A Fourier-cosine method for risk quantification and allocation of credit portfolios

Fang Fang[1] Xiaoyu Shen[2]

Nov. 1, 2023

[1]FF Quant Advisory and TU Delft; f.fang@tudelft.nl
[2]FF Quant Advisory; xiaoyu.shen@ffquant.nl

The Background

Portfolio-level risk quantification and allocation in the factor-copula model framework are needed for a few purposes in practice:

- Economic Capital (EC) for the Banking Book;
- Incremental Risk Charge (IRC) under Basel III and Default Risk Charge (DRC) introduced in the Fundamental Review of the Trading Book (FRTB) for the Trading Book;
- Collteralized Debt Obligation (CDO);
- Including wrong way risk in credit valuation adjustment (CVA).

Existing Numerical Methods (1/3)

- Monte Carlo simulation:
  - Pros: Easy to implement and flexible to cope with exotic features;
  - Cons:
    - Slow in convergence, especially low accuracy at regulatory or industry-standard high quantiles;
    - Risk allocation is particularly not just time consuming, but also instable;
    - The computational complexity is linear in the number of obligors, meaning slow speed for large portfolios.

Existing Numerical Methods (2/3)

- Faster alternatives in the literature about 10 or more years ago:
  - Asymptotic approximations based on simplified model assumptions, which is considered not realistic by regulators;
  - FFT-based methods: the computation is still heavy, as the inverse Fourier transform is still based on low-order discretization of the Fourier integral.
  - Wavelets-based methods: no thorough error analysis was provided and is difficult for the industry to embrace due to interpretability.

Existing Numerical Methods (3/3)

- A faster alternative in more recent literature ([2] in 2019):
  - Re-formulating the problem to the Fourier domain and then converting the found solution thereof back to the real domain using the COS method.
  - Pros: the method is faster than MC method and covers single-factor (later also multi-factor) Guassuan and t copula model.
  - Cons:
    - The risk allocation problem, mathematically defined as the Euler decomposition of risk quatified in terms of e.g. Expected Shortfall (ES), and is much more difficult to solve in practice than risk quantification, is still not sovled;
    - It does not cover the hybrid copula structure, such as the Gaussian-t hybrid structure proposed as a benchmark model in 2019 ECB Guide to Internal Models;
    - A rigorous theoretical derivation is lacking for the error convergence, since the portoflio loss distribution is discrete and there is Gibbs phenomenon.

To Summarize - What We Miss in the Literature

- An efficient method to allocate the portfolio-level risk measure, e.g., Expected Shortfal (ES), across sub-portfolio or single credits.
- The existing approach to risk allocation is the importance sampling method; however how to find the alternative sampling distribution can be a problem numerically difficult to sovle when the portfolio is not homogeneous and reported uncertain is high based on simulated confidence intervals.
- Outdated statements: In the discussion of pricing CDO tranches, O'kane (2008) states that "*there has been a general trend away from Fourier methods towards recursion methods*"; and "*recursion is generally faster than Fourier methods*" (i.e., methods based on the Fast Fourier Transform (FFT) algorithm).

Our Goal and Approach

- The goal: developing a fast and generic solution method for risk quantification and allocation (i.e. Euler decomposition of the quantified risk) in the factor-copula model framework.
- The approach: we also take the Fourier apporoach, but
  - we extend the existing literature to tackle the risk allocation problem in particular, in the multi-factor-copula framework. Risk allocation helps identify risk concentration, e.g., identifying the top contributors of a risk measure and quantifying their contribution, or measuring risk contributions of all the obligors in a specific industrial sector.
  - we propose a generic approach how to use this method to cope with hybrid copula structures, such as the Gaussian-t hybrid;
  - we provide a rigorous derivation of the error convergence of applying a filter to Fourier-cosine series expansions to recover the distribution function of a discrete random variable. The theoretical result defines a tighter upper bound of the error than reported in the existing literature and is tested to agrees with numerical results.

The Factor-copula Framework (1/3)

The random variables $x_{n,1\leq n\leq N}$ represent the creditworthiness of the obligors in a reference portfolio of N risky obligors, and x_n's are correlated via the common factors $\mathbf{Z} = [Z_1, Z_2, \cdots, Z_d]^d$, i.e. for $n = 1, \cdots, N$:

$$x_n = \beta_n^T \mathbf{Z} + b_n \varepsilon_n, \tag{1}$$

where

- $\beta_n = [\beta_{n,1}, \beta_{n,2}, \cdots, \beta_{n,d}]^T$ are the weighting coefficients of the common factors $\mathbf{Z}$;
- $b_n = \sqrt{1 - \sum_{i=1}^d \beta_{n,i}^2}$ are the weighting coefficients of the idiosyncratic factor ε_n;
- The common factors are assumed to be independent to the idiosyncratic factor, i.e. $\mathbf{Z} \perp \varepsilon_n$.

The Factor-copula Framework (2/3)

$$x_n = \beta_n^T \mathbf{Z} + b_n \varepsilon_n,$$

- The obligor n defaults if and only if x_n is less than the default threshold, given as the inverse of its CDF;
- The common factors $\mathbf{Z}$ can be independent to each other or governed by a copula, which can impose a different marginal distribution for than that of ε_n;
- In the old literature, Gaussian copula is usually used, i.e. $\mathbf{Z}, \varepsilon_n \sim \mathrm{N}(0,1)$;
- In recent years, Student-t copula has been adopted in the industry at least for benchmark purpose, i.e. $\mathbf{Z}, \varepsilon_n \sim \mathrm{F}_t(0,1)$;

The Factor-copula Framework (3/3)

- There are factor-copula models where the systematic factors and the idiosyncratic factors follow different distributions, e.g.
  - Oh and Patten 2017
  - 2019 ECB Guide to Internal Models
- To illustrate the generic method, in our example we assume the systematic factors follow a d-variate t distribution. The idiosyncratic factors remain Gaussian. (It can be made the other way around or assumed to follow other distributions)
  - A more convenient formulation of the hybrid copula model:

$$x_n = \sqrt{W}\beta_n^T \mathbf{Z} + b_n \varepsilon_n. \tag{2}$$

where W has an inverse gamma distribution, i.e, $W \sim I_g\left(\nu/2, \nu/2\right)$. ν is the degrees of freedom.

Mathematical Formulation of the Problem-to-solve

The portfolio loss L is defined as

$$L = \sum_{n=1}^{N} \mathbf{1}_{x_n \leq \xi_n} \cdot l_n \tag{3}$$

Then the problem-to-solve comprise two parts

- Risk quantification: Valut-at-Risk (VaR) and ES of L.
- Risk allocation: Euler risk allocation of VaR and ES.
  - A risk measure is decomposed as the sum of risk contributions of the obligors/sub-portfolios in the reference portfolio.
  - Homogenous: scaling the risk measure by a constant changes the risk decomposition by the same scale. E.g., increasing the loss-at-default of all the obligors by 10% would increase VaR/ES by 10% and the risk contribution of a certain obligor should also increase by 10%.

Recall the COS Method (1/2)

The essence of the COS method is that, a probability density function can be recovered from a truncated Fourier cosine series, of which the coefficients can be extracted from the characteristics function (ch.f), and thus, are readily available.

- That is, within the truncation range $[a, b]$ of a density function f, we have

$$f(x) \approx {\sum_{k=0}^{K}}' A_k \cos\left(k\pi \frac{x-a}{b-a}\right), \tag{4}$$

where

$$A_k = \frac{2}{b-a} \mathrm{Re}\left\{\varphi\left(\frac{k\pi}{b-a}\right) \cdot \exp\left(-i\frac{ka\pi}{b-a}\right)\right\}$$

with $\varphi(\cdot)$ being the ch.f. of $f(x)$, and $\sum'$ indicates that the first term in the sum is weighted by one-half.

Recall the COS Method (2/2)

To apply COS to the portfolio loss distribution of a multifactor copula model, we

1. first numerically evaluate the ch.f at a grid of points in the Fourier domain, i.e., $k\pi\frac{x-a}{b-a}, 0 \leq k \leq K$.
2. then reconstruct the CDF function of the loss by COS, i.e.

$$F(y) = \int_a^y f(x)dx = \frac{A_0}{2}(y-a) + \sum_{k=1}^{K} A_k \frac{b-a}{k\pi} \sin\left(k\pi\frac{y-a}{b-a}\right) \tag{5}$$

It does not reply on the assumption of Gaussian distributions!

The COS Approach for Risk Quantification

The key is to solve the characteristic function (ch.f.) of the portolio loss L. The ch.f. is derived as follows:

1. Conditional on the common factors, defaults of the obligors are independent Bernoulli random variables. Thus the conditional ch.f of the total loss $L = \sum l_n \mathbf{1}_{x_n \leq \xi_n}$ is

$$\mathbb{E}\left[\varphi_L(\omega)|Z\right] = \Pi_{n=1}^{N} \mathbb{E}\left[e^{i\omega l_n \cdot \mathbf{1}_{\varepsilon_n \leq \alpha_n(\mathbf{z}_n)}}\right] \tag{6}$$

   where $\alpha_n(\mathbf{z}_n) = \frac{\xi_n - \beta_n^T \mathbf{z}}{b_n}$

2. Each expectation $\mathbb{E}\left[e^{i\omega l_n \cdot \mathbf{1}_{\varepsilon_n \leq \alpha_n(\mathbf{z}_n)}}\right]$ in the product is simply given by the analytical expression of the Bernoulli ch.f.
3. Finally, the ch.f of the portfolio loss distribution can be obtained from the conditional ch.f $\mathbb{E}\left[\varphi_L(\omega)|Z\right]$ by numerically integrating out the common factors $\mathbf{Z}$.

Risk Measures via the COS-recovered CDF

- VaR : Very simple! Given the recovered CDF of the portfolio loss, the q-th quantile can be solved numerically, e.g., solving $P(L \leq \theta) = q$ via a root-searching algorithm
- ES: An analytical expression for ES is available, by integrating the loss with respect to the Fourier series expansion of CDF.

Euler Risk Allocation of ES

Conditional ES decomposes the ES by the Euler principle for risk allocation. We consider the following definition:

$$\mathrm{CES}_n = \mathbb{E}\left[\mathbf{1}_{x_n \leq \xi_n} \cdot l_n \middle| L \geq \mathrm{VaR}_\alpha\right].$$

Such that

$$\mathrm{ES} = \mathbb{E}\left[\sum_n \mathbf{1}_{x_n \leq \xi_n} \cdot l_n \middle| L \geq \mathrm{VaR}_\alpha\right] = \sum_n \mathrm{CES}_n$$

Our Solution for Euler Risk Allocation of ES (1/2)

Apply Bayes law, we yield that

$$\begin{aligned}\mathrm{CES}_n &= l_n \cdot P\left(x_n \leq \xi_n \,|\, L \geq \mathrm{VaR}_\alpha\right) \\ &= l_n \cdot \frac{P\left(x_n \leq \xi_n, L \geq \mathrm{VaR}_\alpha\right)}{P\left(L \geq \mathrm{VaR}_\alpha\right)} \\ &= \frac{l_n \cdot p_n}{\alpha} \cdot P\left(L \geq \mathrm{VaR}_\alpha \,|\, x_n \leq \xi_n\right).\end{aligned} \tag{7}$$

$P\left(L \geq \mathrm{VaR}_\alpha \,|\, x_n \leq \xi_n\right)$ can be solved by the COS approach again as for VaR and ES!

QUANT TU Delft

Our Solution for Euler Risk Allocation of ES (2/2)

The coding of the COS calculation for $P\left(L \geq \mathrm{VaR}_\alpha \middle| x_n \leq \xi_n\right)$ can be easily integrated into the coding of the COS calculation for the portfolio loss distribution.

$$\begin{aligned}
\varphi_{n,L}(\omega) &= \mathbb{E}\left[e^{i\omega L}\middle| x_n \leq \xi_n\right] \\
&= \frac{\mathbb{E}\left[e^{i\omega L}\cdot \mathbf{1}_{x_n \leq \xi_n}\right]}{P\left(x_n \leq \xi_n\right)} \\
&= \frac{1}{p_n}\mathbb{E}\left[\mathbb{E}\left[e^{i\omega L}\cdot \mathbf{1}_{x_n \leq \xi_n}\middle| \mathbf{Z} = \mathbf{z}\right]\right] \\
&= \frac{1}{p_n}\mathbb{E}\left[\left(\prod_{j\neq n}\mathbb{E}\left[e^{i\omega l_j\cdot \mathbf{1}_{\varepsilon_j \leq \alpha_j(\mathbf{z}_j)}}\middle| \mathbf{Z} = \mathbf{z}\right]\right)\right. \\
&\qquad \left.\cdot\, \mathbb{E}\left[e^{i\omega l_n\cdot \mathbf{1}_{\varepsilon_n \leq \alpha_n(\mathbf{z}_n)}}\cdot \mathbf{1}_{\varepsilon_n \leq \alpha_n(\mathbf{z}_n)}\middle| \mathbf{Z} = \mathbf{z}\right]\right] \qquad (8)
\end{aligned}$$

Euler Risk Allocation

Conditional VaR decomposes the VaR by the Euler principle for risk allocation.

$$\mathrm{CVaR}_n = \mathbb{E}\left[\mathbf{1}_{x_n \leq \xi_n} \cdot l_n \middle| L = \mathrm{VaR}_\alpha\right].$$

Such that

$$\mathrm{VaR} = \mathbb{E}\left[\sum_n \mathbf{1}_{x_n \leq \xi_n} \cdot l_n \middle| L = \mathrm{VaR}_\alpha\right] = \sum_n \mathrm{CVaR}_n$$

The rest of the calculation follows the same steps as for the Euler risk allocation of ES, except that we need to evaluate the expectation conditional on a small neighborhood around VaR_α.

Our Solution for Conditional VaR

Similar to Conditional ES, Conditional VaR decomposes the VaR by the Euler principle for risk allocation. We consider the following similar definition:

$$\mathrm{CVaR}_n = \mathbb{E}\left[\mathbf{1}_{x_n \le \xi_n} \cdot l_n \middle| L = \mathrm{VaR}_\alpha\right].$$

So that

$$\mathrm{VaR} = \mathbb{E}\left[\sum_n \mathbf{1}_{x_n \le \xi_n} \cdot l_n \middle| L = \mathrm{VaR}_\alpha\right] = \sum_n \mathrm{CVaR}_n$$

It then follows that

$$\begin{aligned}
\mathrm{CVaR}_n &= l_n \cdot P\left(x_n \le \xi_n \,|L = \mathrm{VaR}_\alpha\right) \\
&= l_n \cdot \frac{P\left(x_n \le \xi_n, L = \mathrm{VaR}_\alpha\right)}{P\left(L = \mathrm{VaR}_\alpha\right)} \\
&\approx l_n \cdot \frac{P\left(x_n \le \xi_n, \mathrm{VaR}_\alpha - \epsilon \le L \le \mathrm{VaR}_\alpha + \epsilon\right)}{P\left(\mathrm{VaR}_\alpha - \epsilon \le L \le \mathrm{VaR}_\alpha + \epsilon\right)} \\
&= l_n \cdot p_n \cdot \frac{P\left(\mathrm{VaR}_\alpha - \epsilon \le L \le \mathrm{VaR}_\alpha + \epsilon \middle| x_n \le \xi_n\right)}{P\left(\mathrm{VaR}_\alpha - \epsilon \le L \le \mathrm{VaR}_\alpha + \epsilon\right)}
\end{aligned} \tag{9}$$

Gibb's Phenomenon

However, there is an issue - L is a discrete random variable, which gives rise to the Gibb's phenomenon in the Fourier-series.

- Gibb's phenomenon: very slow or no convergence of the series due to discontinuities in the function.
- Appears as overshooting and undershooting close to the discontinuous points.
- It is an issue for all eigen decomposition based methods.

Example of Bernoulli Distribution

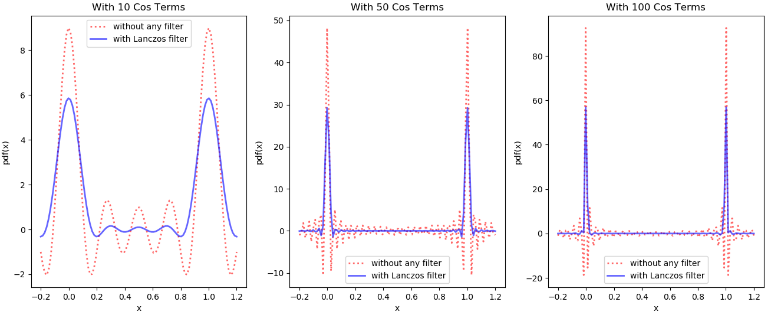

Solutions to Gibb's Phenomenon

- Solutions in existing literature:
  - Fourier space filters: enhancing the decay rate of the given Fourier coefficients without reducing the accuracy.
    - The Lanczos filter: $\sigma(\eta) = \frac{\sin(\pi\eta)}{\pi\eta}$
    - Higher order filters, such as raised cosine filter, exponential filter, Daubechics filter, etc.
  - Filters in physical space: localizing the information that determines the Fourier coefficients by means of convolution.
  - In essence, these two types of solutions are equivalent.

Adjusted Formulas with Filters

- Portfolio loss density function is a discrete function and the CDF is a piece-wise constant function. Thus, Gibb's phenomenon can have impact on the accruacy when VaR level is close to the discountinuous points.
- We chose Fourier space filters, as the only modification needed is on the Fourier coefficients.
- The adjusted COS formula for CDF of portfolio loss:

$$F(y) \approx \frac{A_0}{2}(y-a)+\sum_{k=1}^{K} A_k \sigma\left(\frac{k}{K}\right)\frac{b-a}{k\pi}\sin\left(k\pi\frac{y-a}{b-a}\right) \quad (10)$$

Error Analysis - 1/7

- Denote the possible realizations of L by $\{0 \leq L_0 \leq L_1, \cdots, L_m, \cdots, \leq L_M \leq \pi\}$.
- Applying the COS expansion to have

$$f_L(x) = \sum_{k=0}^{\prime\infty} A_k \cos kx$$

with

$$A_k = \frac{2}{\pi}\mathrm{Re}\{\varphi(k)\} = \frac{2}{\pi}\mathrm{Re}\left\{\sum_{m=0}^{M} e^{ikL_m} p_m\right\} = \frac{2}{\pi}\sum_{m=0}^{M}\cos(kL_m)p_m \qquad (11)$$

where p_m is the probability of $L = L_m$.

Error Analysis - 2/7

- Thus the Fourier cosine expansion of the probability density of L is

$$
\begin{aligned}
f_L(x) &= {\sum_{k=0}^{\infty}}' \frac{2}{\pi} \sum_{m=0}^{M} \cos(kL_m) p_m \cos kx \\
&= \sum_{m=0}^{M} p_m {\sum_{k=0}^{\infty}}' \frac{2}{\pi} \cos(kL_m) \cos kx \\
&= \sum_{m=0}^{M} p_m f_m(x) \qquad (12)
\end{aligned}
$$

where

$$
f_m(x) = {\sum_{k=0}^{\infty}}' \frac{2}{\pi} \cos(kL_m) \cos kx
$$

Error Analysis - 3/7

- Integrating f_L from 0 gives the COS CDF of L

$$
\begin{aligned}
F_L(x) &= \sum_{m=0}^{M} p_m \left[\frac{1}{\pi} x + \sum_{k=1}^{\infty} \frac{2}{k\pi} \cos(kL_m) \sin kx \right] \\
&= \sum_{m=0}^{M} p_m F_m(x)
\end{aligned}
\tag{13}
$$

where

$$
F_m(x) = \frac{1}{\pi} x + \sum_{k=1}^{\infty} \frac{2}{k\pi} \cos(kL_m) \sin kx
$$

Error Analysis - 4/7

Then we truncate the number of series term to K and modify the series coefficients by the filter to have

$$f_L^\sigma(x) = \sum_{m=0}^{M} p_m f_m^\sigma(x) \tag{14}$$

and

$$F_L^\sigma(x) = \sum_{m=0}^{M} p_m F_m^\sigma(x) \tag{15}$$

where

$$f_m^\sigma(x) = {\sum_{k=0}}'^{K} \frac{2}{\pi}\sigma(k/K)\cos(kL_m)\cos kx$$

and

$$F_m^\sigma(x) = \frac{1}{\pi}x + \sum_{k=1}^{K} \frac{2}{k\pi}\sigma(k/K)\cos(kL_m)\sin kx$$

Error Analysis - 4/7

The key insight here is that on $[-\pi, \pi]$,

$$F_0(x) = \frac{1}{\pi}x + \sum_{k=1}^{\infty} \frac{2}{k\pi} \sin kx$$

is the Fourier series expansion of the function

$$H_0(x) = \begin{cases} 1 & \text{if } 0 \leq x \leq \pi \\ -1 & \text{if } -\pi \leq x < 0 \end{cases}$$

and that

$$F_m(x) = \frac{1}{\pi}x + \sum_{k=1}^{\infty} \frac{2}{k\pi} \cos(kL_m) \sin kx$$

is the Fourier series expansion of the function

$$H_m(x) = \begin{cases} 1 & \text{if } L_m \leq x \leq \pi \\ 0 & \text{if } -L_m < x < L_m \\ -1 & \text{if } -\pi \leq x < L_m \end{cases}$$

Error Analysis - 6/7

The convergence speed of the Fourier series expansion with spectral filter for a piecewise constant function is governed by the convergence order of the filter, as proven in [Vandeven 1991]:

- If we have a function $f \notin C^{p-1}$, ie, if $f(y)$ has a jump discontinuity at one or more points of order smaller than, or equal to, $p-1$, the following estimate holds:

$$f_N^{\sigma}(y) - f(y) \sim O\left(N^{1/2-p}\right).$$

Given that the CDF of L is a linear combination of H_m, weighted by p_m, it follows that F_L^{σ} converges to the true CDF of L at the speed as described above.

Error Analysis - 7/7

Recall that there is one extra layer of approximation: the cosine coefficients A_k is obtained via numerical integration based on Clenshaw-Curtis rule after we truncate the integration range with a truncation error at the level of TOL. Let us denote the error term from this numerical integration part as $\epsilon(N, TOL)$, which depends on the the number of integration points N and the range truncation tolerance TOL. Then it can be shown that this error term propogates in our approximation of the CDF as follows:

$$\hat{F}^{\sigma}_{m}(x) = F^{\sigma}_{m}(x) + O(K) \cdot \epsilon(N, TOL).$$

Numerical Example 1: A Small portfolio

To observe the behavior of the COS method, we first constructed a simple portfolio with 10 obligors, one of which creates name concentration. We consider a two-factor Gaussian copula, and a hybrid copula with Student-t distribution for the systematic factors and Gaussian distribution for the idiosyncratic factors.

- Number of obligors: 10
- $\beta_{n,1} = 0.8, \beta_{n,2} = 0.4$
- $p_1 = 0.01, p_n = 0.001, n = 2, \cdots, N$
- $l_1 = 10, l_n = 1, n = 2, \cdots, N$
- Degree of freedom in the t Copula: 8.

CDF Conditional on Default of One Name

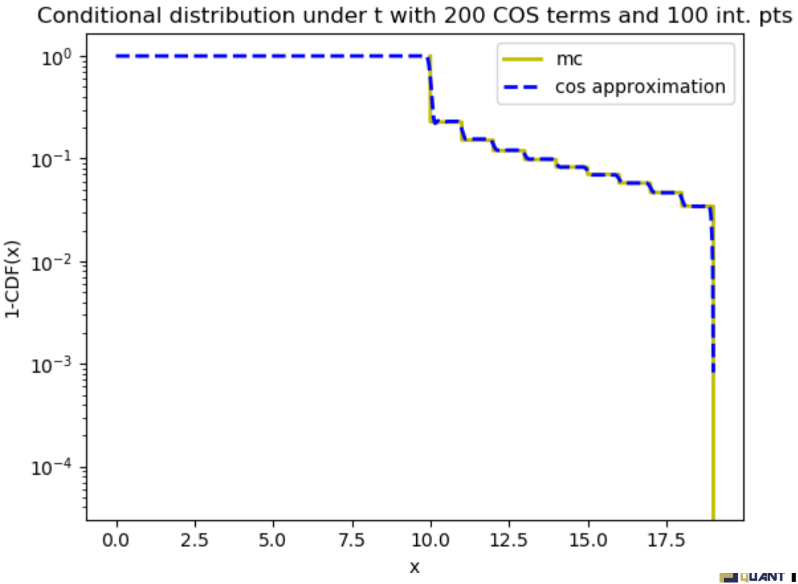

Numerical Example 2: A Large portfolio

- Number of obligors: 1000
- Ratings are uniformly sampled from AAA, AA, A, BBB, BB, B and CCC.
- SP PDs.
- Losses are uniformly sampled from $[10, 1000]$.
- Create a few name concentration of CCC obligors by multiplying the losses by a factor of 50 or 10.
- Factor loadings $\beta_{n,1}, \beta_{n,2}$ are randomly drawn from $[0, 1]$.
- Degree of freedom in the t copula: 8

CDF of Portfolio Loss

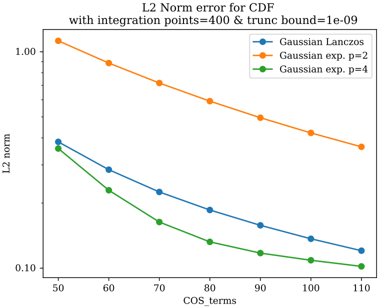

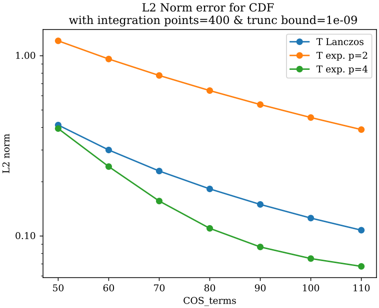

Left: Gaussian copula. Right: T copula.

VaR of Portfolio Loss

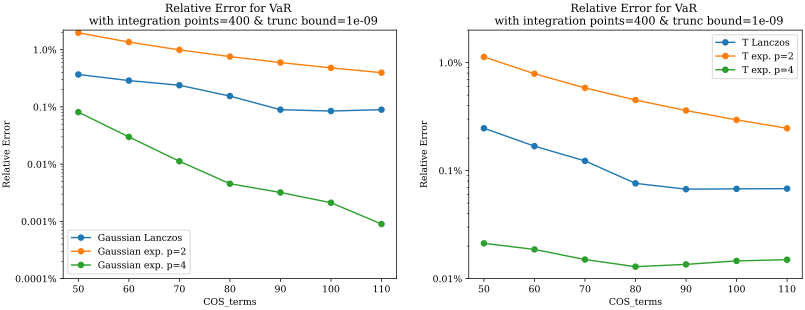

Left: Gaussian copula. Right: T copula.

ES of Portfolio Loss

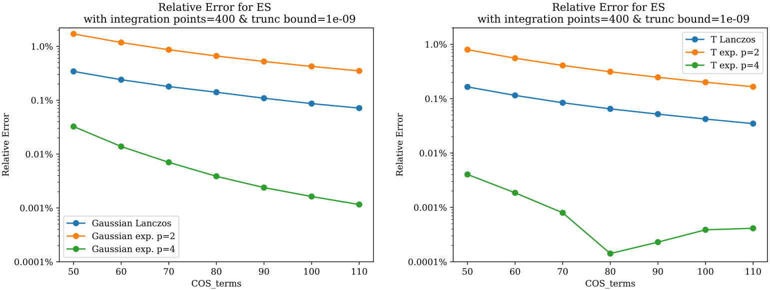

Left: Gaussian copula. Right: T copula.

CES of Portfolio Loss

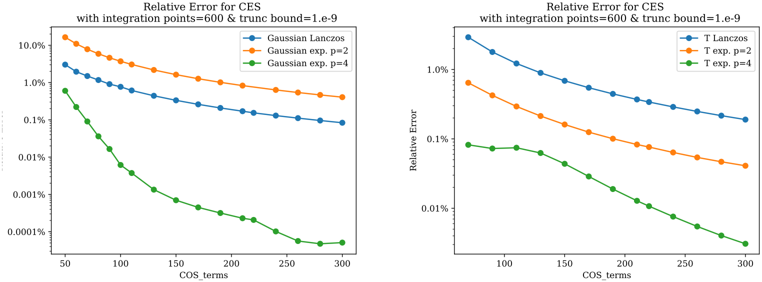

Left:Gaussian copula. Right: T copula.

Performance

Copula	Method	Time (sec.)	VaR Val.	ES Val.	Abs. Err. VaR	Abs. Err. ES	Rel. Err. VaR	Rel. Err. ES
Gaussian-t hybrid	COS (Nr. COS terms = 110, exp. filter p=4)	28	2,333.61	2,810.32	0.55	3.85	0.02%	0.14%
	MC	413	2,340.97	2,822.40	7.91	8.23	0.34%	0.29%
Gaussian	COS (Nr. COS terms = 110, exp. filter p=4)	20	1,967.53	2,215.72	0.16	0.79	0.01%	0.04%
	MC	268	1,971.86	2,224.60	4.49	9.67	0.23%	0.44%

Conclusions

- Key insight - we can solve the problem of both risk quantification and allocation in the Fourier domain with the help of the COS method.
- For dimension less than 4, this method is possible for real-time calculations, such as loan pricing.
- Current and future research: Tackle the curse of dimension via various techniques.

Reference

Dominic O'kane
Modelling single-name and multi-name cedit derivatives
The Wiley Finance Series, (2008).

G.C. Papiol, L.O.Gracia and C.W. Oosterlee
Quantifying credit portfolio losses under multif-factor models
International J of Computer Mathematics, 96(11): 2135-2156, 2019. https://doi.org/10.1080/00207160.2018.1447666.

F. Fang and C.W. Oosterlee
A Novel Pricing Method for European Options Based on Fourier-Cosine Series Expansions
SIAMJournal on Scientific Computing, vol. 31, no. 2, pp. 826–848, 2009..