A Copula-Based Simulation Model for Supply Portfolio Risk

Halis Sak*
Yeditepe University

Çağrı Haksöz†
Sabancı University

May 24, 2011

Abstract

We introduce a copula-based simulation model for supply portfolio risk in the presence of dependent breaches of contracts. We demonstrate our method for a supply chain contract portfolio of commodity metals traded at the London Metal Exchange (LME). The analysis of spot price data of six LME commodity metals gives us the motive to use a t-copula dependence structure with t and generalized hyperbolic marginals for the log-returns. We also provide efficient simulation algorithms using importance sampling for the normal and t-copula dependence structure to quantify risk measures, supply-at-risk (SaR) and conditional supply-at-risk (cSaR). Numerical examples on a portfolio of

*Assistant Professor of Systems Engineering, Kayışdağı, 34755 İstanbul, Turkey. Phone:+90-216-578 04 50. Email:halis.sak@gmail.com.

†Assistant Professor of Operations Management, Sabancı School of Management, Orhanlı, Tuzla, 34956 İstanbul, Turkey. Phone:+90-216-483-9683. Email:cagrihaksoz@sabanciuniv.edu.

six commodity metals demonstrate that our proposed method succeeds in decreasing the variance of the simulations. A numerical sensitivity analysis for the choice of the copula function is also provided. To our knowledge, this is the first paper proposing efficient simulation algorithms on a supply chain contract portfolio having a copula-based dependence structure with generalized hyperbolic marginals.

Keywords: Breach of contract risk, Supply chain contracts, Procurement, Copula, Dependence, Importance sampling, Commodity metals

1 Introduction

Managing risks in global supply chains is getting more difficult due to increasing volatility and interdependence. Commodity price risk is significantly important for firms that consume various commodity metals in their operations. A recent McKinsey CEO survey report by Gyorey et al. (2010) notes that 37% of the CEO respondents state that in the next five year period, they would not be prepared for the increasing volatility of commodity prices. Moreover, commodity price risk is exacerbated in the presence of breach of contract risk (Haksöz and Şimşek (2010)). A breach of contract risk is a fundamental operational risk classified under "Clients, Products, and Business Processes" as well as the "Execution, Delivery, and Process Management" categories of the Basel II framework (See for example Cruz (2002), Chernobai et al. (2002), Haksöz and Kadam (2009) and Haksöz and Şimşek (2010) for details on this type of operational risk). A breach of contract may occur due to several reasons. It may be intentional such that a supplier may prefer to take advantage of favorable spot market price instead of selling via fixed-price contract. Surely, firms do pay penalty charges in case they breach contracts, which may somewhat compensate the financial loss for the other party. Yet, reputations are tarnished and strategic alliances are broken. There is certainly a need to assess the potential severity of breach of contract risk.

In contrast to the interest of the practice, there is scant amount of research activity in this area of operational risk that addresses the breach of contract risk and methods to assess and hedge it. In a single buyer-single supplier model, Haksöz and Seshadri (2007) valued an American type abandonment option, which models the breach of contract for a supplier, written in a long term contract with a fixed penalty. To assess the portfolio risk of various commodity supply chain contracts, Haksöz and Kadam (2009) provided a supply portfolio risk measurement tool based on the celebrated CreditRisk+ model. Haksöz and Kadam (2009) coined the term supply-at-risk (SaR) and also presented risk metric computations for a supplier portfolio of petrochemicals. In Haksöz and Kadam (2009), breach of contract was due to spot price evolution. However, Haksöz and Kadam (2009) has not addressed the dependency issue among multiple breach of contracts in the portfolio. On the other hand, Wagner et al. (2009) presented a model for correlated supplier defaults (due to many financial-economic factors, not only breach of contracts) with a copula dependence structure. Most recently, Haksöz and Şimşek (2010) provided a model to price bundled options (abandonment and price renegotiation option) in a supply chain contract. This type of bundled option is shown to be valuable to mitigate the breach of contract risk.

In this paper, building on the setting of Haksöz and Kadam (2009), we contribute by providing an efficient simulation method for supply portfolio risk assessment, where the supply chain contracts in the portfolio have a dependence structure. The efficient simulation method is borrowed from Sak et al. (2010) and it is modified for our problem. Moreover, the given algorithm in Sak. et. al. (2010) is designed for the t-copula dependence structure, we modify it so that it works for the normal copula as well. To demonstrate the value of our method, we study in detail a supply contract portfolio with a copula dependence structure, which is composed of a number of commodity metals that are traded at London Metal Exchange (LME). We further provide efficient algorithms in order to compute risk metrics such as supply-at-risk (SaR) and conditional supply-at-risk (CSaR).

Our paper is organized as follows. Section 2 presents the mathematical details on the supply portfolio model. Section 3 conducts the marginal distribution and copula fitting to commodity metal data. Sections 4 and 5 present the efficient simulation algorithms with importance sampling for tail loss probabilities and conditional expectations which are used for calculating SaR and CSaR respectively. Then, we present our numerical results and managerial insights on the commodity metal portfolio in Section 6. Finally, we conclude in Section 7.

2 The Model

We assume that a buyer procures a variety of metal commodities from different global suppliers using long term fixed price contracts. These commodity metals are also traded at London Metal Exchange (LME). The market prices are known and there are liquid spot markets. During the contract duration, the suppliers can breach their supply contracts for any reason. Actual breach of contract event is exogenous and not modeled in this paper.[1] Moreover, multiple dependent breaches of contracts can occur at the same time. We also assume that the buyer has to go to the spot market in case there is breach of contract by the suppliers. That is, the buyer does not have alternative suppliers for the specific commodities purchased in this portfolio apart from the spot market option. Even if there are potential backup suppliers, the price quoted for such emergency orders would closely follow the spot market price at that particular time. To that end, the buyer will be exposed to multiple spot market price risks under several dependent breaches of contracts. Hence, the buyer needs to assess the supply-at-risk (SaR) and conditional supply-at-risk (CSaR) for such supply contract portfolios in order to better manage its breach of contract risks.

Following the mathematical setting of Haksöz and Kadam (2009), we assume that the long term contract price is equal to the median metal spot price in this supply portfolio without loss of generality. Besides, we also assume that there is a fixed penalty paid by the suppliers in case of breach of contract and this penalty covers the transaction costs required to purchase commodities from the spot market for the buyer.[2] Basically, these assumptions help

[1] In practice, one also may need to determine the explicit breach of contract probabilities. These probabilities may be affected by internal and external factors. Environmental disruptions such as earthquakes, floods, hurricanes may cause suppliers breach their contracts. On the other hand, a firm may go bankrupt and cannot fulfill the contract, thus has to breach it. This nontrivial problem is left for future research.

[2] As pointed out by a referee, buyers can also design floating penalty contracts that may mitigate

us delineate the impact of spot price risk in a portfolio of contracts without considering the actual penalty and transaction cost data. Thus, we can write the individual risk exposure for the buyer at breach of contract as follows:

$$\epsilon_i = \max\{0, Q_i(P_i - \bar{P}_i)\}, \tag{1}$$

where Q_i is the contracted quantity of the metal i, P_i is the spot price for metal i, and $\bar{P}_i$ is the median spot price for metal i. We use the median spot price as a proxy for the long term contract price of the metal procured.

For a number of contracts in the portfolio, $i = 1, ..., n$, the total risk exposure for a number of potential breach of contracts can be expressed as follows:

$$R = \sum_{i=1}^{n} \epsilon_i, \tag{2}$$

where ϵ_i is given in (1).

In this expression, we only have the financial impact at breach, that is the severity of the breach events. Note that this severity is driven only by the spot price risks. We assume that the log-returns of n metals over a day follow an elliptical copula and its dependence structure is described by the positive definite matrix Σ; L denotes the (lower triangular) Cholesky factor of Σ satisfying $L\,L' = \Sigma$. We consider only the normal and t-copula alternatives for the elliptical copula function. We give the model, algorithms and numerical results for the normal and t-copula together for saving space. While writing the model and algorithms, we only give the differences between the normal and the t-copula. Classical random return vector generation algorithm from the normal and t-copula starts with a vector Z of d iid. standard normal variates that is then transformed into the correlated normal vector $\tilde{Z} = L\,Z$ (for a different generation algorithm for the t-copula, see Hörmann and Sak (2010)). For the t-copula, we obtain the vector T from the multivariate t-distribution by generating a random variate Y from chi-square distribution

the price risk. However enforcing such a contract would be harder.

with ν degrees of freedom (χ^2_ν) and calculating $T = \tilde{Z}/\sqrt{Y/\nu}$. The log-return vector $S = (S_1, S_2, \ldots, S_n)'$ is then the result of the component-wise transform

$$S_i = c_i \, G_i^{-1}(F(V_i)), \tag{3}$$

where F denotes the cumulative distribution function (CDF) of standard normal distribution for the normal copula and CDF of t-distribution with ν degrees of freedom for the t-copula. We have $V_i = \tilde{Z}_i$ for the normal copula and $V_i = T_i$ for the t-copula. G_i denotes the CDF of the marginal distribution of the return of the i-th metal and c_i the volatility scaling parameter of the i-th metal defined in (5) below. Then, given that P_{i0} is the spot price for metal i at time 0, spot price of metal i at the end of time horizon of m days is

$$P_i = P_{i0} \prod_{j=1}^{m} e^{S_i^{(j)}} \qquad \text{with} \qquad S_i^{(j)} = c_i \, G_i^{-1}(F(V_i^{(j)}))\,, j = 1, ..., m, \tag{4}$$

where c_i denotes a scaling factor related to the daily volatility σ_i and the variance var_i of the i-th marginal distribution by the formula

$$c_i = \sigma_i \sqrt{\frac{1}{var_i}}\,. \tag{5}$$

First, we are interested in the supply-at-risk (SaR) which is the quantile of the total risk exposure given in (2) for a probability level. To compute SaR for a time horizon, we require an efficient simulation algorithm to compute $P(R > x)$ for different values of x (may be simultaneously). Then inversion of tail loss probability distribution can be used to compute SaR. Second, we also provide an algorithm to compute the conditional expectations which can be used to compute supply-at-risk (CSaR) given that we have SaR.

3 Marginal Distribution and Copula Fitting to Commodity Metal Data

We use Inference Functions for Margins method to fit a dependence structure between log-returns of LME daily metal cash price data[3] as it is a simple and efficient method (for other parametric and nonparametric possibilities, see Malevergne and Sornette (2006) and Karadağ (2008)). In this method, as the first step, the parameters for marginal distributions are estimated using likelihood maximization, then the parameters of the copula are estimated using again maximum likelihood method and the estimated marginal distributions in the first step.

Given that T daily log-returns are available for metal i, the log-likelihood maximization problem for the first step is

$$\max_{\beta_i} \sum_{t=1}^{T} \ln(f_i(x_i^t; \beta_i)), i = 1, \ldots, n, \tag{6}$$

where f_i is the probability density function and β_i is the parameter vector for candidate distribution. As candidate distributions, we try three different alternatives: Gaussian, t distribution with location and scale, and the generalized hyperbolic distribution. The number of parameters that needs to be estimated for these continuous distributions are 2, 3, and 5 in the given order. t distribution is a natural candidate distribution as it is a simple extension of Gaussian distribution. On the other hand, our motive to try the generalized hyperbolic distribution comes from the field of finance where the most flexible and best fitting distribution to financial data seems to be generalized hyperbolic distribution (See Aas and Haff (2006), Behr and Pötter (2009), Prause (1997)).

[3]LME historical metal price data for the current year is freely available at www.lme.com.

For the second step, the log-likelihood function of the copula is

$$\max_{\alpha} \sum_{t=1}^{T} \ln c(F_1(x_1^t; \hat{\beta}_1), \ldots, F_n(x_n^t; \hat{\beta}_n); \alpha), \tag{7}$$

where F_i is the cumulative distribution function of marginal distribution for metal i, c is the density of copula function and α is the parameter vector for the copula. We consider only the normal and t-copula alternatives as Kole et al. (2007) conclude that the t-copula is the best fitting copula for the risk management of linear asset portfolios in finance among other alternatives such as the normal and Gumbel copula.

We use R (R Development Core Team, 2008) as a convenient working environment for solving (6) and (7), and for carrying out our simulations in the next sections. We use R packages fitdistrplus Marie Laure Delignette-Muller and Dutang (2010) for fitting Gaussian and t distribution, ghyp Breymann and Lüthi (2008) for fitting the generalized hyperbolic distribution and copula Yan and Kojadinovic (2010) for fitting copulas.

The best fitting criteria for marginal distributions and copulas is the magnitude of log-likelihood values. However, since it does not account for the estimated number of parameters, we also look at AKAIKE Information Criterion (AIC) values. The higher the log-likelihood value and the smaller the AIC value, the better the fit is. AIC value is calculated in fitdistrplus, and ghyp package as in Matteis (2001):

$$AIC = -2 \times LogLik + 2 \times NE, \tag{8}$$

where $LogLik$ denotes the log-likelihood value and NE denotes the number of estimated parameters.

We use daily LME spot prices of Copper (Cu), Aluminum (Al), Nickel (Ni), Zinc (Zn), Lead (Pb) and Tin (Sn) for the year 2010 given in Figure 1. First eleven months of the data were used to fit the dependence structure and marginal distributions. Last month's data are used for measuring the goodness

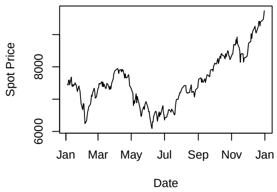

(a) Copper

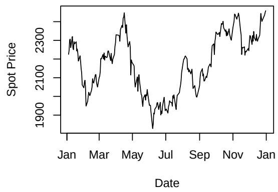

(b) Aluminum

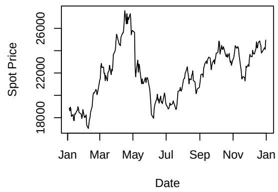

(c) Nickel

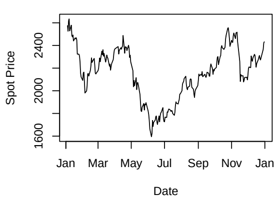

(d) Zinc

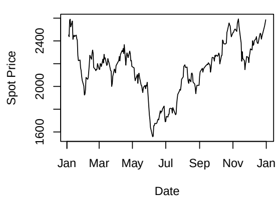

(e) Lead

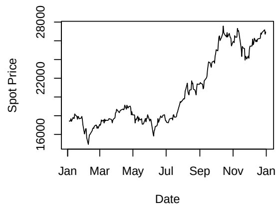

(f) Tin

Figure 1: Commodity metal spot prices in 2010.

of fit of the marginal distributions. The correlation matrix of the log-returns is given in Table 1. The maximum linear correlation is 0.844, which is between Copper and Zinc. The minimum one is 0.496, which is between Lead and Tin.

We should note that the number of metals and duration of the data is quite limited to derive conclusions for all of the metal spot markets. In this section, our aim is to see whether there is a tendency of metal data to deviate from normal distribution. If so, what may be the best distribution fit? And, is the t-copula dependence structure fitting well?

Shapiro-Francia (SF), Anderson-Darling (AD), Cramer-Von Mises (CVM), Lilliefors, and Pearsons chi-square tests are applied to test the normality of metal data using built-in functions in nortest Gross (2006) R package (see, Ricci (2005)). The estimated p-values of the test statistics are given in Table 2. Based on a significance level of 0.05 null hypothesis stating that data comes from normal distribution is rejected in most of the tests. Shapiro-Francia, Anderson-Darling, and Cramer-Von Mises tests give more weight to the tails than does Lilliefors, and Pearsons chi-square tests. p-values suggest that log-returns for Copper could be assumed normal although Shapiro-Francia, Anderson-Darling, Cramer-Von Mises oppose this. However, semi-heavy tails of log-returns of the other metals assert that the log-returns do not follow normal distribution.

Table 1: Correlation matrix of the metal log-returns

	Cu	Al	Ni	Zn	Pb	Sn
Cu	1.000	0.785	0.697	0.844	0.595	0.678
Al	0.785	1.000	0.685	0.753	0.606	0.621
Ni	0.697	0.685	1.000	0.664	0.553	0.583
Zn	0.844	0.753	0.664	1.000	0.679	0.657
Pb	0.595	0.606	0.553	0.679	1.000	0.496
Sn	0.678	0.621	0.583	0.657	0.496	1.000

Table 2: p-values for five normality tests for metal log-returns

Metal	SF	AD	CVM	Lilliefors	Pearson
Cu	0.043	0.026	0.028	0.093	0.241
Al	0.001	0.010	0.020	0.118	0.408
Ni	0.000	0.002	0.002	0.011	0.067
Zn	0.002	0.001	0.001	0.003	0.130
Pb	0.000	0.000	0.000	0.002	0.509
Sn	0.000	0.001	0.002	0.002	0.179

Motivated by the non-normality of the log-returns, we fit t and generalized hyperbolic distributions. Log-likelihood and AIC values for Gaussian, t and the generalized hyperbolic distribution fits are given in Table 3. In all the cases, the generalized hyperbolic distribution produces the highest log-likelihood value. However, it is a five-parameter distribution. Thus, AIC values are not always the minimum. It is better to use the generalized hyperbolic distribution for first the four metals (Cu, Al, Ni, Zn) and for the last two (Pb, Sn), it is better to use t distribution. Estimated parameters of the fitted marginal distributions for the log-returns and p-values for Kolmogorov-Smirnov test (denoted as KS) for measuring the goodness of fits are given in Table 4. We use log-return data for the month December to compute Kolmogorov-Smirnov test statistics since it can not be used when parameters of the distributions need to be estimated from the data. Anderson-Darling, Cramer-Von Mises, and Lilliefors are all modifications of Kolmogorov-Smirnov test. Computed p-values for Kolmogorov-Smirnov test are greater than 0.05 which leads to the conclusion that we cannot reject the null hypothesis that log-returns in December come from those marginal distributions. We use alpha/delta parametrization (see Breymann and Lüthi (2008)) for the generalized hyperbolic distribution to print the estimated parameters as we use this parametrization in our simulation functions. For t distribution, there are three parameters estimated, i.e., location, scale and degrees of freedom (df).

The histograms of the log-returns with the fitted t and generalized hyperbolic distributions, and Q-Q plots for only Copper and Aluminum are given in the Appendix.[4] Log-returns for Copper seems to be quite close to normal distribution. This visual observation is consistent with the tabulated p-values in Table 2. However, for the other five metals, t and the generalized hyperbolic distribution capture the high kurtosis and fat tails of the data.

We fit elliptical copulas, the normal and t-copula, to the log-returns of the metal data using the estimated marginals. As the dimension of the portfolio increases, the expression of the probability density functions for the Archimedean copulas become more complex and thus the probability density function is not available due to intensive computing involved in differentiating the cumulative distribution function (see Yan (2007) and Karadağ (2008)). Moreover, the t-copula is preferred to Gaussian and Gumbel copulas because of capturing the dependence better in the non-extremes and extremes (tails) of financial returns.

[4]Plots for the rest of the commodity metals are available from the authors upon request.

Table 3: Log-likelihood (LogLik) and AIC values for Gaussian, t and generalized hyperbolic (GH) distributions for metals.

Metal	Gaussian		t		GH		Which to use?
	LogLik	AIC	LogLik	AIC	LogLik	AIC	
Cu	606.93	−1209.86	607.14	−1208.28	610.24	−1210.49	GH
Al	615.82	−1227.65	619.04	−1232.08	622.57	−1235.15	GH
Ni	543.16	−1082.33	549.12	−1092.24	555.03	−1100.05	GH
Zn	544.95	−1085.90	545.79	−1085.59	551.08	−1092.15	GH
Pb	525.11	−1046.21	539.64	−1073.28	541.25	−1072.49	t
Sn	587.92	−1171.84	592.63	−1179.25	594.51	−1179.02	t

Table 4: Parameters of the fitted marginal distributions for metal log-returns and p-values for Kolmogorov-Smirnov test (KS).

Metal	GH					t			KS
	λ	α	δ	β	μ	Location	Scale	df	
Cu	7.1683	254.62	0.0002	−83.39	0.0212				0.051
Al	−2.8473	121.55	0.0399	−62.05	0.0146				0.365
Ni	0.6901	269.48	0.0421	−188.01	0.0479				0.630
Zn	1.5894	123.22	0.0259	−48.81	0.0192				0.532
Pb						0.00084	0.0187	4.69	0.148
Sn						0.00228	0.0154	5.46	0.294

For copula fitting we use built-in functions in copula Yan and Kojadinovic (2010) R package. Fitting results are summarized in Table 5. The correlation matrix and standard error of the point estimates are given for normal and t-copulas in Tables 6 and 7 in the given order. Numerical results given in Table 5 suggest that the t-copula is better than the normal copula in capturing the dependence structure of metal log-returns.

Table 5: Results of copula fitting for a portfolio consisting of all metals

Copula	Parameter(s)	SE	LogLik	AIC
Normal	ρ_{norm}	$SE_{\rho_{norm}}$	505.64	−981.28
Student-t	ρ_t, υ=11.53	SE_{ρ_t}, SE_υ=2.45	523.86	−1015.72

Table 6: Correlation matrix of the fitted normal copula (ρ_{norm}) for metal log-returns. Standard errors ($SE_{\rho_{norm}}$) are given in parentheses.

	Cu	Al	Ni	Zn	Pb	Sn
Cu	1.000	0.774(0.021)	0.683(0.029)	0.834(0.015)	0.619(0.034)	0.670(0.030)
Al	0.774	1.000	0.667(0.030)	0.741(0.023)	0.620(0.034)	0.611(0.034)
Ni	0.683	0.667	1.000	0.641(0.032)	0.569(0.038)	0.573(0.037)
Zn	0.834	0.741	0.641	1.000	0.703(0.027)	0.657(0.031)
Pb	0.619	0.620	0.569	0.703	1.000	0.524(0.041)
Sn	0.670	0.611	0.573	0.657	0.524	1.000

To conclude this section, according to the empirical results from the limited data, the t-copula with t and the generalized hyperbolic marginals seems to be an adequate model to capture the dependencies and explain the semi heavy tails of the returns of commodity metal data.

Table 7: Correlation matrix of the fitted t-copula (ρ_t) for metal log-returns. Standard errors (SE_{ρ_t}) are given in parentheses.

	Cu	Al	Ni	Zn	Pb	Sn
Cu	1.000	0.779(0.023)	0.700(0.030)	0.833(0.017)	0.675(0.033)	0.659(0.034)
Al	0.779	1.000	0.672(0.032)	0.741(0.026)	0.662(0.033)	0.602(0.039)
Ni	0.700	0.672	1.000	0.653(0.034)	0.599(0.039)	0.574(0.041)
Zn	0.833	0.741	0.653	1.000	0.747(0.026)	0.647(0.035)
Pb	0.747	0.662	0.599	0.747	1.000	0.559(0.042)
Sn	0.659	0.602	0.574	0.647	0.559	1.000

4 Simulating Tail Loss Probabilities with Importance Sampling (IS)

Algorithm 1 gives all the details of the naive simulation algorithm necessary to evaluate the tail loss probability $P(R > x)$ for time horizon of m days.

We modify the importance sampling (IS) technique described in Sak et al. (2010) for our problem. To summarize the technique: we add a mean shift vector with positive entries to the normal vector Z for the normal copula and additionally use a scale parameter θ less than two for the Gamma (scale parameter two corresponds to chi-square distribution) random variate Y in order to increase the probability of very high returns for the t-copula. The main practical problem in the application of IS is the choice of the parameters of the IS distribution. We only give the algorithms here. For a better understanding of the technique refer to Sak et al. (2010).

We use R package Runuran Leydold and Hörmann (2008) for evaluating quantiles from the generalized hyperbolic distribution (see, Section 3 of Sak et al. (2010)). Runuran uses a numerical inversion algorithm that requires only the probability density function instead of the cumulative density function (see, Derflinger et al. (2010) and Derflinger et al. (2009)).

Algorithm 3 returns the optimal mean-shift μ for the normal copula and also the optimal mode y_0 of the IS density for Y for the t-copula. Note that function $R()$ used in Step 2 of Algorithm 2 denotes for the total risk exposure defined in (2). Following equation gives the optimal scale parameter θ of the gamma IS density for Y

$$\theta = \frac{y_0}{\nu/2 - 1}. \tag{9}$$

The likelihood ratio for the normal copula is

$$W_\mu(Z) = \exp\left(-\mu' Z + \mu'\mu/2\right), \tag{10}$$

and for the t-copula it is

$$W_{\mu,\theta}(Z, Y) = \exp\left(-\mu' Z + \mu'\mu/2 - Y/2 + Y/\theta + \log(\theta/2)\nu/2\right), \tag{11}$$

where $\exp(-\mu' Z + \mu'\mu/2)$ accounts for the mean shift we have added to the normal vector and the term $\exp(-Y/2 + Y/\theta + \log(\theta/2)\nu/2)$ relates the density of chi-square distribution with degrees of freedom ν to that of gamma distribution with shape parameter $\nu/2$ and scale parameter θ. The final IS algorithm is presented as Algorithm 4.

The SaR associated with probability $1 - \alpha$ is the quantile

$$SaR_\alpha = \inf\{x : P(R > x) \leq \alpha\}. \tag{12}$$

Algorithm 4 can be used to simulate tail loss probabilities for various threshold levels to derive probability distribution of total risk exposure. Then a regression algorithm can be used to calculate SaR_α.

5 Simulating Conditional Expectations with Importance Sampling (IS)

In this section, we tackle the problem of simulating conditional expectation $E[R|R > x]$.

If we assume that $P(R > x) > 0$, $E[R|R > x]$ can be written as

$$r = E[R|R > x] = \frac{E\left[R\,\mathbf{1}\{R > x\}\right]}{P(R > x)}\,, \tag{13}$$

where $\mathbf{1}\{\}$ is an indicator function.

The naive simulation estimate for this ratio is

$$\hat{r}^{naive} = \frac{\sum_{k=1}^{N} R^{(k)}\mathbf{1}\{R^{(k)} > x\}}{\sum_{k=1}^{N}\mathbf{1}\{R^{(k)} > x\}}\,. \tag{14}$$

To estimate the accuracy of (14), we use $\delta\%$ confidence interval (see, Glasserman (2005) and Glasserman (2004))

$$\hat{r}^{naive} \pm z_{\delta/2}\frac{\hat{\sigma}^{naive}}{\sqrt{N}} \tag{15}$$

where

$$\hat{\sigma}^{naive} = \left(\frac{N\sum_{k=1}^{N}\left(R^{(k)} - \hat{r}^{naive}\right)^2\mathbf{1}\{R^{(k)} > x\}}{\left(\sum_{k=1}^{N}\mathbf{1}\{R^{(k)} > x\}\right)^2}\right)^{1/2} \tag{16}$$

and $z_{\delta/2}$ denotes the quantile of the standard normal distribution for the probability level of $\delta/2$.

Algorithm 1 gives all the details of how to use the naive simulation estimate.

Following Glasserman (2005) and Sak and Hörmann (2011), we use the importance sampling distribution computed for the problem ($P(R > x)$) in simulating $E[R|R > x]$. The IS simulation estimate and its $\delta\%$ confidence interval (see, Glasserman (2005) and Glasserman (2004)) is as follows:

$$\hat{r}^{IS} = \frac{\sum_{k=1}^{N} R^{(k)}W^{(k)}\mathbf{1}\{R^{(k)} > x\}}{\sum_{k=1}^{N} W^{(k)}\mathbf{1}\{R^{(k)} > x\}} \tag{17}$$

and

$$\hat{r}^{IS} \pm z_{\delta/2}\frac{\hat{\sigma}^{IS}}{\sqrt{N}} \tag{18}$$

where

$$\hat{\sigma}^{IS} = \left(\frac{N \sum_{k=1}^{N} \left(R^{(k)} W^{(k)} - \hat{r}^{IS} W^{(k)} \right)^2 \mathbf{1}\{R^{(k)} > x\}}{\left(\sum_{k=1}^{N} W^{(k)} \mathbf{1}\{R^{(k)} > x\} \right)^2} \right)^{1/2} \tag{19}$$

and $z_{\delta/2}$ denotes the quantile of the standard normal distribution for the probability level of $\delta/2$.

The details of how to use this estimate are presented as Algorithm 4. Using this algorithm we can compute $cSaR_\alpha = E[R|R > SaR_\alpha]$.

6 Numerical Results

We use the importance sampling algorithms given in Sections 4 and 5 for simulating the total risk exposure of the metal portfolio analyzed in Section 3. The log-returns following the t-copula with generalized hyperbolic or t-marginals is the dependence structure the data suggest. We use the fitting results summarized in Tables 5 and 7 for the t-copula. The correlation matrices (Σ that we use in simulation algorithms) for the normal and t-copula are given in Tables 6 and 7. For the numerical results presented we use $Q_i = 1, i = 1, ..., n$.

The efficiency of a simulation method is inversely proportional to the product of the sampling variance and the required simulation time. We therefore report as a main result of our comparison efficiency ratio (E.R.), the ratio of the product of the sampling variance and the execution time of the naive (NA) and the importance sampling method (IS).

We use IS algorithm given in Algorithm 4 to compute tail loss probabilities for various threshold levels. Then we fit a cubic smoothing spline to tail loss probabilities versus thresholds data in order to compute SaR_α for a number of α values in Table 8. We use R package fields Furrer et al. (2010) for fitting a cubic smoothing spline to the data. We also provide $cSaR_\alpha$ values and the half length of the 95% confidence intervals in percent (95% C.I.) for $cSaR_\alpha$. E.R.'s indicate the relative efficiency of the IS with respect to the naive simulation in computing $cSaR_\alpha$ values. E.R.'s increase as the event simulated becomes rarer. This is an attribute of importance sampling. Execution times are 17.0, 19.7 seconds for naive (NA) and the IS in the given order for computing $cSaR_\alpha$ values for $N = 100,000$. Furthermore, efficiency of the IS in computing tail loss probabilities with respect to the naive simulation is quite similar to the presented E.R.'s for $cSaR_\alpha$.

Table 8: Over a one day horizon SaR_α, $cSaR_\alpha$ values and $cSaR_\alpha$'s 95% confidence interval as percentage of the point estimates for the naive and the IS.

α	SaR_α	$cSaR_\alpha$	95% C.I.(NA)	95% C.I.(IS)	E.R.
0.05	$9,839.5$	$10,341.6$	$\pm 0.14\%$	$\pm 0.04\%$	10
0.01	$10,635.7$	$11,182.3$	$\pm 0.34\%$	$\pm 0.05\%$	46
0.005	$10,989.1$	$11,573.8$	$\pm 0.48\%$	$\pm 0.05\%$	75
0.002	$11,489.6$	$12,147.0$	$\pm 0.74\%$	$\pm 0.06\%$	148
0.001	$11,902.7$	$12,638.7$	$\pm 1.31\%$	$\pm 0.06\%$	382

It is important to asses the sensitivity of the numerical results given in Table 8 to the choice of the copula function while keeping the marginal distributions and the correlation matrix the same as suggested in Johnson and Tenenbein (1981) for a similar problem. We use different degrees of freedom for the t-copula and normal copula to see how the choice of the copula-based joint distribution affects the simulated results for the SaR_α and $cSaR_\alpha$ measures. Over a one day horizon, SaR_α and $cSaR_\alpha$ values for sets of α and degrees of freedom of the t-copula ($\nu = \infty$ is the normal copula case) are provided in Table 9. In particular simulated results change very little; the maximum difference in these results is 3.2% for the ten cases considered. The fact that these results change very little adds credibility to the measures that we developed.

As the degrees of freedom for the t-copula increase (approaching to the normal copula), SaR_α and $cSaR_\alpha$ decrease for the tails ($\alpha < 0.05$). This is an expected result since the tail dependence between contracts is lessening (probability distribution of total risk is less fat). However, as we approach to the center of the distribution, SaR_{α} and $cSaR_{\alpha}$ values increase as the dependence structure gets stronger. This is observed for $SaR_{\alpha=0.05}$.

To give a rough idea about how the tail loss probabilities and $cSaR_{\alpha}$ change in time, we draw tail loss probabilities and $cSaR_{\alpha}$ of total exposures for time horizons of one day and one week simultaneously in Figures 2 and 3. We use the IS and naive simulation for one day horizon and only the IS for one week horizon in computing the tail loss probabilities, $cSaR_{\alpha}$'s, and their confidence intervals. Efficiency of IS for one day horizon could be easily observed when we compare it with the naive simulation. Although we use greater number of replications for naive simulation, it gives wider confidence intervals and stops giving sensible confidence intervals for thresholds greater than 12, 500.

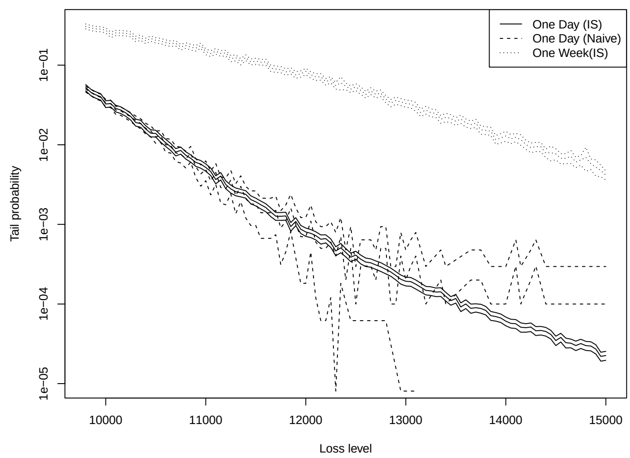

Figure 2: Tail loss probabilities of total risk exposure for time horizons of one day and one week using the IS (using 1, 000 replications) and naive (using 10, 000 replications). The three curves show the sample mean and its 95% confidence interval.

When we compare the tail loss probabilities for one day and one week hori-

Table 9: Over a one day horizon SaR_α and $cSaR_\alpha$ values for sets of α and degrees of freedom of the t-copula ($\nu = \infty$ is the normal copula case).

α	ν	SaR_α	$cSaR_\alpha$
	3	$9,826.7$	$10,374.5$
	5	$9,833.4$	$10,362.8$
0.05	10	$9,837.7$	$10,349.6$
	15	$9,839.6$	$10,345.2$
	∞	$9,840.0$	$10,326.1$
	3	$10,688.9$	$11,287.2$
	5	$10,666.7$	$11,247.4$
0.01	10	$10,640.1$	$11,194.8$
	15	$10,630.1$	$11,167.1$
	∞	$10,602.1$	$11,109.6$
	3	$11,078.3$	$11,715.9$
	5	$11,039.4$	$11,660.8$
0.005	10	$10,998.2$	$11,588.1$
	15	$10,979.6$	$11,559.2$
	∞	$10,933.5$	$11,475.7$
	3	$11,619.2$	$12,322.2$
	5	$11,567.3$	$12,260.6$
0.002	10	$11,499.1$	$12,160.4$
	15	$11,470.5$	$12,118.3$
	∞	$11,397.0$	$12,007.9$
	3	$12,065.3$	$12,845.7$
	5	$12,000.5$	$12,766.7$
0.001	10	$11,914.4$	$12,662.8$
	15	$11,876.2$	$12,591.3$
	∞	$11,774.6$	$12,448.6$

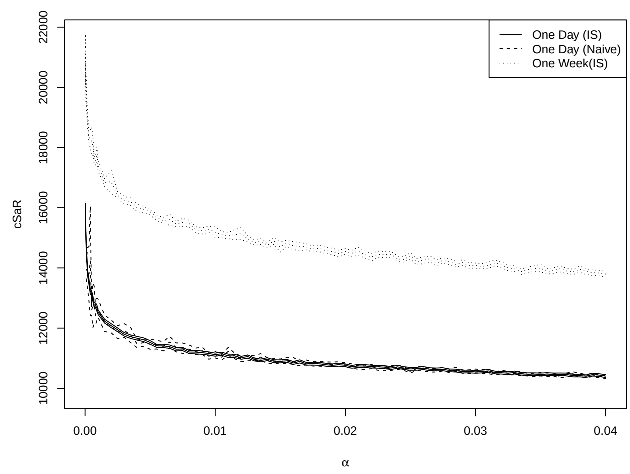

Figure 3: $cSaR_\alpha$'s for time horizons of one day and one week using the IS (using $1,000$ replications) and naive (using $10,000$ replications). The three curves show the sample mean and its 95% confidence interval.

zons, we observe that the tail loss probabilities increase as we increase the time horizon. For time horizon of one day, we use $m = 1$ in Algorithm 4. For one week horizon, we use $m = 5$ (five working days is equivalent to 1 week). Due to the complicated return function in Step 2 of Algorithm 2, as we extend the time horizon, the efficiency of the IS decreases as it can be observed in Figures 2 and 3. Although the wideness of confidence intervals for time horizons of one day and one week seem to be nearly the same in Figure 2, the tail loss probability axis is given in logarithmic scale in Figure 2. Indeed, the correct way of simulating the tail loss probabilities and conditional expectations for long horizons like one week horizon is to fit marginal distributions and copula for weekly instead of daily log-returns. Then, after adjusting the scaling factor (σ_i should be this time weekly volatility), presented algorithms could be used to simulate the tail loss probabilities and conditional expectations of total risk exposures for one week time horizon using $m = 1$.

For the same α level, one week horizon $cSaR_\alpha$ is greater than for one day horizon $cSaR_\alpha$ as in line with our observations in Figure 2. Furthermore, confidence intervals get worse in computing one day and one week horizons $cSaR_\alpha$'s as α decreases as it can be observed in Figure 3. As α decreases, the likelihood ratios for the IS decrease to make the losses equal to the threshold on average. This decrease in the likelihood ratios is responsible for degradation in the quality of confidence intervals for the IS.

7 Conclusion

In this paper, we introduced an efficient simulation model for quantifying the risk measures for supply portfolio risk in the presence of dependent breaches of contracts. The model is based on a copula dependence structure. For assessing model parameters, we analyzed a limited data set of London Metal Exchange commodity metal spot prices. This process revealed a better fit for the t-copula dependence structure with t or generalized hyperbolic marginal distributions for the log-returns of the metals. Furthermore, we adopted the importance sampling strategy given in Sak et al. (2010) to compute SaR and cSaR under the normal and t-copula dependence structure. Our numerical results showed that the proposed method is much more efficient than naive simulation for computing tail loss probabilities and conditional expectations. We also provided a numerical sensitivity analysis for the choice of the copula function.

We think that the method proposed in this paper could very well assist supply chain, procurement, and operational risk executives while assessing supply portfolio risk with dependence structure. Surely, future research is needed in developing models with explicit breach of contract probabilities and more sophisticated dependence structures. Furthermore, using real life penalty and transaction cost data would be useful to quantify the aggregate supply risk. We hope that our paper motivates more research in this growing field.

Acknowledgements

The authors are grateful to the editor-in-chief, associate editor, three anonymous referees, and Wolfgang Hörmann for their valuable suggestions that substantially improved the paper. They also acknowledge Institute for Statistics and Mathematics of the Vienna University of Economics and Business for letting them use their computing cluster.

References

Aas, K., Haff, I. H., 2006. The generalized hyperbolic skew student's t-distribution. Journal of Financial Econometrics 4 (2), 275–309.

Behr, A., Pötter, U., 2009. Alternatives to the normal model of stock returns: Gaussian mixture, generalised logf and generalised hyperbolic models. Annals of Finance 5 (1), 49–68.

Breymann, W., Lüthi, D., 2008. ghyp: A Package on Generalized Hyperbolic Distributions. R package version 1.5.4.

Chernobai, A. S., Rachev, S. T., Fabozzi, F. J., 2002. Operational Risk: A Guide to Basel II Capital Requirements, Models, and Analysis. John Wiley and Sons, Hoboken, NJ.

Cruz, M. G., 2002. Modeling, Measuring and Hedging Operational Risk. John Wiley and Sons, Hoboken, NJ.

Derflinger, G., Hörmann, W., Leydold, J., 2010. Random variate generation by numerical inversion when only the density function is known. ACM Transactions on Modeling and Computer Simulation 20 (4), 18:1–18:25.

Derflinger, G., Hörmann, W., Leydold, J., Sak, H., 2009. Efficient numerical inversion for financial simulations. In: L'Ecuyer, P., Owen, A. B. (Eds.), Monte Carlo and Quasi-Monte Carlo Methods 2008. Springer-Verlag, Heidelberg, pp. 297–304.

Furrer, R., Nychka, D., Sain, S., 2010. fields: Tools for spatial data. R package version 6.3.

Glasserman, P., 2004. Monte Carlo Methods in Financial Engineering. Springer.

Glasserman, P., 2005. Measuring marginal risk contributions in credit portfolios. Journal of Computational Finance 9 (2), 1–41.

Gross, J., 2006. nortest: Five Omnibus Tests for the Composite Hypothesis of Normality. R package version 1.0.

Gyorey, T., Jochim, M., Norton, S., 2010. Mckinsey global survey results: The challenges ahead for supply chains. McKinsey and Company.

Haksöz, Ç., Şimşek, K. D., 2010. Modeling Breach of Contract Risk Through Bundled Options. The Journal of Operational Risk 5 (3), 3–20.

Haksöz, Ç., Kadam, A., 2009. Supply Portfolio Risk. The Journal of Operational Risk 4 (1), 59–77.

Haksöz, Ç., Seshadri, S., 2007. Supply Chain Operations in the Presence of a Spot Market: A Review with Discussion. Journal of Operational Research Society 58 (11), 1412–1429.

Hörmann, W., Sak, H., 2010. t-copula generation for control variates. Mathematics and Computers in Simulation 81, 782–790.

Johnson, M. E., Tenenbein, A., 1981. A bivariate distribution family with specified marginals. Journal of the American Statistical Association 76 (373), 198–201.

Karadağ, D. T., 2008. Portfolio risk calculation and stochastic portfolio optimization by a copula based approach. Master's thesis, Boğaziçi University Istanbul.

Kole, E., Koedijk, K., Verbeek, M., 2007. Selecting copulas for risk management. Journal of Banking and Finance 31 (8), 2405–2423.

Leydold, J., Hörmann, W., 2008. Runuran – R interface to the UNU.RAN random variate generators, Version 0.8. Department of Statistics and Mathematics, WU Wien, A-1090 Wien, Austria.

Malevergne, Y., Sornette, D., 2006. Extreme Financial Risks: From Dependence to Risk Management. Springer.

Marie Laure Delignette-Muller, Regis Pouillot, J.-B. D., Dutang, C., 2010. fitdistrplus: Help to Fit of a Parametric Distribution to Non-censored or Censored data. R package version 0.1-4.

Matteis, R., 2001. Fitting copulas to data. Diploma thesis, Institute of Mathematics of the University of Zurich.

Prause, K., 1997. Modelling financial data using generalized hyperbolic distributions. FDM preprint 48, University of Freiburg.

R Development Core Team, 2008. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, `http://www.R-project.org`.

Ricci, V., 2005. Fitting Distributions with R. R Foundation for Statistical Computing, Vienna, Austria.

Sak, H., Hörmann, W., 2011. Fast simulations in Credit Risk, Quantitative Finance (accepted for publication).

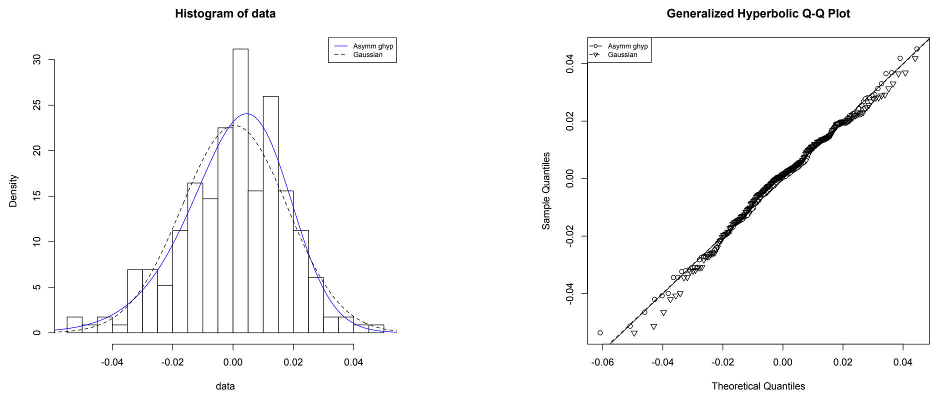

(a) (b)

Figure 4: Histogram and Q-Q plot for the log-returns of Copper.

Sak, H., Hörmann, W., Leydold, J., 2010. Efficient risk simulations for linear asset portfolios in the t-copula model. European Journal of Operational Research 202 (3), 802–809.

Wagner, S. M., Bode, C., Koziol, P., 2009. Supplier default dependencies: Empirical evidence from the automotive industry. European Journal of Operational Research 199, 150–161.

Yan, J., 2007. Enjoy the joy of copulas: With a package copula. Journal of Statistical Software 21, 1–21.

Yan, J., Kojadinovic, I., 2010. copula: Multivariate Dependence with Copulas. R package version 0.9-7.

8 Appendix

Algorithm 1 Computation of $P(R > x)$ and $E[R|R > x]$ using naive simulation for the normal and t-copula.

0. Initialization.
   (a) Compute Cholesky factor L of Σ, i.e., $L\,L' = \Sigma$.
   (b) Compute c_i, for $i = 1, \ldots, n$ using (5).
1. Repeat for replications $k = 1, \ldots, N$:
   (a) Repeat for replications $j = 1, \ldots, m$:
      (i) Generate independent standard normal variates Z then compute $\tilde{Z} = L\,Z$.
      (ii) Generate Y from χ^2_ν distribution for the t-copula.
      (iii) $V_i = \tilde{Z}_i$ for the normal copula and $V_i = \tilde{Z}_i/\sqrt{Y/\nu}$ for the t-copula, for $i = 1, \ldots, n$.
      (iv) Calculate $S_i^{(j)}$, for $i = 1, \ldots, n$ using using (3).
   (b) Calculate P_i, for $i = 1, \ldots, n$ using (4) then total risk exposure $R^{(k)}$ using (2).
2. Return $\frac{1}{N}\sum_{k=1}^{N} \mathbf{1}\{R^{(k)} > x\}$ for computing $P(R > x)$ and return $\hat{r}^{naive}$ using (14) for computing $E[R|R > x]$.

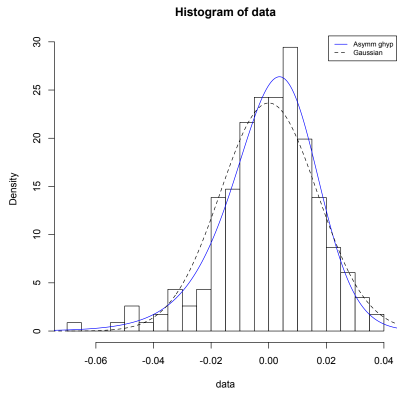

(a)

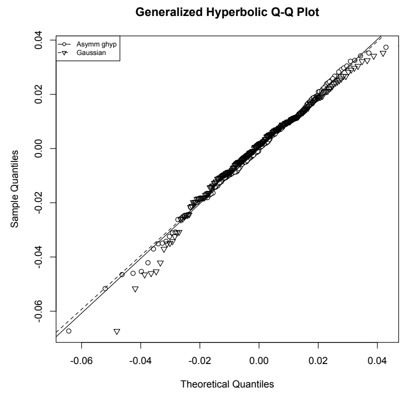

(b)

Figure 5: Histogram and Q-Q plot for the log-returns of Aluminum.

Algorithm 2 Computation of z_0, y_0, and o_f for a given direction z_d.

---

1. Set $z_d^1 = z_d/||z_d||$.
2. Compute r_0 by solving $(R(z = r_0 z_d^1) - (x + \Delta) = 0)$ numerically (Use, e.g., $\Delta = 10^{-5}$) for the normal copula and by solving $(R(z = r_0 z_d^1, y = \nu) - (x + \Delta) = 0)$ for the t-copula.
3. Return vector $z_0 = r_0\, z_d^1$, and objective function value $o_f = -r_0^2$ for the normal copula and vector $z_0 = r_0 \sqrt{y_0/\nu}\, z_d^1$, $y_0 = (\nu - 2)/(1 + r_0^2/\nu)$, and objective function value $o_f = (\nu/2 - 1)(\log(y_0) - 1)$ for the t-copula.

Algorithm 3 Computation of the mean shift vector μ and y_0.

0. Initialization.
   (a) Compute Cholesky factor L of Σ, i.e., $L L' = \Sigma$.
   (b) Compute c_i, for $i = 1, \ldots, n$ using (5).
1. Compute $z_d = L'c$.
2. Call an optimization algorithm with starting direction z_d, objective function as given in Algorithm 2, and non-negativity constraints for all components of z_d (we used a quasi-Newton method with constraints). Get optimal direction $_{opt}z_d$.
3. Call Algorithm 2 with direction $_{opt}z_d$ and get the optimal vector z_0 for the normal copula, z_0 and y_0 for the t-copula.
4. Return the optimal mean shift $\mu = z_0$ for the normal copula and $\mu = z_0$ and optimal mode y_0 for Y for the t-copula.

Algorithm 4 Computation of $P(R > x)$ and $E[R|R > x]$ using importance sampling for the normal and t-copula.

0. Initialization.
   (a) Compute Cholesky factor L of Σ, i.e., $L\,L' = \Sigma$.
   (b) Compute c_i, for $i = 1, \ldots, n$ using (5).
   (c) Compute μ for the normal copula and μ and y_0 for the t-copula using Algorithm 3.
   (d) Compute $\theta = y_0/(\nu/2 - 1)$ for the t-copula.
1. Repeat for replications $k = 1, \ldots, N$:
   (a) Repeat for replications $j = 1, \ldots, m$:
      (i) Generate $Z_i \sim N(\mu_i, 1), i = 1, .., n$, independently then compute $\tilde{Z} = L\,Z$.
      (ii) Generate Y from gamma distribution with shape parameter $\nu/2$ and scale parameter θ for the t-copula.
      (iii) Calculate $W_{\mu}^{(j)}$ for the normal copula as in (10) and $W_{\mu,\theta}^{(j)}$ as in (11) for the t-copula.
      (iv) $V_i = \tilde{Z}_i$ for the normal copula and $V_i = \tilde{Z}_i/\sqrt{Y/\nu}$ for the t-copula, for $i = 1, \ldots, n$.
      (v) Calculate $S_i^{(j)}$, for $i = 1, \ldots, n$ using using (3).
   (b) Calculate P_i, for $i = 1, \ldots, n$ using (4) then total risk exposure $R^{(k)}$ using (2).
   (c) Calculate $W^{(k)} = \prod_{j=1}^{m} W_{\mu}^{(j)}$ for the normal copula and $W^{(k)} = \prod_{j=1}^{m} W_{\mu,\theta}^{(j)}$ for the t-copula.
2. Return $\frac{1}{N}\sum\limits_{k=1}^{N} W^{(k)}\mathbf{1}\{R^{(k)} > x\}$ for computing $P(R > x)$ and return $\hat{r}^{IS}$ using (17) for computing $E[R|R > x]$.