Portfolio Losses in Factor Models: Term Structures and Intertemporal Loss Dependence

Leif Andersen[1]
Bank of America

First version: March 24, 2006

This version: September 09, 2006

Abstract

Due to their computational efficiency, simple factor models remain popular in the pricing of credit portfolio derivatives. In this paper, we continue the elaboration on the fundamental structure of factor models initiated in [3], with a special focus on term structure effects. We describe a number of techniques to understand, and to improve control over, portfolio loss distribution term structures and intertemporal loss correlation. As part of our analysis, we introduce an extension of the RFL model ([2]) to incorporate jumps in the systematic factor and in firm residuals. We also numerically test the dependence of forward-starting synthetic CDOs on the correlation of losses across time. Finally, our analysis highlights the fact that several of the models suggested in the literature are essentially equivalent.

1 Introduction

For the typical credit basket with 100 or more underlying risky obligors, co-dependence modeling and synthetic CDO valuation constitute difficult problems, with many modeling frameworks demanding very significant computing resources for calibration, pricing, and risk generation. In the interest of computational efficiency, it is very convenient to assume the existence of a low-dimensional set of random variables – or *factors* – conditional on which obligor defaults on some time interval become mutually independent. As shown in, e.g., [7] and [4], the structure of the resulting model allows for the application of recursion arguments that lead to large gains in the efficiency of CDO computations, especially for risk measures. Many standard models (e.g. the Gaussian and Student t copulas) are special cases of the general factor framework, as are several newer models designed to better match observable prices in the quoted markets for index tranches. For examples, see [2], [11], [12], [14], [9], [6], and [19], among many others. [1] lists unifying results for Levy-type default drivers, and [3] contains a more general survey and further references, as well as a discussion of a number of generic issues associated with the factor setup, most notably in the computation of meaningful hedge information.

[1] `leif.andersen@bofasecurities.com`. The author is indebted to Jakob Sidenius and Victor Piterbarg for their insights and suggestions. The author is also grateful for useful comments made by participants at the May 2006 ICBI Global Derivatives Conference in Paris.

Most factor models appearing in the literature are at heart one-period models that aim primarily to generate portfolio loss distributions at a single fixed maturity, typically coinciding with that of a particular CDO tranche under consideration. Loss distributions at earlier dates can in many cases be implied from the resulting model, but there is little control over these distributions which, as discussed in [3], may look quite different from the term structure of loss distributions that can be implied from, say, CDX and I-Traxx CDO tranches quoted at multiple maturities[1]. [13] and many others have made similar observations.

As discussed further later in the paper, a regular synthetic CDO tranche (and a CDO-squared contract, for that matter) can be considered a one-period product, so for these products the inability of simple factor models to control loss distribution term structures is often (but not always) only a minor nuisance. Recently, however, several new credit basket securities with strong timing features have emerged. Examples include forward starting tranches, options on CDO tranches, leveraged super-senior tranches, and reset tranches. While liquidity in most of these exotic tranches is currently quite low, their emergence has nevertheless triggered research into true dynamic loss models, exemplified for instance in [17] and [18]. Such models, however, are quite data-demanding and remain somewhat speculative at the present time. Given this, it is quite tempting to consider whether the application domain of simple factor models possibly could be stretched to cover securities with a stronger temporal component than that of regular CDOs. Indeed, some banks appear to promote application of standard factor models for a variety of quite complicated exotic structures; see e.g. [5]. As we shall see, this practice will require considerable care.

In this brief paper we shall analyze several fundamental issue of factor models, most notably their behavior in modeling term structures of portfolio losses. Our approach is primarily descriptive, but we do also offer certain techniques and ideas aimed to alleviate some of the issues we uncover. The paper is organized as follows. In Section 1, we briefly provide notation for the standard one-period factor model and then proceed, in Section 2, with an analysis of the multi-factor case. Based on whether a model generates a true default time copula, we distinguish between complete and incomplete multi-period factor models; both cases are illustrated with several examples. To demonstrate common modeling procedures – and the snags associated with them – Section 2 develops a new extension of the radnom factor loading (RFL) model of [2] to include jumps in both residuals and factors. This exercise, which may be of independent interest, demonstrates how easy it is to develop new factor models (as is evident from the steady stream of papers introducing new twists on the factor model idea), but also highlights the fact that many existing models ultimately share common characteristics, to the extent that several existing models can be classified as special cases of other models. In Section 3, we proceed to the problem of matching a given term structure of marginal loss distributions. Merely matching marginal loss distributions is, however, of limited value when pricing securities that depend on the dynamic evolution of losses. In Section 4 we focus on this

[1]The trader-friendly base correlation framework introduced in [15] in fact often implies arbitrages (= positive probabilities of negative losses) in pre-maturity loss distributions.

issue by specifically introducing intertemporal correlation parameters. Numerical tests on a forward-starting CDO highlights the practical importance of controlling loss correlation across time. Finally, Section 5 concludes the paper.

Finally, a few words about market reality. Standard credit default swaps generally trade with good liquidity, and can be used to uncover risk-neutral survival probability curves using standard methods (e.g. bootstrapping or similar). We throughout the paper assume that risk-neutral survival probabilities are known with certainty for all firms. Market-prices for some tranches – namely those on certain indices such as CDX, I-Traxx, etc. – are observable at selected maturites (typically 3, 5, 7 and 10 years), and are obvious calibration targets for portfolio loss models in cases where we are considering non-standard tranches on portfolios identical (or near-identical) to those of the stadard indices. In practice, calibration to index tranches typically is done cross sectionally (to multiple tranches) but often only a single maturity at a time; the methods in this paper will be of use for calibration to multiple maturities simultaneously. We should not pretend, however, that good calibration targets exist for all CDO tranches. Many bespoke portfolios allows for very little tranche price discovery, in which case models will have to be calibrated in more direct fashion (by outright specification of model parameters) or by calibration against prices which are somehow "extrapolated" from known index tranches. WDiscussion of loss term structures is, of course, most relevant when we can actually observe loss distributions in the market, and much of the discussion in this paper is most easily quantified in practice for tranches for which there is good price discovery. Nevertheless, all of the issues we indentify in this paper apply to bespoke tranches, even though parameterization of models for these instruments remains a challenging exercice.

2 Factor models: one-period setting

Let us consider a general factor framework, to be applied on a horizon $[0,T]$. We have identified a systematic d-dimensional random variable Z_T, with (generalized) density $\phi_T : \mathscr{D} \to \mathbb{R}_+$, where $\mathscr{D}$ is some subset of $\mathbb{R}^d$. Consider an N-firm portfolio, and use τ_i to denote the default time for firm i. Z_T-conditional survival on $[0,T]$ for the firms in the basket are assumed independent events, with risk-neutral probabilities

$$\Pr\left(\tau_i > T | Z_T = z\right) = f_T^{(i)}(z), \quad i = 1, \ldots, N, \tag{1}$$

where $f_T^{(i)} : \mathscr{D} \to [0,1]$. We obviously have

$$\Pr\left(\tau_i > T\right) = \int_{\mathscr{D}} f_T^{(i)}(z)\phi_T(z)\, dz, \tag{2}$$

which puts consistency constraints on $f_T^{(i)}(z)$, as the left-hand side of this equation is observable in the market for default swaps. Often, we set

$$f_T^{(i)}(z) = f_T^{(i)}\left(z; H_T^{(i)}\right)$$

for some free scalar parameter $H_T^{(i)}$ to be calibrated such that (2) is satisfied. Obviously, such calibration is particularly convenient if (2) can be evaluated in closed form.

The form of the factor model setup allows for efficient computation of the distribution of the portfolio loss

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

where we for simplicity have assumed \$1 loss-given-default for each firm[2]. Specifically, we can build the Z_T-conditional N-name portfolio loss distributions incrementally, adding firms one at a time to an initially empty basket. If $L^{(j)}(T)$ is the loss associated with a basket containing firms $i = 1, \dots, j$, we get, from conditional independence,

$$\Pr\left(L^{(j+1)}(T) = n | Z_T = z\right) = \Pr\left(L^{(j)}(T) = n | Z_T = z\right) f_T^{(j+1)}(z) \\ + \Pr\left(L^{(j)}(T) = n-1 | Z_T = z\right)\left(1 - f_T^{(j+1)}(z)\right),$$

where n can take on the values $0, 1, \dots, j+1$. The recursion starts at $\Pr\left(L^{(0)}(T) = n\right) = 1_{n=0}$. The unconditional loss distribution can be recovered as

$$\Pr\left(L(T) = n\right) = \int_{\mathscr{D}} \Pr\left(L^{(N)}(T) = n | Z_T = z\right) \phi_T(z)\, dz.$$

Efficient techniques for computation of sensitivities exist in this setting; see [4]. See also [8] for a transform-based method.

3 Factor models: multi-period setting

3.1 Basic Setup

In many applications – such as the pricing of synthetic CDO tranches – we require knowledge of the distribution of $L(t)$ on either a continuum of dates, or at least on a discrete time line $\{T_j\}_{j=1}^{M}$; see [4] for details. Focusing on the discrete case, to complete pricing in a factor setting, we can introduce M different d-dimensional variables $Z_{T_1}, Z_{T_2}, \dots, Z_{T_M}$ with M separate densities $\phi_{T_1}, \dots, \phi_{T_M}$. For each firm we then need M conditional survival probabilities

$$f_{T_j}^{(i)}(z), \quad j = 1, \dots, M,$$

where we impose independence of the events $1_{\tau_i > T_j}$ and $1_{\tau_k > T_j}$ for all $i \neq k$ when we condition on Z_{T_j}. We also still impose the consistency requirement

$$\Pr\left(\tau_i > T_j\right) = \int_{\mathscr{D}} f_{T_j}^{(i)}(z) \phi_{T_j}(z)\, dz, \quad j = 1, \dots, M,$$

[2] Quantization methods to relax this assumption are straightforward; see e.g. [4] for a simple approach.

where we implicitly have assumed that all Z_{T_j}'s have the same domain $\mathscr{D}$. For the resulting model to make sense (and to be void of arbitrage), we must further insist that, for all values of $n = 0, 1, \ldots,$,

$$T_k > T_j \Rightarrow \Pr(L(T_k) \geq n) \geq \Pr\left(L(T_j) \geq n\right), \tag{3}$$

where $L(T_k)$ and $L(T_j)$ can be computed from the recursion discussed earlier. This latter condition can be quite constraining – and rather difficult to check – for arbitrary model specifications, but may sometimes be naturally satisfied for models that originate out of, say, a structural default framework.

3.2 Shortcomings

While the multi-period factor specification in Section 3.1 above does lock down marginal loss probabilities and allows for pricing of a "static" security such as a CDO tranche, it generally fails to uniquely fix the joint distribution of default times τ_i, $i = 1, \ldots, N$. While sometimes classified otherwise, the specification in Section 3.1 does thereby generally *not* constitute a copula representation of default times. To expand on this, consider just the case where $M = 2$, and let us contemplate the computation of $\Pr\left(\tau_i > T_1, \tau_j > T_2\right)$. We note that the simpler task of computing, say, $\Pr\left(\tau_i > T_1, \tau_j > T_1\right)$, $i \neq j$, is easily accomplished by simply writing

$$\begin{aligned}
\Pr\left(\tau_i > T_1, \tau_j > T_1\right) &= \int_{\mathscr{D}} \Pr\left(\tau_i > T_1, \tau_j > T_1 | Z_{T_i} = z\right) \phi_{T_1}(z)\, dz \\
&= \int_{\mathscr{D}} f_{T_1}^{(i)}(z)\, f_{T_1}^{(j)}(z) \phi_{T_1}(z)\, dz,
\end{aligned}$$

by the assumed independence of the events $1_{\tau_i > T_1}$ and $1_{\tau_j > T_1}$ when conditioned on Z_{T_i}. Without further structure on the problem, however, we cannot extend this methodology to two horizons simultaneously – as required for computation of $\Pr(\tau_i > T_1, \tau_j > T_2)$ – as we have not described how, say, the event $1_{\tau_i > T_1}$ may be affected by Z_{T_2}. A natural way forward would be to make the assumption that the events $1_{\tau_i > T_1}$ and $1_{\tau_j > T_2}$ are independent when conditioned on *both* Z_{T_1} and Z_{T_2}, such that

$$\begin{aligned}
\Pr\left(\tau_i > T_1, \tau_j > T_2\right) &= \int_{\mathscr{D}^2} \Pr\left(\tau_i > T_1, \tau_j > T_2 | Z_{T_i} = z_1, Z_{T_2} = z_2\right) \phi_{T_1,T_2}(z_1, z_2)\, dz \\
&= \int_{\mathscr{D}^2} \Pr\left(\tau_i > T_1 | Z_{T_i} = z_1, Z_{T_2} = z_2\right) \Pr\left(\tau_j > T_2 | Z_{T_i} = z_1, Z_{T_2} = z_2\right) \phi_{T_1,T_2}(z_1, z_2)\, dz
\end{aligned}$$

where ϕ_{T_1,T_2} is the joint density of Z_{T_1} and Z_{T_2}. This computation, however, cannot be accomplished from our standard factor setup in Section 3.1, which neither prescribes the joint density of the Z_{T_i}'s, nor does it specify conditional probabilities of the type

$$\Pr\left(\tau_i > T_1 | Z_{T_i} = z_1, Z_{T_2} = z_2\right).$$

3.3 A Generalization

The lack of a default time copula is a rather severe practical drawback, as Monte Carlo simulation of joint default times of portfolio obligors becomes impossible. This, in turn, precludes usage of the model for most non-vanilla securities, including CDO-squared contracts. To ensure that our factor model allows for construction of a default time copula, we need to add more structure to the rather loose setup in Section 3.1. For this, imagine aggregating all Z_{T_i}'s into a state vector $S = (Z_{T_1}, \ldots, Z_{T_M})$ of dimension[3] $M \times d$ for which we now assume that the joint density ϕ_S is known. Let $\mathscr{D}^*$ be the domain for S, and assume that conditional on S, the default times of all firms are independent. Also, for all $t \geq 0$, let conditional survival probabilities be characterized by functions $q_i : \mathbb{R}_+ \times \mathscr{D}^* \to [0,1]$, given by

$$\Pr(\tau_i > t | S = s) = q_i(t,s), \quad i = 1, \ldots, N, \quad t \geq 0,$$

and subject to

$$\Pr(\tau_i > t) = \int_{\mathscr{D}^*} q_i(t,s) \phi_S(s)\, ds.$$

In this setting – which we can call a *complete factor model* specification – a default time copula automatically arises, as

$$\Pr(\tau_1 > t_1, \ldots, \tau_N > t_N) = \int_{\mathscr{D}^*} \prod_{i=1}^{N} q_i(t_i,s) \phi_S(s)\, ds.$$

The default time copula arising from this can be written down easily explicitly from Sklar's Theorem (an exercise that we omit).

For the specification above to be meaningful, we impose restrictions which were not needed in Section 3.1. Specifically, we want

1. For all $i = 1, \ldots, N$ and all $s \in \mathscr{D}^*$, $t' > t \Rightarrow q_i(t',s) < q_i(t,s)$
2. For all $i = 1, \ldots, N$ and all $s \in \mathscr{D}^*$, $q_i(0,s) = 1$ and $q_i(\infty,s) = 0$.

These two requirements ensures that a) the portfolio loss distribution will satisfy the monotonicity constraint (3); and b) $\Pr(\tau_i = 0) = \Pr(\tau_i = \infty) = 0$. We also ensure that Monte Carlo simulation of joint default times $\tau_1, \ldots, \tau_N$ is possible, by the following algorithm:

1. Draw a random sample S, as prescribed by the density ϕ_S
2. Draw N uniform samples $U_1, \ldots, U_N$
3. For $i = 1, \ldots, N$, set $\tau_i = q_i^{-1}(U_i, S)$, where we have defined q_i^{-1} from the relation $q_i\left(q_i^{-1}(x,S), S\right) = x$.

[3] Extensions to the case where the different Z_{T_i}'s have different dimension are trivial, of course.

The inversion in Step 3 can here always be carried out, due to the requirements 1 and 2 on q_i listed above.

The reader may reasonably ask how one goes about specifiying the N issuer-specific functions q_i in practice. There are several approaches, some of which involve functional forms with one or more issuer-specific parameters and some of which arise from a more structural approach. The next few sections give examples; more can be found in [3]. Our focus here is on relatively simple models that are have sufficient computational efficiency for practical tasks such as calibration and hedge computations.

3.4 Example: An incomplete factor model

A multi-period factor model that has arisen many times in the literature involves variations on a setting where default is generated by driver processes

$$X_i(t) = \beta_i Y(t) + e_i(t), \quad i = 1,\ldots,N. \tag{4}$$

Here β_i is a firm-specific constant, $Y(t)$ is a one-dimensional[4] systematic process shared by all drivers, and $e_i(t)$ is a residual process idiosyncratic to firm i and independent of Y and e_j, $j \neq i$. $Y(t)$ is often loosely considered a proxy for the state of the "market", giving (4) some likeness to the CAPM setup. Assume that the density of $Y(t)$ is $\phi_Y(\cdot;t)$ and all the $e_i(t)$ have (time-dependent) cumulative distribution function $F_i(\cdot;t)$. The driver process $X_i(t)$ will trigger default at the first passage-time of a (curved) barrier $h_i(t)$, i.e.

$$\tau_i = \inf\{t : X_i(t) \leq h_i(t)\}, \quad i = 1,\ldots,N. \tag{5}$$

The location of $h_i(t)$ must (often at significant numerical expense) be calibrated such that the first hitting time τ_i has the prescribed cumulative probability distribution.

While the setup above constitutes a valid dynamic model with a well-defined joint default time copula, calibration of the barrier $h_i(t)$ is often quite involved, and generation of default-times from (5) can be computationally expensive. To improve efficiency and tractability, it is tempting to alter the model to

$$\Pr(\tau_i \leq t) = \Pr(X_i(t) \leq H_i(t)), \quad i = 1,\ldots,N, \tag{6}$$

with a new function $H_i(t)$ to be calibrated against given default probabilities; see e.g. [6] and [19] for examples of this procedure. With (6) we get

$$\Pr(\tau_i \leq t | Y(t) = y) = \Pr(e_i(t) \leq H_i(t) - \beta_i y) = F_i(H_i(t) - \beta_i y; t)$$

and, when discretizing on a time grid, end up in the framework of Section 3.1 with $Z_{T_j} = Y(T_j)$ and

$$f_{T_j}^{(i)}(z) = 1 - F_i\left(H_i(t) - \beta_i z; T_j\right).$$

[4] In this and many other examples, we shall work with a one-dimensional systematic factor $(d = 1)$. This choice is made for convenience only, and it should generally be transparent how to extend to a case where $d > 1$.

The resulting model allows for generation of marginal loss distributions, but contains no information about the joint distribution of default times and cannot be used to simulate paths of default times. This should not be a surprise, as (6) amounts to suppressing all transition dynamics from the original model (5), leading to a setup with no well-defined default dynamics and no method of sampling default times.

3.5 Example: A complete factor model

Consider a setup similar to (4), but now remove time-dependence from all terms, such that we write

$$X_i = \beta_i Y + e_i. \tag{7}$$

A specification such as (5) is obviously not possible anymore, but we can still use (6):

$$\Pr(\tau_i \leq t) = \Pr(X_i \leq H_i(t)), \quad i = 1, \ldots, N.$$

It is clear that

$$\Pr(\tau_i > t | Y = y) = 1 - F_i(H_i(t) - \beta_i y)$$

which now depends only on time through $H_i(t)$. Making the additional assumption that all firm default times are independent when we condition on Y, the joint distribution of default times is here recoverable:

$$\Pr(\tau_1 > t_1, \ldots, \tau_N > t_N) = \int_{\mathbb{R}} \prod_{i=1}^{N} (1 - F_i(H_i(t_i) - \beta_i y)) \, \phi_Y(y) \, dy. \tag{8}$$

In the setting of Section 3.3, the setup above has $Z_{T_1} = Z_{T_2} = \ldots = Y$, such that the state vector S degenerates into a single variable, $S = Y$. This, in turn, implies that the functions q_i become known from the given f_{T_j}, allowing for construction of the joint default time distribution. It should be obvious, however, that the "dynamics" associated with the specification of (7) are rather suspect.

We should note that the standard (one-factor) Gaussian copula (Gaussian Y and Gaussian e_i's) fall into this setting of this example, as does, say, the RFL model (see [2]) and the NIG copula (see [14] and [9]). For concreteness, the latter two models are characterized by Gaussian e_i's and

- RFL: $Y = a(U)U$, where U is a Gaussian distribution and a some function[5] $a : \mathbb{R} \to \mathbb{R}$.
- NIG: Y follows a Normal inverse Gaussian distribution

It is clear that the NIG model amounts to a special case of the RFL model, as we can always, by trivial means, construct the function a such that the factor density ϕ_Y of the RFL model matches that of the NIG model; we leave this exercise to the reader. From

[5] The function a can be firm-specific, as long as the Gaussian argument U is shared among the firms. Extensions to multi-dimensional U can be found in [2].

(8) it follows that the NIG-implied joint distribution function of default times can always be replicated by a special case of the RFL model. Indeed, if residuals e_i are Gaussian, *all* models of the type (7) fall in the RFL class[6]. This includes the model in [19], say.

3.6 Example: RFL Model with Jumps

There has recently been interest in using jump-style processes for the purpose of generating realistic conditional default specifications; see e.g. [6], [19], and [1], among several others. We can use such a jump-setting to make the discussion in Sections 3.4 and 3.5 more concrete. Specifically, we wish to develop an extension of the RFL model in [2], to allow for jumps in both residuals and systematic variables. By incorporating both jumps and state-dependent correlation, the resulting model is quite rich in its dependence structure, and as such may be of independent interest.

As a starting point, consider a setting where default for a specific firm i is governed by a driver process X_i of the form

$$dX_i(t) = a(M(t))\,dM(t) + \sigma_i dW_i(t) + J_i(t)dN_i(t), \quad i = 1,\ldots,N, \tag{9}$$

$$dM(t) = \sigma dW(t) + J(t)dN(t). \tag{10}$$

In this setup, we have introduced:

- $M(t)$: A systematic factor process, representing general market conditions
- $a(M(t))$: A function $\mathbb{R} \to \mathbb{R}$ determining the strength of the coupling of X_i to the systematic process M
- $W(t)$: A Brownian motion, representing smooth evolution of general market conditions
- $N(t)$: A Poisson process with intensity λ, representing sudden jumps in general market conditions. $J(t)$ represents the jump in M upon a jump in N, and is assumed Gaussian with distribution $J \sim N(\mu, \gamma^2)$.
- $W_i(t)$: A Brownian motion, representing smooth evolution of idiosyncratic credit factors, unique to firm i.
- $N_i(t)$: A Poisson process with intensity λ_i, representing jumps in idiosyncratic credit factors, unique to firm i. $J_i(t)$ represents the jump in $X_i(t)$ upon a jump in N_i, and is assumed Gaussian with distribution $J_i \sim N(\mu_i, \gamma_i^2)$.

[6] Despite the equivalence mentioned in the text, we highlight a "philosophical" difference between RFL and models such as NIG: in the former, the distribution of the market factor Y is universal (Gaussian for basic RFL, but see Section 3.6) with calibration of an individual portfolio to market accomplished by altering the specific way the firms in the portfolio are affected by the market; in the latter, calibration of each basket to market involves altering the distribution of Y in a basket-specific way, ultimately resulting in a model with one different market variable per portfolio. For multi-portfolio products (such as CDO-squared contracts), the RFL viewpoint seems more natural, we believe.

We assume that $W(t)$, $N(t)$, $J(t)$, $W_i(t)$, $N_i(t)$, and $J_i(t)$ (for all $i = 1,\ldots,N$) are independent. Default for firm i is assumed to take place at the first hitting-time

$$\tau_i = \inf\{t : X_i(t) \le h_i(t)\}, \quad i = 1,\ldots,N,$$

where h_i is a (firm-specific) deterministic function of time. While we for simplicity of notation do not assume so here, the function a can be allowed to be firm-specific with no additional complications.

3.6.1 One-Period Factor Model

Calibration of the firm-specific barriers $H_i(t)$ to match default probabilities is generally tricky, given the complexity of the full-blown process. To simplify, let us just consider a single-period model on some horizon $[0,T]$. As a proxy, we can, for instance, write (assuming arbitrarily that $M(0) = X_i(0) = 0$)

$$X_i(T) = a(M(T))M(T) + \sigma_i W_i(T) + Q_i(T), \tag{11}$$

$$M(T) = \sigma W(T) + Q(T), \tag{12}$$

where we have defined variables

$$Q(T) = \sum_{j=1}^{N(T)} J^{(j)}; \quad Q_i(T) = \sum_{j=1}^{N_i(T)} J_i^{(j)},$$

with each (independent) $J^{(j)} \sim N(\mu,\gamma^2)$ and each (independent) $J_i^{(j)} \sim N(\mu_i,\gamma_i^2)$. Then, we proceed to write (as in (6)), as an approximation,

$$\Pr(\tau_i < T) = \Pr\{X_i(T) \le H_i(T)\},$$

for some firm-specific barrier $H_i(T)$. In associating our model with a conditional survival probability function (as in Section 2), we can set our systematic factor Z_T to either $Z_T = M(T)$ or $Z_T = (W(T), Q(T))$. Choosing the former representation, we note that $\sigma W(T) + Q(T) \sim N(k\mu, k\gamma^2 + \sigma^2 T)$ when conditioned on the event $N(T) = k$. Thereby, we can recover the systematic factor density

$$\phi_T(z) = \sum_{k=0}^{\infty} p(T,k) \frac{\exp\left(-\frac{1}{2}\left(\frac{z-k\mu}{v(k)}\right)^2\right)}{v(k)\sqrt{2\pi}}, \tag{13}$$

where we have defined

$$p(T,k) = \Pr(N(T) = k) = e^{-\lambda T}\frac{(\lambda T)^k}{k!},$$

$$v(k) = \sqrt{\sigma^2 T + k\gamma^2}.$$

Using the same technique, for $i = 1,\ldots,N$ we can compute conditional survival probabilities:

$$\begin{aligned} f_T^{(i)}(z) &= \Pr(\tau_i > T | M(T) = z) = 1 - \Pr(\sigma_i W_i(T) + Q_i(T) \le H_i(T) - a(z)z) \\ &= 1 - \sum_{k=0}^{\infty} p_i(T,k) \Phi\left(\frac{H_i(T) - a(z)z - k\mu_i}{\sqrt{\sigma_i^2 T + k\gamma_i^2}}\right), \end{aligned} \quad (14)$$

with

$$p_i(T,k) = \Pr(N_i(T) = k) = e^{-\lambda_i T} \frac{(\lambda_i T)^k}{k!}. \quad (15)$$

We note in passing that it may sometimes be useful to abandon the pure Poisson setup, for a simpler setup where we specify directly the jump probabilities[7] $p_i(T,k)$ and $p(T,k)$. This approach – which is akin to simply using a Gaussian mixture for the distribution of the systematic factor – would, for instance, allow us to specifically limit the number of jumps possible, thereby ensuring that the sum in (14) is finite[8]. That said, for typical parameter values, the infinite sums in (14) and (13) will typically converge very quickly, requiring the evaluation of only a small (say, 3 or 4) number of terms.

Establishing the function $H_i(T)$ to satisfy (2) must generally be done numerically. For some specifications of the function a, however, an analytical approach is possible. For instance, if we consider the case where $a(z)$ is piecewise flat, the cumulative distribution function of $X_i(T)$ can be written down in closed-form – see Appendix A for the details – and $H_i(T)$ can be determined by straightforward inversion. The assumption of piecewise flat a entails little loss of generality, as any choice of a can be approximated to arbitrary precision by piecewise flat basis functions.

3.6.2 Multi-Period Extension

Extension of the model above to multiple horizons could proceed as in Section 3.1, by direct application of (14) and (13) on a grid of dates $\{T_j\}_{j=1}^{M}$. For reasons already discussed, the resulting model would, however, not constitute a complete factor model and would be of limited use beyond generation of marginal loss distributions. To circumvent this, we can proceed as in Section 3.5 and remove time dependence from the process X_i. The scaling of the individual parameters in this exercise is somewhat arbitrary, but we can, for instance, simply assume that the default driving process X_i is

$$X_i = X_i^*(T_M),$$

where X_i^* is generated according to the prescription (9)-(10), with T set equal to the final date in the time-line, $T = T_M$. In the notation of Section 3.3, we would have $S = M(T_M)$,

[7] Indeed, for even more flexibility, we can use a different μ (or μ_i) and γ (or γ_i) for each value of k. The physical interpretation of the model obviously becomes a little less clear with such extensions.

[8] Taken to the extreme, this approach allows us to specify the distribution of Z as being completely discrete with exogenously specified values and probabilities.

and

$$q_i(t,s) = \Pr(\tau_i > t | M(T_M) = s) = 1 - \sum_{k=0}^{\infty} p_i(T_M,k)\Phi\left(\frac{H_i(t) - a(s)s - k\mu_i}{\sqrt{\sigma_i^2 T_M + k\gamma_i^2}}\right). \quad (16)$$

The relevant factor density ϕ_s can be computed directly from (13), with $T = T_M$.

3.6.3 Some Comments

While the model outlined above serves here primarily to illustrate general concepts, let us nevertheless briefly contemplate practical parameterizations. To avoid having to estimate an excessive amount of parameters, it would probably often be preferable to have all λ_i, μ_i, and γ_i be firm-independent constants. As a further simplification, it may often be reasonable to simply set, for all i, $\mu_i = \mu$ and $\gamma_i = \gamma$, where we recall that μ and γ are parameters for $J(t)$, the systematic jump magnitude. It would also be convenient to reparameterize, and write $\lambda = r\lambda_{tot}$ and $\lambda_i = (1-r)\lambda_{tot}$ for two non-negative constants λ_{tot} and r, where $0 \le r \le 1$. λ_{tot} would then represent the total arrival rate of jumps in the default driver X_i, and r would represent a measure of ”jump-correlation”.

Note that, in this reduced-parameter setting discussed above, the only firm-specific parameters are the barrier functions H_i. While this is standard practice – firm-specific default co-dependence parameters are invariable hard to parameterize – we are obviously free to relax this. Tests of various parameterization strategies against actual tranche market prices is an interesting subject, but one that we must leave for future research.

Consistent with earlier comments (in Section 3.5), we should note that the specification (16) only truly generalizes the standard jump-free RFL setup if at least one λ_i is non-zero, i.e. if at least one residual has jumps (and thereby is non-Gaussian). Inclusion of jumps in the systematic factor M is, strictly speaking not necessary, as a suitable modification of the function a can replicate the effect of having jumps in M. Nevertheless, it may be convenient to work with jumps in the systematic factor, as, conceptually speaking, we can let the jump-part take care of inducing extreme scenarios (as is required, for instance, to match typical super-senior tranche prices), and then let the function a be used for fine-tuning of the model against observed CDO market prices. With a jump-process, it also becomes more natural to have a multi-peaked distribution of the common factor, which often helps in calibration of long-dated (10 year) index tranches where the market is often displaying strong market segmentation between the traded tranches[9]. In multi-basket applications (such as CDO-squared), the ability to independently specify a basket-specific factor loading function (the function a) while still retaining flexibility in setting the distribution parameters of basket-independent “global” market parameter (M) might also offer benefits; see Footnote 6.

[9] Such segmentation can, of course, also be emulated by having a spiky a-function in an ordinary RFL model (in which case the distribution of $a(Z)Z$ becomes multi-peaked)

4 Term Structures of Loss Distributions

The model class outlined in Example 3.5 (and further exemplified in Section 3.6.2) leads to a usable model where default times of individual firms can be simulated. However, as we already discussed there are several problems associated with the style of modeling employed. For one, by having time-independent random variables govern default behavior at all horizons, the model has a strong one-period flavor, and lacks the ability to control portfolio loss distributions at multiple horizons, both marginally and intertemporally. To avoid misunderstandings, let us emphasize that the model, through the calibration to individual credit spread curves, certainly matches *expected* portfolio losses at all horizons. We are, however, interested in the entire loss distribution, not just its mean.

Let us, for now, not worry about intertemporal correlation, and simply consider the notion of extending a single-factor model, such as the one in Example 3.5, to better match given term structures of marginal loss distributions. For several CDS indexes, synthetic CDO tranche prices are quoted with decent liquidity at multiple maturities (typically 3, 5, 7, and 10 years); in such cases, the target loss distribution term structures could be implied directly from market observations. To proceed, an obvious idea would involve introducing horizon-dependence in the model specification, for instance by making the parameter β_i horizon-specific. In the setting of Section 3.3, this would lead to conditional survival probabilities of

$$q_i(s;t) = 1 - F_i\left(H_i(t) - \beta_i(t)s\right), \quad i = 1, \ldots, N, \tag{17}$$

where F_i again is the distribution function of firm i residuals. To test whether this is a reasonable specification, we notice that

$$\begin{aligned} sign\left(q_i(s;t') - q_i(s;t)\right) &= sign\left(F_i\left(H_i(t) - \beta_i(t)s\right) - F_i\left(H_i(t') - \beta_i(t')s\right)\right) \\ &= sign\left(H_i(t) - H_i(t') + s\left(\beta_i(t') - \beta_i(t)\right)\right), \end{aligned}$$

where the second equation follows from the monotonicity of F_i. The requirement 1 in Section 3.3 then translates into

$$H_i(t) - H_i(t') + s\left(\beta_i(t') - \beta_i(t)\right) \leq 0, \text{ for all } s \text{ and } t' > t. \tag{18}$$

Unless β_i is constant, this restriction, however, is always violated for models where the domain of the systematic factor variable is the real line (as is typically the case), as s can then take arbitrarily high or low values. Even for models where s may be bounded, it will often be hard to prevent violation of (18).

We note in passing that a Gaussian copula with horizon-dependent correlation parameter is of the type (17), and will not lead to a consistent arbitrage-free model. Despite this fact, the resulting "model" appears to be in practical use as a way to equip the so-called *base correlation* framework (see [15]) with a mechanism to manipulate the term-structure of loss distributions. The resulting horizon-dependent base correlation curves obviously have limited physical meaning[10].

[10] An alternative approach to work with time sequences of base correlation curves could be based on

4.1 Forward Construction of Factor Model

To circumvent the inconsistencies associated with (17), we can contemplate using the default driver $\beta_i(t)Y + e_i$ in a strictly "forward" manner. Working (for simplicity) on some grid $\{T_j\}_{j=1}^M$, let us interpret $\beta_i(T_j)$ as the future loading on the systematic variable Y on the interval $[T_{j-1}, T_j]$,and let $X^{(j)} = \beta_i(T_j)Y + e_i$ be the driver process to be applied on this same interval. Note that Y and the various e_i's do *not* vary with the time interval (we relax this later, in Section 5). Following standard procedures, we first consider $0 \le t \le T_1$ and write, for some barrier function H_i,

$$\begin{aligned}\Pr(\tau_i > t|Y = y) &= \Pr(\beta_i(T_1)Y + e_i > H_i(t)|Y = y)\\ &= 1 - F_i(H_i(T_1) - \beta_i(T_1)y), \quad t \in [0, T_1].\end{aligned}$$

and solve for $H_i(T_1)$ to match a given probability $\Pr(\tau_i > T_1)$. Moving on to the period $(T_1, T_2]$, we let ourselves be inspired by a proper first-passage time approach, and write

$$\begin{aligned}\Pr(T_1 < \tau_i \le t) &= \Pr(\tau_i > T_1) - \Pr(\tau_i > t) = \Pr\left(X^{(1)} > H_i(T_1), X^{(2)} \le H_i(t)\right)\\ &= \Pr(\beta_i(T_1)Y + e_i > H_i(T_1), \beta_i(T_2)Y + e_i \le H_i(t)), \quad t > T_1.\end{aligned}$$

Conditional on Y, we thus get

$$\begin{aligned}\Pr(\tau_i > T_1|Y = y) - \Pr(\tau_i > t|Y = y) &= \Pr(\beta_i(T_1)y + e_i > H_i(T_1), \beta_i(T_2)y + e_i \le H_i(t))\\ &= \Pr(H_i(T_1) - \beta_i(T_1)y < e_i \le H_i(t) - \beta_i(T_2)y)\\ &= (F_i(H_i(T_2) - \beta_i(T_2)y) - F_i(H_i(t) - \beta_i(T_1)y))^+,\end{aligned}$$

where F_i is the distribution function of firm i residuals e_i, and where we use the notation $x^+ = \max(x, 0)$. If we as before write $\Pr(\tau_i > t|Y = y) = q_i(t, y)$, we thus have

$$q_i(t, y) = q_i(T_1, y) - (F_i(H_i(t) - \beta_i(T_2)y) + F_i(H_i(T_1) - \beta_i(T_1)y))^+, \quad t \in (T_1, T_2].$$

In more generality, for the interval $t \in (T_{k-1}, T_k]$ we can write

$$\begin{aligned}&\Pr(T_{k-1} < \tau_i \le t) =\\ &\quad \Pr(\beta_i(T_1)Y + e_i > H_i(T_1), \ldots, \beta_i(T_{k-1})Y + e_i > H_i(T_{k-1}), \beta_i(T_k)Y + e_i \le H_i(t)),\end{aligned}$$

such that

$$\begin{aligned}&\Pr(\tau_i > T_{k-1}|Y = y) - \Pr(\tau_i > t|Y = y)\\ &\qquad = \Pr\left(e_i > \max_{j=1,\ldots,k-1}(H_i(T_j) - \beta_i(T_j)y), e_i \le H_i(t) - \beta_i(T_k)y\right)\end{aligned}$$

---

the "forward" correlation technique in Section 4.1, applied to the Gaussian copula. As is everything that involves base correlations, the resulting model would still have a strong ad-hoc flavor or, for $t \in (T_{k-1}, T_k]$

$$q_i(t,y) = q_i(T_{k-1},y) - \left(F_i\left(\max_{j=1,\ldots,k-1} \left(H_i(T_j) - \beta_i(T_j)y\right)\right) - F_i\left(H_i(t) - \beta_i(T_k)y\right)\right)^+ . \quad (19)$$

The barrier function $H_i(\cdot)$ can be found iteratively, in "bootstrapping" fashion. Specifically, assuming that all $H_i(T_j),\ j = 0,1,\ldots,k-1$ have been determined, $H_i(t)$ can be found from the usual condition

$$\Pr(\tau_i > T_k) = \int_{\mathbb{R}} q_i(t,y)\,\phi_Y(y)dy,$$

where ϕ_Y is the density of Y. The right-hand side of this equation depends on $H_i(t)$ and all $H_i(T_j),\ j = 0,1,\ldots,k-1$, and will generally require numerical integration.

We notice that $q_i(t,y)$ defined as in (19) will, by construction[11], satisfy the constraints (1)-(2). We also notice that for the special case where β_i is independent of time, (19) reduces to a standard factor model with specification

$$q_i(t,y) = \Pr(\tau_i > t | Y = y) = 1 - F_i\left(H_i(t) - \beta_i y\right).$$

This follows immediately from the fact that $H_i(T)$ would necessarily be increasing in T, such that, for $t \in (T_{k-1}, T_k]$,

$$\max_{j=1,\ldots,k-1} \left(H_i(T_j) - \beta_i y\right) = H_i(T_{k-1}) - \beta_i y < H_i(t) - \beta_i y.$$

4.2 Forward Construction with Independent Residuals

There are alternative ways of "chaining" together one-period models across time. To show one alternative to the method above, consider writing for the default driver $X^{(j)} = \beta_i(T_j)Y + e_i^{(j)}$, where all $e_i^{(j)}$'s have identical cumulative distribution F_i, but $e_i^{(k)}$ and $e_i^{(j)}$ are *independent* for $k \neq j$. Again starting with the idea that, for $t \in (T_{k-1}, T_k]$,

$$\begin{aligned}&\Pr(T_{k-1} < \tau_i \leq t)\\ &= \Pr\left(\beta_i(T_1)Y + e_i^{(1)} > H_i(T_1), \ldots, \beta_i(T_{k-1})Y + e_i^{(k-1)} > H_i(T_{k-1}), \beta_i(T_k)Y + e_i^{(k)} \leq H_i(t)\right),\end{aligned}$$

it is easy to show that, for $t \in (T_{k-1}, T_k]$,

$$q_i(t,y) = q_i(T_{k-1},y) - \left(\Pi_{j=1}^{k-1}\left(1 - F_i\left(H_i(T_j) - \beta_i(T_j)y\right)\right)\right) F_i\left(H_i(t) - \beta_i(T_k)y\right) \quad (20)$$

Note, however, that this model does *not* reduce to a simple factor model $q_i(t,y) = 1 - F_i(H_i(t) - \beta_i y)$ when β_i is made independent of time.

[11] To complete the model for t beyond the final grid date T_M, we can simply assume that $q_i(t,y)$ for $t > T_M$ is identical to $q_i(t,y)$, $t \in (T_{M-1}, T_M]$.

4.3 Example: RFL Model

The technique behind (19) and (20) generalizes easily to more complicated setups, including the case where the loading on the systematic factor is state-dependent. For instance, in a standard RFL model with firm-independent loading function a, we write

$$X_i = a(Y)Y + e_i, \quad i = 1, \ldots, N,$$

where the idiosyncratic variables e_i are Gaussian with standard deviation v and mean m. We could introduce time-dependence in the systematic term several ways, say by using

$$X_i = a\left(Y \cdot w(T_k)\right) Y + e_i$$

on the interval $[T_{k-1}, T_k]$. If $w(T) = \sqrt{T}$, say, we could mimic the behavior of a diffusion model; see [4] for some numerical experiments suggesting that this specification is reasonably consistent with market observations. Alternatively, we could use $w(T)$ non-parametrically to calibrate against known term structures of portfolio loss distributions. Applying the idea of Section 4.1, arguments identical to those leading to (19), we get, for $t \in (T_{k-1}, T_k]$,

$$q_i\left(t, y\right) = q_i\left(T_{k-1}, y\right) - \left(\Phi\left(\frac{\max_{j=1}^{k-1}\left(H_i(T_j) - a\left(y \cdot w(T_j)\right) y\right) - m}{v}\right) - \Phi\left(\frac{H_i(t) - a\left(y \cdot w(T_k)\right) y - m}{v}\right)\right)^+.$$

We can easily extend this to the jump-enhanced RFL model in Section 3.6. The techniques of Section 4.2 also carry over in straightforward manner to the RFL model; we omit the obvious details.

4.4 Discussion and Tests

With the chaining techniques outlined above, we can improve – at the expense of losing some tractability – the ability of the model to match marginal loss distributions at multiple horizons. An interesting question is to which extent this will make a difference in the pricing and hedging of standard securities such as CDO tranches. As we have argued earlier, regular CDO tranches depend primarily on the marginal loss distribution at terminal maturity, but there will be cases where relatively significant effects can be seen. To understand this, we cannot avoid at this point a brief excursion into the mechanics of CDO tranche payouts (for more details, consult any credit derivatives textbook).

We recall that the protection seller on a CDO tranche is paid (by the protection buyer) either a lump sum upfront, a periodic fee for the life of the trade, or a combination of the two. The most prevalent payment form is the periodic fee structure, except in cases of high-risk tranches – typically equity tranches – where the fee would be impractically large. In the most common type of trade, the periodic fee payment is equal to a fixed spread s multiplied by the current notional of the tranche; this notional is subject to downward amortizations as losses take place in the underlying portfolio and eventually eat into the tranche in question. Ignoring minor discounting effects, it is easy to see that the upfront value of buying protection on a CDO contract only depends on the distribution of cumulative losses over the life of the trade – the exact timing of these losses is irrelevant. The value of the periodic fee payments, however, is different: all things equal, if tranche losses tend to take place early rather than late, the protection seller will receive fewer fee payments (as the notional will amortize down quicker) and the value of the fee payments will decrease. Another way of saying the same thing: the more front-loaded tranche losses are, the lower the duration of the fee payments.

For CDO tranches where the protection buyer pays a periodic fee, the market quote will be a fair spread, computed as the ratio between the upfront value of protection and the value of the *risky annuity factor*, the latter being defined as the present value of the periodic fee assuming a unity spread $s = 100\%$. It is worth emphasizing that the market quote reveals neither the upfront value of tranche protection nor the risky annuity in separation, only the ratio is known. Suppose now that two models – models I and II – are calibrated to the market tranche spread of a given CDO, but model I tends to produce tranche losses earlier than those of model II (so, model I frontloads losses and model II backloads them). All things equal, for reasons discussed above model I will return a lower risky annuity factor than will model II, and consequently will end up assigning a lower upfront value to the CDO tranche than will model I – despite the fact that both models agree on the market quoted spread. In a valuation system, we would see model I would yield a lower price than model II.

The size of CDO pricing discrepancies driven by the fee-effect discussed above is often quite small ($< 0.5\%$ of the upfront value), as risky annuity factors for all medium- and low-risk tranches are generally quite close to risk-free annuity factors. As the riskiest tranches typically trade predominatly up-front (limiting the fee-effect), it is a comparatively small set of actual CDOs that in practice are markedly affected by loss term structure effects. To get a feel for rough magnitudes, we consider a small example. Specifically, we use the “chained” RFL model in Section 4.3 with the function $w(T)$ piecewise flat at levels set to best-fit 3-year, 5-year, 7-year, and 10-year tranche prices in the CDX-NA market (as of mid-July 2006). The 10-year prices produced by the calibrated model are then used as calibration targets for a standard (non-chained) RFL model. Table 1 below illustrates the pricing differences produced by the two models, for various 10-year tranches. (Results for shorter-dated tranches are omitted, as effects are less pronounced).

The results in the table reflect the fact that the standard RFL model has a known propensity to front-load scenarios with many losses, which in turn lowers the duration of senior tranche fee streams. By the argument given earlier, this lowers the upfront value of senior tranches relative to the market. The standard RFL model also (by a parity argument) will tend to produce relatively many scenarios with negligible early losses, which will increase the duration of the equity tranche fee streams and, thereby, increase the upfront value. Overall, however, the effects in Table 1 are relatively subdued. We note in passing that had the 0%-3% equity tranche traded entirely in running-fee form[12] with

[12] The break-even spread would in this case be around 17%.

Tranche	Difference in Upfront Value
0%-3%	2.34%
3%-7%	0.92%
7%-10%	-0.72%
10%-15%	-0.80%
15%-30%	-0.85%

Table 1: Relative difference (non-chained minus chained) in 10-year CDX tranche upfront values produced by non-chained vs. chained RFL models. The model differences in par spreads were smaller than 0.2 basis points for all tranches. Note that the 0%-3% tranche trades with a 5% fixed running spread, with the remaining value paid upfront.

no upfront, the 2.34% relative upfront pricing discrepancy would increase to 7% for this tranche.

In situations where there is already considerable valuation uncertainty (such as bespoke CDO tranches), the value of attempting to match a term structure of loss distributions is likely negligible and may not be worth the effort, at least for standard trades quoted with moderate running fees. That said, setting up a model to match term structures may in some cases have side benefits, including a stabilization of model parameters across deals with different maturities and an improved ability to perform inter- and extrapolation in maturity space. Hedging performance may be affected too, of course, but we remind the reader that hedging in factor models is in itself a rather delicate topic (see [3]), one that we opt not to pursue further here.

5 Intertemporal Correlation. Forward-Starting CDOs.

While certainly useful, the ability to control and fit term structures of portfolio loss distributions is, as discussed, often of somewhat secondary importance in pricing of standard portfolio derivatives. This, however, is generally not true for non-standard CDO tranches with a strong timing component. The prime example is probably the *forward starting CDO* (FCDO) tranche, a CDO tranche characterized by having a start date T^* set in the future ($T^* > 0$). As a consequence, in the FCDO structure portfolio losses that take place in the pre-start time interval $[0, T^*]$ do *not* count in determining the cumulative losses in a given tranche on the underlying portfolio. With the terminal maturity denoted T, it is clear that proper pricing of the FCDO requires, as a minimum, that the underlying model gives the user control over marginal loss distributions (at least) to the two separate horizons $[0, T^*]$ and $[0, T]$. Controlling marginal loss distributions on $[0, T^*]$ and $[0, T]$ can be accomplished by the techniques discussed in Section 4, and would ensure that the expected portfolio losses on $[T^*, T]$ are as desired. Merely pinning down this expectation in the model is, however, inadequate for pricing of the FCDO, which depends on the entire loss distribution on $[T^*, T]$. This loss-distribution, in turn, depends strongly on the *intertemporal* co-dependence between cumulative portfolio losses on $[0, T^*]$ and $[0, T]$.

We shall see examples of this effect shortly; for now, we just notice that the models outlined in Section 4 have essentially *no* control over intertemporal correlation, as the factor variable used in an arbitrary each time slot $[T_{k-1},T_k]$ is perfectly correlated to (in fact, identical to) the one used in any other time slot $[T_{l-1},T_l]$, $T_l \neq T_k$.

5.1 Two-period Setting

Let us start out with a factor model of the type (7), and contemplate how to extend the approach in Section 4 to control intertemporal correlation between default drivers. For simplicity (and with an eye on the FCDO contract), we initially work with a time line $\{T_j\}_{j=1}^M$ with only two dates, $M=2$. We now start out by defining, for $i=1,\ldots,N$,

$$\begin{aligned} X_i^{(1)} &= \beta_i(T_1)Y^{(1)} + e_i^{(1)}, \\ X_i^{(2)} &= \beta_i(T_2)Y^{(2)} + e_i^{(2)} \end{aligned}$$

where the density of $Y^{(j)}$ is $\phi_Y^{(j)}(y)$, $j=1,2$, and the joint density of $Y^{(1)}$ and $Y^{(2)}$ is $\phi_Y^{(1,2)}(y_1,y_2)$. Also let $F_i^{(j)}(e)$ denote the cumulative distribution function for $e_i^{(j)}$, $j=1,2$, and let $F_i^{(1,2)}(e_1,e_2)$ be the joint cumulative distribution function for $e_i^{(1)}$ and $e_i^{(2)}$. Often, we would have $\phi_Y^{(1)}=\phi_Y^{(2)}$ and $F_i^{(1)}=F_i^{(2)}$, but this is not required. Mimicking the construction used in Section 4.1, we introduce a barrier[13] H_i, and write

$$\begin{aligned} \Pr(\tau_i > t) &= \Pr\left(\beta_i(T_1)Y^{(1)} + e_i^{(1)} > H_i(t)\right), \quad t\in[0,T_1], \\ \Pr(T_1 < \tau_i \le t) &= \Pr\left(\beta_i(T_1)Y^{(1)} + e_i^{(1)} > H_i(T_1), \beta_i(T_2)Y^{(2)} + e_i^{(2)} \le H_i(t)\right), \quad t\in(T_1,T_2]. \end{aligned}$$

Setting our total systematic factor vector to $S=\left(Y^{(1)},Y^{(2)}\right)$, we therefore get the following result for the conditional survival function $q(t,s)=q(t,s_1,s_2)$:

$$q_i(t,s_1,s_2) = 1 - F_i^{(1)}\left(H_i(t) - \beta_i(T_1)s_1\right), \quad t\in[0,T_1], \tag{21}$$

$$\begin{aligned} q_i(t,s_1,s_2) &= q(T_1,s_1,s_2) - F_i^{(2)}\left(H_i(t) - \beta_i(T_2)s_2\right) \\ &\quad + F_i^{(1,2)}\left(H_i(T_1) - \beta_i(T_1)s_1, H_i(t) - \beta_i(T_2)s_2\right), \quad t\in(T_1,T_2]. \end{aligned} \tag{22}$$

In the last equation, we have used the fact that for arbitrary x,y

$$\Pr\left(e_i^{(1)} > x, e_i^{(2)} \le y\right) + \Pr\left(e_i^{(1)} \le x, e_i^{(2)} \le y\right) = \Pr\left(e_i^{(2)} \le y\right).$$

Calibration of $H_i(t)$ is accomplished, as usual, from the equation

$$\Pr(\tau_i > t) = \int_{\mathbb{R}^2} q(t,s_1,s_2)\phi_Y^{(1,2)}(s_1,s_2)ds_1 ds_2. \tag{23}$$

[13] While we for notational convenience use a single barrier-function H, we note that this function is not continuous across time T_1 . In fact, just as $H_i(0)=-\infty$, we would have $H_i(T_2+)=-\infty$, reflecting the fact that the model "refreshes" at time T_2.

For some models (see e.g. Section 5.2 below), the right-hand side may be available in closed form.

In the setup (21)-(22), control of intertemporal co-dependence of losses on $[0,T_1]$ and $[T_1,T_2]$ is accomplished through specification of the joint distribution function $F_i^{(1,2)}(e_1,e_2)$ and the joint density $\phi_Y^{(1,2)}(y_1,y_2)$. We note that the model considered in Section 4.1 corresponds to the perfect dependence case where $Y^{(1)}=Y^{(2)}$ and $e_i^{(1)}=e_i^{(2)}$. The approach in Section 4.2 also uses $Y^{(1)}=Y^{(2)}$, but sets $F_i^{(1,2)}(e_1,e_2)=F_i(e_1)F_i(e_2)$. As mentioned earlier, the intertemporal co-dependence structure of the models in Sections 4.1 and 4.2 are both highly suspect, obviously, as is also the case for any model of the type (7).

In general, determination of the "best" specification of $F_i^{(1,2)}$ and $\phi_Y^{(1,2)}(y_1,y_2)$ can be difficult, but sometimes we can let ourselves be guided by results from a full-blown first-passage time model (see (5)). We consider this idea in Examples 5.2 and 5.4 below. Without going into details, we also note that it sometimes might be convenient to use a copula approach to "stitch" the known marginals $\phi_Y^{(1)},\phi_Y^{(2)},F_i^{(1)}$, and $F_i^{(2)}$ together into joint distributions.

5.2 Example: Gaussian copula

To give a concrete example of (21)-(22), consider a Gaussian copula setup where we use a correlation of ρ_1 in the first period $[0,T_1]$ and a correlation of ρ_2 in the second period. That is, we write

$$X_i^{(1)} = \sqrt{\rho_1}Y^{(1)}+e_i^{(1)}, \tag{24}$$
$$X_i^{(2)} = \sqrt{\rho_2}Y^{(2)}+e_i^{(2)}, \tag{25}$$

where $e_i^{(1)}$ and $e_i^{(2)}$ are Gaussian with mean 0 and variance $1-\rho_1$ and $1-\rho_2$, respectively; and $Y^{(1)}$ and $Y^{(2)}$ are Gaussian with mean 0 and variance 1. We assume that $Y^{(1)}$ and $Y^{(2)}$ are correlated with correlation coefficient ρ_Y; and that $e_i^{(1)}$ and $e_i^{(2)}$ are correlated with correlation ρ_e. It follows from (21)-(22) that

$$q(t,s_1,s_2)=1-\Phi\left(\frac{H_i(t)-\sqrt{\rho_1}s_1}{\sqrt{1-\rho_1}}\right), \quad t\in[0,T_1], \tag{26}$$

$$q(t,s_1,s_2)=1-\Phi\left(\frac{H_i(T_1)-\sqrt{\rho_1}s_1}{\sqrt{1-\rho_1}}\right)-\Phi\left(\frac{H_i(t)-\sqrt{\rho_2}s_2}{\sqrt{1-\rho_2}}\right)$$
$$+\Phi_2\left(\frac{H_i(T_1)-\sqrt{\rho_1}s_1}{\sqrt{1-\rho_1}},\frac{H_i(t)-\sqrt{\rho_2}s_2}{\sqrt{1-\rho_2}};\rho_e\right), \quad t\in(T_1,T_2]. \tag{27}$$

We note that if $\rho_Y=\rho_e=1$ and $\rho_1=\rho_2$, the model reduces to a standard Gaussian copula.

Calibration of the barrier level $H_i(t)$ can, for the model above, be done quickly, as (unconditional) survival probabilities are available in closed form. To see this, we notice that both $X_i^{(1)}$ and $X_i^{(2)}$ are standard Gaussian variables, with correlation $\rho_X=\rho_Y\sqrt{\rho_1\rho_2}+$

$\rho_e\sqrt{(1-\rho_1)(1-\rho_2)}$. Thereby

$$\begin{aligned}\Pr(\tau_i > t) &= \Pr\left(X_i^{(1)} > H_i(t)\right) = \Phi(-H_i(t)), \quad t \in [0, T_1], \\ \Pr(\tau_i > t) &= \Pr\left(X_i^{(1)} > H_i(T_1)\right) - \Pr\left(X_i^{(1)} > H_i(T_1), X_i^{(2)} \le H_i(t)\right) \\ &= \Phi(-H_i(T_1)) - \Phi(H_i(t)) + \Phi_2(H_i(T_1), H_i(t); \rho_X), \quad t \in (T_1, T_2].\end{aligned}$$

In the setup above, we have complete freedom in selecting the various correlations ρ_1, ρ_2, ρ_e, and ρ_Y, giving us significant flexibility in controlling the intertemporal correlation of portfolio losses. If we wish to base our correlation estimates on a dynamic model, we can consider a full-blown first-passage time model [14] of the type

$$\tau_i = \inf\left\{t : \sqrt{r}W_Y(t) + \sqrt{1-r}W_i(t) = h_i(t)\right\}, \quad i = 1, \ldots, N, \tag{28}$$

where $0 \le r \le 1$ is a correlation parameter, W_Y is a systematic Brownian motion affecting all firms, $W_i(t)$ is an idiosyncratic Brownian motion unique to firm i (and independent of W_Y and W_j, $j \ne i$), and the barrier level $h_i(t)$ is calibrated to match survival probabilities for firm i. It is well-known (see for instance [16]) that this model produces CDO tranche prices that are very similar to a standard Gaussian copula with correlation r, i.e. a model such as the one above but in the special setting $\rho_Y = \rho_e = 1$ and $\rho_1 = \rho_2 = r$. This is not surprising: while the two models have different dynamics and intertemporal correlation structures, standard CDOs depend only on marginal loss distributions which here intuitively are quite close as long as both models are calibrated to the individual firm default probability curves. For a *forward-starting* CDO, however, we would need to carefully contemplate what setting of ρ_e and ρ_Y may be most suitable to make the simple model (24)-(25) behave as much as possible like (28) for FCDO pricing purposes. An educated guess – the validity of which we have confirmed in Monte Carlo studies[15] – is to set

$$\rho_e = \rho_Y = \operatorname{corr}(W_Y(T_2), W_Y(T_1)) = \operatorname{corr}(W_e(T_2), W_e(T_1)) = \sqrt{T_1/T_2}. \tag{29}$$

5.3 Generic pricing of an FCDO in a factor model

In preparation for numerical tests in Section 5.4, we provide a brief interlude about the pricing of FCDO tranches. Let us assume that we are in a complete factor setting, as in Section 3.3. Assuming a \$1 loss-given default for each firm in an N-dimensional basket (extensions to non-homogeneous loss-given-defaults are straightforward), the "effective" cumulative portfolio loss hitting an FCDO tranche is given by

$$L_{eff}(t) = \sum_{i=1}^{N} 1_{T^* \le \tau_i \le t}, \quad t \in (T^*, T].$$

[14] See e.g. [10] and [16]. We should clarify that we are not endorsing (28) for general use as a dynamic model (the model is too non-stationary in its credit spread dynamics for such a recommendation), but merely use (28) as an example. First-passage time models with Poisson jumps (e.g. *a la* the model in Section 3.6) in the default driver are generally a better choice in this regard, but also computationally more demanding.

[15] We omit detailed test data here as the result is intuitive and its validity peripheral to our main discussion. Numerical test data is available from the author upon request.

Iterative construction of this effective loss can proceed by standard algorithms for loss construction on the horizon $[0,t]$, by simply replacing conditional default probabilities $q_i(t,s)$ with $1-\Pr(T^* \le \tau_i \le t | S=s) = 1+q_i(t,S)-q_i(T^*,s)$, for $t \ge T^*$. (For $t < T^*$, $L_{eff}(t)=0$ always, by definition). In turn, pricing of an FCDO tranche can reuse existing factor-model algorithms, provided that we for all $t \ge 0$ make the substitution $q_i(t,s) \mapsto \min(1, 1+q_i(t,s)-q_i(T^*,s))$.

5.4 Example: FCDO Pricing with Gaussian Copula

We consider FCDO pricing on a CDX-like portfolio with 125 names, with average 5-year spread of 75bps, and typically upward-sloping credit spread term structures. As described earlier, the start date of the FCDO is T^*, and the deal maturity is T. In our examples, we let $T-T^*=5$ always, and consider two trades: 1) $T^*=2$; and 2) $T^*=5$. We work with the model in Example 5.2 above (with $T_1=T^*$ and $T_2=T$) and consider two correlation scenarios: $\rho_1=\rho_2=20\%$ and $\rho_1=\rho_2=80\%$. The tables below illustrates the effects on tranche break-even spreads from changing the "intertemporal correlations" ρ_e and ρ_Y. (In the tables we have $\rho_e=\rho_Y=r$ for some varying constant r; using $\rho_e=\rho_Y$ involves the assumption that the systematic factor and the residuals have identical auto-correlation, as is the case, say, in the Brownian motion interpretation (29)). Pricing is done according to the technique outlined in Section 5.3 above.

		$T^*=2$				$T^*=5$			
Tranche	Spot Trade	$r=1$	$r=\sqrt{T^*/T}$	$r=0$	$r=-1$	$r=1$	$r=\sqrt{T^*/T}$	$r=0$	$r=-1$
0%-3%	20.564%	33.876%	31.949%	31.151%	30.818%	41.142%	34.838%	31.759%	30.897%
3%-7%	5.288%	9.315%	9.080%	8.983%	8.990%	10.643%	9.835%	9.425%	9.333%
7%-10%	1.644%	3.242%	3.341%	3.409%	3.447%	3.033%	3.526%	3.743%	3.824%
10%-15%	0.504%	1.046%	1.198%	1.272%	1.310%	0.547%	1.169%	1.454%	1.552%
15%-30%	0.047%	0.078%	0.134%	0.158%	0.169%	0.006%	0.103%	0.196%	0.228%

Table 2: Break-even spreads for various tranches, in the scenario $\rho_1=\rho_2=20\%$. In all cases $T-T^*$=5 years. The "Spot Trade" column lists break-even spreads for a regular 5-year CDO ($T^*=0$) in the standard Gaussian copula. Numbers are based on 100,000 MC simulations.

Let us lend some intuition to the results above. In a Gaussian model where intertemporal correlation is high, crash-scenarios with multiple defaults tend to contain many of these defaults inside the "safe" period $[0,T^*]$, with the net result that senior FCDO tranches rarely suffers a loss. As the intertemporal correlation is lowered, however, there is an increased likelihood of a crash isolated to the period $[T^*,T]$ (corresponding to a bad outcome of the systematic driver $Y^{(2)}$) without a bad outcome in the period $[0,T^*]$ (corresponding to a good outcome of the systematic driver $Y^{(1)}$). As a result, there may be a sufficient amount of defaults located inside $[T^*,T]$ to cause losses in the senior tranche. Consistent with this explanation, break-even spread of senior tranches in the tables above increase markedly when r is lowered.

		$T^*=2$				$T^*=5$			
Tranche	Spot Trade	$r=1$	$r=\sqrt{T^*/T}$	$r=0$	$r=-1$	$r=1$	$r=\sqrt{T^*/T}$	$r=0$	$r=-1$
0%-3%	9.579%	14.464%	12.647%	11.655%	11.136%	18.591%	12.847%	9.878%	8.709%
3%-7%	3.490%	5.384%	5.012%	4.931%	4.943%	7.024%	5.583%	4.978%	4.827%
7%-10%	2.211%	3.491%	3.313%	3.307%	3.353%	4.649%	3.756%	3.530%	3.536%
10%-15%	1.578%	2.525%	2.430%	2.439%	2.467%	3.232%	2.721%	2.674%	2.719%
15%-30%	0.851%	1.366%	1.357%	1.386%	1.407%	0.893%	1.428%	1.570%	1.634%

Table 3: Break-even spreads for various tranches, in the scenario $\rho_1=\rho_2=80\%$. In all cases $T-T^*$=5 years. The "Spot Trade" column lists break-even spreads for a regular 5-year CDO ($T^*=0$) in the standard Gaussian copula. Numbers are based on 100,000 MC simulations.

For the equity investor, as usual, things work in exactly the opposite way than for the senior investor. As explained above, a low intertemporal correlation will tend to produce more "clustered" loss paths in the period $[T^*,T]$ which, all things equal, will make it more likely that there will be outcomes where very few defaults take place in $[T^*,T]$. In other words, lowering intertemporal correlation will make the equity tranche less risky. This is reflected in the tables above, where the equity break-even spreads decrease markedly when r is lowered.

For mezzanine tranches in the middle of the capital structure, the effects of changes to the intertemporal correlation r is somewhere between that of the equity and senior tranches, with some tranches having quite limited sensitivity to r.

5.5 Example: RFL Model

Retaining most notation from Example 5.2, we consider a two-period RFL model of the form

$$\begin{aligned} X_i^{(1)} &= a^{(1)}(Y^{(1)})Y^{(1)}+e_i^{(1)}, \\ X_i^{(2)} &= a^{(2)}(Y^{(2)})Y^{(2)}+e_i^{(2)}, \end{aligned}$$

where $a^{(1)}$ and $a^{(2)}$ are – possibly different – factor loading functions shared among all firms. (Again, extensions to firm-specific a are trivial). The residuals $e_i^{(1)}$ and $e_i^{(2)}$ are correlated Gaussian variables with means $m^{(1)}$ and $m^{(2)}$, and standard deviations $v^{(1)}$ and $v^{(2)}$. The conditional survival function is easily seen to become (ρ_e is the correlation between $e_i^{(1)}$ and $e_i^{(2)}$)

$$q(t,s_1,s_2) = 1 - \Phi\left(\frac{H_i(t) - a^{(1)}(s_1) - m^{(1)}}{v^{(1)}}\right), \quad t \in [0,T_1],$$

$$q(t,s_1,s_2) = 1 - \Phi\left(\frac{H_i(T_1) - a^{(1)}(s_1) - m^{(1)}}{v^{(1)}}\right) - \Phi\left(\frac{H_i(t) - a^{(2)}(s_2) - m^{(2)}}{v^{(2)}}\right)$$
$$+ \Phi_2\left(\frac{H_i(T_1) - a^{(1)}(s_1) - m^{(1)}}{v^{(1)}}, \frac{H_i(t) - a^{(2)}(s_2) - m^{(2)}}{v^{(2)}}; \rho_e\right), \quad t \in (T_1,T_2].$$

Under the assumption of piecewise flat a_i, for $t \le T_1$ the unconditional survival probability $\Pr(\tau_i > t)$ can, as we know from [2], be computed in closed form as a function of $H_i(t)$. For $t > T_1$ the calibration expression (23) (with $\phi_Y^{(1,2)}$ here being a bi-variate Gaussian density) will involve higher-dimensional Gaussian integrals and is probably best done by a numerical quadrature scheme.

5.6 Generalization to multiple periods

The two-period approach in Section 5.1 is primary of interest for securities with a single "important" future date (such as FCDOs), but remains rather crude as a true model of dynamic evolution of losses. Extensions to multiple time-periods will improve the realism of the setup, but it is clear, however, that computational effort will start to increase rapidly in the number of time periods, unless we (as in Section 4) limit ourselves to special cases of intertemporal co-dependence. This is not surprising, as increasing the number of time-steps will, with proper parameterization, take us closer and closer to a true dynamic model of the first-passage time type, the numerical difficulties of which are well-known.

6 Conclusion

While an abundant number of factor-type models have been proposed, reality is that these models generally are quite similar in their structure and in their limitations. A particular weakness of the standard set of models is their inability to generate a meaningful description of portfolio loss distributions over time, both statically and dynamically. In this paper, we have analyzed some of the rather thorny issues surrounding factor models and temporal loss evolution. In certain cases – those where only a fit to marginal loss distributions at multiple times is required – a "chaining" technique can assist in extending existing models to better match observed term structure market data. This technique, however, is inadequate in pricing of more exotic tranche securities with a pronounced timing element, in which case explicit control over the dynamic evolution of portfolio losses is required. Small steps in this direction are possible through introduction of an intertemporal correlation element into the factor framework, but the resulting framework is generally rather limited in scope and quickly becomes numerically intractable. The importance of the issue has been demonstrated through numerical simulations: even straightforward CDO variations such as forward starting CDOs can have strong sensitivity to intertemporal loss correlation.

It is clear that much work remains. As part of our analysis, several new models and techniques were suggested; pursuing detailed pricing and hedging performance of concrete implementations would be an interesting undertaking. Given the practical limits of factor models – which we have spent considerable time on both in this paper and in [3] – a broader task, however, is the search for practical extensions that would solidify further our handle on the pricing and hedging of securities with exposure to loss dynamics. Complicating matters considerable here is the fact that some securities (such as leveraged super-senior tranches, see [18]) depend on joint dynamics of credit spreads and portfolio losses, adding another layer of complexity. While abstract models such as [18] and [17] offer a potential road map for model development, there will likely be room for pragmatic compromises.

7 Appendix A

We consider the special case where the loading function $a(z)$ is a step-function

$$a(z) = \begin{cases} \overline{a}, & z > \theta, \\ \underline{a} & z \leq \theta. \end{cases}$$

Extensions to multiple steps in the a function (which is often useful in application) is a straightforward extension that we omit in the interest of brevity. Let us define a function

$$\begin{aligned} w_i(k,l,T) &= \Pr\left(N(T)=k, N_i(T)=l\right) = \Pr\left(N(T)=k\right)\Pr\left(N_i(T)=l\right) \\ &= e^{-(\lambda+\lambda_i)T}\frac{\lambda^k \lambda_i^l T^{k+l}}{l!k!}. \end{aligned}$$

Conditional on $N(T)=k$ and $N_i(T)=l$, we have

$$X_i(T) = a\left(\Omega\right)\Omega + \Psi,$$

where $\Omega \sim N(k\mu, \sigma^2 T + k\gamma^2)$ and $\Psi \sim N(l\mu_i, \sigma_i^2 T + l\gamma_i^2)$. Define $m(k) = k\mu$, $m_i(l) = l\mu_i$, $v(k)^2 = \sigma^2 T + k\gamma^2$, $v_i(l)^2 = \sigma_i^2 T + l\gamma_i^2$, and $\theta^*(k) = [\theta - m(k)]/v(k)$. Conditional on

$N(T) = k$ and $N_i(T) = l$, we can write (dropping arguments k and l for brevity)

$$\begin{aligned}
\Pr\left(a(\Omega)\Omega + \Psi \le x\right) &= \mathrm{E}_\Omega\left[\Pr\left(a(\Omega)\Omega + \Psi \le x \,|\, \Omega\right)\right] \\
&= \mathrm{E}_\Omega\left[\Pr\left(\Psi \le x - \overline{a}\Omega 1_{\Omega>\theta} - \underline{a}\Omega 1_{\Omega\le\theta} \,|\, \Omega\right)\right] \\
&= \mathrm{E}_\Omega\left[\Phi\left(\frac{x - m_i - \overline{a}\Omega 1_{\Omega>\theta} - \underline{a}\Omega 1_{\Omega\le\theta}}{v_i}\right)\right] \\
&= \int_{-\infty}^{\infty} \Phi\left(\frac{x - m_i - \overline{a}(\omega v + m) 1_{\omega>\theta^*} - \underline{a}(\omega v + m) 1_{\omega\le\theta^*}}{v_i}\right)\phi(\omega)\,d\omega \\
&= \int_{-\infty}^{\theta^*} \Phi\left(\frac{x - m_i - \underline{a}(\omega v + m)}{v_i}\right)\phi(\omega)\,d\omega \\
&\quad + \int_{\theta^*}^{\infty} \Phi\left(\frac{x - m_i - \overline{a}(\omega v + m)}{v_i}\right)\phi(\omega)\,d\omega
\end{aligned}$$

where Φ is the cumulative Gaussian distribution function, and $\phi(x) = (2\pi)^{-1/2}\exp(-x^2/2)$ is the standard Gaussian density. A result in [2] establishes that for arbitrary constants C, D, ζ

$$\int_{-\infty}^{\zeta} \Phi(Cy + D)\,\phi(y)\,dy = \Phi_2\left(\frac{D}{\sqrt{1+C^2}}, \zeta; \frac{-C}{\sqrt{1+C^2}}\right)$$

where $\Phi_2(\cdot,\cdot;\rho)$ is the bi-variate cumulative Gaussian distribution function with correlation ρ. Thereby, and still conditional on $N(T) = k$ and $N_i(T) = l$, we have

$$\begin{aligned}
\Pr\left(a(\Omega)\Omega + \Psi \le x\right) &= \Phi_2\left(\frac{D_i(k,l,x,\underline{a})}{\sqrt{1+\underline{a}^2 C_i(k,l)^2}}, \theta^*(k); \frac{\underline{a}C_i(k,l)}{\sqrt{1+\underline{a}^2 C_i(k,l)^2}}\right) \\
&\quad + \Phi_2\left(\frac{D_i(k,l,x,\overline{a})}{\sqrt{1+\overline{a}^2 C_i(k,l)^2}}, -\theta^*(k); \frac{-\overline{a}C_i(k,l)}{\sqrt{1+\overline{a}^2 C_i(k,l)^2}}\right) \\
&\equiv A(k,l,x)
\end{aligned}$$

where

$$\begin{aligned}
D_i(k,l,x,y) &= \frac{x - m_i(l) - y m(k)}{v_i(l)} = \frac{x - l\mu_i - yk\mu}{\sqrt{\sigma_i^2 T + l\gamma_i^2}}; \\
C_i(k,l) &= \frac{\sqrt{\sigma^2 T + k\gamma^2}}{\sqrt{\sigma_i^2 T + l\gamma_i^2}}.
\end{aligned}$$

Removing the conditioning on $N(T)$ and $N_i(T)$ then finally yields

$$\begin{aligned}
\Pr\left(X_i(T) \le x\right) &= \sum_{k=0}^{\infty}\sum_{j=0}^{\infty} w_i(k,l,T) A(k,l,x) \\
&= \sum_{k=0}^{\infty}\sum_{l=0}^{\infty} e^{-(\lambda+\lambda_i)T} \frac{\lambda^k \lambda_i^l T^{k+l}}{l!k!} A(k,l,x). \qquad (30)
\end{aligned}$$

In practice, we often can get away with evaluating less than 10 or 15 of the terms in the double-sum in (30).

While we do not bother listing the results here, we note that several summary statistics are available in closed form for the $X_i(T)$. This includes the mean and variance, quantities that are useful should we wish to normalize the distribution for the $X_i(T)$ in some way (see [2] for a discussion).

References

[1] H Albrecher, S Ladoucette, and W Schoutens. A generic one-factor Levy model for pricing synthetic cdos. Working paper, May 2006.

[2] L Andersen and J Sidenius. Extensions of the Gaussian copula. *Journal of Credit Risk*, 1(1):29–70, 2004/2005.

[3] L Andersen and J Sidenius. CDO pricing with factor models: Survey and comments. *Journal of Credit Risk*, 1(3), 2005.

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

[5] P Baheti. Pricing exotic tranches. Presentation for quant congress 2006, ny, Lehman Brothers, July 2006.

[6] M Baxter. Dynamic modelling of single-name credits and CDO tranches. Working paper, Nomura Fixed Income, March 2006.

[7] J Brasch. A note on efficient pricing and risk calculation of credit basket products. Working paper, TD Securities, 2004.

[8] J Gregory and J-P Laurent. I will survive. *RISK*, June, 2003.

[9] D Guegan and J Houdain. Collateralized debt obligations pricing and factor models: A new methodology using normal inverse gaussian distributions. Working paper, Ecole Normale Superieure, June 2005.

[10] J Hull and A White. Valuing credit default swaps II: Modeling default correlations. *Journal of Derivatives*, 3(5):12–21, 2001.

[11] J Hull and A White. Valuation of a CDO and an Nth-to-default CDS without Monte Carlo simulation. Working paper, University of Toronto, 2004.

[12] J Hull and A White. The perfect copula. Working paper, University of Toronto, 2005.

[13] M Joshi and A Stacey. Intensity gamma: A new approach to pricing portfolio credit derivatives. Working paper, Royal Bank of Scotland, 2005.

[14] A Kalemanova, B Schmid, and R Werner. The normal inverse gaussian distribution for synthetic cdo. Working paper, August 2005.

[15] L McGinty, E Beinstein, R Ahluwalia, and M Watts. Credit correlation: A guide. Technical paper, JP Morgan, March 2004.

[16] L Overbeck and W Schmidt. Modeling default dependence with threshold models. Working paper, March 2003.

[17] P Schonbucher. Portfolio losses and the term structure of loss transition rates: A new methodology for the pricing of portfolio credit derivatives. Working paper, ETH Zurich, September 2005.

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

[19] S Willeman. Fitting the CDO correlation skew: A tractable structural jump model. Working paper, Aarhus Business School, 2005.