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

Fang Fang[1] Xiaoyu Shen[2] Chujun Qiu[3]

May 31, 2023

[1]FF Quant Advisory B.V. and Delft University of Technology ; f.fang@tudelft.nl
[2]FF Quant Advisory BV., the Netherlands; xiaoyu.shen@ffquant.nl
[3]Tsinghua University, China

The problems to solve

We aim at developing a much faster solution method for risk quantification and allocation under the factor-copula model. Examples of such problems include:

- Economic Capital (EC) and Risk-adjusted Return on Capital (ROROC) for the banking books;
- Incremental Risk Charge (IRC) and Default Risk Charge (DRC) for the trading books;
- Collteralized Debt Obligation (CDO);
- Wrong way risk of the correlation between exposure and default in credit valuation adjustment (CVA).

Existing methods

- Monte-Carlo simulation:
  - Easy to implement and flexible in coping with exotic features;
  - But slow convergence, especially low accuracy at standard high quantiles;
  - Risk allocation is particularly time consuming;
  - The computational complexity is linear in the number of obligors, meaning slow speed for large portfolios.
- Faster alternatives:
  - Asymptotic approximations based on simplified model assumptions, which might not be realistic in practice;
  - FFT-based methods: the computation is still heavy.
  - Wavelets-based methods: no thorough error analysis was provided and is difficult for the industry to embrace due to interpretability...

Why revisting Fourier methods? (1/3)

- 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).
- To speed up it is natural to consider the Fourier cosine method (COS) instead of FFT; COS can directly sample the Fourier coffiecients from the Fourier transform.

Why revisiting Fourier methods? (2/3)

- In the literature there have been applications of COS to solve the credit loss distributions under the factor copula models. E.g., [2].
- However, a solid theory is lacking for the error convergence, since the portoflio loss distribution is discrete and Gibbs phenomenon appears.

Why revisiting Fourier methods? (3/3)

- What is also missing in the literature, is a Fourier-based 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.

The contribution of this work

In this work, we

- provide the theoretical convergence proof of applying the COS method to recover the distribution of discrete random variables.
- derive the COS formula for risk allocation via conditional VaR, which is lacking in literature.
- as a by-product of the derivations, we yield a fast alternative to MC for risk quantification.

Multifactor Gaussian Copula

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 a standard $d-$variate Gaussian random vector $\mathbf{Z} = [Z_1, Z_2, \cdots, Z_d]^d$. I.e., for $n = 1, \cdots, N$:

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

where

- $\beta_n = [\beta_{n,1}, \beta_{n,2}, \cdots, \beta_{n,d}]^T$, $b_n = \sqrt{1 - \sum_{i=1}^{d} \beta_{n,i}^2}$, and $\varepsilon_n \sim \mathrm{N}(0, 1)$.
- $\mathbf{Z} \perp \varepsilon_n$, and, $\mathbf{Z}$ and ε_n represent the systematic factors and idiosyncratic factors, respectively.
- The obligor n defaults if and only if x_n is less than the default threshold, given as the inverse of the CDF of the standard normal random variable at the default probability p_n, i.e.

$$\xi_n = \mathrm{N}^{-1}(p_n).$$

Factor Copulas with Mixed Distributions

There are factor copula models where the systematic factors and the idiosyncratic factors follow different distributions (e.g. Oh and Patten 2017 and the recent ECB Guide to Internal Models). Particularly, we consider that the systematic factors follow a d-variate t distribution. The idiosyncratic factors remain Gaussian.

Re-parametrize the equation in the Gaussian copula as

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

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.

The COS Method

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

$$f(x) \approx {\sum_{k=0}^{K}}' A_k \cos\left(k\pi \frac{x-a}{b-a}\right), \tag{3}$$

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.

The COS Method - Contd.

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.

$$F(y) = \int_a^y f(x)dx = \frac{A_0}{2}(y-a) + \sum_{k=1}^{K} A_k \frac{b-a}{k\pi} \sin\left(k\pi\frac{y-a}{b-a}\right) \tag{4}$$

It does not reply on the assumption of Gaussian distributions!

Characteristic Function of the Portfolio Loss

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

$$\mathbb{E}\left[\varphi_L(\omega)|Z\right] = \Pi_{n=1}^N \mathbb{E}\left[e^{i\omega l_n \cdot \mathbf{1}_{\varepsilon_n \le \alpha_n(\mathbf{z}_n)}}\right] \tag{5}$$

   where $\alpha_n(\mathbf{z}_n) = \frac{\xi_n - \beta_n^T \mathbf{z}}{b_n}$
2. Each expectation $\mathbb{E}\left[e^{i\omega l_n \cdot \mathbf{1}_{\varepsilon_n \le \alpha_n(\mathbf{z}_n)}}\right]$ 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 $\mathbf{Z}$.

TU Delft

Risk Measures via the COS Method

- 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.

Risk Allocation

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.

Some properties of the Euler risk allocation principle:

- 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%.

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

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

Euler Risk Allocation of ES

It then follows that

$$\begin{aligned}\mathrm{CES}_n &= l_n \cdot P\left(x_n \leq \xi_n \,|\, L \geq \mathrm{VaR}_\alpha\right)\\ &= l_n \cdot \frac{P\left(x_n \leq \xi_n, L \geq \mathrm{VaR}_\alpha\right)}{P\left(L \geq \mathrm{VaR}_\alpha\right)}\\ &= \frac{l_n \cdot p_n}{\alpha} \cdot P\left(L \geq \mathrm{VaR}_\alpha \,|\, x_n \leq \xi_n\right). \end{aligned} \tag{6}$$

$P\left(L \geq \mathrm{VaR}_\alpha \,|\, x_n \leq \xi_n\right)$ can be solved by the COS Method!

Euler Risk Allocation of ES

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.

$$\begin{aligned}
\varphi_{n,L}(\omega) &= \mathbb{E}\left[e^{i\omega L}\middle| x_n \leq \xi_n\right] \\
&= \frac{\mathbb{E}\left[e^{i\omega L}\cdot \mathbf{1}_{x_n\leq \xi_n}\right]}{P\left(x_n \leq \xi_n\right)} \\
&= \frac{1}{p_n}\mathbb{E}\left[\mathbb{E}\left[e^{i\omega L}\cdot \mathbf{1}_{x_n\leq \xi_n}\middle| \mathbf{Z}=\mathbf{z}\right]\right] \\
&= \frac{1}{p_n}\mathbb{E}\left[\left(\prod_{j\neq n}\mathbb{E}\left[e^{i\omega l_j\cdot \mathbf{1}_{\varepsilon_j\leq \alpha_j(\mathbf{z}_j)}}\middle| \mathbf{Z}=\mathbf{z}\right]\right)\right. \\
&\qquad \left.\cdot\, \mathbb{E}\left[e^{i\omega l_n\cdot \mathbf{1}_{\varepsilon_n\leq \alpha_n(\mathbf{z}_n)}}\cdot \mathbf{1}_{\varepsilon_n\leq \alpha_n(\mathbf{z}_n)}\middle| \mathbf{Z}=\mathbf{z}\right]\right] \qquad (7)
\end{aligned}$$

TU Delft

Euler Risk Allocation of VaR

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

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

Such that

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

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.

TU Delft

Conditional VaR

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

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

So that

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

It then follows that

$$\begin{aligned}
\mathrm{CVaR}_n &= l_n \cdot P\left(x_n \leq \xi_n \,|L = \mathrm{VaR}_\alpha\right) \\
&= l_n \cdot \frac{P\left(x_n \leq \xi_n, L = \mathrm{VaR}_\alpha\right)}{P\left(L = \mathrm{VaR}_\alpha\right)} \\
&\approx l_n \cdot \frac{P\left(x_n \leq \xi_n, \mathrm{VaR}_\alpha - \epsilon \leq L \leq \mathrm{VaR}_\alpha + \epsilon\right)}{P\left(\mathrm{VaR}_\alpha - \epsilon \leq L \leq \mathrm{VaR}_\alpha + \epsilon\right)} \\
&= l_n \cdot p_n \cdot \frac{P\left(\mathrm{VaR}_\alpha - \epsilon \leq L \leq \mathrm{VaR}_\alpha + \epsilon \middle| x_n \leq \xi_n\right)}{P\left(\mathrm{VaR}_\alpha - \epsilon \leq L \leq \mathrm{VaR}_\alpha + \epsilon\right)}
\end{aligned} \tag{8}$$

Gibb's Phenomenon and the Solution

- 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.
- Solutions:
  - 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.

Example of Bernoulli Distribution

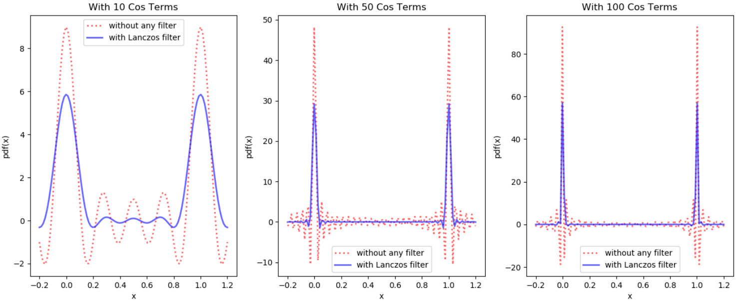

Adjusted Formulas with Filters

- 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:

$$F(y) \approx \frac{A_0}{2}(y-a) + \sum_{k=1}^{K} A_k \sigma\left(\frac{k}{K}\right) \frac{b-a}{k\pi} \sin\left(k\pi \frac{y-a}{b-a}\right) \quad (9)$$

TU Delft

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

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

with

$$A_k = \frac{2}{\pi}\mathrm{Re}\{\varphi(k)\} = \frac{2}{\pi}\mathrm{Re}\left\{\sum_{m=0}^{M} e^{ikL_m} p_m\right\} = \frac{2}{\pi}\sum_{m=0}^{M}\cos(kL_m)p_m \tag{10}$$

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

$$
\begin{aligned}
f_L(x) &= {\sum_{k=0}^{\infty}}' \frac{2}{\pi} \sum_{m=0}^{M} \cos(kL_m) p_m \cos kx \\
&= \sum_{m=0}^{M} p_m {\sum_{k=0}^{\infty}}' \frac{2}{\pi} \cos(kL_m) \cos kx \\
&= \sum_{m=0}^{M} p_m f_m(x) \qquad (11)
\end{aligned}
$$

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

$$F_L(x) = \sum_{m=0}^{M} p_m \left[\frac{1}{\pi} x + \sum_{k=1}^{\infty} \frac{2}{k\pi} \cos(kL_m) \sin kx \right] = \sum_{m=0}^{M} p_m F_m(x) \tag{12}$$

where

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

TU Delft

Error Analysis - 4/7

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

$$f_L^\sigma(x) = \sum_{m=0}^{M} p_m f_m^\sigma(x) \tag{13}$$

and

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

where

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

and

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

TU Delft

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

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

TU Delft

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:

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

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.

TU Delft

CDF Conditional on Default of One Name

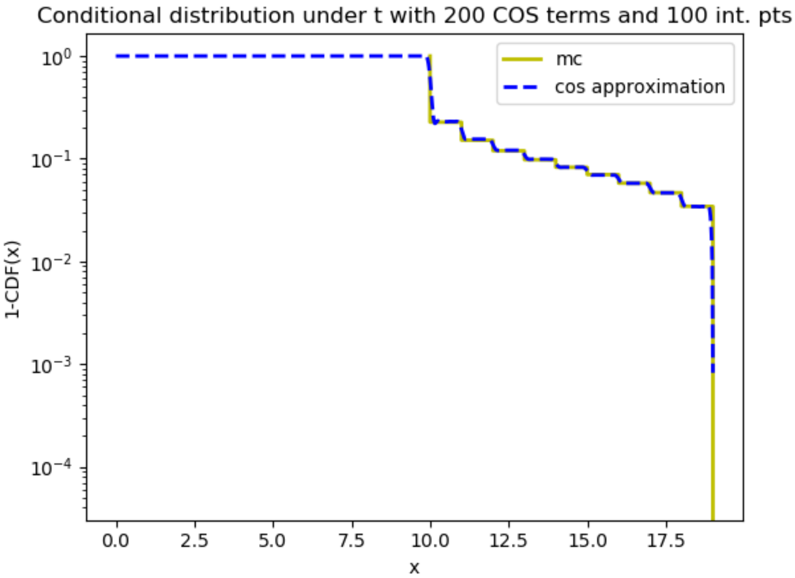

Numerical Example 2: A Large portfolio

- 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

TU Delft

CDF of Portfolio Loss

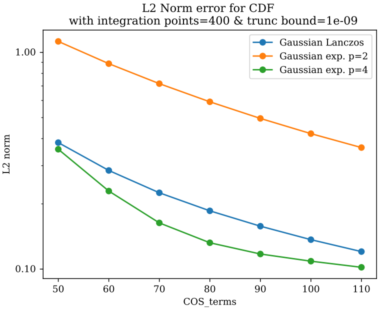

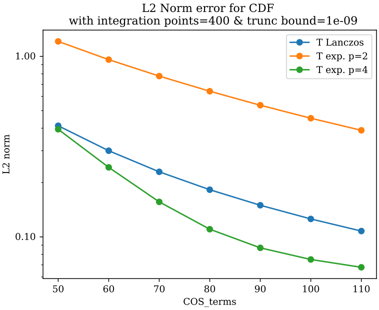

Left: Gaussian copula. Right: T copula.

VaR of Portfolio Loss

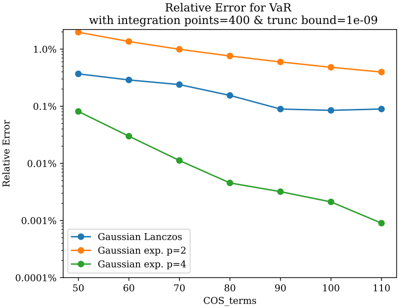

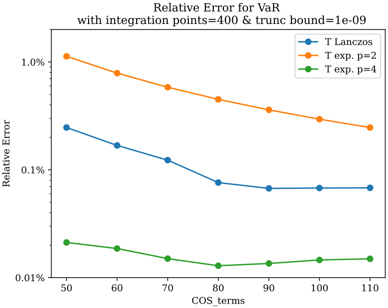

Left: Gaussian copula. Right: T copula.

ES of Portfolio Loss

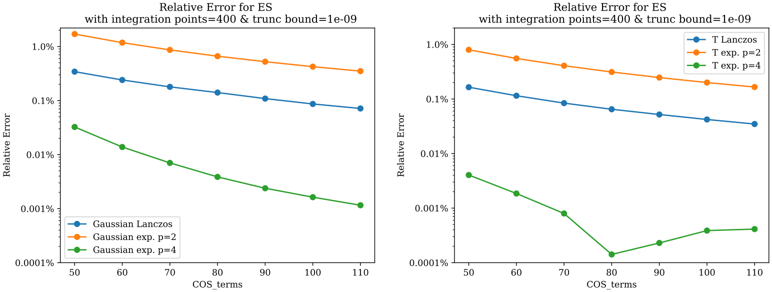

Left: Gaussian copula. Right: T copula.

CES of Portfolio Loss

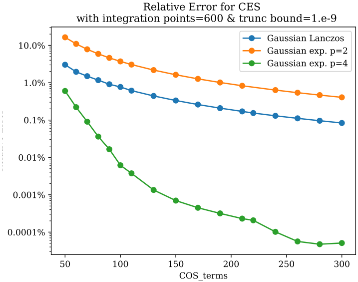

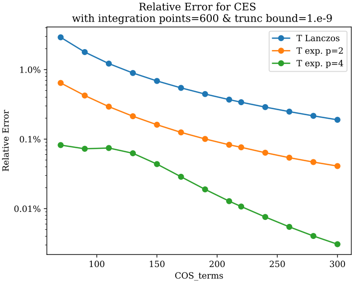

Left:Gaussian copula. Right: T copula.

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 in the Fourier domain with the help of the COS method.
- For dimension less than 4, this method is possible for real-time calculation usage, 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.