{13g}$$

Proof: Introduce $Y_i = \mu_i + b_i \cdot Z + \xi_i$ which is obviously Gaussian with mean μ_i and standard deviation σ_i. The statement (13a) follows from

$$\mathsf{Prob}(R_i \le x) = \mathsf{Prob}\left(Y_i \le \Phi^{-1}(x)\right) = \mathsf{Prob}\left(\frac{Y_i - \mu_i}{\sigma_i} \le \frac{\Phi^{-1}(x) - \mu_i}{\sigma_i}\right) = \Phi\left(\frac{\Phi(x) - \mu_i}{\sigma_i}\right).$$

Differentiation of this expression results in (13b). As for statements (13c) and (13d), we note that

$$\begin{aligned} \mathsf{E}(R_i) &= \int_{-\infty}^{\infty} \Phi(\mu_i + \sigma_i x)\varphi(x)dx; \\ \mathsf{V}(R_i) &= \int_{-\infty}^{\infty} \Phi(\mu_i + \sigma_i x)^2\varphi(x)dx - \mathsf{E}(R_i)^2. \end{aligned}$$

Expressions for these integrals can be found in Appendix A, yielding the stated results. Statement (13e) follows directly from Proposition 1, and statement (13f) follows from (2) and (13c), after observing that $Y_i | Z$ is Gaussian with mean $\mu_i + b_i \cdot Z$ and variance $\sigma_{\xi_i}^2$. To prove (13g) note that the correlation between Y_i and X_i is ρ_i and that, conditional on

$X_i = x$, the distribution of Y_i is $N\left(\mu_i + \rho_i\sigma_i x, \sigma_i^2\sqrt{1-\rho_i^2}\right)$. It then follows easily from (13c) that

$$\mathsf{E}(R_i|X_i = x) = \Phi\left(\frac{\mu_i + \rho_i\sigma_i x}{\sqrt{1+\sigma_i^2(1-\rho_i^2)}}\right),$$

whereby

$$\mathsf{E}(R_i|\tau_i \leq T) = \mathsf{E}(R_i|X_i \leq c_i) = p_i^{-1}\int_{-\infty}^{c_i}\Phi\left(\frac{\mu_i + \rho_i\sigma_i x}{\sqrt{1+\sigma_i^2(1-\rho_i^2)}}\right)\varphi(x)dx.$$

Using the result (30c) in Appendix A proves the proposition. ∎

Calibration of recovery model parameters could be done by matching moments and rank correlation properties of the recovery distribution to empirical data. The results in Propositions 2 and 3 are useful for this. Importantly, in any calibration scheme we must ensure that the model remains consistent with the default swap market. For this we use (13g) in conjunction with (3) and (4). We note that if, for instance, recovery on default swaps (ie, the expectation of recovery conditioned on default) is assumed time-independent, then it follows from (13g) that one or more of the parameters σ, μ or c must in general be time-dependent. Alternatively, if the parameters are assumed time-independent, the recovery expectation must be time-dependent. In either case it is straightforward to fit the parameters in a bootstrap procedure.

The specific model for random recovery just described is tractable and capable of producing a variety of distributions of recovery rates. The use of a Gaussian distribution may be unconventional, but we note that the more traditional choice of a specification in terms of the beta-distribution leads to a less tractable model. Furthermore, the range of possible recovery densities does not seem to differ much between the two specifications[8]. For illustration, Figure 3 graphs the (unconditional) recovery densities for various parameter specifications for our model. Finally, we note that extensions to a Student's t copula are straightforward.

3.3 Large Portfolio Limit

To gain intuition, we shall now consider some useful limit results arising when the size of the underlying portfolio becomes large, in the sense that $N \to \infty$. Such results are convenient and, as documented in Schonbucher (2001) and Vasicek (2003), large-portfolio limit loss distributions are often remarkably accurate approximations for finite-sized portfolios, especially in the upper tail. Moreover, for empirical estimations done on economy-wide data, the large-portfolio approximation is often natural to consider.

[8]In fact, when the first two recovery rate moments are set to fit the moments of a given beta distribution, the model density is nearly indistinguishable from the beta density.

Figure 3: Recovery rate densities in cumulative Gaussian model

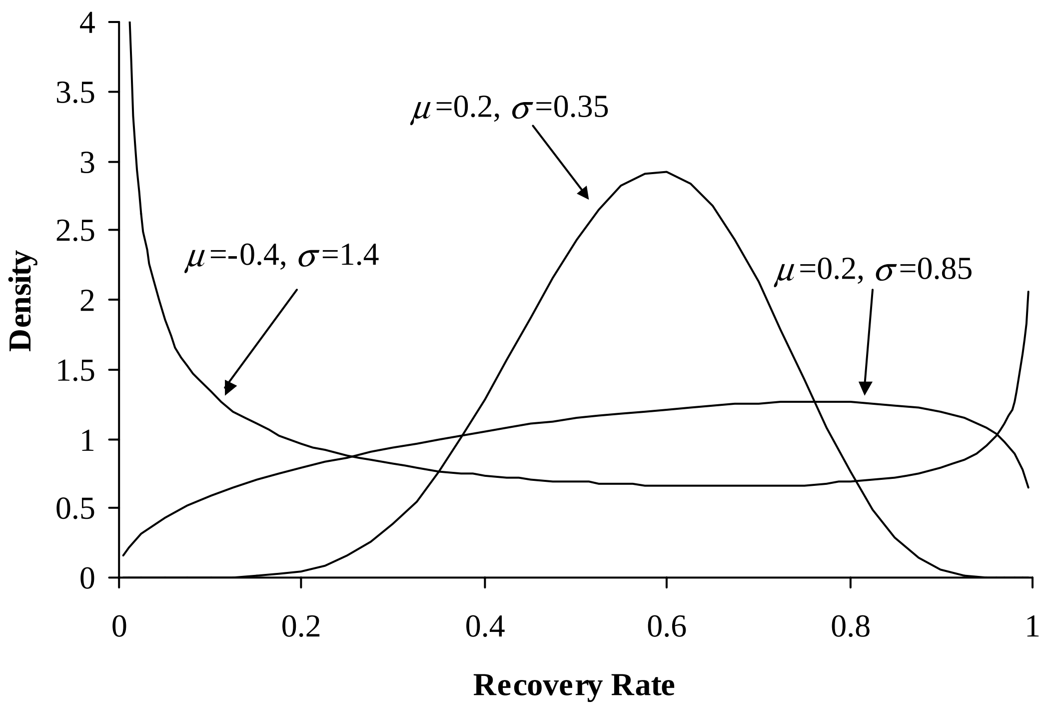

Notes: The densities is given by formula (13b) in Proposition 3. In the graph, $\sigma = \sqrt{\|b\|^2 + \sigma_\xi^2}$ (and we have suppressed the obligor index i).

To ensure the existence of a large-portfolio limit, we shall need to impose some structure on the portfolio:

Assumption 1 *There exists systematic factors Z and a function h such that (pointwise)*

$$\lim_{N\to\infty} N^{-1} \sum_{i=1}^{N} p_i(Z) l_i(Z) = h(Z). \tag{14}$$

Moreover, the composition of the portfolio satisfies

$$\lim_{N\to\infty} N^{-2} \sum_{i=1}^{N} (l_i^{\max})^2 = 0. \tag{15}$$

Proposition 4 *Let the distribution function for the (independent) systematic factors Z be denoted F^Z. Under the terms of Assumption 1, we have*

$$\lim_{N\to\infty} \mathsf{Prob}(L/N \le y) = \mathsf{Prob}\,(h(Z) \le y) = \int_{h(z)\le y} dF^Z(z).$$

Proof:

First notice that

$$\mathsf{E}(L/N|Z) = \mathsf{E}\left(N^{-1}\sum_{i=1}^{N} l_i 1_{\tau_i<T}|Z\right) = N^{-1}\sum_{i=1}^{N}\mathsf{E}(l_i|Z)\mathsf{E}(1_{\tau_i<T}|Z) = N^{-1}\sum_{i=1}^{N} p_i(Z)l_i(Z),$$

where we have relied on conditional independence in the second equality. Conditional on Z, the variance of the average loss L/N is bounded by $N^{-2}\sum_i (l_i^{\max})^2$ which, by assumption, approaches 0 for large N. We can use Chebychev's inequality to show that the random variable $L/N|Z$ converges in probability to $\lim_{N\to\infty} N^{-1}\sum_{i=1}^{N} p_i(Z)l_i(Z) = h(Z)$. This, in turn, implies convergence in distribution:

$$\lim_{N\to\infty} \mathsf{Prob}(L/N \leq y|Z) = 1_{h(Z)\leq y}.$$

Now,

$$\begin{aligned}\lim_{N\to\infty} \mathsf{Prob}(L/N \leq y) &= \lim_{N\to\infty} \mathsf{E}\left(\mathsf{Prob}(L/N \leq y|Z)\right) \\ &= \mathsf{E}\left(\lim_{N\to\infty} \mathsf{Prob}(L/N \leq y|Z)\right) = \mathsf{E}(1_{h(Z)\leq y}),\end{aligned}$$

where interchange of limit and expectation is justified by the fact that $\mathsf{Prob}(L/N \leq y|Z)$ is bounded. ∎

The conditions in Assumption 1 are quite weak. For instance, (15) is satisfied by any portfolio where the $l_i^{\max}$ are bounded for all i. Equation (14) is trivially satisfied by homogeneous portfolios where $p_i(Z) \equiv p(Z)$ and $l_i(Z) \equiv l(Z)$ are independent of i; in this case, obviously $h(Z) = p(Z)l(Z)$. A more general type of portfolio satisfying the condition splits the portfolio into M equal-sized homogeneous subgroups, with M some fixed positive integer.

The result in Proposition 4 simplifies for the popular special case of a one-dimensional systematic factor structure. Specifically, assuming that $d = 1$ and that conditional loss-weighted default probability function $h : \mathbb{R} \to \mathbb{R}$ is strictly decreasing in its argument, we just get

$$\lim_{N\to\infty} \mathsf{Prob}(L/N \geq y) = \mathsf{Prob}\left(h(Z) \geq y\right) = \mathsf{Prob}\left(Z \leq h^{-1}(y)\right) = F^Z\left(h^{-1}(y)\right). \quad (16)$$

We note that the assumption that $h(z)$ is decreasing in z is natural and consistent with our earlier discussion of loss-given-default being positively related to portfolio default frequency. Indeed, interpreting Z as a proxy for the state of the economic environment, it is reasonable to assume that all $p_i(Z)$'s and $l_i(Z)$'s are adversely affected (that is, increase) by a low outcome of Z.

For a concrete model parameterization, consider the following proposition, valid for a homogeneous portfolio (so we omit obligor indices on parameters):

Proposition 5 *Consider an homogeneous portfolio in a one-factor version of the model in Section 3.2, with the identical default probability of all obligors denoted $\bar{p}$. We have*

$$\lim_{N\to\infty} \mathsf{Prob}(L/N \le y) = \Phi\left(-h^{-1}(y)\right),$$

$$h(z)/l^{\max} = \Phi\left(\frac{\Phi^{-1}(\bar{p}) - az}{\sqrt{1\text{-}a^2}}\right)\left(1 - \Phi\left(\frac{\mu + bz}{\sqrt{1+\sigma_\xi^2}}\right)\right)$$

for constants a, μ, b, σ_ξ, where $0 < a < 1$ and $b \ge 0$.

Proof:

The function h is defined in Assumption 1 and can be computed explicitly by combining equation (12) and (13e). For the given restrictions on a and b, it is easy to see that h is strictly decreasing. As the model in Section 3.2 has standard Gaussian systematic factors, the proposition then follows directly from (16). ∎

The result in Proposition 5 is straightforward to implement and gives an attractive alternative to the constant-recovery result in Vasicek (2003) (which is the special case of Proposition 5, for $b = 0$). Determination of the parameters of the recovery model (a, μ, b, σ_ξ) can be done straightforwardly by matching moments and rank correlation, using the expressions in Propositions 1 and 2. For illustration, in Figure 4, we have locked the values of $\mathsf{E}(R)$, $\mathsf{V}(R)$, and $\rho(X_i, X_j) = a^2$ and consider the effect on the large portfolio limit distribution of varying the recovery/default rank correlation $\tau(R_i, X_i) = 2\pi^{-1}\sin^{-1}(ba/\sigma)$, $\sigma^2 \equiv \sigma_\xi^2 + b^2$. As one would expect, increasing the rank correlation will increase the likelihood of having high-loss scenarios, making the upper distribution tail significantly fatter.

To support empirical estimation, we round off this section with a few results on the rank correlation between default frequency and average recovery. Let $\Delta_N := \sum_{i=1}^N 1_{\tau_i \le T}$ be the total number of defaults and define the default frequency $f_N := \Delta_N/N$ and average recovery

$$r_N := \begin{cases} L/\Delta_N & , \quad \Delta_N > 0 \\ 0 & , \quad \Delta_N = 0 \end{cases}$$

Proposition 6 *Consider a homogeneous portfolio with non-zero default probability and let $p_i(Z) \equiv p(Z)$ and $l_i(Z) \equiv l(Z)$ for all i. Then*

$$\tau(f_\infty, r_\infty) = \tau\left(p(Z), l(Z)\right).$$

If the factor setting of equation 8) holds and: a) systematic factors Z are jointly elliptical; b) the recovery mapping function C is strictly increasing; and c) the distribution function of residuals ε_i is strictly increasing; then

$$\tau(f_\infty, l_\infty) = \tau(a \cdot Z, b \cdot Z) = 2\pi^{-1}\sin^{-1}\left(b \cdot a/\sqrt{\|a\|^2 + \|b\|^2}\right).$$

Figure 4: Large-portfolio distribution, cumulative Gaussian recovery model

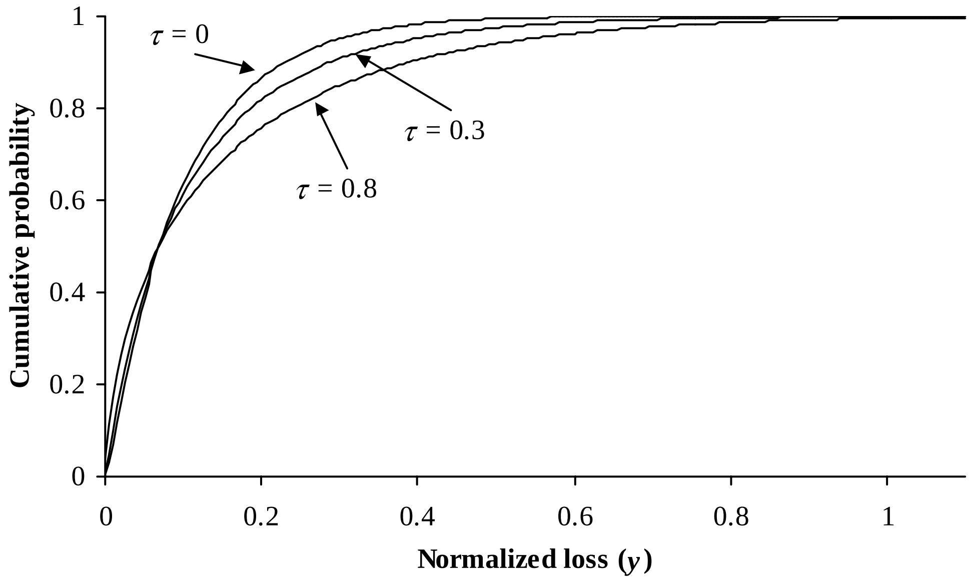

Notes: Homogeneous portfolio in one-factor Gaussian copula, with cumulative Gaussian recovery. The graph shows $\mathsf{Prob}\left([Nl^{\max}(1-\mathsf{E}(R))]^{-1}L \leq y\right)$. Model parameters are set such that $\mathsf{E}(R) = 0.5$ (i.e. $\mu = 0$) and $\rho(X_i, X_j) = 0.25$ (i.e. $a = 0.25$). Parameters b and σ_ξ are determined such that $\mathsf{V}(R) = 1/15$ and Kendall's Tau, $\tau(X_i, R_i)$ is as listed on the graphs.

Proof:

Notice that $\mathsf{E}(f_N|Z) = N^{-1}\sum_{i=1}^N \mathsf{E}(1_{\tau_i \leq T}|Z) = N^{-1}Np(Z) = p(Z)$. As the variance of $f_N|Z$ is bounded from above by N^{-1}, f_N converges to the random variable $p(Z)$ in probability. Assuming that at least one default takes place, we notice that

$$\begin{aligned} \mathsf{E}(r_N|Z) &= \mathsf{E}\left(\sum_{i=1}^N 1_{\tau_i \leq T} l_i/\Delta_N \middle| Z\right) \\ &= \mathsf{E}\left(\mathsf{E}\left(\sum_{i=1}^N 1_{\tau_i \leq T} l_i/\Delta_N \middle| Z, 1_{\tau_1 \leq T}, \ldots, 1_{\tau_N \leq T}\right)\right), \end{aligned}$$

where the outer expectation in the last expression is with respect to the N default indicator functions. Due to conditional independence, we have

$$\begin{aligned} \mathsf{E}\left(\sum_{i=1}^N 1_{\tau_i \leq T} l_i/\Delta_N \middle| Z, 1_{\tau_1 \leq T}, \ldots, 1_{\tau_N \leq T}\right) &= \sum_{i=1}^N 1_{\tau_i \leq T}\mathsf{E}(l_i|Z)/\Delta_N \\ &= l(Z). \end{aligned}$$

As the probability of at least one default taking place approaches 1 for $N \to \infty$, and the variance of $f_N|Z$ is bounded from above by $(l^{\max})^2N^{-1}$, r_N converges to the random variable $l(Z)$ in probability. The first part of the proposition then follows. To prove the second part notice that c) and Proposition 1 ensures that $p(Z)$ is strictly decreasing in $a \cdot Z$. Similarly, b) ensures that $l(Z)$ is strictly decreasing in $b \cdot Z$. By the invariance of Kendall's Tau under strictly increasing transformations, we get that $\tau(f_\infty, l_\infty) = \tau(-a \cdot Z, -b \cdot Z) = \tau(a \cdot Z, b \cdot Z)$, where the second equality follows from the definition of Kendall's Tau. The statement $\tau(a \cdot Z, b \cdot Z) = 2\pi^{-1} \sin^{-1}(b \cdot a/\sqrt{\|a\|^2 + \|b\|^2})$ follows from the elliptical distribution of Z and the proof of Proposition 2. ∎

4 Random factor loadings

In this section we consider the special case of the general framework of Section 2 where the factor loadings depend deterministically on the factors. More specifically, let $Z_j, \ j = 1, \dots, d$, and $\varepsilon_i, \ i = 1, \dots, N$, be independent random variables with zero means and unit variances. Then we consider

$$\left.\begin{aligned} X_i &= a_i(Z) \cdot Z + v_i \varepsilon_i + m_i \\ l_i &= l_i^{\max}(1 - R_i) \end{aligned}\right\} \ i = 1, \dots, N, \tag{17}$$

where $R_i \in [0,1]$ are fixed, $v_i := \sqrt{1 - \mathsf{V}[a_i(Z) \cdot Z]}$ and $m_i := -\mathsf{E}[a_i(Z) \cdot Z]$, such that X_i has zero mean and unit variance.

Note that

$$v_i = \sqrt{1 - \int_{\mathbb{R}^d} (a_i(z) \cdot z)^2 dF^Z(z) + m_i^2} \tag{18}$$

$$m_i = -\int_{\mathbb{R}^d} a_i(z) \cdot z dF^Z(z) \tag{19}$$

where F^Z is the distribution function for Z.

For this model we can build the conditional loss distribution by the simple recursion (6) used for the Gaussian copula model. The requisite conditional default probabilities are given by

$$p_i(Z) = F_i^\varepsilon\left(\frac{c_i - a_i \cdot Z - m_i}{v_i}\right), \tag{20}$$

where F_i^ε is the cumulative distribution of ε_i. The computation of sensitivities closely follows the Gaussian copula model as described in Section 2.1.

In the model (17) the probability of default before some horizon T is

$$\begin{aligned} p_i = \mathsf{Prob}(\tau_i \leq T) &= \mathsf{Prob}(X_i \leq c_i) = \mathsf{E}\left(\mathsf{Prob}\left(\varepsilon_i \leq \frac{c_i - a_i(Z) \cdot Z - m_i}{v_i} | Z\right)\right) \\ &= \int_{\mathrm{R}^d} F_i^\varepsilon\left(\frac{c_i - a_i(z) \cdot z - m_i}{v_i}\right) dF^Z(z). \end{aligned} \tag{21}$$

This expression—computed either numerically or, if possible, by a closed-form solution—allows us to calibrate the model by determining the trigger levels c_i from default probabilities.

In practice, it is advantageous to work with a "separable" structure

$$a_i(Z) = \left(a_{i,1}(Z_1), \ldots, a_{i,d}(Z_d)\right), \tag{22}$$

which will simplify variance and mean computations. For instance.

$$\mathsf{V}\left(a_i(Z)Z\right) = \sum_{j=1}^{d} \mathsf{V}\left(a_{ij}(Z_j)Z_j\right); \tag{23}$$

$$\mathsf{Cov}(a_i(Z)\cdot Z, a_j(Z)\cdot Z) = \sum_{k=1}^{d} \mathsf{Cov}(a_{i,k}(Z_k)Z_k, a_{j,k}(Z_k)Z_k), \tag{24}$$

and so forth.

4.1 Specific model: two-point loadings distribution with Gaussian factors

Above we defined a fairly rich class of random factor loading models differing in the functional relationship between factors and loadings and in the choice of distributions for the factors and residuals.

For illustration, we shall now study in more detail a simple model, building on the standard Gaussian copula specification. Specifically, in this model $Z_j,\ j = 1, \ldots d$, and $\varepsilon_i,\ i = 1, \ldots N$, are standard Gaussian variates and the factor loadings are given by a *two-point distribution*:

$$a_{ij}(Z_j) = \begin{cases} \alpha_{ij}, & Z_j \leq \theta_{ij} \\ \beta_{ij}, & Z_j > \theta_{ij} \end{cases} \tag{25}$$

where α_{ij}, β_{ij} are positive constants and $\theta_{ij} \in \mathbb{R}$. Note that we have assumed the separable structure of equation (22). Loosely, we can think of this as a regime-switching model where loading j takes value α_{ij} with probability $\Phi(\theta_{ij})$ and value β_{ij} with probability $1-\Phi(\theta_{ij})$. If $\alpha_{ij} > \beta_{ij}$ then the factor loadings decrease in Z_j and thus, intuitively, asset values couple more strongly to "the economy" in bad times than in good. Clearly, the special case of constant, factor-independent a_{ij}'s is the Gaussian copula model, but in general the default drivers, X_i, (and thus the copula) will *not* be Gaussian.

While analytical tractability is not particularly important[9] we note that more general specifications than 25) are also tractable. For instance, extending the 2-point specification to an N-point specification is trivial and allows one to approximate a given smooth function to arbitrary precision. A piecewise linear (rather than piecewise flat) factor loading

[9] Non-tractable factor loading functions can be handled by numerical quadrature coupled with some type of function caching

function is also tractable. We also note in passing that (25) is a special case of the smooth and analytically tractable specification

$$\begin{aligned} a_i(Z) &= \left(\chi_{i,1} + \gamma_{i,1}\Phi\left(\mu_{i,1}Z_1 + \nu_{i,1}\right), \ldots, \chi_{i,d} + \gamma_{i,d}\Phi\left(\mu_{i,d}Z_d + \nu_{i,d}\right)\right) \\ &= \chi_i + \mathsf{diag}(\gamma_i)\Phi\left(\mu_i \cdot Z + \nu_i\right), \end{aligned} \tag{26}$$

where we introduced the d-dimensional vectors $\chi_i \geq 0$, $0 \leq \gamma_i \leq 1 - \chi_i$, μ_i, and ν_i. In the last equation, we understand that $\Phi(\cdot)$ is to be applied componentwise. We note that $0 \leq a_i(Z) \leq \chi_i + \gamma_i$, and must thus require that

$$||\chi_i + \gamma_i||^2 \leq 1.$$

We would typically expect μ_i to be negative, such that factor loadings increase when the factors decrease.

Although the specification (26) is tractable it requires additional parameters and involves fairly lengthy expressions. For the purposes of this paper the simpler specification (25) suffices and will serve to illustrate the basic idea of random factor loadings.

Moment properties of (25) are straightforward and are listed below.

Proposition 7 *For the model in (25), we have*

$$\begin{aligned} m_i &= -\sum_{j=1}^{d}\left(-\alpha_{ij}\varphi(\theta_{ij}) + \beta_{ij}\varphi(\theta_{ij})\right); \\ v_i &= \sqrt{1 - \sum_{j=1}^{d} V_{ij}}, \end{aligned}$$

where

$$\begin{aligned} V_{ij} = {} & \alpha_{ij}^2\left(\Phi(\theta_{ij}) - \theta_{ij}\varphi(\theta_{ij})\right) + \beta_{ij}^2\left(\theta_{ij}\varphi(\theta_{ij}) + 1 - \Phi(\theta_{ij})\right) \\ & - \left(-\alpha_{ij}\varphi(\theta_{ij}) + \beta_{ij}\varphi(\theta_{ij})\right)^2. \end{aligned}$$

Proof: Using the results of Lemma 5, we immediately get

$$\begin{aligned} \mathsf{E}\left(a_{ij}(Z_j)Z_j\right) &= \mathsf{E}\left(\alpha_{ij}1_{Z_j \leq \theta_{ij}}Z_j + \beta_{ij}1_{Z_j > \theta_{ij}}Z_j\right) = -\alpha_{ij}\varphi(\theta_{ij}) + \beta_{ij}\varphi(\theta_{ij}) \\ \mathsf{E}\left(a_{ij}(Z_j)^2 Z_j^2\right) &= \mathsf{E}\left(\alpha_{ij}^2 1_{Z_j \leq \theta_{ij}}Z_j^2 + \beta_{ij}^2 1_{Z_j > \theta_{ij}}Z_j^2\right) \\ &= \alpha_{ij}^2\left(\Phi(\theta_{ij}) - \theta_{ij}\varphi(\theta_{ij})\right) + \beta_{ij}^2\left(\theta_{ij}\varphi(\theta_{ij}) + \left(1 - \Phi(\theta_{ij})\right)\right). \end{aligned}$$

From this $\mathsf{V}\left(a_{ij}(Z_j)^2 Z_j^2\right) = V_{ij}$, with V_{ij} given above. The result then follows from equation (23). ∎

To characterize the dependence structure of the model we give the following result.

Proposition 8 *Define* $\overline{\theta}_{i,j,k} = \max(\theta_{i,k}, \theta_{j,k})$ *and* $\underline{\theta}_{i,j,k} = \min(\theta_{i,k}, \theta_{j,k})$*. Then*

$$\mathsf{E}(a_{i,j}) = \alpha_{i,j}\Phi(\theta_{i,j}) + \beta_{i,j}\left(1 - \Phi(\theta_{i,j})\right); \tag{27a}$$

$$\mathsf{V}(a_{i,j}) = \alpha_{i,j}^2\Phi(\theta_{i,j}) + \beta_{i,j}^2\left(1 - \Phi(\theta_{i,j})\right) - \mathsf{E}(a_{i,j})^2; \tag{27b}$$

$$\rho\left(X_i, X_j\right) = \sum_{k=1}^{d}(E_{i,j,k} - m_i m_j),\ i \neq j, \tag{27c}$$

where

$$\begin{aligned} E_{i,j,k} &:= \mathsf{E}\left(a_{i,k}(Z_k)a_{j,k}(Z_k)Z_k^2\right) \\ &= \alpha_{i,k}\alpha_{j,k}\left(\Phi\left(\underline{\theta}_{i,j,k}\right) - \underline{\theta}_{i,j,k}\varphi(\underline{\theta}_{i,j,k})\right) \\ &\quad + \beta_{i,k}\beta_{j,k}\left(1 - \Phi\left(\overline{\theta}_{i,j,k}\right) + \overline{\theta}_{i,j,k}\varphi(\overline{\theta}_{i,j,k})\right) \\ &\quad + \alpha_{i,k}\beta_{j,k}1_{\theta_{i,k}\geq\theta_{j,k}}\left(\Phi(\theta_{i,k}) - \Phi(\theta_{j,k}) + \theta_{j,k}\varphi(\theta_{j,k}) - \theta_{i,k}\varphi(\theta_{i,k})\right) \\ &\quad + \beta_{i,k}\alpha_{j,k}1_{\theta_{j,k}\geq\theta_{i,k}}\left(\Phi(\theta_{j,k}) - \Phi(\theta_{i,k}) + \theta_{i,k}\varphi(\theta_{i,k}) - \theta_{j,k}\varphi(\theta_{j,k})\right). \end{aligned}$$

Proof: The first equation is obvious. The second follows from

$$\begin{aligned} \mathsf{E}\left(\left(\alpha_{i,j}1_{Z_j\leq\theta_{i,j}} + \beta_{i,j}1_{Z_j>\theta_{i,j}}\right)^2\right) &= \mathsf{E}\left(\alpha_{i,j}^2 1_{Z_j\leq\theta_{i,j}} + \beta_{i,j}^2 1_{Z_j>\theta_{i,j}} + 2\alpha_{i,j}\beta_{i,j}1_{Z_j>\theta_{i,j}}1_{Z_j\leq\theta_{i,j}}\right) \\ &= \mathsf{E}\left(\alpha_{i,j}^2 1_{Z_j\leq\theta_{i,j}} + \beta_{i,j}^2 1_{Z_j>\theta_{i,j}}\right) \\ &= \alpha_{i,j}^2\Phi(\theta_{i,j}) + \beta_{i,j}^2\left(1 - \Phi(\theta_{i,j})\right), \end{aligned}$$

$$\begin{aligned} \overline{v}_{i,j} &= \mathsf{V}\left(\chi_{i,j} + \theta_{i,j}\Phi\left(\alpha_{i,j}Z_j + \beta_{i,j}\right)\right) = \theta_{i,j}^2\mathsf{V}\left(\Phi\left(\alpha_{i,j}Z_j + \beta_{i,j}\right)\right) \\ &= \theta_{i,j}^2\mathsf{E}\left(\Phi\left(\alpha_{i,j}Z_j + \beta_{i,j}\right)^2\right) - \theta_{i,j}^2\mathsf{E}\left(\Phi\left(\alpha_{i,j}Z_j + \beta_{i,j}\right)\right)^2 \\ &= \theta_{i,j}^2\Phi_2\left(\frac{\beta_{i,j}}{\sqrt{1+\alpha_{i,j}^2}}, \frac{\beta_{i,j}}{\sqrt{1+\alpha_{i,j}^2}}; \frac{\alpha_{i,j}}{\sqrt{1+\alpha_{i,j}^2}}\right) - \theta_{i,j}^2\Phi\left(\frac{\beta_{i,j}}{\sqrt{1+\alpha_{i,j}^2}}\right), \end{aligned}$$

where we used (30b) from Appendix A.

To prove the third equation we notice that

$$\begin{aligned} \mathsf{E}\left(a_{i,k}(Z_k)a_{j,k}(Z_k)Z_k^2\right) &= \mathsf{E}\left(\alpha_{i,k}\alpha_{j,k}1_{Z_k\leq\theta_{i,k}}1_{Z_k\leq\theta_{j,k}}Z_k^2 + \beta_{i,k}\beta_{j,k}1_{Z_k>\theta_{i,k}}1_{Z_k>\theta_{j,k}}Z_k^2\right) \\ &\quad + \mathsf{E}\left(\alpha_{i,k}\beta_{j,k}1_{Z_k\leq\theta_{i,k}}1_{Z_k>\theta_{j,k}}Z_k^2 + \beta_{i,k}\alpha_{j,k}1_{Z_k>\theta_{i,k}}1_{Z_k\leq\theta_{j,k}}Z_k^2\right) \\ &= \mathsf{E}\left(\alpha_{i,k}\alpha_{j,k}1_{Z_k\leq\min(\theta_{i,k},\theta_{j,k})}Z_k^2 + \beta_{i,k}\beta_{j,k}1_{Z_k>\max(\theta_{i,k},\theta_{j,k})}Z_k^2\right) \\ &\quad + \mathsf{E}\left(\alpha_{i,k}\beta_{j,k}1_{\theta_{j,k}<Z_k\leq\theta_{i,k}}Z_k^2 + \beta_{i,k}\alpha_{j,k}1_{\theta_{i,k}<Z_k\leq\theta_{j,k}}Z_k^2\right). \end{aligned}$$

From Lemma 5 in Appendix A, the value of this expression is $E_{i,j,k}$, as given in the proposition. The final result now follows from

$$\rho(X_i, X_j) = \mathsf{E}(X_i X_j) = \mathsf{Cov}(a_i(Z)\cdot Z, a_j(Z)\cdot Z)$$

and

$$\mathsf{Cov}(a_i(Z)\cdot Z, a_j(Z)\cdot Z) = \sum_{k=1}^{d} \mathsf{Cov}(a_{i,k}(Z_k)Z_k, a_{j,k}(Z_k)Z_k). \qquad \blacksquare$$

If we restrict ourselves to low-dimensional models, the specification (25) is simple enough to allow for closed-form default probabilities. For instance:

Proposition 9 *Assume that the dimension of Z is $d = 1$. Then*

$$\begin{aligned} \mathsf{Prob}\left(\tau_i \le T\right) &= \mathsf{Prob}\left(X_i \le c_i\right) \\ &= \Phi_2\left(\frac{c_i - m_i}{\sqrt{v_i^2 + \alpha_i^2}}, \theta_i; \frac{\alpha_i}{\sqrt{v_i^2 + \alpha_i^2}}\right) + \Phi\left(\frac{c_i - m_i}{\sqrt{v_i^2 + \beta_i^2}}\right) \\ &\quad -\Phi_2\left(\frac{c_i - m_i}{\sqrt{v_i^2 + \beta_i^2}}, \theta_i; \frac{\beta_i}{\sqrt{v_i^2 + \beta_i^2}}\right). \end{aligned} \tag{28}$$

Proof: We have

$$\begin{aligned} \mathsf{Prob}\left(\tau_i \le T\right) &= \mathsf{Prob}\left(\alpha_i 1_{Z \le \theta_i} Z + \beta_i 1_{Z > \theta_i} Z + \varepsilon_i v_i + m_i \le c_i\right) \\ &= \mathsf{E}\left(\mathsf{Prob}\left(\varepsilon_i \le \frac{c_i - (\alpha_i 1_{Z \le \theta_i} Z + \beta_i 1_{Z > \theta_i} Z) - m_i}{v_i} | Z\right)\right) \\ &= \mathsf{E}\left(\Phi\left(\frac{c_i - (\alpha_i 1_{Z \le \theta_i} Z + \beta_i 1_{Z > \theta_i} Z) - m_i}{v_i}\right)\right) \\ &= \int_{-\infty}^{\theta_i} \Phi\left(\frac{c_i - \alpha_i Z - m_i}{v_i}\right) \varphi(Z) dZ + \int_{\theta_i}^{\infty} \Phi\left(\frac{c_i - \beta_i Z - m_i}{v_i}\right) \varphi(Z) dZ \end{aligned}$$

and the proposition now follows from results in appendix A. $\blacksquare$
We note that the density of X_i can be obtained by differentiating (28) wrt c_i.

4.2 Large-portfolio limits of two-point distribution model

With the specification (25), we are not only hoping to (crudely) mimic an empirical dependence of correlations on broad market conditions, we also hope to generate a base correlation skew when $\alpha_{ij} > \beta_{ij}$. To elaborate, consider the view of a senior tranche investor. This investor will only experience losses on his position when several names default together, that is, in scenarios where the systematic factors Z likely take on low values. If Z is low, however, the factor loading will be high, making it appear to the senior investor that correlations are quite high. For the equity investor, who is likely to experience losses even in scenarios where Z is not low, the effective factor loading will appear as a weighted average between α_{ij} and β_{ij}. To the equity investor, the world will thus look as if correlations are of average magnitude. In fact, the convexity of tranche values in correlation will further reduce the “effective” correlation seen by the equity investor and increase that seen by the senior investor.

To lend some credence to the above intuitive argument, and to provide some useful results for large portfolio risk management under (25), we shall now briefly examine the large-portfolio limit distribution of the model[10]. A direct examination of base correlation skews for a finite portfolio can be found in Section 6.2. To proceed, assume for simplicity that $d = 1$ and that the portfolio is homogeneous[11]. Reflecting the homogeneity of the portfolio, we omit subscripts in the following result ($\beta_{ij} = \beta$, $\alpha_{ij} = \alpha$, $m_i = m$, and so forth).

Proposition 10 *Consider a homogeneous portfolio with a 1-dimensional factor structure. Under (25), the following holds*

$$\begin{aligned} \lim_{N\to\infty} \mathsf{Prob}\left(L/N \geq ly\right) &= \Phi\left(\min\left(\Omega(y)/\alpha, \theta\right)\right) + 1_{\Omega/\beta>\theta}\left(\Phi\left(\Omega(y)/\beta\right) - \Phi(\theta)\right), \\ \Omega(y) &:= c - v\Phi^{-1}(y) - m. \end{aligned}$$

Proof: Consider

$$\begin{aligned} h(Z) &= N^{-1}l^{-1}\mathsf{E}(L|Z) = N^{-1}\mathsf{E}\left(\sum_{i=1}^{N} 1_{\tau_i \leq T}|Z\right) \\ &= N^{-1}\sum_{i=1}^{N}\mathsf{E}\left(1_{\tau_i\leq T}|Z\right) = \mathsf{Prob}\left(X_1 \leq c|Z\right) \\ &= \mathsf{Prob}\left(\varepsilon_1 \leq \frac{c - a(Z)Z - m}{v}|Z\right) = \Phi\left(\frac{c - 1_{Z\leq\theta}\alpha Z - 1_{Z>\theta}\beta Z - m}{v}\right). \end{aligned}$$

Then, by the usual diversification arguments,

$$\begin{aligned} \lim_{N\to\infty} \mathsf{Prob}\left(L/N \geq ly\right) &= \mathsf{Prob}\left(h(Z) \geq y\right) \\ &= \mathsf{Prob}\left(a(Z)Z \leq \Omega(y)\right) \\ &= \mathsf{Prob}\left(a(Z)Z \leq \Omega(y), Z \leq \theta\right) + \mathsf{Prob}\left(a(Z)Z \leq \Omega(y), Z > \theta\right) \\ &= \mathsf{Prob}\left(\alpha Z \leq \Omega(y), Z \leq \theta\right) + \mathsf{Prob}\left(\beta Z \leq \Omega(y), Z > \theta\right) \\ &= \Phi\left(\min\left(\Omega(y)/\alpha, \theta\right)\right) + 1_{\Omega(y)/\beta>\theta}\left(\Phi\left(\Omega(y)/\beta\right) - \Phi(\theta)\right). \quad \blacksquare \end{aligned}$$

In Figure 5 below, we use the result in Proposition 10 to examine the loss distributions associated with different settings β and α. We note that when $\alpha > \beta$, the loss distribution has qualitative properties consistent with the existence of a base skew: a fat upper tail, yet a reduced probability of generating small losses, c.f. Figure 5. (The somewhat non-smooth form of the large-portfolio loss distribution in Figure 5 is caused by the simple two-point specification used and the absence of residual noise in the large-portfolio limit).

---

[10] Large-portfolio results for more general models than (25) are obviously possible, given the conditional independence of the model. The required procedure to obtain such results is similar to that given in the text.

[11] Extensions to non-homogeneous portfolios are trivial; see Section 3.3 for the required arguments and assumption.

Figure 5: Large-portfolio distribution, two-point random factor loading model

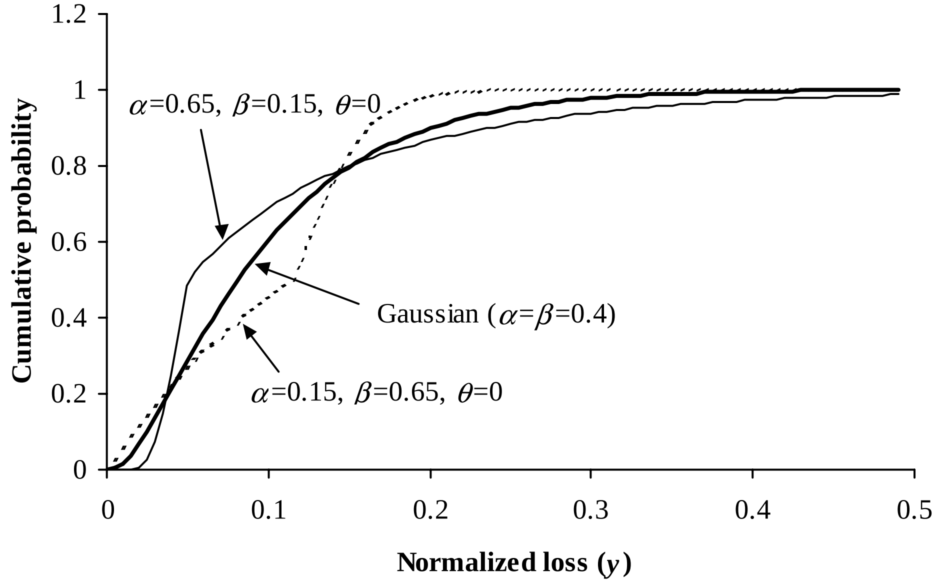

Notes: Homogenous portfolio in one-factor Gaussian copula, with two-point random factor loading model. The graphs shows $\text{Prob}(L/(Nl) < y)$. The probability of default was 10% per name; parameters m and v were computed according to Proposition 7.

5 Alternative correlation skew models.

While the method presented in Section 4 is attractive for its simplicity and close relation to an established and well-understood framework, it is just one of many ways to generate a correlation skew. At the end of the day, any model that can generate a loss distribution similar to that of Figure 2 will produce a correlation skew. To briefly comment on a few alternative approaches, consider first an outright move to a copula different from the Gaussian[12]. For instance, among the better-known alternatives to the Gaussian copula, the Marshall-Olkin copula (see e.g. Giesecke (2003) and references therein) can be verified to have enough flexibility to produce correlation skews. As the model apparently enjoys some popularity with practitioners, let us discuss it in more detail.

We first recall that the Marshall-Olkin is a Poisson-type model where the default intensity of credit i, λ_i, is broken into an idiosyncratic component and a number of systematic components. Say that we have M systematic default shocks, modeled as independent Poisson processes $\bar{N}_k(t)$ with intensities $\bar{\lambda}_k$, $k = 1,\ldots,M$; these systematic shocks can potentially affect several firms simultaneously. To define whether a particular firm is sensitive to a particular systematic shock, one introduces 0- or 1-valued indicator variables I_{ik}

[12] The random factor loading approach obviously makes the copula non-Gaussian, but the change of copula is, in a sense, implicit rather than explicit.

such that firm i will default with certainty the first time $\bar{N}_k(t)$ jumps if and only if $I_{ik}=1$. If $I_{ik}=0$, firm i is assumed completely insensitive to $\bar{N}_k(t)$. Define also N idiosyncratic Poisson processes $N_i^{\mathsf{f}}(t)$, $i=1,\ldots,N$, ("f" for "firm") independent of each other and of the $\bar{N}_k(t)$, $k=1,\ldots,M$. Let the intensities of these Poisson processes be $\lambda_i^{\mathsf{f}}(t)$, $i=1,\ldots,N$. In the Marshall-Olkin copula model, we then express the aggregate default Poisson processes for the firms as

$$N_i(t)=\sum_{k=1}^{M}\bar{N}_k(t)I_{ik}+N_i^{\mathsf{f}}(t),\quad i=1,\ldots,N,$$

where the right-hand side is obviously Poisson (as a thinned sum of independent Poisson processes is again Poisson) with intensity

$$\lambda_i(t)=\sum_{k=1}^{M}\bar{\lambda}_k I_{ik}+\lambda_i^{\mathsf{f}}(t),\quad i=1,\ldots,N.$$

The dependence structure of the Marshall-Olkin model is determined by the matrix $\{I_{ik}\}$ and the intensities $\bar{\lambda}_k$, $k=1,\ldots,M$. By having a construct such as the $\{I_{ik}\}$ matrix, the Marshall-Olkin copula is quite flexible and grants tight control of the relative probabilities of various combinations of firm defaults, and thereby of the total loss distribution. We have verified that this model is easily capable of generating a correlation skew. On the other hand, the exact parameterization of the Marshall-Olkin copula is a rather formidable problem, given the abstract nature (and sheer number) of its parameters. We note that to make the model consistent across different CDOs, one really must calibrate a single matrix and a single set of intensities for all credits in the universe of traded credit default swaps. Such a calibration would likely be difficult to make robust, and strong assumptions will be needed to make it feasible. We also point out that even if a calibration method could be constructed, the model remains quite unwieldy and involve a number of non-trivial operational and computational issues, particularly in the computation of hedges. Similar issues arise with many other copulas and, to an even greater extent, dynamic portfolio models (such as the one in Duffie and Garleanu (2001)).

Stepping back to the Gaussian copula set-up, it is well-known that a Gaussian copula equipped with a non-constant correlation matrix will produce a non-flat base correlation skew. Some authors (see eg, Gregory and Laurent (2004)) have proposed correlation structures where correlations are high within certain 'sectors' (arbitrary subportfolios) and low between sectors. However, when—as is perhaps natural—the sectors are taken to represent industrial or geographical categories, the resulting skew is far weaker than that observed in practice. One could imagine taking this approach one step further and attempting to "imply" a correlation matrix that will reproduce the observed smile. This approach is perhaps best framed in a factor set-up, where we make factor loadings functions of the credit spreads. To be more specific, consider replacing equation (5) with[13]

$$X_i=g(s_i)Z+\varepsilon_i\sqrt{1-||g(s_i)||^2},\ i=1,\ldots,N, \tag{29}$$

[13] Or perhaps a perturbative specification: $X_i=a_i g(s_i)Z+\varepsilon_i\sqrt{1-||a_i g(s_i)||^2},\ i=1,\ldots,N$, where a_i may be an empirical factor loading, for instance.

where s_i represents the input credit spread of firm i, $i = 1,\ldots,N$, and $g : \mathrm{R} \to [0,1]^d$ is some mapping function. We stress that the $g(s_i)$ are to be interpreted as constants, not as stochastic quantities; $g(\cdot)$ can be "implied" (parametrically or non-parametrically) to best-fit the correlation skew. We note that to produce a skew consistent with that seen in the market, g needs to produce a correlation matrix such that the senior investor (who is affected more by low-spread credits than is the equity investor) will "see" a higher correlation than the equity investor. For $d = 1$ this simply means that g must be decreasing.

While the approach (29)—even in its one-factor version—is certainly capable of generating strong correlation skews, it has a number of drawbacks. First, there appears to be little empirical support for the basic conjecture, namely that firms with high default risk are less correlated to the overall market than firms with low default risk; second, to reproduce the correlation skews observed in practice, the function g turns out to have be quite extreme, producing correlation matrices with elements that get unrealistically close to 0 or 1; and third, the model predicts that near-homogeneous baskets should have no correlation skew (which is unlikely). Furthermore, depending on the details of the chosen specification, one will have the problem that the correlations change daily when spreads are updated. Finally, the extreme correlations required to match observed skews lead to hedges in terms of single-obligor positions which are radically different from those computed with other models.

6 Numerical results

Having previously examined large-portfolio limits, we now turn to a numerical examination of tranches in finite-sized portfolios, using the models and algorithms developed in Sections 3 and 4. While we shall conclude this section with some data for a realistic portfolio (I-Boxx) under real market conditions, to facilitate replication of our results we use a simplified portfolio for most of our tests. Results for more realistic portfolios are qualitatively similar.

We consider a test portfolio consisting of $N = 25$ obligors each with a notional of 10 million for a total portfolio notional of 250 million. On this portfolio we consider 5-year tranches 0–3%, 3–7%, 7–12%, 12–20%, and 20–30%. All models will be calibrated to a market where default-free interest rates are zero, where the default swap spread of the i'th obligor is $i \cdot 10$bps, ie, spreads run from 10 to 250bps, and where the expectation of recovery conditioned on default ($\mathsf{E}(R_i|\tau_i = T)$) is assumed to be 40% for all obligors and all T. To ease comparison all parameterized models will be required to price the 0–7% base tranche at 20% flat correlation.

To ease the notation we shall refer to the Gaussian copula model with flat correlation ρ as 'G_ρ', to the specific random recovery model of Section 3.2 as 'RR' and to the specific random factor loading model of Section 4.1 as 'RFL'.

6.1 Random recovery model

For the specific RR model of Section 3.2 we consider a generic model with two factors in order to allow control over the codependence between defaults and recovery rates. We assume that the parameters, a, b, μ and σ_ξ are obligor-independent and we suppress the obligor index i. For given values of a, b and σ_ξ we calibrate μ such that $\mathsf{E}(R_i|\tau_i = T) =$ 40%, for all T, as required. Note that this requires μ to depend on T. Next, for given b and σ_ξ we choose the value of a such that the requirement on the valuation of the 0–7% tranche is satisfied[14].

We shall look at the six different sets of parameter values listed in Table 1. From the table we see that the value of a is smaller in all cases than the $\sqrt{20\%} \approx 0.447$ required to satisfy the 0–7% tranche price requirement in the absence of random recovery. Going into more detail, we note that the required value of a increases with σ_ξ and decreases with default-recovery correlation (as measured by $a \cdot b$). Intuitively, this is because tranches are actually sensitive to the covariance between single-obligor *losses* (not just defaults); a large value of b results in large recovery covariance and this leads to high loss correlation even when default correlations and default-recovery correlations are low. When tranche values are translated into a Gaussian copula framework the loss covariance has to be represented solely in terms of default correlation, which will therefore appear higher. These observations hold also for more realistic tranched portfolios such as I-Boxx.

Name	a	b	σ_ξ
RR1	(0.225,0)	(0.8,0)	1.0
RR2	(0.305,0)	(0.4,0.4)	1.0
RR3	(0.355,0)	(0,0.8)	1.0
RR4	(0.195,0)	(0.8,0)	0.01
RR5	(0.3,0)	(0.4,0.4)	0.01
RR6	(0.355,0)	(0,0.8)	0.01

Table 1: The six named sets of parameter values for the RR model. Note that a and b are two-dimensional vectors reflecting the fact that the model has two factors.

As pointed out in the introduction, a skew in tranche values arises from a loss distribution which differs from the loss distribution of a Gaussian copula model by having lower probability of zero loss and a fat upper tail. Based on our intuitive understanding of the impact of random recovery in generating loss covariance, we would expect that its effect on the loss distribution can largely be compensated by a reduction in default correlation. Thus we would expect a properly based comparison between models with and without random recovery to show little difference.

To examine this in detail Figure 6 shows the *difference* between the cumulative loss distributions for the RR model with the parameter values of Table 1 and the distribution

[14] By rotational invariance we can take the second component of a to be zero.

of the $G_{20\%}$ model. As expected, the differences are quite small and do not exhibit the characteristics required to generate a skew in tranche values.

Figure 6: Loss distributions for random recovery model

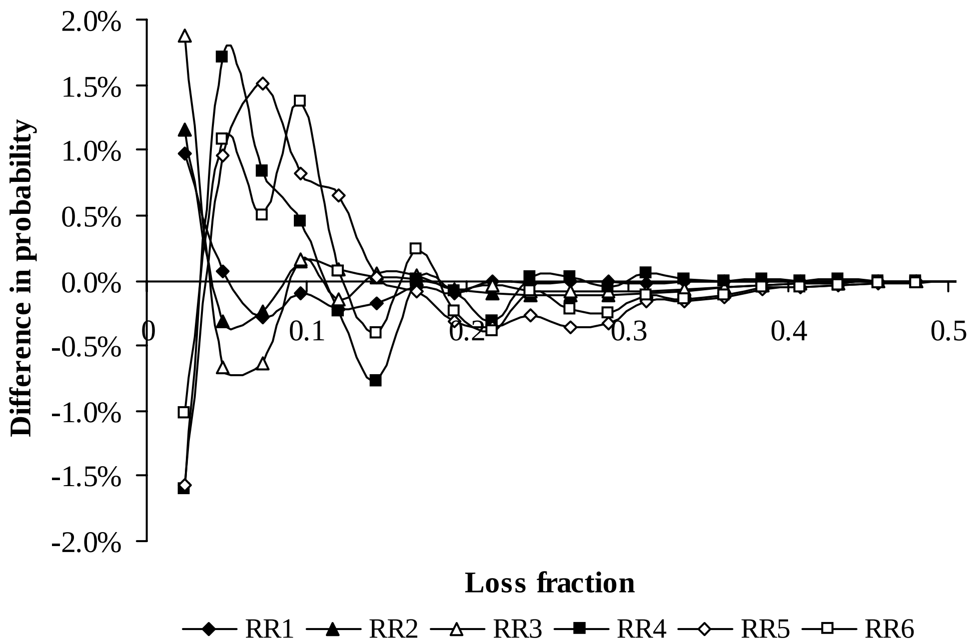

Notes: Parameter sets RR1-RR6 are in Table 1. The graphs show the *differences* between the model distributions and the distribution of the $G_{20\%}$ model.

To investigate the impact on tranche values in more detail, Figure 7 shows the implied base correlations for the RR model. We observe that, as expected from the loss distributions, the RR model is not capable of producing strong skews[15]. This finding also applies to tranches of traded index portfolios.

6.2 Random factor loadings model

For the specific RFL model of Section 4.1 we shall consider a generic one-factor model where the parameters α, β and θ are identical for all obligors. We shall look at the six different sets of parameter values listed in Table 2. For given values of β and θ we have determined α such that the requirement on the 0–7% tranche is satisfied. Note that when β is small and θ is low (negative) a high value of α is required. Thus, intuitively, this gives a model where factor loadings are "normally" low, but, for very low factor values, they become very high. In this model senior tranches derive a great deal of value from the many expected defaults in the "disaster state" with high factor loadings, whilst junior

[15] For the special case of independence between default and recovery, this conclusion can also be drawn from numerical results in Gregory and Laurent (2004).

Figure 7: Implied base correlations for random recovery model

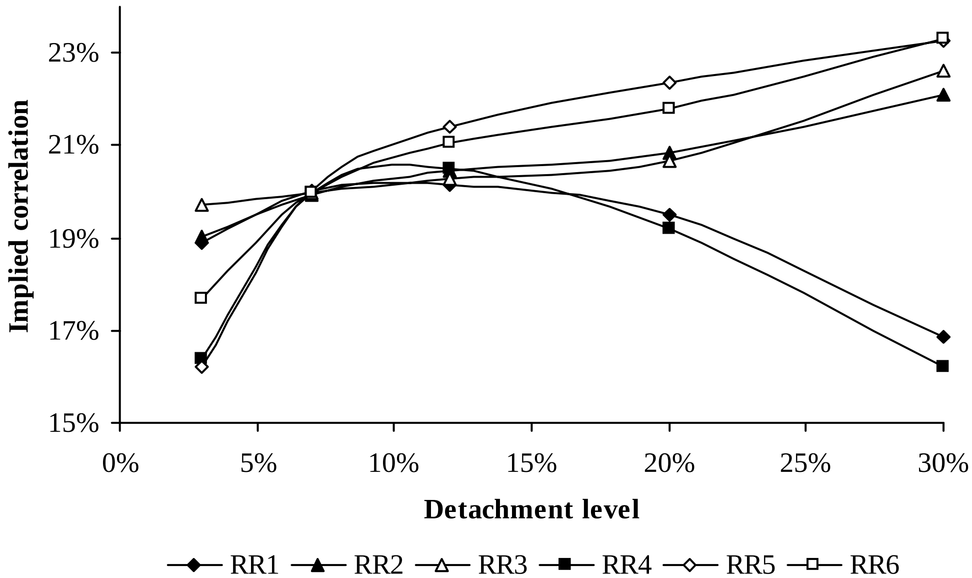

Notes: Parameter sets RR1-RR6 are in Table 1.

tranches derive value from the low effective correlation in the normal state; intuitively, this should produce a strong skew. On the other hand, when θ is high (and positive), the difference $|\alpha - \beta|$ cannot be too great and this leads to more "peaceful" and less skewed models. We first look at the cumulative loss distributions generated by the model with

Name	α	β	θ
RFL1	0.62	0.03	0
RFL2	0.425	0.5	0
RFL3	0.54	0.03	1
RFL4	0.43	0.54	1
RFL5	0.9	0.269	-2
RFL6	0.425	0.485	-1

Table 2: The six named sets of parameter values for the RR model.

these parameters. To show the differences from the Gaussian copula model more clearly, Figure 8 actually shows the *difference* between the model distributions and the distribution of the $G_{20\%}$ model. As expected on intuitive grounds, the RFL model can give rise to loss distributions which differ markedly from the distribution of the Gaussian copula model in the way required to generate a skew: lower probability of zero loss and a fat upper tail.

Next, Figure 9 shows the implied base correlations for the same sets of parameters. As is clear from Figure 9, the RFL model is capable of producing strong skews. In particular,

Figure 8: Loss distributions for random factor loading model

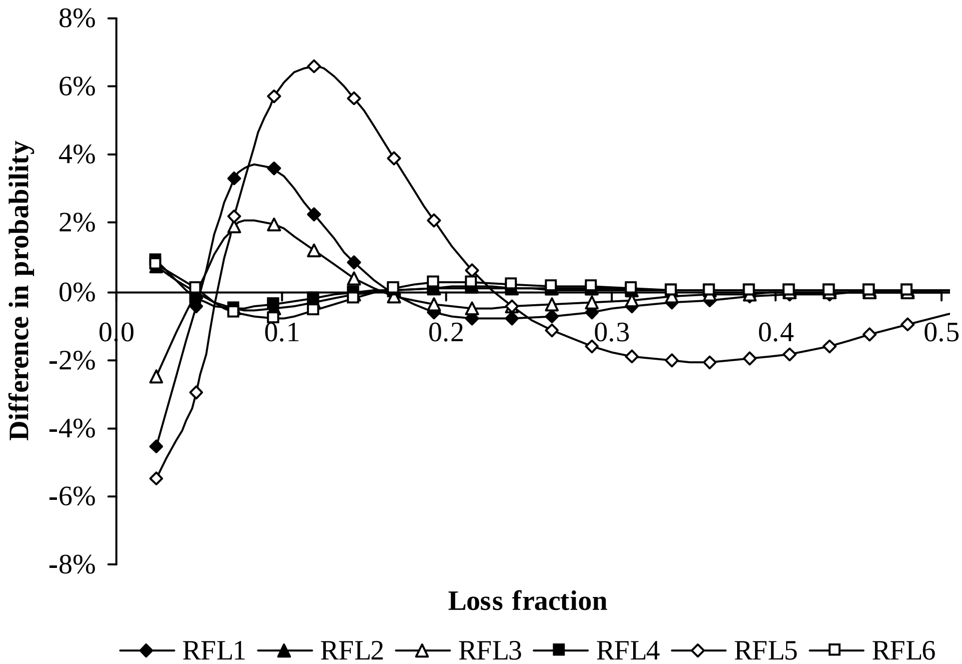

Notes: Parameter sets RFL1-RFL6 are in Table 2. The graphs show the *differences* between the model distributions and the distribution of the $G_{20\%}$ model.

we note that the skew is upward (downward) sloped for α greater (smaller) than β and that the steepest skew is produced with the smallest value (-2) of θ. This is not unexpected, given the interpretation of models with small β and low θ given above.

We now turn to the properties of the skew, in particular its dynamics under spread changes. The question of skew dynamics is important because it affects the sensitivities of tranche values to default swap spreads. More precisely, when spreads go up, tranche values move for two reasons: first, because of the increase in overall expected loss, second, because (implied) correlations change when the skew changes. Of course, it is possible that the skew should be insensitive to spreads. This seems very counterintuitive, however, when we consider that the relative riskiness of tranches clearly depends on the overall portfolio spread. For example, in a high spread environment the 3–7% tranche will look much more like an equity tranche than it does when spreads are low. Thus—rather than expecting a static skew—we would expect the skew to move to the right, ie, for an upward sloping skew, implied correlations to move down, when, *ceteris paribus*, spreads go up.

To illustrate the dynamics of the skew in the RFL model, we give in Figure 10 the skews produced by the model with parameters RFL5 when all default swap spreads are increased by 50bps. When compared to the skew produced for the initial spreads using the same parameters, we see that the effect of an increase in spreads is to shift the skew to the right. For an upward sloping skew this implies that the base correlation for given

Figure 9: Implied base correlations for random factor loading model

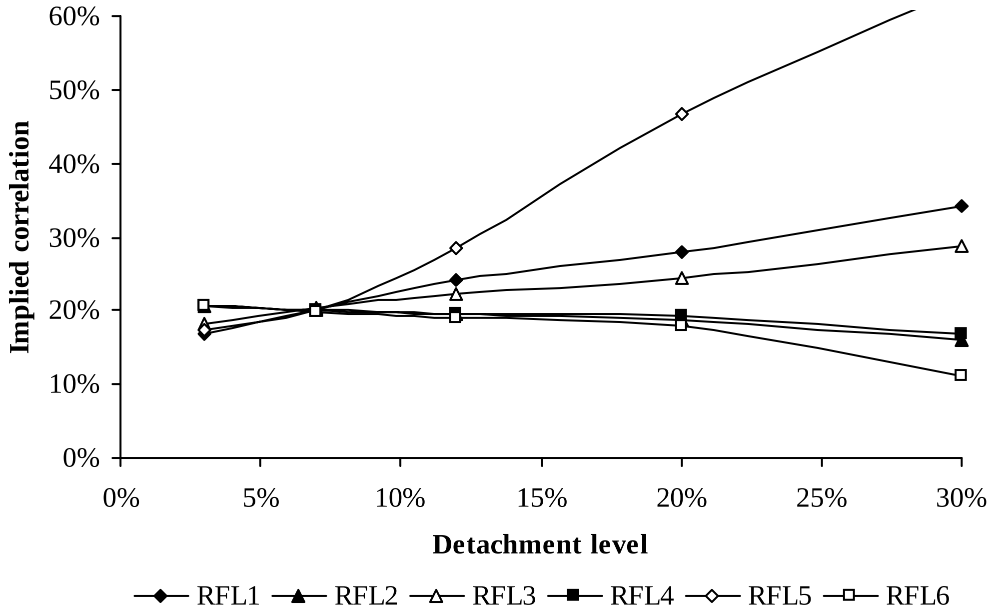

Notes: Parameter sets RFL1-RFL6 are in Table 2.

detachment level decreases.

Since all base tranche values decrease with increasing correlation, base tranche spread sensitivities in the RFL model (with upward sloping skew) must be *greater* than those computed with a static skew. For tranches with non-zero attachment levels, sensitivities may be computed as differences between base tranche sensitivities; for such tranches the contribution to sensitivities coming from the skew dynamics can have either sign.

6.3 A real-life example: I-Boxx

For a more realistic application of the RFL model we here present the base correlation skews produced for the I-Boxx NA portfolio (of May 2004) with a range of model parameters. The parameter sets are listed in Table 3 and the resulting base correlations are shown in Figure 11. Evidently, the RFL model—even in its simplest form with only three parameters—is capable of generating skews of a wide range of slopes and curvatures for I-Boxx. We note that the market skew is compatible with a low, negative threshold and a large difference in factor loadings between the low-factor and high-factor regimes. It is thus possible that the relatively high premiums of senior tranches are due to a market perception of the existence of a "disaster state" in which not only default probabilities, but also effective correlations are high. On the flip side, the relatively high price of insurance on the equity tranche could be due to the perception that correlations are "normally"

Figure 10: Skew dynamics in random factor loading model

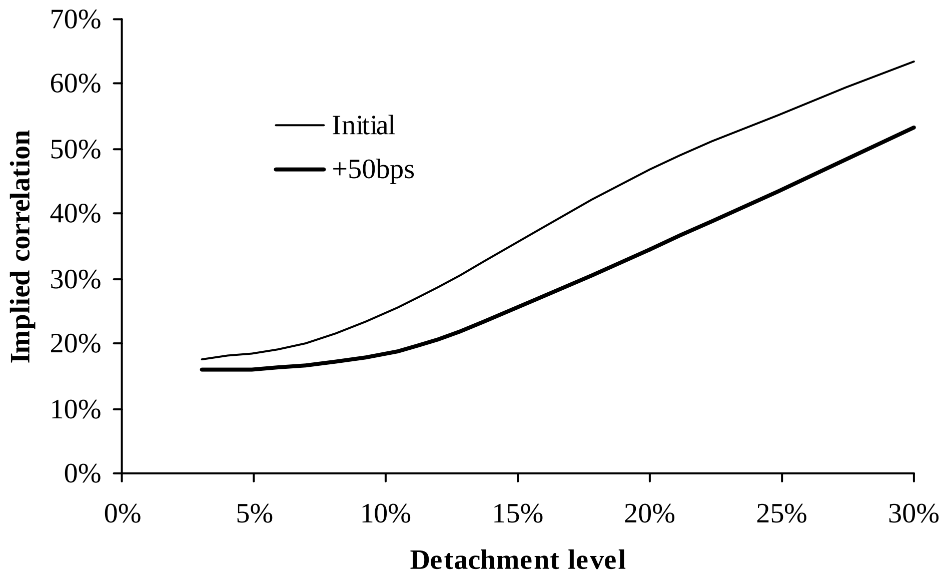

Notes: Base correlation for the RFL model with parameters RFL5 after a 50bps additive shift to all default swap spreads (fat curve). The correlations produced with the initial spreads is given for reference (thin curve).

Name	α	β	θ
RFL7	0.58	0.38	-1.6
RFL8	0.85	0.45	-3
RFL9	1.01	0.41	-2.5
RFL10	1.26	0.36	-2.4
RFL11	0.63	0.01	-0.6

Table 3: The six named sets of RFL model parameters applied to the tranches of I-Boxx NA.

rather low[16].

7 Conclusion

This paper has introduced two tractable extensions of the Gaussian copula, both aimed at making the model conform better with observed phenomena. The first of these mecha-

[16]This perception does not necessarily conflict with the correlation values computed from historical data (such as, eg, the KMV data sets); if the sampling period spans several business cycles, the resulting correlations will be weighted averages of correlations in bear and bull markets.

Figure 11: Base correlation skews for random factor loading model
(I-Boxx NA portfolio, May 2004)

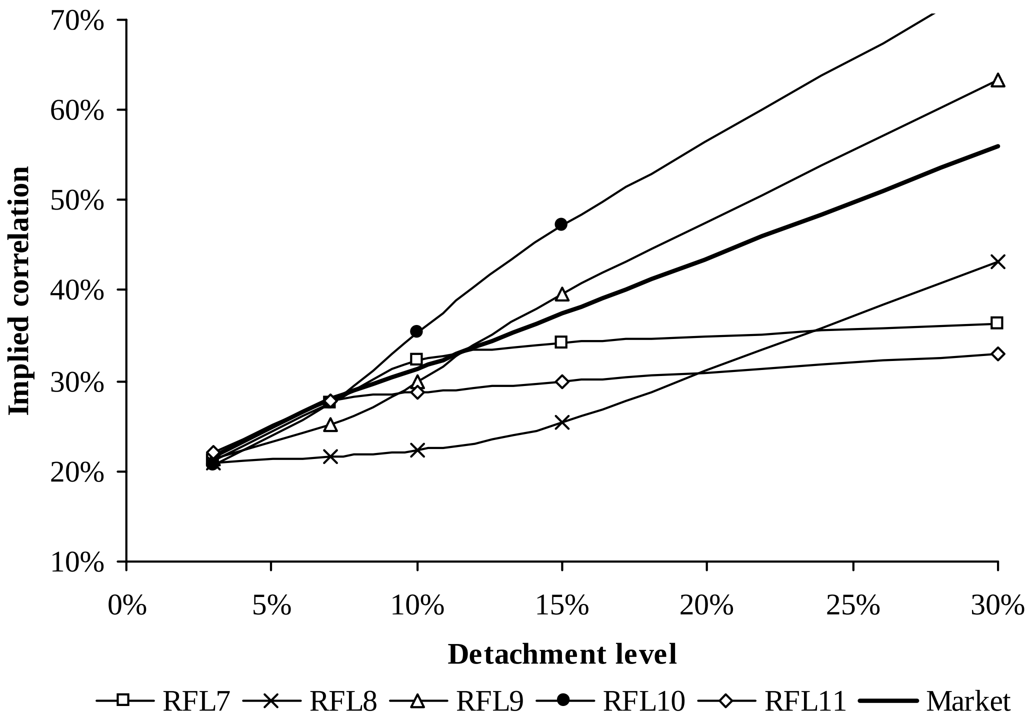

Notes: Parameter sets RFL7-RFL11 are in Table 3. The market correlation skew is given for reference (fat curve).

nisms – *randomized recovery* – is shown to produce a heavy upper tail in portfolio loss distributions and as such should be of interest in the tail risk management of default-sensitive portfolios. Convenient large-portfolio results useful for this exercise were produced. For finite portfolios, efficient numerical techniques were developed to handle random recovery; we use these techniques, along with a carefully developed calibration mechanism, to demonstrate that it is unlikely that random recovery is the prime mechanism behind the correlation skews observed in CDO tranches on standardized index portfolios. Our second extension – *randomized factor loadings* – does a much better job in this respect and, at reasonable parameter levels, is capable of producing correlation skews similar to those observed in the market. The resulting model is straightforward to parameterize and numerically efficient.

The specific parameterizations and model examples used in this paper were rather simplistic, and work remains in uncovering parameters and functional forms that best describe the market. We note that while we, for clarity, have treated the two model extensions in separation, combining both into a single model is an obvious possibility that may further increase realism (at the expense of more parameters). Other topics for future research involves a deeper examination of the hedges produced by various models, and possible further model extensions to better capture observed term structures of loss distributions.

We conclude the paper with a word of warning about implied correlation skew modeling. Here, it should be kept in mind that the market is in its absolute infancy and all modeling attempts are necessarily based on a very limited set of observations that may not even be fully representative. For instance, it is not inconceivable that market imperfections and misbalances currently contribute effects that are transitory and will abate as the market matures. Parameters and models will then obviously require revisions over time. In the same vein, as time progresses more information will be revealed about the dynamics of the correlation skew and its dependence on spread levels. This, in turn, will allow for evaluation of the realism and hedge performance of models, and will undoubtedly lead to more sophisticated models. At the current time, however, one is probably best served with relatively simple mechanisms, such as those presented in this paper.

A Some Gaussian integrals

In our notation φ, Φ and Φ_2 denote respectively the standard Gaussian density function, the standard Gaussian cumulative distribution function and the standard bivariate Gaussian cumulative distribution function.

Lemma 1 For arbitrary real constants a, b and c

$$\int_{-\infty}^{\infty} \Phi(ax+b)\varphi(x)dx = \Phi\left(\frac{b}{\sqrt{1+a^2}}\right); \tag{30a}$$

$$\int_{-\infty}^{\infty} \Phi(ax+b)^2\varphi(x)dx = \Phi_2\left(\frac{b}{\sqrt{1+a^2}}, \frac{b}{\sqrt{1+a^2}}; \frac{a^2}{\sqrt{1+a^2}}\right); \tag{30b}$$

$$\int_{-\infty}^{c} \Phi(ax+b)\varphi(x)dx = \Phi_2\left(\frac{b}{\sqrt{1+a^2}}, c; \frac{-a}{\sqrt{1+a^2}}\right). \tag{30c}$$

Proof: Consider a Gaussian random variable $\tilde{z}_1 \sim N(0, \sqrt{1+a^2})$. We write $\tilde{z}_1 = -a\tilde{x} + \delta_1$ where $\tilde{x}, \tilde{\delta}_1 \sim N(0,1)$ are independent. We notice that for some arbitrary constant b, $\mathsf{Prob}(\tilde{z}_1 \leq b) = \Phi(b/\sqrt{1+a^2})$. But, by the law of iterated expectations,

$$\begin{aligned} \mathsf{Prob}(\tilde{z}_1 \leq b) &= \mathsf{E}(1_{\tilde{z}_1 \leq b}) = \mathsf{E}\left(\mathsf{E}(1_{\tilde{z}_1 \leq b}|\tilde{x})\right) \\ &= \mathsf{E}\left(\mathsf{Prob}\left(\tilde{\delta}_1 \leq a\tilde{x} + b\right)\right) = \int_{-\infty}^{\infty} \Phi(ax+b)\varphi(x)dx. \end{aligned}$$

This proves (30a). Introduce now a second variable $\tilde{z}_2 = N(0, \sqrt{1+a^2})$ and write $\tilde{z}_2 = -a\tilde{x} + \delta_2$, where $\tilde{x}$ is defined above and $\tilde{\delta}_2 \sim N(0,1)$ is independent of both $\tilde{x}$ and $\tilde{\delta}_1$. The correlation between $\tilde{z}_1$ and $\tilde{z}_2$ is $a^2/(1+a^2)$, such that

$$\mathsf{Prob}(\tilde{z}_1 \leq b, \tilde{z}_2 \leq b) = \Phi_2\left(\frac{b}{\sqrt{1+a^2}}, \frac{b}{\sqrt{1+a^2}}; \frac{a^2}{1+a^2}\right),$$

where $\Phi_2(\cdot, \cdot; \rho)$ is the bivariate cumulative Gaussian distribution function for correlation ρ. But also

$$\mathsf{Prob}(\tilde{z}_1 \le b, \tilde{z}_2 \le b) = \mathsf{E}\left(\mathsf{Prob}(\tilde{z}_1 \le b, \tilde{z}_2 \le b|\tilde{x})\right) = \mathsf{E}\left(\mathsf{Prob}(\tilde{z}_1 \le b|\tilde{x})\mathsf{Prob}(\tilde{z}_2 \le b|\tilde{x})\right)$$

where we have relied on conditional independence in the second equality. Recognizing that $\mathsf{Prob}(\tilde{z}_1 \le b|\tilde{x}) = \mathsf{Prob}(\tilde{z}_2 \le b|\tilde{x}) = \Phi(a\tilde{x} + b)$ proves (30b).

Finally, to prove (30c) note that

$$\mathsf{Prob}(\tilde{z}_1 \le b, \tilde{x} \le c) = \int_{-\infty}^{c} \mathsf{Prob}(\tilde{z}_1 \le b|\tilde{x} = x)\varphi(x)dx = \int_{-\infty}^{c} \Phi(ax+b)\varphi(x)dx$$

On the other hand,

$$\mathsf{Prob}(\tilde{z}_1 \le b, \tilde{x} \le c) = \Phi_2\left(\frac{b}{\sqrt{1+a^2}}, c; \frac{-a}{\sqrt{1+a^2}}\right)$$

where we have used the fact that the correlation between $\tilde{z}_1$ and $\tilde{x}$ is $-a/\sqrt{1+a^2}$. ∎

An expression for the integral in (30b), but with a finite upper integration limit is also possible, but involves the three-dimensional Gaussian distribution and is likely easier to compute numerically.

Lemma 2 Let x be a standard Gaussian variate. For constants a and b define $\omega := a/\sqrt{1+a^2}$. Then

$$\mathsf{E}\left(x\Phi(ax+b)\right) = \frac{\omega e^{-\frac{1}{2}b^2(1-\omega^2)}}{\sqrt{2\pi}} = \omega\varphi\left(b\sqrt{1-\omega^2}\right). \tag{31}$$

Proof: Note that

$$\begin{aligned}
\mathsf{E}\left(x\Phi(ax+b)\right) &= \int_{-\infty}^{\infty} x\varphi(x)\Phi(ax+b)dx = -\int_{-\infty}^{\infty} \varphi'(x)\Phi(ax+b)dx \\
&= a\int_{-\infty}^{\infty} \varphi(x)\varphi(ax+b)dx - [\varphi(x)\Phi(ax+b)]_{-\infty}^{\infty} \\
&= \frac{a}{2\pi}\int_{-\infty}^{\infty} e^{-\frac{1}{2}(x^2+a^2x^2+2abx+b^2)}dx \\
&= \frac{a}{2\pi}e^{-\frac{1}{2}b^2(1-\omega^2)}\int_{-\infty}^{\infty} e^{-\frac{1}{2}\left(x\sqrt{1+a^2}+b\omega\right)^2}dx.
\end{aligned}$$

Clearly, then,

$$\mathsf{E}\left(x\Phi(ax+b)\right) = \frac{a}{\sqrt{2\pi}}e^{-\frac{1}{2}b^2(1-\omega^2)}\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sqrt{1+a^2}}e^{-\frac{1}{2}y^2}dy.$$ ∎

Lemma 3 Let x be a standard Gaussian variate. For constants a and b

$$\mathsf{E}\left(x^2\Phi(ax+b)\right) = \frac{-b\omega}{\sqrt{1+a^2}}\mathsf{E}\left(x\Phi(ax+b)\right) + \Phi(b/\sqrt{1+a^2}). \tag{32}$$

Proof: We have first

$$\mathsf{E}\left(x^2\Phi(ax+b)\right) = \int_{-\infty}^{\infty} x^2\varphi(x)\Phi(ax+b)dx = \int_{-\infty}^{\infty}\left(\varphi''(x)+\varphi(x)\right)\Phi(ax+b)dx.$$

Here

$$\begin{aligned}\int_{-\infty}^{\infty}\varphi''(x)\Phi(ax+b)dx &= -a\int_{-\infty}^{\infty}\varphi'(x)\varphi(ax+b)dx + \left[\varphi'(x)\Phi(ax+b)\right]_{-\infty}^{\infty} \\ &= a\int_{-\infty}^{\infty} x\varphi(x)\varphi(ax+b)dx\end{aligned}$$

and

$$\begin{aligned} a\int_{-\infty}^{\infty} x\varphi(x)\varphi(ax+b)dx &= \frac{a}{2\pi}e^{-\frac{1}{2}b^2(1-\omega^2)}\int_{-\infty}^{\infty} xe^{-\frac{1}{2}\left(x\sqrt{1+a^2}+b\omega\right)^2}dx \\ &= \frac{a}{\sqrt{2\pi}}e^{-\frac{1}{2}b^2(1-\omega^2)}\int_{-\infty}^{\infty}\left(\frac{y-b\omega}{\sqrt{1+a^2}}\right)\frac{1}{\sqrt{2\pi}\sqrt{1+a^2}}e^{-\frac{1}{2}y^2}dy \\ &= \frac{-b\omega^2\varphi\left(b\sqrt{1-\omega^2}\right)}{\sqrt{1+a^2}}.\end{aligned}$$

The result now follows from Lemmas 1 and 2. ∎

Lemma 4 Let x be a standard Gaussian variate. For constants a and b define $\omega := a/\sqrt{1+a^2}$ and $\vartheta := \omega/\sqrt{1+\omega^2}$. Then

$$\begin{aligned}\mathsf{E}\left(x^2\Phi(ax+b)^2\right) &= \Phi_2\left(\frac{b}{\sqrt{1+a^2}},\frac{b}{\sqrt{1+a^2}};\frac{a^2}{\sqrt{1+a^2}}\right) \\ &\quad +\frac{2\mathsf{E}\left(x\Phi(ax+b)\right)}{\sqrt{1+a^2}}\left[\vartheta\varphi\left(b(1-\omega^2)\sqrt{1-\vartheta^2}\right)-b\omega\Phi\left(\frac{b(1-\omega^2)}{\sqrt{1+\omega^2}}\right)\right]\end{aligned}$$

Proof: First

$$\mathsf{E}\left(x^2\Phi(ax+b)^2\right) = \int_{-\infty}^{\infty} x^2\varphi(x)\Phi(ax+b)^2dx = \int_{-\infty}^{\infty}\left(\varphi''(x)+\varphi(x)\right)\Phi(ax+b)^2dx.$$

Here

$$\begin{aligned}\int_{-\infty}^{\infty}\varphi''(x)\Phi(ax+b)^2dx &= -2a\int_{-\infty}^{\infty}\varphi'(x)\Phi(ax+b)\varphi(ax+b)dx + \left[\varphi'(x)\Phi(ax+b)^2\right]_{-\infty}^{\infty} \\ &= 2a\int_{-\infty}^{\infty} x\varphi(x)\Phi(ax+b)\varphi(ax+b)dx \\ &= \frac{2a}{2\pi}e^{-\frac{1}{2}b^2(1-\omega^2)}\int_{-\infty}^{\infty} x\Phi(ax+b)e^{-\frac{1}{2}\left(x\sqrt{1+a^2}+b\omega\right)^2}dx \\ &= \frac{2\mathsf{E}\left(x\Phi(ax+b)\right)}{\sqrt{1+a^2}}\int_{-\infty}^{\infty}(y-b\omega)\Phi\left(y\omega+b(1-\omega^2)\right)\varphi(y)dy.\end{aligned}$$

We can now use Lemma 2 to see that

$$\int_{-\infty}^{\infty} y\Phi\left(y\omega + b(1-\omega^2)\right)\varphi(y)dy = \vartheta\varphi\left(b(1-\omega^2)\sqrt{1-\vartheta^2}\right).$$

and the lemma follows by application of Lemma 1. ∎

Lemma 5 For a standardized Gaussian variable x and arbitrary constants a and b, we have

$$\mathsf{E}\left(1_{a<x\le b}x\right) = 1_{b\ge a}\left(\varphi(a)-\varphi(b)\right); \tag{33a}$$
$$\mathsf{E}\left(1_{a<x\le b}x^2\right) = 1_{b\ge a}\left(\Phi(b)-\Phi(a)\right) + 1_{b\ge a}\left(a\varphi(a)-b\varphi(b)\right). \tag{33b}$$

In particular, $\mathsf{E}\left(1_{x\le b}x\right) = -\varphi(b)$; $\mathsf{E}\left(1_{x>a}x\right) = \varphi(a)$; $\mathsf{E}\left(1_{x>a}x^2\right) = a\varphi(a) + (1-\Phi(a))$; *and* $\mathsf{E}\left(1_{x\le b}x^2\right) = \Phi(b) - b\varphi(b)$.

Proof: If $b < a$, the expectations (33a)-(33b) are obviously zero, consistent with the given expressions. For $b \ge a$, we have, trivially,

$$\mathsf{E}\left(1_{a<x\le b}x\right) = 1_{b>a}\int_a^b x\varphi(x)dx = -1_{b>a}\int_a^b \varphi'(x)dx = 1_{b>a}\left(\varphi(a)-\varphi(b)\right),$$

and

$$\begin{aligned}\mathsf{E}\left(1_{a<x\le b}x^2\right) &= 1_{b>a}\int_a^b x^2\varphi(x)dx = 1_{b>a}\int_a^b \left(\varphi''(x)+\varphi(x)\right)dx \\ &= 1_{b>a}\left(\Phi(b)-\Phi(a)\right) + 1_{b>a}\left(a\varphi(a)-b\varphi(b)\right).\end{aligned}$$ ∎

References

Acharya, V., Bharath, S., and A, S. (2003). Understanding the recovery rates on defaulted securities. Working paper, London Business School.

Altman, E. I., Brady, B., Resti, A., and Sironi, A. (2003). The Link between Defaults and Recovery Rates: Theory, Empirical Evidence, and Implications. Working paper, Stern School of Business.

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

Duffie, D. and Garleanu, N. (2001). Risk and valuation of collateralized debt obligations. *Financial Analysts Journal*, 57(1): 41–59.

Embrechts, P., Lindskog, F., and McNeil, A. (2001). Modeling dependence with copulas and applications to risk management. Working paper, Department of Mathematics, ETHZ.

Frey, R. and McNeil, A. J. (2003). Dependent defaults in models of portfolio credit risk. *Journal of Risk*, (forthcoming).

Frye, J. (2000). Collateral damage. *RISK*, April: 91–94.

Frye, J. (2003). A false sense of security. *RISK*, August: 63–67.

Giesecke, K. (2003). A simple exponential model for dependent defaults. Working paper, Cornell University.

Gregory, J. and Laurent, J.-P. (2003). I will survive. *RISK*, June.

Gregory, J. and Laurent, J.-P. (2004). In the core of correlation. Working paper, BNP Paribas.

Gupton, G. and Stein, R. (2002). Losscalc: Moody's model for predicting loss given default (lgd). Working paper, Moody's Investor Services.

Hu, Y. and Perraudin, W. (2002). The dependence of recovery rates and defaults. Working paper, BirkBeck College.

Jarrow, R. (2001). Default parameter estimation using market prices. *Financial Analysts Journal*, 57(5): 75–92.

Jokivuolle, E. and Peura, S. (2000). A model for estimating recovery rates and collateral haircuts for bank loans. Working paper, Bank of Finland.

Li, D. X. (2000). On default correlation: a copula function approach. Working Paper 99-07, The Riskmetrics Group.

Lindskog, F., McNeil, A., and Schmock, u. (2003). Kendall's tau for elliptical distributions. In *Credit Risk: Measurement, Evaluation, and Management*, 149–156. Physica-Verlag, Heidelberg.

Merton, R. (1974). On the pricing of corporate debt: the risk structure of interest rates. *Journal of Finance*, 29: 449–470.

Pykhtin, M. (2003). Unexpected recovery risk. *RISK*, August: 74–78.

Schonbucher, P. (2001). Factor models: Portfolio credit risks when defaults are correlated. *Journal of Risk Finance*, 3(1): 45–56.

Vasicek, O. (2003). Loan portfolio value. *RISK*, December: 160–162.