Gaussian copula approximations and their applications

Philippos Papadopoulos*

February 27, 2008

DRAFT

Abstract

We examine the standard Gaussian copula model for correlated defaults (also called the survival copula) and its relationship with the theoretically richer model based on diffusion processes and default thresholds. We show that in a discrete time framework the Gaussian copula can be seen as a simple global approximation to the Brownian copula implied by correlated diffusions. More precisely, the Gaussian copula is *larger*, in the sense that it induces systematically higher dependency. The result helps clarify some of the peculiar aspects of the standard model. Turning the argument around, the framework allows us to design further applications, capitalizing on the model's notable tractability. We show that a multi-period, correlated migrations model based on discrete diffusions can for some purposes be replicated surprisingly well by very simple analytic expressions. As a practical example we show how one can readily compute the distributions of the "weighted average rating factor" of a pool of credits for any future time period.

1 Introduction

The last ten years have witnessed significant advances in the theory and practise of credit risk modeling, encouraged by a confluence of regulatory initiatives, internal risk management developments in financial institutions and the explosive development of credit as an asset class. Even as all three contributing agents seem to be facing a period of introspection, some key advances have become (and will likely continue as) mainstays of the industry. Arguably one of the most prominent of those advances is the improved understanding of the basic building block of credit risk modeling, namely the construction of loss distributions via some specification of correlated defaults.

A key idea in this direction has been the "Gaussian copula" or "survival copula" [1] which correlates default times in a mathematically very tractable manner. The basic model has been used widely in practise, extended in various ways, see e.g., [2] and has reinforced the trend of studying and using copula functions in quantitative finance [3]. On the other hand it has been criticized for various shortcomings, especially its *static* aspect, which turned out to be difficult to overcome and hence reduced the model usefulness in a continuous trading context. As a static model it has been recognized [4] as equivalent to the one-period structural type models such as the CreditMetrics model [5] used widely in portfolio management. This equivalence is then the meeting point with another long line of academic and applied research which emanates from structural models of corporations and endogenous assessment of default risk. Notable in this line of work has been the

*Group Risk Management, ABN AMRO B.V., Netherlands. All opinions expressed in this article are the responsibility of the author only and do not necessarily reflect those of his employer. This is a draft version. Comments are welcome at philippos.o.papadopoulos@gmail.com

discussion of default correlation between firms [6]. Establishing the relationships of various portfolio models has been the theme of early and influential studies such as [7] and [8]. Structural approaches to portfolio credit risk in a multi-period context continue to be popular in practise, see e.g., [9], [10] but the precise relationships (and possible advantages) of multi-period diffusion based models in comparison to the survival copula model remained somewhat unclear. In a comparative study [11] pointed out the close numerical match of the two approaches when applied to CDO valuation.

An exploration of the relationship of the standard Gaussian copula model and the discrete multi-period diffusion will be the starting point and first contribution of this paper: We will examine and clarify the relationship in the context of multivariate normal copula functions. The outcome of the first section motivates the next contribution, which is an extension of the standard Gaussian copula framework so as to model credit migrations. Although an approximation by design, it turns out that for many applications the results are extremely close to an equivalently calibrated diffusion model. We illustrate this fact more clearly by means of the third contribution, which is to derive easily computable expressions for the distribution of the weighted average rating of a credit portfolio (commonly abbreviated as WARF by practitioners). We compare the accuracy of those expressions against detailed simulation results.

2 The standard Gaussian (survival) copula model as an approximation

It will be useful to imagine already a population of M reference *names*, e.g., corporate issuers, indexed by $i = 1, \ldots, M$. We demonstrate here that the copula underlying the survival copula can be seen as a certain approximation to a copula with a more "natural" dependency structure. We recap first some aspects of the standard Gaussian copula [1]: The model aims to model directly the default time τ^i of the i-th name. The credit status of each name i is governed by the single random variable W^i, distributed like $N(0,1)$. Given the realization of the credit variable (also called latent credit factors) W^i and a default curve $q(t) = P(\tau \leq t)$, the default time τ^i is given by $\tau^i = q^{-1}(N(W^i))$. A correlation structure is specified, e.g., using factor decompositions such as $W^i = \rho Z + \sqrt{1-\rho^2}\epsilon^i$ which imply dependency via the copula C_S.

$$C_S(t_1, \ldots, t_M; \Sigma) = N_\Sigma(N^{-1}(q(t_1)), \ldots, N^{-1}(q(t_M)))\,, \tag{1}$$

where N_Σ is the multivariate normal distribution with correlation matrix Σ given by $\Sigma_{ij} = \rho^2 + \delta^{ij}(1-\rho^2)$[1] and $\rho \in [0,1]$. The model is pleasantly simple. One of its perplexing characteristics is that the realization of the set of random variables (Z, ϵ^i) etc. does not seem to be localized in any particular point in time. To make this more precise we will first turn to the more familiar family of Wiener processes. We will work in a discrete time framework with the time horizon of interest denoted as T, the present time as $t_0 = 0$ and N equal time intervals $\delta_k = t_k - t_{k-1} = \delta$ between the timepoints t_k, $k = 1, \ldots, N$. Each name (and eventually its credit quality status) will be associated with a discrete credit process W_k^i, loosely linked to asset or equity value, which is built up by adding independent normal increments:

$$W_k^i = \sum_{l=1}^{k} \Delta W_l^i \tag{2}$$

We have thus (for each i) the marginal distributions $P(W_k \leq x) = F_k(x) = N(x; t_k)$, where $N(x; \sigma^2)$ denotes a Gaussian distribution of zero mean and variance σ^2 whereas the dependency (again, for each i) is described by the "discrete time" Brownian copula C_B (the continuous time version first discussed in [12]):

$$C_B(u_1, \ldots, u_N) = N_b(F_1^{-1}(u_1), \ldots, F_N^{-1}(u_N)) \tag{3}$$

where N_b is the multivariate normal distribution with correlation matrix b given by $b_{kl} = \frac{t_k \wedge t_l}{\sqrt{t_k t_l}}$. [2]

We introduce dependency among the M issuers using a single factor correlation structure:

$$W_k^i = \rho Z_k + \sqrt{1-\rho^2}\epsilon_k^i \tag{4}$$

[1]The expression δ^{ij} is unity when $i = j$ and zero otherwise

[2]The expression $t_k \wedge t_l$ denotes the minimum of the variables t_k, t_l

where Z_k and ϵ_k^i are discrete Wiener processes, as per equation (2) above, with the standard marginal distributions. The dependency, including now both the time and name dimensions, is given by the "Generalized Brownian" copula (C_{GB})

$$C_{GB}((u_1^1,\ldots,u_N^1),\ldots,(u_1^M,\ldots u_N^M);\hat{\rho}) = N_{\hat{\rho}}((F_1^{-1}(u_1^1),\ldots,F_N^{-1}(u_N^1)),\ldots,(F_1^{-1}(u_1^M),\ldots,F_N^{-1}(u_N^M))) \quad (5)$$

where the correlation matrix $\hat{\rho}$ is given by

$$\hat{\rho}_{kl}^{ij} = \mathrm{Corr}(W_k^i, W_l^j) = (\rho^2 + \delta^{ij}(1-\rho^2))\frac{t_k \wedge t_l}{\sqrt{t_k t_l}} \,. \quad (6)$$

In contrast with the survival copula (1), the flow of information encapsulated in (4) is more transparent: The increments (innovations) of process W_k^i can be intuitively associated with the time points t_k. In order to understand better the structural relationship, we are motivated to think in terms of the well known Karhunen-Loeve expansion, which (for any given time horizon T) represents the Wiener process via an infinite trigonometric series expansion. Whereas the time-domain representation (2) builds the Wiener process via the successive addition of independent random increments, the spectral (also called frequency domain) representation builds the process by the successive addition of sinusoidal patterns of random amplitude. Clearly the incremental information added by every higher harmonic component is now "global" in nature since new increments modify the process paths at *all* times $t \in [0,T]$. Unfortunately the KL expansions are not well suited for our purposes, because the marginal distributions of any finite truncation of the series are different. Hence it becomes more difficult to separate out and compare the copula function. Instead, we ask whether we can construct finite approximations that have the same marginal distributions as the discrete process (2). I.e., we look for a finite sum of the form

$$\bar{W}_t(n) = \sum_{l=1}^{n} \sqrt{t}\beta_l(t) Y_l \,, \quad (7)$$

where Y_l are independent random variables and the expansion coefficients $\beta_l(t)$ satisfy the relation $\sum_{l=1}^{n} \beta_l(t)^2 = 1$ for all times t. We are not aware of the solution to this problem for general n, but already the simplest expansion ($n=1$) seems particularly useful. In the $n=1$ case, we have $\bar{W}_t(1) = \sqrt{t}Y$ and we recognize the crude global representation of the Wiener process whereby the path is simply the graph of the square-root function multiplied by a random amplitude of unit variance (see also Fig. 1). In discrete time we have the corresponding process $\bar{W}_k = \sqrt{t_k}Y$ as the simplest *global approximation* to the discrete Wiener process. With a suitable calibration of thresholds (see next section) we will reproduce precisely the discrete version of the survival copula model. For now we concentrate on the dependency structure. The marginal distributions at each timepoint t_k are by construction correct. The dependency (in the time dimension) is clearly given by the maximum copula (or upper Fretchet bound)

$$C^+(u_1,\ldots,u_N) = \min(u_1,\ldots,u_N) \,, \quad (8)$$

consistent with 100% dependency among the variables $\bar{W}_k$, given that there is really only one independent random variable underlying the entire path.

Continuing with the single factor correlation framework we introduce the multi-name process

$$\bar{W}_k^i = \rho\sqrt{t_k}Z + \sqrt{1-\rho^2}\sqrt{t_k}\epsilon^i \,, \quad (9)$$

with dependency described by the "approximation" copula

$$C_A((u_1^1,\ldots,u_N^1),\ldots,(u_1^M,\ldots u_N^M);\bar{\rho}) = N_{\bar{\rho}}((F_1^{-1}(u_1^1),\ldots,F_N^{-1}(u_N^1)),\ldots,(F_1^{-1}(u_1^M),\ldots,F_N^{-1}(u_N^M))) \quad (10)$$

where the correlation matrix is given by

$$\bar{\rho}_{kl}^{ij} = \mathrm{Corr}(\bar{W}_k^i, \bar{W}_l^j) = \rho^2 + \delta^{ij}(1-\rho^2) \quad (11)$$

Already from the intuitive motivation of process (9) and the comparison of the correlation matrices (6) and (11) we can surmise that C_A is a copula which relative to the Brownian copula C_{GB} enhances dependency along the time dimension.

Proposition: More precisely, it is easy to show that $C_{GB} < C_A$, or that the Generalized Brownian copula is *smaller* than the approximation copula, in the sense that $\forall ((u_1^1, \ldots, u_N^1), \ldots, (u_1^M, \ldots u_N^M)) \in \mathbf{I}^{N\times M}$,

$$C_{GB}((u_1^1, \ldots, u_N^1), \ldots, (u_1^M, \ldots u_N^M); \hat{\rho}) \leq C_A((u_1^1, \ldots, u_N^1), \ldots, (u_1^M, \ldots u_N^M); \bar{\rho}) . \tag{12}$$

Proof: In compact form we represent both copula functions as

$$C(\bar{u}; \rho_i) = N_{\rho_i}(\bar{X}) , \tag{13}$$

with $i \in (GB, A)$ and $\bar{X} = \{F^{-1}(u)\}$ is shorthand notation for the collection of arguments. First we recall that the corresponding correlation matrices (6) and (11) obey the inequality $(\rho_{GB} = \hat{\rho}) \leq (\rho_A = \bar{\rho})$ given that $t_k \wedge t_l / \sqrt{t_k t_l} \leq 1$ for all pairs of time-points. Next, for any fixed point $\bar{u}$, the copulas (13) are functions only of their correlation argument. Hence can always link the two copulas via a line integral in "correlation space"

$$C(\bar{u}; \rho_A) = \int_{\rho_{GB}}^{\rho_A} \nabla_\rho N_\rho(\bar{X}) \cdot d\rho + C(\bar{u}; \rho_{GB}) . \tag{14}$$

Next, given the ordering of the correlation matrices, we can construct a path with exclusively positive correlation increments $(d\rho)$ leading from ρ_{GB} to ρ_A. Finally, it suffices to show that the gradient of the copula function with respect to any of the correlation matrix elements is positive. This follows from explicit differentiation of the multivariate normal distribution function (see e.g. the appendix of [13]). Hence the copula ordering follows.

To summarize, we have shown that starting from the copula of the "natural" discrete, correlated, Brownian process we can obtain a *larger* copula by essentially eliminating any small scale inter-temporal structure of the underlying processes. This relationship clarifies in an intuitive manner the peculiarity of the information flow implied by the survival copula model.

Copula Name	Symbol Used	Dimensions	Correlation matrix
Survival	C_S	M	$\rho^2 + \delta^{ij}(1-\rho^2)$
Brownian	C_B	N	$t_k \wedge t_l / \sqrt{t_k t_l}$
Frechet copula	C^+	N	N/A
Generalized Brownian	C_{GB}	$M \times N$	$(\rho^2 + \delta^{ij}(1-\rho^2))\frac{t_k \wedge t_l}{\sqrt{t_k t_l}}$
Approximation	C_A	$M \times N$	$\rho^2 + \delta^{ij}(1-\rho^2)$

Table 1: Summary of the various copula functions and their relationships. All but the Frechet copula belong to the family of multivariate normal copulas, but have different dimensionality and parametrization

3 Threshold based migration processes

Whereas the survival copula directly links the timing of *default* events, with the current reinterpretation of the model we can envisage using this approximation for modeling more general credit states (ratings) given that we have an approximation for the credit factor at all intermediate times. We turn now our attention to this task. A name is at any time described by a *rating state* S_k, taking values in $[1, \ldots, D]$. D is the default state, which is assumed to be absorbing, i.e., no re-emergence from default. Hence, denote S_k^i the rating state of the i-th name during the k-th period. We aim to link S_k^i to the set of continuous random variables W_k^i in order to utilize the machinery of continuous copulae.

A key empirical input will be the set of *multi-period transition matrices* T_{0l}^{mn}. Depending on the application, this set of matrices would either be obtained entirely from historical multi-year migration data [14] or based on pricing information. For simplicity we assume homogeneity, i.e., all names are governed by the same transition matrix (but can obviously have different initial states S_0^i). This set of matrices encapsulates the probabilities that a name will move from rating state m at time 0 to state n at time l, i.e.,

$$T_{0l}^{mn} = P(S_l^i = n | S_0^i = m)\,. \tag{15}$$

In line with this probabilistic interpretation, we require that each matrix T_{0l}^{mn} is an $D \times D$ dimensional real matrix with strictly positive entries and rows summing to unity. The existence of an absorbing state implies that $T_{0l}^{Dn} = 0$ and the requirement that the default transition rates are strictly increasing in time:

$$k < l \Rightarrow T_{0k}^{mD} < T_{0l}^{mD}\,. \tag{16}$$

We do not require that the rating process S_k^i follows a Markov process. In fact, for a multi-period threshold model that is based on an underlying process W_k^i it appears difficult to enforce the Markov property at the rating state level. This is because time t_k realizations of W_k^i (which determine the rating state) carry more information than the rating state itself and this additional information also depends on the previous time period realizations of W_{k-1}^i. The best that can be achieved in this direction is to mimic Markov transitions by setting the multi-period transition matrices equal to the respective powers of the one-period transition matrix, i.e., $T_{0k} = (T_{01})^k$, but this procedure does not guarantee e.g., that $T_{kl} = (T_{01})^{l-k}$.

We introduce some additional shorthand notation for the migrations to default: Incremental default probabilities during period k, given an initial rating of m, are denoted by p_k^m. Cumulative probabilities up to period k, given an initial rating of m, i.e., $q_k^m = T_{0k}^{mD}$, are then given by $q_k^m = \sum_{l=1}^k p_k^m$.

3.1 Calibration of multi-period rating migration thresholds

Our aim here is to calibrate suitable *thresholds* in the range of the process W_k^i governed by Eq. 4 so that cumulative probabilities of crossing those thresholds reproduce the probabilities embedded in the multi-period transition matrix. More precisely, we aim to derive the rating state S_k, during period k, of given name, with an initial rating state $S_0^i = m$, by examining the realization of the random variables $\{W_l, l = 1, \dots k\}$ in comparison with a set of thresholds A_k^{mn}. The "natural" specification is to posit the following link

$$\begin{aligned} S_k &= \mathbb{1}_{\left\{(W_k > A_k^{m1}) \bigcap_{l=1}^{k-1} (W_l > A_l^{mD})\right\}} + \sum_{n=2}^{m-1} n \mathbb{1}_{\left\{(A_k^{mn} < W_k < A_k^{m\,n-1}) \bigcap_{l=1}^{k-1} (W_l > A_l^{mD})\right\}} \\ &+ m \mathbb{1}_{\left\{(A_k^{m\,m+1} < W_k < A_k^{mm-1}) \bigcap_{l=1}^{k-1} (W_l > A_l^{mD})\right\}} + \sum_{n=m+1}^{D-1} n \mathbb{1}_{\left\{(A_k^{m\,n+1} < W_k < A_k^{mn}) \bigcap_{l=1}^{k-1} (W_l > A_l^{mD})\right\}} \\ &+ D \mathbb{1}_{\left\{\bigcup_{l=1}^{k} (W_l < A_l^{mD})\right\}}\,, \end{aligned} \tag{17}$$

which simply means that the thresholds are ordered (with the lowermost A_l^{mD} always associated with default). Figure 1 illustrates a Wiener path and the arrangement of the thresholds. A non-default rating state during some period is obtained by exceeding the default threshold at all previous periods and falling within the appropriate "bucket" during the current period. The thresholds are obtained by solving the equations:

$$\begin{aligned} T_{0k}^{mD} &= P(\bigcup_{l=1}^{k} (W_l < A_l^{mD})) \\ T_{0k}^{mn} &= P((A_k^{mn} < W_k < A_k^{m\,n-1}) \bigcap_{l=1}^{k-1} (W_l > A_l^{mD})) \qquad : \quad m < n \\ T_{0k}^{mm} &= P((A_k^{m\,m+1} < W_k < A_k^{m\,m-1}) \bigcap_{l=1}^{k-1} (W_l > A_l^{mD})) \\ T_{0k}^{mn} &= P((A_k^{m\,n+1} < W_k < A_k^{mn}) \bigcap_{l=1}^{k-1} (W_l > A_l^{mD})) \qquad : \quad m > n \\ T_{0k}^{m1} &= P((W_k > A_k^{m1}) \bigcap_{l=1}^{k-1} (W_l > A_l^{mD})) \end{aligned} \tag{18}$$

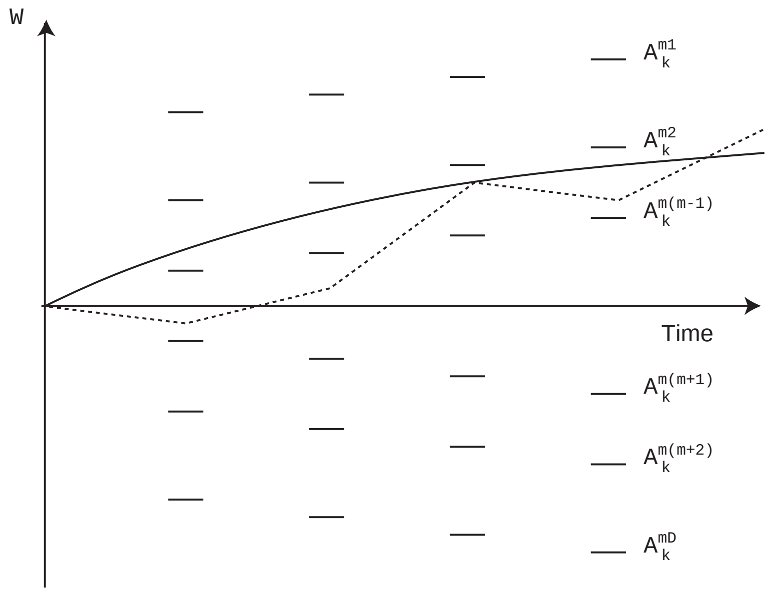

Figure 1: Visual illustration of the processes W_k and $\bar{W}_k$ and associated migration thresholds as a function of time. The default thresholds are the lowest (most negative) at each time period. The dotted line represents a realization of the discrete Wiener process, whereas the continuous line represents a realization of the $n = 1$ truncated model. The thresholds are arranged to illustrate an upward migration followed up by a steady-state evolution for the approximate model along with a slightly different rating path for the full model. For clarity only a single set of thresholds is depicted but those of-course would be in general be different for the two models.

where $n = 2, \ldots, D-1$. For each initial rating m, this is a system of $k \times D$ equations for the same number of unknowns. A numerical solution algorithm for the *default* thresholds A_k^{mD} was described in [15]. In the appendix we show how to extend the method to calibrate migration thresholds (a straightforward exercise, since the non-default states are all non-absorbing).

3.2 Migration thresholds in the survival copula approximation

First we discuss the calibration of the default thresholds: The name i defaults at the first time point t_k such that $\bar{W}_k^i \leq \hat{\alpha}_k$. We calculate thresholds $\hat{\alpha}_k$ satisfying

$$P(\bar{W}_k^i < \hat{\alpha}_k) = q_k \,, \tag{19}$$

which leads to

$$\hat{\alpha}_k = \sqrt{t_k} N^{-1}(q_k) \,. \tag{20}$$

Following the intuition that the simplified credit process $\bar{W}_k^i$ is an approximation to the more natural process W_k^i we can now easily extend the calibration to incorporate rating migration. The rating state formula in this case is given by

$$\begin{aligned} S_k &= \mathbb{1}_{\left\{\bar{W}_k > A_k^{m1}\right\}} + \sum_{n=2}^{m-1} n \mathbb{1}_{\left\{A_k^{mn} < \bar{W}_k < A_k^{m\,n-1}\right\}} \\ &+ m \mathbb{1}_{\left\{A_k^{m\,m+1} < \bar{W}_k < A_k^{mm-1}\right\}} + \sum_{n=m+1}^{D-1} n \mathbb{1}_{\left\{A_k^{m\,n+1} < \bar{W}_k < A_k^{mn}\right\}} \\ &+ D \mathbb{1}_{\left\{\bar{W}_k < A_k^{mD}\right\}} \end{aligned} \tag{21}$$

The formula is considerably simplified because for deriving the state at any given period k there is no need to condition on survival up to period $k-1$; such survival is implicit given the known path structure. Figure 1 illustrates again a path and corresponding thresholds (generally not the same as the Wiener thresholds!). The migration thresholds are now obtained by inverting equations

$$\begin{aligned} T_{0k}^{mD} &= P(\bar{W}_k < A_k^{mD}) \\ T_{0k}^{mn} &= P(A_k^{m\,n+1} < \bar{W}_k < A_k^{m\,n}) \qquad : \quad m < n \\ T_{0k}^{mm} &= P(A_k^{m\,m+1} < \bar{W}_k < A_k^{m\,m-1}) \\ T_{0k}^{mn} &= P(A_k^{m\,n} < \bar{W}_k < A_k^{m\,n-1}) \qquad : \quad m > n \\ T_{0k}^{m1} &= P(\bar{W}_k > A_k^{m1}) \end{aligned} \tag{22}$$

which can be inverted considerably easier than Eq. 18 with the threshold given in fact explicitly as

$$A_k^{mn} = \begin{cases} \sqrt{t_k} N^{-1}(\sum_{r=n+1}^{D} T_{0k}^{mr}) & : \quad m > n \\ \sqrt{t_k} N^{-1}(\sum_{r=n}^{D} T_{0k}^{mr}) & : \quad m < n \end{cases} \tag{23}$$

We note that while the calibration of migration thresholds for each individual time period t_k can be seen independently as a single step migration model, the approximate model implies a very specific temporal structure of migrations. E.g., conditioning on a certain final rating state, the intermediate period rating states must follow a monotonic path. This feature is the extreme opposite of the sequential application of a single step migration model, which does not exhibit any serial dependence in the rating paths. For some applications this rigidity of the rating paths might be unacceptable, while elsewhere the tractability advantage may be a reasonable tradeoff.

4 Portfolio rating distributions, the large pool limit and applications

In the previous two sections we have achieved the following two results: In Section 2 we found that the survival copula model can be seen as a global approximation to correlated diffusions. In Section 3 we utilized this approximation to define multi-period migration in a very tractable manner, (at the possible expense of losing accuracy in certain circumstances). Now we continue developing this simplified correlated migrations framework with an eye towards practical applications but we will also study some aspects of the quality of the approximation versus the underlying diffusion model. The focus here will be in marginal rating distributions at a fixed time period, so we might be expecting this to be a regime with reasonable performance.

First, some additional notation. Most portfolio credit risk research focuses on the credit loss distribution L_k. In a credit migration context (at portfolio level) we have instead multiple distributions N_k^r, i.e., a D-dimensional family of correlated processes capturing the number of names having the r-th state during the k-th period:

$$N_k^r = \sum_{i=1}^{M} \mathbb{1}_{\{S_k^i = r\}} \,. \tag{24}$$

The initial rating state per name is assumed to be known, and given by S_0^i (which determines N_0^r). More generally, we introduce the general rating dependent process R_k^r

$$R_k^r = \sum_{i=1}^{M} w_k^i(r) \mathbb{1}_{\{S_k^i = r\}} \,, \tag{25}$$

where $w_k^i(r)$ are known weights per name, period and rating, which can be chosen so as to make the rating process sensitive to notional exposure, valuation per rating etc.

Two key practical questions are how to compute the distributions of such processes and the accuracy of the survival copula approximation (21) in this respect. It is clear that the full model (17) requires a simulation approach because the dimensionality of the associated integrals is at least equal to the number of periods considered.[3] While a perfectly viable option on a stand-alone basis, detailed simulations of migrations can be cumbersome e.g., if the calculation must be embedded in some larger computation. We may suspect instead that the excellent tractability properties of the survival copula model which have been used extensively in default risk modeling can also be extended in the area of rating migrations with minimal effort.

4.1 Large pool limit

We seek to compute the limiting case

$$P(\tilde{N}_k^r \le x) = P\left(\lim_{M\to\infty} \frac{1}{M}\sum_{i=1}^{M} \mathbb{1}_{\{S_k^i = r | S_0^i\}} < x\right) \tag{26}$$

utilizing the standard techniques of conditional independence and the law of large numbers. We use a tilde to denote that in the limit we need the portfolio *fraction* rather than the absolute number of names in a given rating state. It is useful to partition the summation according to the initial rating

$$\frac{1}{M}\sum_{i=1}^{M} \mathbb{1}_{\{S_k^i = r | S_0^i\}} = \frac{1}{M}\sum_{m=1}^{D}\sum_{i=1}^{N_0^m} \mathbb{1}_{\{S_k^i = r | S_0 = m\}} \,. \tag{27}$$

Conditional on the systematic factor Z, the random variables $\mathbb{1}_{\{S_k^i = r | S_0^i, Z\}}$ are independent and identically distributed and hence their limiting sum converges to their expected value

$$\tilde{N}_0^m \mathbb{E}\left[\mathbb{1}_{\{S_k = r | S_0 = m, Z\}}\right] \,. \tag{28}$$

Putting things together,

$$P(\tilde{N}_k^r \le x) = \int P\left(\sum_{m=1}^{D} \tilde{N}_0^m \mathbb{E}\left[\mathbb{1}_{\{S_k = r | S_0 = m, Z\}}\right] \le x\right) dN(z) \tag{29}$$

$$= \int \mathbb{1}_{\{\sum_{m=1}^{D} \tilde{N}_0^m \mathbb{E}\left[\mathbb{1}_{\{S_k = r | S_0 = m, Z\}}\right] \le x\}} dN(z) \,. \tag{30}$$

To conclude the derivation we need to write down the conditional migration probabilities

$$Q_{0k}^{mr}(z) = \mathbb{E}\left[\mathbb{1}_{\{S_k = r | S_0 = m, Z\}}\right] \,. \tag{31}$$

Those are derived by simply integrating the conditional probability of hitting any of the intervals between thresholds A_k^{mr}. Explicitly, the three distinct expressions are

$$Q_{0k}^{mD}(z) = N\left(\frac{N^{-1}(T_{0k}^{mD}) - \rho z}{\sqrt{1-\rho^2}}\right) \tag{32}$$

$$Q_{0k}^{mr}(z) = N\left(\frac{N^{-1}(\sum_{s=r}^{D} T_{0k}^{ms}) - \rho z}{\sqrt{1-\rho^2}}\right) - N\left(\frac{N^{-1}(\sum_{s=r+1}^{D} T_{0k}^{ms}) - \rho z}{\sqrt{1-\rho^2}}\right) \tag{33}$$

$$Q_{0k}^{m1}(z) = N\left(\frac{N^{-1}(\sum_{s=2}^{D} T_{0k}^{ms}) - \rho z}{\sqrt{1-\rho^2}}\right) \tag{34}$$

In conclusion, the cumulative probability distribution of names rated in the r-th rating class by period k is given by the expression

$$P(\tilde{N}_k^r < x) = \int \mathbb{1}_{\{\sum_{m=1}^{D} \tilde{N}_0^m Q_{0k}^{mr}(z) \le x\}} dN(z) \,, \tag{35}$$

[3] In the best case scenario, namely a single-factor model where idiosyncratic contributions can be assumed diversified away within an infinitely large pool

which is easy to compute numerically. The derivation generalizes for any linear function (Eq. 25). In fact, while we restricted our focus here to the aggregation of same period statistics, practically *any* measure of interest that can be constructed from the rating state variables S_k is reduced in the large pool limit to a one-dimensional integral.

4.2 Calculation of the weighted average rating distribution

One simple and useful rating based measure is the weighted average rating of the portfolio. While the weights can be chosen according to any desired principles, one well established convention is to use the long-term default probabilities of various rating classes as weights. Estimating the WAR at any time point t_k for the portfolio is thus giving an immediate indication of future losses (say at time $t_k + t_N$) as seen at time t_k (assuming of-course that the future term structure of default probabilities is the same as at $t = 0$). More precisely, lets define for each period

$$A_k = \sum_{r=1}^{D} w_r N_k^r \tag{36}$$

where

$$w_r = \frac{q_T^r}{\sum_{r=1}^{D} q_N^r} \tag{37}$$

and q_T^r is the cumulative default probability per rating class for the final period N. We choose to normalize the weights in the above manner so that in combination with normalized population counts, ($\tilde{N}_k^r = N_k^r/M$), the range of WAR is in $[0, 1]$.

The distribution of WAR is derived analytically with the obvious re-application of the results of the previous subsection:

$$P(A_k \leq x) = \int \mathbb{1}_{\{\sum_{r=1}^{D}\sum_{m=1}^{D} w_r \tilde{N}_0^m Q_{0k}^{mr}(z) \leq x\}} dN(z) \tag{38}$$

This expression leads to useful analytic formulas for the volatility of WAR as function of correlation, which provides even simpler approaches to quantifying migration risk, but we don't pursue this further and instead we turn now to perform a numerical analysis for stylized credit portfolios.

4.3 Numerical Results

We now turn to a quantitative comparison of the two approaches. For the simulation model, we'll assume a homogeneous portfolio of $M = 400$ names, e.g., we will assume the same initial rating $S_0^i = r$ for all names. We fix the maximum period to five years ($N = 5$) and use annual steps. The multi-year transition matrices are obtained from a stylized one-year transition matrix by multiplication (See appendix). The correlation parameter is set to 30% unless otherwise noted. The analytic approximation uses always identical calibration, with the exception ofcourse of the number of names which is implicitly infinite.

In the first set of results we set all names initially at the $S_0 = 4$ state. The initial WAR is thus computed as 1.53% whereas the *expected* WAR in periods 1 to 5 (as computed directly from the transition matrices) is respectively $\{2.01\%, 2.53\%, 3.09\%, 3.67\%, 4.27\%\}$. The full distributions are given in Fig. 2, derived both by simulation and using the analytic approximation. The plots are for five successive periods (annual), with the leftmost distribution corresponding to the first period. We note the heavy right tail of the WAR distribution, i.e., the increased likelihood of downward migration. In contrast the likelihood of a significant reduction is extremely small. We notice the excellent agreement between approaches which is quantitatively capture in Table 2.

In the next set of results we examine how the initial rating of the pool affects the distribution and the model accuracy. In Fig. 3 we show the dependency of the WAR distribution on the initial pool rating. Again both simulation and approximate distributions are shown. The results are for initial rating states ($S_0 = 1, 2, 3, 4, 5$), with the leftmost distributions corresponding to the lowest state (better ratings). The correlation parameter is left unchanged and the plots focus on final (5yr) period only.

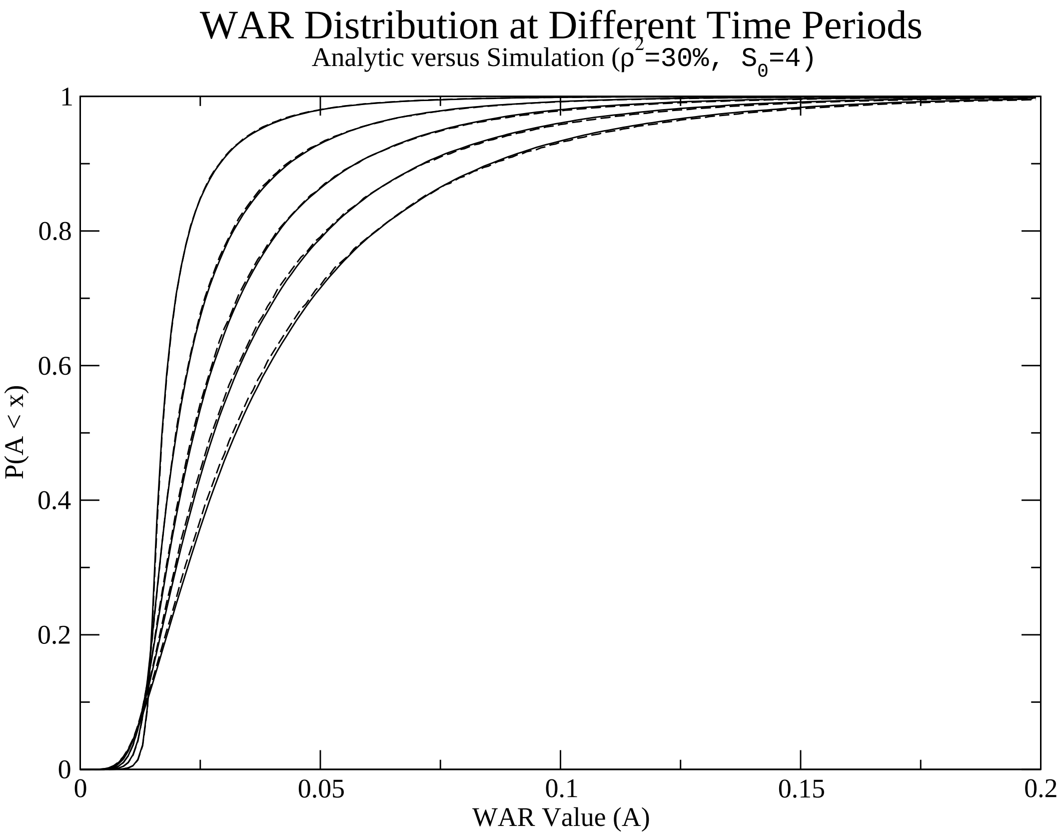

Figure 2: Illustration of the Weighted Average Rating distribution obtain by simulation and analytically. Simulation results are indicated by solid lines whereas the analytic approximation is indicated with dashed lines. Five periods are shown, with the leftmost pair of distributions corresponding to the first period (for clarity only the fraction up to 20% of the range of A_k is shown). The initial rating state is set to $S_0 = 4$ whereas correlation is set at 30%.

Computation of $P(A < x)$				
x	**Sim.**	**Approx.**	**Diff**	**Diff (%)**
0.005	0.0008	0.0008	0.0000	-0.99%
0.010	0.0296	0.0307	-0.0011	3.57%
0.020	0.2457	0.2545	-0.0088	3.47%
0.030	0.4584	0.4681	-0.0097	2.08%
0.040	0.6089	0.6180	-0.0091	1.48%
0.050	0.7152	0.7191	-0.0040	0.55%
0.060	0.7904	0.7911	-0.0007	0.09%
0.070	0.8425	0.8439	-0.0013	0.16%
0.080	0.8832	0.8811	0.0021	-0.24%
0.090	0.9125	0.9100	0.0025	-0.27%
0.100	0.9337	0.9320	0.0017	-0.18%
0.110	0.9497	0.9475	0.0022	-0.24%
0.120	0.9619	0.9591	0.0027	-0.28%
0.130	0.9712	0.9686	0.0026	-0.26%
0.140	0.9783	0.9762	0.0020	-0.21%
0.150	0.9836	0.9817	0.0019	-0.19%
0.160	0.9878	0.9858	0.0020	-0.21%
0.170	0.9910	0.9893	0.0017	-0.17%
0.180	0.9931	0.9918	0.0013	-0.13%
0.190	0.9949	0.9938	0.0011	-0.11%
0.198	0.9959	0.9949	0.0010	-0.10%

Table 2: Illustrative example of the numerical differences between simulation based estimates and the analytic approximation. The results are refer to the same setup as Fig. 2, with a portfolio set at the initial rating state $S_0 = 4$ and correlation at 30%. Only the 5yr results are shown. The fist column denotes the WAR evaluation points (more or less in regular increments). The second and third columns are the Monte Carlo and analytically obtained distributions at those sample points. The fourth column is the numerical difference of the distributions, whereas the final column is the relative difference. Generally the approximation is uniformly good. The relative error increases when the target probability is very small.

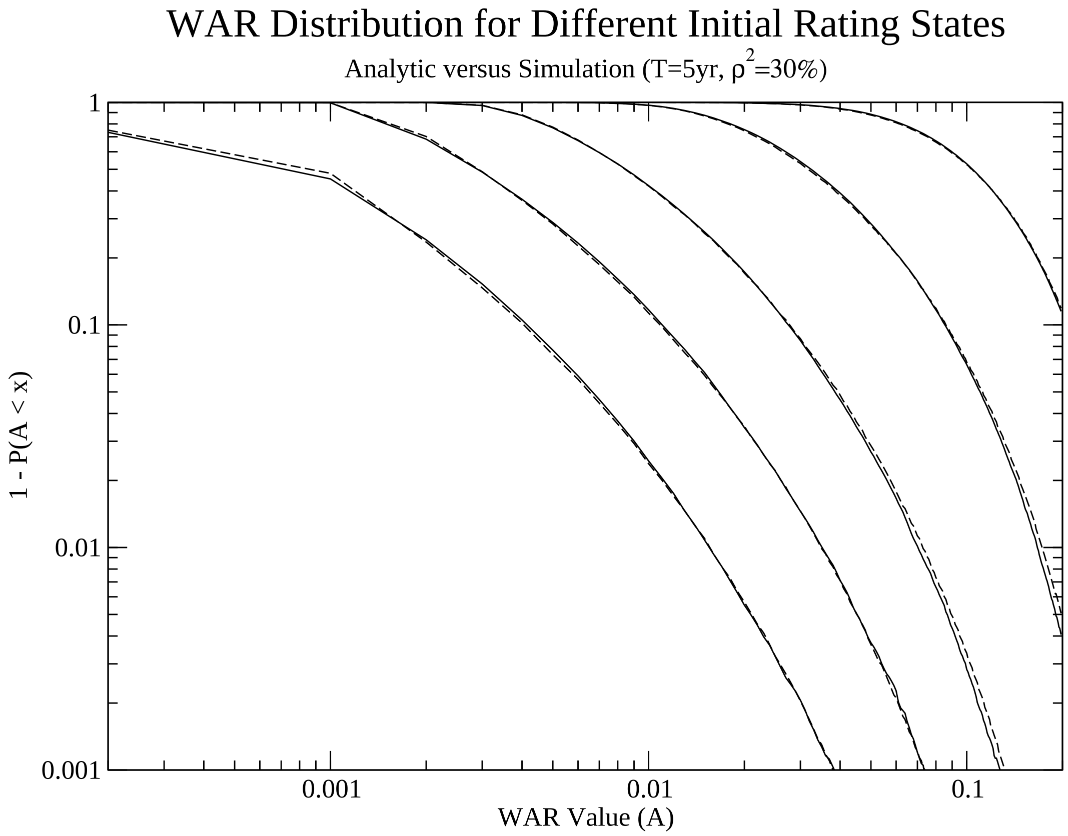

Figure 3: Illustration of the dependence of the WAR distribution on initial rating. Simulation results are indicated by solid lines whereas the analytic approximation is indicated with dashed lines. Five different initial rating state choices are shown, with the leftmost pair of distributions corresponding to $S_0 = 1$. The use of log-log scale illustrates that the distributions for different initial ratings are largely homologous. The kink at the low end of the WAR range is due to the limited set of observation points of the distributions. Importantly, the good match between methods is not (numerically) sensitive to the pool rating

Finally we turn to examine the dependence on asset correlation. In Fig. 4 we get synoptic results using the standard parameters and varying the asset correlation. Again both simulation and approximate distributions are shown. The results are for correlation levels ($\rho^2 = 5\%, 15\%, 30\%, 45\%$), with the leftmost distributions corresponding to the lowest correlation. The insert shows the complementary distributions in logarithmic scale in order to highlight the right tail dependence. The agreement is seen to be uniformly good.

5 Summary and Discussion

Credit risk applications based on either the survival copula or discrete diffusions continue to be widely used by institutions exposed to portfolios of credit risk. The relationship of those two strands has not been clarified, despite various pieces of numerical evidence suggesting their linkages. As as starting point in this work, we found that once we ensure identical marginal distributions of the latent credit factors, the dependency structure of the two models is ordered in a simple and intuitive way. Both dependencies are examples of multi-variate normal copulas and the survival copula is seen to be an approximation to the generalized Brownian copula, in the sense that it is derived by the maximization of the inter-temporal dependencies of the Brownian increments.

The positive impact of the above find is that the simpler and more tractable survival copula framework can be used also in the context of modeling correlated credit migrations. In this direction we constructed,

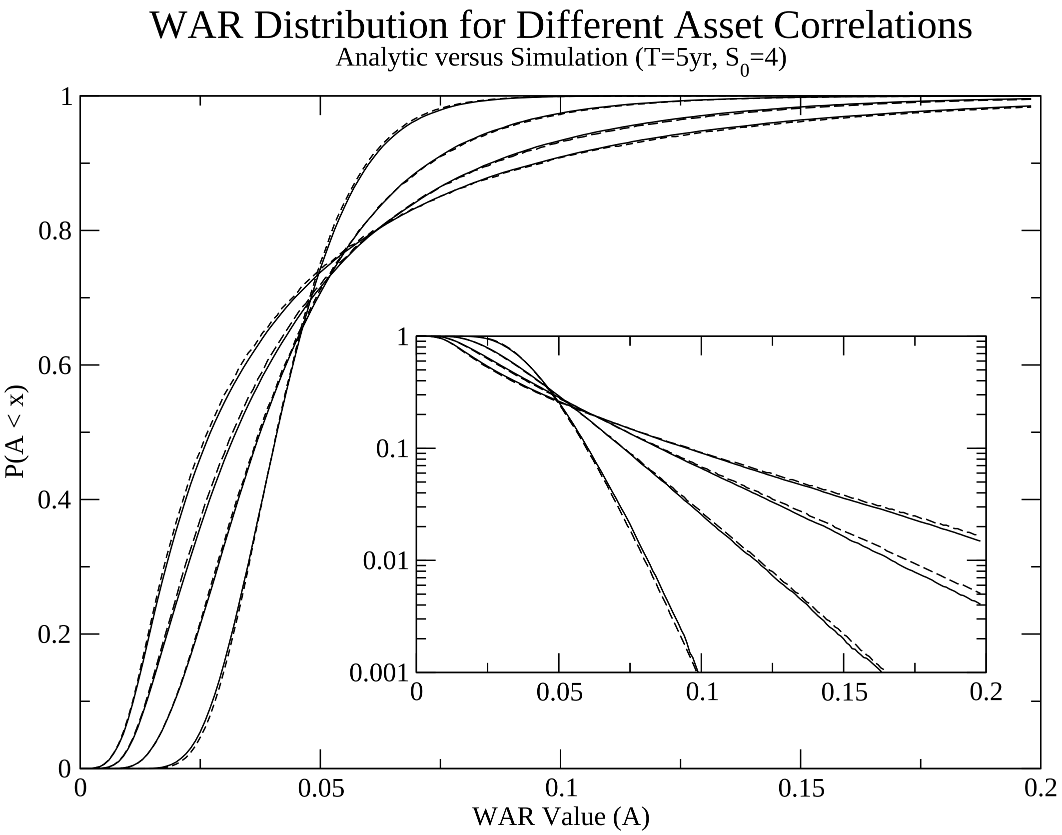

Figure 4: Illustration of the dependence of the WAR distribution on asset correlation. Four different asset correlation levels are shown, with the leftmost pair of distributions corresponding to $\rho^2 = 5\%$. The insert shows the complementary distributions in logarithmic scale and thus highlights the different right tail slopes of the distributions for different correlations. The agreement between methods is evident here as well.

side by side, threshold based migration models. The diffusion based model is a straightforward extension of [15] and shares its non-trivial implementation burden, both in the calibration and simulation phase. The survival copula based model is in contrast very easy to implement.

The approximate model sacrifices some realism in rating migration paths, as it eliminates, by construction, paths of high variance. Hence it would be expected to be a poor approximation for products or measures that depend sensitively to such variance. But this is also a regime where it is by no means assured that diffusion driven migrations are a good model fit. Overall the approximate model would appear to be the simplest way to obtain multi-period portfolio rating distributions that are i) consistent with a set of multi-period transition matrices and ii) exhibit controlled and well understood migration correlation.

We derived an analytic formula for the weighted average rating distribution of a credit pool. We demonstrated the closeness of numerical results (for a reasonably generic credit portfolio) to the simulation based diffusion model. The quality of agreement is more striking given the significantly different computational overhead of the two models: on the one hand a semi-analytic formula, on the other hand a multi-period simulation model.

Throughout the paper we restricted the discussion to single factor correlation models. This was purely a choice for brevity of presentation in Sections 2 and 3 but obviously required for the semi-analytic results of Section 4. In similar vein we did not discuss alternative functional forms for the copula (i.e., besides the Gaussian). To the extend that empirical migrations exhibit higher tail dependency this is a straightforward next step to consider. More generally, the stated good agreement between the two specific models considered here is not intended to imply that actual observed migrations of credit pools exhibit an empirical copula that is close to any of the two. Nevertheless it is likely that the approximation model introduced here is the simplest Gaussian model that can be calibrated to joint multi-period rating migration data (i.e., determining a single correlation parameter), which would be an interesting alternative to the study in [14].

While not shown here, the results do not depend on the discretization interval δ_k in any non-obvious way. It is easy to see, for example, that the norm of the differences between the two correlation matrices will limit to a finite approximation error in the continuous time limit. The discussion is unfortunately somewhat occluded in the context of migration results, given that one needs to recalibrate for each subdivision of the time grid, thereby continuously introducing further assumptions about the new (shorter period) transition matrices. In particular, the threshold computations near $t = 0$ are problematic in the continuous time limit (it is a well known issue of structural models with deterministic thresholds that they produce zero instantaneous, off-diagonal, transition probabilities).

A Calibration of the diffusion model to a migration matrix

Here we both suppress the individual name index i, and to lighten up the notation we denote the default thresholds by $\alpha_k = A_k^{mD}$. Default events are signaled by the first exit times of W_k from suitable levels α_k. The calculation of α_k beyond the first period requires the forward evolution of conditional (on non-default) probability density of the process W_k.

Let $f_k(x)$ denote the probability density that given name has not defaulted until period $k-1$. Then

$$f_k(x) = \int_{\alpha_{k-1}}^{+\infty} du \frac{1}{\sqrt{2\pi\delta_k}} f_{k-1}(u) e^{-\frac{(x-u)^2}{2\delta_k}} \tag{39}$$

(recall that $\delta_k = t_k - t_{k-1}$). The first period density is

$$f_1(x) = \frac{1}{\sqrt{2\pi\delta_1}} e^{-\frac{x^2}{2\delta_1}} \tag{40}$$

The default threshold α_k for each period k is calculated iteratively. An efficient approach is based on

Newton-Raphson. Let α_k^n denote the $n-th$ iterate. The recursion reads:

$$\alpha_k^{n+1} = \alpha_k^n - \frac{G(\alpha_k^n) - p_k}{f_k(\alpha_k^n)} \tag{41}$$

$$G(\alpha_k^n) = \int_{-\infty}^{\alpha_k^n} dx \int_{\alpha_{k-1}}^{+\infty} du \frac{1}{\sqrt{2\pi\delta_k}} f_k(u) e^{-\frac{(x-u)^2}{2\delta_k}} \tag{42}$$

$$f_k(\alpha_k^n) = \int_{\alpha_{k-1}}^{+\infty} du \frac{1}{\sqrt{2\pi\delta_k}} f_k(u) e^{-\frac{(\alpha_k^n-u)^2}{2\delta_k}} \tag{43}$$

The procedure is repeated until successive iterations produces changes smaller than a threshold. The numerical implementation requires a suitable discretization of the one-dimensional integrals appearing above (note that the double integral involves the cumulative normal distribution which is assumed available). It is most convenient to use a grid with variable offset and scale so that the new cut-off level always coincides with a grid point. The probability density $f_k(x)$ satisfies the normalization condition

$$q_k + \int_{\alpha_k}^{\infty} dx f_k(x) = 1 \tag{44}$$

which can be used to construct an error norm for the numerical calculation. The initial condition for the iteration is taken from the previous time period ($\alpha_k^0 = \alpha_{k-1}$). First period threshold is given by

$$\alpha_1 = \sqrt{\delta_1} N^{-1}(q_1) \tag{45}$$

The extension to obtain migration thresholds is straight-forward. We replace in equation above the period-k default probability p_k with the appropriate transition probability T_k.

$$\alpha_k^{n+1} = \alpha_k^n - \frac{G(\alpha_k^n) - T_k}{f_k(\alpha_k^n)} \,. \tag{46}$$

B Numerical Data

From/To	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**D**
1	92.039%	7.090%	0.630%	0.150%	0.060%	0.020%	0.010%	0.001%
2	0.620%	90.840%	7.760%	0.590%	0.060%	0.100%	0.020%	0.010%
3	0.050%	2.090%	91.380%	5.790%	0.440%	0.160%	0.040%	0.050%
4	0.040%	0.210%	4.100%	89.360%	4.820%	0.860%	0.240%	0.370%
5	0.030%	0.080%	1.400%	5.530%	82.250%	8.150%	1.110%	1.450%
6	0.010%	0.040%	0.570%	1.340%	5.390%	81.140%	4.920%	6.590%
7	0.001%	0.020%	0.290%	0.580%	1.550%	10.540%	52.879%	34.140%
D	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	100.000%

Table 3: One year transition matrix used for the computations. This is a stylized matrix originally based on historical data.

From/To	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**D**
1	66.393%	25.014%	6.539%	1.399%	0.345%	0.185%	0.043%	0.082%
2	2.212%	63.443%	27.256%	5.389%	0.820%	0.523%	0.101%	0.256%
3	0.298%	7.405%	66.844%	19.920%	3.287%	1.239%	0.244%	0.764%
4	0.178%	1.409%	14.553%	60.699%	13.895%	4.949%	0.909%	3.406%
5	0.122%	0.599%	6.066%	16.734%	41.982%	19.583%	2.980%	11.934%
6	0.050%	0.281%	2.686%	5.910%	13.165%	39.762%	5.909%	32.238%
7	0.014%	0.116%	1.078%	2.173%	4.261%	12.642%	5.563%	74.154%
D	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	100.000%

Table 4: Five year transition matrix. The matrix is simply the fifth power of the one-year matrix presented in the previous table. The last column of this matrix is used to derive weights for the Weighted Average Rating calculation

Rating (r)	**q[r]**	**W[r]**
1	0.08%	0.04%
2	0.26%	0.11%
3	0.76%	0.34%
4	3.41%	1.53%
5	11.93%	5.36%
6	32.24%	14.47%
7	74.15%	33.28%
D	100.00%	44.88%

Table 5: Weight factors per rating state used in the calculation of the Weighted Average Rating.

References

[1] D.Li. On default correlation: A copula function approach. *Riskmetrics*, 2000.

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

[3] U.Cherubini, E.Luciano, and W.Vecchiato. *Copula methods in finance*. John Wiley & Sons Ltd, 2004.

[4] P.Schonbucher. *Credit derivatives pricing models*. John Wiley & Sons Ltd, 2003.

[5] G.Gupton, C.Finger, and M.Bhatia. Creditmetrics. *JPMorgan Technical Document*, 1997.

[6] C.Zhou. An analysis of default correlations and multiple defaults. *Review of Financial Studies*, 14(2), 2001.

[7] M.Gordy. A comparative anatomy of credit risk models. *Board of Governors of the Federal Reserve System FEDS Paper*, (98-47), 1998.

[8] C.Finger. A comparison of stochastic default rate models. *Riskmetrics*, 2000.

[9] W.Morokoff. Simulation of risk and return profiles for portfolios of cdo tranches. *Proceedings of the 2005 Winter Simulation Conference*, 2005.

[10] K.Thompson, A.McLeod, P.Teklos, and S.Gupta. Time for multi-period capital models. *RISK*, (74-78), 2005.

[11] J.Hull, M.Predescu, and A.White. The valuation of correlation-dependent credit derivatives using a structural model. *Preprint*, 2005.

[12] W.Darsow, B.Nguyen, and W.Olsen. Copulas and markov processes. *Illinois Journal of Mathematics*, 36(4).

[13] O.Vasicek. A series expansion of the bivariate normal integral. *KMV*, 1998.

[14] L.Carty. Moody's rating migration and credit quality correlation, 1920-1996. *Moody's Investors Service*, 1997.

[15] J.Hull and A.White. Valuing credit default swaps II: Modeling default correlations. *The Journal of Derivatives*, Spring 2001.