

# Monte Carlo and Quasi-Monte Carlo Density Estimation via Conditioning

Pierre L'Ecuyer

Département d'Informatique et de Recherche Opérationnelle, Pavillon Aisenstadt, Université de Montréal, C.P. 6128, Succ. Centre-Ville, Montréal, Québec, Canada H3C 3J7, lecuyer@iro.umontreal.ca

Florian Puchhammer

Basque Center of Applied Mathematics, Alameda de Mazarredo 14, 48009 Bilbao, Basque Country, Spain; and Département d'Informatique et de Recherche Opérationnelle, Université de Montréal, fpuchhammer@bcamath.org

Amal Ben Abdellah

Département d'Informatique et de Recherche Opérationnelle, Pavillon Aisenstadt, Université de Montréal, C.P. 6128, Succ. Centre-Ville, Montréal, Québec, Canada H3C 3J7, amal.ben.abdellah@umontreal.ca

Estimating the unknown density from which a given independent sample originates is more difficult than estimating the mean, in the sense that for the best popular density estimators, the mean integrated square error converges more slowly than at the canonical rate of  $\mathcal{O}(1/n)$ . When the sample is generated from a simulation model and we have control over how this is done, we can do better. We examine an approach in which conditional Monte Carlo permits one to obtain a smooth estimator of the cumulative distribution function, whose sample derivative is an unbiased estimator of the density at any point, and therefore converges at a faster rate than the usual density estimators, under mild conditions. By combining this with randomized quasi-Monte Carlo to generate the sample, we can achieve an even faster rate.

*Key words:* density estimation; conditional Monte Carlo; quasi-Monte Carlo

---

## 1. Introduction

Simulation is commonly used to generate independent realizations of a random variable  $X$  that may represent a payoff, a cost, or a performance of some kind, and then to estimate from this sample the unknown expectation of  $X$  together with a confidence interval (Asmussen and Glynn 2007, Law 2014). However, one can extract from the sample much more information than just a confidence interval on the mean. When  $n$  is reasonably large, one may in fact estimate the *entire distribution of  $X$* . Simple ways of showing this information is to plot the empirical *cumulative distribution function* (cdf) or (more commonly) a histogram of the observations. The histogram provides a better visual insight on the distribution. When  $X$  is a continuous random variable, it can be interpreted as a rough sketch (or estimate) of the density of  $X$ .

There are more refined and efficient density estimation methods, the leading one being the *kernel density estimator* (KDE). Given  $n$  independent realizations of  $X$ , the mean integrated square error (MISE) of the KDE converges as  $\mathcal{O}(n^{-4/5})$  in the best case, while a histogram can only achieve

$\mathcal{O}(n^{-2/3})$ . Both rates are slower than the canonical  $\mathcal{O}(n^{-1})$  rate for the variance of the sample average as an unbiased estimator of the mean. The slower rate for the KDE, histogram, and other similar density estimators stems from the presence of bias. For a histogram, taking wider rectangles reduces the variance but increases the bias by flattening out the short-range density variations, and a compromise must be made to minimize the MISE. The same happens with the KDE, with the rectangle width replaced by the bandwidth of the kernel.

Density estimation has several applications (Van der Vaart 2000, Scott 2015). A common setting is when a set of observations is given and one wishes to estimate the density from which they come. A different one is when we have a simulation (or generative) model from which observations can be generated and we want to estimate the density of the model output. This is the focus of our paper. Density estimates are also important in various other settings. For instance when computing a confidence interval for a quantile using the central-limit theorem (CLT), one needs a density estimator at the quantile to estimate the variance (Serfling 1980, Asmussen and Glynn 2007, Nakayama 2014a,b, Peng et al. 2017). Another application is for maximum likelihood estimation when the likelihood does not have a closed-form expression, so to maximize it with respect to some parameter  $\theta$ , the likelihood function (which in the continuous case is a density at any value of  $\theta$ ) must be estimated (Van der Vaart 2000, Peng et al. 2020). A related application is the estimation of the posterior density of  $\theta$  given some data, in a Bayesian model (Efron and Hastie 2016).

The most common density estimators are biased and their MISE converges slower than  $\mathcal{O}(n^{-1})$ . Can we construct density estimators that avoid bias and whose MISE converges as  $\mathcal{O}(n^{-1})$  or faster? When the  $n$  observations of  $X$  are given and nothing else is known, as traditionally assumed in classical non-parametric statistics, no known practical estimator does better than the KDE (Scott 2015). But if the observations of  $X$  are generated by simulation and we have control of how this is done, there are opportunities to improve the density estimators. This is the subject of the present paper. We study two main ideas and their combination.

The first general idea is to build a smooth estimator of the cdf via conditional Monte Carlo (CMC), and take the sample derivative of this estimator to estimate the density. We call it a *conditional density estimator* (CDE). Under appropriate conditions, the CDE is unbiased and its variance is bounded uniformly by a constant divided by  $n$ , so its MISE is  $\mathcal{O}(n^{-1})$ . This idea of using CMC was mentioned by Asmussen and Glynn (2007), page 146, Example 4.3, and further studied in Asmussen (2018), both for the special case where the goal is to estimate the density of a sum of continuous random variables having a known distribution from which we can sample exactly. Asmussen (2018) simply “hides” the last term of the sum, meaning that this term is not generated or its actual value is erased, and he takes the conditional distribution of the sum given the other terms, to estimate the cdf, the density, the value at risk (VaR), and the conditional value

at risk (CVaR) of the sum. Smoothing by CMC has also been studied for estimating the derivative of an expectation (L'Ecuyer and Perron 1994, Fu and Hu 1997) and the derivative of a quantile (Fu et al. 2009) with respect to a model parameter.

In this paper, we study the CDE method to estimate the cdf and density in a much more general setting than Asmussen (2018). We provide several examples in which  $X$  is *not* defined as a sum of random variables and where we hide (or erase) more than just one random variable to do the conditioning. We give conditions under which we can prove that the density estimator is unbiased. A key condition is that the conditional cdf must be a continuous function of the point  $x$  at which we estimate the density. Sometimes, this can be achieved only by hiding several variables and not only one. The variance of the density estimator may depend strongly on which variables we hide, i.e., on what we are conditioning. We illustrate this with several examples.

Once we have a smooth density estimator, the second strategy to further improve the convergence rate is to replace the independent uniform random numbers that drive the simulation by *randomized quasi-Monte Carlo* (RQMC) points. We show in this paper that by combining these two strategies, under some conditions, we can obtain a density estimator whose MISE converges at a faster rate than  $\mathcal{O}(n^{-1})$ , for instance  $\mathcal{O}(n^{-2+\epsilon})$  for any  $\epsilon > 0$  in some situations. We also observe this fast rate empirically on numerical examples.

In related work, Ben Abdellah et al. (2019) studied a different approach which combines RQMC with an ordinary KDE. They were able to prove a faster rate than  $\mathcal{O}(n^{-4/5})$  for the MISE when the RQMC points have a small number of dimensions, and they observed this faster rate empirically on examples. They also observed that the MISE reduction provided by RQMC degrades rapidly when the bandwidth is reduced (to reduce the bias) and that this degradation effect gets stronger when the dimension increases, to the point that the gain is usually quite limited in more than 5 dimensions or so. The CDE+RQMC approach proposed in the present paper avoids this problem (there is no bias and no bandwidth) and is generally more effective than the KDE+RQMC combination. We provide some numerical comparisons in our examples.

Other Monte Carlo density estimators were proposed recently, based on the idea of estimating the derivative of the cdf using a likelihood ratio (LR) method. This general approach permits one to estimate the derivative of the expectation of a random variable with respect to some parameter of the underlying distribution when this random variable is discontinuous (Glynn 1987, L'Ecuyer 1990). Laub et al. (2019) proposed an approach that combines a clever change of variable with the LR method to estimate the density of a sum of random variables as in Asmussen (2018), but in a setting where the random variables are dependent. Peng et al. (2018) proposed a generalized version of the LR gradient estimator (GLR) to estimate the derivative of an expectation with respect to a more general model parameter. Lei et al. (2018) then sketched out how GLR could

be used to estimate a density. Formulas for these GLR density estimators are given in Theorem 1 of Peng et al. (2020). A referee pointed out these last two papers to us. We compare our method with these GLR-based estimators in our numerical illustrations.

The remainder is organized as follows. In Section 2, we define our general setting, recall some key facts about density estimators, introduce the general CMC method to build the CDEs considered in this paper, prove some of their properties, and give small examples to provide insight on the key ideas. In Section 3, we explain how to combine the CDE with RQMC and discuss the convergence properties for this combination. Section 4 reports experimental results with various examples. Some of the examples feature creative ways of conditioning to boost the effectiveness of the method. A conclusion is given in Section 5.

## 2. Model and conditional density estimator

### 2.1. Density estimation

We have a real-valued random variable  $X$  that can be simulated from its exact distribution, but we do not know the cdf  $F$  and density  $f$  of  $X$ . Typically,  $X$  will be an easily computable function of several other random variables with known densities. Our goal is to estimate  $f$  over a finite interval  $[a, b]$ . Let  $\hat{f}_n$  denote an estimator of  $f$  based on a sample of size  $n$ . We measure the quality of  $\hat{f}_n$  by the *mean integrated square error* (MISE), defined as

$$\text{MISE} = \text{MISE}(\hat{f}_n) = \int_a^b \mathbb{E}[\hat{f}_n(x) - f(x)]^2 dx. \quad (1)$$

The MISE is the sum of the *integrated variance* (IV) and the *integrated square bias* (ISB):

$$\text{MISE} = \text{IV} + \text{ISB} = \int_a^b \mathbb{E}(\hat{f}_n(x) - \mathbb{E}[\hat{f}_n(x)])^2 dx + \int_a^b (\mathbb{E}[\hat{f}_n(x)] - f(x))^2 dx.$$

A standard way of constructing  $\hat{f}_n$  when  $X_1, \dots, X_n$  are  $n$  independent realizations of  $X$  is via a KDE, defined as follows (Parzen 1962, Scott 2015):

$$\hat{f}_n(x) = \frac{1}{nh} \sum_{i=1}^n k\left(\frac{x - X_i}{h}\right),$$

where the *kernel*  $k$  is a probability density over  $\mathbb{R}$ , usually symmetric about 0 and non-increasing over  $[0, \infty)$ , and the constant  $h > 0$  is the *bandwidth*, whose role is to stretch [or compress] the kernel horizontally to smooth out [or unsmooth] the estimator  $\hat{f}_n$ . The KDE was developed initially for the setting in which  $X_1, \dots, X_n$  are given a priori, and it is still the most popular one for this situation. It can be used in exactly the same way when  $X_1, \dots, X_n$  are independent observations produced by simulation from a generative model, but in that case there is an opportunity to do better, as we now explain.

## 2.2. Conditioning and the stochastic derivative as an unbiased density estimator

Since the density of  $X$  is the derivative of its cdf,  $f(x) = F'(x)$ , a natural idea would be to take the derivative of an estimator of the cdf as a density estimator. The simplest candidate for a cdf estimator is the *empirical cdf*

$$\hat{F}_n(x) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}[X_i \leq x],$$

but  $d\hat{F}_n(x)/dx = 0$  almost everywhere, so this one *cannot* be a useful density estimator. Here,  $\hat{F}_n(x)$  is an unbiased estimator of  $F(x)$  at each  $x$ , but its derivative is a biased estimator of  $F'(x)$ . That is, because of the discontinuity of  $\hat{F}_n$ , we cannot exchange the derivative and expectation:

$$0 = \mathbb{E} \left[ \frac{d\hat{F}_n(x)}{dx} \right] \neq \frac{d\mathbb{E}[\hat{F}_n(x)]}{dx} = F'(x).$$

A continuous estimator of  $F$  can be constructed by CMC, as follows. Replace the indicator  $\mathbb{I}[X \leq x]$  by its *conditional cdf* given filtered (reduced) information  $\mathcal{G}$ :  $F(x | \mathcal{G}) \stackrel{\text{def}}{=} \mathbb{P}[X \leq x | \mathcal{G}]$ , where  $\mathcal{G}$  is a sigma-field that contains not enough information to reveal  $X$  but enough to compute  $F(x | \mathcal{G})$ . Here, knowing the realization of  $\mathcal{G}$  means knowing the realizations of all  $\mathcal{G}$ -measurable random variables. Under the following assumption, we prove that  $F'(x | \mathcal{G}) = dF(x | \mathcal{G})/dx$  is an unbiased estimator of  $f(x)$  whose variance is bounded uniformly in  $x$ . Note that since  $F(\cdot | \mathcal{G})$  cannot decrease,  $F'(\cdot | \mathcal{G})$  is never negative.

**ASSUMPTION 1.** *For all realizations of  $\mathcal{G}$ ,  $F(x | \mathcal{G})$  is a continuous function of  $x$  over the interval  $[a, b]$ , and is differentiable except perhaps at a countable set of points  $D(\mathcal{G}) \subset [a, b]$ . There is also a random variable  $\Gamma$  defined over the same probability space as  $F(x | \mathcal{G})$ , such that  $\mathbb{E}[\Gamma^2] \leq K_\gamma$  for some constant  $K_\gamma < \infty$ , and for which  $\sup_{x \in [a, b] \setminus D(\mathcal{G})} F'(x | \mathcal{G}) \leq \Gamma$ .*

**THEOREM 1.** *Under Assumption 1, for all  $x \in [a, b] \setminus D(\mathcal{G})$ , we have  $\mathbb{E}[F'(x | \mathcal{G})] = f(x)$  and  $\text{Var}[F'(x | \mathcal{G})] \leq K_\gamma$ .*

**PROOF.** We adapt the proof of Theorem 1 of L'Ecuyer (1990). By the mean value inequality theorem of Dieudonné (1969), Theorem 8.5.3, which is a form of mean value theorem for non-differentiable functions, for every  $x \in [a, b]$  and  $\delta > 0$ , with probability 1, we have

$$0 \leq \frac{\Delta(x, \delta, \mathcal{G})}{\delta} \stackrel{\text{def}}{=} \frac{F(x + \delta | \mathcal{G}) - F(x | \mathcal{G})}{\delta} \leq \sup_{y \in [x, x + \delta] \setminus D(\mathcal{G})} F'(y | \mathcal{G}) \leq \Gamma.$$

Then, by the dominated convergence theorem,

$$\mathbb{E} \left[ \lim_{\delta \rightarrow 0} \frac{\Delta(x, \delta, \mathcal{G})}{\delta} \right] = \lim_{\delta \rightarrow 0} \mathbb{E} \left[ \frac{\Delta(x, \delta, \mathcal{G})}{\delta} \right],$$

which shows the unbiasedness. Moreover,  $\text{Var}[F'(x | \mathcal{G})] \leq \mathbb{E}[\Gamma^2] \leq K_\gamma$ . □

Suppose now that  $\mathcal{G}^{(1)}, \dots, \mathcal{G}^{(n)}$  are  $n$  independent realizations of  $\mathcal{G}$ , so  $F(x | \mathcal{G}^{(1)}), \dots, F(x | \mathcal{G}^{(n)})$  are independent realizations of  $F(x | \mathcal{G})$ , and consider the CDE

$$\hat{f}_{\text{cde},n}(x) = \frac{1}{n} \sum_{i=1}^n F'(x | \mathcal{G}^{(i)}). \quad (2)$$

It follows from Theorem 1 that  $\text{ISB}(\hat{f}_{\text{cde},n}) = 0$  and  $\text{MISE}(\hat{f}_{\text{cde},n}) = \text{IV}(\hat{f}_{\text{cde},n}) \leq (b-a)K_\gamma/n$ . An unbiased estimator of this IV is given by

$$\widehat{\text{IV}} = \widehat{\text{IV}}(\hat{f}_{\text{cde},n}) = \frac{1}{n} \int_a^b \sum_{i=1}^n \left[ F'(x | \mathcal{G}^{(i)}) - \hat{f}_{\text{cde},n}(x) \right]^2 dx. \quad (3)$$

In practice, this integral can be approximated by evaluating the integrand at a finite number of points over  $[a, b]$  and taking the average, multiplied by  $(b-a)$ .

It is well known that in general, when estimating  $\mathbb{E}[X]$ , a CMC estimator never has a larger variance than  $X$  itself, and the more information we hide, the smaller the variance. That is, if  $\mathcal{G} \subset \tilde{\mathcal{G}}$  are two sigma-fields such that  $\mathcal{G}$  contains only a subset of the information of  $\tilde{\mathcal{G}}$ , then

$$\text{Var}[\mathbb{E}[X | \mathcal{G}]] \leq \text{Var}[\mathbb{E}[X | \tilde{\mathcal{G}}]] \leq \text{Var}[X]. \quad (4)$$

Noting that  $F(x | \mathcal{G}) = \mathbb{E}[\mathbb{I}[X \leq x] | \mathcal{G}]$ , we also have

$$\text{Var}[F(x | \mathcal{G})] \leq \text{Var}[F(x | \tilde{\mathcal{G}})] \leq \text{Var}[\mathbb{I}[X \leq x]] = F(x)(1 - F(x)).$$

Thus, (4) applies as well to the (conditional) cdf estimator.

However, applying it to the CDE, which is the derivative of the conditional cdf, is less straightforward. It is obviously not true that  $\text{Var}[F'(x | \mathcal{G})] \leq \text{Var}[d\mathbb{I}[X \leq x]/dx]$  because the latter is zero almost everywhere. Nevertheless, we can prove the following.

**THEOREM 2.** *If  $\mathcal{G} \subset \tilde{\mathcal{G}}$  both satisfy Assumption 1, then for all  $x \in [a, b]$ , we have  $\text{Var}[F'(x | \mathcal{G})] \leq \text{Var}[F'(x | \tilde{\mathcal{G}})]$ .*

**PROOF.** The result does not follow directly from (4) because  $F'$  is not an expectation; this is why our proof does a little detour. For an arbitrary  $x \in [a, b]$  and a small  $\delta > 0$ , define the random variable  $I = I(x, \delta) = \mathbb{I}[x < X \leq x + \delta]$ . We have  $\mathbb{E}[I | \mathcal{G}] = F(x + \delta | \mathcal{G}) - F(x | \mathcal{G})$ , as in the proof of Theorem 1, and similarly for  $\tilde{\mathcal{G}}$ . Using (4) with  $I$  in place of  $X$  gives

$$\text{Var}[\mathbb{E}[I | \mathcal{G}]] \leq \text{Var}[\mathbb{E}[I | \tilde{\mathcal{G}}]]. \quad (5)$$

We have

$$F'(x | \mathcal{G}) = \lim_{\delta \rightarrow 0} \frac{F(x + \delta | \mathcal{G}) - F(x | \mathcal{G})}{\delta} = \lim_{\delta \rightarrow 0} \mathbb{E}[I(x, \delta)/\delta | \mathcal{G}]$$

and similarly for  $\tilde{\mathcal{G}}$ . Combining this with (5), we obtain

$$\begin{aligned} \text{Var}[F'(x | \mathcal{G})] &= \text{Var}[\lim_{\delta \rightarrow 0} \mathbb{E}[I(x, \delta)/\delta | \mathcal{G}]] = \lim_{\delta \rightarrow 0} \text{Var}[\mathbb{E}[I(x, \delta)/\delta | \mathcal{G}]] \\ &\leq \lim_{\delta \rightarrow 0} \text{Var}[\mathbb{E}[I(x, \delta)/\delta | \tilde{\mathcal{G}}]] = \text{Var}[\lim_{\delta \rightarrow 0} \mathbb{E}[I(x, \delta)/\delta | \tilde{\mathcal{G}}]] = \text{Var}[F'(x | \tilde{\mathcal{G}})], \end{aligned}$$

in which the exchange of “Var” with the limit (at two places) can be justified by a similar argument as in Theorem 1. More specifically, we need to apply the dominated convergence theorem to  $\mathbb{E}[I(x, \delta)/\delta | \mathcal{G}]$ , which is just the same as in Theorem 1, and also to its square, which is also valid because the square is bounded uniformly by  $\Gamma^2$ . This completes the proof.  $\square$

When none of  $\mathcal{G}$  or  $\tilde{\mathcal{G}}$  is a subset of the other, Theorem 2 does not help, but the variances of the corresponding conditional density estimators may differ significantly. Generally speaking, we want to choose  $\mathcal{G}$  so that the density of  $X$  conditional on  $\mathcal{G}$  is spread out rather than being concentrated in a narrow peak. We give examples of this in the following.

### 2.3. Small examples to provide insight

This subsection provides simple examples to illustrate the key ideas. For these examples, we take  $X = h(Y_1, \dots, Y_d)$  where  $Y_1, \dots, Y_d$  are independent continuous random variables, each  $Y_j$  has cdf  $F_j$  and density  $f_j$ , and we condition on  $\mathcal{G} = \mathcal{G}_{-k}$  defined as the information that remains after erasing the value taken by  $Y_k$ . We can write  $\mathcal{G}_{-k} = (Y_1, \dots, Y_{k-1}, Y_{k+1}, \dots, Y_d)$ . The CDE  $F'(x | \mathcal{G}_{-k})$  will be related to the density  $f_k$  and will depend on the form of  $h$ . In Section 4, we consider more elaborate forms of conditioning.

EXAMPLE 1. A very simple situation is when  $X = h(Y_1, \dots, Y_d) = Y_1 + \dots + Y_d$ , a sum of  $d$  independent continuous random variables. By hiding  $Y_k$  for an arbitrary  $k$ , we get

$$F(x | \mathcal{G}_{-k}) = \mathbb{P}[X \leq x | S_{-k}] = \mathbb{P}[Y_k \leq x - S_{-k}] = F_k(x - S_{-k}),$$

where  $S_{-k} \stackrel{\text{def}}{=} \sum_{j=1, j \neq k}^d Y_j$ , and the density estimator becomes  $F'(x | \mathcal{G}_{-k}) = f_k(x - S_{-k})$ . This form also works if the  $Y_j$  are not independent, provided that we are able to compute the density of  $Y_k$  conditional on  $\mathcal{G}_{-k}$ . It suffices to replace  $f_k$  by this conditional density. The model studied by Asmussen (2018) was a special case of this example, with independent variables and  $k = d$ .

EXAMPLE 2. To show how we can verify Assumption 1, let  $X$  be the sum of two independent normal random variables,  $X = Y_1 + Y_2$ , where  $Y_1 \sim \mathcal{N}(0, \sigma_1^2)$ ,  $Y_2 \sim \mathcal{N}(0, \sigma_2^2)$ , and  $\sigma_1^2 + \sigma_2^2 = 1$ , so  $X \sim \mathcal{N}(0, 1)$ . With  $\mathcal{G} = \mathcal{G}_{-2}$ , we have  $F(x | \mathcal{G}_{-2}) = \mathbb{P}[Y_2 \leq x - Y_1] = \Phi((x - Y_1)/\sigma_2)$  and the CDE is

$F'(x | \mathcal{G}_{-2}) = \phi((x - Y_1)/\sigma_2)/\sigma_2$ . Assumption 1 is easily verified with  $\Gamma = \phi(0)/\sigma_2$  and  $K_\gamma = \Gamma^2$ , so this estimator is unbiased for  $f(x) = \phi(x)$ . Its variance is

$$\begin{aligned} \text{Var}[\phi((x - Y_1)/\sigma_2)/\sigma_2] &= \mathbb{E}[\exp[-(x - Y_1)^2/\sigma_2^2]/(2\pi\sigma_2^2)] - \phi^2(x) \\ &= \frac{1}{\sigma_2^2\sqrt{2\pi}}\mathbb{E}[\phi(\sqrt{2}(x - Y_1)/\sigma_2)] - \phi^2(x) \\ &= \frac{1}{\sigma_2\sqrt{2\pi(1 + \sigma_1^2)}}\phi\left(\frac{\sqrt{2}x}{\sqrt{1 + \sigma_1^2}}\right) - \phi^2(x). \end{aligned} \quad (6)$$

EXAMPLE 3. Let  $X = (Y_1 + Y_2^2)/Y_3$  where  $Y_3$  is always positive. If  $k = 3$ , then  $F(x | \mathcal{G}_{-3}) = \mathbb{P}(X \leq x | Y_1, Y_2) = \mathbb{P}(Y_3 \geq (Y_1 + Y_2^2)/x) = 1 - F_3((Y_1 + Y_2^2)/x)$  and the density estimator at  $x$  is

$$F'(x | \mathcal{G}_{-3}) = f_3((Y_1 + Y_2^2)/x)(Y_1 + Y_2^2)/x^2.$$

If  $k = 2$  instead, then  $F(x | \mathcal{G}_{-2}) = \mathbb{P}(X \leq x | Y_1, Y_3) = \mathbb{P}(|Y_2| \leq (Y_3x - Y_1)^{1/2}) = F_2(Z(x)) - F_2(-Z(x))$  where  $Z(x) = (Y_3x - Y_1)^{1/2}$ , and the density estimator at  $x$  is

$$F'(x | \mathcal{G}_{-2}) = (f_2(Z(x)) + f_2(-Z(x)))dZ(x)/dx = (f_2(Z(x)) - f_2(-Z(x)))Y_3/(2Z(x)).$$

EXAMPLE 4. The following tiny example provides further insight into the choice of  $\mathcal{G}$ . Suppose  $X$  is the sum of two independent uniform random variables:  $X = Y_1 + Y_2$  where  $Y_1 \sim \mathcal{U}(0, 1)$  and  $Y_2 \sim \mathcal{U}(0, \epsilon)$  where  $0 < \epsilon < 1$ . The exact density of  $X$  here is  $f(x) = x/\epsilon$  for  $0 \leq x \leq \epsilon$ ,  $f(x) = 1$  for  $\epsilon \leq x \leq 1$ , and  $f(x) = (1 + \epsilon - x)/\epsilon$  for  $1 \leq x \leq 1 + \epsilon$ . Figure 1 illustrates this density.

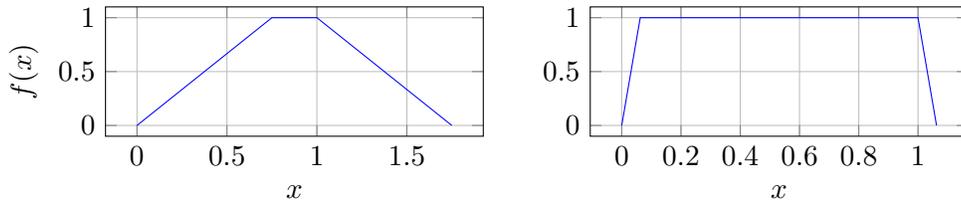


Figure 1 Exact density of  $X$  for the model in Example 4 with  $\epsilon = 3/4$  (left) and  $\epsilon = 1/16$  (right).

Using the notation of Example 1, with  $\mathcal{G} = \mathcal{G}_{-1}$ , we have  $F(x | \mathcal{G}_{-1}) = \mathbb{P}[X \leq x | Y_2] = \mathbb{P}[Y_1 \leq x - Y_2 | Y_2] = x - Y_2$  and the density estimator is  $F'(x | \mathcal{G}_{-1}) = 1$  for  $Y_2 \leq x \leq 1 + Y_2$ , and 0 elsewhere. If  $\mathcal{G} = \mathcal{G}_{-2}$  instead, then  $F(x | \mathcal{G}_{-2}) = \mathbb{P}[Y_2 \leq x - Y_1 | Y_1] = (x - Y_1)/\epsilon$  and the density estimator is  $F'(x | \mathcal{G}_{-2}) = 1/\epsilon$  for  $Y_1 \leq x \leq \epsilon + Y_1$ , and 0 elsewhere. In both cases, the density estimator with one sample is a uniform density, but the second one is over a narrow interval if  $\epsilon$  is small. Assumption 1 holds in both cases, since  $F(\cdot | \mathcal{G}_{-j})$  is continuous and differentiable for  $j = 1$  and 2. For  $\mathcal{G}_{-1}$  we can take  $\Gamma = K_\gamma = 1$ , and for  $\mathcal{G}_{-2}$  we can take  $\Gamma = \epsilon^{-1}$  and  $K_\gamma = \epsilon^{-2}$ . When  $\epsilon$  is small,  $\mathcal{G} = \mathcal{G}_{-2}$  gives a density estimator  $\hat{f}_{\text{cde},n}$  which is a sum of high narrow peaks and has much

larger variance. The constants  $K_\gamma$  give upper bounds on the variance at any given  $x$ . For this simple example, we can also derive exact formulas for the IV of the CDE under MC. For  $\mathcal{G} = \mathcal{G}_{-1}$ ,  $F'(x | \mathcal{G}_{-1}) = \mathbb{I}[Y_2 \leq x \leq 1 + Y_2]$  is a Bernoulli random variable with mean  $\mathbb{P}[x - 1 \leq Y_2 \leq x] = f(x)$ , so its variance is  $f(x)(1 - f(x))$ . Integrating this over  $[0, 1 + \epsilon]$  gives  $\text{IV} = \epsilon/3$  for one sample. For a sample of size  $n$ , this gives  $\text{IV} = \epsilon/(3n)$ . For  $\mathcal{G} = \mathcal{G}_{-2}$ ,  $F'(x | \mathcal{G}_{-2}) = \mathbb{I}[Y_1 \leq x \leq \epsilon + Y_1]/\epsilon$  has also mean  $f(x)$ , but its variance is  $\epsilon^{-1}f(x)(1 - \epsilon f(x))$ . Integrating over  $[0, 1 + \epsilon]$  gives  $\text{IV} = 1/\epsilon - 1 + \epsilon/3$  for one sample, which is much larger than  $\epsilon/3$  when  $\epsilon$  is small.

EXAMPLE 5. If  $X$  is the min or max of two or more continuous random variables, then in general  $F(\cdot | \mathcal{G}_{-k})$  is not continuous, so Assumption 1 does not hold. To illustrate this, let  $X = \max(Y_1, Y_2)$  where  $Y_1$  and  $Y_2$  are independent. With  $\mathcal{G} = \mathcal{G}_{-2}$  (we hide  $Y_2$ ), we have

$$\mathbb{P}[X \leq x | Y_1 = y] = \begin{cases} \mathbb{P}[Y_2 \leq x | Y_1 = y] = F_2(x) & \text{if } x \geq y; \\ 0 & \text{if } x < y. \end{cases}$$

If  $F_2(y) > 0$ , this function is discontinuous at  $x = y$ .

Likewise, if  $X = \min(Y_1, Y_2)$  instead, with  $\mathcal{G} = \mathcal{G}_{-1}$ , we have

$$\mathbb{P}[X \leq x | Y_1 = y] = \begin{cases} F_2(x) & \text{if } x < y; \\ 1 & \text{if } x \geq y, \end{cases}$$

If  $F_2(y) < 1$ , this function is discontinuous at  $x = y$ , so Assumption 1 is not satisfied.

Of course, there is no need to estimate  $f$  for this simple example, because it can be computed exactly. For  $X = \max(Y_1, Y_2)$ , we have  $F(x) = \mathbb{P}[\max(Y_1, Y_2) \leq x] = F_1(x)F_2(x)$  and then  $f(x) = F_1(x)f_2(x) + F_2(x)f_1(x)$ . See Section 4.5 for a generalization.

## 2.4. Convex combination of conditional density estimators

When there are many possible choices of  $\mathcal{G}$  for a given problem, one may try to pick the best one, but sometimes a better approach is to select more than one and take a convex linear combination of the corresponding CDEs as the final density estimator. This idea is well known for general mean estimators (Bratley et al. 1987). More specifically, suppose  $\hat{f}_{0,n}, \dots, \hat{f}_{q,n}$  are  $q + 1$  distinct unbiased density estimators. Typically, these estimators will be dependent and will be based on the same simulations. They could be all CDEs based on different choices of  $\mathcal{G}$  (so they will not hide the same information), but there could be non-CDEs as well. A convex combination can take the form

$$\hat{f}_n(x) = \beta_0 \hat{f}_{0,n}(x) + \dots + \beta_q \hat{f}_{q,n}(x) = \hat{f}_{0,n}(x) - \sum_{\ell=1}^q \beta_\ell (\hat{f}_{0,n}(x) - \hat{f}_{\ell,n}(x)) \quad (7)$$

for all  $x \in \mathbb{R}$ , where  $\beta_0 + \dots + \beta_q = 1$ . This combination is equivalent to choosing  $\hat{f}_{0,n}(x)$  as the main estimator, and taking the  $q$  differences  $\hat{f}_{0,n}(x) - \hat{f}_{\ell,n}(x)$  as control variables (Bratley et al.

1987). With this interpretation, the optimal coefficients  $\beta_\ell$  can be estimated via standard control variate theory (Asmussen and Glynn 2007) by trying to minimize the IV of  $\hat{f}_n(x)$  w.r.t. the  $\beta_\ell$ 's. More precisely, if we denote  $\text{IV}_\ell = \text{IV}(\hat{f}_{\ell,n}(x))$  and

$$\text{IC}_{\ell,k} = \int_a^b \text{Cov}[\hat{f}_{\ell,n}(x), \hat{f}_{k,n}(x)] dx,$$

we obtain

$$\text{IV} = \text{IV}(\hat{f}_n(x)) = \sum_{\ell=0}^q \beta_\ell^2 \text{IV}_\ell + 2 \sum_{0 \leq \ell < k \leq q} \beta_\ell \beta_k \text{IC}_{\ell,k}.$$

Given the  $\text{IV}_\ell$ 's and  $\text{IC}_{\ell,k}$ 's (or good estimates of them), this IV is a quadratic function of the  $\beta_\ell$ 's, which can be minimized exactly as in standard least-squares linear regression, to obtain estimates of the optimal coefficients  $\beta_j$ . Estimating the density and coefficients from the same data yields biased but consistent density estimators, and the bias is rarely a problem. We did this for some of the examples in Section 4.

A more refined approach is to allow the coefficients  $\beta_j$  to depend on  $x$ :

$$\hat{f}_n(x) = \beta_0(x) \hat{f}_{0,n}(x) + \cdots + \beta_q(x) \hat{f}_{q,n}(x) = \hat{f}_{0,n}(x) - \sum_{\ell=1}^q \beta_\ell(x) (\hat{f}_{0,n}(x) - \hat{f}_{\ell,n}(x)), \quad (8)$$

where  $\beta_0(x) + \cdots + \beta_q(x) = 1$  for all  $x \in \mathbb{R}$ . We can estimate the optimal coefficients by standard control variate theory at selected values of  $x$ , then for each  $\ell \geq 1$ , we can fit a smoothing spline to these estimated values, by least squares. This provides estimated optimal coefficients that are smooth functions of  $x$ , which can be used to obtain a final CDE. This type of strategy was used in L'Ecuyer and Buist (2008) to estimate varying control variate coefficients in a different setting. The additional flexibility can provide much more variance reduction in some situations.

## 2.5. A GLR density estimator

The *generalized likelihood ratio* (GLR) method, originally developed by Peng et al. (2018) to estimate the derivative of an expectation with respect to some model parameter, can be adapted to density estimation, as shown in Peng et al. (2020). We explain how this density estimator works to estimate  $f(x)$ , the density at  $x$ , in our setting. The assumptions stated below differ slightly from those in Peng et al. (2020). In particular, here we do not have a parameter  $\theta$ , the conditions on the estimator are required only in the area where  $X \leq x$ , and we add a condition to ensure finite variance. As in Section 2.3, we assume here that  $X = h(\mathbf{Y}) = h(Y_1, \dots, Y_d)$  where  $Y_1, \dots, Y_d$  are independent continuous random variables, and  $Y_j$  has cdf  $F_j$  and density  $f_j$ . Let  $P(x) = \{\mathbf{y} \in \mathbb{R}^d : h(\mathbf{y}) \leq x\}$ . For  $j = 1, \dots, d$ , let  $h_j(\mathbf{y}) := \partial h(\mathbf{y}) / \partial y_j$ ,  $h_{jj}(\mathbf{y}) := \partial^2 h(\mathbf{y}) / \partial y_j^2$ , and

$$\Psi_j(\mathbf{y}) = \frac{\partial \log f_j(y_j) / \partial y_j - h_{jj}(\mathbf{y}) / h_j(\mathbf{y})}{h_j(\mathbf{y})}. \quad (9)$$

ASSUMPTION 2. The Lebesgue measure of  $h^{-1}((x - \epsilon, x + \epsilon))$  in  $\mathbb{R}^d$  goes to 0 when  $\epsilon \rightarrow 0$  (this means essentially that the density is bounded around  $x$ ).

ASSUMPTION 3. The set  $P(x)$  is measurable, the functions  $h_j$ ,  $h_{jj}$ , and  $\Psi_j$  are well defined over it, and  $\mathbb{E}[\mathbb{I}[X \leq x] \cdot \Psi_j^2(\mathbf{Y})] < \infty$ .

PROPOSITION 1. Under Assumptions 2 and 3,  $\mathbb{I}[X \leq x] \cdot \Psi_j(\mathbf{Y})$  is an unbiased and finite-variance estimator of the density  $f(x)$  at  $x$ .

PROOF. See Peng et al. (2020). □

### 3. Combining RQMC with the CMC density estimator

We now discuss how RQMC can be used with the CDE, and under what conditions it can provide a convergence rate faster than  $\mathcal{O}(n^{-1})$  for the IV of the resulting unbiased estimator. For this, we first recall some basic facts about QMC and RQMC. More detailed coverages can be found in Niederreiter (1992), Dick and Pillichshammer (2010), and L'Ecuyer (2009, 2018), for example.

For a function  $g : [0, 1]^s \rightarrow \mathbb{R}$ , the integration error by the average over a point set  $P_n = \{\mathbf{u}_1, \dots, \mathbf{u}_n\} \subset [0, 1]^s$  is defined by

$$E_n = \frac{1}{n} \sum_{i=1}^n g(\mathbf{u}_i) - \int_{[0,1]^s} g(\mathbf{u}) d\mathbf{u}. \quad (10)$$

Classical QMC theory bounds this error as follows. Let  $\mathbf{v} \subseteq \mathcal{S} := \{1, \dots, s\}$  denote an arbitrary subset of coordinates. For any point  $\mathbf{u} = (u_1, \dots, u_s) \in [0, 1]^s$ ,  $\mathbf{u}_{\mathbf{v}}$  denotes the projection of  $\mathbf{u}$  on the coordinates in  $\mathbf{v}$  and  $(\mathbf{u}_{\mathbf{v}}, \mathbf{1})$  is the point  $\mathbf{u}$  in which  $u_j$  is replaced by 1 for each  $j \notin \mathbf{v}$ . Let  $g_{\mathbf{v}} := \partial^{|\mathbf{v}|} g / \partial \mathbf{u}_{\mathbf{v}}$  denote the partial derivative of  $g$  with respect to all the coordinates in  $\mathbf{v}$ . When  $g_{\mathbf{v}}$  exists and is continuous for  $\mathbf{v} = \mathcal{S}$  (i.e., for all  $\mathbf{v} \subseteq \mathcal{S}$ ), the *Hardy-Krause (HK) variation* of  $g$  can be written as

$$V_{\text{HK}}(g) = \sum_{\emptyset \neq \mathbf{v} \subseteq \mathcal{S}} \int_{[0,1]^{|\mathbf{v}|}} |g_{\mathbf{v}}(\mathbf{u}_{\mathbf{v}}, \mathbf{1})| d\mathbf{u}_{\mathbf{v}}. \quad (11)$$

On the other hand, the *star-discrepancy* of  $P_n$  is

$$D^*(P_n) = \sup_{\mathbf{u} \in [0,1]^s} \left| \frac{|P_n \cap [\mathbf{0}, \mathbf{u}]|}{n} - \text{vol}[\mathbf{0}, \mathbf{u}] \right|$$

where  $\text{vol}[\mathbf{0}, \mathbf{u}]$  is the volume of the box  $[\mathbf{0}, \mathbf{u}]$ . The classical *Koksma-Hlawka (KH) inequality* bounds the absolute error by the product of these two quantities, one that involves only the function  $g$  and the other that involves only the point set  $P_n$ :

$$|E_n| \leq V_{\text{HK}}(g) \cdot D^*(P_n). \quad (12)$$

There are explicit construction methods (e.g., digital nets, lattice rules, and polynomial lattice rules) of deterministic point sets  $P_n$  for which  $D^*(P_n) = \mathcal{O}((\log n)^{s-1}/n) = \mathcal{O}(n^{-1+\epsilon})$  for all  $\epsilon > 0$ . This means that functions  $g$  for which  $V_{\text{HK}}(g) < \infty$  can be integrated by QMC with a worst-case error that satisfies  $|E_n| = \mathcal{O}(n^{-1+\epsilon})$ . There are also known methods to randomize these point sets  $P_n$  in a way that each randomized point  $\mathbf{u}_i$  has the uniform distribution over  $[0, 1]^s$ , so  $\mathbb{E}[E_n] = 0$ , and the  $\mathcal{O}(n^{-1+\epsilon})$  discrepancy bound is preserved, which gives

$$\text{Var}[E_n] = \mathbb{E}[E_n^2] = \mathcal{O}(n^{-2+\epsilon}). \quad (13)$$

The classical definitions of variation and discrepancy given above are in fact only one pair among an infinite collection of possibilities. There are other versions of (12), with different definitions of the discrepancy and the variation, such that there are known point set constructions for which the discrepancy converges as  $\mathcal{O}(n^{-\alpha+\epsilon})$  for  $\alpha > 1$ , but the conditions on  $g$  to have finite variation are more restrictive (more smoothness is required) Dick and Pillichshammer (2010).

From a practical viewpoint, getting a good estimate or an upper bound on the variation of  $g$  that can be useful to bound the RQMC variance is a notoriously difficult problem. Even just showing that the variation is finite is not always easy. However, finite variation is not a necessary condition. In many realistic applications in which variation is known to be infinite, RQMC can nevertheless reduce the variance by a large factor (L'Ecuyer 2009, L'Ecuyer and Munger 2012, He and Wang 2015). The appropriate explanation for this depends on the application. In many cases, part of the explanation is that the integrand  $g$  can be written as a sum of orthogonal functions (as in an ANOVA decomposition) and a set of terms in that sum have a large variance contribution and are smooth low-dimensional functions for which RQMC is very effective (L'Ecuyer and Lemieux 2000, L'Ecuyer 2009, Lemieux 2009). Making such a decomposition and finding the important terms is usually difficult for realistic problems, but to apply RQMC in practice, this is not needed. The usual approach in applications is to try it and compare the RQMC variance with the MC variance empirically. We will do that in Section 4. To estimate the RQMC variance, we usually replicate the RQMC scheme  $n_r$  times independently, using the same point set but with  $n_r$  independent randomizations, then we compute the empirical mean and variance of the  $n_r$  independent realizations of the RQMC density estimator  $(1/n) \sum_{i=1}^n g(\mathbf{U}_i)$ .

To combine the CDE with RQMC, we must be able to write  $F(x | \mathcal{G}) = \tilde{g}(x, \mathbf{u})$  and  $F'(x | \mathcal{G}) = \tilde{g}'(x, \mathbf{u}) = d\tilde{g}(x, \mathbf{u})/dx$  for some function  $\tilde{g} : [a, b] \times [0, 1]^s$ . The function  $\tilde{g}'(x, \cdot)$  will act as  $g$  in (10). The combined CDE+RQMC estimator  $\hat{f}_{\text{cde-rqmc},n}(x)$  will be defined by

$$\hat{f}_{\text{cde-rqmc},n}(x) = \frac{1}{n} \sum_{i=1}^n \tilde{g}'(x, \mathbf{U}_i), \quad (14)$$

which is the RQMC version of (2). To estimate the RQMC variance, we can perform  $n_r$  independent randomizations to obtain  $n_r$  independent realizations of  $\hat{f}_{\text{cde},n}$  in (2) with MC and  $n_r$  independent realizations of  $\hat{f}_{\text{cde-rqmc},n}$  in (14) with RQMC, and compute the empirical IV. We will do this to estimate the IV in our experiments.

Even though it is rarely done in practice, we think it is instructive to examine how the HK variation of  $g \equiv \tilde{g}'(x, \cdot)$  can be bounded in our CDE setting, at least for small examples. This will prove that the variance bound (13) holds for the CDE for these examples. To prove the bounded HK variation, we need to take the partial derivative of  $\tilde{g}'(x, \mathbf{u})$  with respect to each subset of coordinates of  $\mathbf{u}$  and show that the integral of each such partial derivative is finite. We show how this can be done for some of our earlier examples.

EXAMPLE 6. Consider a sum of random variables as in Example 1, with  $\mathcal{G} = \mathcal{G}_{-k}$  summarized by the single real number  $S_{-k}$ . We have  $F(x | \mathcal{G}) = F_k(x - S_{-k})$  and  $F'(x | \mathcal{G}) = f_k(x - S_{-k})$ . Without loss of generality, let  $k = d$ . Suppose that each  $Y_j$  is generated by inversion from  $U_j \sim \mathcal{U}(0, 1)$ , so  $Y_j = F_j^{-1}(U_j)$  and  $S_{-d} = F_1^{-1}(U_1) + \dots + F_s^{-1}(U_s)$  with  $s = d - 1$ . This gives  $\tilde{g}(x, \mathbf{U}) = F_d(x - S_{-d}) = F_d(x - F_1^{-1}(U_1) - \dots - F_s^{-1}(U_s))$  and  $\tilde{g}'(x, \mathbf{U}) = f_d(x - S_{-d}) = f_d(x - F_1^{-1}(U_1) - \dots - F_s^{-1}(U_s))$ . The partial derivatives of this last function are

$$\tilde{g}'_{\mathbf{v}}(x, \mathbf{U}_{\mathbf{v}}, \mathbf{1}) = f_d^{(|\mathbf{v}|)}(x - S_{-d}) \prod_{j \in \mathbf{v}} \frac{\partial(F_j^{-1}(U_j))}{\partial U_j}.$$

So the functions  $F_j^{-1}$  must be differentiable over  $(0, 1)$  for  $j = 1, \dots, d - 1$ , the density  $f_d$  must be  $s$  times differentiable, and the integral of  $|\tilde{g}'_{\mathbf{v}}(x, \mathbf{u}_{\mathbf{v}}, \mathbf{1})|$  with respect to  $\mathbf{u}_{\mathbf{v}}$  must be finite. Under these conditions, the HK variation is finite.

For Example 2, if  $\mathcal{G} = \mathcal{G}_{-2}$ , we have  $Y_1 = \sigma_1 \Phi^{-1}(U_1)$  where  $U_1 \sim \mathcal{U}(0, 1)$ . Then,  $F(x | \mathcal{G}_{-2}) = F_2(x - Y_1) = \Phi((x - Y_1)/\sigma_2)$  and  $F'(x | \mathcal{G}_{-2}) = \phi((x - \sigma_1 \Phi^{-1}(U_1))/\sigma_2)/\sigma_2 = \tilde{g}'(x, U_1)$ . Taking the derivative with respect to  $u$  and noting that  $d\Phi^{-1}(u)/du = 1/(\phi(\Phi^{-1}(u)))$  yields

$$\tilde{g}'_{\mathbf{v}}(x, u) = \frac{\phi'((x - \sigma_1 \Phi^{-1}(u))/\sigma_2) \sigma_1}{\sigma_2^2 \phi(\Phi^{-1}(u))}$$

for  $\mathbf{v} = \{1\} = \mathcal{S}$  (the only subset in this case). Integrating this with respect to  $u$  by making the change of variable  $z = \Phi^{-1}(u)$  gives

$$\int_0^1 \tilde{g}'_{\mathbf{v}}(x, u) du = \frac{\sigma_1}{\sigma_2^2} \int_{-\infty}^{\infty} |\phi'((x - \sigma_1 z)/\sigma_2)| dz < \infty,$$

because  $|\phi'(\cdot)|$  is bounded by  $\phi(\cdot)$  multiplied by the absolute value of a polynomial of degree 1. So we have finite HK variation.

For Example 4, with  $\mathcal{G} = \mathcal{G}_{-2}$  and  $Y_1 = U_1$ , we have  $\tilde{g}'(x, \mathbf{u}) = \tilde{g}'(x, U_1) = \mathbb{I}[U_1 \leq x \leq \epsilon + U_1]/\epsilon = \mathbb{I}[x - \epsilon \leq U_1 \leq x]/\epsilon$ . This function is not continuous, but its HK variation (not given by (11) in

this case) is  $2/\epsilon < \infty$ , because it is piecewise constant with only two jumps, each one of size  $1/\epsilon$ . Thus, the HK variation is unbounded when  $\epsilon \rightarrow 0$ , but it is finite for any fixed  $\epsilon$ . The behavior with  $\mathcal{G} = \mathcal{G}_{-1}$  is similar and the HK variation is 2 in that case, which is much better.

For all these examples in which the HK variation is finite, we have a proof that the MISE of the CDE+RQMC density estimator converges as  $\mathcal{O}(n^{-2+\epsilon})$  for any  $\epsilon > 0$ .

On the other hand, the GLR estimator in Proposition 1 is typically discontinuous because of the indicator function, and therefore its HK variation is usually infinite.

## 4. Examples and numerical experiments

We now examine larger examples, summarize the results of our numerical experiments with the CDE and CDE+RQMC, and make comparisons with KDE+RQMC and with GLR.

### 4.1. Experimental setting

Since the CDE is unbiased, we measure its performance by the IV, which equals the MISE in this case. To approximate the IV estimator (3) for a given  $n$ , we first take a stratified sample  $e_1, \dots, e_{n_e}$  of  $n_e$  *evaluation points* at which the empirical variance will be computed. We sample  $e_j$  uniformly in  $[a + (j-1)(b-a)/n_e, a + j(b-a)/n_e]$  for  $j = 1, \dots, n_e$ . Then we use the unbiased IV estimator

$$\widehat{\text{IV}} = \frac{(b-a)}{n_e} \sum_{j=1}^{n_e} \widehat{\text{Var}}[\hat{f}_n(e_j)],$$

where  $\widehat{\text{Var}}[\hat{f}_n(e_j)]$  is the empirical variance of the CDE at  $e_j$ , obtained as follows. We repeat the following  $n_r$  times, independently: Generate  $n$  observations of  $X$  from the density  $f$  with the given method (MC or RQMC), and compute the CDE at each evaluation point  $e_j$ . Then, compute  $\widehat{\text{Var}}[\hat{f}_n(e_j)]$  as the empirical variance of the  $n_r$  density estimates at  $e_j$ , for each  $j$ .

To estimate the convergence rate of the IV as a function of  $n$  with the different methods, we fit a model of the form  $\text{IV} \approx Kn^{-\nu}$ . For the CDE with independent points (no RQMC), this model holds exactly with  $\nu = 1$ . We hope to observe  $\nu > 1$  with RQMC. The parameters  $K$  and  $\nu$  are estimated by linear regression in log-log scale, i.e., by fitting the model  $\log \text{IV} \approx \log K - \nu \log n$  to data. Since  $n$  is always taken as a power of 2, we report the logarithms in base 2. We estimated the IV for  $n = 2^{14}, \dots, 2^{19}$  (6 values) to fit the regression model. We also report the observed  $-\log_2 \text{IV}$  for  $n = 2^{19}$  and use *e19* as a shorthand for this value in the tables. We use exactly the same procedure for the GLR. For the KDE, these values are for the MISE instead of the IV. We report results with the following types of point sets:

- (1) independent points (MC);
- (2) a randomly-shifted lattice rule (Lat+s);
- (3) a randomly-shifted lattice rule with a baker's transformation (Lat+s+b);
- (4) Sobol' points with a left random matrix scramble and random digital shift (Sobol'+LMS).

The short names in parentheses are used in the plots and tables. For the definitions and properties of these RQMC point sets, see L'Ecuyer and Lemieux (2000), Owen (2003), L'Ecuyer (2009, 2018). They are implemented in SSJ (L'Ecuyer 2016), which we used for our experiments. The parameters of the lattice rules were found with the Lattice Builder software of L'Ecuyer and Munger (2016), using a fast-CBC construction method with the  $\mathcal{P}_2$  criterion and order dependent weights  $\gamma_v = \rho^{|v|}$ , with  $\rho$  ranging from 0.05 to 0.8, depending on the example (a larger  $\rho$  was used when the dimension  $s$  was smaller). The baker's transformation sometimes improves the convergence rate by making the integrand periodic (Hickernell 2002), but it usually also increases the variation of the integrand, so its impact on the variance can go either way.

For some examples, we tried CDEs based on different choices of  $\mathcal{G}$  and a convex combination as in Section 2.4. For comparison, we provide results for the GLR and for a KDE with a normal kernel and a bandwidth  $h$  selected by the methodology described in Ben Abdellah et al. (2019).

## 4.2. A sum of normals

We start with a very simple example in which the density  $f$  is known beforehand, so there is no real need to estimate it, but this type of example is very convenient for testing the performance of our density estimators. Let  $Z_1, \dots, Z_d$  be independent standard normal random variables, i.e., with mean 0 and variance 1, and define

$$X = (a_1 Z_1 + \dots + a_d Z_d) / \sigma, \quad \text{where} \quad \sigma^2 = a_1^2 + \dots + a_d^2.$$

Then  $X$  is also standard normal, with density  $f(x) = \phi(x) \stackrel{\text{def}}{=} \exp(-x^2/2) / \sqrt{2\pi}$  and cdf  $\mathbb{P}[X \leq x] = \Phi(x)$  for  $x \in \mathbb{R}$ . The term  $a_j Z_j$  in the sum has variance  $a_j^2$ . We pretend we do not know this and we estimate  $f(x)$  over the interval  $[-2, 2]$ , which contains slightly more than 95% of the density. We also tried larger intervals, such as  $[-5, 5]$ , and IVs for the CDE were almost the same.

To construct the CDE, we define  $\mathcal{G}_{-k}$  as in Example 1, for any  $k = 1, \dots, d$ . That is, we hide  $Z_k$  and estimate the cdf by

$$F(x | \mathcal{G}_{-k}) = \mathbb{P} \left[ a_k Z_k \leq x\sigma - \sum_{j=1, j \neq k}^d a_j Z_j \mid \mathcal{G}_{-k} \right] = \Phi \left( \frac{x\sigma}{a_k} - \frac{1}{a_k} \sum_{j=1, j \neq k}^d a_j Z_j \right).$$

The CDE becomes

$$F'(x | \mathcal{G}_{-k}) = \phi \left( \frac{x\sigma}{a_k} - \frac{1}{a_k} \sum_{j=1, j \neq k}^d a_j Z_j \right) \frac{\sigma}{a_k} = \phi \left( \frac{x\sigma}{a_k} - \frac{1}{a_k} \sum_{j=1, j \neq k}^d a_j \Phi^{-1}(U_j) \right) \frac{\sigma}{a_k} \stackrel{\text{def}}{=} \tilde{g}'(x, \mathbf{U})$$

for  $x \in \mathbb{R}$ , where  $\mathbf{U} = (U_1, \dots, U_{k-1}, U_{k+1}, \dots, U_d)$ ,  $Z_j = \Phi^{-1}(U_j)$ , and the  $U_j$  are independent  $\mathcal{U}(0, 1)$  random variables. Assumption 1 is easily verified, so this CDE is unbiased.

For CMC+MC (independent sampling), we get an exact formula for the variance of the CDE directly from Example 2, by taking in that example  $Y_2 = a_k Z_k / \sigma$  and  $Y_1 = X - Y_2$ , whose variances are  $\sigma_2^2 = (a_k / \sigma)^2$  and  $\sigma_1^2 = 1 - \sigma_2^2$ , and plugging these values into (6). To prove that RQMC gives a better convergence rate for the variance than MC, it suffices to show that  $V_{\text{HK}}(\tilde{g}'(x, \cdot)) < \infty$  for any  $x$ . This can be done by the same argument as in the second part of Example 6. Then we expect to observe a convergence rate near  $\mathcal{O}(n^{-2})$ , at least when  $d$  is small.

For the GLR method, with  $Y_j = Z_j a_j / \sigma \sim \mathcal{N}(0, a_j^2 / \sigma^2)$ , we obtain  $\partial(\log f_j(y_j)) / \partial y_j = -y_j \sigma^2 / a_j^2$ ,  $h_j(y_j) = 1$ ,  $h_{jj}(y_j) = 0$ , and then  $\Psi_j = -Y_j \sigma^2 / a_j^2 = -Z_j \sigma / a_j$ . Note that we could also replace  $Y_j$  by  $Z_j$  and  $f_j$  by  $\phi_j$  (the standard normal density), which would give  $\partial(\log \phi_j(z_j)) / \partial z_j = -z_j$ ,  $h_j(z_j) = a_j / \sigma$ ,  $h_{jj}(y_j) = 0$ , and again  $\Psi_j = -Z_j \sigma / a_j$ .

In our first experiment, we take  $a_j = 1$  for all  $j$ , and  $k = d$ . By symmetry, the true IV is the same for any other  $k$ . Table 1 reports the estimated rate  $\hat{\nu}$  and the estimated value of  $\text{e19} = -\log_2(\text{IV})$  for  $n = 2^{19}$ , for various values of  $d$  and sampling methods. The rows marked CDE-1 give the results for  $k = d$ , while those labeled CDE-Avg are for a convex combination (7) with equal weights  $\beta_\ell = 1/d$  for all  $\ell = k - 1$ , after computing the CDE for each  $k$  from the same simulations.

For MC, the rates  $\hat{\nu}$  agree with the (known) exact asymptotic rates of  $\nu = 1$  for the CDE and GLR, and  $\nu = 0.8$  for the KDE. By looking at  $\text{e19}$ , we see that the MISE with MC is much smaller for the CDE than for the GLR and KDE, for example for  $d = 2$  by a factor of about 32 for CDE-1 and about 70 for CDE-avg. For  $d = 20$ , the gains are more modest. RQMC methods provide huge improvements for small  $d$  with the CDE. We observe rates  $\hat{\nu}$  larger than 2 for  $d = 2$  and 3, and by looking at the exponents  $\text{e19}$ , we see that for  $d = 3$ , for example, the MISE goes from  $2^{-17}$  for the GLR and KDE to about  $2^{-42}$  for CDE-1 with Sobol' points with LMS. This is a reduction factor of about  $2^{25} \approx 33$  millions for  $n = 2^{19}$ . The large values of  $\hat{\nu}$  imply of course that this factor is smaller for smaller  $n$ . When  $d$  is large, such as  $d = 20$ , RQMC brings only a small gain. The values of  $\hat{\nu}$  are sometimes noisy. For GLR with Lat+s and  $d = 5$ , for example, the large  $\hat{\nu} = 1.45$  comes from the fact that the IV for  $n = 2^{14}$  (not shown) is unusually large (an outlier). Looking at  $\text{e19}$  gives a more robust assessment of the performance. GLR performs better than the KDE under RQMC for small  $d$ , but not comparable to the CDE. Under MC, GLR is slightly worse than the KDE.

In our second experiment, we take  $a_j^2 = 2^{1-j}$  for  $j = 1, \dots, d$ . Now, the choice of  $k$  for the CDE makes a difference, and the best choice will obviously be  $k = 1$ , i.e., hide the term that has the largest variance. Note that with MC,  $\text{Var}[X] = 2 - 2^{-d}$ , and when we apply CMC by hiding  $a_k Z_k$  from the sum, we hide a term of variance  $a_k^2 = 2^{1-k}$  and generate a partial sum  $S_{-k}$  of variance  $2 - 2^{1-k} - 2^{-d}$ . Both terms have a normal distribution with mean 0. The results of Example 2 hold with these variances. Table 2 reports the numerical results for  $d = 11$  and  $k = 1, 2, 5, 11$ .

**Table 1** Values of  $\hat{\nu}$  and e19 for a CDE, and a convex combination of CDEs, and a KDE, for a sum of  $d = k$  normals with  $a_j = 1$ , over  $[-2, 2]$ .

		$\hat{\nu}$					e19				
		$d = 2$	$d = 3$	$d = 5$	$d = 10$	$d = 20$	$d = 2$	$d = 3$	$d = 5$	$d = 10$	$d = 20$
CDE-1	MC	0.99	0.98	1.02	1.00	1.02	22.1	21.4	20.8	19.8	19.2
	Lat+s	2.83	2.00	1.85	1.40	1.04	52.3	39.8	32.1	23.6	19.7
	Lat+s+b	2.69	2.11	1.69	1.14	1.05	50.5	41.5	31.1	21.8	20.0
	Sob+LMS	2.62	2.10	1.81	1.04	1.04	49.3	40.7	31.1	21.3	19.7
CDE-avg	MC	1.06	0.92	1.03	1.01	1.01	23.4	22.1	21.6	20.6	19.8
	Lat+s	2.79	1.84	1.33	1.19	1.05	53.3	39.8	32.2	23.0	20.6
	Lat+s+b	2.65	1.90	1.71	1.05	1.08	51.6	41.4	32.3	23.4	21.3
	Sob+LMS	2.60	2.10	1.92	1.02	1.03	49.8	42.0	33.0	22.7	20.5
GLR	MC	0.98	0.95	1.03	1.05	1.00	17.0	16.1	15.9	14.9	14.1
	Lat+s	1.51	1.56	1.45	0.94	1.06	28.2	24.9	22.1	17.8	17.2
	Lat+s+b	1.49	1.41	1.05	1.06	1.04	27.3	23.9	20.4	18.8	17.6
	Sob+LMS	1.49	1.33	1.15	0.99	1.16	27.5	24.0	21.0	18.3	17.4
KDE	MC	0.79	0.80	0.76	0.75	0.77	17.0	17.0	16.9	16.9	17.0
	Lat+s	1.08	1.39	0.92	0.97	0.76	25.1	22.4	19.4	18.2	17.4
	Lat+s+b	1.23	0.94	0.72	0.73	0.74	24.1	20.1	18.1	17.3	17.2
	Sob+LMS	1.18	0.98	0.83	0.74	0.77	24.4	20.8	17.9	17.2	17.1

**Table 2** Values of  $\hat{\nu}$  and e19 with a CDE for selected choices of  $\mathcal{G}_{-k}$ , for a linear combination of  $d = 11$  normals with  $a_j^2 = 2^{1-j}$ .

	$\hat{\nu}$				e19			
	$k = 1$	$k = 2$	$k = 5$	$k = 11$	$k = 1$	$k = 2$	$k = 5$	$k = 11$
MC	1.00	1.02	1.01	1.00	22.2	21.0	18.8	15.5
Lat+s	1.43	1.48	1.34	1.04	30.3	28.5	22.8	15.6
Lat+s+b	1.57	1.65	1.28	1.02	33.5	30.8	22.1	15.6
Sob+LMS	1.78	1.56	1.21	1.02	34.1	30.4	21.7	15.7

The MC rates  $\hat{\nu}$  agree again with the theory, but here the IV depends very much on the choice of  $k$ , and this effect is more significant when  $k$  is smaller. For example, for Sobol' points, the IV with  $k = 1$  is about 300,000 times smaller than with  $k = 11$ . The reason is that with  $k = 11$ , we hide only a variable having a very small variance, so the CDE for one sample is a high narrow peak, and the HK variation of  $\tilde{g}'(x, \mathbf{u})$  is very large. For  $k = 1$  or 2, we have the opposite and the integrand is much more RQMC-friendly.

### 4.3. Displacement of a cantilever beam

We consider the following model for the displacement  $X$  of a cantilever beam with horizontal and vertical loads, taken from Bingham (2017):

$$X = h(Y_1, Y_2, Y_3) = \frac{4\ell^3}{Y_1 w t} \sqrt{\frac{Y_2^2}{w^4} + \frac{Y_3^2}{t^4}} \quad (15)$$

in which  $\ell = 100$ ,  $w = 4$  and  $t = 2$  are constants (in inches), while  $Y_1$  (Young's modulus),  $Y_2$  (the horizontal load), and  $Y_3$  (the vertical load), are independent normal random variables,  $Y_j \sim$

$\mathcal{N}(\mu_j, \sigma_j^2)$ , i.e., normal with mean  $\mu_j$  and variance  $\sigma_j^2$ . The parameter values are  $\mu_1 = 2.9 \times 10^7$ ,  $\sigma_1 = 1.45 \times 10^6$ ,  $\mu_2 = 500$ ,  $\sigma_2 = 100$ ,  $\mu_3 = 1000$ ,  $\sigma_3 = 100$ . We will denote  $\kappa = 4\ell^3/(wt) = 5 \times 10^5$ . The goal is to estimate the density of  $X$  over the interval  $[3.1707, 5.6675]$ , which covers about 99% of the density (it clips 0.5% on each side). It is possible to have  $X < 0$  in this model, but the probability is  $\mathbb{P}[Y_1 < 0] = \Phi(-20) = 2.8 \times 10^{-89}$ , which is negligible. This example fits the framework of Section 2.3, with  $d = 3$ . We can hide any of the three random variables for the conditioning, and we will examine each case.

Conditioning on  $\mathcal{G}_{-1}$  means hiding  $Y_1$ . We have

$$X = \frac{\kappa}{Y_1} \sqrt{\frac{Y_2^2}{w^4} + \frac{Y_3^2}{t^4}} \leq x \quad \text{if and only if} \quad Y_1 \geq \frac{\kappa}{x} \sqrt{\frac{Y_2^2}{w^4} + \frac{Y_3^2}{t^4}} \stackrel{\text{def}}{=} W_1(x).$$

Note that  $W_1(x) > 0$  if and only if  $x > 0$ . For  $x > 0$ ,

$$F(x | \mathcal{G}_{-1}) = \mathbb{P}[Y_1 \geq W_1(x) | W_1(x)] = 1 - \Phi((W_1(x) - \mu_1)/\sigma_1)$$

which is continuous and differentiable in  $x$ , and

$$F'(x | \mathcal{G}_{-1}) = -\phi((W_1(x) - \mu_1)/\sigma_1) W_1'(x)/\sigma_1 = \phi((W_1(x) - \mu_1)/\sigma_1) W_1(x)/(x\sigma_1).$$

If we condition on  $\mathcal{G}_{-2}$  instead, i.e., we hide  $Y_2$ , we have  $X \leq x$  if and only if

$$Y_2^2 \leq w^4 \left( (xY_1/\kappa)^2 - Y_3^2/t^4 \right) \stackrel{\text{def}}{=} W_2(x).$$

If  $W_2(x) \leq 0$ , then  $F'(x | \mathcal{G}_{-2}) = F(x | \mathcal{G}_{-2}) = \mathbb{P}[X \leq x | W_2(x)] = 0$ . For  $W_2(x) > 0$ , we have

$$\begin{aligned} F(x | \mathcal{G}_{-2}) &= \mathbb{P}[X \leq x | W_2(x)] = \mathbb{P}\left[-\sqrt{W_2(x)} \leq Y_2 \leq \sqrt{W_2(x)} | W_2(x)\right] \\ &= \Phi((\sqrt{W_2(x)} - \mu_2)/\sigma_2) - \Phi(-(\sqrt{W_2(x)} + \mu_2)/\sigma_2), \end{aligned}$$

which is again continuous and differentiable in  $x$ , and

$$F'(x | \mathcal{G}_{-2}) = \frac{\phi((\sqrt{W_2(x)} - \mu_2)/\sigma_2) + \phi(-(\sqrt{W_2(x)} + \mu_2)/\sigma_2)}{(\sigma_2 \sqrt{W_2(x)})/(w^4 x (Y_1/\kappa)^2)} > 0.$$

If we condition on  $\mathcal{G}_{-3}$ , the analysis is the same as for  $\mathcal{G}_{-2}$ , by symmetry, and we get

$$F'(x | \mathcal{G}_{-3}) = \frac{\phi((\sqrt{W_3(x)} - \mu_3)/\sigma_3) + \phi(-(\sqrt{W_3(x)} + \mu_3)/\sigma_3)}{(\sigma_3 \sqrt{W_3(x)})/(t^4 x (Y_1/\kappa)^2)} > 0$$

for  $W_3(x) > 0$ , where  $W_3(x)$  is defined in a similar way as  $W_2(x)$ . In addition to testing these three ways of conditioning, we also tested a convex combination of the three, as explained in Section 2.4, with coefficients  $\beta_\ell$  that do not depend on  $x$ .

**Table 3** Values of  $\hat{\nu}$  and e19 with a CDE for each choice of  $\mathcal{G}_{-k}$ , for the best convex combination, for the GLR, and for the KDE, for the cantilever beam model.

	$\hat{\nu}$						e19					
	$\mathcal{G}_{-1}$	$\mathcal{G}_{-2}$	$\mathcal{G}_{-3}$	comb.	GLR	KDE	$\mathcal{G}_{-1}$	$\mathcal{G}_{-2}$	$\mathcal{G}_{-3}$	comb.	GLR	KDE
MC	0.97	0.98	0.99	0.98	1.02	0.76	19.3	14.5	22.8	22.5	14.1	15.8
Lat+s	1.99	1.95	2.06	2.04	1.38	1.03	39.8	25.2	41.6	41.9	23.4	21.9
Lat+s+b	2.24	2.08	2.27	2.25	1.37	0.93	44.5	23.7	46.8	47.0	23.3	21.0
Sob+LMS	2.21	2.03	2.21	2.21	1.32	0.97	44.0	23.6	45.7	46.1	23.4	21.5

For GLR using  $Y_1$ , let  $C = C(Y_2, Y_3) = (4\ell^3/wt)\sqrt{Y_2^2/w^4 + Y_3^2/t^4}$ . Then, we have  $X = h(\mathbf{Y}) = C/Y_1$ ,  $h_1(\mathbf{Y}) = -CY_1^{-2}$ ,  $h_{11}(\mathbf{Y}) = 2CY_1^{-3}$ ,  $\partial \log f_1(Y_1)/\partial Y_1 = (Y_1 - \mu_1)/\sigma_1^2$ , and

$$\Psi_1 = \frac{Y_1}{C} (Y_1(Y_1 - \mu_1)/\sigma_1^2 - 2).$$

Table 3 summarizes the results. The MISE is about  $2^{-47}$  for the best CDE+RQMC compared with  $2^{-15.8}$  for usual KDE+MC, a gain by a factor of over  $2^{31} \approx 2$  billions. With RQMC, the convergence rate  $\hat{\nu}$  is around 2 in all cases with the CDE methods, and much less for GLR and KDE. GLR benefits significantly from RQMC, more than the KDE, but cannot compete with the CDE. For the lattice rules, the baker's transformation helps significantly for the CDE.

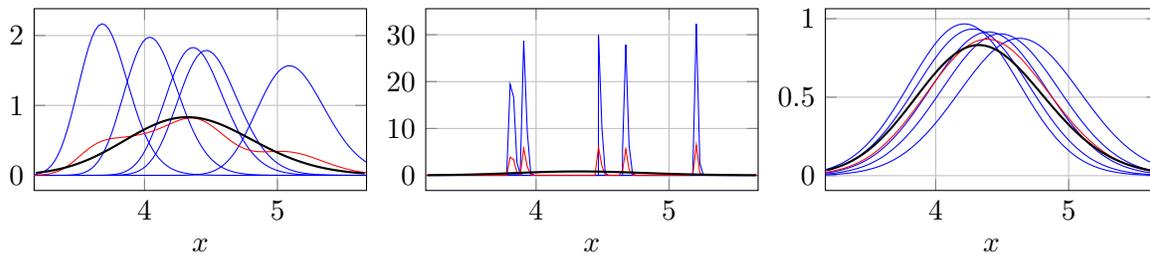
We also see that conditioning on  $\mathcal{G}_{-2}$  does not give as much reduction than for the other choices. The reason is that the conditional density in this case is a high narrow peak, similar to what we saw for  $k = 11$  at the end of Example 4.2. To provide visual insight, Figure 2 shows plots of five realizations of the conditional density for  $\mathcal{G}_{-1}$ ,  $\mathcal{G}_{-2}$ , and  $\mathcal{G}_{-3}$ . The realizations of  $F'(\cdot | \mathcal{G}_{-2})$  have high narrow peaks. The average of the five realizations is shown in red and the true density in black. In Figure 3, we zoom in on part of the estimated densities to show the difference between MC and RQMC. In each panel one can see the CDE using MC (in red), using RQMC (in green), and the "true density" (black, dashed) estimated with RQMC using a very large number of samples. We have  $\mathcal{G}_{-1}$  with  $n = 2^{10}$  on the left and  $\mathcal{G}_{-2}$  with  $n = 2^{16}$  on the right. In both cases, the RQMC estimate is closer to the true density, and on the right it oscillates less. If we repeat this experiment several times, the red curve would vary much more than the green one across the realizations.

#### 4.4. Buckling strength of a steel plate

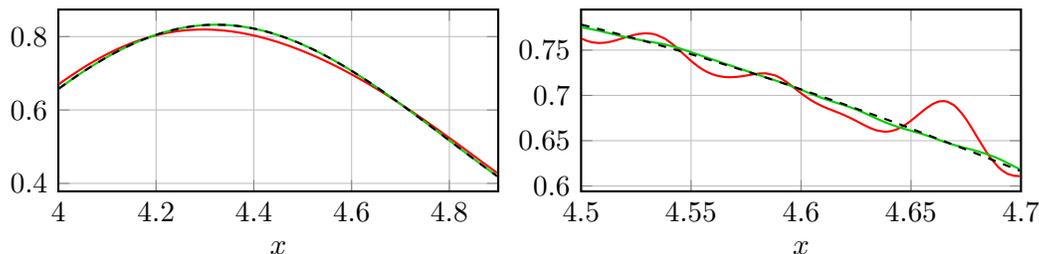
This is a higher-dimensional example, with  $d = 6$ , taken from Schields and Zhang (2016). It models the buckling strength of a steel plate by

$$X = \left( \frac{2.1}{\Lambda} - \frac{0.9}{\Lambda^2} \right) \left( 1 - \frac{0.75Y_5}{\Lambda} \right) \left( 1 - \frac{2Y_6Y_2}{Y_1} \right), \tag{16}$$

where  $\Lambda = (Y_1/Y_2)\sqrt{Y_3/Y_4}$ , and  $Y_1, \dots, Y_6$  are independent random variables whose distributions are given in Table 4. Each distribution is either normal or lognormal, and the table gives the mean



**Figure 2** Five realizations of the density conditional on  $\mathcal{G}_{-k}$  (blue), their average (red), and the true density (thick black) for  $k=1$  (left),  $k=2$  (middle), and  $k=3$  (right), for the cantilever example.



**Figure 3** The CDE under MC (red), under RQMC (green) and the true density (black, dashed) for  $\mathcal{G}_{-1}$  with  $n=2^{10}$  (left) and for  $\mathcal{G}_{-2}$  with  $n=2^{16}$  (right), for the cantilever example.

and the coefficient of variation (cv), which is the standard deviation divided by the mean. We estimate the density of  $X$  over  $[a, b] = [0.5169, 0.6511]$ , which contains about 99% of the density (leaving out 0.5% on each side). There is a nonzero probability of having  $Y_4 \leq 0$ , in which case  $X$  is undefined, but this probability is extremely small and this has a negligible impact on the density estimator over  $[a, b]$ , so we just ignore it (alternatively we could truncate the density of  $Y_4$ ). There are also negligible probabilities that the density estimates below are negative and we ignore this.

**Table 4** Distribution of each parameter for the buckling strength model.

parameter	distribution	mean	cv
$Y_1$	normal	23.808	0.028
$Y_2$	lognormal	0.525	0.044
$Y_3$	lognormal	44.2	0.1235
$Y_4$	normal	28623	0.076
$Y_5$	normal	0.35	0.05
$Y_6$	normal	5.25	0.07

For this example, computing the density of  $X$  conditional on  $\mathcal{G}_{-5}$  or  $\mathcal{G}_{-6}$  (i.e., when hiding  $Y_5$  or  $Y_6$ ) is relatively easy, so we will try and compare these two choices. If we hide one of the variables that appear in  $\Lambda$ , the CDE would be harder to compute (it would require to solve a polynomial equation of degree 4 for each sample), and we do not do it. Let us define

$$V_1 = \frac{2.1}{\Lambda} - \frac{0.9}{\Lambda^2}, \quad V_2 = 1 - \frac{2Y_6Y_2}{Y_1}, \quad \text{and} \quad V_3 = 1 - \frac{3Y_5}{4\Lambda}.$$

Then we have

$$X \leq x \iff Y_5 \geq \left(1 - \frac{x}{V_1 V_2}\right) \frac{4\Lambda}{3}$$

and

$$F'(x | \mathcal{G}_{-5}) = f_5 \left( \left(1 - \frac{x}{V_1 V_2}\right) \frac{4\Lambda}{3} \right) \frac{4\Lambda}{3V_1 V_2} = \phi \left( \frac{(1 - x/(V_1 V_2)) 4\Lambda/3 - 0.35}{0.0175} \right) \frac{4\Lambda}{0.0525 \cdot V_1 V_2}.$$

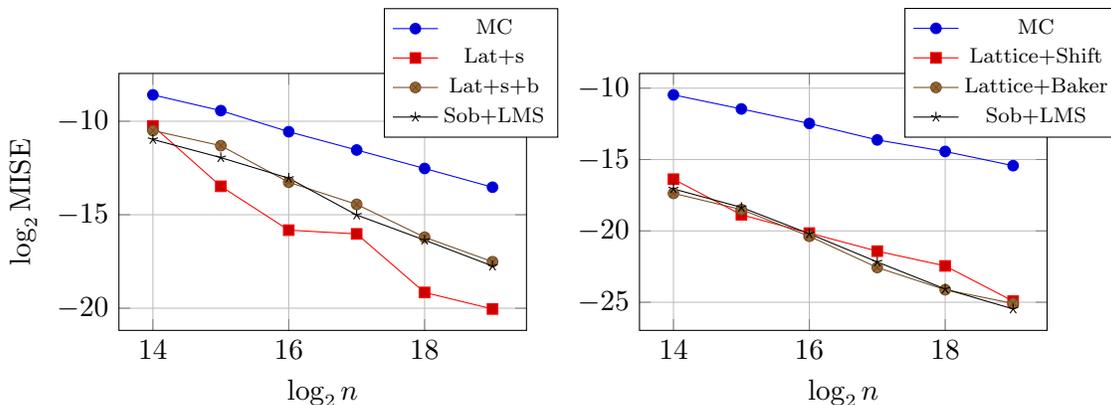
Similarly,

$$F'(x | \mathcal{G}_{-6}) = f_6 \left( \left(1 - \frac{x}{V_1 V_3}\right) \frac{Y_1}{2Y_2} \right) \frac{Y_1}{2Y_2 V_1 V_3} = \phi \left( \frac{(1 - x/(V_1 V_3)) Y_1/(2Y_2) - 5.25}{0.3675} \right) \frac{Y_1}{0.735 \cdot Y_2 V_1 V_3}.$$

For GLR using  $Y_6$ , let  $C = (2.1/\Lambda - 0.9/\Lambda^2)(1 - 0.75Y_5/\Lambda)$ . We have  $X = h(\mathbf{Y}) = C(1 - 2Y_6 Y_2/Y_1)$ ,  $h_6(\mathbf{Y}) = 2CY_2/Y_1$ ,  $h_{66}(\mathbf{Y}) = 0$ ,  $\partial \log f_6(Y_6)/\partial Y_6 = -(Y_6 - \mu_6)/\sigma_6^2$ , and  $\Psi_6 = Y_1(Y_6 - \mu_6)/(2CY_2\sigma_6^2)$ .

**Table 5** Values of  $\hat{\nu}$  and e19 with a CDE for  $\mathcal{G}_{-5}$ ,  $\mathcal{G}_{-6}$ , and their combination, for the buckling strength model.

	$\hat{\nu}$					e19				
	$\mathcal{G}_{-5}$	$\mathcal{G}_{-6}$	comb.	GLR	KDE	$\mathcal{G}_{-5}$	$\mathcal{G}_{-6}$	comb.	GLR	KDE
MC	1.00	1.00	1.00	0.98	0.76	13.5	15.4	15.4	10.2	11.7
Lat+s	1.89	1.56	1.56	1.29	0.81	20.0	24.9	24.9	16.6	13.7
Lat+s+b	1.46	1.65	1.60	1.19	0.85	17.5	25.1	25.1	15.9	12.7
Sob+LMS	1.40	1.75	1.75	1.16	0.81	17.7	25.5	25.5	15.9	12.4

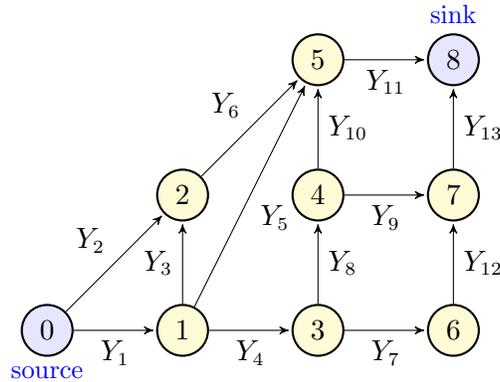


**Figure 4** MISE vs  $n$  in log-log scale for the  $\mathcal{G} = \mathcal{G}_{-5}$  (left) and  $\mathcal{G} = \mathcal{G}_{-6}$  (right) for the buckling strength model.

Table 5 summarizes the results. We see again that the CDE with RQMC performs very well and much better than the GLR and KDE, that it is much better to condition on  $\mathcal{G}_{-6}$  than on  $\mathcal{G}_{-5}$ , and that combining the two provides no significant improvement. The GLR is also better than the KDE under RQMC, but not under MC. Figure 4 displays the IV as a function of  $n$  in a log-log-scale for the CDE with  $\mathcal{G}_{-5}$  and  $\mathcal{G}_{-6}$ . It unveils a slightly more erratic behavior of the MISE for the shifted lattice rule (Lat+s) than for the other methods; the performance depends on the choice of parameters of the lattice rule and their interaction with the particular integrand.

#### 4.5. A stochastic activity network

In this example, the conditioning must hide more than one random variable. The example is taken from Avramidis and Wilson (1998) and was also used in L'Ecuyer and Lemieux (2000) and L'Ecuyer and Munger (2012). We consider the stochastic activity network of Fig. 5, where arc  $j$  has random length  $Y_j$  for  $j = 1, \dots, 13$ . The  $Y_j$  are independent with continuous cdf's  $F_j$  as given in L'Ecuyer and Munger (2012), and are generated by inversion:  $Y_j = F_j^{-1}(U_j)$  where  $U_j \sim U(0, 1)$ . The goal is to estimate the density of  $X$  which represents the length of the longest path from the source to the sink, over  $[a, b] = [22, 106.24]$ , which covers about 95% of the density.



**Figure 5** A stochastic activity network

Here,  $X$  is defined by a maximum (over all paths from 0 to 8), and if we hide only a single random variable  $Y_j$ , we run into the same problem as in Example 5: Assumption 1 does not hold, because  $F(\cdot | \mathcal{G})$  has a jump. We must hide more (condition on less). We apply the conditioning proposed by Avramidis and Wilson (1998) by hiding  $\{Y_j, j \in \mathcal{L} \stackrel{\text{def}}{=} \{5, 6, 7, 9, 10\}\}$ , and let  $\mathcal{G}$  represent  $\{Y_j, j \notin \mathcal{L}\}$ . The corresponding conditional cdf is

$$F(x | \mathcal{G}) = \mathbb{P}[X < x | \{Y_j : j \notin \mathcal{L}\}] = \prod_{j \in \mathcal{L}} \mathbb{P}[Y_j \leq x - P_j] = \prod_{j \in \mathcal{L}} F_j(x - P_j) \quad (17)$$

where  $P_j$  is the length of the longest path that goes through arc  $j$  when we exclude  $Y_j$  from that length. The conditional density is

$$F'(x | \mathcal{G}) = \frac{d}{dx} F(x | \mathcal{G}) = \sum_{j \in \mathcal{L}} f_j(x - P_j) \prod_{l \in \mathcal{L}, l \neq j} F_l(x - P_l).$$

Under this conditioning, since the  $Y_j$ 's are continuous variables with bounded variance, Assumption 1 holds, so  $F'(x | \mathcal{G})$  is an unbiased density estimator with uniformly bounded variance.

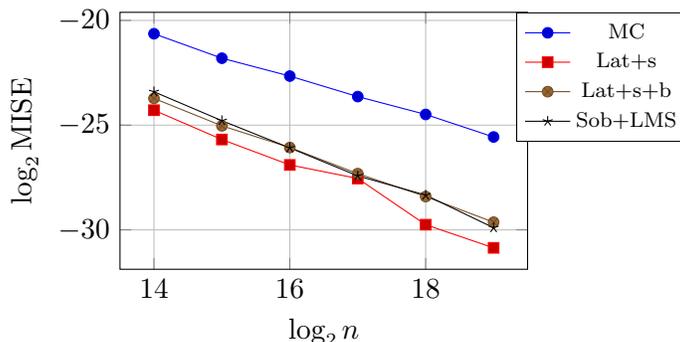
The GLR method described in Section 2.5 does not work for this example. Indeed, with  $X = h(\mathbf{Y})$  defined as the length of the longest path, for any  $j$ , the derivative  $h_j(\mathbf{Y})$  is zero whenever arc  $j$  is

not on the longest path, so we would need to select an arc  $j$  that is guaranteed to be always on the longest path. But there is no such arc. We could apply a modified GLR that selects a cut instead of a single coordinate  $Y_j$ , but this is beyond the scope of this paper.

**Table 6** Values of  $\hat{\nu}$  and e19 for the SAN example.

		$\hat{\nu}$	e19
CDE	MC	0.96	25.6
	Lat+s	1.31	30.9
	Lat+s+b	1.17	29.6
	Sob+LMS	1.27	29.9
KDE	MC	0.78	20.9
	Lat+s	0.95	22.7
	Lat+s+b	0.93	22.0
	Sob+LMS	0.74	21.9

Table 6 and Figure 6 summarize our results. We see that for  $n = 2^{19}$ , the CDE outperforms the KDE by a factor of about 20 with MC, and by a factor of about  $2^8 \approx 250$  with RQMC. Interestingly, here the lattice rules work better without the baker's transformation.



**Figure 6** MISE vs  $n$  in log-log scale, for the SAN example.

#### 4.6. Density of the failure time of a system

We consider a  $d$ -component system in which each component starts in the operating mode (state 1) and fails (jumps to state 0) at a certain random time, to stay there forever. Let  $Y_j$  be the failure time of component  $j$  for  $j = 1, \dots, d$ . For  $t \geq 0$ , let  $W_j(t) = \mathbb{I}[Y_j > t]$  be the state of component  $j$  and  $\mathbf{W}(t) = (W_1(t), \dots, W_d(t))^t$  the system state, at time  $t$ . The system is in the failed mode at time  $t$  if and only if  $\Phi(\mathbf{W}(t)) = 0$ , where  $\Phi : \{0, 1\}^d \rightarrow \{0, 1\}$  is called the *structure function*. Let  $X = \inf\{t \geq 0 : \Phi(\mathbf{W}(t)) = 0\}$  be the random time when the system fails. We want to estimate the density of  $X$ . A straightforward way of simulating a realization of  $X$  is to generate the component lifetimes  $Y_j = \inf\{t \geq 0 : W_j(t) = 0\}$  for  $j = 1, \dots, d$ , and then compute  $X$  from that.

As in Section 4.5, the GLR method of Section 2.5 does not work for this example, because  $h_j(\mathbf{Y}) \neq 0$  only when  $X = Y_j$ , and there is no  $j$  for which this is sure to happen.

If the  $Y_j$  are independent and exponential, one can construct a CMC estimator of the cdf  $F(x) = \mathbb{P}[X \leq x]$  as follows (Gertsbakh and Shpungin 2010, Botev et al. 2013, 2016). Generate all the  $Y_j$ 's and sort them in increasing order. Then, erase their values and retain only their order, which is a permutation  $\pi$  of  $\{1, \dots, d\}$ . Compute the critical number  $C = C(\pi)$ , defined as the number of component failures required for the system to fail (that is, the system fails at the  $C$ th component failure, for the given  $\pi$ ). Note that  $C$  can also be computed by starting with all components failed and resurrecting them one by one in reverse order of their failure, until the system becomes operational. Computing  $C$  using this reverse order is often more efficient (Botev et al. 2016). Then compute the conditional cdf  $\mathbb{P}[X \leq x \mid \pi]$ , where  $X$  is the time of the  $C$ th component failure. This is an unbiased estimator of  $F(x)$  with smaller variance than the indicator  $\mathbb{I}[X \leq x]$ . It can also be shown that in an asymptotic regime in which the component failure rates converge to 0 so that  $1 - F(x) \rightarrow 0$ , the relative variance of this CMC estimator of  $1 - F(x)$  remains bounded whereas it goes to infinity with the conventional estimator  $\mathbb{I}[X > x]$ ; i.e., the CMC estimator has bounded relative error (Botev et al. 2013, 2016).

When the lifetimes are independent and exponential,  $X$  is a sum of  $C$  independent exponentials, so it has a hypoexponential distribution, whose cdf has an explicit formula that can be written in terms of a matrix exponential, and developed explicitly as a sum of products in terms of the rates of the exponential lifetimes, as explained in Botev et al. (2016). By taking the derivative of the conditional cdf formula with respect to  $x$ , one obtains the conditional density.

More specifically, let component  $j$  have an exponential lifetime with rate  $\lambda_j > 0$ , for  $j = 1, \dots, d$ . For a given realization, let  $\pi(j)$  be the  $j$ th component that fails and let  $C(\pi) = c$  for the given  $\pi$ , let  $A_1$  be the time until the first failure, and let  $A_j$  be the time between the  $(j-1)$ th and  $j$ th failures, for  $j > 1$ . Conditional on  $\pi$ , we have  $X = A_1 + \dots + A_c$  where the  $A_j$ 's are independent and  $A_j$  is exponential with rate  $\Lambda_j$  for all  $j \geq 1$ , with  $\Lambda_1 = \lambda_1 + \dots + \lambda_d$ , and  $\Lambda_j = \Lambda_{j-1} - \lambda_{\pi(j)}$  for all  $j \geq 2$ . The conditional distribution of  $X$  is then hypoexponential with cdf

$$\mathbb{P}[X \leq x \mid \pi] = \mathbb{P}[A_1 + \dots + A_c \leq x \mid \pi] = 1 - \sum_{j=1}^c p_j e^{-\Lambda_j x},$$

where

$$p_j = \prod_{k=1, k \neq j}^c \frac{\Lambda_k}{\Lambda_k - \Lambda_j}.$$

See Gertsbakh and Shpungin (2010), Appendix A, and Botev et al. (2016), for example. Taking the derivative with respect to  $x$  gives the CDE

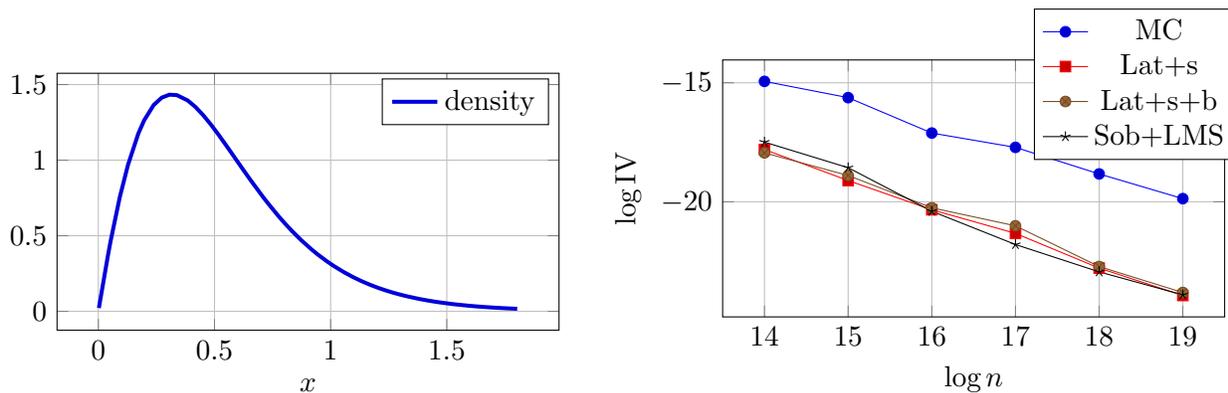
$$f(x \mid \pi) = \sum_{j=1}^c \Lambda_j p_j e^{-\Lambda_j x},$$

in which  $c$ , the  $\Lambda_j$  and the  $p_j$  depend on  $\pi$ . This conditional density is well defined and computable everywhere in  $[0, \infty)$ . There are instability issues for computing  $p_j$  when  $\Lambda_k - \Lambda_j$  is close to 0 for some  $k \neq j$ , but this can be addressed by a stable numerical algorithm of Higham (2009). All of this can be generalized easily to a model in which the lifetimes are dependent, with the dependence modeled by a Marshall-Olkin copula (Botev et al. 2016). In that model, the  $Y_j$  represent the occurrence times of shocks that can take down one or more components simultaneously.

It is interesting to note that although  $f(x | \pi)$  is an unbiased estimator of the density  $f(x)$  at any  $x$ , this estimator is a function of the permutation  $\pi$  only, so it takes its values in a finite set, which means that the corresponding  $\tilde{g}(\mathbf{u})$  is a piecewise constant function, which is not RQMC-friendly. Therefore, we do not expect RQMC to bring a very large gain.

**Table 7** Values of  $\hat{\nu}$  and e19 with the CDE, for the network reliability example.

	$\hat{\nu}$	e19
MC	1.00	19.9
Lat+s	1.22	23.9
Lat+s+b	1.19	23.8
Sob+LMS	1.33	23.9



**Figure 7** Density (left) and log IV as a function of  $\log n$  (right) for the network failure time.

For a numerical illustration, we take the same graph as in Section 4.5. For  $j = 1, \dots, 13$ ,  $Y_j$  is exponential with rate  $\lambda_j$  and the  $Y_j$  are independent. The system fails as soon as there is no path going from 0 to 8. For simplicity, here we take  $\lambda_j = 1$  for all  $j$ , although taking different  $\lambda_j$ 's brings no significant additional difficulty. We estimate the density over the interval  $(a, b] = (0, 1.829]$ , which cuts off roughly 1% of the probability on the right side. Table 7 and Figure 7 give the results. The density of  $X$  estimated with  $n = 2^{20}$  random samples is shown on the left and the IV plots are on the right. Despite the discontinuity of  $\tilde{g}$ , RQMC outperforms MC in terms of the IV by a factor of about  $2^4 = 16$  for  $n = 2^{19}$ , and also by improving the empirical rate  $\hat{\nu}$  to about  $-1.2$  for lattices and even better with Sobol' points. The Sobol' points used here were constructed using LatNet

Builder (Marion et al. 2020) with a CBC search based on the  $t$ -value of all projections up to order 6, with order-dependent weights  $\gamma_k = 0.8^k$  for projections of order  $k$ .

#### 4.7. Density of waiting times in a single queue

**4.7.1. Model with independent days.** We consider a single-server FIFO queue in which customers arrive from an arbitrary arrival process (not necessarily stationary Poisson) and the service times are independent, with continuous cdf  $G$  and density  $g$ . If  $W$  denotes the waiting time of a “random” customer, we want to estimate  $p_0 = \mathbb{P}[W = 0]$  and the density  $f$  of  $W$  over  $(0, \infty)$ .

We first consider a system that starts empty and evolves over a fixed time horizon  $\tau$ , which we call a *day*. Let  $T_j$  be the arrival time of the  $j$ th customer,  $T_0 = 0$ ,  $A_j = T_j - T_{j-1}$  the  $j$ th interarrival time,  $S_j$  the service time of customer  $j$ , and  $W_j$  the waiting time of customer  $j$ . Since the system starts empty, we have  $W_1 = 0$ , and the Lindley recurrence gives us that  $W_j = \max(0, W_{j-1} + S_{j-1} - A_j)$  for  $j \geq 2$ . At time  $\tau$ , the arrival process stops, but service continues until all customers already arrived are served. The number of customers handled in a day is the random variable  $N = \max\{j \geq 1 : T_j < \tau\}$ . The cdf of  $W$  can be written as  $F(0) = p_0$  and for  $x > 0$ ,  $F(x) = \mathbb{P}[W \leq x] = \mathbb{E}[\mathbb{I}(W \leq x)]$ . Note that the sequence of waiting times of all customers over an infinite number of independent successive days is a regenerative process that regenerates at the beginning of each day. We can then apply the renewal reward theorem, which gives

$$F(x) = \mathbb{E}[\mathbb{I}(W \leq x)] = \frac{\mathbb{E}[\mathbb{I}[W_1 \leq x] + \cdots + \mathbb{I}[W_N \leq x]]}{\mathbb{E}[N]}. \quad (18)$$

Since  $\mathbb{E}[N]$  does not depend on  $x$ , we see that for  $x > 0$ , the density  $f(x)$  is the derivative of the numerator with respect to  $x$ , divided by  $\mathbb{E}[N]$ .

To obtain a differentiable cdf estimator, we want to replace each indicator in the numerator by a conditional expectation. One simple way of doing this is to hide the service time  $S_{j-1}$  of the previous customer; that is, replace  $\mathbb{I}[W_j \leq x]$  by

$$P_j(x) = \mathbb{P}[W_j \leq x \mid W_{j-1} - A_j] = \mathbb{P}[S_{j-1} \leq x + A_j - W_{j-1}] = G(x + A_j - W_{j-1}) \quad \text{for } x \geq 0.$$

This gives  $P_j(0) = G(A_j - W_{j-1})$  (there is a probability mass at 0), whereas for  $x > 0$ , we have  $P'_j(x) = dP_j(x)/dx = g(x + A_j - W_{j-1})$  and then, since  $N$  does not change when we change  $x$ ,

$$f(x) = \frac{\mathbb{E}[D(x)]}{\mathbb{E}[N]} \quad \text{where } D(x) = \sum_{j=1}^N g(x + A_j - W_{j-1}). \quad (19)$$

Note that we are not conditioning on the same information for all terms of the sum, so what we do is not exactly CMC, but *extended CMC*. It nevertheless provides the required smoothing and an unbiased density estimator for the numerator of (18). The denominator  $\mathbb{E}[N]$  can be estimated in

the usual way. We are then in the standard setting of estimating a ratio of expectations (Asmussen and Glynn 2007), for which we have unbiased estimators for the numerator and the denominator.

We simulate  $n$  days, independently (with MC) or with  $n$  RQMC points, to obtain  $n$  realizations of  $(N, D(x))$ , say  $(N_1, D_1(x)), \dots, (N_n, D_n(x))$ . The ratio estimator (CDE) of  $f(x)$  is

$$\hat{f}(x) = \frac{\sum_{i=1}^n D_i(x)}{\sum_{i=1}^n N_i}. \quad (20)$$

It can be computed at any  $x \in [0, \infty)$ . For independent realizations (with MC), the variance of  $\hat{f}(x)$  can be estimated using the delta method for ratio estimators (Asmussen and Glynn 2007):

$$n \text{Var}[\hat{f}(x)] \rightarrow \frac{\text{Var}[D_i(x)] + \text{Var}[N_i]f^2(x) - 2\text{Cov}[D_i(x), N_i]f(x)}{\mathbb{E}^2[N_i]}$$

asymptotically, when  $n \rightarrow \infty$ . This variance can be estimated by replacing the unknown quantities in this expression by their empirical values. This is consistent because the  $n$  pairs  $(D_i(x), N_i)$ ,  $i = 1, \dots, n$ , are independent. Alternatively, a confidence interval on  $f(x)$  can also be computed with a bootstrap approach (Choquet et al. 1999).

In the RQMC case, the pairs  $(D_i(x), N_i)$  are no longer independent. Then, to obtain an estimator of  $f(x)$  for which we can estimate the variance, we make  $n_r$  independent replicates of the RQMC estimator of the pair  $(\mathbb{E}[D(x)], \mathbb{E}[N])$ , say  $(\bar{D}_1(x), \bar{N}_1), \dots, (\bar{D}_{n_r}(x), \bar{N}_{n_r})$ , where each  $(\bar{D}_j(x), \bar{N}_j)$  is the average of  $n$  pairs  $(D_i(x), N_i)$  sampled by RQMC. We estimate the density  $f(x)$  by the ratio of the two grand sums

$$\hat{f}_{\text{rqmc}, n_r}(x) = \frac{\sum_{j=1}^{n_r} \bar{D}_j(x)}{\sum_{j=1}^{n_r} \bar{N}_j}.$$

To estimate the variance, we use that

$$\text{Var}[\hat{f}_{\text{rqmc}, n_r}(x)] \approx \frac{\text{Var}[\bar{D}_j(x)] + \text{Var}[\bar{N}_j]f^2(x) - 2\text{Cov}[\bar{D}_j(x), \bar{N}_j]f(x)}{n_r(\mathbb{E}[N])^2}$$

and we replace all the unknown quantities in this expression by their empirical values.

Here, the required dimension of the RQMC points is the (random) total number of inter-arrival times  $A_j$  and service times  $S_j$  that we need to generate during the day. It is approximately twice the number of customers that arrive during the day. This number is unbounded, so the RQMC points must have unbounded (or infinite) dimension. Recurrence-based RQMC point sets have this property; they can be provided for instance by ordinary or polynomial Korobov lattice rules (L'Ecuyer and Lemieux 1999, 2000, 2002, Lemieux and L'Ecuyer 2003), which are available in the hups package of SSJ (L'Ecuyer 2016).

**4.7.2. Steady-state model.** In a slightly different setting, we can assume that the single queue evolves in steady-state over an infinite time horizon, under the additional assumptions that the  $A_j$ 's are i.i.d. and the  $S_j$ 's are also i.i.d. Again, we want to estimate the density of the waiting time  $W$  of a random customer. In this case, the system regenerates whenever a new customer arrives in an empty system. The regenerative cycles can be much shorter on average than for the previous case, unless the day is very short or the utilization factor of the system is close to 1. The CDE has exactly the same form, apart from the different definition of regenerative cycle. In this case  $n$  represents the number of regenerative cycles,  $N_i$  is the number of customers in the  $i$ th cycle and  $D_i(x)$  is the realization of  $D(x)$  over the  $i$ th cycle.

In both settings, one could also hide  $A_j$  instead of  $S_{j-1}$ . The density estimator is similar and easy to derive. Intuition says that this should be a better choice if  $A_j$  has more variance than  $S_{j-1}$ .

**4.7.3. The GLR estimator** Peng et al. (2020), Section 4.2.2., show how to construct a GLR estimator for the density of the *sojourn time* of customer  $j$  in this single-queue model. If the service times  $S_j$  are lognormal with parameters  $(\mu, \sigma^2)$ , we can write

$$X = W_j = \max(0, W_{j-1} + S_{j-1} - A_j) = \max(0, W_{j-1} + \exp[\sigma Z_{j-1} + \mu] - A_j) =: h(\mathbf{Y})$$

where  $Z_{j-1}$  has the standard normal density  $\phi$ , and  $\mathbf{Y} = (Y_1, Y_2, Y_3) = (Z_{j-1}, A_j, W_{j-1})$ . When  $W_j > 0$ , taking the derivative of  $h$  with respect to  $Y_1 = Z_{j-1}$  gives  $h_1(\mathbf{Y}) = \exp[\sigma Z_{j-1} + \mu] \sigma = S_{j-1} \sigma$ ,  $h_{11}(\mathbf{Y}) = S_{j-1} \sigma^2$ , and these derivatives are 0 when  $W_j = 0$ . We also have  $\partial \log \phi(x) / \partial x = -x$ , and therefore for  $x > 0$ ,  $f(x) = \mathbb{E}[L(x)] / \mathbb{E}[N]$  where  $L(x) = \sum_{j=1}^N \mathbb{I}[W_j \leq x] \cdot \Psi_j$  and  $\Psi_j = -(Z_{j-1} + \sigma) / (S_{j-1} \sigma)$ . We can do  $n$  runs to estimate each of the two expectations in the ratio. This provides a very similar density estimator as with the CDE in (19), but here  $L(x)$  is discontinuous in  $x$ , whereas  $D(x)$  in (19) is continuous.

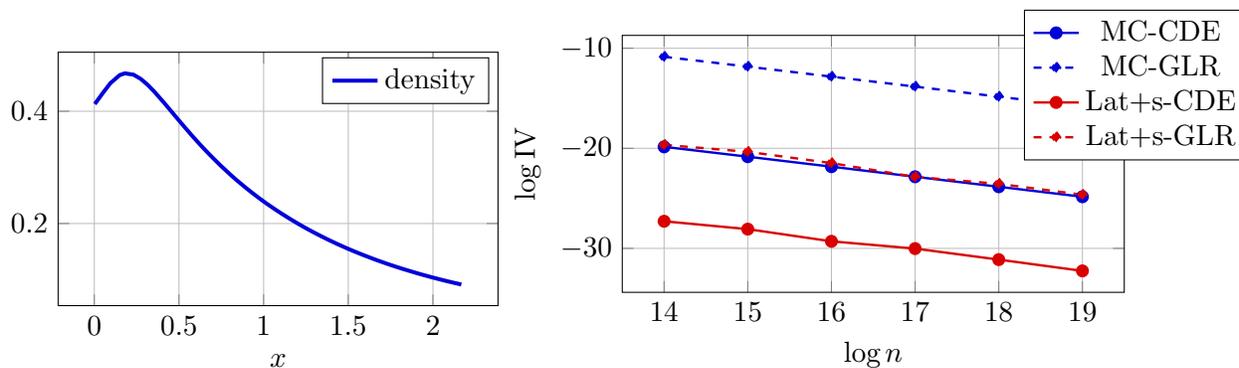
**4.7.4. Numerical results.** For a numerical illustration, let the arrival process be Poisson with constant rate  $\lambda = 1$ , and the service times  $S_j$  lognormal with parameters  $(\mu, \sigma^2) = (-0.7, 0.4)$ . This gives  $\mathbb{E}[S_j] = e^{-0.5} \approx 0.6065$  and  $\text{Var}[S_j] = e^{-1}(e^{0.4} - 1) \approx 0.18093$ . For RQMC, we use infinite-dimensional RQMC points defined by Korobov lattice rules (L'Ecuyer and Lemieux 2000) selected with *Lattice Builder* (L'Ecuyer and Munger 2016) using order-dependent weights  $\gamma_k = 0.005^k$  for projections of order  $k$ . We do not use Sobol' points because they are finite-dimensional.

*Finite-horizon case.* For the finite-horizon case, take  $\tau = 60$ . If the time is in minutes, this means that the “day” is one hour. The results for  $(a, b] = (0, 2.2]$  are in Table 8 and Figure 8. An important observation is that CDE+MC provides an unbiased density estimator all over  $[0, \infty)$ . Due to the large and random dimensionality of the required RQMC points, and more importantly the discontinuity of the derivative of the estimator with respect to the underlying uniforms (because

of the max, the HK variation is infinite), it was unclear if RQMC could bring any significant gain for this example. We see that although RQMC does not improve  $\hat{\nu}$  significantly, it improves the IV itself by a factor of about  $2^{8.5} \approx 180$  for  $n = 2^{19}$ , which is quite significant. We also see that CDE beats GLR by a factor of about 500 with MC and about 200 with RQMC.

**Table 8** Values of  $\hat{\nu}$  and e19 for the single queue example, finite-horizon case.

		$\hat{\nu}$	e19
CDE	MC	1.00	24.8
	Lat+s	0.99	32.3
	Lat+s+b	1.02	32.3
GLR	MC	1.00	15.8
	Lat+s	1.03	24.6
	Lat+s+b	1.08	25.0



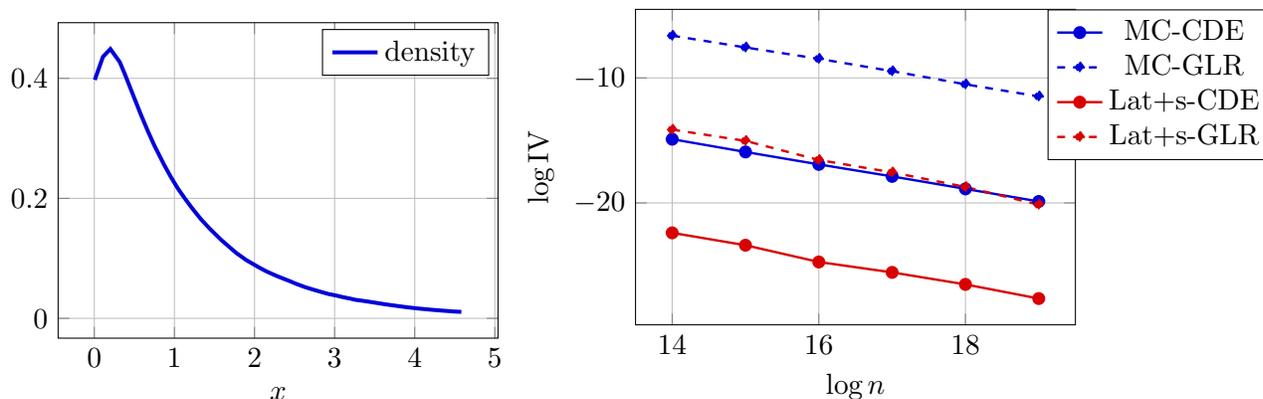
**Figure 8** Estimated density (left) and log IV as a function of log  $n$  (right) for the single queue over a finite-horizon.

*Steady-state case.* We performed a similar experiment using regenerative simulation for the steady-state model. The density is similar but not exactly the same as in the finite-horizon case. The results are in Table 9 and Figure 9. They are very similar to those of the finite-horizon case, with similar empirical convergence rates, and the IV for  $n = 2^{19}$  is again about 180 times smaller with CDE+RQMC compared to CDE+MC. The IV for GLR with  $n = 2^{19}$  is roughly 1000 times larger than with CDE with MC and 250 times larger than with CDE with RQMC. The only important difference is that here, the IV is about 30 times *larger* than in the finite-horizon case, for all the methods. The explanation is that in the finite-horizon, we simulate  $n$  runs with about 60 customers per run, whereas in the steady-state case, we have about 2.5 customers per regenerative cycle on average, so we simulate about 25 times fewer customers. Interestingly, the fact that we use much more coordinates of the RQMC points in the finite-horizon case (on average) makes no significant difference. A similar observation was made by L'Ecuyer and Lemieux (2000), Section 10.3, who compared finite-horizon runs of 5000 customers each on average, with regenerative simulation, in the context of estimating the probability of a large waiting time using RQMC. The reason why

RQMC performs well even for a very large time horizon is that the integrand has low effective dimension in the successive-dimensions sense (L'Ecuyer and Lemieux 2000).

**Table 9** Values of  $\hat{\nu}$  and e19 for the single queue example, steady-state case.

		$\hat{\nu}$	e19
CDE	MC	0.99	19.9
	Lat+s	1.04	27.6
	Lat+s+b	1.08	27.8
GLR	MC	0.99	11.5
	Lat+s	1.20	20.1
	Lat+s+b	1.21	20.4



**Figure 9** Estimated density (left) and log IV as a function of  $\log n$  (right) for the single queue in steady-state.

#### 4.8. A function of a multivariate normal vector

Suppose we want to estimate the density of  $X = h(\mathbf{Y})$  over some interval  $(a, b)$ , where  $h$  is a continuously differentiable function  $h: \mathbb{R}^s \rightarrow \mathbb{R}$ , and  $\mathbf{Y} = (Y_1, \dots, Y_s)^t$  has a multivariate normal distribution with mean vector  $\boldsymbol{\mu}$  and covariance matrix  $\mathbf{C}$ . To generate realizations of  $\mathbf{Y}$ , we can first decompose  $\mathbf{C}$  as  $\mathbf{C} = \mathbf{A}\mathbf{A}^t$  (once for all), then for each realization we generate a vector  $\mathbf{Z}$  of  $s$  independent standard normals (e.g., by inversion), and put  $\mathbf{Y} = \boldsymbol{\mu} + \mathbf{A}\mathbf{Z}$ . There are many ways of decomposing  $\mathbf{C} = \mathbf{A}\mathbf{A}^t$ . The choice of decomposition does not matter for ordinary MC, but for the CDE and/or RQMC, it does matter (Glasserman 2004, L'Ecuyer 2009).

Denoting  $F(x) = \mathbb{P}[X \leq x] = \mathbb{P}[h(\mathbf{Y}) \leq x]$ , the density can be written as

$$f(x) = F'(x) = \lim_{\delta \rightarrow 0} \frac{F(x + \delta) - F(x)}{\delta} = \lim_{\delta \rightarrow 0} \frac{\mathbb{P}[x < h(\mathbf{Y}) \leq x + \delta]}{\delta}.$$

For the CDE, we will hide one component of the vector  $\mathbf{Z} = (Z_1, \dots, Z_s)^t$ , say  $Z_j$ , and condition on the other ones. Let  $\mathbf{Z}_{-j}$  denote the vector  $\mathbf{Z}$  with coordinate  $j$  removed, and let  $\tilde{\mathbf{Y}}(z) = \tilde{\mathbf{Y}}(z; \mathbf{Z}_{-j})$

and  $\gamma(z) = \gamma(z; \mathbf{Z}_{-j}) = h(\tilde{\mathbf{Y}}(z))$  denote the values of  $\mathbf{Y}$  and of  $h(\mathbf{Y})$  as functions of  $Z_j = z$  when  $\mathbf{Z}_{-j}$  is fixed. We will assume from here that for any realization of  $\mathbf{Z}_{-j}$ ,  $\gamma(z; \mathbf{Z}_{-j})$  is a monotone non-decreasing function of  $z$ , so that  $\gamma^{-1}(x) = \inf\{z \in \mathbb{R} : \gamma(z) \geq x\}$  is well defined for any  $x$ . Conditional on  $\mathbf{Z}_{-j}$ , we then have

$$\mathbb{P}[x < h(\mathbf{Y}) \leq x + \delta \mid \mathbf{Z}_{-j}] = \mathbb{P}[x < \gamma(Z_j) \leq x + \delta \mid \mathbf{Z}_{-j}] = \mathbb{P}[z < Z_j \leq z + \Delta \mid \mathbf{Z}_{-j}] \approx \phi(z)\Delta$$

where  $z = \gamma^{-1}(x)$ ,  $z + \Delta = \gamma^{-1}(x + \delta)$ , and  $\phi$  is the standard normal density. Taking the limit gives

$$f(x \mid \mathbf{Z}_{-j}) = \lim_{\delta \rightarrow 0} \frac{\mathbb{P}[z < Z_j \leq z + \Delta \mid \mathbf{Z}_{-j}]}{\delta} = \lim_{\delta \rightarrow 0} \frac{\phi(z)\Delta}{\delta} = \frac{\phi(z)}{\gamma'(z)} = \frac{\phi(\gamma^{-1}(x))}{\gamma'(\gamma^{-1}(x))}.$$

Note that the function  $\gamma$  and its inverse  $\gamma^{-1}$  depend on  $\mathbf{Z}_{-j}$ . This means that in general, computing  $z = \gamma^{-1}(x)$  for any given  $x$  means inverting a different function for each sample realization.

**Arithmetic average for a GBM.** One class of applications (among others) leading to this multivariate normal model is the following. Consider a geometric Brownian motion (GBM) observed at times  $0 = t_0 < t_1 < \dots < t_s$ . The GBM is defined as  $\{S(t) = \exp(Y(t)), t \geq 0\}$  where  $\{Y(t), t \geq 0\}$  is a Brownian motion with drift and variance parameters  $\mu$  and  $\sigma^2$ , so we have  $Y(t_0) = 0$  and

$$Y(t_j) = Y(t_{j-1}) + (t_j - t_{j-1})\mu + \sqrt{t_j - t_{j-1}}\sigma Z_j, \quad (21)$$

where  $Z_1, \dots, Z_s$  are i.i.d.  $N(0, 1)$ . The vector  $\mathbf{Y} = (Y_1, \dots, Y_s)^t = (Y(t_1), \dots, Y(t_s))^t$  has a multivariate normal distribution with mean vector  $\boldsymbol{\mu} = (t_1, \dots, t_s)\mu$  and covariance matrix  $\mathbf{C}$  with elements  $c_{i,j} = t_i\sigma^2$  for  $i < j$ .

One example of payoff function in this setting is  $X = h(\mathbf{Y}) = \max(0, \bar{S} - K)$  where  $\bar{S} = (S_1 + \dots + S_s)/s$ ,  $S_j = S(t_j) = e^{Y_j}$  for all  $j$ , and  $K > 0$  is a constant. This  $X$  can represent the payoff of an Asian option (we ignore the discount factor). The density of  $X$  over the interval  $(0, b)$  for  $b > 0$  is the same as the density of  $\bar{S}$  over the interval  $(K, K + b)$ , so we will focus on estimating the density of  $\bar{S}$ . Thus, in our numerical examples, we will simply take  $X = h(\mathbf{Y}) = \bar{S}$ .

**Hiding the last observation.** One simple way to apply the CDE is to simulate the process  $Y$  sequentially as in (21) but hide  $Z_s$  and take the conditional expectation instead. This is equivalent to using the Cholesky decomposition for  $\mathbf{C}$ . We have

$$\begin{aligned} \mathbb{P}[\bar{S} \leq x \mid \mathbf{Z}_{-s}] &= \mathbb{P}[S_s \leq sx - (S_1 + \dots + S_{s-1}) \mid \mathbf{Z}_{-s}] = \mathbb{P}[Y_s \leq \ln[sx - (S_1 + \dots + S_{s-1})] \mid \mathbf{Z}_{-s}] \\ &= \mathbb{P}[Z_s \leq W(x)] = \Psi(W(x)), \end{aligned}$$

where

$$W(x) = \frac{\ln[sx - (S_1 + \dots + S_{s-1})] - Y_{s-1} - (t_s - t_{s-1})\mu}{\sqrt{t_s - t_{s-1}}\sigma}.$$

Taking the derivative with respect to  $x$  gives the unbiased density estimator

$$\frac{\partial}{\partial x} \mathbb{P}[\bar{S} \leq x \mid \mathbf{Z}_{-s}] = \psi(W(x))W'(x) = \frac{\psi(W(x))s}{[sx - (S_1 + \dots + S_{s-1})]\sqrt{t_s - t_{s-1}}\sigma}. \quad (22)$$

Unfortunately, this *sequential* CDE is ineffective, because hiding only the last value does not remove much information, so the conditional density turns out to have a large variance.

**Brownian bridge construction.** A more effective approach uses a Brownian bridge construction to sample the process, as follows (Glasserman 2004). The goal is to be able to hide one standard normal  $Z_s$  that has much more impact on the payoff  $X$  than in the previous sequential setting. To simplify the notation, we assume here that  $s$  is a power of 2, say  $s = 2^k$ . First sample  $Y_s = Y(t_s)$  directly from its unconditional normal distribution, using  $Z_s$ . This gives  $Y_s = t_s\mu + \sqrt{t_s}\sigma Z_s$ . Then, given  $Y_s = y_s$ , we sample  $Y_{s/2} = Y(t_{s/2})$  from its normal distribution conditional on  $Y_s = y_s$ , which is normal with mean  $y_s/2$  and variance  $(t_s - t_{s/2})(t_{s/2} - t_0)\sigma^2/(t_s - t_0)$ , then sample  $Y_{s/4}$  conditionally on  $Y_{s/2}$ , then  $Y_{3s/4}$  conditionally on  $(Y_{s/2}, Y_s)$ , and so on. For the CDE, we can hide again  $Z_s$ , but this time it means that we hide the impact of  $Z_s$  (or equivalently of  $Y_s$ ) on the entire trajectory. To do that, we can sample the intermediate values  $Y_1, \dots, Y_{s-1}$  conditional on  $Z_s = z_s = 0$ , which gives say  $Y_1^0, \dots, Y_{s-1}^0$ , and then write  $X$  as a function of  $z_s$  conditional on these values, that is, conditional on  $\mathbf{Z}_{-s} = (Z_1, \dots, Z_{s-1})$ . This gives the function  $X = \gamma(z_s)$ . By putting  $Y_s^0 = t_s\mu$ , we have  $Y_s = Y_s^0 + \sigma Z_s \sqrt{t_s}$  and  $Y_j = Y_j^0 + (t_j/t_s)\sigma Z_s \sqrt{t_s} = Y_j^0 + \sigma Z_s t_j t_s^{-1/2}$ . Then,

$$X = \bar{S} = \frac{1}{s} \sum_{j=1}^s e^{Y_j} = \frac{1}{s} \sum_{j=1}^s e^{Y_j^0 + Z_s \sigma t_j t_s^{-1/2}}.$$

This gives

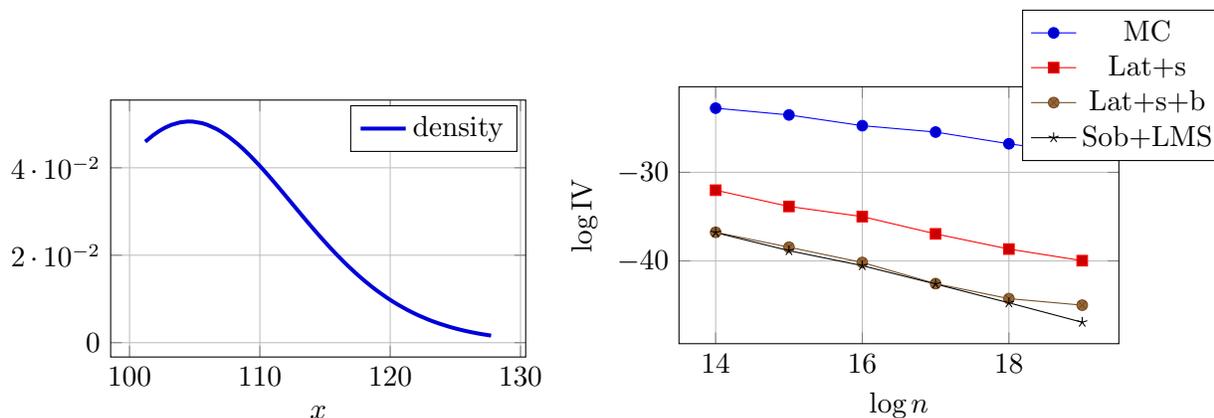
$$\gamma(z) = \frac{1}{s} \sum_{j=1}^s e^{Y_j^0 + z \sigma t_j t_s^{-1/2}}.$$

and

$$\gamma'(z) = \frac{1}{s} \sum_{j=1}^s e^{Y_j^0 + z \sigma t_j t_s^{-1/2}} \sigma t_j t_s^{-1/2}.$$

The CDE at  $x = \gamma(z)$  is then  $f(x \mid \mathbf{Z}_{-j}) = \phi(z)/\gamma'(z)$ . We call it the *bridge* CDE.

To compute this density at a specified  $x$ , we need to compute  $z = \gamma^{-1}(x)$ . Unfortunately, we have no explicit formula for  $\gamma^{-1}$  in this case, so it seems we have to compute a root of  $\gamma(z) - x = 0$  numerically. To evaluate the density at the  $n_e$  evaluation points  $e_1, \dots, e_{n_e}$  in  $(a, b)$ , we can proceed as follows. We first compute  $x_* = \gamma(0)$  and let  $j_*$  be the smallest  $j$  for which  $e_j \geq x_*$ . We compute  $z = w_{j_*}$  such that  $\gamma(w_{j_*}) = e_{j_*}$ . This can be done via Newton iteration,  $z_k = z_{k-1} - (\gamma(z_{k-1}) - e_{j_*})/\gamma'(z_{k-1})$ , starting with  $z_0 = 0$ . Then, for  $j = j_* + 1, \dots, n_e$ , we use again Newton iteration to find  $z = w_j$  such that  $\gamma(w_j) = e_j$ , starting at  $z_0 = w_{j-1}$ . We do the same to find  $z = w_j$  such that  $\gamma(w_j) = e_j$  for  $j = j_* - 1, \dots, 1$ , starting at  $z_0 = w_{j+1}$ . This provides the point  $w_j$  required to evaluate



**Figure 10** Estimated density (left) and log IV as a function of log  $n$  (left) for the Asian option.

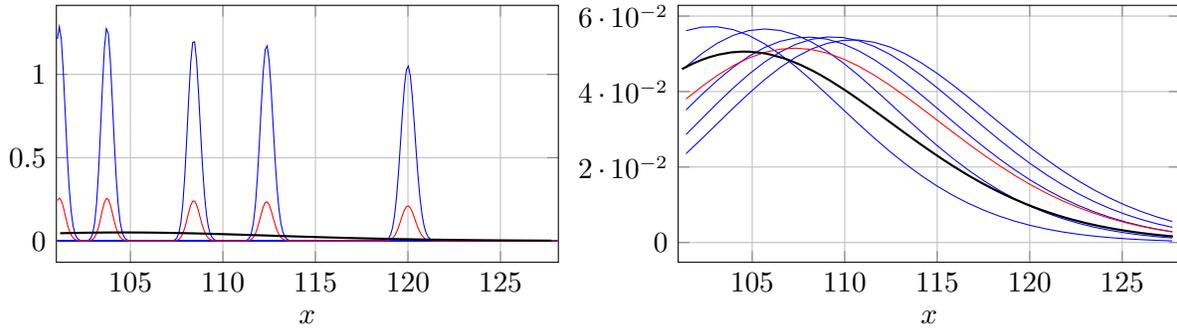
the conditional density at  $e_j$ , for each  $j$ . We have to repeat this procedure for each realization of  $\mathbf{Z}_{-j}$ , because the function  $\gamma$  depends on  $\mathbf{Z}_{-j}$ . This implies additional computations. However, the gain in accuracy can be quite significant.

**Table 10** Values of  $\hat{\nu}$  and e19 for the Asian option, with sequential and bridge CDE constructions.

		$\hat{\nu}$	e19
sequential KDE	MC	-0.78	-20.4
	Sob+LMS	-0.76	-20.6
sequential CDE	MC	-1.00	-19.9
	Lat+s	-1.07	-20.3
	Lat+s+b	-1.01	-20.1
	Sob+LMS	-1.00	-20.0
bridge CDE	MC	-1.04	-27.9
	Lat+s	-1.60	-40.0
	Lat+s+b	-1.74	-45.0
	Sob+LMS	-2.01	-46.9

For a numerical illustration, we take  $S_0 = 100$ ,  $r = 0.1$ ,  $\sigma = 0.12136$ ,  $\mu = r - \sigma^2/2$ ,  $K = 101$ , and  $X = \max(0, \bar{S} - K)$ . We estimate the density over  $[a, b] = [101, 128.13]$ . To approximate the root of  $\gamma(z) - x = 0$  for the bridge CDE, we used five Newton iterations. Doing more iterations made no significant difference. The results are in Table 10 and Figure 10. We find that RQMC with the bridge CDE performs extremely well. For example, for Sob+LMS, the MISE with  $n = 2^{19}$  is approximately  $2^{-46.9}$ , which is about  $2^{19}$  (half a million) times smaller than for the same CDE with MC, and it decreases as  $n^{-2}$ . To illustrate the differences in performance between the two different strategies of conditioning, we include in Figure 11 a plot of 5 single realizations of the sequential CDE (left) and of the bridge CDE (right), similar to what we did for the cantilever example. The blue lines depict these individual realizations, whereas the red line shows their average and the black line the true density. We see that, for the sequential CDE, what we hide does not have much

variance, which results in narrow high peaks of the estimator. In contrast, the realizations for the bridge CDE are much wider, smoother curves, which explains why RQMC works better for this strategy. With a KDE, the MISE with  $n = 2^{19}$  is about  $2^{21} \approx 2$  million times larger with the same Sobol' points and  $2^{26} \approx 67$  million times larger with MC. With the sequential CDE, RQMC is ineffective and the IV of the MC estimator is also quite large, as expected.



**Figure 11** Five realizations of the density estimator (blue), their average (red), and the true density (thick black) for the sequential CDE (left) and the bridge CDE (right), for the Asian option example.

#### 4.9. Estimating a quantile with a confidence interval

For  $0 < q < 1$ , the  $q$ -quantile of the distribution of  $X$  is defined as  $\xi_q = F^{-1}(q) = \inf\{x : F(x) \geq q\}$ . Given  $n$  i.i.d. observations of  $X$ , a standard (consistent) estimator of  $\xi_q$  is the  $q$ -quantile of the empirical distribution, defined as  $\hat{\xi}_{q,n} = X_{(\lceil nq \rceil)}$ , where  $X_{(1)}, \dots, X_{(n)}$  are the  $n$  observations sorted in increasing order (the order statistics). We assume that the density  $f(x)$  is positive and continuously differentiable in a neighborhood of  $\xi_q$ . Then we have the central limit theorem (CLT):

$$\sqrt{n}(\hat{\xi}_{q,n} - \xi_q)/\sigma_\xi \Rightarrow \mathcal{N}(0, 1) \quad \text{for } n \rightarrow \infty,$$

where  $\sigma_\xi^2 = q(1-q)/f^2(\xi_q)$  (Serfling 1980). This provides a way to compute a confidence interval on  $\xi_q$ , but requires the estimation of  $f(\xi_q)$ , which is generally difficult. Some approaches for doing this include finite differences with the empirical cdf, batching, and sectioning (Asmussen and Glynn 2007, Nakayama 2014a,b).

In our setting, one can do better by taking the  $q$ -quantile  $\hat{\xi}_{\text{cmc},q,n}$  of the conditional cdf

$$\hat{F}_{\text{cmc},n}(x) = \frac{1}{n} \sum_{i=1}^n F(x | \mathcal{G}^{(i)}).$$

That is,  $\hat{\xi}_{\text{cmc},q,n} = \inf\{x : \hat{F}_{\text{cmc},n}(x) \geq q\}$ . This idea was already suggested by Nakayama (2014b), who pointed out that this estimator obeys a CLT just like  $\hat{\xi}_{q,n}$ , but with the variance constant  $\sigma_\xi^2$

replaced by  $\sigma_{\text{cmc},\xi}^2 = \text{Var}[F(\xi_q | \mathcal{G})]/f^2(\xi_q) \leq \sigma_\xi^2$ . This is an improvement on the quantile estimator itself. Our CDE approach also provides an improved estimator of the density  $f(\xi_q)$  which appears in the variance expression. We estimate  $f(\xi_q)$  by  $\hat{f}_{\text{cde},n}(\hat{\xi}_{\text{cmc},q,n})$ . This provides a more accurate confidence interval of  $\xi_q$ .

Further improvements on the variances of both the quantile and density estimators can be obtained by using RQMC to generate the realizations  $\mathcal{G}^{(i)}$ . In particular, if  $\tilde{g}(\xi_q, \mathbf{u}) = F(\xi_q | \mathcal{G})$  is a sufficiently smooth function of  $\mathbf{u}$ ,  $\text{Var}[\hat{\xi}_{\text{cmc},q,n}]$  can converge at a faster rate than  $\mathcal{O}(n^{-1})$ . When using RQMC with  $n_r$  randomizations to estimate a quantile, the quantile estimator will be the empirical quantile of all the  $n_r \times n$  observations.

A related quantity is the *expected shortfall*, defined as  $c_q = \mathbb{E}[X | X > \xi_q] = \xi_q - \mathbb{E}[(\xi_q - X)^+]/q$  which is often estimated by its empirical version (Hong et al. 2014)

$$\hat{c}_{q,n} = \hat{\xi}_{q,n} - \frac{1}{nq} \sum_{i=1}^n (\hat{\xi}_{q,n} - X_i)^+.$$

This estimator obeys the CLT  $\sqrt{n}(\hat{c}_{q,n} - c_q)/\sigma_c \Rightarrow \mathcal{N}(0, 1)$  for  $n \rightarrow \infty$ , where  $\sigma_c^2 = \text{Var}[(\xi_q - X)^+]/q^2$ , if this variance is finite (Hong et al. 2014). By improving the quantile estimator, CDE+RQMC can also improve the expected shortfall estimator as well as the estimator of the variance constant  $\sigma_c^2$  and the quality of confidence intervals on  $c_q$ . We leave this as a topic for future work.

## 5. Conclusion

We have examined a novel approach for estimating the density of a random variable generated by simulation from a stochastic model, by conditioning. The resulting CDE is unbiased and its MISE converges at a faster rate than for other popular density estimators such as the KDE. We have also shown how to further reduce the IV, and even improve its convergence rate, by combining the CDE with RQMC sampling. Our numerical examples show that this combination can be very efficient. It sometimes reduces the MISE by factors over a million (cf. the cantilever example). Our CDE approach also outperforms the recently proposed GLR method, and CDE+RQMC outperforms both GLR+RQMC and KDE+RQMC, in all our examples.

Suggested future work includes experimenting this methodology on larger and more complicated stochastic models, designing and exploring different types of conditioning, and perhaps adapting the Monte Carlo sampling strategies to make the method more effective (e.g., by changing the way  $X$  is defined in terms the basic input random variates). Its application to quantile and expected shortfall estimation also deserves further study.

## Acknowledgments

This work has been supported by an IVADO Research Grant, an NSERC-Canada Discovery Grant, a Canada Research Chair, and an Inria International Chair, to P. L'Ecuyer. We acknowledge the financial support by the Ministerio de Economía y Competitividad (MINECO) of the Spanish Government through BCAM Severo Ochoa accreditation SEV-2017-0718. This work was supported by the BERC 2018e2021 Program and by ELKARTEK Programme, grants KK-2018/00054 and KK-2019/00068, funded by the Basque Government. The research is partially supported by the Basque Government under the grant “Artificial Intelligence” in BCAM number EXP. 2019/00432. This work has been possible thanks to the support of the computing infrastructure of the i2BASQUE academic network and the technical and human support provided by IZO-SGI SGIker of UPV/EHU and European funding (ERDF and ESF). We thank Julien Keutchayan for pointing out mistakes in a previous version.

## References

- Asmussen, S. 2018. Conditional Monte Carlo for sums, with applications to insurance and finance. *Annals of Actuarial Science* **12**(2) 455–478.
- Asmussen, S., P. W. Glynn. 2007. *Stochastic Simulation*. Springer-Verlag, New York.
- Avramidis, A. N., J. R. Wilson. 1998. Correlation-induction techniques for estimating quantiles in simulation experiments. *Operations Research* **46**(4) 574–591.
- Ben Abdellah, A., P. L'Ecuyer, A. Owen, F. Puchhammer. 2019. Density estimation by randomized quasi-Monte Carlo. Manuscript.
- Bingham, Derek. 2017. Virtual library of simulation experiments. URL <https://www.sfu.ca/~ssurjano/canti.html>.
- Botev, Z. I., P. L'Ecuyer, G. Rubino, R. Simard, B. Tuffin. 2013. Static network reliability estimation via generalized splitting. *INFORMS Journal on Computing* **25**(1) 56–71.
- Botev, Z. I., P. L'Ecuyer, B. Tuffin. 2016. Static network reliability estimation under the Marshall-Olkin copula. *ACM Transactions on Modeling and Computer Simulation* **26**(2) Article 14, 28 pages.
- Bratley, P., B. L. Fox, L. E. Schrage. 1987. *A Guide to Simulation*. 2nd ed. Springer-Verlag, New York, NY.
- Choquet, D., P. L'Ecuyer, C. Léger. 1999. Bootstrap confidence intervals for ratios of expectations. *ACM Transactions on Modeling and Computer Simulation* **9**(4) 326–348.
- Dick, J., F. Pillichshammer. 2010. *Digital Nets and Sequences: Discrepancy Theory and Quasi-Monte Carlo Integration*. Cambridge University Press, Cambridge, U.K.
- Dieudonné, J. 1969. *Foundations of Modern Analysis*. 2nd ed. Academic Press, New York.
- Efron, B., T. Hastie. 2016. *Computer Age Statistical Inference*. Cambridge University Press, New York.
- Fu, M., J.-Q. Hu. 1997. *Conditional Monte Carlo*. Kluwer Academic, Boston.
- Fu, Michael C., L. Jeff Hong, Jian-Qiang Hu. 2009. Conditional Monte Carlo estimation of quantile sensitivities. *Management Science* **55**(12) 2019–2027.

- Gertsbakh, I. B., Y. Shpungin. 2010. *Models of Network Reliability*. CRC Press, Boca Raton, FL.
- Glasserman, P. 2004. *Monte Carlo Methods in Financial Engineering*. Springer-Verlag, New York.
- Glynn, P. W. 1987. Likelihood ratio gradient estimation: an overview. *Proceedings of the 1987 Winter Simulation Conference*. IEEE Press, Piscataway, NJ, 366–375.
- He, Z., X. Wang. 2015. On the convergence rate of randomized quasi-monte carlo for discontinuous functions. *SIAM Journal on Numerical Analysis* **53**(5) 2488–2503.
- Hickernell, F. J. 2002. Obtaining  $O(N^{-2+\epsilon})$  convergence for lattice quadrature rules. K.-T. Fang, F. J. Hickernell, H. Niederreiter, eds., *Monte Carlo and Quasi-Monte Carlo Methods 2000*. Springer-Verlag, Berlin, 274–289.
- Higham, N. J. 2009. The scaling and squaring method for the matrix exponential revisited. *SIAM Review* **51**(4) 747–764.
- Hong, L. J., Z. Hu, G. Liu. 2014. Monte Carlo methods for value-at-risk and conditional value-at-risk: A review. *ACM Transactions on Modeling and Computer Simulation* **24**(4) Article 22.
- Laub, P. J., R. Salomone, Z. I. Botev. 2019. Monte Carlo estimation of the density of the sum of dependent random variables. *Mathematics and Computers in Simulation* **161** 23–31.
- Law, A. M. 2014. *Simulation Modeling and Analysis*. 5th ed. McGraw-Hill, New York.
- L'Ecuyer, P. 1990. A unified view of the IPA, SF, and LR gradient estimation techniques. *Management Science* **36**(11) 1364–1383.
- L'Ecuyer, P. 2009. Quasi-Monte Carlo methods with applications in finance. *Finance and Stochastics* **13**(3) 307–349.
- L'Ecuyer, P. 2016. SSJ: Stochastic simulation in Java. <http://simul.iro.umontreal.ca/ssj/>.
- L'Ecuyer, P. 2018. Randomized quasi-Monte Carlo: An introduction for practitioners. P. W. Glynn, A. B. Owen, eds., *Monte Carlo and Quasi-Monte Carlo Methods: MCQMC 2016*. Springer, Berlin, 29–52.
- L'Ecuyer, P., E. Buist. 2008. On the interaction between stratification and control variates, with illustrations in a call center simulation. *Journal of Simulation* **2**(1) 29–40.
- L'Ecuyer, P., C. Lemieux. 1999. Quasi-Monte Carlo via linear shift-register sequences. *Proceedings of the 1999 Winter Simulation Conference*. IEEE Press, 632–639.
- L'Ecuyer, P., C. Lemieux. 2000. Variance reduction via lattice rules. *Management Science* **46**(9) 1214–1235.
- L'Ecuyer, P., C. Lemieux. 2002. Recent advances in randomized quasi-Monte Carlo methods. M. Dror, P. L'Ecuyer, F. Szidarovszky, eds., *Modeling Uncertainty: An Examination of Stochastic Theory, Methods, and Applications*. Kluwer Academic, Boston, 419–474.
- L'Ecuyer, P., D. Munger. 2012. On figures of merit for randomly-shifted lattice rules. H. Woźniakowski, L. Plaskota, eds., *Monte Carlo and Quasi-Monte Carlo Methods 2010*. Springer-Verlag, Berlin, 133–159.

- L'Ecuyer, P., D. Munger. 2016. Algorithm 958: Lattice builder: A general software tool for constructing rank-1 lattice rules. *ACM Transactions on Mathematical Software* **42**(2) Article 15.
- L'Ecuyer, P., G. Perron. 1994. On the convergence rates of IPA and FDC derivative estimators. *Operations Research* **42**(4) 643–656.
- Lei, L., Y. Peng, M. C. Fu, J.-Q. Hu. 2018. Applications of generalized likelihood ratio method to distribution sensitivities and steady-state simulation. *Discrete Event Dynamic Systems* **28**(1) 109–125.
- Lemieux, C. 2009. *Monte Carlo and Quasi-Monte Carlo Sampling*. Springer-Verlag.
- Lemieux, C., P. L'Ecuyer. 2003. Randomized polynomial lattice rules for multivariate integration and simulation. *SIAM Journal on Scientific Computing* **24**(5) 1768–1789.
- Marion, P., M. Godin, P. L'Ecuyer. 2020. An algorithm to compute the  $t$ -value of a digital net and of its projections. *Journal of Computational and Applied Mathematics* **371**(June) 112669. To appear.
- Nakayama, M. K. 2014a. Confidence intervals for quantiles using sectioning when applying variance-reduction techniques. *ACM Transactions on Modeling and Computer Simulation* **24**(4) Article 9.
- Nakayama, M. K. 2014b. Quantile estimation when applying conditional Monte Carlo. *2014 International Conference on Simulation and Modeling Methodologies, Technologies, and Applications (SIMULTECH)*. IEEE, 280–285.
- Niederreiter, H. 1992. *Random Number Generation and Quasi-Monte Carlo Methods, SIAM CBMS-NSF Reg. Conf. Series in Applied Mathematics*, vol. 63. SIAM.
- Owen, A. B. 2003. Variance with alternative scramblings of digital nets. *ACM Transactions on Modeling and Computer Simulation* **13**(4) 363–378.
- Parzen, E. 1962. On estimation of a probability density function and mode. *Annals of Mathematical Statistics* **33**(3) 1065–1076.
- Peng, Y., M. C. Fu, P. W. Glynn, J.-Q. Hu. 2017. On the asymptotic analysis of quantile sensitivity estimation by Monte Carlo simulation. *Proceedings of the 2017 Winter Simulation Conference*. IEEE Press, Piscataway, NJ, 2336–2347.
- Peng, Y., M. C. Fu, B. Heidergott, H. Lam. 2020. Maximum likelihood estimation by Monte Carlo simulation: Towards data-driven stochastic modeling. *Operations Research* To appear.
- Peng, Y., M. C. Fu, J.-Q. Hu, B. Heidergott. 2018. A new unbiased stochastic derivative estimator for discontinuous sample performances with structural parameters. *Operations Research* **66**(2) 487–499.
- Schiels, M. D., J. Zhang. 2016. The generalization of Latin hypercube sampling. *Reliability Engineering and System Safety* **148** 96–108.
- Scott, David W. 2015. *Multivariate Density Estimation*. 2nd ed. Wiley.
- Serfling, R. J. 1980. *Approximation Theorems for Mathematical Statistics*. Wiley, New York, NY.
- Van der Vaart, A. W. 2000. *Asymptotic Statistics*. Cambridge University Press.