

EFFICIENT ESTIMATION OF LARGE PORTFOLIO LOSS PROBABILITIES IN t -COPULA MODELS

Joshua C.C. Chan* Dirk P. Kroese†

The University of Queensland‡

November, 2009

Abstract

We consider the problem of accurately measuring the credit risk of a portfolio consisting of loans, bonds and other financial assets. One particular performance measure of interest is the probability of large portfolio losses over a fixed time horizon. We revisit the so-called t -copula that generalizes the popular normal copula to allow for extremal dependence among defaults. By utilizing the asymptotic description of how the rare event occurs, we derive two simple simulation algorithms based on conditional Monte Carlo to estimate the probability that the portfolio incurs large losses under the t -copula. We further show that the less efficient estimator exhibits bounded relative error. An extensive simulation study demonstrates that both estimators outperform existing algorithms. We then discuss a generalization of the t -copula model that allows the multivariate defaults to have an asymmetric distribution. Lastly, we show how the estimators proposed for the t -copula can be modified to estimate the portfolio risk under the skew t -copula model.

Keywords: Credit risk; copula models; rare-event simulation; cross-entropy method; conditional Monte Carlo

*Correspondent author. Department of Mathematics, University of Queensland, Brisbane, QLD 4072, Australia. Email: c.chan5@uq.edu.au.

†Email: kroese@maths.uq.edu.au

‡This work is supported by the Australian Research Council (Discovery Grant DP0985177).

1 Introduction

Losses resulting from the failure of an obligor to make a contractual payment, generally referred to as *credit risk*, are one of the major concerns of financial institutions. Consequently, the problem of accurately measuring the credit risk of a portfolio consisting of various financial assets has received considerable attention in the literature. Each obligor in the portfolio is subject to possible default, and such an event is often captured by the so-called *threshold models*, where a default occurs when a latent variable exceeds a given threshold. In order to model the dependence of simultaneous defaults observed empirically, a dependence structure is often imposed on the multivariate default distribution. The most popular choice of such a structure is the multivariate normal distribution. This gives rise to the celebrated *normal copula model*, which is widely used in the financial industry and forms the basis of the J. P. Morgan's CreditMetrics and other management systems (Gupton et al., 1997; Li, 2000). See also the monographs by Bluhm et al. (2002) and McNeil et al. (2005). Under the normal copula framework, dependence is often induced via a set of common factors affecting multiple obligors. These factors are typically interpreted as economy-wide risks, to which all the obligors are exposed, though to varying degrees. Conditional on these factors, the obligors then become independent. In terms of performance measures of credit risk, one that is of particular importance is the probability of large portfolio losses over a fixed time horizon. Since this probability is typically not available analytically, Monte Carlo methods are required to estimate this quantity.

To generate more scenarios with large losses in simulation, one common approach is to shift the factor mean via *importance sampling* (IS) (see, e.g., Rubinstein and Kroese, 2007), as suggested in, e.g., Kalkbrenner et al. (2004), Joshi (2004) and Egloff et al. (2005). Although this heuristic works well empirically in the context of single-factor normal copula models, there is little theoretical support, and consequently, the procedure might fail for certain sets of parameter values. Glasserman and Li (2005) derive logarithmic limits for the tail of the loss distribution associated with single-factor homogeneous portfolios. In particular, they show that for the regime with moderately high correlation among the obligors, the occurrence of large losses is determined primarily by the common factor, thus justifying the heuristics of shifting the factor mean. Moreover, they propose the following two-step IS procedure: first apply IS to shift the factor mean, then apply IS again conditional on the common factor affecting multiple obligors. They further show that the proposed estimator is logarithmically efficient. Although the utility of this two-step procedure is supported by both theoretical and numerical results, it is difficult to generalize the procedure to the general multi-factor model. In view of this difficulty, Glasserman et al. (2007) analyze the general multi-factor normal copula setting and derive logarithmic asymptotics for the loss distribution. The asymptotic results are later exploited in Glasserman et al. (2008) to develop logarithmically efficient IS techniques to estimate the tail probabilities of large portfolio losses. We refer the readers to the recent review in Grundke (2009) for other related approaches.

Despite its popularity, the normal copula model does not capture various stylized facts about financial variables brought forth by recent empirical research. In particular, one of the most prominent features of financial variables is that they exhibit *extremal dependence*, i.e., they are asymptotically dependent. Loosely speaking, the variables take on large values (in absolute terms) simultaneously with non-negligible probability, and it is not captured by the correlation structure implied by the multivariate normal distribution. In view of this inadequacy of the normal copula, Bassamboo et al. (2008) propose the *t-copula model*, based on the multivariate

t -distribution, that attempts to capture the relatively frequent occurrences of extremal comovements of financial variables. They argue that in many instances it is a more adequate way to model dependencies than the normal copula. The authors derive sharp asymptotics for the loss distribution, and show that under the t copula model, large portfolio losses occur primarily when the so-called *shock variable* takes on small values, while other random variables, including the common factors, are relatively unimportant in determining the occurrence of large losses. In other words, shifting the factor mean alone, as suggested by the aforementioned IS procedures, would not significantly increase the number of scenarios with large losses, and consequently, substantial variance reduction might not be achieved. Therefore, the authors propose two IS algorithms to estimate the probability of large portfolio losses. The first estimator uses IS based on an *exponential change of measure* (ECM) (see, e.g., Asmussen and Glynn, 2007) and has bounded relative error; the second uses a variant of *hazard rate twisting* (HRT) (Juneja and Shahabuddin, 2002), which is shown to be logarithmically efficient. An extensive simulation study shows that while both estimators offer substantial variance reduction, the former provides 6 to 10 times higher variance reduction than the latter. Nevertheless, the more efficient ECM algorithm involves generating random variables from a nonstandard distribution via rejection sampling, which takes on average three times more time compared to naive Monte Carlo simulation. In addition, the normalizing constant of the proposal density is not known, and has to be computed by numerical routines in order to be used in the likelihood ratio evaluation.

Instead of the two IS algorithms, we propose two novel estimators based on *conditional Monte Carlo* (see, e.g., Asmussen and Glynn, 2007; Rubinstein and Kroese, 2007) to estimate the probability of large portfolio loss under the t -copula model. We prove that the less efficient estimator has bounded relative error. A simulation study similar to that in Bassamboo et al. (2008) further shows that the proposed estimators outperform (in terms of variance reduction) both ECM and HRT algorithms. Moreover, the new algorithms involve only generating random variables from standard distributions, and consequently they are as efficient as naive simulation in terms of random variable generation effort. An additional advantage is that the new algorithms require trivial programming effort and are easier to implement than those proposed in Bassamboo et al. (2008), as the latter require generating random variables from nonstandard distributions. We then consider a generalization of the t -copula model to an asymmetric default distribution, as opposed to the symmetric distribution implied by the normal copula and t -copula models. This generalization is relevant and potentially important as it incorporates the well-documented observation that in practice financial variables are highly skewed (Fernandez and Steel, 1998; Franses and van Dijk, 2000). In a credit risk setting, for instance, there is relatively little potential gain when the underlying economic conditions improve, but there is a substantial downside risk when the market condition worsens. Consequently, the multivariate default distribution is expected to be positively skewed (since a large positive draw of the latent variable represents a default). Failure of taking this asymmetry into account might result in underestimation of the credit risk of the portfolio.

The rest of this article is organized as follows. In Section 2 we formulate the problem of estimating large portfolio losses and introduce the normal copula model. We then discuss the t -copula model in Section 3. Section 4 discusses two estimation methods based on conditional Monte Carlo for estimating the probability of large portfolio loss under the t -copula model. The performance of these estimators are demonstrated via an extensive simulation study in Section 5. Finally, we consider the skew t -copula model that accommodates an asymmetric default distribution. There we also study how the skewness of the multivariate default distribution affects the probability of large portfolio loss.

2 Problem Formulation

Consider a lender owning a portfolio of loans consisting of n obligors, each of whom has a positive, albeit small, probability of defaulting. Let the probability of default for the i th obligor be $p_i \in (0, 1)$, which we take as given. In practice, these probabilities can often be estimated by various econometrics models using historical data and other observed characteristics of the current obligors. We further assume that the monetary loss associated with the default of the i th obligor, denoted as c_i , is known. We introduce a vector of underlying latent variables $\mathbf{X} = (X_1, \dots, X_n)$ so that the i th obligor defaults if X_i exceeds some given threshold level x_i . More specifically, let $f_{X_i}(x)$ denote the (marginal) probability density function (pdf) of X_i . Given the probability of default p_i , the threshold x_i is determined implicitly by

$$\mathbb{P}_{f_{X_i}}(X_i > x_i) = \int_{x_i}^{\infty} f_{X_i}(u) du = p_i .$$

One could interpret the latent variable X_i as the underlying financial condition of the i th obligor, which is not directly observable to the lender. However, when the obligor's financial condition become worse than a critical level (x_i), she goes bankrupt and the lender observes a default in the i th loan. Our main interest is to learn about the distribution of the loss incurred from defaults

$$L(\mathbf{X}) = c_1 I_{\{X_1 > x_1\}} + \dots + c_n I_{\{X_n > x_n\}}, \quad (1)$$

where $I_{\{\cdot\}}$ denotes the indicator function. In particular, we wish to estimate accurately the probability of large losses of the form

$$\ell(\gamma) = \mathbb{P}_f(L(\mathbf{X}) > \gamma) \quad (2)$$

where $\mathbf{X} \sim f(\mathbf{x})$ and $\gamma = bn$ for some $b > 0$. In order to estimate the above probability, one needs to specify the joint distribution of the latent variables $\mathbf{X} = (X_1, \dots, X_n)$. It is obvious that the usefulness of the model depends critically on the distributional assumptions of the vector \mathbf{X} . On the one hand, the researcher wishes to make as few assumptions about the joint distribution as possible, since imposing restrictive but unrealistic assumptions often lead to misleading conclusions. On the other hand, a parameter-rich model often makes the analysis intractable. How this trade-off between flexibility and tractability is made is therefore of vital importance.

One popular model that is widely used in the financial industry is the normal copula model that forms the basis of the CreditMetrics and other related models. The normal copula model attempts to capture the dependence among obligors while maintaining mathematical tractability

by assuming the vector of latent variables follows a multivariate normal distribution. More specifically, the underlying correlations are often specified through a linear factor model

$$X_i = w_{i1}Z_1 + \cdots + w_{id}Z_d + w_i\eta_i, \quad i = 1, \dots, n, \quad (3)$$

where Z_1, \dots, Z_d are independent and identically distributed (iid) standard normal variables known as *factors* that capture the systemic risk common to all the obligors, w_{i1}, \dots, w_{id} are *factor loadings* that measure the sensitivity of i th obligor's financial position to each factor, η_i is a zeros mean normal random variable independent of the factors that captures the idiosyncratic risk of the i th obligor. Since each latent variable X_i is a linear combination of normally distributed random variables, the latent vector \mathbf{X} has a multivariate normal distribution where the correlation structure is induced by the linear factor model in (3).

3 The t -copula Model

One of the potential problems of the normal copula model is that it might assign too little probability to the event of many simultaneous defaults. In view of this, Bassamboo et al. (2008) introduce the t -copula model by assuming the underlying latent variables \mathbf{X} follow a multivariate t distribution. Following Bassamboo et al. (2008) we restrict our attention to the single-factor model ($d=1$) to keep the notations simple. It is important to note that there is no difficulty generalizing the model to a general d -factor model (see the discussion in Section 4). As in the normal copula model, the factors and the individuals' idiosyncratic risks are modeled as independent normally distributed random variables, i.e.,

$$Z \sim \mathbf{N}(0, 1), \quad \eta_i \sim \mathbf{N}(0, \sigma_\eta^2), \quad i = 1, \dots, n. \quad (4)$$

To induce a t structure, let us introduce a *shock variable* $\lambda > 0$ that is independent of Z and $\boldsymbol{\eta} = (\eta_1, \dots, \eta_n)$ such that $\lambda^2 \sim \text{Gamma}(\nu/2, \nu/2)$ for some $\nu > 0$. It is easy to check that the pdf of λ is given by

$$f_\lambda(x) = \frac{2\nu^{\nu/2}}{2^{\nu/2}\Gamma(\nu/2)} x^{\nu-1} e^{-\nu x^2/2}, \quad x > 0. \quad (5)$$

Next, define

$$X_i = \left(\rho Z + \sqrt{1 - \rho^2} \eta_i \right) \lambda^{-1}, \quad i = 1, \dots, n. \quad (6)$$

It is well known that if λ has the pdf in (5), then *marginally* $\mathbf{X} = (X_1, \dots, X_n)$ follows a multivariate t distribution with degree of freedom ν (see, e.g., Geweke, 1993). The shock variable λ is a pure mathematical construct to induce a t structure, but one might interpret it as an economy-wide shock variable. When it takes on small values, we observe many simultaneous defaults.

The following theorem states an asymptotic result for the probability of large portfolio losses. To set the stage, let $h(x)$ be a function that increases at a subexponential rate such that $h(x) \rightarrow \infty$ as $x \rightarrow \infty$. Let n denote the number of obligors in the portfolio and set the default thresholds for the i th obligor to be $x_i = a_i h(n)$, where $a_i > 0$ is a positive constant. Put differently, the i th obligor defaults when $X_i > a_i h(n)$. Recall that c_i denotes the monetary loss associated with the i th obligor's default. Lastly, write the overall portfolio as

$$L_n = c_1 I_{\{X_1 > a_1 h(n)\}} + \cdots + c_n I_{\{X_n > a_n h(n)\}}.$$

Theorem 1 (Bassamboo et al., 2008). *Let the sequence $((c_i, a_i) : i \geq 0)$ take values in a finite set \mathcal{W} . Further assume that the proportion of each element $(c_i, a_i) \in \mathcal{W}$ in the portfolio converges to $q_i > 0$. Under the distributional assumptions in (4) and (5)*

$$\lim_{n \rightarrow \infty} h(n)^\nu \mathbb{P}(L_n > nb) = K,$$

for $0 < b < \bar{c}$, where K is a positive constant and \bar{c} is the limiting average loss when all the obligors default.

4 Conditional Monte Carlo Simulation

In this section we present two algorithms based on conditional Monte Carlo to estimate $\mathbb{P}(L(\mathbf{X}) > \gamma)$ —the rare-event probability of the occurrence of large losses—under the t -copula introduced in Section 3. We emphasize that the proposed algorithms can be easily extended to cover the general d -factor model of the form

$$X_i = (w_{i1}Z_1 + \cdots + w_{id}Z_d + w_i\eta_i)\lambda^{-1}, \quad i = 1, \dots, n, \quad (7)$$

where Z_1, \dots, Z_d are iid standard normal random variables, as this model has exactly the same asymptotics as the single-factor model. In fact, it can be shown that the loss $L(\mathbf{X})$ is large primarily when the shock variable λ is small. In Chan and Kroese (2010) we consider a slightly more general d -factor model where the factors and idiosyncratic risks are modeled as independent t random variables. More specifically, by introducing independent shock variables $\tau_j^2 \sim \text{Gamma}(\nu_z/2, \nu_z/2)$, $j = 1, \dots, d$, and $\lambda_i^2 \sim \text{Gamma}(\nu_\eta/2, \nu_\eta/2)$, $i = 1, \dots, n$, we can equivalently write the model as

$$X_i = w_{i1}Z_1\tau_1^{-1} + \cdots + w_{id}Z_d\tau_d^{-1} + w_i\eta_i\lambda_i^{-1}, \quad i = 1, \dots, n. \quad (8)$$

Comparing to the specification in (7), the model in (8) is more general in the sense that the former is a restricted version of the latter model by assuming $\tau_1 = \cdots = \tau_d = \lambda_1 = \cdots = \lambda_n$. Regarding the tail asymptotics under the specification (8), the rare event $\{L(\mathbf{X}) > \gamma\}$ occurs primarily when one of the shock variables τ_1, \dots, τ_d is small while the other random variables are ‘typical’. By utilizing this asymptotic description, we are able to construct an efficient conditional Monte Carlo algorithm based on an estimator first developed in Asmussen and Kroese (2006).

Despite the similarity of tail asymptotics for the loss distribution, the techniques developed previously do not directly apply to the current model. However, the same approach still proves to be fruitful. More specifically, the rare event $\{L(\mathbf{X}) > \gamma\}$ happens primarily when the shock variable λ takes small values, while Z and $\boldsymbol{\eta} = (\eta_1, \dots, \eta_n)$ have little influence on the occurrence of the rare event. Put differently, only λ is significantly affected by conditioning on the rare event, while all the other variables are not. This suggests that substantial variance reduction could be achieved simply by ‘integrating out’ λ analytically (or evaluating the integral using fast routines). Therefore, let us consider the following simple algorithm: first simulate Z and $\boldsymbol{\eta}$ from the nominal distributions, and then ‘integrate out’ λ given $(Z, \boldsymbol{\eta})$. To this end, define

$$R_i = \frac{\rho Z + \sqrt{1 - \rho^2} \eta_i}{a_i}, \quad i = 1, \dots, n.$$

Recall that the individual threshold of defaulting, x_i , is related to a_i via $x_i = a_i h(n)$. Let $R_{(1)} \leq \dots \leq R_{(n)}$ be the order statistics of R_1, \dots, R_n , and let $c_{(i)}$ denote the corresponding monetary loss associated with $R_{(i)}$. Then, it is easy to check that the event $\{L(\mathbf{X}) > \gamma\}$ happens if and only if $\lambda < R_{(k)}/h(n)$, where $k = \min\{l : \sum_{i=l+1}^n c_{(i)} \leq \gamma\}$. In particular, if $c_i = c$ for all $i = 1, \dots, n$, then $k = n - \lfloor \gamma/c \rfloor$, where $\lfloor \cdot \rfloor$ represents the integer part. For notational convenience, let $\mathbf{Y} = (Z, \boldsymbol{\eta})$ and denote the nominal density of \mathbf{Y} as $f(\mathbf{y}; \mathbf{u})$, where \mathbf{u} is a vector of parameters. In other words,

$$f(\mathbf{y}; \mathbf{u}) = f_{\mathbf{N}}(z; 0, 1) \prod_{i=1}^n f_{\mathbf{N}}(\eta_i; 0, \sigma_{\boldsymbol{\eta}}^2), \quad (9)$$

where $f_{\mathbf{N}}(\cdot; \mu, \sigma^2)$ is the density of the $\mathbf{N}(\mu, \sigma^2)$ distribution. Since $\lambda^2 \sim \text{Gamma}(\nu/2, \nu/2)$, we have

$$\mathbb{P}(L(\mathbf{X}) > \gamma | Z, \boldsymbol{\eta}) = \mathbb{P}\left(\lambda < \frac{R_{(k)}}{h(n)} \mid Z, \boldsymbol{\eta}\right) = F_{\mathbf{G}}\left(\frac{r^2}{h(n)^2}\right) \equiv S(\mathbf{Y}), \quad (10)$$

where $r = \max(R_{(k)}, 0)$ and $F_{\mathbf{G}}$ is the cumulative distribution function (cdf) of the $\text{Gamma}(\nu/2, \nu/2)$ distribution. Finally, the rare-event probability (2) can be estimated by the conditional Monte Carlo (CondMC) estimator:

$$\frac{1}{N} \sum_{i=1}^N S(\mathbf{Y}^{(i)}), \quad (11)$$

where $\mathbf{Y}^{(i)} = (Z^{(i)}, \boldsymbol{\eta}^{(i)})$, $i = 1, \dots, N$ are obtained from the nominal pdf $f(\mathbf{y}; \mathbf{u})$ in (9). We summarize the procedure below.

Algorithm 1 (Conditional Monte Carlo for the t -copula Model).

1. Obtain N samples $\mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(N)}$ from the nominal pdf $f(\mathbf{y}; \mathbf{u})$ defined in (9).
2. Use the samples $\mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(N)}$ to compute the CondMC estimator in (11).

We now show that the CondMC estimator has bounded relative error. Let $F_{\mathbf{Y}}$ denote the distribution of \mathbf{Y} , and let \mathbb{E} be the corresponding expectation operator.

Theorem 2. *Under the same assumptions as in Theorem 1, we have*

$$\limsup_{n \rightarrow \infty} \frac{\mathbb{E}S^2(\mathbf{Y})}{\mathbb{P}(L_n > nb)^2} < \infty.$$

In other words, the CondMC estimator in (11) has bounded relative error.

The proof of the above theorem is given in the appendix. Despite the good computational and theoretical properties of the CondMC estimator (see Section 5), one can improve its efficiency by adding ideas from the *cross-entropy method* (Rubinstein and Kroese, 2004, 2007). Specifically, instead of the naive Monte Carlo estimator in (11), we consider the IS estimator

$$\frac{1}{N} \sum_{i=1}^N S(\mathbf{Y}^{(i)}) \frac{f(\mathbf{Y}^{(i)}; \mathbf{u})}{f(\mathbf{Y}^{(i)}; \mathbf{v}^*)}, \quad (12)$$

where $\mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(N)} \sim f(\cdot; \mathbf{v}^*)$ and $f(\cdot; \mathbf{v}^*)$ is a proposal density chosen within the same parametric family as that of the nominal pdf $f(\cdot; \mathbf{u})$. In our case the parametric family under consideration is

$$\mathcal{F} = \left\{ f(\mathbf{y}; \mathbf{v}) = f_{\mathbf{N}}(z; \mu_z, V_z) \prod_{i=1}^n f_{\mathbf{N}}(\eta_i; \mu_{\boldsymbol{\eta}}, V_{\boldsymbol{\eta}}) \right\},$$

where $\mathbf{v} = (\mu_z, V_z, \mu_{\boldsymbol{\eta}}, V_{\boldsymbol{\eta}})$ is a vector of parameters. The optimal proposal density is obtained by locating the member $f \in \mathcal{F}$ that minimizes its cross-entropy distance to the zero-variance proposal density

$$g^*(\mathbf{y}) \propto S(\mathbf{y})f(\mathbf{y}; \mathbf{u}).$$

In our case, minimization of the cross-entropy is equivalent to solving the following maximization problem (for details see Rubinstein and Kroese, 2007, pp. 136–141):

$$\max_{\mathbf{v}} \int S(\mathbf{y}) \log f(\mathbf{y}; \mathbf{v}) f(\mathbf{y}; \mathbf{u}) d\mathbf{y}. \quad (13)$$

Since most often an analytical solution to the above maximization problem is not available, we consider instead its stochastic counterpart

$$\max_{\mathbf{v}} \frac{1}{M} \sum_{i=1}^M S(\mathbf{Y}^{(i)}) \log f(\mathbf{Y}^{(i)}; \mathbf{v}), \quad (14)$$

where $\mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(M)} \sim f(\cdot; \mathbf{u})$. The maximization problem in (14) is easy to solve. In fact, the solutions are available analytically:

$$\mu_z^* = \frac{\sum_{i=1}^M S(\mathbf{Y}^{(i)}) Z^{(i)}}{\sum_{i=1}^M S(\mathbf{Y}^{(i)})}, \quad V_z^* = \frac{\sum_{i=1}^M S(\mathbf{Y}^{(i)}) (Z^{(i)} - \mu_z^*)^2}{\sum_{i=1}^M S(\mathbf{Y}^{(i)})}, \quad (15)$$

$$\mu_{\boldsymbol{\eta}}^* = \frac{\sum_{i=1}^M S(\mathbf{Y}^{(i)}) \sum_{j=1}^n \eta_j^{(i)}}{n \sum_{i=1}^M S(\mathbf{Y}^{(i)})}, \quad V_{\boldsymbol{\eta}}^* = \frac{\sum_{i=1}^M S(\mathbf{Y}^{(i)}) \sum_{j=1}^n (\eta_j^{(i)} - \mu_{\boldsymbol{\eta}}^*)^2}{n \sum_{i=1}^M S(\mathbf{Y}^{(i)})}. \quad (16)$$

Once we obtain the optimal proposal density $f(\mathbf{y}; \mathbf{v}^*)$, we simply compute the CondMC-CE estimator in (12). We summarize the above procedure as follows.

Algorithm 2 (Conditional Monte Carlo with CE for the t -copula Model).

1. Obtain M samples $\mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(M)}$ from the nominal pdf $f(\mathbf{y}; \mathbf{u})$ defined in (9).
2. Use the sample to solve the maximization program in (14). That is, compute $\mathbf{v}^* = (\mu_z^*, V_z^*, \mu_{\boldsymbol{\eta}}^*, V_{\boldsymbol{\eta}}^*)$ using (15)–(16) to obtain the proposal density $f(\mathbf{y}; \mathbf{v}^*)$.
3. Obtain N samples $\mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(N)}$ from $f(\mathbf{y}; \mathbf{v}^*)$ and compute the CondMC-CE estimator in (12).

It is worth noting that CondMC involves generating random variables only from the nominal distribution $f(\mathbf{y}; \mathbf{u})$. Hence, it is as efficient as naive simulation in terms of random variable generation time. In addition, it does not require computing any likelihood ratio that occurs in any IS estimator. In contrast, CondMC-CE requires a preliminary run to estimate the optimal

parameters \mathbf{v}^* for the proposal density, and at a first glance seems to be less efficient. However, the sample size needed for the preliminary run is small (in practice $M = 1000$ is enough), and the effort for computing the optimal parameter \mathbf{v}^* in (15)–(16) is trivial. Moreover, as the simulation experiments in the next section show, the variance reduction achieved is well worth the additional computational effort.

5 Numerical Results

In this section we demonstrate the performance of the proposed estimators, CondMC and CondMC-CE, via simulation studies similar to those in Bassamboo et al. (2008). The broad conclusions drawn from these experiments are that both algorithms offer substantial variance reduction compared with naive simulation, and they compare favorably to the two IS estimators, called ECM and HRT, proposed in Bassamboo et al. (2008). Specifically, while CondMC performs similarly to the more efficient ECM algorithm, offering 1.2 to 4 times higher variance reduction, CondMC-CE performs much better, providing 20 to 100 times higher variance reduction. Another factor that is in favor of the proposed algorithms is that both involve generating random variables only from standard distributions (normal distributions to be specific). In contrast, the ECM estimator involves rejection sampling, which takes on average three times longer than naive simulation. Moreover, the normalizing constant of the proposal density, which is not known, has to be computed by numerical routines, thus making the algorithm slower and more difficult to implement.

For comparison purposes, we utilize the same sets of parameter values as those in Bassamboo et al. (2008) Tables 1–4,¹ where only homogeneous portfolios are considered. Nevertheless, it is important to emphasize that the empirical performance of the proposed algorithms as well as the fact that the estimators have bounded relative error are not affected by assuming an inhomogeneous credit portfolio, as the asymptotics under an inhomogeneous credit portfolio are exactly the same. In fact, Theorem 1 is proved by assuming an inhomogeneous credit portfolio. In all the experiments in this subsection we set $\sigma_\eta^2 = 9$, $x = \sqrt{n} \times 0.5$, $l = b \times n$ and $c = 1$.

To access the accuracy of the estimators, we use the notion of relative error, which is simply the ratio of the estimator’s standard deviation to the true probability of the rare event. More precisely, for an unbiased estimator $\hat{\ell} = N^{-1} \sum_{i=1}^N H(\mathbf{X}_i)$ of ℓ , its relative error is defined as $\text{RE} = \sqrt{\text{Var}(\hat{\ell})}/\ell$. However, in practice we do not know the value of the true probability ℓ , but the relative error can be estimated by the consistent estimator

$$\frac{S/\sqrt{N}}{\hat{\ell}},$$

where S^2 is the sample variance of $H(\mathbf{X})$.

For each set of specified parameters, we generate 50000 samples for both algorithms. Specifically, for CondMC-CE we use a sample size of $M = 1000$ for estimating the optimal parameters via

¹In Bassamboo et al. (2008) Tables 3–4, the authors actually computed $\mathbb{P}(L(\mathbf{X}) \geq \gamma)$ instead of $\mathbb{P}(L(\mathbf{X}) > \gamma)$ as stated. As a result, the estimated rare-event probabilities there are slightly larger than those we report in the corresponding tables.

the CE method, and a sample of size $N = 49000$ for the main run. Table 1 shows the estimated relative errors (in %) of the proposed estimators, as well as those of the ECM and HRT, for various values of the degree of freedom parameter ν . We also report the variance reduction achieved by the proposed estimators compared with naive simulation. The estimated rare-event probabilities $\widehat{\ell}(\gamma)$ are obtained by the more accurate CondMC-CE estimator. Other model parameters are chosen to be $n = 250$, $\rho = 0.25$ and $b = 0.25$. In Table 2 we perform the same comparison but now we vary the correlation parameter ρ while keeping ν fixed at 12. As is clear from the tables, both algorithms offer substantial variance reduction compared with naive simulation, and they compare favorably to both the ECM and HRT estimators.

Table 1: Performance of the CondMC and CondMC-CE estimators for the t -copula model for various values of ν . Variance reduction is compared with naive simulation.

ν	$\widehat{\ell}(\gamma)$	Relative Error (%)				Variance Reduction	
		CondMC	CondMC-CE	ECM	HRT	CondMC	CondMC-CE
4	8.13×10^{-3}	0.3	0.1	0.6	1.1	271	2440
8	2.42×10^{-4}	0.7	0.2	0.9	1.8	1690	20656
12	1.07×10^{-5}	1.2	0.3	1.7	2.6	12980	2.08×10^5
16	6.16×10^{-7}	2.0	0.5	2.8	3.6	81170	1.30×10^6
20	4.38×10^{-8}	3.3	0.6	3.7	5.4	4.19×10^5	1.27×10^7

Table 2: Performance of the CondMC and CondMC-CE estimators for the t -copula model for various values of ρ . Variance reduction is compared with naive simulation.

ρ	$\widehat{\ell}(\gamma)$	Relative Error (%)				Variance Reduction	
		CondMC	CondMC-CE	ECM	HRT	CondMC	CondMC-CE
0.1	8.58×10^{-6}	0.8	0.4	0.9	1.8	32520	1.77×10^5
0.2	9.83×10^{-6}	1.0	0.4	1.2	2.3	18370	1.34×10^5
0.3	1.19×10^{-5}	1.3	0.4	1.7	3.2	9112	1.68×10^5
0.4	1.46×10^{-5}	1.7	0.3	3.1	4.0	4472	1.56×10^5

It is also worth noting that as ν decreases (see Table 1), the probability of large portfolio loss increases several orders of magnitude, highlighting the importance of correctly modeling the tail behavior of the latent variables \mathbf{X} . This observation suggests that the t -copula model might be a better choice than the normal copula model, as the former offers more flexible modeling of tail behavior, and includes the latter as a limiting case. Another feature that is worth commenting is that when ρ increases, the performance of CondMC deteriorates while that of CondMC-CE is essentially unchanged (see Table 2). This should not be a surprise, as for CondMC we only ‘integrate out’ λ , while generating other variables from their nominal distributions. The underlying rationale for this procedure is that the rare event $\{L(\mathbf{X}) > \gamma\}$ happens primarily when λ is small, and other variables have little influence on the occurrence of the rare event $\{L(\mathbf{X}) > \gamma\}$. As ρ increases, the factor Z becomes relatively more important in determining the occurrence of the rare event, and ignoring this contribution of Z deteriorates the performance of the algorithm. But for CondMC-CE, in addition to ‘integrating out’ λ , we also twist the

distributions of all the other variables. Therefore its performance is essentially unchanged as we increase ρ .

In Table 3 we report the relative errors (in %) of CondMC and CondMC-CE as well as those of ECM and HRT for various values of n , the number of obligors. Other model parameters are chosen to be $\nu = 12$, $\rho = 0.25$ and $b = 0.25$. Table 4 shows the results of a similar analysis but now we vary b , the proportion of defaults in the portfolio, while keeping n fixed at 250. The results suggest that both algorithms perform remarkably well even when n is large, where the model contains hundreds of random variables.

Table 3: Performance of the CondMC and CondMC-CE estimators for the t -copula model for various values of n . Variance reduction is compared with naive simulation.

n	$\widehat{\ell}(\gamma)$	Relative Error (%)				Variance Reduction	
		CondMC	CondMC-CE	ECM	HRT	CondMC	CondMC-CE
100	1.83×10^{-3}	1.2	0.5	1.6	1.8	73	550
250	1.07×10^{-5}	1.2	0.3	1.7	2.6	12980	20660
500	1.51×10^{-7}	1.1	0.3	1.5	3.4	1.07×10^6	1.93×10^7
1000	2.28×10^{-9}	1.0	0.2	1.6	3.6	8.98×10^7	2.72×10^9

Table 4: Performance of the CondMC and CondMC-CE estimators for the t -copula model for various values of b . Variance reduction is compared with naive simulation.

b	$\widehat{\ell}(\gamma)$	Relative Error (%)				Variance Reduction	
		CondMC	CondMC-CE	ECM	HRT	CondMC	CondMC-CE
0.1	3.47×10^{-3}	0.5	0.2	0.9	1.6	232	1107
0.2	7.37×10^{-5}	0.9	0.3	1.2	2.5	3553	32950
0.3	1.12×10^{-6}	1.7	0.4	2.0	3.4	59980	9.15×10^5

6 The Skew t -copula Model

Although the t -copula model allows a more flexible dependence structure that accommodates the extremal dependence among the defaults, it is not flexible enough to allow for an asymmetric default distribution. In this section we consider a skew t -copula model based on the *skew-normal distribution* (Azzalini, 1985; Azzalini and Valle, 1996) that allows for the aforementioned asymmetry. Specifically, given a normally distributed random variable $U \sim \mathcal{N}(\mu, \sigma^2)$, if we define $Y = U + \delta W$, where δ is a fixed parameter and W follows a standard normal distribution left truncated at 0, then Y has a skew-normal distribution with parameter δ controlling the skewness of the distribution: if $\delta > 0$ ($\delta < 0$), then Y is positively (negatively) skewed; if $\delta = 0$, it reduces to a symmetric normally distributed random variable. To introduce a skew t structure, the

latent variables X_i are modeled as follows:

$$X_i = \left(\rho Z + \delta W + \sqrt{1 - \rho^2} \eta_i \right) \lambda^{-1}, \quad i = 1, \dots, n, \quad (17)$$

$$W \sim \text{TN} \left(-\sqrt{\frac{2}{\pi}}, 1 \right), \quad (18)$$

where $\text{TN}(\mu, \sigma^2)$ is a normal distribution with mean μ and variance σ^2 left truncated at $-\sqrt{2/\pi}$.

The main reason why the random variable W is assumed to have mean $-\sqrt{2/\pi}$ (and left truncated at the same value) rather than 0 is that in this way its expected value is 0. Hence, the expected value of X_i would not be affected by adding the term δW . This simple formulation generalizes the t -copula model to allow \mathbf{X} to have an asymmetric distribution, where the skewness is controlled by the parameter δ . It is obvious that if we set $\delta = 0$, it reduces to the t -copula model introduced in Section 3. Another advantage of this construction is that it is very parsimonious as it introduces only one extra parameter compared to the t copula model. Moreover, the two conditional Monte Carlo algorithms developed for estimating the probability of large losses under the t -copula can be easily modified to cover the skew t -copula case.

We now draw our attention to the problem of estimating $\mathbb{P}(L(\mathbf{X}) > \gamma)$ under the skew t -copula model. For notational convenience, let $\tilde{\mathbf{Y}} = (\mathbf{Y}, W) = (Z, \boldsymbol{\eta}, W)$ and denote the nominal distribution of $\tilde{\mathbf{Y}}$ by $\tilde{f}(\tilde{\mathbf{y}}; \tilde{\mathbf{u}})$, where $\tilde{\mathbf{u}}$ is a vector of parameters. In other words,

$$\tilde{f}(\tilde{\mathbf{y}}; \tilde{\mathbf{u}}) = f_{\text{TN}} \left(-\sqrt{2/\pi}, 1 \right) f_{\text{N}}(z; 0, 1) \prod_{i=1}^n f_{\text{N}}(\eta_i; 0, \sigma_{\boldsymbol{\eta}}^2), \quad (19)$$

where $f_{\text{TN}}(\mu, \sigma^2)$ is the pdf of a $\text{N}(\mu, \sigma^2)$ distributed random variable left truncated at $-\sqrt{2/\pi}$. Define

$$\tilde{R}_i = \frac{\rho Z + \sqrt{1 - \rho^2} \eta_i + \delta W}{a_i}, \quad i = 1, \dots, n,$$

and let $\tilde{R}_{(1)}, \dots, \tilde{R}_{(n)}$ be the corresponding order statistics. Under the skew t -copula model, the rare event $\{L(\mathbf{X}) > \gamma\}$ happens if and only if $\lambda < \tilde{R}_{(k)}/h(n)$, where $k = \min\{l : \sum_{i=l+1}^n c_{(i)} \leq \gamma\}$. Therefore, we can estimate the rare-event probability $\ell(\gamma) = \mathbb{P}(L(\mathbf{X}) > \gamma)$ by the CondMC estimator

$$\frac{1}{N} \sum_{i=1}^N \tilde{S}(\tilde{\mathbf{Y}}^{(i)}), \quad (20)$$

where $\tilde{Y}_1, \dots, \tilde{Y}_n \sim \tilde{f}(\cdot; \tilde{\mathbf{u}})$, $\tilde{S}(\tilde{\mathbf{Y}}) = F_{\text{G}}(\tilde{r}^2/h(n)^2)$ and $\tilde{r} = \max(\tilde{R}_{(k)}, 0)$. Truncated normal random variables can be obtained by the inverse-transform method or various efficient rejection algorithms (e.g., Geweke, 1991; Robert, 1995).

For the CondMC-CE estimator of $\ell(\gamma)$ under the skew t -copula model, we locate the proposal density $\tilde{f}(\cdot; \tilde{\mathbf{v}}^*)$ within the parametric family

$$\tilde{\mathcal{F}} = \left\{ \tilde{f}(\tilde{\mathbf{y}}; \tilde{\mathbf{v}}) = f_{\text{TN}}(\mu_w, 1) f_{\text{N}}(z; \mu_z, V_z) \prod_{i=1}^n f_{\text{N}}(\eta_i; \mu_{\boldsymbol{\eta}}, V_{\boldsymbol{\eta}}) \right\},$$

where $\tilde{\mathbf{v}} = (\mu_w, \mu_z, V_z, \mu_{\boldsymbol{\eta}}, V_{\boldsymbol{\eta}})$. Notice that for the random variable W we only consider tilting its mean but not its variance. To obtain the optimal parameter vector $\tilde{\mathbf{v}}^*$ via the CE method, we

solve a maximization program similar to (14). Specifically, we obtain M samples $\tilde{\mathbf{Y}}^{(1)}, \dots, \tilde{\mathbf{Y}}^{(M)}$ from the nominal pdf $\tilde{f}(\cdot; \tilde{\mathbf{u}})$ and use them to compute μ_z^* , V_z^* , μ_η^* , and V_η^* via (15)–(16), replacing $S(\cdot)$ by $\tilde{S}(\cdot)$ and \mathbf{Y} by $\tilde{\mathbf{Y}}$. Lastly, μ_w^* can be found by locating the root of the following univariate function in μ_w (e.g., by bisection method):

$$-\frac{\sum_{i=1}^M \tilde{S}(\tilde{\mathbf{Y}}^{(i)}) W^{(i)}}{\sum_{i=1}^M \tilde{S}(\tilde{\mathbf{Y}}^{(i)})} + \mu_w - \frac{\phi(\mu_w + \sqrt{2/\pi})}{\Phi(\mu_w + \sqrt{2/\pi})} = 0, \quad (21)$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are respectively the pdf and cdf of the standard normal distribution. Once the proposal density $\tilde{f}(\cdot; \tilde{\mathbf{v}}^*)$ is obtained, the CondMC-CE estimator for the skew t -copula model can be computed similar to (12).

Finally, we investigate how the skewness parameter δ of the multivariate default distribution affects the probability of large portfolio losses under the skew t -copula model. For the following simulated experiment, we set $n = 250$, $\nu = 12$, $\sigma_\eta^2 = 9$, $c = 1$, $x = \sqrt{n} \times 0.5$, $\rho = 0.25$, $b = 0.25$, and $l = b \times n$. The simulation budget is again 50000. Table 5 reports the relative errors (in %) of the proposed estimators, as well as the variance reduction obtained compared with naive simulation. The estimated rare-event probabilities $\hat{\ell}(\gamma)$ are obtained by the more accurate CondMC-CE estimator.

Table 5: Performance of the CondMC and CondMC-CE estimators for the skew t -copula model for various values of n . Variance reduction is compared with naive simulation.

δ	$\hat{\ell}(\gamma)$	Relative Error (%)		Variance Reduction	
		CondMC	CondMC-CE	CondMC	CondMC-CE
0	1.07×10^{-5}	1.2	0.5	12980	70630
0.5	2.74×10^{-5}	2.1	0.3	1509	80200
1	1.93×10^{-4}	3.6	0.3	88	15120
1.5	1.03×10^{-3}	3.5	0.3	16	2542

As is clear from the table, the skewness of the multivariate default distribution has a profound effect on the the probability of incurring a large portfolio loss, thus highlighting the importance of correctly choosing a credit risk model. It is also of interest to note that as δ gets larger and the event $\{L(\mathbf{X}) > \gamma\}$ becomes less rare, the performance of the CondMC estimator actually deteriorates. This at first sight might seem surprising. But when one takes a closer look, one realizes that for the CondMC estimator, only λ is ‘integrated out’ while all the other variables are generated from their nominal distributions. Therefore, its performance depends critically on the asymptotic description of the way in which the event $\{L(X) > \gamma\}$ occurs. As the event becomes less rare, the asymptotic description is less accurate. Consequently, its performance also deteriorates and approaches to that of the crude Monte Carlo.

7 Concluding Remarks

In this article we first propose two new simulation algorithms based on conditional Monte Carlo to estimate large portfolio losses under the t -copula model. Through an extensive simulation

study, we demonstrate that both estimators offer substantial variance reduction and outperform existing algorithms. Next, we generalize the t -copula model to allow an asymmetric default distribution. A simulation study shows that the skewness parameter of the default distribution has a large impact on the probability of large portfolio losses. This illustrates the importance of correctly specifying the joint default distribution, and since the skew t -copula model includes both t - and normal copulas as special cases, it is arguably a more adequate way to model the default correlation.

Appendix

We will first introduce some notations that are needed for the proofs. Recall that

$$R_i = \frac{\rho Z + \sqrt{1 - \rho^2} \eta_i}{a_i},$$

$R_{(i)}$ is the i th order statistics of R_1, \dots, R_n , and $c_{(i)}$ the corresponding monetary loss. Define $r_n = R_{(k)}$, where $k = \min\{l : \sum_{i=l+1}^n c_{(i)} \leq \gamma\}$ and $\gamma = bn$. Let $Q_i = \sqrt{1 - \rho^2} \eta_i / a_i$, and define $Q_{(i)}$ and q_n similarly. We will need the following lemma for the proof of Theorem 2.

Lemma 1. *Under the assumptions of Theorem 1, we have*

$$\limsup_{n \rightarrow \infty} \mathbb{E} r_n^{2\nu} < \infty.$$

Proof. For notational simplicity, we will only prove the case where $c_i = c$ and $a_i = a$; the general case follows analogously. Since $Q_i \sim \mathcal{N}(0, (1 - \rho^2)\sigma_\eta^2/a^2)$ and $q_n = Q_{(k)}$ is the k th order statistics with $\lim_{n \rightarrow \infty} k/n = (1 - b/a)$, by the Central Limit Theorem q_n is asymptotically normal with mean q and variance σ^2/n , where q and σ^2 are some constants not depending on Z (for more details see van der Vert, 1998, Chapter 21). Therefore, given $Z = z$, $r_n = q_n + \rho Z/a$ is asymptotically normal with mean $\rho z/a + q$ and variance σ^2/n . Consequently, $\mathbb{E}[r_n^{2\nu} | Z = z]$ is asymptotically a polynomial of order 2ν in z , with leading term $(\rho z/a + q)^{2\nu}$. Finally, by the law of iterated expectation, asymptotically $\mathbb{E} r_n^{2\nu} = \mathbb{E}[(\rho Z/a + q)^{2\nu}] + O(n)$. Now the result follows immediately. \square

Proof of Theorem 2. The theorem follows from the following computations:

$$\begin{aligned} \limsup_{n \rightarrow \infty} h(n)^{2\nu} \mathbb{E} S^2(\mathbf{Y}) &= \limsup_{n \rightarrow \infty} h(n)^{2\nu} \mathbb{E} \left(\int_0^{\frac{r_n}{h(n)}} k_1 t^{\nu-1} e^{-\nu t^2/2} dt \right)^2 \\ &\leq \limsup_{n \rightarrow \infty} h(n)^{2\nu} \mathbb{E} \left(\int_0^{\frac{r_n}{h(n)}} k_1 t^{\nu-1} dt \right)^2 \\ &= \limsup_{n \rightarrow \infty} h(n)^{2\nu} \mathbb{E} \left(\frac{k_1 r_n^\nu}{\nu h(n)^\nu} \right)^2 \\ &= k_2 \limsup_{n \rightarrow \infty} \mathbb{E} r_n^{2\nu} < \infty, \end{aligned}$$

where k_1 and k_2 are some unimportant constants not depending on n . By Theorem 1, which provides the asymptotic for the probability of large portfolio losses $\mathbb{P}(L_n > nb)$, we conclude that the CondMC estimator has bounded relative error.

□

References

- S. Asmussen and P. W. Glynn. *Stochastic simulation: algorithms and analysis*. Springer, New York, 2007.
- S. Asmussen and D. P. Kroese. Improved algorithms for rare event simulation with heavy tails. *Advances in Applied Probability*, 38:545–558, 2006.
- A. Azzalini. A class of distributions which includes the normal ones. *Scandinavian Journal of Statistics*, 12:171–178, 1985.
- A. Azzalini and A. Dalla Valle. The multivariate skew-normal distribution. *Biometrika*, 83:715–726, 1996.
- A. Bassamboo, S. Juneja, and A. Zeevi. Portfolio credit risk with extremal dependence: Asymptotic analysis and efficient simulation. *Operations Research*, 56(3):593–606, 2008.
- C. Bluhm, L. Overbeck, and C. Wagner. *An Introduction to Credit Risk Modeling*. Chapman & Hall/CRC, 2002.
- J. C. C. Chan and D. P. Kroese. Rare-event probability estimation with Monte Carlo. *Annals of Operations Research*, 2010. Forthcoming.
- D. Egloff, M. Leippold, S. Jöhri, and C. Dalbert. Optimal importance sampling for credit portfolios with stochastic approximation, SSRN working paper, 2005.
- C. Fernandez and M. F. J. Steel. On Bayesian modeling of fat tails and skewness. *Journal of the American Statistical Association*, 93(441):359–371, 1998.
- P. H. Franses and D. van Dijk. *Non-linear Time Series Models in Empirical Finance*. Cambridge University Press, Cambridge, 2000.
- J. Geweke. Efficient simulation from the multivariate normal and student- t distributions subject to linear constraints and the evaluation of constraint probabilities. In E. Keramidas and S. Kaufman, editors, *Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface*, pages 571–578, 1991.
- J. Geweke. Bayesian treatment of the independent Student- t linear model. *Journal of Applied Econometrics*, 8:S19–S40, 1993.
- P. Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer, New York, 2004.
- P. Glasserman and J. Li. Importance sampling for portfolio credit risk. *Management Science*, 51(11):1643–1656, 2005.

- P. Glasserman, W. Kang, and P. Shahabuddin. Large deviations of multifactor portfolio credit risk. *Mathematical Finance*, 17:345–379, 2007.
- P. Glasserman, W. Kang, and P. Shahabuddin. Fast simulation of multifactor portfolio credit risk. *Operations Research*, 56(5):1200–1217, 2008.
- P. Grundke. Importance sampling for integrated market and credit portfolio models. *European Journal of Operational Research*, 194(1):206–226, 2009.
- G. Gupton, C. Finger, and M. Bhatia. Creditmetrics technical document. Technical report, J. P. Morgan & Co., New York, 1997.
- M. S. Joshi. Applying importance sampling to pricing single tranches of CDOs in a one-factor Li model. Technical report, QUARC, Group Risk Management, Royal Bank of Scotland, 2004.
- S. Juneja and P. Shahabuddin. Simulating heavy tailed processes using delayed hazard rate twisting. *ACM Trans. Model. Comput. Simul.*, 12(2):94–118, 2002.
- M. Kalkbrener, H. Lotter, and L. Overbeck. Sensible and efficient capital allocation for credit portfolios. *Risk*, 17(1):S19–S24, 2004.
- D. Li. On default correlations: a copula function approach. *Journal of Fixed Income*, 9:43–54, 2000.
- A. McNeil, R. Frey, and P. Embrechts. *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press, Princeton, New Jersey, 2005.
- C. P. Robert. Simulation of truncated normal variables. *Statistics and Computing*, 5:121–125, 1995.
- R. Y. Rubinstein and D. P. Kroese. *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization Monte-Carlo Simulation, and Machine Learning*. Springer-Verlag, New York, 2004.
- R. Y. Rubinstein and D. P. Kroese. *Simulation and the Monte Carlo Method Second Edition*. John Wiley & Sons, New York, 2007.
- A. W. van der Vert. *Asymptotic Statistics*. Cambridge University Press, 1998.