Research Article

Open Access

Damien Ackerer and Thibault Vatter*

Dependent defaults and losses with factor copula models

<https://doi.org/10.1515/demo-2017-0022>

Received September 13, 2017; accepted December 10, 2017

Abstract: We present a class of flexible and tractable static factor models for the term structure of joint default probabilities, the factor copula models. These high-dimensional models remain parsimonious with pair-copula constructions, and nest many standard models as special cases. The loss distribution of a portfolio of contingent claims can be exactly and efficiently computed when individual losses are discretely supported on a finite grid. Numerical examples study the key features affecting the loss distribution and multi-name credit derivatives prices. An empirical exercise illustrates the flexibility of our approach by fitting credit index tranche prices.

Keywords: credit portfolio, credit derivatives, discrete Fourier transform, factor copula, random loss, survival models

MSC: 60E05, 60E10, 62H05, 62H20, 65T50, 91G20, 91G40, 91G60

1 Introduction

The modeling of dependent events over time, and of random losses, is a common challenging task in insurance, quantitative risk management, financial engineering, and reliability engineering. In this work we suggest a parsimonious and tractable framework to model potentially inhomogeneous and dependent default times in high dimensions. Built upon bivariate and factor copulas, our framework allows us to compute the loss distribution for portfolios of contingent claims whose losses are discretely supported on a finite grid. We numerically illustrate how the risk profiles of such portfolios vary when changing the copula specification, and empirically verify the flexibility of our framework by fitting credit derivative prices.

Factor copulas have been widely used to directly model dependence structure in many different fields. While classical static credit risk models use latent variables conditional on which the default times are independent, such approaches recover the copula “indirectly” by assuming simple functional relationships (e.g., often linear). In this paper, we use copulas to directly link such variables to default probabilities. This more direct approach generalizes virtually all standard copula models in the credit risk literature and offers two important advantages. First, the dependence between default times and the latent factor can be heterogeneous, that is different for each entity. Second, numerous tractable and flexible models can be constructed with parsimony using mixtures and cascades of pair copulas.

The two standard approaches to price credit derivatives with static factor models are the exact but slow recursive method, and the fast but approximate Fourier inversion method. We combine the best of both approaches and show how the loss distribution of complex credit portfolios can be exactly and efficiently computed. We assume that the realized individual losses take values on a finite grid and that, conditional on the latent factor, they are independent from each others and from the default times. We discuss factor dependent Beta-binomial distributions as a flexible way to model individual loss amounts. The aggregate loss then takes

Damien Ackerer: Swissquote Bank, E-mail: damien.ackerer@swissquote.ch

***Corresponding Author: Thibault Vatter:** Department of Statistics, Columbia University, E-mail: tv2233@columbia.edu

values on a finite grid whose size increases linearly with the number of firms. This allows us, for example, to compute the exact payoff distribution of credit portfolio derivatives such as tranches, collateralized debt obligations (CDO) squared, and credit index swaptions.

We numerically explore the performance and the flexibility of our framework. We first show that the discrete Fourier method is significantly faster than the recursive method of [7, 29] especially as the number of latent factors or the loss support size increase. We then study the impact of various dependence assumptions on the loss distribution. Several examples on tranches and CDOs squared illustrate that the loss distribution of repackaged structured products can have dramatically different risk profiles. We also show that the distribution choice of individual firm losses, fixing the average recovery value at default, can critically affect the portfolio loss distribution as well.

The practice of inconsistently pricing complex structured credit derivatives has been blamed for partially causing the subprime financial crisis. We study this challenging exercise within our framework. More precisely, we calibrate several models to market tranche prices from the North America investment grade credit index series 21. Some of the models include stochastic loss amounts which have been fitted on the historical realized recovery rates. We suggest a mixture of two copulas as a flexible specification mimicking two regimes. In a static analysis, we find that the mixture outperforms the other models, as it is the only one reproducing the prices of both the junior and senior tranches. Calibrating this model for all days in our sample, we further find that the parameters are stable over time. Furthermore, one of the correlations being almost always equal to one, we repeat the exercise by fixing it to 0.999. Interestingly, we find similar results, therefore achieving an almost perfect calibration to all tranches using only two parameters (i.e., the other correlation and the weight).

In summary, the contributions of this paper are to

- construct high-dimensional and inhomogeneous credit risk models in a parsimonious and tractable fashion using factor copulas,
- show that any static credit risk model equipped with the conditional independence property can be equivalently rewritten as a factor copula model,
- model random and factor dependent individual losses on a finite grid,
- compute efficiently the exact loss distribution of complex portfolios, such as portfolio of repackaged tranches, potentially conditional on the realization of some defaults,
- numerically investigate how portfolio loss distributions are affected by some dependence and loss assumptions, and
- empirically illustrate that CDX tranches can be consistently priced with a simple model.

The proposed framework is well suited to the consistent pricing and risk-management of insurance claims and financial credit derivatives. It could for example be used to efficiently stress test large portfolios under various scenarios, such as realized firm defaults and/or changes in the dependence structure of default times.

We now review some of the related literature. Our approach builds on recent advances on the high-dimensional modeling of random variables. When dealing with multivariate data, copulas are attractive, allowing to model separately the marginal distributions and the dependence structure. Unfortunately, few copulas remain practically useful in high-dimensional settings, because common parametric families are often either too flexible, or not enough. An example of the former is the elliptical family, whose members have a number of parameters that grows quadratically with the dimension. Conversely, members of the Archimedean family have a small and fixed number of parameters, independently of the dimension. Recently, high-dimensional copulas using a factor structure have been constructed independently by [45, 46] and [35, 36]. Such approaches alleviate the curse of dimensionality by considering a smaller set of latent variables, conditional upon which the random variables of interest are assumed independent. Arguably the main difference between the methods presented in [45, 46] and [35, 36] is that copulas proposed in the former can only be simulated, whereas those in the latter admit closed form expressions. In fact, it can be shown the factor copulas from [35, 36] are a special case of pair-copula constructions (PCCs). One of the hot topics

of multivariate analysis over the last couple of years, PCCs are flexible representations of the dependence structure underlying a multivariate distribution. Introduced by [8, 9] and popularized by [1], PCCs are decompositions of a joint distribution by considering pairs of conditional random variables. For a given joint distribution, such a construction is not unique, but all possible decompositions can be organized as graphical structures, the so-called vines. Assuming the copula linking default times as in [35, 36], an interesting aspect of our approach is that it nests the standard models described for instance in [13, 28, 38, 51] as special cases.

Fourier transform techniques have been considered in [24, 37] to approximate the loss distribution, and exact recursive algorithms have been derived in [7, 29] with finite and discrete support. To the best of our knowledge, our work is the first to combine discrete Fourier transform and losses with finite and discrete support, allowing us to study structured product such as CDO squared without simulations, unlike [25, 30]. The Fourier techniques have also been applied to price stock options and insurance claims for many years, see [15, 18, 19]. Alternatively, in some special cases such as large homogeneous portfolios, explicit expressions have been derived to approximate the portfolio loss distribution, see [49, 51].

The accurate modeling of dependent defaults is particularly important for credit derivatives pricing, we refer to [16] for a description of standard structured credit products. In particular, the calibration of tranches on credit portfolios is a daunting task, which is often solved in an ad-hoc way, that is by considering a specific model for each tranche. Significant effort have been made to develop consistent models, see [23] for a comparison between top down and bottom up approaches and [32] for a comparison of different copulas in the CreditRisk+ framework. Standard copula models had however limited empirical success and other frameworks have been developed, see [3, 11, 14, 17, 21, 22, 27, 31, 33, 40]. In this paper, we develop bottom-up models that are both simple to calibrate and successful at reproducing all the tranche spreads. Furthermore, while the valuation of CDO squared has been considered with simulations in [25, 30], this work is the first to derive explicitly the loss distribution of a CDO squared in a factor copula framework. We refer to [12] for a technical analysis of valuation methods for structured credit products. Note that the same defaults dependence as in our static framework may be obtained for some homogeneous portfolio in a dynamic framework where the default times are driven by a stochastic process, see the survey by [41] on dynamic models.

The realized loss at default on corporate loans and bonds is known to be stochastic, volatile, and negatively correlated with the business cycle. The recovery rates volatility and correlation with default risk is studied, for example, in [4]. The valuation of credit derivatives with random losses has also been investigated in [5, 6, 34].

The remainder of the paper is structured as follows. Section 2 presents the factor copula framework. Section 3 describes the construction of the individual loss amounts and the exact computation of the loss distributions. Sections 4 and 5 contain respectively the numerical analysis and an application to financial market data. Section 6 concludes and Appendix A contains the proofs.

2 The factor copula framework

We first recall how dependent default times can be constructed using a copula when the marginal default probabilities are deterministic. We then combine factor copulas and bivariate copulas to construct high-dimensional models in which the dependence between the default times and the latent factors can easily be made inhomogeneous. We finally show that standard factor models in the credit risk literature can always be rewritten, sometimes in simpler terms, as factor copula models.

2.1 Default Times Construction

We consider N entities. For each $j = 1, \dots, N$, let $p_{j,t}$ be a non-decreasing deterministic function for all $0 < t < \infty$ with $p_{j,0} = 0$ and $\lim_{t \rightarrow \infty} p_{j,t} = 1$. We define the default time τ_j of entity j as follows

$$\tau_j := \inf\{t \geq 0 : U_j \leq p_{j,t}\},$$

where U_j is a uniform random variable on the unit interval. This is a standard construction of default times, see [10, 42]. The function $p_{j,t}$ is equivalent to the marginal default probability of entity j before any default is observed

$$\mathbb{P}[\tau_j \leq t] = \mathbb{P}[U_j \leq p_{j,t}] = p_{j,t}. \quad (1)$$

When $p_{j,t}$ is absolutely continuous with respect to time then it is given by $p_{j,t} = 1 - e^{-\int_0^t \lambda_{j,s} ds}$ for some non-negative default intensity function $\lambda_{j,s}$. Note that, in this setup, the random vector $U = (U_1, \dots, U_N)$ is the only stochastic object whose probability distribution is given by a copula C_U . In other words, if for any vector $(u_1, \dots, u_N) \in [0, 1]^N$ the random vector $U \in [0, 1]^N$ is such that $\mathbb{P}[U_j \leq u_j] = u_j$ for each j , then its joint distribution is called a copula and we write

$$C_U(u_1, \dots, u_N) = \mathbb{P}[U_1 \leq u_1, \dots, U_N \leq u_N]. \quad (2)$$

The following well-known lemma shows that for $(t_1, \dots, t_N) \in \mathbb{R}_+^N$ there exists a simple expression linking joint to marginal default probabilities using the copula C_U of U .

Lemma 2.1. *The joint default probability is given by*

$$\mathbb{P}[\tau_1 \leq t_1, \dots, \tau_N \leq t_N] = C_U(p_{1,t_1}, \dots, p_{N,t_N}). \quad (3)$$

A direct construction of high-dimensional copulas amounts at trading-off model complexity and tractability. This is somewhat problematic, because the usual parametric families contain either too many (e.g., in the case of implicit copulas extracted from known multivariate distributions), or too few (e.g., in the case of Archimedean copulas built using a continuous and nonincreasing N -monotone generator) parameters. Furthermore, as we will show in Section 3 when pricing complex financial derivatives, the notion of conditional independence (on a set of latent factors) allows us to obtain a flexible yet tractable class of models. Hereinafter we therefore focus on the so-called factor copulas of [35, 36].

2.2 One-factor copulas

A one-factor copula model is constructed by assuming that there exists a latent factor V uniformly distributed on the unit interval such that, conditional on the realization of V , the coordinates of the random vector U are independent. In other words, we have

$$\mathbb{P}[U_1 \leq u_1, \dots, U_N \leq u_N \mid V = v] = \prod_{j=1}^N \mathbb{P}[U_j \leq u_j \mid V = v] \quad (4)$$

for any vector $(u_1, \dots, u_N) \in [0, 1]^N$ and $v \in [0, 1]$. Note that we directly consider a uniformly distributed factor. This is without loss of generality using the probability integral transform for factors having alternative continuous distributions. While factor models in the credit risk literature usually build upon real-valued random variables, this creates an unnecessary layer of complexity layer when computing expressions such as the joint default probability (3).

The following proposition shows that our assumption about V yields a simple decomposition in terms of bivariate copulas for C_U , also known as one-factor copula.

Proposition 2.2 (One-factor copula [35, 36]). For $j = 1, \dots, N$, let $C_{U_j, V}$ denote the joint distribution of U_j and V , that is $\mathbb{P}[U_j \leq u_j, V \leq v] = C_{U_j, V}(u_j, v)$. If the coordinates of U are independent conditionally on V , then

$$C_U(u_1, \dots, u_N) = \int_{[0,1]} \prod_{j=1}^N C_{U_j|V}(u_j | v) dv, \quad (5)$$

where, for all $j = 1, \dots, N$, $C_{U_j|V}(u_j | v) = \frac{\partial C_{U_j, V}(u_j, v)}{\partial v}$ are the so-called h -functions.

The h -functions have been introduced by [1] while studying the pair-copula decomposition of a general multivariate distribution: if $C_{U_j, V}(u_j, v) = \mathbb{P}[U_j \leq u_j, V \leq v]$, then $C_{U_j|V}(u_j | v) = \mathbb{P}[U_j \leq u_j | V = v]$.

Note that $C_{U_j, V}(u, v) = uv$ implies $C_U(u_1, \dots, u_N) = \prod_{j=1}^N u_j$. In other words, if U_j is independent from V , then it is also independent from U_k for all $k \neq j$, which means that the coordinates of U depend on each other only through the factor V .

Example 2.3. The Gaussian model of [38] is a one-factor copula model obtained by using for all j the bivariate copula $C_{U_j, V}(u_j, v; \rho) = \Phi_2(\Phi^{-1}(u_j), \Phi^{-1}(v); \rho)$ which implies $C_{U_j|V}(u_j | v; \rho) = \Phi\left(\frac{\Phi^{-1}(u_j) - \rho\Phi^{-1}(v)}{1-\rho^2}\right)$ and $C_U(u_1, \dots, u_N; \rho) = \int_0^1 \prod_{j=1}^N \Phi\left(\frac{\Phi^{-1}(u_j) - \rho\Phi^{-1}(v)}{1-\rho^2}\right) dv$ where $\Phi(\cdot)$ is the standard normal distribution and $\Phi_2(\cdot, \cdot; \rho)$ is the bivariate normal distribution with correlation ρ .

The specification in Proposition 2.2 is more flexible than standard factor models from the credit risk literature. The reason is that it is straightforward to build non-homogeneous models with firm-specific dependence between the default time and the common factor. For instance, one could build a model using a different bivariate copula for each conditional probability $\mathbb{P}[U_j \leq u_j | V = v]$. Furthermore, there exists many well-studied parametric families for bivariate copulas, see [48].

Beyond such parametric families, a simple way to increase the modeling flexibility while preserving analytical tractability is to combine different bivariate copulas. Hence, we suggest mixture distributions as a natural and simple extension, and which considerably enrich the class of one-factor copulas models.

Definition 2.4. Let K be a positive integer, $C_{U_j, V}$ is a mixed bivariate copula if there exists K copulas $C_{U_j, V}^k$, K positive weights $w_k > 0$ such that $\sum_{k=1}^K w_k = 1$, and

$$C_{U_j, V}(u_j, v) = \sum_{k=1}^K w_k C_{U_j, V}^k(u_j, v). \quad (6)$$

One way to interpret this expression is Bayesian, namely assuming that the dependence between the random variable U_j and the factor V is uncertain and follows the distribution $C_{U_j, V}^k$ with probability w_k . The corresponding h -function still has a simple expression, as we have $C_{U_j|V}(u_j | v) = \sum_{k=1}^K w_k C_{U_j|V}^k(u_j | v)$. While a mixture of Gaussian copulas was studied in [39], the definition above can accommodate a different parametric family for each mixture component.

The loss distribution of a credit portfolio conditional on the default time of a specific entity is of particular interest for risk management applications. For instance, it is a necessary input to compute a Credit Valuation Adjustment (CVA), namely the expected loss on a bilateral position resulting from the default risk of this entity (see [52] for an introduction). As it turns out, the joint distribution of default times conditional on a subset of realized default times is obtained as a simple modification of Equation (5).

Let $\mathcal{J} = \{1, \dots, N\}$ and $\mathcal{D} \subset \mathcal{J}$ denote respectively the entire set and a subset of entities. The following proposition shows that the joint default distribution conditional on the defaults of all the entities in $\mathcal{D}$ also has a simple representation.

Proposition 2.5. In a one-factor copula model, the joint default distribution conditional on $\tau_k = t_k$ for $k \in \mathcal{D}$ is

$$\mathbb{P}[\tau_1 \leq t_1, \dots, \tau_N \leq t_N \mid \tau_k = t_k : k \in \mathcal{D}] = \frac{\int_{[0,1]} \prod_{j \in \mathcal{I} \setminus \mathcal{D}} C_{U_j|V}(p_{j,t_j} \mid v) \prod_{k \in \mathcal{D}} c_{U_k,V}(p_{k,t_k}, v) dv}{\int_{[0,1]} \prod_{k \in \mathcal{D}} c_{U_k,V}(p_{k,t_k}, v) dv} \quad (7)$$

where $c_{U_j,V}(u, v) = \frac{\partial^2 C_{U_j,V}(u, v)}{\partial u \partial v}$ is the density of the bivariate copula $C_{U_j,V}$.

Although the default times are correlated, conditioning on a subset of defaulted entities does not significantly complexify the expression for the joint distribution of the surviving entities. The denominator on the right hand side in (7) is the copula density of the defaulted entities evaluated at the default times.

2.3 Multi-factor copulas

We generalize the framework of Section 2.2 by considering a d -dimensional random vector of latent factors $V = (V_1, \dots, V_d)$. We assume that V takes values on the hypercube $[0, 1]^d$ and has uniform marginal distributions. The joint distribution of V is by definition a copula that we denote C_V . The following proposition shows that the one-factor framework extends to a multi-factor one.

Proposition 2.6 (Multi-factor copula). *For $j = 1, \dots, N$, let $C_{U_j,V}$ denote the joint distribution of U_j and V , that is $\mathbb{P}[U_j \leq u_j, V \leq v] = C_{U_j,V}(u_j, v)$. If the coordinates of U are independent conditionally on V , then*

$$C_U(u) = \int_{[0,1]^d} \prod_{j=1}^N C_{U_j|V}(u_j \mid v) dC_V(v) \quad (8)$$

where, for all $j = 1, \dots, N$, $C_{U_j|V}(u_j \mid v) = \frac{\partial^d C_{U_j,V}(u_j, v)}{\partial v_1 \dots \partial v_d}$.

Note that the representation in (8) is not equivalent to the one found in [35, 36] since the coordinates of V are not necessarily independent.

The representation (8) is arguably more complicated than (5) despite their apparent similarity. The reason is that, instead of being bivariate, each $C_{U_j,V}$ has dimension $d + 1$. However, the multi-factor framework simplifies under the assumption of independent latent factors V as shown in the following proposition. We denote the function composition with the symbol $\circ$, that is $f \circ g(x) = f(g(x))$ for any real valued functions f and g .

Corollary 2.7 (Copulas with independent factors [35, 36]). *If $C_V(v) = \prod_{j=1}^d v_j$, then*

$$C_U(u_1, \dots, u_n) = \int_{[0,1]^d} \prod_{j=1}^N C_{U_j|V_1}(\cdot \mid v_1) \circ \dots \circ C_{U_j|V_d}(u_j \mid v_d) dv, \quad (9)$$

where C_{U_j,V_k} is a bivariate copula for $j \in 1, \dots, d$ and $k = 1, \dots, d$.

Note that the recursive decomposition (9) is a particular case of pair-copula constructions [35, 36]. This construction is interesting for several reasons. First, it is a parsimonious way to model a complex multivariate dependencies. Second, the hierarchical structure, which can be represented as a graphical model, has an intuitive interpretation. Third, because the integrand in (9) is a simple recursion, it can be vectorized in a computationally efficient manner.

Finally, it should be noted that the number of latent factors is also the dimension of the hypercube on which the product of conditional copulas has to be integrated to compute the joint default probability. One should therefore balance model flexibility and computational cost.

2.4 Comparison with standard factor models

We show that standard static factor models can be rewritten explicitly as factor copula models. In this context, one usually considers a random vector $Y = (Y_1, \dots, Y_N) \in \mathbb{R}^N$ along with a deterministic and component-wise non-decreasing vector $y_t = (y_{1,t}, \dots, y_{N,t}) \in \mathbb{R}^N$. For instance, Y and y_t (or their exponentials) can represent the firm values and corresponding default barriers. The default time τ_j of firm j is then defined as the first time its value is below its default barrier, that is $\tau_j = \inf\{t \geq 0 : Y_j \leq y_{j,t}\}$. Additionally, standard factor models are constructed by decomposing the stochastic behavior of the firm value into a systemic and an idiosyncratic component. In other words, one assumes the existence of a random vector $X \in \mathbb{R}^d$ and N variables ϵ_j for $j \in \{1, \dots, N\}$, such that Y_j is a function X and ϵ_j , that is $Y_j = f_j(X, \epsilon_j)$ for some $(d+1)$ -dimensional function f_j taking values on $\mathbb{R}_+$.

Let F_{Y_j} , F_X , respectively $F_{Y_j}^{-1}$, F_X^{-1} , denote the distributions of Y and X , respectively their inverse, and $F_{Y_j|X}$ denote the conditional distribution of Y_j given X . The following proposition shows that any standard factor model is equivalent to a specific factor copula model.

Theorem 2.8. *A standard factor model is a factor copula model with marginal default probabilities $p_{j,t} = F_{Y_j}(y_{j,t})$ and conditional copulas*

$$C_{U_j|V}(u | v) = F_{Y_j|X}(F_{Y_j}^{-1}(u) | (F_{X_1}^{-1}(v_1), \dots, F_{X_N}^{-1}(v_N))),$$

for $j = 1, \dots, N$, and where the copula of V is given by

$$C_V(v) = F_X(F_{X_1}^{-1}(v_1), \dots, F_{X_N}^{-1}(v_N)).$$

Furthermore, if the functions F_X and F_{Y_j} for all $j = 1, \dots, N$ are continuous, then the copulas C_V and $C_{U_j|V}$ for all $j = 1, \dots, N$ are unique.

While $C_{U_j|V}$ and C_V sometimes admit closed-form expressions, it is clear that the marginal distributions are irrelevant. Instead, working directly with copulas offers more modeling flexibility while ensuring tractability.

Example 2.9. *The Gaussian model described in Example 2.3 is obtained by writing, for $j \in \{1, \dots, N\}$, $Y_j = \rho X + \sqrt{1 - \rho^2} Z_j$ and $y_{j,t} = \Phi(p_{j,t})$ where $X, Z_1, \dots, Z_N$ are i.i.d. $N(0, 1)$ random variables.*

The d -factors Gaussian model as well as the Student t and commonly encountered Archimedean copulas lead to similarly simple explicit expressions.

3 Discrete loss distributions

We first define a specific finite grid on which losses will take values. Then, we show that portfolio loss distributions can be computed in quasi-closed form using discrete Fourier inversion, and that this methodology also applies to the underlying portfolio of complex credit derivatives such as CDO squared. The section concludes by presenting a flexible approach to model random and factor dependent loss amounts.

3.1 Loss Amounts Specification

We define the time- t loss L_t on a portfolio composed of securities written on N different entities,

$$L_t = \sum_{j=1}^N \ell_j \mathbb{1}_{\{\tau_j \leq t\}} = \sum_{j=1}^N \ell_j \mathbb{1}_{\{U_j \leq p_{j,t}\}}, \quad (10)$$

where ℓ_j is the possibly random loss amount experienced when entity j defaults, and $\mathbb{1}_{\{\tau_j \leq t\}}$ is the default indicator of entity j . In this section, we make two assumptions on ℓ_j to preserve the tractability of the portfolio loss distribution, and to enable efficient numerical techniques.

First, we assume that ℓ_j is V -conditionally independent of both ℓ_k for $k \neq j$ and U (or equivalently τ), that is

$$\mathbb{P}[U \leq u, \ell \leq x \mid V = v] = \prod_{j=1}^N C_{U_j|V}(u_j \mid V = v) \mathbb{P}[\ell_j \leq l_j \mid V = v],$$

with $\ell = (\ell_1, \dots, \ell_N)$, and for any $u \in [0, 1]^N$, $v \in [0, 1]^d$ and $l \in \mathbb{R}_+^N$. As for the joint distribution of default times, V -conditional probabilities can be arbitrarily specified. Hence, this assumption does not preclude dependent default times and losses given default. Note that, in the literature, loss amounts are commonly assumed to be independent from each others and from the default times, and often set to be constant. Our framework goes beyond such limitations, which can be of practical importance as shown in Section 4

Second, we assume that losses have a discrete distribution with finite support. More specifically, we let $\delta \in \mathbb{R}_+$ be the common loss unit, such that each ℓ_j has a discrete and finite support starting at zero and with mesh δ , that is

$$\ell_j \in \{0, \delta, 2\delta, \dots, m_j\delta\}, \quad j = 1, \dots, N$$

for some integer $m_j \in \mathbb{N}$. Hence, the portfolio loss distribution also has a discrete support with the same mesh δ , that is

$$L_t \in \{0, \delta, 2\delta, \dots, M\delta\}$$

where $M = \sum_{i=1}^N m_i$. Although δ is an arbitrary constant, it can be as fine as required in order to mimic the discreteness of real-world prices. For instance, assuming that the granularity of prices is in cents (i.e., $\delta = 0.01\$$) and that the notional of each contract is 1\$, then $m_j = 100$ and $M = N \times 100$. This setup can also be understood as a specific discretization of continuously distributed random losses. The use of a common loss unit can be traced back at least to [7, 29].

3.2 Portfolio loss distribution

We show that the portfolio loss distribution has an almost closed-form expression that can be efficiently computed numerically. Recall that, for a discrete and finitely supported random variable $X \in \{0, 1, \dots, M\}$ admitting a characteristic function $\phi_X(u) = \mathbb{E}[e^{iuX}]$, its distribution can be represented as a finite sum

$$\mathbb{P}[X = k] = \frac{1}{M+1} \sum_{m=0}^M \phi_X\left(\frac{2\pi m}{M+1}\right) e^{-\frac{2\pi i k m}{M+1}},$$

where i denotes the imaginary unit. Therefore, if the characteristic function of the loss distribution admits a closed-form expression, so does the loss distribution itself. Using the V -conditional independence, the following proposition shows that the characteristic function of the loss admits a simple expression. To improve the clarity of the formulas, we work with the normalized losses $\ell_j\delta^{-1} \in \{0, 1, \dots, m_j\}$ and normalized portfolio loss $L_t\delta^{-1} \in \{0, 1, \dots, M\}$.

Proposition 3.1. *The characteristic function of the normalized portfolio loss $L_t\delta^{-1}$ is given by*

$$\phi_{L_t}(u) = \mathbb{E}[e^{iuL_t\delta^{-1}}] = \int_{[0,1]^d} \prod_{j=1}^N (1 - p_{j,t}(v) + p_{j,t}(v)\phi_{\ell_j}(u, v)) dC_V(v), \quad (11)$$

for any time $t \geq 0$ and for $u \in \mathbb{R}$, where $p_{j,t}(v) = C_{U_j|V}(p_{j,t} \mid v)$ is the conditional default probability of j , $p_{j,t}$ is the unconditional default probability of j in (1), and

$$\phi_{\ell_j}(u, v) = \sum_{k=0}^{n_j} \mathbb{P}[\ell_j = \delta k \mid V = v] e^{iuk} \quad (12)$$

the V -conditional characteristic function of $\ell_j\delta^{-1}$.

The characteristic function is therefore explicit, up to the integral over the compact set $[0, 1]^d$ which can be efficiently computed using, for example, Legendre quadrature. Denoting by $\{w_i, x_i\}_{i=1}^n$ n pairs of quadrature weights and nodes to approximate an integral over $[0, 1]$, they can be combined in a straightforward way to perform multi-dimensional integration by using a product rule. In other words, Equation (11) can be approximated by

$$\phi_{L_t}(u) \approx \sum_{i_1, \dots, i_d} \prod_{l=1}^d w_{i_l} \prod_{j=1}^N (1 - p_{j,t}(v_{i_1, \dots, i_d}) + p_{j,t}(v_{i_1, \dots, i_d}) \phi_{\ell_j}(u, v_{i_1, \dots, i_d})) dC_V(v_{i_1, \dots, i_d}),$$

where $v_{i_1, \dots, i_d} = (x_{i_1}, \dots, x_{i_d})$. Even though the number of quadrature nodes and weights can be chosen small, such an approach works well only with a small number of factors, as the total number of grid points is n^d . If d is larger than 4 or 5, sparse grids [see e.g., 26] can be used to reduce the number of grid points required.

The following Lemma is a reminder that, since the support of L_t is discrete and finite, its distribution is equal to the discrete Fourier transform (DFT) of its characteristic function.

Lemma 3.2. *The probability distribution of the portfolio loss is given by*

$$\mathbb{P}[L_t = k\delta] = \frac{1}{M+1} \sum_{m=0}^M \phi_{L_t}(\mu m) e^{-i\mu k m} \text{ for } k \in \{0, \dots, M\}, \quad (13)$$

with $\mu = 2\pi/(M+1)$ and $\phi_{L_t}(\cdot)$ is the characteristic function of $L_t \delta^{-1}$.

Calculating directly $\mathbb{P}[L_t = k\delta]$ is in general a combinatorial problem whose complexity is increasing exponentially fast with M . However the complexity of computing the DFT in Lemma 3.2 is $\mathcal{O}(M^2)$ which is of significant practical importance as long as evaluating the characteristic function is also fast. Note that the Fast Fourier Transform (FFT) algorithm, whose complexity is only $\mathcal{O}(M \log(M))$, can also be used.

The assumption of loss unit and discretely supported portfolio losses appears already in [6, 7, 29], where the distribution is computed without approximation by a recursive algorithm. However, as will be shown in Section 4.1, the computational cost of this recursion increases much faster with both the support size and the number of factors than that of our approach. Note also that the discrete Fourier inversion in Lemma 3.2 differs from the continuous Fourier inversion described in [14, 37] which aims to approximate a continuous loss distribution. Since our approach provides quasi-closed expressions for the loss distribution, its scope is much wider.

Whereas tranches and CDO squared can be priced using the portfolio loss distribution only (see Section 3.3), other products necessitate the joint distribution of the total number of defaulted entities and of the total loss. This is the case for credit index swaptions, namely options on indices paying realized losses in exchange for premium payments proportional to the number of non-defaulted entities, the market for which is currently booming. We now derive this joint distribution. Let the number of defaulted entities at time t be

$$N_t = \sum_{j=1}^N \mathbb{1}_{\{\tau_j \leq t\}}. \quad (14)$$

The following Proposition gives a generic expression for the joint distribution of (N_t, L_t) .

Proposition 3.3. *The joint distribution of (N_t, L_t) is given by*

$$\mathbb{P}[N_t = n, L_t = \delta k] = \sum_{j=0}^N \sum_{l=0}^M \frac{\phi_{N_t, L_t}(\mu j, \nu l) e^{-i\mu n j} e^{-i\nu k l}}{(1+N)(1+M)} \quad (15)$$

with $\mu = 2\pi/(M+1)$, $\nu = 2\pi/(N+1)$, and

$$\phi_{N_t, L_t}(x, y) = \int_{[0,1]^d} \prod_{j=1}^N (1 - p_{j,t}(v) + p_{j,t}(v) \phi(x, y, v)) dC_V(v)$$

where $p_{j,t}(v)$ is as in Proposition 3.1, and $\phi(x, y, v) = \sum_{k=0}^{n_j} \mathbb{P}[\ell_j = \delta k \mid V = v] e^{i(x+yk)}$.

The computation of (15) thus boils down a two-dimensional discrete Fourier transform inversion. Note that $\phi(x, y, v)$ is the V -conditional characteristic function of $x + y\ell_j\delta^{-1}$ evaluated at one, that is $\phi(x, y, v) = \mathbb{E}[\exp(i(x + y\ell_j\delta^{-1})) \mid V = v]$ and $\phi_{N_t, L_t}(x, y)$ is the characteristic function of $(N_t, L_t\delta^{-1})$ evaluated at (x, y) , that is $\phi_{N_t, L_t}(x, y) = \mathbb{E}[\exp(i(xN_t + yL_t\delta^{-1}))]$.

When the loss amounts ℓ_j are homogeneous and independent from V , then the following more direct calculation can be applied

$$\mathbb{P}[N_t = n, L_t = \delta k] = \mathbb{P}[L_t = k\delta \mid N_t = n] \mathbb{P}[N_t = n]$$

where $\mathbb{P}[N_t = n]$ can be computed as in Lemma 3.2, and $\mathbb{P}[L_t = k\delta \mid N_t = n] = \mathbb{P}[\sum_{j=1}^n \ell_j = k\delta]$ can also be derived using the discrete Fourier transform.

Similar expressions can be derived for the default probabilities when the default intensities are driven by a stochastic process, see for example [50]. However these expressions involve expectations of copula functions in random marginal default probabilities whose computations generally require costly numerical techniques. Yet, combining the linear credit risk models presented in [2] and polynomial factor copulas would result in tractable polynomial models with dependent default times and stochastic default intensities. In that case, the joint default probability rewrites as an integral over the expectation of a polynomial in a polynomial diffusion which is an analytical expression, see [20]. Some examples of polynomial copulas, such as the Farlie-Gumbel-Morgenstern copula, can be found in [44] and the Bernstein copulas, which can approximate any copula, are studied in [47].

3.3 Tranches and CDO squared

We describe how to compute loss distributions for two structured credit products: tranches and portfolios of tranches, also known as CDO squared. To the best of our knowledge, this is the first time that a methodology to retrieve the exact loss distribution of a CDO squared is presented for an inhomogeneous bottom-up model.

A tranche on a credit portfolio is a derivative that pays a fraction of the realized portfolio losses above the attachment point a and below the detachment point b with $0 \leq a < b$, in exchange of regular payments on the remaining size of the tranche. Define the tranche loss as

$$\mathcal{T}_t^{a,b} := \min \{ \max \{ L_t - a, 0 \}, b - a \}, \quad (16)$$

and denote $\epsilon_a := \delta - (a \bmod \delta)$. The following expresses the relationship between the probability distribution of L_t and that of $\mathcal{T}_t^{a,b}$.

Proposition 3.4. *The tranche loss $\mathcal{T}_t^{a,b}$ probability mass function is given by*

$$\begin{aligned} \mathbb{P}[\mathcal{T}_t^{a,b} = \epsilon_a + k\delta] &= \mathbb{P}[L_t = (k + \lceil a/\delta \rceil)\delta], \quad \text{for any } k \in \mathbb{N} \text{ such that } 0 < \epsilon_a + k\delta < b - a, \\ \mathbb{P}[\mathcal{T}_t^{a,b} = 0] &= \sum_{m=0}^{\lfloor a/\delta \rfloor} \mathbb{P}[L_t = m\delta], \quad \text{and} \quad \mathbb{P}[\mathcal{T}_t^{a,b} = b - a] = \sum_{m=\lceil b/\delta \rceil}^M \mathbb{P}[L_t = m\delta], \end{aligned}$$

where $\lfloor x \rfloor$ (respectively $\lceil x \rceil$) denotes the closest integer smaller (respectively larger) than x .

The CDO squared loss distribution can also be computed explicitly, even when the default times of entities composing the different portfolios are dependent. Let us consider K tranches on portfolios written on different entities. For each portfolio k we denote by $\mathcal{T}_{k,t}^{a_k,b_k}$ and $L_{k,t}$ the k -th tranche and portfolio loss, with a_k and b_k the k -th tranche attachment and detachment points.

Denote $\mathcal{L}_t = \sum_{k=1}^K \mathcal{T}_{k,t}^{a_k,b_k}$ the CDO-squared loss, at time t . Assume that for all k , we have

$$a_k \bmod \delta = 0 \quad \text{and} \quad b_k \bmod \delta = 0. \quad (17)$$

Then, each of the tranche as well as the CDO squared have discrete and finite supports, that is $\mathcal{T}_{kt}^{a_k, b_k} \in \{0, \delta, 2\delta, \dots, b_k - a_k\}$ for each k and $\mathcal{L}_t \in \{0, \delta, 2\delta, \dots, M_K \delta\}$ where $M_K = \sum_{k=1}^K (b_k - a_k)/\delta$.

Corollary 3.5. *If (17) holds for $k = 1, \dots, K$, then the characteristic function of the squared loss is*

$$\phi_{\mathcal{L}_t}(u) = \int_{\mathbb{R}^d} \prod_{k=1}^K \phi_{\mathcal{T}_{kt}}(u, v) dC_V(v),$$

where

$$\phi_{\mathcal{T}_{kt}}(u, v) = \sum_{n=1}^{(b_k - a_k)/\delta} \mathbb{P} \left[\mathcal{T}_{kt}^{a_k, b_k} = n\delta \mid V = v \right] e^{iun},$$

is the V -conditional characteristic function of $\mathcal{T}_{kt}\delta^{-1}$.

To compute $\phi_{\mathcal{T}_{kt}}$ one may use Proposition 3.4 applied to the V -conditional portfolio loss distribution, namely $\mathbb{P}[L_{kt} = m\delta \mid V = v]$. Applying Lemma 3.2 with $\mathcal{L}_t$ replacing L_t , one finally obtains the distribution of the squared loss. With the distribution of the squared loss, one can then price derivatives such as tranches on a portfolio of tranches.

3.4 Beta-binomial loss amounts

We describe how to use the Beta-binomial in order to model dependent loss amounts. More specifically, we allow the distribution of the loss amounts for each entity to be dependent on the default times and others loss amounts. For each $j = 1, \dots, N$, we let the loss amount ℓ_j take value in a set of the form

$$\ell_j \in \{b_j\delta, (a_j + b_j)\delta, \dots, (n_j a_j + b_j)\delta\} \subset \{0, \delta, 2\delta, \dots, m_j\delta\}$$

with the integers $a_j, b_j, n_j \in \mathbb{N}$ such that $n_j a_j + b_j = m_j > 0$. Note that the two sets are equivalent when $a_j = 1$ and $b_j = 0$. The Beta-binomial model is obtained by assuming that the V -conditional distribution of the loss amount increment $(\ell_j \delta^{-1} - b_j)/a_j$ is a Beta-binomial random variable. We recall that the Beta function is defined by $B(\alpha, \beta) = \Gamma(\alpha)\Gamma(\beta)/\Gamma(\alpha + \beta)$ with the Gamma function $\Gamma(z) = \int_0^\infty x^{z-1} e^{-x} dx$.

Definition 3.6 (The Beta-binomial loss model). *The V -conditional probability of loss is*

$$\mathbb{P}[\ell_j = (a_j k + b_j)\delta \mid V = v] = \int_{[0,1]} \mathbb{P}[Z = k \mid p, n_j] \pi_j(p \mid V = v) dp$$

for any $k = 0, \dots, n_j$, where $\mathbb{P}[Z = k \mid p, n_j] = \binom{n_j}{k} p^k (1-p)^{n_j-k}$, and with the Beta distribution $\pi_j(p \mid V = v) = p^{\alpha(v)-1} (1-p)^{\beta(v)-1} / B(\alpha(v), \beta(v))$ for some functions $\alpha : [0, 1]^d \rightarrow \mathbb{R}_{+*}$ and $\beta : [0, 1]^d \rightarrow \mathbb{R}_{+*}$.

Conditional on V the number of loss units experienced upon default is the sum of a constant b_j and of k units a_j where k follows a Binomial distribution with parameter p and support $0, \dots, n_j$. In addition, the probability p is random and distributed according to a Beta distribution with parameters $\alpha(v)$ and $\beta(v)$. Note that the functions α and β may be entity specific.

The Beta-binomial is a flexible distribution that nests a large spectrum of distributions such as the Bernoulli (see below), the discrete uniform (when $\alpha = \beta = 1$), and asymptotically the binomial (for large α and β). Its probability mass function is given by

$$\mathbb{P}[\ell_j = (a_j k + b_j)\delta \mid V = v] = \binom{n_j}{k} \frac{B(\alpha(v) + k, \beta(v) + n_j - k)}{B(\alpha(v), \beta(v))}.$$

for any $k = 0, \dots, n_j$ and where Γ denotes the gamma function. The V -conditional loss amount mean and variance therefore also have an explicit expression $\mathbb{E}[\ell_j \mid V = v] = (a_j \frac{n_j \alpha(v)}{\alpha(v) + \beta(v)} + b_j)\delta$ and $\text{Var}[\ell_j \mid V = v] =$

$\frac{n_j \alpha(v) \beta(v) (\alpha(v) + \beta(v) + n_j)}{(\alpha(v) + \beta(v))^2 (\alpha(v) + \beta(v) + 1)} \alpha_j^2 \delta^2$. Remark that the mean loss amount is positively correlated with V when the function $v \mapsto \alpha(v)/(\alpha(v) + \beta(v))$ is increasing on $[0, 1]$.

Example 3.7 (Bernoulli model). *The loss amount distribution reduces to a Bernoulli when $n_j = 1$ with probability $p(v) = \frac{\Gamma(\alpha(v) + \beta(v))}{\Gamma(\alpha(v)) \Gamma(\beta(v))} \frac{\Gamma(1 + \alpha(v))}{\Gamma(1 + \alpha(v) + \beta(v))}$ which can take any value in $(0, 1)$ and thus also be arbitrary close to the Dirac delta function.*

Example 3.8 (Linear Beta-binomial model). *Assume that $d = 1$ and that the functions α , β are linear such that $\alpha(v) = m_1 + m_2 v$ and $\beta(v) = m_3 + m_4 v$ where $m_i > 0$ for all $i = 1, \dots, 4$. This specification is discussed in further details in Section 4.4.*

4 Numerical analysis

In this section, we compare the computational costs of the discrete Fourier transform (DFT) and recursive methods, and numerically study the loss distribution of selected factor copula models with different features.

4.1 Computational performance

We show that the DFT method presented in Section 3 is significantly more efficient than the recursive methods suggested in [7, 29]. We consider the standard one-factor and two-factor copula models, Figure 1 displays the computing time necessary to retrieve the probability mass function with the DFT and with the recursive method. The calculations have been performed on a single CPU from a standard personal computer in the R programming language. The DFT method is significantly faster than the recursive method in both cases: it takes roughly the same amount of time to retrieve a distribution with 1000 points with DFT and a 100 points with recursion.

4.2 Dependent defaults with a mixed copula

We investigate the joint default probability and the total number of defaults density in a one-factor copula model with a mixed bivariate copula specification as defined in Equation (6). Fix $K = 2$ and assume that $p_{j,t} = 1 - e^{-\lambda t}$ for $j = 1, 2$ with $\lambda = 0.05$. Consider the following copula mixture

$$C_{U_j, V}(u_j, v) = w C_{U_j, V}^C(u_j, v) + (1 - w) C_{U_j, V}^G(u_j, v)$$

for $j = 1, \dots, N$, for some $w \in [0, 1]$, and where C^C denotes the Clayton copula with parameter 5 and C^G the Gaussian copula with parameter 0.25. Figure 2 (a) displays the probability and cumulative density functions of joint defaults of two entities for the times $0 \leq t \leq 20$, and for the weights $w \in \{0, 0.5, 1\}$. The two limit cases therefore correspond to the Gaussian and Clayton copulas. We observe that the joint probability of default is also a mixture of the two limit cases. With $N = 125$, Figure 2 (b) displays the total number of defaults at a 5-years horizon. It is clear that the distribution of the number of defaults is a mixture of the two limit components: it has the bump of the Gaussian with parameter 0.25 and the fat tail of the Clayton with parameter 5.

4.3 Credit derivatives

We explore the loss distribution of a large portfolio, a tranche on this portfolio, and a portfolio of tranches when the underlying tranches are independent and when they depend on a common factor. Let $N = 1000$ and assume that $\ell_j = 1$ and $\lambda_{jt} = 0.01$ for all $j \in \mathcal{J}$ and $t \geq 0$. The reference model is the standard one-factor

Figure 1: Computation performance.

The time in seconds to compute the loss probability mass function is displayed against the loss support size M for the discrete Fourier transform (black line) and recursive (grey line) methods. The one-factor (left panel) and two-factor (right panel) standard Gaussian copula have been used under the assumption of constant loss given default.

Gaussian copula with correlation parameter equal to 0.25. All the tranches have attachment point $a_k = 100$ and detachment point $b_k = 200$. The CDO squared is composed of 10 tranches to have the same loss support as the portfolio.

Figure 3 displays the probability and cumulative mass functions of the portfolio, tranche, and portfolio of tranches at the 5-year horizon. The tranche loss distribution has two masses at the beginning and end of its support corresponding the probabilities of no loss and full loss respectively. These concentrated masses combine to create a spiky pattern in the portfolio of tranches loss distribution.

The CDO squared loss distribution computed under the assumption of unique factor and tranche specific factor have dramatically different profiles. With independent factors the CDO squared appears less risky than the initial portfolios. For example, senior tranches on the pooled portfolio are virtually riskless. On the other hand, with a unique common factor, the CDO squared has a fat tailed loss distribution and a large probability, about 91%, of having zero losses: when the risk driver behind all tranches is the same, the diversification benefit almost completely disappears. [30] discussed similar results using Monte Carlo simulations.

4.4 Stochastic and correlated loss amounts

In this section, we investigate the impact on the portfolio loss distribution of stochastic losses correlated with the common factor. We consider the linear Beta-Binomial model presented in Section 3.4 with $a_j = 1$, $b_j = 0$, $n_j \delta = 1$, $m_3 = m_1$ and $m_4 = m_2$. Since

$$\mathbb{E}[\ell_j | V = v] = \frac{(m_1 + m_2(1 - v))}{2m_1 + m_2},$$

this specification implies that the expected loss is always equal to 0.5. Figure 4 shows that the V -conditional distribution of the loss amount exhibits various shapes, even-though the expected loss is constant.

Consider the standard one-factor Gaussian copula with parameter equal to 0.25, $N = 125$, $\lambda_j = 0.05$ for all $j \in \mathcal{J}$, and with the same loss amount model as above having an expected loss one half. Figure 5 shows that the loss distribution is significantly affected by the choice of the parameters and the value of the common factor. Compared to the benchmark case of independent and equi-distributed loss amounts (i.e., when $m_2 = 0$), increasing the dependence on the factor V also increases the portfolio average loss and tail risk.

(a) Joint default probability.

(b) Number of defaults.

Figure 2: Dependent defaults with a copula mixture.

The probability (left) and cumulative (right) density functions of the joint default of two homogeneous entities are displayed on Figure 2 (a) for time horizons ranging from 1 week to 20 years for three different one-factor models. The probability (left) and cumulative (right) density functions of the total number of defaults on a portfolio of 125 homogeneous entities are displayed on Figure 2 (b) at a 5-years horizon. The marginal default probability is given by $p_{j,t} = 0.05$ for all $j \in \mathcal{J}$ and the reference factor copula models is an equiweighted copula mixture (black line) between a Gaussian copula with $\rho = 0.25$ (light-grey line) and a Clayton copula with parameter equal to 5 (grey line).

Figure 3: Multi-name credit derivatives losses.

The probability (first row) and cumulative (second row) density functions of the loss distribution are displayed for three different derivatives. The first column is concerned with a portfolio of 1000 entities, the second column with a tranche on this portfolio with attachment point 100 and detachment point 200, and the third column with a portfolio of 10 such tranches coming from different portfolios with a unique risk factor (black line) and with independent risk factors (grey line).

4.5 Number of defaults and loss dependence

We investigate how individual losses affect the portfolio loss distribution given a number of realized defaults. We consider the one-factor homogeneous Gaussian copula with parameter equal to 0.25, default intensities $\lambda_{jt} = 0.05$, $N = 125$ entities, and for a 5-year horizon. Let the V -conditional loss amounts be equal to 1 with probability $1 - v$ and equal to 0 with probability v so that $\mathbb{E}[\ell_j] = 0.5$ for all j .

The left panel of Figure 6 displays the density of the joint distribution of the number of defaults and loss (N_t, L_t) computed as described in Lemma 2.1. We observe that most of the probability mass is concentrated on a diagonal band near the origin, and that there is little to no mass on the off diagonal parts. This is intuitive as more defaults implies larger losses.

The right panel of Figure 6 displays the expected loss given a certain number of default, that is $\mathbb{E}[L_t | N_t = n]$ for $n = 0, \dots, 125$, with the black line corresponding to the model above and the grey line to V -independent loss amounts with $\mathbb{P}[\ell_j = 0] = \mathbb{P}[\ell_j = 1] = 0.5$. This value is increasing with $N_t = n$ as the losses are expected to increase with the total number of defaults. Several interesting observations can be made. First, the marginal rate of losses starts from almost zero at the origin and increases rapidly. Second, the conditional expected loss converges to the maximal possible loss. This is contrasted with V -independent loss amounts where the relation between N_t and L_t is linear.

5 Empirical analysis

In this section, we illustrate our approach by calibrating various factor copula models to credit index tranche prices.

Figure 4: Beta-Binomial distribution.

The distribution of the Beta-Binomial random variables with $n = 10$ is displayed for different parameters choices. The Beta-Binomial loss distribution for a single entity: $m_1 = 1$ and $m_2 = 0$ (top left), $m_1 = 1$ and $m_2 = 1$ (top right), $m_1 = 3$ and $m_2 = 1$ (bottom left), $m_1 = 3$ and $m_2 = 5$ (bottom right). In each panel, the distribution is represented for different values of the systematic factor.

5.1 Data

We focus on tranches of the CDX.NA.IG index, which is composed of 125 investment grade North American companies. Historically, all tranches (except the most junior) were unfunded meaning that they were traded without upfront payments. However, since 2009 and the resolution of the financial crisis, these products are now traded with standardized quarterly payments and the variations in financial risks are reflected in the upfront payments. Based on liquidity and risk, new series with tenors of 3, 5, 7, and 10 years are determined in March and September.

The series 21, issued in September 2013 with a tenor of 5 years, came along with four standardized tranches, whose spreads and attachments/detachments points are detailed in Table 1 (a). Our sample contains 405 daily upfront payments for the four tranches, which we summarize in Table 1 (b) and display in Figure 7. This series was selected because it was the most recent with more than a year of existence in our sample. By convention, the market quotes upfronts in percentage of the corresponding tranche width, which is about thirty times larger for the super-senior than for the equity. Furthermore, the sign of the upfront is also interesting: since it is negative, one most often receives money to buy protection on the super-senior tranche, as well as on the senior tranche at the end of the sample period.

To calibrate recoveries, we looked at all credit events leading to losses (i.e., bankruptcy, failure to pay, and restructuring) between 2005 and 2014 from entities listed in Series 1 to 22 in the CDX.NA.IG, CDX.NA.HY,

Figure 5: Loss distribution and loss amounts dependence.

The probability (left) and cumulative (right) density functions of the loss distribution are displayed for three different loss amounts specifications. With a standard one-factor copula model with parameter equal to 0.25, constant marginal default intensity $\lambda_{jt} = 0.05$, and a 5-year horizon we consider the linear Beta-Binomial loss amounts with $m_1 = m_3$ and $m_2 = m_4$ for the values: $m_1 = 1$ and $m_2 = 0$ (black line), $m_1 = 1$ and $m_2 = 1$ (grey line), $m_1 = 3$ and $m_2 = 1$ (light-grey line), and $m_1 = 3$ and $m_2 = 5$ (dotted light-grey line).

Figure 6: Number of defaults and loss dependence.

The left panel displays the joint probability density of the number of default N_t and the loss L_t when the loss amounts depend on V . The right panel displays the expected loss given n defaults, when the loss amounts are respectively V -independent (grey line) and V -dependent (black line). The reference model is a one-factor homogeneous Gaussian copula with parameter equal to 0.25, default intensities $\lambda_{jt} = 0.05$, a 5-year horizon, and contains $N = 125$ entities. The loss amount ℓ_{jt} is zero or one and has an expected value of 0.5.

iTraxx.Eur, and iTraxx.Eur.Xover. Table 2 presents summary statistics of the 43 credit events leading to losses. One noteworthy observation is that both the mean and median are significantly below the 40% recovery rate commonly assumed when modeling corporate credit portfolios. Furthermore, while about half of the losses are related to CDX.NA.HY entities, which are notoriously more risky, this observation is robust to the removal of such events. In Figure 8, we show the discretized distribution of the recoveries as well as the result of a Beta-binomial fit. The $[0\%, 100\%]$ interval was divided into 10 equal-sized bins and each of the 43 recovery rates was assigned to a bin. Then, Beta-binomial parameters of $\alpha = 0.4$ and $\beta = 1.1$ were obtained by maximizing the likelihood, leading to a fitted mean and standard deviation of 29% and 27%.

Table 1: Structure and summary statistics on CDX.NA.IG.21 tranches.

(a) Attachment-detachment points and the tranche spreads in basis points (or percents) per annum.

Name	Attachement	Detachment	Spread
Equity	0	300	500
Mezzanine	300	700	100
Senior	700	15%	100
Super-senior	15%	100%	25

(b) Statistic on the upfront payments in basis points (or percents) of the tranche width.

	Equity	Mezzanine	Senior	Super-senior
Mean	15.18%	592	-29	-23
Vol	410	274	127	22
Min	859	134	-207	-53
Max	24.87%	13.28%	282	25

Figure 7: CDX.NA.IG.21 upfront values in percents over the sample period.

Figure 7 displays the upfronts for the equity (black), mezzanine (grey), senior (light-grey) and super-senior (dotted light-grey) tranches.

5.2 Calibration

Let P^{a_i, b_i} , a_i and b_i for $i \in \{1, \dots, 4\}$ denote the quoted upfronts, and attachments/detachments points of each tranches. For such derivatives, the contract buyer pays predefined coupons to the seller at the payments dates $0 = T_0 \leq \dots \leq T_n = T$ where T is the contract maturity, we call this series of cash-flow the premium leg V_{prem} . The contract seller pays default contingent cash-flows to the buyer at the defaults dates when losses materialize, we call this series of cash-flow the protection leg V_{prot} . The contract value for the buyer is then given by $V_{\text{prot}} - V_{\text{prem}}$.

For a model parametrized with $\theta \subseteq \Theta \subseteq \mathbb{R}^l$ (i.e., l is the number of parameters), we denote by $P^{a_i, b_i}(\theta)$ the model price, that is the quantity satisfying

$$P^{a_i, b_i}(\theta)(b_i - a_i) + V_{\text{prem}}^{a_i, b_i}(\theta) = V_{\text{prot}}^{a_i, b_i}(\theta),$$

Table 2: Summary statistics for the recoveries

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.	Std
%	0	6	16	28	42	94	28

Figure 8: Recovery data and Beta-binomial fit.

The data points represent the number of credit events leading to losses in each 10% interval divided by the total number of events, and the solid line is the result of a Beta-binomial fit.

where $b_i - a_i$ is the tranche width, and the premium and protection legs are defined as

$$V_{\text{prem}}^{a_i, b_i}(\boldsymbol{\theta}) = S^{a_i, b_i} \mathbb{E}_{\boldsymbol{\theta}} \left[\sum_{j=1}^n e^{-\int_0^{T_j} r_s ds} \int_{T_{j-1}}^{T_j} (b - a - \mathcal{T}_t^{a, b}) dt \right], \quad V_{\text{prot}}^{a_i, b_i}(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{\theta}} \left[\int_0^T e^{-\int_0^t r_s ds} d\mathcal{T}_t^{a_i, b_i} \right],$$

with S^{a_i, b_i} the tranche spread, the time- t risk-free rate r_t , and the risk-free bond price $B(t)$ with maturity t and notional equal to one. In practice, we fix a time grid $0 = t_0 < t_1 < \dots < t_m = T$ with constant and small time step $t_i - t_{i-1} = \Delta t$ and approximate the leg values as follows,

$$V_{\text{prem}}^{a_i, b_i}(\boldsymbol{\theta}) \approx S^{a, b} \sum_{j=1}^n B(T_j) \left((T_j - T_{j-1})(b - a) - \sum_{T_{j-1} \leq t_k \leq T_j} \Delta t E_{\boldsymbol{\theta}} [\mathcal{T}_{t_k}^{a, b}] \right),$$

$$V_{\text{prot}}^{a_i, b_i}(\boldsymbol{\theta}) \approx \sum_{j=1}^m B\left(\frac{t_j + t_{j-1}}{2}\right) \left(\mathbb{E}_{\boldsymbol{\theta}} [\mathcal{T}_{t_j}^{a, b} - \mathcal{T}_{t_{j-1}}^{a, b}] \right),$$

where the discretization assumes that independent short-rate and the default time, as in [43].

Assuming $r_t = 0$ and a homogeneous portfolio with no recovery (i.e., $\ell_j = 1$), we let the default probability be $p_{j,t} = 1 - e^{-\lambda t}$, $j \in \{1, \dots, 125\}$ where λ is the credit index swap spread. When using Beta-binomial recoveries, we use $p_{j,t} = 1 - e^{-\lambda t/(1-R)}$, where R is the expected recovery. Note that we could use individual spreads instead of the index spread. Because we have no access to such data, and, as will be shown in the next section, good calibration to tranches data does not warrant it, such an approach is left as direction of future research.

Finally, the model is calibrated by minimizing the squared pricing error, that is

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta} \in \Theta}{\operatorname{argmin}} \sum_{i=1}^4 \left(P^{a_i, b_i} - P^{a_i, b_i}(\boldsymbol{\theta}) \right)^2. \quad (18)$$

We solve (18) in two steps. First, we explore the parameter space to find a good starting value via a differential evolution algorithm. Second, we use the Nelder-Mead algorithm to refine the solution, enforcing the bounds by means of a parameter transformation.

Figure 9: CDX.NA.IG.21 upfront fits (left) in percents and absolute errors (right) in basis points on January 6th, 2014.

Fits for the one-factor Gaussian (dotted), the one-factor t copula (dashed), the two-factors Gaussian-Clayton copula (grey), and the one-factor mixture with two Gaussians (black). The shaded area is the bid-ask spread.

5.3 Results

Figure 9 displays the fitted upfronts for various copulas models on January 6th, 2014. While the one-factor Gaussian (dotted line) is completely off, both the one-factor t copula (dashed line) and the two-factors Gaussian-Clayton copula (grey line) perform better but miss the senior tranche. The only model achieving a perfect fit (i.e., the black line) is the following one-factor two-Gaussians mixture

$$C_{U_j, V}(u_j, v) = wC_{U_j, V}^{\rho_1}(u_j, v) + (1 - w)C_{U_j, V}^{\rho_2}(u_j, v), \quad j = 1, \dots, 125$$

with $w \in [0, 1]$ and C^{ρ_i} is a Gaussian copula with parameter ρ_i (i.e., $\theta = (w, \rho_1, \rho_2)$ and $\Theta = [0, 1] \times [-1, 1] \times [-1, 1]$).

Repeating (18) for the mixture models each day of the sample, we obtain time-series of calibrated parameters that we display as the plain lines in Figure 10. There are two interesting observations that can be made. First, the parameters do not vary much over time, which indicates that the model is not over-parametrized and can be reliably estimated. Second, whether we assume a Beta-binomial or zero recovery, the parameters are very similar. Third, the parameter ρ_2 is close to 1 whether the recovery is zero or Beta-binomial, which means the second component of the mixture describes a comonotonic relationship between the factor and the uniform random variables for each entity. By fixing $\rho_2 = 0.99$ and calibrating w and ρ_1 , we obtained similar fits in terms of parameter values, with corresponding parameters time series also reported in Figure 10.

Figure 11 displays a model diagnostic for each of the four tranches. For each day in the sample period the pricing errors, namely $P^{a_i, b_i} - P^{a_i, b_i}(\theta_i)$ with $i = 1, 2$, are displayed as well as the bid-ask spread, $P_{ask}^{a_i, b_i} - P_{bid}^{a_i, b_i}$. As the pricing errors are much lower than the bid-ask spread, the equity and mezzanine tranches are perfectly calibrated by both models. For the senior tranche with $\rho_2 = 1$, the end of the sample for the senior tranche and Beta-binomial recovery, and the super-senior tranche however, the pricing errors and the bid-ask spread have the same order of magnitude. To alleviate this issue, we could switch the target of the minimization in the right-hand side of (18) from percentage of the tranche width to dollar amount. In other words, by weighting each term of the sum by $(b_i - a_i)^2$, we would increase the relative importance of the super-senior tranche in the objective function. Nonetheless, the pricing error (and the bid-ask spread) are between 10 and 30 times smaller than the upfront itself.

To summarize, we achieve an almost perfect calibration to all tranches with only two parameters that remain stable over time. Furthermore, the assumptions on the recovery distribution have only a small impact on the overall fit quality. Using this model, one could then price tranches with non-standard attachment and detachment points, or study the impact of some distressed companies using Proposition 2.5.

Figure 10: CDX.NA.IG.21 time series of Gaussian mixture parameters.

Displays the calibrated w (plain), ρ_1 (dashed), and ρ_2 (dotted). Three parameter sets: $\theta_1 = (w, \rho_1, \rho_2)$ with Beta-binomial recovery (top), $\theta_2 = (w, \rho_1, \rho_2)$ with zero recovery (middle) and $\theta_2 = (w, \rho_1, 0.99)$ with zero recovery (bottom).

6 Conclusion

We described a flexible and tractable class of copula-based models for dependent default times and losses. We showed that common standard models are nested as special cases, and that many other models can easily be constructed. With losses taking values on a finite grid, we show how to efficiently compute the exact loss distribution of portfolios using discrete Fourier transform techniques. This allows us to study without simulations the loss distribution of complex portfolios such as portfolio of tranches, also known as CDO-squared. We numerically study how various model features affect the portfolio loss distribution. We calibrate multiple models to credit index tranche prices and show that a particular specification achieve almost perfect calibration to all tranches using only two parameters which appear stable over time. Our framework is therefore a reliable solution for risk-management and pricing applications. Potential future research directions include the estimation of bottom-up models using firm-level data, and the exploration of models with stochastic default intensities.

Acknowledgement: The authors would like to thank the two referees and the associate editor for their helpful suggestions to strengthen the presentation of our work. We also would like to thank for useful comments and discussions Valérie Chavez-Demoulin, Pierre Collin-Dufresne, Damir Filipović, Monique Jeanblanc, and Benjamin Junge, as well as participants from the 2015 CEQURA conference on Advances in Financial and Insurance Risk Management in Munich, the 2016 conference on Dependence Modeling in Finance, Insurance and Environmental Science in Munich, and the 2017 Actuarial and Financial Mathematics conference in Brussels. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 307465-POLYTE.

A Proofs

This appendix contains proofs for theoretical results which, to the best of our knowledge, cannot be found in the copula literature.

Figure 11: Diagnostic of models calibrated on CDX.NA.IG.21 tranches.

Time-series are displayed in basis points with the bid-ask spread (black) and the pricing errors (grey) for the model with either three parameters (w, ρ_1, ρ_2) and Beta-Binomial (top) or zero (middle) recovery, or two parameters ($w, \rho_1, 0.99$) and zero-recovery (bottom).

Proof of Proposition 2.5

We denote U_I the vector which contains the coordinates of U which are in $\mathcal{I} \setminus \mathcal{D}$, and U_J contains the ones which are in $\mathcal{D}$. For readability we assume that the coordinates are ordered according to $U = (U_I, U_J)$. Similarly we group the marginal default probabilities into two vectors p_I and p_J . The size of U_I and U_J are respectively given by N_I and $N_J = N - N_I$. We directly write the proof for the multivariate case with $V \in [0, 1]^d$. The joint default probability conditional on the default of the $k \in \mathcal{D}$ entities then rewrites

$$\begin{aligned} \mathbb{P} [\tau_1 \leq t_1, \dots, \tau_N \leq t_N \mid \tau_k = t_k : k \in \mathcal{D}] &= \mathbb{P} [U_1 \leq p_{1,t_1}, \dots, U_N \leq p_{N,t_N} \mid U_k = p_{k,t_k} : k \in \mathcal{D}] \\ &= \frac{\int_{[0,p_I]} c_{U_I, U_J}(u_I, p_J) du_I}{\int_{[0,1]^{N_I}} c_{U_I, U_J}(u_I, p_J) du_I} = \frac{\int_{[0,1]^d} c_{U_I|V}(p_I \mid v) c_{U_J, V}(p_J, v) dC_V(v)}{\int_{[0,1]^d} c_{U_J, V}(p_J, v) dC_V(v)}. \end{aligned}$$

The second equality comes from the definition of the conditional probability for measures. The third equality follows from the definition of the factor copula and Fubini's theorem,

$$\int_{[0,x]} c_{U_1, U_j}(u_1, p_j) du_1 = \int_{[0,1]^d} \prod_{U_j \in U_1} \int_0^{x_j} c_{U_j, V}(u, v) du \prod_{U_j \in U_1} c_{U_j, V}(u_j, v) dC_V(v),$$

along with $\int_0^{x_j} c_{U_j, V}(u, v) du = C_{U_j|V}(x_j | v)$ and $C_{U_j|V}(1 | v) = 1$.

Alternatively, the same result can be proved by first showing that

$$\begin{aligned} \mathbb{P}[\tau_1 \leq t_1, \dots, \tau_N \leq t_N | \{\tau_k = t_k : k \in \mathcal{D}\} \cup \{V = v\}] &= \frac{\int_{[0, p_1]} c_{U_1, V}(x, v) dx c_{U_j, V}(p_j, v) c_V(v)}{\int_{[0, 1]^{N_1}} c_{U_1, V}(x, v) dx c_{U_j, V}(p_j, v) c_V(v)} \\ &= \prod_{j \in \mathcal{J} \setminus \mathcal{D}} C_{U_j|V}(p_{j, t_j} | v) \end{aligned}$$

and then integrating with respect to the following conditional density

$$\mathbb{P}[V \leq v | \{U_k = p_{k, t_k} : k \in \mathcal{D}\}] = \frac{\int_{[0, v]} \prod_{j \in \mathcal{D}} c_{U_j, V}(p_{j, t_j}, x) dC_V(x)}{\int_{[0, 1]^d} \prod_{j \in \mathcal{D}} c_{U_j, V}(p_{j, t_j}, x) dC_V(x)}.$$

Proof of Proposition 3.1

The default times and the loss amounts being independent conditional on V we have

$$\mathbb{E} \left[e^{iu L_t \delta^{-1}} | V = v \right] = \mathbb{E} \left[e^{iu \sum_{j=1}^N \mathbb{1}_{\{\tau_j \leq t\}} \ell_j \delta^{-1}} | V = v \right] = \prod_{j=1}^N \mathbb{E} \left[e^{iu \mathbb{1}_{\{\tau_j \leq t\}} \ell_j \delta^{-1}} | V = v \right]$$

Furthermore, by independence of the random variables $\mathbb{1}_{\{\tau_j \leq t\}}$ and ℓ_j conditional on V we have

$$\mathbb{E} \left[e^{iu \mathbb{1}_{\{\tau_j \leq t\}} \ell_j \delta^{-1}} | V = v \right] = 1 - \mathbb{P}[\tau_j \leq t | V = v] + \mathbb{P}[\tau_j \leq t | V = v] \phi_{\ell_j}(u, v)$$

where $\phi_{\ell_j}(u, v) := \mathbb{E} \left[e^{iu \ell_j \delta^{-1}} | V = v \right]$ denotes the V -conditional characteristic function of $\ell_j \delta^{-1}$. We finally apply the tower property

$$\begin{aligned} \phi_{L_t}(u) &= \mathbb{E} \left[\mathbb{E} \left[e^{iu L_t \delta^{-1}} | V = v \right] \right] = \int_{[0, 1]^d} \mathbb{E} \left[e^{iu L_t \delta^{-1}} | V = v \right] dC_V(v) \\ &= \int_{[0, 1]^d} (1 - p_{j, t}(v) + p_{j, t}(v) \phi_{\ell_j}(u, v)) dC_V(v) \end{aligned}$$

where C_V is the density of X , and $p_{j, t}(v) = C_{U_j|V}(p_{j, t} | v)$.

Proof of Lemma 3.2

The proof is an application of discrete Fourier transform inversion. Observe that the random variable $L_t \delta^{-1}$ has state space $\{0, 1, \dots, M\}$. Its discrete Fourier transform is given by

$$F_m = \sum_{k=0}^M \mathbb{P} \left[L_t \delta^{-1} = k \right] e^{-i \frac{2\pi m k}{M+1}} = \phi_{L_t} \left(\frac{-2\pi m}{(M+1)} \right)$$

where ϕ_{L_t} as in Proposition 3.1 is the characteristic function of $L_t \delta^{-1}$. The probability mass function can be recovered as follows

$$\mathbb{P} [L_t = k\delta] = \frac{1}{M+1} \sum_{m=0}^M F_m e^{i \frac{2\pi m k}{M+1}}.$$

Equation (13) follows by observing that the signs can equivalently be switched between the complex weights.

Proof of Proposition 3.3

By construction we have

$$\begin{aligned} F_{x,y} &:= \phi_{N_t, L_t}(\mu x, \nu y) = \mathbb{E} \left[\prod_{j=1}^N \exp \left\{ i \mathbb{1}_{\{\tau_j \leq t\}} (\mu x + \nu y \ell_j \delta^{-1}) \right\} \right] \\ &= \mathbb{E} \left[\exp \left\{ \sum_{j=1}^N i \mathbb{1}_{\{\tau_j \leq t\}} (\mu x + \nu y \ell_j \delta^{-1}) \right\} \right] = \mathbb{E} \left[e^{i \mu x N_t + i \nu y L_t \delta^{-1}} \right] \end{aligned}$$

Using this last expectation and the explicit expressions for μ and ν we obtain

$$F_{x,y} = \sum_{j=0}^N \sum_{k=0}^M \mathbb{P}[N_t = j, L_t = \delta k] e^{i \frac{2\pi j}{N+1} x} e^{i \frac{2\pi k}{M+1} y}.$$

This last expression is the two dimensional discrete Fourier transform of the density of the variable $(N_t, L_t \delta^{-1})$. The density can then immediately be retrieved by applying the inverse two-dimensional discrete Fourier transform inversion as follows

$$\mathbb{P}[N_t = j, L_t = \delta k] = \sum_{x=0}^N \sum_{y=0}^M F_{x,y} e^{-i \frac{2\pi x}{N+1} j} e^{-i \frac{2\pi y}{M+1} k}.$$

References

- [1] Aas, K., C. Czado, A. Frigessi, and H. Bakken (2009). Pair-copula constructions of multiple dependence. *Insurance Math. Econom.* 44(2), 182–198.
- [2] Ackerer, D. and D. Filipović (2016). Linear credit risk models. *Swiss Finance Institute Research Paper* 16-34. Available at <https://ssrn.com/abstract=2782455>.
- [3] Albrecher, H., S. A. Ladoucette, and W. Schoutens (2007). A generic one-factor Lévy model for pricing synthetic CDOs. In *Advances in Mathematical Finance*, pp. 259–277. Birkhäuser, Boston MA.
- [4] Altman, E., A. Resti, and A. Sironi (2004). Default recovery rates in credit risk modelling: A review of the literature and empirical evidence. *Economic Notes* 33(2), 183–208.
- [5] Amraoui, S. and S. G. Hitier (2008). Optimal stochastic recovery for base correlation. Available at <https://ssrn.com/abstract=2719672>.
- [6] Andersen, L. and J. Sidenius (2004). Extensions to the Gaussian copula: random recovery and random factor loadings. *J. Credit Risk* 1(1), 29–70.
- [7] Andersen, L., J. Sidenius, and S. Basu (2003). All your hedges in one basket. *RISK* 16 (11), 67–72.
- [8] Bedford, T. and R. M. Cooke (2001). Probability density decomposition for conditionally dependent random variables modeled by vines. *Ann. Math. Artif. Intell.* 32(1–4), 245–268.
- [9] Bedford, T. and R. M. Cooke (2002). Vines - a new graphical model for dependent random variables. *Ann. Statist.* 30(4), 1031–1068.
- [10] Bielecki, T. R. and M. Rutkowski (2013). *Credit Risk: Modeling, Valuation and Hedging*. Springer-Verlag, Berlin.
- [11] Brigo, D., A. Pallavicini, and R. Torresetti (2007). Calibration of CDO tranches with the dynamical generalized-Poisson loss model. *RISK* 20(6), 70–75.
- [12] Brigo, D., A. Pallavicini, and R. Torresetti (2010). *Credit Models and the Crisis: A Journey Into CDOs, Copulas, Correlations and Dynamic Models*. John Wiley & Sons, Chichester.
- [13] Burtschell, X., J. Gregory, and J.-P. Laurent (2007). Beyond the Gaussian copula: stochastic and local correlation. *J. Credit Risk* 3(1), 31–62.
- [14] Burtschell, X., J. Gregory, and J.-P. Laurent (2009). A comparative analysis of CDO pricing models. *J. Derivatives* 16(4), 9–37.
- [15] Carr, P. and D. Madan (1999). Option valuation using the fast Fourier transform. *J. Comput. Financ.* 2(4), 61–73.
- [16] Collin-Dufresne, P. (2009). A short introduction to correlation markets. *J. Financ. Economet.* 7(1), 12–29.
- [17] Cousin, A. and J.-P. Laurent (2008). An overview of factor models for pricing CDO tranches. In R. Cont (Ed.), *Frontiers In Quantitative Finance: Volatility and Credit Risk Modeling*, pp. 185–216. John Wiley & Sons, Hoboken.
- [18] Dufresne, D., J. Garrido, and M. Morales (2009). Fourier inversion formulas in option pricing and insurance. *Methodol. Comput. Appl. Probab.* 11(3), 359–383.

- [19] Embrechts, P., R. Grübel, and S. Pitts (1993). Some applications of the fast Fourier transform algorithm in insurance mathematics. *Statist. Neerlandica* 47(1), 59–75.
- [20] Filipović, D. and M. Larsson (2016). Polynomial diffusions and applications in finance. *Finance Stoch.* 20(4), 931–972.
- [21] Filipović, D., L. Overbeck, and T. Schmidt (2011). Dynamic CDO term structure modeling. *Math. Finance* 21(1), 53–71.
- [22] Fouque, J.-P., R. Sircar, and K. Sølna (2009). Multiname and multiscale default modeling. *Multiscale Model. Simul.* 7(4), 1956–1978.
- [23] Giesecke, K. (2008). Portfolio credit risk: Top-down versus bottom-up approaches. In R. Cont (Ed.), *Frontiers In Quantitative Finance: Volatility and Credit Risk Modeling*, pp. 251–267. John Wiley & Sons, Hoboken.
- [24] Gregory, J. and J.-P. Laurent (2003). I will survive. *Risk* 16(6), 103–107.
- [25] Guillaume, F., P. Jacobs, and W. Schoutens (2009). Pricing and hedging of CDO-squared tranches by using a one factor Lévy model. *Int. J. Theor. Appl. Finance* 12(5), 663–685.
- [26] Heiss, F. and V. Winschel (2008). Likelihood approximation by numerical integration on sparse grids. *J. Econometrics* 144(1), 62–80.
- [27] Herberthsson, A. (2008). Pricing synthetic CDO tranches in a model with default contagion using the matrix-analytic approach. *J. Credit Risk* 4(4), 3–35.
- [28] Hofert, M. and M. Scherer (2011). CDO pricing with nested Archimedean copulas. *Quant. Finance* 11(5), 775–787.
- [29] Hull, J. and A. White (2004). Valuation of a CDO and an n-th to default CDS without Monte Carlo simulation. *J. Derivatives* 12(2), 8–23.
- [30] Hull, J. and A. White (2010). The risk of tranches created from mortgages. *Financ. Analysts J.* 66(5), 54–67.
- [31] Hull, J. C. and A. D. White (2006). Valuing credit derivatives using an implied copula approach. *J. Derivatives* 14(2), 8–28.
- [32] Jakob, K. and M. Fischer (2014). Quantifying the impact of different copulas in a generalized CreditRisk+ framework. *Depend. Model.* 2(1), 1–21.
- [33] Kalemánova, A., B. Schmid, and R. Werner (2007). The normal inverse Gaussian distribution for synthetic CDO pricing. *J. Derivatives* 14(3), 80–94.
- [34] Krekel, M. (2008). Pricing distressed CDOs with base correlation and stochastic recovery. *Rev. Derivatives Res.* 13(3), 219–244.
- [35] Krupskii, P. and H. Joe (2013). Factor copula models for multivariate data. *J. Multivariate Anal.* 120, 85–101.
- [36] Krupskii, P. and H. Joe (2015). Structured factor copula models: theory, inference and computation. *J. Multivariate Anal.* 138, 53–73.
- [37] Laurent, J.-P. and J. Gregory (2005). Basket default swaps, CDOs and factor copulas. *J. Risk* 7(4), 1–20.
- [38] Li, D. X. (2000). On default correlation: A copula function approach. *J. Fixed Income* 9(4), 43–54.
- [39] Li, D. X. and M. H. Liang (2005). CDO squared pricing using Gaussian mixture model with transformation of loss distribution. Available at <https://ssrn.com/abstract=890766>.
- [40] Mai, J.-F., P. Olivares, S. Schenk, and M. Scherer (2014). A multivariate default model with spread and event risk. *Appl. Math. Finance* 21(1), 51–83.
- [41] Mai, J.-F., M. Scherer, and R. Zagst (2012). CIID frailty models and implied copulas. In P. Jaworski, F. Durante, and W. K. Härdle (Eds.), *Copulae in Mathematical and Quantitative Finance*, pp. 201–230. Springer, Heidelberg.
- [42] McNeil, A. J., R. Frey, and P. Embrechts (2015). *Quantitative Risk Management. Concepts, Techniques and Tools*. Revised edition. Princeton University Press.
- [43] Mortensen, A. (2006). Semi-analytical valuation of basket credit derivatives in intensity-based models. *J. Derivatives* 13(4), 8–26.
- [44] Nelsen, R.B. (2006). *An Introduction to Copulas*. Second edition. Springer, New York.
- [45] Oh, D. H. and A. J. Patton (2013). Simulated method of moments estimation for copula-based multivariate models. *J. Amer. Statist. Assoc.* 108(502), 689–700.
- [46] Oh, D. H. and A. J. Patton (2017). Modeling dependence in high dimensions with factor copulas. *J. Bus. Econom. Statist.* 35(1), 139–154.
- [47] Sancetta, A. and S. Satchell (2004). The Bernstein copula and its applications to modeling and approximations of multivariate distributions. *Econometric Theory* 20(3), 535–562.
- [48] Schepsmeier, U. and J. Stöber (2014). Derivatives and Fisher information of bivariate copulas. *Statist. Papers* 55(2), 525–542.
- [49] Schloegl, L. and D. O’Kane (2005). A note on the large homogeneous portfolio approximation with the Student-t copula. *Finance Stoch.* 9(4), 577–584.
- [50] Schönbucher, P. J. and D. Schubert (2001). Copula-dependent default risk in intensity models. Available at <https://ssrn.com/abstract=301968>.
- [51] Vasicek, O. (2002). The distribution of loan portfolio value. *RISK* 15(12), 160–162.
- [52] Zhu, S. H. and M. Pykhtin (2007). A guide to modeling counterparty credit risk. *GARP Risk Review July-August*. Available at <https://ssrn.com/abstract=1032522>.