

Conditional Monte Carlo for sums, with applications to insurance and finance

Søren Asmussen*

Department of Mathematics, Aarhus University, Ny Munkegade, 8000 Aarhus C, Denmark

Abstract

Conditional Monte Carlo replaces a naive estimate Z of a number z by its conditional expectation given a suitable piece of information. It always reduces variance and its traditional applications are in that vein. We survey here other potential uses such as density estimation and calculations for Value-at-Risk and/or expected shortfall, going in part into the implementation in various copula structures. Also the interplay between these different aspects comes into play.

Keywords

Archimedean copula; Density estimation; Expected shortfall; Lognormal sums; Value-at-Risk

1. Introduction

Let z be a number represented as an expectation $z = \mathbb{E}Z$. The crude Monte Carlo (CrMC) method for estimating z proceeds by simulating R replications Z_1, \dots, Z_R of Z and returning the average $\bar{z} = (Z_1 + \dots + Z_R) / R$ as a point estimate. The uncertainty is reported as an asymptotic confidence interval based on the central limit theorem (CLT); e.g., the two-sided 95% confidence interval is $\bar{z} \pm 1.96s / R^{1/2}$ where s^2 is the empirical variance of the sample Z_1, \dots, Z_R .

The more refined conditional Monte Carlo (CdMC) method uses a piece of information collected in a σ -field \mathcal{F} and is implemented by performing CrMC with Z replaced by $Z_{\text{Cond}} = \mathbb{E}[Z \mid \mathcal{F}]$. It is traditionally classified as a variance reduction method but it can also be used for smoothing, though this is much less appreciated.

Both aspects are well illustrated via the problem of estimating $\mathbb{P}(S_n \leq x)$ where $S_n = X_1 + \dots + X_n$ is a sum of r.v.'s. The obvious choice for CrMC is $Z = Z(x) = \mathbb{1}(S_n \leq x)$. For CdMC, a simple possibility is to take $\mathcal{F} = \sigma(X_1, \dots, X_{n-1})$. In the case where X_1, X_2, \dots are i.i.d. with common distribution F one then has

$$Z_{\text{Cond}} = \mathbb{P}(S_n \leq x \mid X_1, \dots, X_{n-1}) = F(x - S_{n-1}) \quad (1.1)$$

This estimator has two noteworthy properties:

- for a fixed x its variance is smaller than that of $\mathbb{1}(S_n \leq x)$ used in the CrMC method; and
- when averaged over the number R of replications, it leads to estimates of $\mathbb{P}(S_n \leq x)$ which are smoother as function of $x \in (-\infty, \infty)$ than the more traditional empirical c.d.f. of R simulated replicates of S_n .

*Correspondence to: Søren Asmussen, Department of Mathematics, Aarhus University, Ny Munkegade, 8000 Aarhus C, Denmark. Tel: +45-8715 5756; E-mail: asmus@math.au.dk

This last property is easily understood for a continuous F , where $Z_{\text{Cond}}(x) = F(x - S_{n-1})$ is again continuous and therefore averages are also. In contrast, the empirical c.d.f. always has jumps. It also suggests that $f(x - S_{n-1})$ may be an interesting candidate for estimating the density $f_n(x)$ of S_n when F itself admits a density $f(x)$. In fact, density estimation is a delicate topic where traditional methods such as kernel smoothing or finite differences often involve tedious and ad hoc tuning of parameters like choice of kernel, window size, etc.

The variance reduction property holds in complete generality by the general principle (known as Rao–Blackwellization in statistics) that conditioning reduces variance:

$$\text{Var} Z = \mathbb{E}[\text{Var}[Z \mid \mathcal{F}]] + \text{Var}[\mathbb{E}[Z \mid \mathcal{F}]] \geq \text{Var}[\mathbb{E}[Z \mid \mathcal{F}]] = \text{Var} Z_{\text{Cond}}$$

In view of the huge literature on variance reduction, this may appear appealing but it also has some caveats inherent in the choice of \mathcal{F} : $\mathbb{E}[Z \mid \mathcal{F}]$ must be computable and have a variance that is *substantially* smaller than that of Z . Namely, if CdMC reduces the variance on Z of Z_{Cond} by a factor of $\tau < 1$, the same variance on the average could be obtained by taking $1/\tau$ as many replications in CrMC as in CdMC, see Asmussen & Glynn (2007: 126). This point is often somewhat swept under the carpet!

The present paper discusses such issues related to the CdMC method via the example of inference on the distribution of a sum $S_n = X_1 + \dots + X_n$. Here the X_i are assumed i.i.d. in sections 2–7, but we look into dependence in some detail in section 8, whereas a few comments on different marginals are given in section 9.

The motivation comes, to a large extent, from problems in insurance and finance such as assessing the form of the density of the loss distribution, estimating the tail of the aggregated claims in insurance, calculating the Value-at-Risk (VaR) or expected shortfall of a portfolio, etc. In many such cases, the tail of the distribution of S_n is of particular interest, with the relevant tail probabilities being of order 10^{-2} – 10^{-4} (but note that in other application areas, the relevant order is much lower, say 10^{-8} – 10^{-12} in telecommunications). By “tail” we are not just thinking of the *right* tail, i.e., $\mathbb{P}(S_n > x)$ for large x , which is relevant for the aggregated claims and portfolios with short positions. Also the *left* tail $\mathbb{P}(S_n \leq x)$ for small x comes up in a natural way, in particular for portfolios with long positions, but has received much less attention until the recent studies by Asmussen *et al.* (2016) and Gulisashvili & Tankov (2016).

The most noted use of CdMC in the insurance/finance/rare-event area appears to be the algorithm of Asmussen & Kroese (2006) for calculating the right tail of a heavy-tailed sum. A main application is ruin probabilities. We give references and put this in perspective to the more general problems of the present paper in section 4. Otherwise, the use of CdMC in insurance and finance seem to be remarkably few compared to other MC-based tools such as importance sampling (IS), stratification, simulation-based estimation of sensitivities (Greeks), just to name a few (see Glasserman, 2004, for these and other examples). Some exceptions are Fu *et al.* (2009) who study an CdMC estimator of a sensitivity of a quantile (not the quantile itself!) with respect to a model parameter, and Chu & Nakayama (2012) who give a detailed mathematical derivation of the CLT for quantiles estimated in a CdMC set-up, based on methodology from Bahadur (1966) and Ghosh (1971) (see also Nakayama, 2014).

1.1. Conventions

Throughout the paper, $\Phi(x)$ denotes the standard normal c.d.f., $\bar{\Phi}(x) = 1 - \Phi(x)$ its tail and $\varphi(x) = e^{-x^2/2} / \sqrt{2\pi}$ the standard normal p.d.f. For the $\gamma(\alpha, \lambda)$ distribution, α is the shape parameter and λ the rate so that the density is $x^{\alpha-1} \lambda^\alpha e^{-\lambda x} / \Gamma(\alpha)$.

Because of the financial relevance, an example that will be used frequently is X to be Lognormal $(0, 1)$, i.e., the summands in S_n to be of the form $X = e^V$ with $V \text{ Normal}(0, 1)$, and $n = 10$. Note that the mean of V is just a scaling factor and hence unimportant. In contrast, the variance (and the value of n) matters quite a lot for the shape of the distribution of S_n , but to be definite, we took it to be one. We refer to this set of parameters as our recurrent example, and many other examples are taken as smaller or larger modifications.

2. Density Estimation

If F has a density f , then S_n has density f_n given as an integral over a hyperplane:

$$f_n(x) = f^{*n}(x) = \int_{x_1 + \dots + x_n = x} f(x_1) \cdots f(x_n) dx_1 \cdots dx_n$$

Such convolution integrals can only be evaluated numerically for rather small n , and we shall here consider the estimator $f(x - S_{n-1})$ of $f_n(x)$. Because of the analogy with (1.1), it seems reasonable to classify this estimator within the CdMC area, but it should be noted that there is no apparent natural unbiased estimator Z of $f_n(x)$ for which $\mathbb{E}[Z \mid X_1, \dots, X_{n-1}] = f(x - S_{n-1})$. Of course, intuitively

$$\mathbb{P}(S_n \in dx \mid X_1, \dots, X_{n-1}) = f(x - S_{n-1})$$

but $\mathbb{I}(S_n \in dx)$ is not a well-defined r.v.! Nevertheless:

Proposition 2.1 *The estimator $f(x - S_{n-1})$ of $f_n(x)$ is unbiased.*

Proof: $\mathbb{E}f(x - S_{n-1}) = \int f_{n-1}(y)f(x - y) dy = f_n(x)$ □

Unbiasedness is in fact quite a virtue in itself, since the more traditional kernel and finite difference estimators are not so! It also implies consistency, i.e. that the average over R replications converges to the correct value $f_n(x)$ as $R \rightarrow \infty$.

Because of the lack of an obvious CrMC comparison, we shall not go into detailed properties of $\text{Var}[f(x - S_{n-1})]$; one expects such a study to be quite similar to the one in section 3 dealing with $\text{Var}[F(x - S_{n-1})]$. Instead, we shall give some numerical examples. Figure 1 illustrates the influence on the R of replications. For each of the four values $R = 2^8, 2^{10}, 2^{12}, 2^{14}$ we performed three sets of simulation, to assess the degree of randomness inherent in R being finite. Obviously, $R = 2^{14} \approx 16,000$ is almost perfect but the user may go for a substantially smaller value depending on how much the random variation and the smoothness is a concern.

A reasonable question is the comparison of CdMC and a kernel estimate of the form $k(x - S_n)$ for small or moderate R . In Figure 2, we considered our recurrent example of sum of lognormals, but took $R = 32$ for both of the estimators $f(x - S_{n-1})$ and $k(x - S_n)$, with k chosen as the $\text{Normal}(0, \sigma^2)$ density. The upper right panel is a histogram of the 32 simulated values of S_n and the upper left the CdMC estimator. The two lower panels are the kernel estimates, with an extreme high value $\sigma^2 = 10^2$ to the left and an extreme low $\sigma^2 = 10^{-2}$ to the right. A high value will produce oversmoothed estimates and a low undersmoothed ones with a marked effect of single observations. However, for R as small as 32 it is hard to assess what is a reasonable value of σ^2 . In fairness, we also admit that the single observation effect is clearly visible for the CdMC estimator and that it leads to estimates

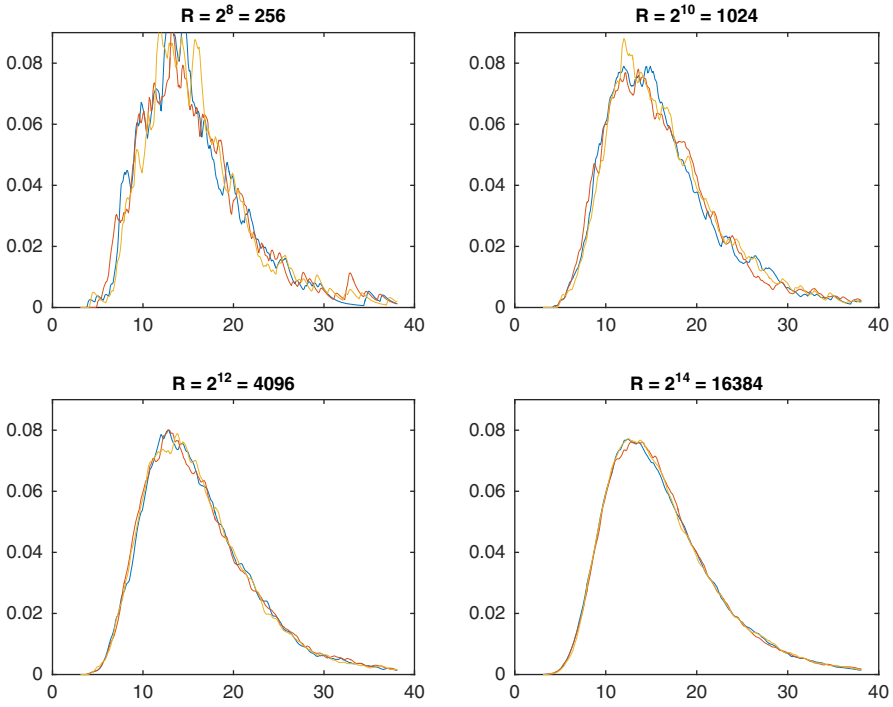


Figure 1. Estimated density of S_n as function of R .




Figure 2. Comparison with kernel smoothing.

which are undersmoothed. By this we mean more precisely that if f is in C^p for some $p = 0, 1, \dots$, then $f_n(x)$ is in C^{np} but $f(x - S_{n-1})$ only in C^p . In contrast, a normal kernel estimate is in C^∞ .

The first example of CdMC density estimation we know of is in Asmussen & Glynn (2007: 146), but in view of the simplicity of the idea, there may well have been earlier instances. We return to some further aspects of the methodology in section 7. For somewhat different uses of conditioning for smoothing, see L'Ecuyer & Perron (1994), Fu & Hu (1997) and L'Ecuyer & Lemieux (2000), sections 10.1–10.2.

3. Variance Reduction for the c.d.f.

CdMC always gives variance reduction. But as argued, it needs to be substantial for the procedure to be worthwhile. Further in many applications the right and/or left tail is of particular interest, so one may pay particular attention to the behaviour there.

Remark 3.1 That CdMC gives variance reduction in the tails can be seen intuitively by the following direct argument without reference to Rao–Blackwellization. The CrMC, respectively, the CdMC, estimators of $\bar{F}_n(x)$ are $\mathbb{1}(S_n > x)$ and $\bar{F}(x - S_{n-1})$, with second moments

$$\mathbb{E}\mathbb{1}(S_n > x)^2 = \mathbb{E}\mathbb{1}(S_n > x) = \int_{-\infty}^{\infty} f_{n-1}(y)\bar{F}(x-y) dy \tag{3.1}$$

$$\stackrel{X \geq 0}{=} \mathbb{P}(S_{n-1} > x) + \int_0^x f_{n-1}(y)\bar{F}(x-y) dy \tag{3.2}$$

$$\mathbb{E}\bar{F}(x - S_{n-1})^2 = \int_{-\infty}^{\infty} f_{n-1}(y)\bar{F}(x-y)^2 dy \tag{3.3}$$

$$\stackrel{X \geq 0}{=} \mathbb{P}(S_{n-1} > x) + \int_0^x f_{n-1}(y)\bar{F}(x-y)^2 dy \tag{3.4}$$

In the right tail (say), these second moments can be interpreted as the tails of the r.v.'s $S_{n-1} + X$, $S_{n-1} + X^*$ where X, X^* are independent of S_{n-1} and have tails \bar{F} and \bar{F}^2 . Since $\bar{F}^2(x)$ is of smaller order than $\bar{F}(x)$ in the right tail, the tail of $S_{n-1} + X^*$ should be of smaller order than that of $S_{n-1} + X$, implying the same ordering of the second moments. However, as n becomes large one also expects the tail of S_{n-1} to more and more dominate the tails of X, X^* so that the difference should be less and less marked. The analysis to follow will confirm these guesses.

A measure of performance which we consider is the ratio $r_n(x)$ of the CdMC variance to the CrMC variance:

$$r_n(x) = \frac{\mathbb{V}ar[\bar{F}(x - S_{n-1})]}{F_n(x)\bar{F}_n(x)} = \frac{\mathbb{V}ar[F(x - S_{n-1})]}{F_n(x)\bar{F}_n(x)} \tag{3.5}$$

(note that the two alternative expressions reflect that the variance reduction, is the same whether CdMC is performed for F itself or the tail \bar{F}).

To provide some initial insight, we examine in Figure 3, $r_n(x_n, z)$ as function of z where x_n, z is the z -quantile of S_n . In Figure 3(a), the underlying F is Pareto with tail $\bar{F}(x) = 1 / (1+x)^{3/2}$ and in

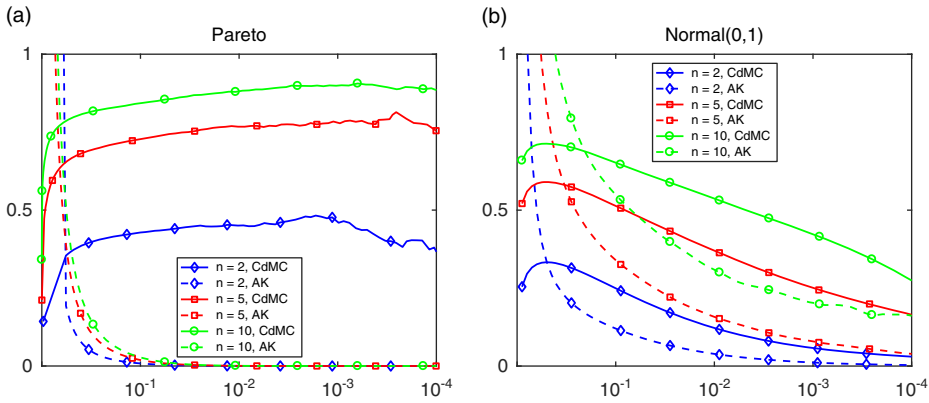


Figure 3. The ratio $r_n(z)$ in (3.5), with F Pareto in (a) and normal in (b).

Figure 3(b), it is standard normal. Both figures consider the cases of a sum of $n = 2, 5$ or 10 terms and use $R = 250,000$ replications of the vector Y_1, \dots, Y_{n-1} (variances are more difficult to estimate than means, therefore the high value of R). The dotted line for AK (the Asmussen-Kroese estimator, see section 4) may be ignored for the moment. The argument z on the horizontal axis is in \log_{10} -scale, and x_n, z was taken as the exact value for the normal case and the CdMC estimate for the Pareto case.

For the Pareto case in Figure 3(a), it seems that the variance reduction is decreasing in both x and n , yet in fact it is only substantial in the left tail. For the normal case, note that there should be symmetry around $x = 0$, corresponding to $z(x) = 1/2$ with base-10 logarithm -0.30 . This is confirmed by the figure (though the feature is of course somewhat disguised by the logarithmic scale). In contrast to the Pareto case, it seems that the variance reduction is very big in the right (and therefore also left) tail but also that it decreases as n increases.

We proceed to a number of theoretical results supporting these empirical findings. They all use formulas (3.3) and (3.4) for the second moments of the CdMC estimators. For the exponential distribution, the calculations are particularly simple:

Example 3.2 Assume $\bar{F}(x) = e^{-x}$, $n = 2$. Then $\mathbb{P}(X_1 + X_2 > x) = xe^{-x} + e^{-x}$ and (5) takes the form

$$\bar{F}(x) + \int_0^x e^{-y} e^{-2(x-y)} dy = e^{-x} + e^{-2x}(e^x - 1) = 2e^{-x} - e^{-2x}$$

and so for the right tail:

$$r_2(x) = \frac{2e^{-x} - e^{-2x} - (xe^{-x} + e^{-x})^2}{(xe^{-x} + e^{-x})(1 - xe^{-x} - e^{-x})}$$

For $x \rightarrow \infty$, this gives

$$r_2(x) = \frac{2e^{-x} + o(e^{-x})}{xe^{-x} + o(xe^{-x})} = \frac{2}{x}(1 + o(1)) \rightarrow 0$$

In the left tail $x \rightarrow 0$, Taylor expansion give that up to the third-order term

$$2e^{-x} - e^{-2x} \sim 1 - x^2 + x^3, \quad xe^{-x} + e^{-x} = 1 - x^2/2 + x^3/3$$

and so

$$\begin{aligned}
 r_2(x) &\sim \frac{1-x^2+x^3-(1-x^2/2+x^3/3)^2}{(1-x^2+x^3)(x^2/2-x^3/6)} \\
 &\sim \frac{1-x^2+x^3-(1-x^2+2x^3/3)}{x^2/2} = \frac{2x}{3} \rightarrow 0 \quad \diamond
 \end{aligned}$$

The relation $r_n(x) \rightarrow 0$ in the left tail (i.e. as $x \rightarrow 0$) in the exponential example is in fact essentially a consequence of the support being bounded to the left:

Proposition 3.3 *Assume $X > 0$ and that the density $f(x)$ satisfies $f(x) \sim cx^p$ as $x \rightarrow 0$ for some $c > 0$ and some $p > -1$. Then $r_n(x) \sim dx^{p+1}$ as $x \rightarrow 0$ for some $0 < d = d(n) < \infty$.*

The following result explains the right tail behaviour in the Pareto example and shows that this extends to other standard heavy-tailed distributions like the lognormal or Weibull with decreasing failure rate (for subexponential distributions, see, e.g. Embrechts *et al.*, 1997):

Proposition 3.4 *Assume $X > 0$ is subexponential. Then $r_n(x) \rightarrow 1 - 1/n$ as $x \rightarrow \infty$.*

For light tails, Example 3.2 features a different behaviour in the right tail, namely $r_n(x) \rightarrow 0$. Here is one more such light-tailed example:

Proposition 3.5 *If X is standard normal, then $r_n(x) \rightarrow 0$ as $x \rightarrow \infty$. More precisely,*

$$r_n(x) \sim \frac{1}{x} \sqrt{\frac{2n-1}{n\pi}} e^{-x^2 / [2n(2n-1)]}$$

The proofs of Propositions 3.3–3.5 are in the Appendix.

To formulate a result of type $r_n(x) \rightarrow 0$ as $x \rightarrow \infty$ in a sufficiently broad class of light-tailed F encounters the difficulty that the general results giving the asymptotics of $\mathbb{P}(S_n > x)$ as $x \rightarrow \infty$ with n fixed are somewhat involved (the standard light-tailed asymptotics is for $\mathbb{P}(S_n > bn)$ as $n \rightarrow \infty$ with b fixed, cf. e.g. Jensen, 1995). It is possible to obtain more general versions of Example 3.2 for close-to-exponential tails by using results of Cline (1986) and of Proposition 3.5 for thinner tails by involving Balkema *et al.* (1993). However, the adaptation of Balkema *et al.* (1993) is rather technical and can be found in Asmussen *et al.* (2017).

One may note that the variance reduction is so moderate in the range of z considered in Figure 3(b) that CdMC may hardly be worthwhile for light tails except for possibly very small n . If variance reduction is a major concern, the obvious alternative is to use the standard IS algorithm which uses exponential change of measure (ECM). The r.v.'s X_1, \dots, X_n are here generated from the exponentially twisted distribution with density $f_\theta(x) = e^{\theta x} f(x) / \mathbb{E}e^{\theta X}$, where θ should be chosen such that $\mathbb{E}_\theta S_n = x$. The estimator of $\mathbb{P}(S_n > x)$ is

$$e^{-\theta S_n} [\mathbb{E}e^{\theta X}]^n \mathbb{1}(S_n > x) \tag{3.6}$$

see Asmussen & Glynn (2007: 167–169) for more detail. Further variance reduction would be obtained by applying CdMC to (3.6) as implemented in the following example.

Table 1. Variance reduction for sum of 10 gamma r.v.'s.

z	x_z	CdMC	IS-ECM	CdMC + IS-ECM
0.95	39.5	0.628	0.121	0.048
0.99	44.2	0.561	0.032	0.010

Note: CdMC, conditional Monte Carlo; IS, importance sampling; ECM, exponential change of measure.

Example 3.6 To illustrate the potential of the IS-ECM algorithm, we consider the sum of $n = 10$ r.v.'s which are $\gamma(3,1)$ at the $z = 0.95, 0.99$ quantiles x_z . The exponentially twisted distribution is $\gamma(3, 1 - \theta)$ and $\mathbb{E}_\theta S_n = x$ means $3/(1 - \theta) = x$, i.e. $\theta = 1 - 3/(x/n)$. With $R = 100,000$ replications, we obtained the values of $r_n(x)$ at the z quantiles for $z = 0.95, 0.99$ given in Table 1. It is seen that IS-ECM indeed performs much better than CdMC, but that CdMC is also moderately useful for providing some further variance reduction. \diamond

A further financially relevant implementation of the IS-ECM algorithm is in Asmussen *et al.* (2016) for lognormal sums. It is unconventional because it deals with the left tail (which is light) rather than the right tail (which is heavy) and because the ECM is not explicit but done in an approximately efficient way. Another IS algorithm for the left lognormal sum tail is in Gulisashvili & Tankov (2016), but the numerical evidence of Asmussen *et al.* (2016) makes its efficiency somewhat doubtful.

4. The AK Estimator

The idea underlying the estimator $Z_{AK}(x)$ of Asmussen & Kroese (2006) for $z = z(x) = \mathbb{P}(S_n > x)$ is to combine an exchangeability argument with CdMC. More precisely (for convenience assuming existence of densities to exclude multiple maxima) one has $z = n \mathbb{P}(S_n > x, M_n = X_n)$, where $M_k = \max_{i \leq k} X_i$. Applying CdMC with $\mathcal{F} = \sigma(X_1, \dots, X_{n-1})$ to this expression the estimator comes out as

$$Z_{AK}(x) = n\bar{F}(M_{n-1} \vee (x - S_{n-1})) \tag{4.1}$$

There has been a fair amount of follow-up work on Asmussen & Kroese (2006) and sharpened versions, see in particular Hartinger and Kortschak (2009), Chan & Kroese (2011), Asmussen *et al.* (2011), Asmussen & Kortschak (2012, 2015), Ghamami & Ross (2012) and Kortschak & Hashorva (2013). In summary, the state-of-the-art is that Z_{AK} not only has bounded relative error (BdRelErr) but in fact vanishing relative error in a wide class of heavy-tailed distributions. Here the relative (squared) error is the traditional measure of efficiency in the rare-event simulation literature, defined as the ratio $r_n^{(2)}(z)$ (say) between the variance and the square of the probability z in question (note that $r_n(x)$ is defined similarly in (3.5) but without the square in the denominator). BdRelErr means $\limsup_{z \rightarrow 0} r_n^{(2)}(z) < \infty$ and is usually considered the most one can hope for, cf. Asmussen & Glynn (2007: VI.1). The following sharp version of the efficiency of Z_{AK} follows, e.g., from Asmussen & Kortschak (2012, 2015).

Theorem 4.1 Assume that the distribution of X is either regularly varying, lognormal or Weibull with tail e^{-x^β} , where $0 < \beta < \log(3/2)/\log 2 \approx 0.585$. Then there exists constants $\gamma > 0$ and $c < \infty$ depending on the distributional parameters such that

$$\text{Var} Z_{AK}(x) \sim cx^{-\gamma} \mathbb{P}(S_n > x)^2 \quad \text{as } x \rightarrow \infty$$

The efficiency of the AK estimator for heavy-tailed F is apparent from Figure 3(a), where it outperforms simple CdMC. For light-tailed F it has been noted that Z_{AK} does not achieve BdRelErr, and

presumably this is the reason it seems to have been discarded in this setting. For similar reasons as in section 3, we shall not go into a general treatment of the efficiency of the AK estimator for light-tailed F , but only present the results for two basic examples when $n = 2$.

Example 4.2 Assume $n = 2$, $f(x) = e^{-x}$. Then $M_{n-1} = X_{n-1} = X_1$ and $M_{n-1} > x - X_{n-1}$ precisely when $X_1 > x/2$. This gives

$$\begin{aligned} \frac{1}{4} \mathbb{E}Z_{AK}(x)^2 &= \int_0^{x/2} e^{-2(x-y)} e^{-y} dy + \int_{x/2}^{\infty} e^{-2y} e^{-y} dy \\ &= e^{-2x} \left(e^{x/2} - 1 \right) + \frac{1}{3} e^{-3x/2} \sim \frac{4}{3} e^{-3x/2}, \quad x \rightarrow \infty \end{aligned}$$

Compared to CrMC, this corresponds to an improvement of the second moment by a factor of order $e^{-x/2}/x$.

Example 4.3 Let $n = 2$ and let F be normal(0,1). Calculations presented in the Appendix then give

$$\mathbb{V}arZ_{AK}(x)^2 \sim \frac{64}{3x^3(2\pi)^{3/2}} e^{-3x^2/8}, \quad x \rightarrow \infty \tag{4.2}$$

Compared to CrMC, this corresponds to an improvement of the error by a factor of order $e^{-5x^2/8}$.

As discussed in section 3, the variance reduction obtained via Z_{AK} is reflected in the improved estimates of the VaR. For the expected shortfall, Z_{AK} -based algorithms are discussed in Hartinger & Kortschak (2009). They assume $VaR_\alpha(S_n)$ to be known, but the discussion of section 5 covers how to give confidence intervals if it is estimated.

Remark 4.4 For rare-event problems similar or related to that of estimating $\mathbb{P}(S_n > x)$, a number of alternative algorithms with similar efficiency as Z_{AK} have later been developed, see, e.g., Dupuis *et al.* (2007), Juneja (2007) and Blanchet & Glynn (2008). Some of these have the advantage of a potentially broader applicability, though Z_{AK} remains the one which is most simple.

5. VaR

The VaR $VaR_\alpha(S_n)$ of S_n at level α is intuitively defined as the number such that the probability of a loss larger than $VaR_\alpha(S_n)$ is $1 - \alpha$. Depending on whether small or large values of S_n mean a loss, there are two forms used, the actuarial $VaR_\alpha(S_n)$ defined as the α -quantile $q_{\alpha, n}$ and the financial $VaR_\alpha(S_n)$ defined as $-q_{1-\alpha, n}$. Typical values of α are 0.95 and 0.99 but smaller values occur in Basel II for certain types of business lines. We use here the actuarial definition and assume F to be continuous to avoid technicalities associated with $\mathbb{P}(S_n = VaR_\alpha(S_n)) > 0$. Also, since α, n are fixed, we write just $q = q_{\alpha, n}$.

The CrMC estimate uses R simulated values $S_n^{(1)}, \dots, S_n^{(R)}$ and is taken as the α -quantile $\hat{q}_{Cr} = \hat{F}_{n;R}^{-1}(\alpha)$ of the empirical c.d.f.

$$\hat{F}_{n;R}(x; S_n) = \frac{1}{R} \sum_{r=1}^R \mathbb{1} \left(S_n^{(r)} \leq x \right)$$

(we ignore here and in the following the issues connected with the ambiguity in the choice of $\widehat{F}_{n;R}^{-1}(\alpha)$ connected with the discontinuity of $\widehat{F}_{n;R}$; asymptotically, these play no role). Thus \widehat{q}_{Cr} is more complicated than an average of i.i.d. r.v.'s but nevertheless there is a CLT

$$\sqrt{R}(\widehat{q}_{Cr}-q) \rightarrow \mathcal{N}(0, \sigma_{Cr}^2) \quad \text{where } \sigma_{Cr}^2 = \frac{\alpha(1-\alpha)}{f_n(q)^2} \tag{5.1}$$

see, e.g., Serfling (1980). Thus confidence intervals require an estimate of $f_n(q)$, an issue about which Glynn (1996) writes that “the major challenge is finding a good way of estimating $f_n(q)$, either explicitly or implicitly” (without providing a method for doing this!) and Glasserman *et al.* (2000) that “estimation of $f_n(q)$ is difficult and beyond the scope of this paper”.

When confidence bands for the VaR are given, a common practice is therefore to use the bootstrap method. However, in our sum setting, CdMC easily gives $f_n(q)$, as outlined in section 2. In addition, the method provides some variance reduction because of its improved estimates of the c.d.f.:

Proposition 5.1 Define \widehat{q}_{Cond} as the solution of

$$\widehat{F}_{n;R}^{Cond}(\widehat{q}_{Cond}) = \alpha \quad \text{where } \widehat{F}_{n;R}^{Cond}(x) = \frac{1}{R} \sum_{r=1}^R F(x - S_{n-1}^{(r)}) \tag{5.2}$$

If F admits a density f that is either monotone or differentiable with f' bounded, then

$$\sqrt{R}(\widehat{q}_{Cond}-q) \xrightarrow{D} \mathcal{N}(0, \sigma_{Cond}^2) \quad \text{where } \sigma_{Cond}^2 = \frac{\text{Var}[F(q-S_{n-1})]}{f_n(q)^2} < \sigma_{Cr}^2$$

Proof: An intuitive explanation on how to as here to deal with CLTs for roots of equations is given in Asmussen & Glynn (2007: III.4). Chu & Nakayama (2012) and Nakayama (2014) give rigorous treatments of problems closely related to the present one but the proofs are quite advanced, building on deep results of Bahadur (1966) and Ghosh (1971). We therefore give a short, elementary and self-contained derivation, even if Proposition 5.1 is a special case of Nakayama (2014).

The key step is to show

$$\widehat{F}_{n;R}^{Cond}(\widehat{q}) - \widehat{F}_{n;R}^{Cond}(q) = (\widehat{q}-q)f_n(q)(1+o(1)) \tag{5.3}$$

In fact, $\widehat{F}_{n;R}^{Cond}(\widehat{q}_{Cond}) = \alpha = F_n(q)$ then gives

$$\begin{aligned} 0 &= \widehat{F}_{n;R}^{Cond}(\widehat{q}) - \widehat{F}_{n;R}^{Cond}(q) + \widehat{F}_{n;R}^{Cond}(q) - F_n(q) \\ &= (\widehat{q}-q)f_n(q)(1+o(1)) + \frac{\sqrt{\text{Var}[F(q-S_{n-1})]}}{\sqrt{R}} V(1+o(1)) \end{aligned}$$

with $V \sim \mathcal{N}(0, 1)$, from which the desired conclusion follows.

For (5.3), note that f_n is automatically continuous and let $\widehat{f}_{n;R}^{Cond}(x) = \sum_{r=1}^R f(x - S_{n-1}^{(r)})/R$ be the CdMC density estimator. Since $\widehat{F}_{n;R}^{Cond}(\cdot)$ is differentiable with derivative $\widehat{f}_{n;R}^{Cond}(\cdot)$, we get

$$\widehat{F}_{n;R}^{Cond}(\widehat{q}) - \widehat{F}_{n;R}^{Cond}(q) = (\widehat{q}-q)\widehat{f}_{n;R}^{Cond}(q^*)$$

Table 2. Value-at-Risk estimates for lognormal example.

	$\alpha=0.95$	$\alpha=0.99$
$n=5$	17.0 ± 0.2	25.4 ± 0.6
$n=10$	29.0 ± 0.2	39.9 ± 0.7
$n=25$	61.0 ± 0.3	74.6 ± 0.8
$n=50$	109.8 ± 0.4	127.2 ± 1.0

for some q^* between \hat{q} and q . Assume first f is monotone, say non-increasing. Since $|\hat{q}-q| \leq \epsilon$ for all large R , we then also have

$$\hat{f}_{n;R}^{\text{Cond}}(q+\epsilon) \leq \hat{f}_{n;R}^{\text{Cond}}(q^*) \leq \hat{f}_{n;R}^{\text{Cond}}(q-\epsilon)$$

for such R , and the consistency of $\hat{f}_{n;R}^{\text{Cond}}(\cdot)$ then gives $\hat{f}_{n;R}^{\text{Cond}}(q^*) \rightarrow f_n(q)$ and (5.3). Assume next f is differentiable with $\sup |f'| < \infty$. Arguing as above, we then get

$$\hat{f}_{n;R}^{\text{Cond}}(q^*) = \hat{f}_{n;R}^{\text{Cond}}(q) + (q^* - q) \frac{1}{R} \sum_{r=1}^R f'(q^{**} - S_{n-1}^{(r)}) = \hat{f}_{n;R}^{\text{Cond}}(q) + (q^* - q)O(1)$$

for some q^{**} between \hat{q} and q , which again gives the desired conclusion. □

Remark 5.2 At a first sight, the more obvious way to involve CdMC would have been to give the VaR estimates as the average over R replications of the α -quantile \tilde{q} in the conditional distribution of S_n given S_{n-1} . However, this does not provide the correct answer and in fact introduces a bias that does not disappear for $R \rightarrow \infty$ as it does for $\hat{q} = \hat{F}_{n;R}^{-1}(\alpha)$ and \hat{q}_{Cond} . For a simple example illustrating this, consider the i.i.d. Normal(0, 1)-setting. Here $\tilde{q} = S_n - 1 + z_\alpha$ where $z_\alpha = \Phi^{-1}(\alpha)$ with expectation z_α but $\sqrt{n}z_\alpha$ is the correct answer! ◇

Example 5.3 As illustration, we used CdMC with $R = 50,000$ replications to compute $\text{VaR}_\alpha(S_n)$ and the associated confidence interval for the sum of $n = 5, 10, 25, 50$ Lognormal(0, 1) r.v.'s. The results are in Table 2.

6. Expected Shortfall

An alternative risk measure receiving much current attention is the expected shortfall (also called conditional VaR). For continuous F , this takes the form (cf. McNeil *et al.*, 2015: 70):

$$\text{ES}_\alpha(S_n) = \mathbb{E}[S_n \mid S_n \geq q] = q + \frac{m_n(q)}{1-\alpha} \tag{6.1}$$

where as above $q = \text{VaR}_\alpha(S_n)$ and

$$m_n(z) = \mathbb{E}[S_n - z]^+ = \int_z^\infty \bar{F}_n(y) dy \tag{6.2}$$

The obvious CrMC algorithm for estimating $\text{ES}_\alpha(S_n)$ is to first compute the estimate $\hat{q}_{\text{Cr}} = \hat{F}_{n;R}^{-1}(\alpha)$ as above and next either (i) perform a new set of simulations with R_1 replications of $Z_1 = [S_n - \hat{q}_{\text{Cr}}]^+$, using the resulting average as estimator of $m_n(q)$ (consistency holds in the limit R ,

$R_1 \rightarrow \infty$), or (ii) use the already simulated $S_n^{(1)}, \dots, S_n^{(R)}$ to estimate $m_n(q)$ as the corresponding empirical value

$$\widehat{m}_{n;R}^{Cr}(\widehat{q}_{Cr}) = \frac{1}{R} \sum_{r=1}^R [S_n^{(r)} - \widehat{q}_{Cr}]^+ \tag{6.3}$$

We shall not pay further attention to (i), but consider a broader class of estimators than in (ii), covering both CdMC and other examples. The issue is how to provide confidence intervals. This is non-trivial already for the CrMC scheme (ii), and since we are not aware of a sufficiently close general reference we shall give some detail here.

In this broader setting, we assume that the simulation generates an estimate $\widehat{F}_{n;R}^*(x)$ of $F_n(x)$ and an estimate $\widehat{m}_{n;R}^*(x)$ of $m_n(x)$ in a x -range asymptotically covering q , such that these estimates are connected by

$$\widehat{m}_{n;R}^*(x) = \int_x^\infty \widehat{F}_{n;R}^*(y) dy \tag{6.4}$$

where $\widehat{F}_{n;R}^*(x) = 1 - \widehat{F}_{n;R}^*(x)$. Precisely as in (ii), we then compute $\widehat{q}_* = \widehat{F}_{n;R}^* - 1(\alpha)$ and estimate $e = \text{ES}_\alpha(S_n)$ as

$$\widehat{e}_* = \widehat{q}_* + \frac{\widehat{m}_{n;R}^*(\widehat{q}_*)}{1 - \alpha} \tag{6.5}$$

Example 6.1 In many main examples, $\widehat{F}_{n;R}^*(x)$ and $\widehat{m}_{n;R}^*(x)$ have the form

$$\frac{1}{R} \sum_{r=1}^R \phi_F(x, V_r), \quad \frac{1}{R} \sum_{r=1}^R \phi_m(x, V_r) \tag{6.6}$$

where V_1, \dots, V_R are i.i.d. replicates of a random vector V simulated from some probability measure $\widetilde{\mathbb{P}}$ and ϕ_F, ϕ_m are functions satisfying

$$\widetilde{\mathbb{E}}\phi_F(x, V) = F_n(x), \quad \widetilde{\mathbb{E}}\phi_m(x, V) = m_n(x) \tag{6.7}$$

The requirement (6.4) then means

$$\phi_m(x, V) = \int_x^\infty \phi_F(y, V) dy \tag{6.8}$$

Special cases:

- a. CrMC where $V = S_n, \widetilde{\mathbb{P}} = \mathbb{P}, \phi_F(x, s) = \mathbb{1}(s \leq x), \phi_m(x, s) = (s - x)^+$.
- b. CdMC where $V = S_{n-1}, \widetilde{\mathbb{P}} = \mathbb{P}, \phi_F(x, s) = F(x - s), \phi_m(x, s) = m(x - s)$ where $m(x) = m_1(x) = \mathbb{E}(X - x)^+$ (typically explicitly available in contrast to $m_n(x)$!).
- c. IS where $V = (S_n, L), \widetilde{\mathbb{P}}$ is the measure w.r.t. which X_1, \dots, X_n are i.i.d. with density $\widetilde{f} \neq f, L = \prod_{i=1}^n f(X_i) / \widetilde{f}(X_i)$ is the likelihood ratio and $\phi_F(x, s, \ell) = \mathbb{1}(s \leq x)\ell, \phi_m(x, s) = (s - x)^+ \ell$.
- d. The AK estimator from section 4, leading to $V = (M_{n-1}, S_{n-1}), \widetilde{\mathbb{P}} = \mathbb{P}, \phi_F(x, z, s) = \overline{F}(z \wedge (x - s)), \phi_m(x, z, s) = t_{AK}(z, x - s)$ where $t_{AK}(z, x) = \mathbb{E}[(X - x)^+ ; X > z]$.

The verification of (6.4) is in all cases an easy consequence of the identity $(v - x)^+ = \int_x^\infty \mathbb{1}(y \leq v) dy$; the most difficult case is (d) where (6.4) follows from

$$t_{AK}(z, x - s) = \mathbb{E}[(X + s - x)^+ ; X > z] = \mathbb{E}\left[\int_{x-s}^\infty \mathbb{1}(X > y, X > z) dy\right]$$

$$= \mathbb{E} \left[\int_x^\infty \mathbb{1}(X > y-s, X > z) dy \right] = \int_x^\infty \bar{F}(z \vee (y-s)) dy$$

The IS in (c) can be combined with (b) or (d) in obvious ways. Examples not covered are: (e) regression-adjusted control variates (Asmussen & Glynn, 2007: V.2) where (6.6) fails; (f) level-dependent IS where the measure $\tilde{\mathbb{P}}$ in (c) depends on x . This may be a quite natural situation, cf. Example 3.6. \diamond

For asymptotics, we assume that

$$\hat{F}_{n;R}^*(x) \sim F_n(x) - Z(x) / \sqrt{R}, \quad \bar{F}_{n;R}^*(x) \sim \bar{F}_n(x) + Z(x) / \sqrt{R} \tag{6.9}$$

as $R \rightarrow \infty$ for a suitable Gaussian process Z . For example, $\mathbb{V}ar Z(x) = F_n(x)\bar{F}_n(x)$ for the empirical c.d.f. For other examples, in particular CdMC, $\mathbb{V}ar Z(x)$ is typically not explicit but must be estimated from the simulation output and varies from case to case. From section 5, one then expects that

$$\sqrt{R}(\hat{q}_* - q) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_*^2(q)) \quad \text{where } \sigma_*^2(q) = \frac{\mathbb{V}ar Z(q)}{f(q)^2} \tag{6.10}$$

as has been verified for CdMC; for IS, see Sun & Hong (2010) and Hong *et al.* (2014). We shall also need

$$\sqrt{R}(m_{n;R}^*(x) - m_n(x)) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_*^2(m_n(x))) \tag{6.11}$$

for some $\sigma_*^2(m_n(x))$; this is obvious in the setting of Example 6.1. We then get the following result, which in particular applies to CdMC for an i.i.d. sum:

Proposition 6.2 *Subject to (6.9), it holds that*

$$\sqrt{R}(\hat{e}_* - e) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_*^2(e)) \quad \text{where } \sigma_*^2(e) = \frac{\sigma_*^2(m_n(x))}{(1-\alpha)^2}$$

Proof: By (6.4),

$$\begin{aligned} \hat{e}_* &= \hat{q}_* + \frac{1}{1-\alpha} \int_{\hat{q}_*}^q \bar{F}_{n;R}^*(x) dx + \frac{1}{1-\alpha} \int_q^\infty \bar{F}_{n;R}^*(x) dx \\ &\approx \hat{q}_* + (q - \hat{q}_*) + \frac{\hat{m}_{n;R}^*(q)}{1-\alpha} = q + \frac{m_n(q)}{1-\alpha} + \frac{\hat{m}_{n;R}^*(q) - m_n(q)}{1-\alpha} \\ &= e + \frac{\hat{m}_{n;R}^*(q) - m_n(q)}{1-\alpha} \end{aligned} \quad \square$$

Summarising the algorithm for the CdMC case:

1. Simulate $S_{n-1}^{(1)}, \dots, S_{n-1}^{(R)}$.
2. Compute \hat{q}_{Cond} as solution of $\frac{1}{R} \sum_{r=1}^R \bar{F}(q - S_{n-1}^{(r)}) = 1 - \alpha$.
3. Let $\hat{e}_{\text{Cond}} = \hat{q}_{\text{Cond}} + \frac{1}{R} \sum_{r=1}^R m(\hat{q}_{\text{Cond}} - S_{n-1}^{(r)})$ where $m(x) = \mathbb{E}(X - x)^+$.
4. Compute the empirical variance s^2 of the $m(\hat{q}_{\text{Cond}} - S_{n-1}^{(r)})$, $r = 1, \dots, R$.
5. Return the 95% confidence interval $\hat{e}_{\text{Cond}} \pm 1.96s / \sqrt{R}$.

7. Averaging

The idea of using $f(x - S_{n-1})$ and $F(x - S_{n-1})$ as estimators of f_n , respectively, F_n , has an obvious asymmetry in that X_n has a special role among X_1, \dots, X_n by being the one that is not simulated but handled via its conditional expectation given the rest. An obvious idea is therefore to repeat the procedure with X_n replace by X_k and average over $k=1, \dots, n$. This leads to the alternative estimators

$$\frac{1}{n} \sum_{k=1}^n f(x - S_n + X_k), \quad \text{respectively,} \quad \frac{1}{n} \sum_{k=1}^n F(x - S_n + X_k) \tag{7.1}$$

Figure 4 illustrates the procedure for our recurrent example of estimating the density of the sum of $n = 10$ lognormals.

It is seen that the averaging procedure has the obvious advantage of producing a smoother estimate. This may be particularly worthwhile for small sample sizes R , as illustrated in Figure 5. Here $R = 32$, the upper panel gives the histogram of the simulated 32 values of S_n and the corresponding simple CdMC estimate, whereas the lower panel has the averaged CdMC estimate (7.1) left. These have been supplemented with the density estimate:

$$\frac{1}{n(n-1)} \sum_{k \neq \ell} f^{*2}(x - S_n + X_k + X_\ell) \tag{7.2}$$

with f^{*2} evaluated by numerical integration (which is feasible for the sum of just $n=2$ r.v.'s). The idea comes from observing that the peaks of f are quite visible in the plot using $f(x - S_{n-1})$, whereas convolution will produce smoother peaks in f^{*2} . The improvement is quite notable. However, we have not pursued this line of thought any further.

The overall performance of the idea involves two further aspects, computational effort and variance.

To asses the performance in terms of variance, consider estimation of the c.d.f. F and let

$$\begin{aligned} \omega^2 &= \text{Var}[F(x - S_{n-1})] = \text{Var}[F(x - S_n + X_k)] \\ \rho &= \text{Corr}(F(x - S_n + X_k), F(x - S_n + X_\ell)), \quad k \neq \ell \end{aligned}$$




Figure 4. $f(x - S_{n-1})$ dotted, (7.1) solid. (a) $R = 128$, (b) $R = 1,024$.

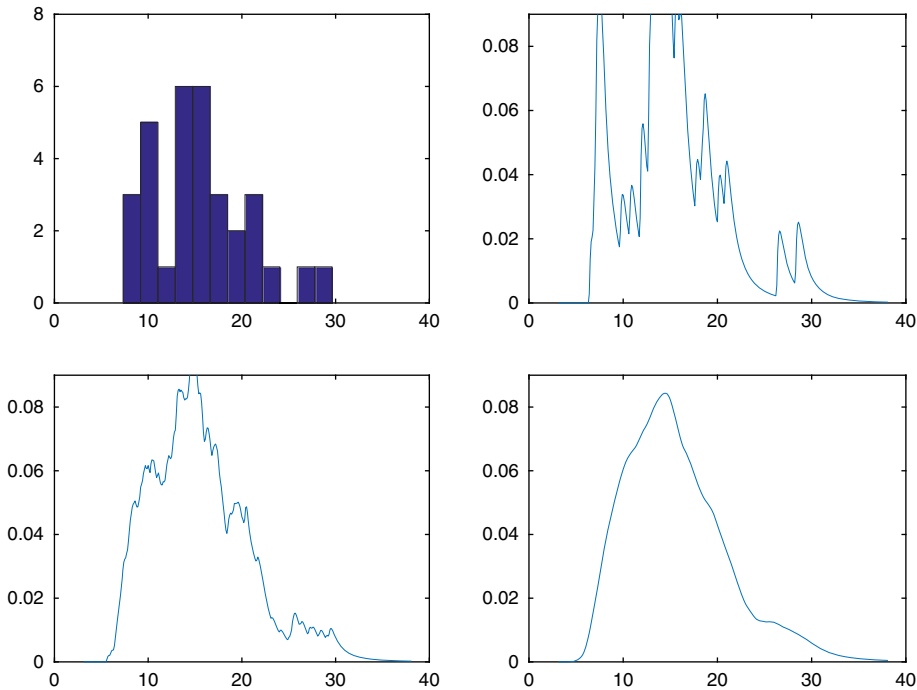


Figure 5. $R=32$. Upper panel simulated data left, $f(x - S_{n-1})$ right. Lower panel (7.1) left, (7.2) right.

Table 3. Comparison of simple and averaged conditional Monte Carlo.

	$n=5$	$n=10$	$n=25$	$n=50$
vf_1	0.66 0.67	0.77 0.72	0.82 0.82	0.90 0.89
vf_n	0.49 0.50	0.64 0.60	0.74 0.74	0.84 0.82
ρ	0.67 0.67	0.81 0.82	0.90 0.89	0.93 0.93

Then ω^2 is the variance of the simple CdMC estimator, whereas that of the averaged one is $\omega^2[1/n + (1 - 1/n)\rho]$. Here $\rho=0$ for $n=2$, but one expects ρ to increase to 1 as n increases. The implication is that there is some variance reduction, but presumably it is only notable for small n . This is illustrated in Table 3, giving some numbers for the sum of Lognormal(0,1) r.v.'s. Within each column, the first entry vf_1 is the variance reduction factor $r_n(q_\alpha)$ for simple CdMC computed at the estimated α -quantile of S_n as given by Example 3.6, the second the same for averaged CdMC. For each entry, the two numbers correspond to the two values of α . The last column gives the empirical estimate of the correlation ρ as defined above.

When assessing the computational efficiency, it seems reasonable to compare with the alternative of using simple CdMC with a larger R than the one used for averaging. The choice between these two alternatives involves, however, features varying from case to case. Averaging has an advantage if computation of densities is less costly than random variate generation, a disadvantage the other way round. In any case, averaging seems only worthwhile if the number R of replications is rather small.

8. Dependence

The current trend in dependence modelling is to use copulas and we shall here show some implementations of CdMC to this point of view. Among the many references in the area, Whelan (2004), McNeil (2006), Cherubini *et al.* (2004), Wu *et al.* (2007) and Mai & Scherer (2012) are of particular relevance for the following.

In the sum setting, we consider again the same marginal distribution F of X_1, \dots, X_n . Let (U_1, \dots, U_n) be a random vector distributed according to the copula in question. Then S_n can be simulated by taking $X_i = F^{-1}(U_i)$, i.e. we can write

$$S_n = F^{-1}(U_1) + \dots + F^{-1}(U_n) \tag{8.1}$$

To estimate the c.d.f. or p.d.f. of S_n we then need the conditional c.d.f. and p.d.f. $F_{n|1:n-1}, f_{n|1:n-1}$ of X_n given a suitable σ -field \mathcal{F} w.r.t. which U_1, \dots, U_{n-1} are measurable. Indeed, then

$$F_{n|1:n-1}(x - S_{n-1}), \quad \text{respectively, } f_{n|1:n-1}(x - S_{n-1}) \tag{8.2}$$

are unbiased estimates (which can of course only be used in practice if $F_{n|1:n-1}, f_{n|1:n-1}$ have accessible expressions).

For a first example where $F_{n|1:n-1}, f_{n|1:n-1}$ are available, we consider Gaussian copulas. This is the case $U_i = \Phi(Y_i)$ where $Y_n = (Y_1 \dots Y_n)^T$ is a multivariate normal vector with standard normal marginals and a general correlation matrix C . In block-partitioned notation, we can write

$$C = \begin{pmatrix} C_{n-1} & C_{n-1,n} \\ C_{n,n-1} & 1 \end{pmatrix}$$

where C_{n-1} is $(n-1) \times (n-1)$, $C_{n-1,n}$ is $(n-1) \times 1$ and $C_{n,n-1} = C_{n-1,n}^T$.

Proposition 8.1 *Consider the Gaussian copula and define*

$$c_{n|1:n-1} = 1 - C_{n,n-1} C_{n-1}^{-1} C_{n-1,n}$$

$$\mu_{n|1:n-1} = C_{n,n-1} C_{n-1}^{-1} (Y_1 \dots Y_{n-1})^T$$

and $\mathcal{F} = \sigma(U_1, \dots, U_{n-1}) = \sigma(Y_1, \dots, Y_{n-1})$. Then

$$F_{n|1:n-1}(y) = \Phi\left(\frac{\Phi^{-1}(F(y)) - \mu_{n|1:n-1}}{\sqrt{c_{n|1:n-1}}}\right) \tag{8.3}$$

$$f_{n|1:n-1}(y) = \frac{\varphi\left(\frac{\Phi^{-1}(F(y)) - \mu_{n|1:n-1}}{\sqrt{c_{n|1:n-1}}}\right)}{\varphi(\Phi^{-1}(F(y))) \sqrt{c_{n|1:n-1}}} f(y) \tag{8.4}$$

Proof: We have

$$\begin{aligned} F_{n|1:n-1}(y) &= \mathbb{P}(X_n \leq y \mid \mathcal{F}) = \mathbb{P}(F^{-1}(\Phi(Y_n)) \leq y \mid \mathcal{F}) \\ &= \mathbb{P}(Y_n \leq \Phi^{-1}(F(y)) \mid \mathcal{F}) \end{aligned}$$

which reduces to (8.3) since the conditional distribution of Y_n given \mathcal{F} is normal with mean $\mu_{n|1:n-1}$ and variance $c_{n|1:n-1}$. (8.4) then follows by differentiation. \square

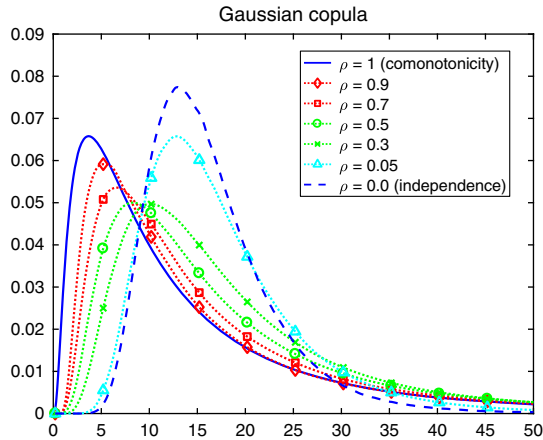


Figure 6. Density of a lognormal sum with an exchangeable Gaussian copula.

Example 8.2 For the Gaussian copula, the procedure of Proposition 8.1 was implemented for the sum S_n of $n=10$ lognormals, with the estimated densities given in Figure 6. The matrix C is taken as exchangeable, meaning that all off-diagonal elements are the same ρ , and various values of ρ are considered. Similar examples produced with somewhat different methods are in Botev *et al.* (2017).

As second example, we shall consider Archimedean copulas

$$\mathbb{P}(U_1 \leq u_1, \dots, U_n \leq u_n) = \psi(\phi(u_1) + \dots + \phi(u_n)) \tag{8.5}$$

where ψ is called the generator and ϕ is its functional inverse. Under the additional condition that the r.h.s. of (8.5) defines a copula for any n , it is known that ψ is the Laplace transform $\psi(s) = \mathbb{E}e^{-sZ}$ of some r.v. $Z > 0$, and we shall consider only that case. A convenient representation is then

$$(U_1, \dots, U_n) = (\psi(V_1/Z), \dots, \psi(V_n/Z)) \tag{8.6}$$

where V_1, \dots, V_n are i.i.d. standard exponential and independent of Z . See, e.g., Marshall & Olkin (1988).

Proposition 8.3 Define $\mathcal{F} = \sigma(V_1, \dots, V_{n-1}, Z)$. Then

$$F_{n|1:n-1}(y) = \exp\{-Z\phi(F(y))\} \tag{8.7}$$

$$f_{n|1:n-1}(y) = -Z\phi'(F(y)) \exp\{-Z\phi(F(y))\}f(y) \tag{8.8}$$

For the survival copula $(1 - \psi(V_1/Z), \dots, 1 - \psi(V_n/Z))$,

$$F_{n|1:n-1}(y) = 1 - \exp\{-Z\phi(\bar{F}(y))\} \tag{8.9}$$

$$f_{n|1:n-1}(y) = -Z\phi'(\bar{F}(y)) \exp\{-Z\phi(\bar{F}(y))\}f(y) \tag{8.10}$$

Proof: Formulas (8.8), (8.10) follow by straightforward differentiation of (8.7), (8.9) and (8.7) from:

$$\begin{aligned}
 F_{n|1:n-1}(y) &= \mathbb{P}(X_n \leq y \mid \mathcal{F}) = \mathbb{P}(F^{-1}(\psi(V_n/Z)) \leq y \mid \mathcal{F}) \\
 &= \mathbb{P}(V_n/Z \geq \phi(F(y)) \mid \mathcal{F}) = \exp\{-Z\phi(F(y))\}
 \end{aligned}$$

Similarly for the survival copula:

$$\begin{aligned}
 F_{n|1:n-1}(y) &= \mathbb{P}(F^{-1}(1-\psi(V_n/Z)) \leq y \mid \mathcal{F}) = \mathbb{P}(\psi(V_n/Z) \geq \bar{F}(y) \mid \mathcal{F}) \\
 &= \mathbb{P}(V_n/Z \leq \phi(\bar{F}(y)) \mid \mathcal{F}) = 1 - \exp\{-Z\phi(\bar{F}(y))\} \quad \square
 \end{aligned}$$

Some numerical results follow for lognormal sums with the two most common Archimedean copulas, Clayton and Gumbel.

Example 8.4 The Clayton copula corresponds to Z being γ with shape parameter α . Traditionally, the parameter is taken as $\theta = 1/\alpha$ and the scale (which is unimportant for the copula) chosen such as that $\mathbb{E}Z = 1$. This means that the generator is $\psi(t) = 1/(1+t\theta)^{1/\theta}$ with inverse $\phi(y) = (y^{-\theta} - 1)/\theta$.

The Clayton copula approaches independence as $\theta \downarrow 0$, i.e. $\alpha \uparrow \infty$, and approaches comonotonicity as $\theta \uparrow \infty$, i.e. $\alpha \downarrow 0$. The density of the sum S_n of $n = 10$ lognormals is in Figure 7(a) for the Clayton copula itself and in Figure 7(b) for the survival copula. The Clayton copula has tail independence in the right tail but tail dependence in the left, implying the opposite behaviour for the survival copula. Therefore, the survival copula may sometimes be the more interesting one for risk management purposes in the Clayton case.

Example 8.5 The Gumbel copula corresponds to Z being strictly α -stable with support $(0, \infty)$. Traditionally, the parameter is taken as $\theta = 1/\alpha$ and the scale chosen such that the generator is $\psi(t) = e^{-t^\alpha} = e^{-t^{1/\theta}}$, with inverse $\phi(y) = (-\log y)^\theta$. The Gumbel copula approaches comonotonicity as $\theta \rightarrow \infty$, i.e. $\alpha \downarrow 0$, whereas independence corresponds to $\theta = \alpha = 1$. It has tail dependence in the right tail but tail independence in the left.

The density of the sum S_n of $n = 10$ lognormals is in Figure 8(a) for the Gumbel copula itself and in Figure 8(b) for the survival copula. ◇

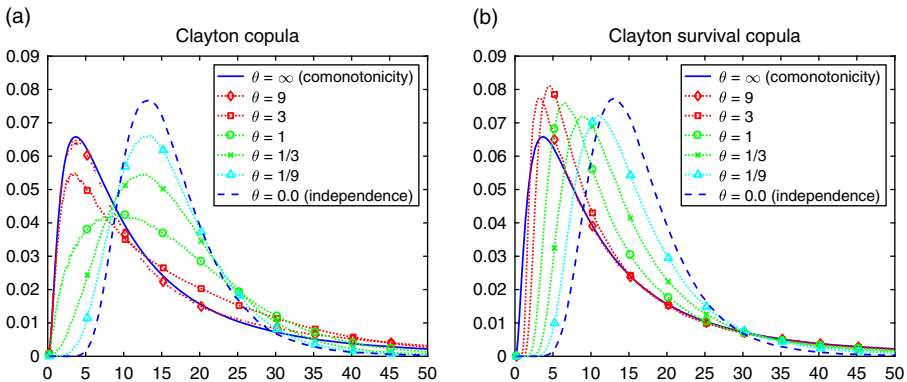


Figure 7. Density of lognormal sum with a Clayton copula.

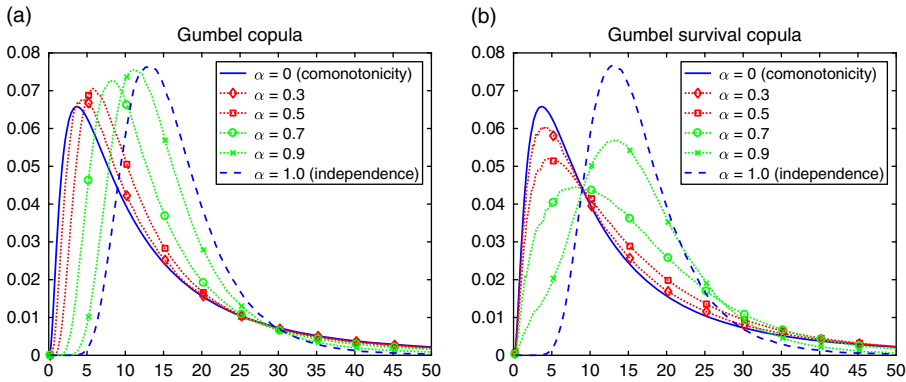


Figure 8. Density of lognormal sum with a Gumbel copula.

Remark 8.6 Despite the simplicity of the Marshall-Olkin representation, much of the literature on conditional simulation in the Clayton (and other Archimedean) copulas concentrates on describing the conditional distribution of U_n given U_1, \dots, U_{n-1} , see, e.g. Cherubini *et al.* (2004). Even with this conditioning, we point out that it may be simpler to just consider the conditional distribution of Z given $U_1 = u_1, \dots, U_{n-1} = u_{n-1}$. Namely, given $Z = z$ the r.v. $W_i = \phi(U_i) = V_i/z$ has density ze^{-zw_i} so that the conditional density must be proportional to the joint density:

$$f_Z(z)ze^{-zw_1} \dots ze^{-zw_{n-1}} \propto z^{1/\theta+n-2} \exp\{-z(1/\theta+w_1+\dots+w_{n-1})\}$$

where the last expression uses $Z \sim \gamma(1/\theta, 1/\theta)$. This gives in particular that for the Clayton copula the conditional distribution of Z given $U_1 = u_1, \dots, U_{n-1} = u_{n-1}$ is $\gamma(\alpha_n, \lambda_n)$ where

$$\alpha_n = 1/\theta + n - 1, \quad \lambda_n = 1/\theta + \phi(u_1) + \dots + \phi(u_{n-1}) \quad \diamond$$

Remark 8.7 Calculation of the VaR follows just the same pattern in the copula context as in the i.i.d. case, cf. the discussion around (5.1). One then needs to replace F with the conditional distribution in (8.2). Also the expected shortfall could be in principle calculated by replacing the $\mathbb{E}(X-x)^+$ from the i.i.d. case with the similar conditional expectation. However, in examples one encounters the difficulty that the form of (8.2) is not readily amenable to such computations. For example, in the Clayton copula with standard exponential marginals the conditional density is

$$\frac{Z}{(1-e^{-x})^{\theta+1}} \exp\left\{-Z\left(\frac{1}{(1-e^{-x})^\theta}-1\right)\right\} e^{-x}$$

The expressions for say a lognormal marginal F are even less inviting! \(\diamond\)

9. Concluding Remarks

The purpose of the present paper has not been to promote the use of CdMC in all the problems looked into, but rather to present some discussion of both the potential and the limitations of the method. Two aspects were argued from the outset to be potentially attractive, variance reduction and smoothing.

Table 4. δ in normal right tail ($n=2$).

CrMC	CdMC	AK	BdRelErr
$\frac{1}{4} = \frac{6}{24}$	$\frac{1}{3} = \frac{8}{24}$	$\frac{3}{8} = \frac{9}{24}$	$\frac{1}{2} = \frac{12}{24}$

Note: CrMC, crude Monte Carlo; CdMC, conditional Monte Carlo; BdRelErr, bounded relative error.

As mentioned in section 4, the traditional measure of efficiency in the rare-event simulation literature is the relative squared error $r_n^{(2)}(z)$, and BdRelErr (BdRelErr is usually considered as the most one can hope for). This and even more is obtained for Z_{AK} . Simple CdMC for the c.d.f. does not achieve BdRelErr, but nevertheless it was found to be worthwhile at least in the right tail of light-tailed sums and in the left tail when the increments are non-negative. For a quantitative illustration, consider estimating $\mathbb{P}(S_2 > x)$ in the Normal(0, 1) case. The variances of the different estimators were all found to be of the form $Cx^\gamma e^{-\beta x^2}$; note that for $\mathbb{P}(S_2 > x)$ itself (A.1) gives $C = 1 / \sqrt{\pi}$, $\gamma = -1$, $\beta = 1/4$. A good algorithm thus corresponds to a large value of β and the values found in the respective results are given in Table 4. Note that estimates similar to those of Asmussen & Glynn (2007: VI.1) show that BdRelErr is in fact obtained by the IS-ECM algorithm sketched at the end of section 3.

Also the smoothing performance of CdMC came out favourably in the examples considered. Averaging as in section 7 seemed too often to be worthwhile. We found that the ease with which CdMC produces plots of densities even in quite complicated models like the Clayton or Gumbel copulas in Figures 7 and 8 is a quite noteworthy property of the method.

In general, one could argue that when CdMC applies to either variance reduction or density estimation, it is at worst harmless and at best improves upon naive methods without involving more than a minor amount of extra computational effort. Some further comments:

1. When moving away from i.i.d. assumptions, we concentrated on dependence. Different marginals F_1, \dots, F_n can, however, also be treated by CdMC. For example, an obvious estimator for $\mathbb{P}(S_n > x)$ in this case is

$$\frac{1}{n} \sum_{i=1}^n \bar{F}_i(x - S_n + X_i)$$

This generalises in an obvious way to Z_{AK} . For discussion and extensions of these ideas, see, e.g., Chan & Kroese (2011) and Kortschak & Harshova (2013).

2. The example we have treated is sums but the CdMC method is not restricted to this case. In general, it is of course a necessary condition to have enough structure that conditional distributions are computable in a simple form.

Acknowledgement

The MSc dissertation work of Søren Høg provided a substantial impetus for the present study. The author is grateful to Patrick Laub for patiently looking into a number of Matlab queries. Both referees provided an extremely careful reading of the manuscript which helped remove various small inaccuracies as well as a more important one.

References

- Asmussen, S., Blanchet, J., Juneja, S. & Rojas-Nandayapa, L. (2011). Efficient simulation of tail probabilities of sums of correlated lognormals. *Annals of Operations Research*, **189**, 5–23.
- Asmussen, S. & Glynn, P.W. (2007). *Stochastic Simulation. Algorithms and Analysis*. Springer-Verlag, New York.
- Asmussen, S., Hashorva, E., Laub, P. & Taimre, T. (2017). Tail asymptotics for light-tailed Weibull-like sums. *Probability and Mathematical Statistics*, **37**, doi:10.19195/0208-4147.37.2.3.
- Asmussen, S., Jensen, J.L. & Rojas-Nandayapa, L. (2016). Exponential family techniques in the lognormal left tail. *Scandinavian Journal of Statistics*, **43**, 774–787.
- Asmussen, S. & Kortschak, D. (2012). On error rates in rare event simulation with heavy tails. Proceedings of the Winter Simulation Conference 2012. Berlin, December 10–12, 2012.
- Asmussen, S. & Kortschak, D. (2015). Error rates and improved algorithms for rare event simulation with heavy Weibull tails. *Methodology and Computing in Applied Probability*, **17**, 441–461.
- Asmussen, S. & Kroese, D.P. (2006). Improved algorithms for rare event simulation with heavy tails. *Advances in Applied Probability*, **38**, 545–558.
- Bahadur, R.R. (1966). A note on quantiles in large samples. *Annals of Mathematical Statistics*, **37**, 577–580.
- Balkema, A.A., Klüppelberg, C. & Resnick, S.I. (1993). Densities with Gaussian tails. *Proceedings of the London Mathematical Society*, **66**, 568–588.
- Blanchet, J. & Glynn, P.W. (2008). Efficient rare-event simulation for the maximum of heavy-tailed random walks. *Annals of Applied Probability*, **18**, 1351–1378.
- Botev, Z., Salomone, R. & MacKinlay, D. (2017). Fast and accurate computation of the distribution of sums of dependent log-normals. arXiv:1705.03196.
- Chan, J. & Kroese, D. (2011). Rare-event probability estimation with conditional Monte Carlo. *Annals of Operations Research*, **189**, 43–61.
- Cherubini, U., Luciano, E. & Vecchiato, W. (2004). *Copula Methods in Finance*. John Wiley & Sons, Chichester.
- Chu, F. & Nakayama, M.K. (2012). Confidence intervals for quantiles when applying variance-reduction techniques. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, **22**, Article 10.
- Cline, D.B.H. (1986). Convolution tails, product tails and domains of attraction. *Probability Theory and Related Fields*, **72**, 529–557.
- Dupuis, P., Leder, K. & Wang, H. (2007). Large deviations and importance sampling for a tandem network with slow-down. *QUESTA*, **57**, 71–83.
- Embrechts, P., Klüppelberg, C. & Mikosch, T. (1997). *Modelling Extremal Events for Finance and Insurance*. Springer-Verlag, Berlin.
- Fu, M.C., Hong, L.J. & Hu, J.-Q. (2009). Conditional Monte Carlo estimation of quantile sensitivities. *Management Science*, **55**, 2019–2027.
- Fu, M.C. & Hu, J.-Q. (1997). *Conditional Monte Carlo, Gradient Estimation and Optimization Applications*. Kluwer Academic Publisher, Springer, Boston.
- Ghamami, S. & Ross, S.M. (2012). Improving the Asmussen–Kroese-type simulation estimators. *Journal of Applied Probability*, **49**, 1188–1193.
- Ghosh, J.K. (1971). A new proof of the Bahadur representation of quantiles and an application. *Annals of Mathematical Statistics*, **42**, 1957–1961.
- Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer-Verlag, New York.
- Glasserman, P., Heidelberger, P. & Shahabuddin, P. (2000). Variance reduction techniques for estimating value-at-risk. *Management Science*, **46**, 1349–1364.

- Glynn, P.W. (1996). Importance sampling for Monte Carlo estimation of quantiles. Proceedings of the 2nd St. Petersburg Workshop on Simulation. Publishing House of St. Petersburg University, St. Petersburg, Russia, June 18-21, pp. 180–185.
- Gulisashvili, A. & Tankov, P. (2016). Tail behavior of sums and differences of log-normal random variables. *Bernoulli*, **22**, 444–493.
- Hartinger, J. & Kortschak, D. (2009). On the efficiency of the Asmussen-Kroese-estimators and its application to stop-loss transforms. *Blätter DGVFM*, **30**, 363–377.
- Hong, J.L., Hu, Z. & Liu, G. (2014). Monte Carlo methods for value-at-risk and conditional value-at-risk: a review. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, **24**, article 22.
- Jensen, J.L. (1995). *Saddlepoint Approximations*. Clarendon Press, Oxford.
- Juneja, S. (2007). Estimating tail probabilities of heavy tailed distributions with asymptotically zero relative error. *QUESTA*, **57**, 115–127.
- Kortschak, D. & Hashorva, E. (2013). Efficient simulation of tail probabilities for sums of log-elliptical risks. *Journal of Computational and Applied Mathematics*, **47**, 53–67.
- L'Ecuyer, P. & Lemieux, C. (2000). Variance reduction via lattice rules. *Management Science*, **46**, 1214–1235.
- L'Ecuyer, P. & Perron, G. (1994). On the convergence rates of IPA and FDC derivative estimators for finite-horizon stochastic simulations. *Operations Research*, **42**, 643–656.
- Mai, J.-F. & Scherer, M. (2012). *Simulating Copulas: Stochastic Models, Sampling Algorithms, and Applications*. Imperial College Press, London.
- Marshall, A.W. & Olkin, I. (1988). Families of multivariate distributions. *JASA*, **83**, 834–841.
- McNeil, A.J. (2006). Sampling nested Archimedean copulas. *Journal of Statistical Computation and Simulation*, **78**, 567–581.
- McNeil, A.J., Frey, R. & Embrechts, P. (2015). *Quantitative Risk Management*, 2nd edition. Princeton University Press, Princeton.
- Nakayama, M.K. (2014). Quantile estimation when applying conditional Monte Carlo. 2014 International Conference on Simulation and Modeling Methodologies, Technologies and Applications (SIMULTECH), IEEE, August 28-30, pp. 280–285.
- Serfling, R. (1980). *Approximation Theorems of Mathematical Statistics*. Wiley, Chichester.
- Sun, L. & Hong, J.L. (2010). Asymptotic representations of importance-sampling estimators for value-at-risk and conditional value-at-risk. *Operations Research Letters*, **38**, 246–251.
- Wu, F., Valdez, E. & Sherris, M. (2007). Simulating from exchangeable Archimedean copulas. *Communications in Statistics – Simulation and Computation*, **36**, 1019–1034.
- Whelan, N. (2004). Sampling from Archimedean copulas. *Quantitative Finance*, **4**, 339–352.

Appendix

A Technical proofs

Proof of Proposition 3.3: Using

$$\int_0^x y^a (x-y)^b dy = B(a+1, b+1) x^{a+b+1}$$

where $B(\cdot, \cdot)$ is the Beta function, we get for suitable constants c_1, c_2, \dots depending on n that

$$f^{*2}(x) = \int_0^x f(x-y)f(y) dy \sim c_1^2 \int_0^x y^p (x-y)^p dy = c_2 x^{2p+1}$$

as $x \rightarrow 0$ and, by induction:

$$f^{*n}(x) \sim c_3 x^{np+n-1}, \quad F^{*n}(x) = \int_0^x f^{*n}(y) dy \sim c_4 x^{np+n}$$

Hence

$$\begin{aligned} \mathbb{V}ar[F(x-S_{n-1})] &\sim \int_0^x c_5 y^{np+n-p-2} F(x-y)^2 dy - c_4^2 x^{2np+2n} \\ &= \int_0^x c_6 y^{np+n-p-2} (x-y)^{2p+2} dy - c_4^2 x^{2np+2n} \\ &\sim c_7 x^{np+n+p+1}, \\ r_n(x) &\sim \frac{c_7 x^{np+n+p+1}}{c_4 x^{np+n} (1-c_4 x^{np+n})} \sim c_8 x^{p+1} \quad \square \end{aligned}$$

Proof of Proposition 3.4: Let Z be a r.v. with tail $\bar{F}(x)^2$. By general subexponential theory, $\mathbb{P}(S_k > x) \sim k\bar{F}(x)$ for any fixed k and $\mathbb{P}(S_{n-1} + Z > x) \sim \bar{F}_{n-1}(x)$ since Z therefore has lighter tail than S_{n-1} . Hence

$$\begin{aligned} \mathbb{E}F(x-S_{n-1})^2 &= \bar{F}_{n-1}(x) + \mathbb{P}(S_{n-1} + Z > x, S_{n-1} \leq x) = \bar{F}_{n-1}(x) + o(\bar{F}(x)) \\ \mathbb{V}ar[F(x-S_{n-1})] &= \bar{F}_{n-1}(x) + o(\bar{F}(x)) - O(\bar{F}(x)^2) \sim (n-1)\bar{F}(x) \quad \square \end{aligned}$$

In the last two proofs, we shall need the Mill's ratio estimate of the normal tail, stating that if V is standard normal, then

$$\mathbb{P}(\sigma V > x) \begin{cases} \leq \frac{\sigma}{x\sqrt{2\pi}} e^{-x^2/2\sigma^2} & \text{for } x > 0 \\ \sim \frac{\sigma}{x\sqrt{2\pi}} e^{-x^2/2\sigma^2} & \text{as } x \rightarrow \infty \end{cases} \quad (\text{A.1})$$

A slightly more general version, proved in the same way via L'Hospital, is

$$\int_{bx}^{\infty} \frac{1}{(y+cx)^k} e^{-ay^2/2} dy \sim \frac{1}{ab(b+c)^k x^{k+1}} e^{-ab^2x^2/2} \quad (\text{A.2})$$

Proof of Proposition 3.5: We have

$$\mathbb{E}\bar{F}(x-S_{n-1})^2 = \int_{-\infty}^{\infty} \bar{\Phi}(x-y)^2 \frac{1}{(2\pi(n-1))^{1/2}} e^{-y^2/2(n-1)} dy \quad (\text{A.3})$$

Let $\bar{H}(x) = \mathbb{P}(V_{n-1} + V_{1/2} > x)$ where $V_{n-1}, V_{1/2}$ are independent mean zero normals with variances $n-1$ and $1/2$, and note that

$$\mathbb{P}(V_{n-1} > x-A) = o(\bar{H}(x)) \quad (\text{A.4})$$

according to (A.1). The $y > x - A$ part of the integral in (A.3) is bounded by (A.4). Noting that

$$\frac{\bar{\Phi}(x)^2}{\mathbb{P}(V_{1/2} > x)} \sim \frac{e^{-x^2}/x^2 2\pi}{1/2 \cdot e^{-x^2}/x\sqrt{2\pi}} = \frac{1}{x} \sqrt{2/\pi}$$

the $y \leq x - A$ part asymptotically becomes

$$\begin{aligned} & \int_{-\infty}^{x-A} \frac{\sqrt{2/\pi}}{x-y} \mathbb{P}(V_{1/2} > x-y) \frac{1}{(2\pi(n-1))^{1/2}} e^{-y^2/2(n-1)} dy \\ & \sim \frac{1}{x} \sqrt{2/\pi} \mathbb{P}(V_{n-1} + V_{1/2} > x, V_{n-1} \leq x-A) \\ & = \frac{1}{x} \sqrt{2/\pi} \bar{H}(x) - o(\bar{H}(x)) \end{aligned}$$

Thus

$$\begin{aligned} r_n(x) & \sim \frac{\bar{H}(x) \sqrt{2/\pi} / x}{\mathbb{P}(S_n > x)} \sim \frac{\sqrt{n-1/2} e^{-x^2/(2n-1)} / \sqrt{2\pi} \sqrt{2/\pi} / x}{\sqrt{ne^{-x^2/2n}} / x \sqrt{2\pi}} \\ & \frac{1}{x} \sqrt{\frac{2n-1}{n\pi}} e^{-x^2/[2n(2n-1)]} \end{aligned}$$

where we used (A.1) two times with $\sigma^2 = n - 1/2$, respectively, $\sigma^2 = n$. □

Proof of (4.2): In the same way as in Example 4.2, $\max(M_{n-1}, x - S_{n-1}) = \max(X_1, x - X_1)$ splits up into $X_1 \leq x/2$ and $X_1 > x/2$ parts. Using (A.3) to estimate $\bar{\Phi}(y)$, the $X_1 > x/2$ part of $\mathbb{E}Z_{\Lambda K}^2(x)$ becomes

$$\frac{4}{\sqrt{2\pi}} \int_{x/2}^{\infty} \bar{\Phi}(y)^2 e^{-y^2/2} dy \sim \frac{4}{(2\pi)^{3/2}} \int_{x/2}^{\infty} \frac{1}{y^2} e^{-3y^2/2} dy \sim \frac{32}{3x^3(2\pi)^{3/2}} e^{-3x^2/8} dy$$

The $X \leq x/2$ part is

$$\begin{aligned} & \frac{4}{\sqrt{2\pi}} \int_{-\infty}^{x/2} \bar{\Phi}(x-y)^2 e^{-y^2/2} dy = \frac{4}{\sqrt{2\pi}} \int_{x/2}^{\infty} \bar{\Phi}(y)^2 e^{-(x-y)^2/2} dy \\ & \sim \frac{4}{(2\pi)^{3/2}} \int_{x/2}^{\infty} \frac{1}{y^2} e^{-y^2-(x-y)^2/2} dy = \frac{4}{(2\pi)^{3/2}} e^{-x^2/3} \int_{x/2}^{\infty} \frac{1}{y^2} e^{-3(y-x/3)^2/2} dy \\ & = \frac{4}{(2\pi)^{3/2}} e^{-x^2/3} \int_{x/6}^{\infty} \frac{1}{(y+x/3)^2} e^{-3y^2/2} dy \\ & = \frac{4}{(2\pi)^{3/2}} e^{-x^2/3} \frac{1}{3 \cdot 1/6 \cdot (1/2)^2 x^3} e^{-3(x/6)^2} = \frac{32}{3x^3(2\pi)^{3/2}} e^{-3x^2/8} \end{aligned}$$

Adding up, the results follows. □