

Importance sampling and its optimality for stochastic simulation models*

Yen-Chi Chen and Youngjun Choe

*University of Washington
Department of Statistics
Seattle, WA 98195
e-mail: yenchic@uw.edu*

*University of Washington
Department of Industrial and Systems Engineering
Seattle, WA 98195
e-mail: ychoe@uw.edu*

Abstract: We consider the problem of estimating an expected outcome from a stochastic simulation model. Our goal is to develop a theoretical framework on importance sampling for such estimation. By investigating the variance of an importance sampling estimator, we propose a two-stage procedure that involves a regression stage and a sampling stage to construct the final estimator. We introduce a parametric and a nonparametric regression estimator in the first stage and study how the allocation between the two stages affects the performance of the final estimator. We analyze the variance reduction rates and derive oracle properties of both methods. We evaluate the empirical performances of the methods using two numerical examples and a case study on wind turbine reliability evaluation.

MSC 2010 subject classifications: Primary 62G20; secondary 62G86, 62H30.

Keywords and phrases: Nonparametric estimation, stochastic simulation model, oracle property, variance reduction, Monte Carlo.

Received October 2018.

1. Introduction

The 2011 Fisher lecture (Wu, 2015) features the landscape change in engineering, where computer simulation experiments are replacing physical experiments thanks to the advance of modeling and computing technologies. An insight from the lecture highlights that traditional principles for physical experiments do not necessarily apply to virtual experiments on a computer. The virtual environment calls for new modeling and analysis frameworks distinguished from those developed under the constraint of physical environment. In this context, this study considers a new problem that emerged along with the recent development of stochastic simulation-based engineering.

A concrete motivating problem of this study is estimating a system failure probability based on stochastic simulations (although our methods are more

*Both authors equally contributed to this work.

generally applicable to the estimation of any expected outcome, as detailed in the next section). A system configuration, X , is randomly sampled from a known density p and passed on to a stochastic simulation model. The simulation model, regarded as a stochastic black-box, produces V that follows an unknown distribution depending on X . When V is greater than a threshold, say ξ , the system fails. Thus, the goal is to estimate the probability $P(V \geq \xi)$ when X is from the density p . Specifically, our case study uses the simulation model that is designed to mimic the real system accurately. Thus, it takes roughly 1-min wall-clock time to simulate 10-min real operation of the system on a computer commonly available nowadays. The engineering goal is estimating the probability of system failure during 50-year operation, which is computationally challenging even with U.S. national labs' supercomputers (Manuel et al., 2013; Graf et al., 2016).

Such computational challenges are commonly observed in engineering simulations. Finite element simulations, which are used widely in various engineering applications, can take hours of computing time to obtain a single data point (e.g., Qian et al., 2006). Despite the computational expense, highly accurate simulations are cost-effective alternatives to physical experiments and used widely in industry (e.g., Ford Motor Company's crash simulation (Wang and Shan, 2007)) and in government (e.g., NASA's rocket booster simulation (Gramacy and Lee, 2012)).

The overarching goal of this study is to propose a framework on analyzing the problem of minimizing the necessary computational burden while maintaining the same level of estimation accuracy. A flip-side of the same problem is minimizing the estimation variance given fixed computational resource. Variance reduction techniques (VRTs) in the simulation literature aim to reduce the variance of estimator in simulation experiments. Traditional VRTs are well studied for the simulation model that outputs V given X in a deterministic fashion, also known as the deterministic simulation model (see Chapter 9 of Kroese et al. (2011) for survey of such VRTs), when the input X is sampled from a known distribution. For stochastic simulation models, if their underlying processes have known properties (e.g., Markovian), Glynn and Iglehart (1989) and Heidelberger (1995) provide VRTs. For black-box stochastic simulation models, few studies (e.g., Choe et al., 2015; Graf et al., 2017) consider VRTs. The research on VRTs for block-box stochastic simulations is still underdeveloped despite the rapid growth of such simulations being used in real systems, for example, chemical systems (Gillespie, 2001), biological systems (Henderson et al., 2012), and engineering systems (Ankenman et al., 2010; Picheny et al., 2013; Plumlee and Tuo, 2014).

Among VRTs, importance sampling (Kahn and Marshall, 1953) is known to be one of the most effective methods and has been used widely in various applications such as communication systems (Heidelberger, 1995; Bucklew, 2004), finance (Owen and Zhou, 2000; Glasserman and Li, 2005), insurance (Blanchet and Lam, 2011), and reliability engineering (Au and Beck, 1999; Lawrence et al., 2013; Choe et al., 2016) to name a few.

In the vast majority of literature, importance sampling takes a parametric

form tailored to a problem at hand for both deterministic simulation model (e.g. Lawrence et al., 2013) and stochastic counterpart (Choe et al., 2015). Nonparametric approaches are also proposed for deterministic simulation models (Zhang, 1996; Neddermeyer, 2009). To the best of our knowledge, no nonparametric approach is developed for stochastic simulation models. This study particularly considers the black-box stochastic simulation model whose output takes an unknown stochastic relationship with the input.

In this paper, we focus on the theoretical aspects of importance sampling with stochastic simulation models. We design two-stage sampling methods that may perform as good as the best sampling method (also known as the oracle property). The main contributions of this paper to the existing body of literature are as follows:

- We introduce an importance sampling approach to estimate the expectation of black-box stochastic simulation output and study the optimal importance sampler (Section 2).
- We design a two-stage procedure that uses a parametric or a nonparametric regression estimator to approximate the optimal importance sampler (Figure 1 and 2).
- We analyze the allocation of the resources in both stages (Theorem 3 and 7) and study the convergence of the two-stage procedure toward the oracle importance sampler (Corollary 4 and 8).
- We conduct an extensive numerical study to evaluate empirical performances of the two-stage importance samplers (Section 4.2.1 and 4.2.2).
- We apply our methods to a case study on wind turbine reliability evaluation (Section 4.3) to validate our results.

This paper is organized as follows. In Section 2 we formulate the stochastic simulation-based estimation problem and introduce a two-stage procedure to estimate the expected simulation output. In Section 3 we study the theoretical performance of the proposed procedure and derive the corresponding oracle properties. In Section 4 we evaluate the empirical performance of our approach using two numerical examples and a wind turbine simulator. We discuss our result in Section 5.

2. Importance sampling for the stochastic simulation model

A stochastic simulation model takes an input configuration value and then returns a random number representing the outcome of this simulation result. The input configuration determines the distribution of the (random) outcome. Thus, the outcome of a stochastic simulation model can be represented by a random variable V conditioned on the input configuration x and the CDF of V is

$$F_{V|\text{config}=x}(v|x) = P(V \leq v|\text{config} = x),$$

where $\text{config} = x$ denotes choosing the configuration to be x . For simplicity, we denote the random variable V conditioned on $\text{config} = x$ as $V(x)$.

In many scientific or engineering problems (e.g., Heidelberger, 1995; Au and Beck, 2003; Bucklew, 2004; Graf et al., 2017), we assume the nature generates the configuration from a known density p and we are interested in evaluating the quantity

$$\mathcal{E} = \mathbb{E}(g(V(X))) = \int \mathbb{E}(g(V)|\text{config} = x)p(x)dx, \quad (1)$$

where g is a known function.

Example 1. A common example for equation (1) is the case when $g(v) = v$, often considered in the literature on two-level nested simulation (Sun et al., 2011), where the outer level generates a scenario (or configuration) according to a known density p and conditioning on the scenario, the inner level simulates a random outcome whose mean is of interest. Applications include decision theory (Brennan et al., 2007), financial engineering (Staum, 2009), and queuing system (Sun et al., 2011).

Example 2. Another example takes $g(v) = 1(v \in S_\xi)$ for some set S_ξ parametrized by ξ . Specifically, Choe et al. (2015) considers a reliability evaluation problem where V stands for an instability measurement of a system that fails when V falls into $S_\xi = [\xi, \infty)$. The goal is to estimate the failure probability when the system is exposed to the nature. In the natural environment, the configuration behaves like a random variable from a density p .

To estimate \mathcal{E} , we can choose several configurations x_1, \dots, x_n and then run the simulation to obtain realizations $v_1 = V(x_1), \dots, v_n = V(x_n)$. However, generating V from a given configuration x is often computationally expensive for stochastic simulation models. Therefore, we would like to run the simulation as few as possible. To put this constraint into consideration, we assume that we run the simulation only n times but we are able to choose the configuration for each simulation. We choose n configurations and evaluate the corresponding value of V . Namely, we only have pairs $(x_1, V_1), \dots, (x_n, V_n)$, where each V_i is a realization of the random variable $V(x_i)$. Such a constraint on the number of simulations, n , is sometimes called a *computational budget*.

Under such situation, a natural question is: *How do we choose the configurations x_1, \dots, x_n ?* Here we use the idea from importance sampling – we choose x_1, \dots, x_n from a density function q . In other words, we first sample X_1, \dots, X_n from q and then use $x_i = X_i$ as the configuration to run the i -th stochastic simulation. The density q is called *sampling density*. Note that each configuration does not necessarily have to be from the same density function.

When we generate X_1, \dots, X_n from q and then obtain V_1, \dots, V_n accordingly, a simple estimator of \mathcal{E} is

$$\widehat{\mathcal{E}}_q = \frac{1}{n} \sum_{i=1}^n g(V_i) \frac{p(X_i)}{q(X_i)}. \quad (2)$$

It is easy to see that $\widehat{\mathcal{E}}_q$ is an unbiased estimator under the assumption that

$q(x) = 0$ implies $g(V(x))p(x) = 0$ for all x , i.e.,

$$\mathbb{E}(\widehat{\mathcal{E}}_q) = \mathcal{E}$$

when the support of q covers the support of $g(V(x))p(x)$. We call this type of estimator an *importance sampling estimator*. Throughout this paper, we will focus on importance sampling estimators.

Using an importance sampling estimator (2), a key problem we want to address is: *what will be the optimal sampling density q^* that minimizes the estimation error?* Because the estimator (2) is unbiased, we only need to find the minimal variance estimator. Thus, the above question is equivalent to: *what will be the optimal sampling density q^* that minimizes the variance of $\widehat{\mathcal{E}}_q$?* The following lemma provides an answer to the above questions:

Lemma 1. *Assume X_1, \dots, X_n are IID from a density function q . Let $r^\dagger(x) = \mathbb{E}(g(V(X))|X = x)$ and $r(x) = \mathbb{E}(g^2(V(X))|X = x)$. Then variance of $\widehat{\mathcal{E}}_q$ equals to*

$$\begin{aligned} \text{Var}(\widehat{\mathcal{E}}_q) &= \frac{1}{n} \left(\mathbb{E}_{X_i \sim q} \left(r(X_i) \frac{p^2(X_i)}{q^2(X_i)} \right) - \mathbb{E}^2(r^\dagger(X^*)) \right) \\ &\geq \frac{1}{n} \left(\mathbb{E}^2(\sqrt{r(X^*)}) - \mathbb{E}^2(r^\dagger(X^*)) \right) \\ &\equiv \frac{1}{n} V_{\min} \end{aligned} \tag{3}$$

where X^* is a random variable from the density p . The equality holds when we choose q to be $q(x) = q^*(x) \propto \sqrt{\mathbb{E}(g^2(V(X))|X = x)} \cdot p(x)$. Namely, the optimal sampling density is $q^*(x)$.

We call the quantity V_{\min} the *oracle variance* of the importance sampling. It is the minimal variance that an importance sampler can achieve. A widely studied special case in engineering is the deterministic simulation model where $\text{Var}(V|X = x) = 0$ for all x , which implies $V_{\min} = 0$ for any nonnegative function $g(v)$ (e.g., Kahn and Marshall, 1953; Au and Beck, 1999; Kroese et al., 2011). The density that leads to the oracle variance, q^* , is called the optimal sampling density. This density is a modification from the natural configuration density p ; q^* puts more weight on the regions with higher $r(x)$ (e.g., higher probability of system failure). As long as r and r^\dagger are uniformly bounded within the support of p , the variance is finite.

However, we cannot directly generate configurations from q^* because it involves the unknown quantity $r(x) = \mathbb{E}(g^2(V(X))|X = x)$. A remedy to this problem is to apply a *two-stage sampling*. In the first stage, we generate part of configurations and evaluate the corresponding values of $V(x)$. Using the sample in the first stage, we obtain a pilot estimator \widehat{r} of r . In the second stage, we generate configurations based on an estimate of q^* using the pilot estimator \widehat{r} and use the remaining computational budget to evaluate values of $V(x)$. Finally, we use samples from both stages to form the final estimator of \mathcal{E} .

Here is a useful insight in estimating $r(x)$. Let Y be a random variable such that $Y = g^2(V(X))$. Then $r(x) = \mathbb{E}(g^2(V(X))|X = x) = \mathbb{E}(Y|X = x)$. Thus, estimating $r(x)$ is equivalent to estimating the regression function with the observations $(X_1, Y_1 = g^2(V_1)), \dots, (X_n, Y_n = g^2(V_n))$ assuming we have n observations.

Thus the two-stage procedure can be summarized as follows. We first generate a size m sample

$$(X_1, V_1), \dots, (X_m, V_m)$$

where $X_i, i = 1, \dots, m$, are from an initial sampling density q_0 . Then we transform V_i into $Y_i = g^2(V_i)$, which leads to a sample

$$(X_1, Y_1), \dots, (X_m, Y_m).$$

Now we estimate the regression function $r(x)$ by a regression estimator $\hat{r}(x)$ and compute the corresponding estimator \hat{q}^* of the oracle sampling density q^* . Finally, we generate the remaining data points

$$(X_{m+1}, V_{m+1}), \dots, (X_n, V_n)$$

from \hat{q}^* and pool both samples together to form the final estimator of the quantity \mathcal{E} . Because \hat{q}^* will tend to be closer to q^* compared to the initial sampling density q_0 , estimating \mathcal{E} using the sample in the second stage is more efficient (lower variance).

The sample size m in the first stage is a crucial quantity in our analysis. The quantity m is called the *allocation*. When m is too small, the estimator of q^* is inaccurate, so that the overall estimation efficiency is suboptimal. When m is too large, we only have a small budget for the second stage so that the overall estimation efficiency is low as well. As a result, to balance the estimation accuracy of q^* and the size of efficient sample in the second stage, there will be an optimal value of m depending on the total sample size n . In what follows we propose two different models to estimate r and q^* and analyze the optimal value of the allocation m .

2.1. Parametric importance sampling

As in the regression analysis, a straightforward approach of estimating the regression function is to assume a parametric model and estimate the corresponding parameters. Namely, we assume $r(x) = r_\theta(x)$ for some $\theta \in \Theta$ and use the first part of the sample to estimate θ .

To estimate $r_\theta(x)$, we use a classical approach – *the least square method*:

$$\hat{\theta}_m = \operatorname{argmin}_{\theta \in \Theta} \sum_{i=1}^m \|Y_i - r_\theta(X_i)\|^2. \quad (4)$$

Then the estimator $\hat{r}(x) = r_{\hat{\theta}_m}(x)$. Note that one can also assume a parametric form for the distribution of $Y_i|X_i$ and then use a maximum likelihood estimator.




FIG 1. Parametric importance sampling for the stochastic simulation model.

Using the estimator $\hat{\theta}_m$, we can then estimate the regression function $r_{\hat{\theta}_m}$ and construct the estimated optimal sampling density

$$q_{\hat{\theta}_m}^*(x) \propto \sqrt{r_{\hat{\theta}_m}(x)} \cdot p(x).$$

For the remaining $(n - m)$ data points, we generate the configurations from $q_{\hat{\theta}_m}^*$, run the simulation, and estimate \mathcal{E} accordingly.

We summarize our Parametric Importance Sampling method in Figure 1. Later in Section 3.1 we will derive the variance of this approach and show that the optimal allocation is to choose $m = O(n^{\frac{2}{3}})$.

2.2. Nonparametric importance sampling

Now we consider estimating $r(x)$ nonparametrically. For simplicity, we use the kernel regression (Nadaraya, 1964; Watson, 1964). Note that other nonparametric regression approach, such as the local polynomial regression (Wasserman, 2006), also works. The kernel regression uses the estimator

$$\hat{r}_h(x) = \frac{\sum_{i=1}^m Y_i K\left(\frac{x-X_i}{h}\right)}{\sum_{i=1}^m K\left(\frac{x-X_i}{h}\right)}, \quad (6)$$

where K is a smooth function (known as the kernel function) such as a Gaussian, and $h > 0$ is the smoothing bandwidth. Similar to the parametric approach, we then use this estimator to construct an estimated optimal sampling density

$$\hat{q}_h^*(x) \propto \sqrt{\hat{r}_h(x)} \cdot p(x),$$




FIG 2. Nonparametric importance sampling for the stochastic simulation model.

generate the remaining data points from it, and construct the final estimator using the procedure described previously.

Figure 2 summarizes the procedure of nonparametric importance sampling. There are two tuning parameters we need to select: the smoothing bandwidth h and the allocation size m . The smoothing bandwidth can be chosen by either cross-validation or a reference rule. In Section 3.3, we will derive the optimal rate for the smoothing bandwidth $h = O\left(\left(\frac{\log m}{m}\right)^{\frac{1}{d+4}}\right)$ and the optimal allocation rate $m = O\left(\left(\frac{n}{\log n}\right)^{\frac{d+4}{d+6}}\right)$.

Remark. We may rewrite $\hat{\mathcal{E}}_h = \hat{\mathcal{E}}_1 + \hat{\mathcal{E}}_2$, where $\hat{\mathcal{E}}_1 = \frac{\sum_{i=1}^m g(V_i) \frac{p(X_i)}{q_0(X_i)}}{n}$ and $\hat{\mathcal{E}}_2 = \frac{\sum_{i=m+1}^n g(V_i) \frac{p(X_i)}{\hat{q}_h^*(X_i)}}{n}$ are the estimator constructed from the two samples. If we know the variance of $\hat{\mathcal{E}}_1$ and $\hat{\mathcal{E}}_2$, then we can further reduce the variance of the final estimator by doing a weighted combination $\hat{\mathcal{E}}_h^\alpha = \alpha \hat{\mathcal{E}}_1 + (1 - \alpha) \hat{\mathcal{E}}_2$, with $\alpha = \frac{\text{Var}(\hat{\mathcal{E}}_2)}{\text{Var}(\hat{\mathcal{E}}_1) + \text{Var}(\hat{\mathcal{E}}_2)}$. One can easily show that $\text{Var}(\hat{\mathcal{E}}_h^\alpha) < \text{Var}(\hat{\mathcal{E}}_h)$ and $\mathbb{E}(\hat{\mathcal{E}}_h^\alpha) = \mathcal{E}$ so $\hat{\mathcal{E}}_h^\alpha$ is a better estimator.

3. Theoretical analysis

Throughout our analysis, we assume that the natural configuration density p has a compact support $\mathbb{K} \subset \mathbb{R}^d$ and the support of the initial sampling density q_0 contains \mathbb{K} .

3.1. Parametric importance sampling

Assumptions.

- (P1) There exists a unique $\theta_0 \in \Theta$ such that $r(x) = r_{\theta_0}(x)$ and $\sup_{x \in \mathbb{K}} \text{Var}(Y_1 - r_{\theta_0}(X_1) | X_1 = x) \leq \sigma_{\max}^2 < \infty$. The support of $r_{\theta_0}(x)$ contains the support of $p(x)$ for every $\theta \in \Theta$ and $r(x) > 0$ for all $x \in \mathbb{K}$. Also, the PDF $q_0(x)$ exists and has non-zero variance.
- (P2) Let $\ell(\theta) = \mathbb{E}\|Y_1 - r_{\theta}(X_1)\|^2$. The Hessian matrix $H(\theta) = \nabla_{\theta} \nabla_{\theta} \ell(\theta)$ is positive definite at $\theta \in B(\theta_0, R_0)$ for some $R_0 < \infty$ and θ_0 is the one in (P1). Note that $B(x, r)$ is a ball centered at x with radius r .
- (P3) There exists a positive $L_0 < \infty$ such that for any $\theta_1, \theta_2 \in B(\theta_0, R_0)$,

$$\sup_{x \in \mathbb{K}} |r_{\theta_1}(x) - r_{\theta_2}(x)| \leq L_0 \cdot \|\theta_1 - \theta_2\|,$$

where θ_0, R_0 are defined in (P2).

(P1) means that the model is correctly specified – the regression function can be parametrized in the parametric model we consider. (P2) is a common assumption in the M-estimation theory (van der Vaart and Wellner, 1996; van der Vaart, 2000) to derive the convergence rate. The extra assumption (P3) is a mild assumption that converts the convergence rate of parameter estimation to the convergence rate of function estimation. As long as $r_{\theta}(x)$ is smooth within an open set around θ_0 , (P3) holds.

The following theorem describes the estimation error when the parametric family contains the true regression function.

Theorem 2. Assume (P1–3). The error rate for the estimator $r_{\hat{\theta}_m}(x)$ is

$$\sup_{x \in \mathbb{K}} \|r_{\hat{\theta}_m}(x) - r(x)\| = O_P \left(\sqrt{\frac{1}{m}} \right).$$

Theorem 2 presents the error rate for estimating $r(x)$ when the model is correct. Based on this error rate, we can further derive the variance of the parametric importance sampler in Figure 1.

Theorem 3. Assume (P1–3). Let $V_{q_0} = \mathbb{E}_{X_i \sim q_0} \left(r(X_i) \frac{p^2(X_i)}{q_0^2(X_i)} \right) - \mathbb{E}^2 \left(\sqrt{r(X^*)} \right)$ be the excess variance from using q_0 compared to q^* . The variance of the estimator $\hat{\mathcal{E}}_{\hat{\theta}_m}$ is

$$\text{Var} \left(\hat{\mathcal{E}}_{\hat{\theta}_m} \right) = \frac{1}{n} V_{\min} + \frac{1}{n^2} \left(m \cdot V_{q_0} + (n - m) \cdot O \left(\sqrt{\frac{1}{m}} \right) \right).$$

Theorem 3 has three components. The first component V_{\min} is the oracle variance we have mentioned previously. It is the minimal variance that can be achieved by an importance sampling estimator. The second component $\frac{1}{n^2} \cdot m \cdot V_{q_0}$ is the excess variance due to the initial sampling density. The third component

$\frac{1}{n^2} \cdot (n - m) \cdot O\left(\sqrt{\frac{1}{m}}\right)$ is the excess variance due to the error of the estimator $r_{\hat{\theta}_m}(x)$.

By optimizing m with respect to the second and third components, we obtain the optimal rate of m as a function of sample size n :

$$\begin{aligned} m \cdot V_{q_0} &= (n - m) \cdot O\left(\sqrt{\frac{1}{m}}\right) \\ \implies m^{\frac{3}{2}} &\asymp n \\ \implies m &\asymp n^{\frac{2}{3}}, \end{aligned}$$

where the notation \asymp means that the two quantities will be of the same order, i.e., $a_n \asymp b_n \Leftrightarrow \lim_{n \rightarrow \infty} \frac{a_n}{b_n} \in (0, \infty)$. Thus, the optimal allocation is to choose $m \asymp n^{\frac{2}{3}}$, which leads to the following:

Corollary 4. *Assume (P1-3). When $m \asymp n^{\frac{2}{3}}$, the variance of the estimator $\hat{\mathcal{E}}_{\hat{\theta}_m}$ is*

$$\text{Var}\left(\hat{\mathcal{E}}_{\hat{\theta}_m}\right) = \frac{1}{n} V_{\min} \left(1 + O\left(n^{-\frac{1}{3}}\right)\right).$$

That is, if the model is correctly specified, the excess variance shrinks at rate $O\left(n^{-\frac{1}{3}}\right)$ under the optimal allocation.

Note that if we add conditions so that

$$\frac{q_{\hat{\theta}_m}^*(x) - q^*(x)}{p(x)} = (\hat{\theta}_m - \theta^*) \cdot C(x) + o_P(\|\hat{\theta}_m - \theta^*\|^2)$$

for some uniformly bounded function $C(x)$, the variance in Theorem 3 can be improved in the sense that the final term will be $(n - m)O\left(\frac{1}{m}\right)$, which leads to the optimal allocation $m \asymp n^{\frac{1}{2}}$. This allocation rate is similar to Li et al. (2013) although we are considering different scenarios. Li et al. (2013) considers mixture importance sampling for a deterministic (as opposed to stochastic) simulation model and study the optimal allocation, where the optimal mixing weight (of each component) is unknown and has to be estimated using a pilot sample. In our case, the parametric model is to approximate the regression function implied by a stochastic simulation model.

The key assumption of the parametric method is (P1): the actual $r(x)$ belongs to the parametric family. However, if this assumption is violated, then the excess variance in the parametric method will never shrink to 0.

Theorem 5. *Assume (P2-3). If $r(x) \neq r_{\theta}(x)$ for all $\theta \in \Theta$, the variance of the parametric estimator*

$$\text{Var}\left(\hat{\mathcal{E}}_{\hat{\theta}_m}\right) \geq \frac{1}{n} V_{\min} + \frac{1}{n} V_{\theta^*}$$

where

$$V_{\theta^*} = \inf_{\theta \in \Theta} \mathbb{E} \left(r(X_{\theta}) \frac{p^2(X_{\theta})}{q_{\theta}^2(X_{\theta})} \right) - \mathbb{E}^2 \left(\sqrt{r(X^*)} \right) > 0,$$

$$X_\theta \sim q_\theta(x) \propto \sqrt{r_\theta(x)} \cdot p(x).$$

The proof of this theorem is trivial, so we omit it. Theorem 5 proves that when the model is incorrectly specified, there is an additional variance V_{θ^*} that never disappears. Thus, the variance of the parametric importance sampler will not converge to the optimal variance. Later we will see that this implies that the parametric importance sampler does not have the oracle inequalities when the model is incorrectly specified.

3.2. Nonparametric importance sampling

In this section, we study the properties of the nonparametric importance sampler in Figure 2. Similarly as the parametric importance sampler, we first derive the convergence rate of estimating $r(x)$, then derive a variance decomposition for $\text{Var}(\widehat{\mathcal{E}}_h)$, and finally study the optimal allocation.

Assumptions.

- (N1) $\sup_{x \in \mathbb{K}} \text{Var}(Y_1 - r(X_1) | X_1 = x) \leq \sigma_{\max}^2 < \infty$ and $r(x) > 0$ for all $x \in \mathbb{K}$.
 (N2) For all x , the function $r(x)$ has bounded second derivative and $q_0(x)$ has bounded first derivative and $\sup_{x \in \mathbb{K}} q_0(x) \geq q_{\min} > 0$.
 (K1) The kernel function $K(x)$ is symmetric and

$$\int K(x) dx = 1, \quad \int \|x\|^2 K(x) dx < \infty, \quad \int K^2(x) dx < \infty.$$

- (K2) The collection of functions

$$\mathcal{K} = \left\{ y \mapsto K\left(\frac{x-y}{h}\right) : x \in \mathbb{K}, h > 0 \right\},$$

is a VC-type class. i.e. there exists constants A, v and a constant envelope b_0 such that

$$\sup_Q N(\mathcal{K}, \mathcal{L}^2(Q), b_0 \epsilon) \leq \left(\frac{A}{\epsilon}\right)^v, \quad (8)$$

where $N(T, d_T, \epsilon)$ is the ϵ -covering number for a semi-metric set T with metric d_T and $\mathcal{L}^2(Q)$ is the L_2 norm with respect to the probability measure Q .

(N1) and (N2) are common assumptions for nonparametric regression; see, e.g., Wasserman (2006) and Györfi et al. (2006). (K1) is a standard condition on kernel function (Wasserman, 2006; Scott, 2015). (K2) regularizes the complexity of kernel functions so that we have a uniform bound on the stochastic variation. This assumption was first proposed in Giné and Guillou (2002) and Einmahl and Mason (2005) and later was used in various studies such as Genovese et al. (2014); Chen et al. (2015b, 2017).

Based on the above assumptions, the uniform convergence rate of the kernel estimator $\widehat{r}_h(x)$ is given by the following.

Theorem 6. Assume (N1-2), (K1-2). The error rate of the kernel estimator $\widehat{r}_h(x)$ is

$$\sup_{x \in \mathbb{K}} \|\widehat{r}_h(x) - r(x)\| = O(h^2) + O_P \left(\sqrt{\frac{\log m}{mh^d}} \right).$$

The error in Theorem 6 can be decomposed into two parts: the bias part $O(h^2)$ and the stochastic variation $O_P \left(\sqrt{\frac{\log m}{mh^d}} \right)$ (which is related to the variance). In many nonparametric studies, similar bounds appear for density estimation; see, e.g., Giné and Guillou (2002); Einmahl and Mason (2005); Genovese et al. (2014); Chen et al. (2015b, 2016).

By Theorem 6, the optimal bandwidth $h^* \asymp \left(\frac{\log m}{m} \right)^{\frac{1}{d+4}}$ leads to the optimal error rate

$$\widehat{r}_{h^*}(x) - r(x) = O_P \left(\left(\frac{\log m}{m} \right)^{\frac{2}{d+4}} \right). \tag{9}$$

Under such an optimal error rate, we again obtain the variance decomposition for the nonparametric importance sampler.

Theorem 7. Assume (N1-2), (K1-2). Let $V_{q_0} = \mathbb{E}_{X_i \sim q_0} \left(r(X_i) \frac{p^2(X_i)}{q_0^2(X_i)} \right) - \mathbb{E}^2 \left(\sqrt{r(X^*)} \right)$ be the excess variance from using q_0 compared to q^* . The variance of the estimator $\widehat{\mathcal{E}}_{h^*}$ under the optimal smoothing bandwidth is

$$\text{Var} \left(\widehat{\mathcal{E}}_{h^*} \right) = \frac{1}{n} V_{\min} + \frac{1}{n^2} \left(m \cdot V_{q_0} + (n - m) \cdot O \left(\left(\frac{\log m}{m} \right)^{\frac{2}{d+4}} \right) \right).$$

Similar to Theorem 3, the variance in Theorem 7 has three components: the oracle variance V_{\min} , the excess variance due to the initial sampling density $\frac{1}{n^2} \cdot m \cdot V_{q_0}$, and the excess variance from the estimator $\widehat{r}_{h^*}(x)$.

To obtain the rate of the optimal allocation, we equate the two excess variances:

$$\begin{aligned} m \cdot V_{q_0} &= (n - m) \cdot O \left(\left(\frac{\log m}{m} \right)^{\frac{2}{d+4}} \right) \\ \implies m \cdot \left(\frac{m}{\log m} \right)^{\frac{2}{d+4}} &\asymp n \\ \implies m &\asymp \left(\frac{n}{\log n} \right)^{\frac{d+4}{d+6}} \quad (\text{ignoring the } \log \log n \text{ and multi-logarithm terms}). \end{aligned}$$

This choice of m yields the following variance reduction rate.

Corollary 8. Assume (N1-2), (K1-2). When $m \asymp \left(\frac{n}{\log n}\right)^{\frac{d+4}{d+6}}$, and $h^* \asymp \left(\frac{\log m}{m}\right)^{\frac{1}{d+4}}$, the variance of the estimator $\widehat{\mathcal{E}}_{h^*}$ is

$$\begin{aligned} \text{Var}\left(\widehat{\mathcal{E}}_{h^*}\right) &= \frac{1}{n}V_{\min} + O\left(\frac{1}{n^2} \cdot n \cdot \log^{\frac{2}{d+4}} n \cdot \left(\frac{n}{\log n}\right)^{\frac{d+4}{d+6} \times \frac{-2}{d+4}}\right) \\ &= \frac{1}{n}V_{\min} \left(1 + O\left(\left(\frac{\log^{(4d+20)/(d+4)} n}{n}\right)^{\frac{2}{d+6}}\right)\right) \\ &= \frac{1}{n}V_{\min} \left(1 + O\left(\left(\frac{\log^5 n}{n}\right)^{\frac{2}{d+6}}\right)\right) \end{aligned}$$

Note that in the last equality in Corollary 8, we use the fact that $a_n = O(\log^{(4d+20)/(d+4)} n)$ implies $a_n = O(\log^5 n)$ to simplify the expression. Corollary 8 shows that under the optimal allocation, the excess variance in the nonparametric importance sampler shrinks at rate $O\left(\left(\frac{\log^5 n}{n}\right)^{\frac{2}{d+6}}\right)$. When the dimension is small, say $d = 1$ or $d = 2$, the nonparametric method has an excess variance at rate $O\left(\left(\frac{\log^5 n}{n}\right)^{\frac{2}{7}}\right)$ and $O\left(\left(\frac{\log^5 n}{n}\right)^{\frac{1}{4}}\right)$, which are just slightly slower than the rate of the parametric importance sampler under correct model (the rate is $O\left(n^{-\frac{1}{3}}\right)$ by Corollary 4).

Although the parametric method enjoys a fast convergence rate, it depends on a very restrictive assumption: the model has to be correctly specified. This assumption is generally not true in most applications. Thus, even if the nonparametric importance sampler has a slower variance reduction rate, the nonparametric approach still has its own merit in applicability.

There is a limitation in the nonparametric approach – the curse of dimensionality. When d is not small, the convergence rate is very slow. So the nonparametric approach may not be applied to scenarios with more than 4 variables. In this situation, even if the parametric model does not converge to the optimal sampler, it may still be a preferred approach. The parametric model will converge to the best possible result under the mis-specified model as long as the conditions of Theorem 2 hold.

Remark 1. Note that the nonparametric rate can be improved if the regression function is very smooth and we use a higher order kernel (Wasserman, 2006). When the regression function is in β -Hölder class with $\beta > 2$, we can boost convergence rate in Theorem 7 to $O\left(\left(\frac{\log m}{m}\right)^{\frac{\beta}{d+2\beta}}\right)$ and under the optimal allocation, the variance of nonparametric importance sampler will be

$$\frac{1}{n}V_{\min} \left(1 + O\left(\left(\frac{\log^{d+2.5\beta} n}{n}\right)^{\frac{\beta}{d+3\beta}}\right)\right).$$

When $\beta \rightarrow \infty$, the variance of nonparametric importance sampler becomes $\frac{1}{n}V_{\min} (1 + O(n^{-1/3}))$ (ignoring the $\log n$ term), which recovers the parametric rate in Corollary 4.

3.3. Oracle properties

In Lemma 1, we see that the oracle sampler $\widehat{\mathcal{E}}_{q^*}$ has the minimal variance. Let

$$\mathcal{G} = \left\{ f : \int f(x)dx = 1, f(x) = 0 \Rightarrow g(V(x))p(x) = 0 \forall x \right\}$$

be the collection of density functions that leads to an unbiased importance sampler and let

$$\Xi = \left\{ \widehat{\mathcal{E}}_q : q \in \mathcal{G} \right\}$$

be the collection of all possible unbiased importance samplers. Because all importance samplers from \mathcal{G} is unbiased, any estimator $\widehat{\mathcal{E}} \in \Xi$ satisfies

$$\mathbb{E} \left\| \widehat{\mathcal{E}} - \mathcal{E} \right\|^2 = \text{Var} \left(\widehat{\mathcal{E}} \right) \tag{10}$$

so the oracle sampler $\widehat{\mathcal{E}}_{q^*}$ satisfies

$$\mathbb{E} \left\| \widehat{\mathcal{E}}_{q^*} - \mathcal{E} \right\|^2 = \text{Var} \left(\widehat{\mathcal{E}}_{q^*} \right) = \inf_{\widehat{\mathcal{E}} \in \Xi} \text{Var} \left(\widehat{\mathcal{E}} \right) = \inf_{\widehat{\mathcal{E}} \in \Xi} \mathbb{E} \left\| \widehat{\mathcal{E}} - \mathcal{E} \right\|^2. \tag{11}$$

Because of equation (11), $\widehat{\mathcal{E}}_{q^*}$ is called the *oracle for Ξ with respect to the mean square error* in nonparametric theory; see, e.g., page 60–61 in Tsybakov (2009).

We say $\widehat{\mathcal{E}}^\dagger \in \Xi$ satisfies the *oracle inequalities* if

$$\frac{\mathbb{E} \left\| \widehat{\mathcal{E}}^\dagger - \mathcal{E} \right\|^2}{\inf_{\widehat{\mathcal{E}} \in \Xi} \mathbb{E} \left\| \widehat{\mathcal{E}} - \mathcal{E} \right\|^2} = \frac{\mathbb{E} \left\| \widehat{\mathcal{E}}^\dagger - \mathcal{E} \right\|^2}{\mathbb{E} \left\| \widehat{\mathcal{E}}_{q^*} - \mathcal{E} \right\|^2} = 1 + o(1). \tag{12}$$

Note that the estimators in the above expressions are all based on a size n sample and we do not include the subscript n to abbreviation. In nonparametric theory, an estimator with the oracle inequalities implies that the estimator is asymptotically as good as the optimal estimator.

The crude Monte Carlo sampler $\widehat{\mathcal{E}}_p$ (which samples only from the natural configuration density p (Kroese et al., 2011)) obviously does not satisfy the oracle inequalities because

$$\frac{\mathbb{E} \left\| \widehat{\mathcal{E}}_p - \mathcal{E} \right\|^2}{\inf_{\widehat{\mathcal{E}} \in \Xi} \mathbb{E} \left\| \widehat{\mathcal{E}} - \mathcal{E} \right\|^2} = \frac{\text{Var} \left(\widehat{\mathcal{E}}_p \right)}{\text{Var} \left(\widehat{\mathcal{E}}_{q^*} \right)} = \frac{V_{\min} + V_p}{V_{\min}} = 1 + \frac{V_p}{V_{\min}} > 1,$$

where $V_p = \mathbb{E}(r(X^*)) - \mathbb{E}^2 \left(\sqrt{r(X^*)} \right) > 0$.

The parametric importance sampler $\widehat{\mathcal{E}}_{\widehat{\theta}_m}$ satisfies the oracle inequalities when the model is correctly specified (i.e. $r(x) = r_\theta(x)$ for some $\theta \in \Theta$). To see this, recall Corollary 4 and equation (10):

$$\frac{\mathbb{E} \left\| \widehat{\mathcal{E}}_{\widehat{\theta}_m} - \mathcal{E} \right\|^2}{\inf_{\widehat{\mathcal{E}} \in \Xi} \mathbb{E} \left\| \widehat{\mathcal{E}} - \mathcal{E} \right\|^2} = \frac{\text{Var} \left(\widehat{\mathcal{E}}_{\widehat{\theta}_m} \right)}{\text{Var} \left(\widehat{\mathcal{E}}_{q^*} \right)} = 1 + O \left(n^{-\frac{1}{3}} \right) = 1 + o(1).$$

However, when the model is incorrect, Theorem 5 proves that $\widehat{\mathcal{E}}_{\widehat{\theta}_m}$ does not have the oracle inequalities:

$$\frac{\mathbb{E} \left\| \widehat{\mathcal{E}}_{\widehat{\theta}_m} - \mathcal{E} \right\|^2}{\inf_{\widehat{\mathcal{E}} \in \Xi} \mathbb{E} \left\| \widehat{\mathcal{E}} - \mathcal{E} \right\|^2} = \frac{\text{Var} \left(\widehat{\mathcal{E}}_{\widehat{\theta}_m} \right)}{\text{Var} \left(\widehat{\mathcal{E}}_{q^*} \right)} = 1 + \frac{V_{\theta^*}}{V_{\min}} > 1.$$

The nonparametric importance sampler has a good advantage that it satisfies the oracle inequalities in most cases. By Corollary 8 and equation (10),

$$\frac{\mathbb{E} \left\| \widehat{\mathcal{E}}_{h^*} - \mathcal{E} \right\|^2}{\inf_{\widehat{\mathcal{E}} \in \Xi} \mathbb{E} \left\| \widehat{\mathcal{E}} - \mathcal{E} \right\|^2} = \frac{\text{Var} \left(\widehat{\mathcal{E}}_{h^*} \right)}{\text{Var} \left(\widehat{\mathcal{E}}_{q^*} \right)} = 1 + O \left(\left(\frac{\log n}{n} \right)^{\frac{2}{d+6}} \right) = 1 + o(1).$$

Thus, without any further information about the structure of $r(x)$, we recommend to use the nonparametric importance sampler since it behaves asymptotically as good as the oracle (optimal) importance sampler.

Remark 2. How we obtain the oracle property is very different from the classical approach. Many estimators with oracle properties are constructed by minimizing an estimated risk (Tsybakov, 2009). That is, for a collection of estimators, the risk of each of them is estimated and the one that minimizes the (estimated) risk is chosen. When the risk is consistently estimated uniformly for all estimators, this procedure leads to an estimator with the oracle property. However, in our case, we do not consider any risk estimator nor do we choose an estimator from many possible candidates, but we still obtain the oracle property.

4. Empirical analysis

To evaluate the empirical performances of the importance samplers, this section presents an implementation guideline, a numerical study, and a case study.

4.1. Implementation guideline

To implement parametric or nonparametric importance sampling, we can follow the procedure in Figure 1 or Figure 2, respectively. In practice, n is typically determined based on the available computational budget. We can choose m according to the optimal allocation rate, $m \asymp n^{\frac{2}{3}}$ in Corollary 4 for parametric

importance sampler, or $m \asymp \left(\frac{n}{\log n}\right)^{\frac{d+4}{d+6}}$ in Corollary 8 for nonparametric importance sampler. In the range of experiments we present below, any choice of multiplicative constant between two and six results in similar empirical performances of the importance samplers.

In the first stage, we can simply choose the natural configuration density p as the initial sampling density q_0 to maintain the same error rate as the crude Monte Carlo sampler. In the second stage, once we build a regression model $\hat{r}(x)$ for the unknown conditional expectation $r(x) = \mathbb{E}(g^2(V(X))|X = x)$, we can exactly sample from $q(x) \propto \sqrt{\hat{r}(x)} \cdot p(x)$ using the acceptance-rejection method (Kroese et al., 2011, p.59) by using $p(x)$ as the proposal (or envelope) density and finding an upper bound on $\sqrt{\hat{r}(x)}$. The upper bound can be computed numerically or even known from physical/engineering knowledge of the system (e.g., the simulation output V can take up to a certain value due to a physical/engineering limit, or $r(x)$ is a probability). As an alternative to the acceptance-rejection method, Markov chain Monte Carlo methods can be used as well.

Also, for an importance sampling estimator (e.g., in (5) or (7)), the normalization constant of $q(x)$ can be calculated to a desired level of accuracy, independent of n , by using a numerical integration such as quadrature for low-dimensional x and Monte Carlo integration for high-dimensional x . Thus, we can ensure that estimation of the normalization constant will not affect the asymptotic optimality and empirical performance of the importance sampling estimator. If desired, one can avoid calculating the normalization constant by using a self-normalized importance sampling estimator at the expense of bias. More in-depth discussion of this trade-off is made by Owen (2018).

Note that in practice, the computational costs of the acceptance-rejection method and the numerical integration are negligible (a matter of hours, if not minutes) compared to running high-fidelity simulation models such as those discussed in Section 1 and the simulation model in our case study. For example, a simulation experiment can take days (if not weeks) even with a supercomputer (Manuel et al., 2013; Graf et al., 2016).

4.2. Numerical study

Our numerical study considers two examples, one with normal distributions for X and $V|X$ and the other with exponential distributions for X and $V|X$. Motivated by our case study, we estimate the probability $\mathcal{E} = P(V > \xi) = \mathbb{E}(g(V(X)))$ for $g(V) = 1(V > \xi)$ and a pre-specified $\xi > 0$. As a baseline, we set ξ such that $P(V > \xi)$ is equal to 0.5 (unless specified otherwise) regardless of the input configuration dimension d because it is known that the performance of importance sampler often depends on the probability being estimated (Heidelberger, 1995; Kroese et al., 2011).

We vary the total sample size $n = 1000, 2000, 4000, 8000$ and the input configuration dimension $d = 1, 2, 4$ to see their impacts on the mean squared error

(MSE)

$$\text{MSE} = \frac{1}{n_{MC}} \sum_{i=1}^{n_{MC}} \left\| \widehat{\mathcal{E}}_i - \mathcal{E} \right\|^2,$$

where $\widehat{\mathcal{E}}_i$ is an estimate of the i th replication and the total number of Monte Carlo replications, n_{MC} , is set as 10,000 to obtain reliable results. We use high-performance computing (Lenovo NextScale E5-2680 v4, total 112 cores with 1TB RAM) for our simulation experiments, and they take several weeks in total. The R scripts for the experiments are available as a supplementary material.

We consider two parametric importance samplers, one with a correct model of

$$r(x) = \mathbb{E}(g^2(V(X)) | X = x) = \mathbb{E}(1(V(X) > \xi) | X = x) = P(V > \xi | X = x)$$

and the other with an incorrect model, and a nonparametric importance sampler. To build the parametric models of $r(x)$, we use the sample of size $m = \lceil 2n^{\frac{2}{3}} \rceil$. For the nonparametric model, we use $m = \left\lceil 6 \left(\frac{n}{\log n} \right)^{\frac{d+4}{d+6}} \right\rceil$.

To sample from the importance sampling density $q(x) \propto \sqrt{\widehat{r}(x)} \cdot p(x)$ using the acceptance-rejection method, we use $p(x)$ as the envelope density because $p(x) \geq \sqrt{\widehat{r}(x)}p(x)$ in the examples: We sample x from $p(x)$ and accept x with the probability $\sqrt{\widehat{r}(x)}$. To compute the normalizing constant of $q(x)$, we use Monte Carlo integration. Since we know the true $r(x)$, which is unknown in practice, in the examples, we calculate the true $\mathcal{E} = P(V > \xi)$ and V_{\min} using Monte Carlo integration and use them to calculate MSEs and demonstrate how empirical results conform to the theoretical predictions made in Section 3.

4.2.1. Example 1: normal-normal data generating model

As a modification of an example in Ackley (1987), we use the data generating model where the d -dimensional input vector $X = (X^{(1)}, \dots, X^{(d)})'$ follows a multivariate normal distribution with zero mean and identity covariance matrix, and the output V at X follows $N(\mu(X), 1)$ with

$$\begin{aligned} \mu(X) = & 20 \left(1 - \exp \left(-0.2 \sqrt{\frac{1}{d} \|X\|^2} \right) \right) \\ & + \left(\exp(1) - \exp \left(\frac{1}{d} \sum_{i=1}^d \cos(2\pi X^{(i)}) \right) \right). \end{aligned}$$

Thus, we have

$$r(x) = P(V > \xi | X = x) = 1 - \Phi(\xi - \mu(x)),$$

where $\Phi(\cdot)$ is the CDF of a standard normal distribution. As parametric models of $r(x)$, we consider two models:

(i) Correct model: We use $r_\theta(x) = 1 - \Phi(\xi - \hat{\mu}(x))$, where

$$\begin{aligned} \hat{\mu}(x) = & 20 \left(\theta_0 - \exp \left(-0.2 \sqrt{\frac{1}{d} \sum_{i=1}^d \theta_i^2 (x^{(i)})^2} \right) \right) \\ & + \left(\theta_0 \exp(1) - \exp \left(\frac{1}{d} \sum_{i=1}^d \theta_i \cos(2\pi x^{(i)}) \right) \right), \end{aligned}$$

such that $r_\theta(x) = r(x)$ for some $\theta = (\theta_0, \dots, \theta_d)' \in \Theta$. For fitting with a least square method, the initial parameters are set at the correct values, i.e., $\hat{\theta}_0 = \dots = \hat{\theta}_d = 1$, in the implementation.

(ii) Incorrect model: We use the logistic regression model $r_\theta(x) = \left(1 + e^{\theta_0 + \theta_1 x^{(1)} + \dots + \theta_d x^{(d)}} \right)^{-1}$ such that $r_\theta(x) \neq r(x)$ for all $\theta \in \Theta$. For least square fitting, the initial parameters are set at $\hat{\theta}_0 = \dots = \hat{\theta}_d = 0$.

As a nonparametric model of $r(x)$, we use the kernel regression model $\hat{r}_h(x)$ with the Gaussian kernel and the smoothing bandwidth h chosen by cross-validation.

Figures 3(a)–3(c) show how $n\text{MSE}$ varies as n increases for $d = 1, 2, 4$. For the correct parametric model, $n\text{Var}(\hat{\mathcal{E}}_{\hat{\theta}_m})$ is predicted to be $V_{\min} \left(1 + O\left(n^{-\frac{1}{3}}\right) \right)$ by Corollary 4. Recalling that MSE is equal to the variance of an importance sampler, because of its unbiasedness, we see that $n\text{MSE}$ approaches V_{\min} as n increases, with roughly the same rate regardless of d . For the incorrect parametric model, as foreseen by Theorem 5, $n\text{MSE}$ fails to approach V_{\min} as n increases, because $n\text{Var}(\hat{\mathcal{E}}_{\hat{\theta}_m}) \geq V_{\min}(1 + O(1))$. As anticipated by Corollary 8, $n\text{MSE}$ for the nonparametric model approaches V_{\min} as n increases, with an apparently slower rate for larger d , because $n\text{Var}(\hat{\mathcal{E}}_{h^*}) = V_{\min} \left(1 + O\left(\left(\frac{\log n}{n}\right)^{\frac{2}{d+6}}\right) \right)$.

Figure 3(d) shows $n\text{MSE}$ against d for fixed $n = 8000$. Regardless of d , $n\text{MSE}$ of the correct parametric importance sampler stays close to V_{\min} . In contrast, $n\text{MSE}$ for the incorrect parametric importance sampler essentially remains the same as d varies in this example, although this observation cannot be taken as a general pattern because the input configuration dimension d impacts how *incorrect* the model is. While the nonparametric importance sampler performs almost as well as the correct parametric importance sampler when $d = 1$, the performance gap widens as d increases since n is fixed.

If X was sampled only from the natural configuration density p instead of an importance sampling density q in the estimator in (2), then this simple baseline approach, commonly called crude Monte Carlo (CMC) (Kroese et al., 2011), results in the estimator having the theoretical $n\text{MSE}$ of $0.25 (= P(V > \xi)(1 - P(V > \xi)))$. In this example, the incorrect parametric importance sampler essentially does not improve over the baseline.

As d increases, V_{\min} approaches the baseline $n\text{MSE}$ of 0.25 in Figure 3(d). This increasing inefficiency of optimal importance sampling with respect to d is regarded as peculiar to this example, because V_{\min} in (3) depends on d only





FIG 3. For the normal-normal data generating model, we compare the three importance samplers in terms of their scaled estimation error, $nMSE$, against the total sample size n for the input dimension $d = 1$ in (a), $d = 2$ in (b), and $d = 4$ in (c). While fixing $n = 8000$, d is varied in (d). An error bar represents the 95% confidence interval based on the Monte Carlo error with 10,000 replications.

through $g(V(X))$. Thus, this observation should not be interpreted as a manifestation of curse of dimensionality known in the importance sampling literature (e.g., Au and Beck, 2003), which may occur when the approximation of optimal density q^* becomes harder as d increases. In contrast, it is known that the optimal importance sampler theoretically attains V_{\min} of zero regardless of d for deterministic simulation models with any nonnegative function $g(v)$ (Kahn and Marshall, 1953).

We also vary the estimand $\mathcal{E} = P(V > \xi) = 0.005, 0.05, 0.5$ to examine its impact on the performance of the importance samplers, while fixing $n = 8000$



(a) $RMSE/\mathcal{E}$ vs. $\mathcal{E} = P(V > \xi)$ in log scale. (b) Computational saving vs. $\mathcal{E} = P(V > \xi)$ in log scale.

FIG 4. For the normal-normal data generating model, we compare the three importance samplers by varying the estimand (or the probability being estimated) $\mathcal{E} = P(V > \xi) = 0.005, 0.05, 0.5$ while fixing the total sample size $n = 8000$ and the input dimension $d = 1$. We examine the normalized root mean squared error (i.e., $RMSE/\mathcal{E}$) in (a) and computational saving against CMC in (b). The computational saving is defined as $(n_{(CMC)} - n)/n_{(CMC)}$, where $n_{(CMC)} = \mathcal{E}(1 - \mathcal{E})/(\text{Standard error})^2$ is the theoretical sample size required for CMC to attain the same standard error as the importance sampler. The experiment results are averaged across 10,000 Monte Carlo replications. For the nonparametric importance sampler when $\mathcal{E} = 0.05$, one estimate (out of 10,000) of \mathcal{E} is observed to be greater than one and thus excluded in the analysis.

and $d = 1$. The initial sampling density q_0 is set as the uniform distribution on $(-5, 5)$ to broadly cover the support where rare events of interest can happen. Figure 4(a) suggests that as the estimand decreases by a factor of 10 (i.e., 0.5, 0.05, 0.005), the normalized root mean squared error (i.e. \sqrt{MSE}/\mathcal{E}) increases more slowly for the correct parametric and nonparametric importance samplers than the incorrect parametric importance sampler. Figure 4(b) shows the well-known phenomenon in the literature that importance samplers save more against CMC as the estimand is the probability of a rarer event. The savings of the correct parametric and nonparametric importance samplers quickly go over 90% as \mathcal{E} decreases to 0.005.

4.2.2. Example 2: exponential-exponential data generating model

Here, we consider a data generating model where both X and $V|X$ follow exponential distributions that have heavier tails than normal distributions and allow analytical calculations of key objects of interest such as the estimand \mathcal{E} , the conditional expectation $r(x) = \mathbb{E}(g^2(V(X))|X = x)$, the optimal sampling density $q^*(x)$, and the oracle variance V_{\min} .

Let $X = (X^{(1)}, \dots, X^{(d)})'$ be a vector of d independent exponential random variables with the identical mean $1/\lambda > 0$ so that the natural configuration density

$$p(x) = \lambda^d e^{-\lambda(x^{(1)} + \dots + x^{(d)})}.$$

Given a configuration X , let V follow an exponential distribution with a mean $1/(X^{(1)} + \dots + X^{(d)})$. In our simulation experiment, we fix $\lambda = 1$.

With the given data generating model, we can analytically calculate

$$\mathcal{E} = P(V > \xi) = \left(\frac{\lambda}{\xi + \lambda} \right)^d,$$

$$r(x) = \mathbb{E}(g^2(V(X)) | X = x) = P(V > \xi | X = x) = e^{-\xi(x^{(1)} + \dots + x^{(d)})},$$

and

$$q^*(x) \propto \sqrt{r(x)} \cdot p(x) \propto e^{-(\frac{\xi}{2} + \lambda)(x^{(1)} + \dots + x^{(d)})},$$

which implies that q^* is the joint density of d independent exponential random variables with the identical mean $1/(\xi/2 + \lambda)$. We determine

$$\xi = \frac{\lambda}{[P(V > \xi)]^{(1/d)}} - \lambda$$

by plugging $P(V > \xi) = 0.5$. We also know

$$r^\dagger(x) = \mathbb{E}(g(V(X)) | X = x) = P(V > \xi | X = x)$$

and calculate


$$\begin{aligned} V_{\min} &= \mathbb{E}^2 \left(\sqrt{r(X)} \right) - \mathbb{E}^2(r^\dagger(X)) \\ &= \left(\frac{\lambda}{\xi/2 + \lambda} \right)^{2d} - \left(\frac{\lambda}{\xi + \lambda} \right)^{2d}. \end{aligned}$$

Similar to the normal-normal example, we consider two parametric models of $r(x)$:

- (i) Correct model: We use $r_\theta(x) = e^{\theta_0 + \theta_1 x^{(1)} + \dots + \theta_d x^{(d)}}$. For least square fitting, the initial parameters are set at $\hat{\theta}_0 = \dots = \hat{\theta}_d = 0$. We set $r_{\hat{\theta}}(x) = 1$ if $r_{\hat{\theta}}(x) > 1$.
- (ii) Incorrect model: We use the logistic regression model, $r_\theta(x) = \left(1 + e^{\theta_0 + \theta_1 x^{(1)} + \dots + \theta_d x^{(d)}} \right)^{-1}$, with the initial parameters $\hat{\theta}_0 = \dots = \hat{\theta}_d = 0$ for least square fitting.

As a nonparametric model, we use the kernel regression model $\hat{r}_h(x)$ as in the normal-normal example.

Figure 5(a) plots $n\text{MSE}$ versus n for $d = 1$ with respect to the three importance samplers. The behaviors of correct and incorrect parametric importance samplers echo what we see in the normal-normal example. In contrast,



(a) $nMSE$ (calculated after removing nonparametric estimates greater than one) vs. n for $d = 1$. (b) Nonparametric estimate $\hat{\mathcal{E}}_h$ (outliers greater than one are presented as numbers at the top) vs. n for $d = 1$.

FIG 5. For the exponential-exponential data generating model, we compare the three importance samplers in terms of $nMSE$ against n for $d = 1$ in (a), where an error bar represents the 95% confidence interval based on the Monte Carlo error with 10,000 replications. In (b), the Tukey box plots are drawn based on the 10,000 nonparametric estimates for each n . The ends of the whiskers represent the most extreme data points which are not exceeding 1.5 times the interquartile range from the box.

the nonparametric importance sampler behaves irregularly (note that $nMSE$ in Figure 5(a) is calculated after discarding the nonparametric estimates exceeding one). Figure 5(b) shows the Tukey box plots of the nonparametric estimates over different n . The interquartile range (i.e., box height) decreases as n increases (note that the interquartile range of estimates may be comparable to the square root of MSE, but not directly to $nMSE$ in Figure 5(a)). The magnitudes of outliers (e.g., estimates greater than one, presented as numbers at the top of Figure 5(b) for each n) suggest that the sampling distribution of the nonparametric estimator might be heavy-tailed for this example.

We calculate effective sample sizes (ESSs) as diagnostic tools of importance sampler performance, where a too small ESS compared to n is interpreted as a sign that the importance sampling is unreliable. An ESS widely used in the importance sampling literature (Kong, 1992; Elvira et al., 2018) is

$$\widehat{ESS} = \frac{(\sum_{i=1}^n w_i)^2}{\sum_{i=1}^n w_i^2}, \tag{13}$$

where w_i is the likelihood ratio or importance weight for X_i . In the first stage ($i = 1, \dots, m$), both parametric and nonparametric importance samplers have $w_i = \frac{p(X_i)}{q_0(X_i)}$. In the second stage ($i = m + 1, \dots, n$), the parametric importance sampler has $w_i = \frac{p(X_i)}{q_{\theta_m}^*(X_i)}$ and the nonparametric importance sampler has $w_i = \frac{p(X_i)}{q_n^*(X_i)}$. While \widehat{ESS} in (13) can be conveniently used for any simulation model or




FIG 6. For the exponential-exponential data generating model, we compare the three importance samplers in terms of their effective sample size (ESS), \widehat{ESS} , in (13) and g -specific ESS, $\widehat{ESS}(g)$, in (14) when $d = 1$ and $n = 1000$. The box plots are created based on 10,000 replications. The nonparametric importance sampler yields $\widehat{ESS}(g)$ close to one, explaining the erratic $nMSE$ observed in Figure 5(a).

$g(V)$, the following ESS specific to $g(V)$ can be a better diagnostic of importance samplers, as demonstrated for deterministic simulation models (where $V(X)$ is a deterministic function of X) (Evans and Swartz, 1995; Owen, 2018):

$$\widehat{ESS}(g) = \frac{(\sum_{i=1}^n \tilde{w}_i(g))^2}{\sum_{i=1}^n \tilde{w}_i(g)^2}, \quad (14)$$

where $\tilde{w}_i(g) = |g(V)| w_i$. For stochastic simulation models, as illustrated in Figure 6(a), a very small \widehat{ESS} (close to one) does not necessarily indicate a poor estimate of \mathcal{E} . On the other hand, $\widehat{ESS}(g)$ is shown to be an effective diagnostic; in Figure 6(b), it is confirmed that when $\widehat{ESS}(g)$ is close to one, the corresponding estimates of nonparametric importance samplers tend to be far from the estimand \mathcal{E} .

The inspection of the g -specific ESS suggests that importance weights in the tail of nonparametric importance sampling density $\hat{q}_h^*(X_i)$ can be unstable. Especially because the kernel regression can poorly estimate the tail part, tail-controlling techniques (Owen and Zhou, 2000; Li et al., 2013) can be helpful to address the issue.

We attribute the erratic $nMSE$ of nonparametric importance sampler in Figure 5(a) to the severe violation of the assumption, made as a basis of our theoretical analysis in Section 3, that the natural configuration density p has a compact support $\mathbb{K} \subset \mathbb{R}^d$. In this example, the tail of $p(x) \sim \exp(-x)$ decays even more slowly than the tail of $p(x) \sim \exp(-x^2)$ in the normal-normal example. The

simulation experiment results for $d = 2$ and 4 are not presented here, as they repeat the same pattern as for $d = 1$. We note that the natural configuration density p in the case study in Section 4.3 has the compact support, indicating that the assumption is still practical.

4.3. Case study

Wind energy is one of the fastest growing renewable energy sources (You et al., 2017). Yet, harvesting wind energy remains expensive, compared with fossil energy sources such as oil, coal, and natural gas, due to the high capital cost in installing wind turbines. A utility-scale wind turbine whose blade diameter is commonly greater than 100 ft typically costs more than one million U.S. dollars. Therefore, wind energy industry pays the utmost attention on ensuring the structural reliability of the wind turbine to prevent its failure (e.g., Moriarty, 2008; Graf et al., 2017).

At the design stage of wind turbine, evaluating its reliability based on physical experiments is very limited due to the associated costs. Alternatively, the international standard, IEC 61400-1 (International Electrotechnical Commission, 2005), requires wind turbine manufacturers to use stochastic simulation models. For this purpose, the most widely used simulation models in the U.S. include TurbSim (Jonkman, 2009) and FAST (Jonkman and Buhl Jr., 2005), which are developed and maintained by the National Renewable Energy Laboratory (NREL) of the U.S. Department of Energy. TurbSim simulates a 3-dimensional time marching wind profile, which becomes an input to FAST that, in turn, simulates a wind turbine's structural response to the wind. This case study focuses on two types of bending moments at the root of a turbine blade, namely, edgewise and flapwise bending moments, which represent two perpendicular structural responses of the blade root due to an external force or moment causing it to bend. We use the same benchmark turbine model (Jonkman et al., 2009) and simulator setup as Moriarty (2008) (see the references for the rest of the settings not provided below).

In this case study, the input configuration X is a 10-min average wind speed (unit: meter per second, m/s), which is fed into TurbSim. X is sampled from the truncated Rayleigh distribution with the support of $[3, 25]$ and the scale parameter $10\sqrt{2/\pi}$. The simulation output of interest, V , is the 10-min maximum bending moment (unit: kilonewton meter, kNm) at the blade root, which is produced by FAST based on the simulated wind from TurbSim. Because V is random even for a fixed X due to the randomness of wind profile generated in TurbSim, we regard TurbSim and FAST together as a black-box stochastic simulation model.

To compare the nonparametric importance sampler proposed in this paper with a parametric importance sampler, we take a parametric method in Choe et al. (2015) as a benchmark, which also approximates the optimal sampling density $q^*(x) \propto \sqrt{r(x)} \cdot p(x)$ in Lemma 1. We use the same simulation experiment setup therein: For the edgewise bending moment, we use $n = 3600$,

$m = 600$, and $\xi = 9300$ kNm; for the flapwise bending moment, we use $n = 9600$, $m = 600$, and $\xi = 14300$ kNm.

To model $r(x)$ using a pilot sample, X is sampled $m = 600$ times from a uniform distribution with the support of $[3, 25]$, and the corresponding V 's are generated from the NREL simulators. To build the parametric model of $r(x)$, the generalized additive model for location, scale and shape (GAMLSS) (Rigby and Stasinopoulos, 2005) is fitted to the pilot sample. Specifically, the GAMLSS model assumes that the conditional distribution of V given X is a generalized extreme value distribution whose location and scale parameters are cubic spline functions of X while the shape parameter is constant over X . The model parameters are estimated using the backfitting algorithm (Rigby and Stasinopoulos, 2005). To build the nonparametric model of $r(x)$, we fit the kernel regression model to the pilot sample with the Gaussian kernel and choose the smoothing bandwidth by cross-validation.

For both parametric and nonparametric importance samplers, we repeat estimating the failure probability $\mathcal{E} = P(V > \xi)$ 50 times (in contrast to 10,000 times in the numerical study in Section 4.2). Recall that running the NREL simulators once takes about 1-min wall-clock time, implying that obtaining the pilot sample of size $m = 600$ takes roughly 10 hours. We use the same pilot sample in all 50 replications, because repeating 50 times of the simulation experiment with $n - m = 3000$ or 9000 itself requires several days of computation even if we use high-performance computing described in Section 4.2.

The parametric importance sampler in Choe et al. (2015) uses the failure probability estimator

$$\tilde{\mathcal{E}}_{\hat{\theta}_m} = \frac{1}{n - m} \sum_{i=m+1}^n g(V_i) \frac{p(X_i)}{q_{\hat{\theta}_m}^*(X_i)}, \quad (15)$$

where the pilot sample is not used, compared with the estimator in (5). In their procedure, the pilot sample is only used to build the model of $r(x)$. For fair comparison, we report both estimation results with and without using the pilot sample in the estimator. Note that nonparametric importance samplers for stochastic simulation models are not reported in the literature.

Table 1 summarizes the simulation experiment results for edgewise bending moments. We see that the standard errors of parametric importance samplers and those of nonparametric importance samplers are not significantly different. Computational savings of both methods against CMC are remarkable and, at the same time, comparable with each other.

Table 2 shows the estimation results for flapwise bending moments, which convey the similar message with the results in Table 1. Note that ξ is set to roughly yield the similar failure probability $\mathcal{E} = P(V > \xi)$ of 0.01 for both structural load types, because the magnitude of \mathcal{E} tends to impact the computational saving. Yet, we see that the computational saving of importance sampling over CMC for flapwise bending moments is, albeit substantial, not as large as that for edgewise bending moment. This is because the natural configuration density p is not very different from the optimal sampling density q^* for

TABLE 1
Estimation of the failure probability $\mathcal{E} = P(V > \xi)$ for the edgewise bending moment V and the threshold $\xi = 9300$ kNm

Method	Without the pilot sample			With the pilot sample		
	Sample mean	Standard error (95% bootstrap CI)	Comp. saving	Sample mean	Standard error (95% bootstrap CI)	Comp. saving
Parametric	0.01005	0.00044 (0.00036, 0.00051)	93%	0.01016	0.00036 (0.00030, 0.00042)	95%
Nonparametric	0.00998	0.00046 (0.00034, 0.00056)	92%	0.01010	0.00038 (0.00029, 0.00047)	95%

Note: The computational (comp.) saving is $(n_{(CMC)} - n)/n_{(CMC)}$, where $n_{(CMC)} = \hat{P}(1 - \hat{P})/(\text{Standard error})^2$ is the theoretical sample size required for CMC to attain the same standard error when the true failure probability is equal to \hat{P} , which is the sample mean of the parametric importance samplers using the pilot sample. The 95% bootstrap confidence interval (CI) is constructed based on 100,000 bootstrap replicates.

TABLE 2
Estimation of the failure probability $\mathcal{E} = P(V > \xi)$ for the flapwise bending moment V and the threshold $\xi = 14300$ kNm

Method	Without the pilot sample			With the pilot sample		
	Sample mean	Standard error (95% bootstrap CI)	Comp. saving	Sample mean	Standard error (95% bootstrap CI)	Comp. saving
Parametric	0.01037	0.00063 (0.00046, 0.00078)	64%	0.01079	0.00059 (0.00043, 0.00073)	69%
Nonparametric	0.01061	0.00075 (0.00057, 0.00090)	49%	0.01101	0.00070 (0.00053, 0.00084)	56%

Note: Refer to the note of Table 1.

flapwise bending moments so that the benefit of changing the sampling density is not enormous.

This case study considers a relatively moderate failure probability $\mathcal{E} = P(V > \xi)$ around 0.01, so we can use the given computational resource to have 50 replications that allow us to construct bootstrap confidence intervals on standard errors and compare the performances of methods. On the other hand, simulation experiments at the scale of a national lab may consider a more extreme event with a smaller probability. For such case, combining the 50 replications can lead to an estimate of a much smaller probability (in the order of 1e-05 or smaller as in Choe et al. (2016)). Still, one can construct a confidence interval on the probability if only using the sample from the second stage (Choe et al., 2017).

5. Discussion

We consider the problem of estimating the average output from a stochastic simulation model and propose two-stage estimators using either a parametric approach or a nonparametric approach. Theoretically, both estimators satisfy the oracle inequalities but they achieve the oracle variance asymptotically under different rates and assumptions. As expected, the parametric approach needs a strong assumption but its variance converges to the oracle variance faster than the nonparametric approach. The nonparametric approach, however, requires

weak assumptions but the variance reduction rate is not as fast as the parametric approach. Empirically, our numerical study confirmed the theoretical results, and our case study indicated that the proposed importance samplers perform well in practice, saving 50%–95% computational resources over a standard Monte Carlo estimator.

We note that Choe et al. (2015) investigated a parametric importance sampler for a stochastic simulation model, which is a special case (with $g(v) = 1(v \geq \xi)$) of our parametric importance sampler. They use implicitly a two-stage method with the pilot sample being given *a priori* for modeling $r(x)$ and do not use the pilot sample in the estimator. They neither consider nonparametric importance samplers nor perform a theoretical analysis on the convergence toward the optimal estimator. Our work focuses more on the formalization of the general two-stage strategy and the theoretical foundation of importance sampling for stochastic simulation models. Because we obtain a concrete convergence rate of the estimator, we enable the optimal choice of the pilot sample size.

In what follows we discuss possible future research directions.

- **Manifold support case.** In general, when the dimension of the configuration d is large, nonparametric importance sampler in Section 2.2 will not work due to the curse of dimensionality (the slow convergence rate). However, if the support of q^* is concentrated around a lower dimensional manifold, the nonparametric sampler may still work because a fast convergence rate of a nonparametric estimator is possible (Balakrishnan et al., 2013; Chen, 2016). So we may be able to design a modified nonparametric importance sampling procedure that achieves oracle variance much faster than the rate in Corollary 8. The construction of such a procedure is left as a future work.
- **Multi-stage sampling.** In this paper we only consider splitting the computational budget into two stages. We can generalize this idea into a k -stage sampling procedure, where at each stage, we use samples in all previous stages to design our estimator and sampler for the current stage (e.g., Choe, 2017). In this case, the allocation problem becomes more complicated since we may assign different sample sizes to different stages. Also, the number of stage k will be another quantity that we want to optimize. Because the two-stage approach is a special case of a multi-stage sampling procedure, the latter will have a higher variance reduction rate than the proposed methods in this paper.
- **Confidence interval.** In the current paper, we focus on the construction of an estimator of \mathcal{E} . In practice, we often report not only a point estimator but also a confidence interval attached to it. Here we briefly describe two potential methods of constructing the confidence interval. The first method is to derive asymptotic normality of $\hat{\mathcal{E}}$ and then find a consistent variance estimator. Note that this is a non-trivial task because when we use a two-stage approach, the observations are no longer IID. Moreover, estimating the variance could be another challenging task. The other approach is to use the bootstrap (Efron, 1982, 1992) to obtain a confidence interval. If

we choose to use the bootstrap, we need to prove the validity of such a bootstrap procedure.

- **Multiple estimands of interest.** This paper considers the single quantity to be estimated, $\mathcal{E} = \mathbb{E}(g(V(X)))$, in (1) based on the output $V(X)$ from a stochastic simulation model. Extending the current work to optimally estimating multiple quantities of interest warrants further investigation. Such extension for stochastic simulation models can build upon this line of research around deterministic simulation models. For example, we can consider importance sampling of the union of target events (Owen et al., 2019). Also, we can recycle existing simulation results (e.g., from the pilot stage or both stages of estimating another quantity) by recomputing importance weights (Cornuet et al., 2012). These approaches can benefit from the classical idea of safe and effective importance sampling (Owen and Zhou, 2000).

Acknowledgements

Y.-C. Chen was partially supported by NSF grant DMS-1810960 and NIH U01-AG016976. Y. Choe was partially supported by NSF grant CMMI-1824681.

Appendix A: Proofs

Proof of Lemma 1. Now by the following variance formula:

$$\text{Var}(Y) = \mathbb{E}(\text{Var}(Y|X)) + \text{Var}(\mathbb{E}(Y|X))$$

and choose $Y = g(V_i) \frac{p(X_i)}{q(X_i)}$ and $X = X_i$, we have

$$\begin{aligned} \text{Var} \left(g(V_i) \frac{p(X_i)}{q(X_i)} \right) &= \mathbb{E} \left(\text{Var} \left(g(V_i) \frac{p(X_i)}{q(X_i)} \middle| X_i \right) \right) + \text{Var} \left(r^\dagger(X_i) \frac{p(X_i)}{q(X_i)} \right) \\ &= \mathbb{E} \left(\left(r(X_i) - r^\dagger(X_i) \right) \frac{p^2(X_i)}{q^2(X_i)} \right) \\ &\quad + \mathbb{E} \left(r^{\dagger 2}(X_i) \frac{p^2(X_i)}{q^2(X_i)} \right) - \mathbb{E}^2 \left(r^\dagger(X_i) \frac{p(X_i)}{q(X_i)} \right) \\ &= \mathbb{E} \left(r(X_i) \frac{p^2(X_i)}{q^2(X_i)} \right) - \mathbb{E}^2(r^\dagger(X^*)), \end{aligned} \tag{16}$$

when $q(x) = 0$ implies $r^\dagger(x)p(x) = 0$. Note that X^* is from density p . Thus, the sampling density q affects the variance only via the quantity $\mathbb{E} \left(r(X_i) \frac{p^2(X_i)}{q^2(X_i)} \right)$.

The quantity $\mathbb{E} \left(r(X_i) \frac{p^2(X_i)}{q^2(X_i)} \right)$ has a lower bound from the Cauchy-Schwarz inequality:

$$\begin{aligned}
 \mathbb{E} \left(r(X_i) \frac{p^2(X_i)}{q^2(X_i)} \right) &= \int r(x) \frac{p^2(x)}{q(x)} dx \\
 &= \int \left(\sqrt{r(x)} \frac{p(x)}{\sqrt{q(x)}} \right)^2 dx \cdot \underbrace{\int (\sqrt{q(x)})^2 dx}_{=1} \\
 &\geq \left(\int \sqrt{r(x)} p(x) dx \right)^2 \\
 &= \mathbb{E}^2 \left(\sqrt{r(X^*)} \right).
 \end{aligned} \tag{17}$$

And the equality holds when $\sqrt{r(x)} \frac{p(x)}{\sqrt{q(x)}} \propto \sqrt{q(x)}$, which implies the optimal sampling density is

$$q^*(x) \propto \sqrt{r(x)} \cdot p(x).$$

Thus, when we choose the sampling density to be q^* , by equation (16) and (17), the variance

$$\begin{aligned}
 \text{Var} \left(\widehat{\mathcal{E}}_{q^*} \right) &= \frac{1}{n} \mathbb{E} \left(r(X_i) \frac{p^2(X_i)}{q^2(X_i)} \right) - \mathbb{E}^2(r^\dagger(X^*)) \\
 &= \frac{1}{n} \left(\mathbb{E}^2 \left(\sqrt{r(X^*)} \right) - \mathbb{E}^2 \left(r^\dagger(X^*) \right) \right). \quad \square
 \end{aligned}$$

Proof of Theorem 2. By assumptions (P1–2) and the M-estimation theory (van der Vaart and Wellner, 1996; van der Vaart, 2000),

$$\|\widehat{\theta}_m - \theta_0\| = O_P \left(\sqrt{\frac{1}{m}} \right),$$

where θ_0 is the parameter such that $r(x) = r_{\theta_0}(x)$.

Now by assumption (P3),

$$\begin{aligned}
 \sup_{x \in \mathbb{K}} \|r_{\widehat{\theta}_m}(x) - r(x)\| &= \sup_{x \in \mathbb{K}} \|r_{\widehat{\theta}_m}(x) - r_{\theta_0}(x)\| \\
 &\leq \sup_{x \in \mathbb{K}} \left\| (\widehat{\theta}_m - \theta_0) \cdot L_0 \right\| \\
 &= O_P \left(\sqrt{\frac{1}{m}} \right),
 \end{aligned}$$

which proves the result. □

Proof of Theorem 3. For our estimator $\widehat{\mathcal{E}}_{\widehat{\theta}_m}$, we decompose it into two parts

$$\widehat{\mathcal{E}}_{\widehat{\theta}_m} = A_m + B_n,$$

where

$$A_m = \frac{1}{n} \sum_{i=1}^m g(V_i) \frac{p(X_i)}{q_0(X_i)},$$

$$B_n = \frac{1}{n} \sum_{i=m+1}^n g(V_i) \frac{p(X_i)}{q_{\hat{\theta}_m}^*(X_i)},$$

Thus,

$$\text{Var}\left(\widehat{\mathcal{E}}_{\hat{\theta}_m}\right) = \text{Var}(A_m + B_n) = \text{Var}(A_m) + \text{Var}(B_n) + 2\text{Cov}(A_m, B_n). \quad (18)$$

Note that

$$\begin{aligned} \mathbb{E}(A_m) &= \frac{m}{n} \mathbb{E}_{X_i \sim q_0} \left(g(V_i) \frac{p(X_i)}{q_0(X_i)} \right) = \frac{m}{n} \mathcal{E} \\ \mathbb{E}(B_n) &= \frac{n-m}{n} \cdot \mathbb{E}_{X_i \sim q_{\hat{\theta}_m}^*} \left(g(V_i) \frac{p(X_i)}{q_{\hat{\theta}_m}^*(X_i)} \right) \\ &= \frac{n-m}{n} \cdot \mathbb{E}_{X_i \sim q_{\hat{\theta}_m}^*} \left(\mathbb{E} \left(g(V_i) \frac{p(X_i)}{q_{\hat{\theta}_m}^*(X_i)} \mid q_{\hat{\theta}_m}^* \right) \right) \\ &= \frac{n-m}{n} \cdot \mathcal{E}. \end{aligned} \quad (19)$$

We first bound the covariance. Let $\mathcal{D}_m = \{(X_1, V_1), \dots, (X_m, V_m)\}$ be the collection of the first part of the data. Then

$$\begin{aligned} \text{Cov}(A_m, B_n) &= \mathbb{E}(A_m B_n) - \mathbb{E}(A_m) \mathbb{E}(B_n) \\ &= \mathbb{E}(A_m \mathbb{E}(B_n \mid \mathcal{D}_m)) - \frac{(n-m) \cdot m}{n^2} \cdot \mathcal{E}^2 \\ &= \mathbb{E}(A_m \mathbb{E}(B_n \mid q_{\hat{\theta}_m}^*)) - \frac{(n-m) \cdot m}{n^2} \cdot \mathcal{E}^2 \\ &= \mathbb{E} \left(A_m \cdot \frac{n-m}{n} \cdot \mathcal{E} \right) - \frac{(n-m) \cdot m}{n^2} \cdot \mathcal{E}^2 \\ &= \frac{(n-m) \cdot m}{n^2} \cdot \mathcal{E}^2 - \frac{(n-m) \cdot m}{n^2} \cdot \mathcal{E}^2 \\ &= 0. \end{aligned} \quad (20)$$

Therefore, we only need to focus on the variance of each part.

Let $V_{\min} = \left(\mathbb{E}^2 \left(\sqrt{r(X^*)} \right) - \mathbb{E}^2(r^\dagger(X^*)) \right)$ be the minimal variance under the optimal sampling density. By Lemma 1,

$$\begin{aligned} \text{Var}(A_m) &= \frac{m}{n^2} \left(\mathbb{E} \left(r(X_i) \frac{p^2(X_i)}{q_0^2(X_i)} \right) - \mathbb{E}^2(r^\dagger(X^*)) \right) \\ &= \frac{m}{n^2} V_{\min} + \frac{m}{n^2} \left(\mathbb{E} \left(r(X_i) \frac{p^2(X_i)}{q_0^2(X_i)} \right) - \mathbb{E}^2 \left(\sqrt{r(X^*)} \right) \right) \\ &= \frac{m}{n^2} V_{\min} + \frac{m}{n^2} V_{q_0}. \end{aligned} \quad (21)$$

And the variance of the second part is

$$\begin{aligned} \text{Var}(B_n) &= \mathbb{E}(\text{Var}(B_n|\mathcal{D}_m)) + \text{Var}(\mathbb{E}(B_n|\mathcal{D}_m)) \\ &= \mathbb{E}(\text{Var}(B_n|\mathcal{D}_m)) + \underbrace{\text{Var}(\mathcal{E})}_{=0} \\ &= \frac{n-m}{n^2} \cdot \mathbb{E} \left(\mathbb{E} \left(r(X_i) \frac{p^2(X_i)}{q_{\hat{\theta}_m}^*(X_i)} \middle| q_{\hat{\theta}_m}^* \right) - \mathbb{E}^2(r^\dagger(X^*)) \right) \\ &= \frac{n-m}{n^2} \cdot \mathbb{E} \left(\mathbb{E} \left(r(X^*) \frac{p(X^*)}{q_{\hat{\theta}_m}^*(X^*)} \middle| q_{\hat{\theta}_m}^* \right) \right) - \frac{n-m}{n^2} \cdot \mathbb{E}^2(r^\dagger(X^*)). \end{aligned} \tag{22}$$

So the key part is in the quantity $\mathbb{E} \left(\mathbb{E} \left(r(X^*) \frac{p(X^*)}{q_{\hat{\theta}_m}^*(X^*)} \middle| q_{\hat{\theta}_m}^* \right) \right)$. By Theorem 2 we have

$$\sup_{x \in \mathbb{K}} \|r_{\hat{\theta}_m}(x) - r(x)\| = O_P \left(\sqrt{\frac{1}{m}} \right),$$

which implies

$$q_{\hat{\theta}_m}^*(x) - q^*(x) = \Delta_m \cdot p(x),$$

where $\Delta_m = O_P \left(\sqrt{\frac{1}{m}} \right)$. Thus,

$$\begin{aligned} \mathbb{E} \left(r(X^*) \frac{p(X^*)}{q_{\hat{\theta}_m}^*(X^*)} \middle| q_{\hat{\theta}_m}^* \right) &= \int r(x) \frac{p^2(x)}{q_{\hat{\theta}_m}^*(x)} dx \\ &= \int r(x) \frac{p^2(x)}{q^*(x) + \Delta_m \cdot p(x)} dx \\ &= \int r(x) \frac{p^2(x)}{q^*(x)} dx + O(\Delta_m) \\ &= \mathbb{E}^2 \left(\sqrt{r(X^*)} \right) + O_P \left(\sqrt{\frac{1}{m}} \right). \end{aligned} \tag{23}$$

Note that in the above equation, the constant in the $O(\Delta_m)$ term is

$$\int r(x)p^3(x)/q^*(x)^2 dx = \int r(x)p(x)/[C^2 \cdot r(x)] dx = \frac{1}{C^2} \int p(x) dx = \frac{1}{C^2},$$

where C is the constant from equating $q^*(x) = C\sqrt{r(x)}p(x)$. The assumption (P1) implies that the ratio $\frac{p(x)}{q^*(x)} = \frac{1}{C \cdot r^{-1/2}(x)}$ is uniformly bounded. Thus, the higher order terms are bounded.

Putting this back to equation (22), we obtain

$$\text{Var}(B_n) = \frac{n-m}{n^2} \left(\mathbb{E}^2 \left(\sqrt{r(X^*)} \right) - \mathbb{E}^2(r^\dagger(X^*)) \right) + \frac{n-m}{n^2} \mathbb{E} \left(O_P \left(\sqrt{\frac{1}{m}} \right) \right)$$

$$= \frac{n-m}{n^2} V_{\min} + \frac{n-m}{n^2} O\left(\sqrt{\frac{1}{m}}\right).$$

Note that $\mathbb{E}\left(O_P\left(\sqrt{\frac{1}{m}}\right)\right) = O\left(\sqrt{\frac{1}{m}}\right)$ because the random quantity in the O_P term is from Δ_m and is from the difference $|\hat{\theta}_m - \theta_0|$, which is the absolute value of an asymptotic normal distribution so the expectation of the O_P leads to the same convergence rate.

Now putting altogether, we obtain

$$\begin{aligned} \text{Var}\left(\widehat{\mathcal{E}}_{\hat{\theta}_m}\right) &= \frac{m}{n^2} V_{\min} + \frac{m}{n^2} V_{q_0} + \frac{n-m}{n^2} V_{\min} + \frac{n-m}{n^2} O\left(\sqrt{\frac{1}{m}}\right) \\ &= \frac{1}{n} V_{\min} + \frac{1}{n^2} \left(m \cdot V_{q_0} + (n-m) \cdot O\left(\sqrt{\frac{1}{m}}\right)\right). \quad \square \end{aligned} \tag{24}$$

Proof of Theorem 6. Recall that $\widehat{r}_h(x) = \frac{\sum_{i=1}^m Y_i K\left(\frac{x-X_i}{h}\right)}{\sum_{i=1}^m K\left(\frac{x-X_i}{h}\right)}$ and $r(x) = \mathbb{E}(Y_i|X_i = x)$. Define the following two quantities

$$\begin{aligned} \widetilde{r}_h(x) &= \frac{\frac{1}{mh^d} \sum_{i=1}^m Y_i K\left(\frac{x-X_i}{h}\right)}{q_0(x)}, \\ \bar{r}_h(x) &= \frac{\mathbb{E}\left(\frac{1}{h^d} Y_i K\left(\frac{x-X_i}{h}\right)\right)}{q_0(x)}. \end{aligned} \tag{25}$$

We can bound the difference $\sup_{x \in \mathbb{K}} \|\widehat{r}_h(x) - r(x)\|$ by

$$\begin{aligned} \sup_{x \in \mathbb{K}} \|\widehat{r}_h(x) - r(x)\| &= \sup_{x \in \mathbb{K}} \|\widehat{r}_h(x) - \widetilde{r}_h(x) + \widetilde{r}_h(x) - \bar{r}_h(x) + \bar{r}_h(x) - r(x)\| \\ &\leq \sup_{x \in \mathbb{K}} \|\widehat{r}_h(x) - \widetilde{r}_h(x)\| + \sup_{x \in \mathbb{K}} \|\widetilde{r}_h(x) - \bar{r}_h(x)\| + \sup_{x \in \mathbb{K}} \|\bar{r}_h(x) - r(x)\|. \end{aligned} \tag{26}$$

Now we separately bound each term.

The first term $\sup_{x \in \mathbb{K}} \|\widehat{r}_h(x) - \widetilde{r}_h(x)\|$ involves the difference between $q_0(x)$ and $\frac{1}{mh^d} \sum_{i=1}^m K\left(\frac{x-X_i}{h}\right) = \widehat{q}_m(x)$, which is the difference between the kernel density estimator (KDE) $\widehat{q}_m(x)$ and $q_0(x)$. By assumption (N2) and (K1-2), it is known in the literature that

$$\sup_{x \in \mathbb{K}} \|\widehat{q}_m(x) - q_0(x)\| = O(h^2) + O_P\left(\sqrt{\frac{\log m}{mh^d}}\right);$$

see, e.g., Lemma 5 in Chen et al. (2015a) and Lemma 10 in Chen et al. (2015b). Thus,

$$\sup_{x \in \mathbb{K}} \|\widehat{r}_h(x) - \widetilde{r}_h(x)\| = O(h^2) + O_P\left(\sqrt{\frac{\log m}{mh^d}}\right). \tag{27}$$

For the second term, it equals to

$$\sup_{x \in \mathbb{K}} \|\widetilde{r}_h(x) - \bar{r}_h(x)\|$$

$$\begin{aligned}
&= \sup_{x \in \mathbb{K}} \left\| \frac{\frac{1}{mh^d} \sum_{i=1}^m Y_i K\left(\frac{x-X_i}{h}\right)}{q_0(x)} - \frac{\mathbb{E}\left(\frac{1}{h^d} Y_i K\left(\frac{x-X_i}{h}\right)\right)}{q_0(x)} \right\| \\
&= \sup_{x \in \mathbb{K}} \left\| \frac{1}{q_0(x)} \left(\frac{1}{mh^d} \sum_{i=1}^m Y_i K\left(\frac{x-X_i}{h}\right) - \mathbb{E}\left(\frac{1}{h^d} Y_i K\left(\frac{x-X_i}{h}\right)\right) \right) \right\|.
\end{aligned}$$

Now using Theorem 2.3 in Giné and Guillou (2002) and assumption (N1) and (K1–2), we can bound

$$\sup_{x \in \mathbb{K}} \left\| \frac{1}{mh^d} \sum_{i=1}^m Y_i K\left(\frac{x-X_i}{h}\right) - \mathbb{E}\left(\frac{1}{h^d} Y_i K\left(\frac{x-X_i}{h}\right)\right) \right\| = O_P\left(\sqrt{\frac{\log m}{mh^d}}\right).$$

Assumption (N2) implies that the density $q_0(x)$ is lower bounded by q_{\min} . Thus, we obtain the bound

$$\begin{aligned}
&\sup_{x \in \mathbb{K}} \|\tilde{r}_h(x) - \bar{r}_h(x)\| \\
&\leq \frac{1}{q_{\min}} \sup_{x \in \mathbb{K}} \left\| \frac{1}{mh^d} \sum_{i=1}^m Y_i K\left(\frac{x-X_i}{h}\right) - \mathbb{E}\left(\frac{1}{h^d} Y_i K\left(\frac{x-X_i}{h}\right)\right) \right\| \\
&= O_P\left(\sqrt{\frac{\log m}{mh^d}}\right). \tag{28}
\end{aligned}$$

The third term $\sup_{x \in \mathbb{K}} \|\bar{r}_h(x) - r(x)\|$ involves the bias in nonparametric regression which is known to be at rate $O(h^2)$ under assumption (N2). Based on this rate and equations (27) and (28), by equation (26) we obtain

$$\sup_{x \in \mathbb{K}} \|\hat{r}_h(x) - r(x)\| = O(h^2) + O_P\left(\sqrt{\frac{\log m}{mh^d}}\right),$$

which is the desired result. \square

Proof of Theorem 7. This proof follows the same way as the proof of Theorem 3: we decompose $\hat{\mathcal{E}}_{h^*} = A_m + B_n$ and control the variance of A_m and B_n and show that the covariance is 0.

The only difference is in the variance of B_n . Because the estimation error now becomes (see equation (9))

$$\sup_{x \in \mathbb{K}} \|\hat{r}_{h^*}(x) - r(x)\| = O_P\left(\left(\frac{\log m}{m}\right)^{\frac{2}{d+4}}\right),$$

the variance

$$\text{Var}(B_n) = \frac{n-m}{n^2} \left(\mathbb{E}^2\left(\sqrt{r(X^*)}\right) - \mathbb{E}^2(r^\dagger(X^*)) \right) + \frac{n-m}{n^2} O\left(\left(\frac{\log m}{m}\right)^{\frac{2}{d+4}}\right)$$

$$= \frac{n-m}{n^2} V_{\min} + \frac{n-m}{n^2} O\left(\left(\frac{\log m}{m}\right)^{\frac{2}{d+4}}\right).$$

Thus, the total variance is

$$\begin{aligned} \text{Var}\left(\widehat{\mathcal{E}}_{h^*}\right) &= \text{Var}(A_m) + \text{Var}(B_n) + \underbrace{2\text{Cov}(A_m, B_n)}_{=0} \\ &= \frac{m}{n^2} V_{\min} + \frac{m}{n^2} V_{q_0} + \frac{n-m}{n^2} V_{\min} + \frac{n-m}{n^2} O\left(\left(\frac{\log m}{m}\right)^{\frac{2}{d+4}}\right) \\ &= \frac{1}{n} V_{\min} + \frac{1}{n^2} \left(m \cdot V_{q_0} + (n-m) \cdot O\left(\left(\frac{\log m}{m}\right)^{\frac{2}{d+4}}\right)\right), \end{aligned}$$

which proves the desired result. \square

Acknowledgement

We thank two referees for very helpful comments on this paper. Dr. Chen is partially supported by NSF grant DMS 1810960 and NIH grant U01 AG016976 and Dr. Choe is partially supported by NSF grant CMMI-1824681.

References

- D. H. Ackley. *A connectionist machine for genetic hillclimbing*. Boston: Kluwer Academic Publishers, 1987.
- B. Ankenman, B. L. Nelson, and J. Staum. Stochastic kriging for simulation metamodeling. *Operations Research*, 58(2):371–382, 2010. [MR2674803](#)
- S. Au and J. L. Beck. A new adaptive importance sampling scheme for reliability calculations. *Structural Safety*, 21(2):135–158, 1999.
- S. K. Au and J. L. Beck. Important sampling in high dimensions. *Structural Safety*, 25(2):139–163, 2003.
- S. Balakrishnan, S. Narayanan, A. Rinaldo, A. Singh, and L. Wasserman. Cluster trees on manifolds. In *Advances in Neural Information Processing Systems*, pages 2679–2687, 2013.
- J. Blanchet and H. Lam. Importance sampling for actuarial cost analysis under a heavy traffic model. In *Proceedings of the Winter Simulation Conference*, pages 3817–3828. Winter Simulation Conference, 2011.
- A. Brennan, S. Kharroubi, A. O’Hagan, and J. Chilcott. Calculating partial expected value of perfect information via Monte Carlo sampling algorithms. *Medical Decision Making*, 27(4):448–470, 2007.
- J. Bucklew. *Introduction to rare event simulation*. New York: Springer-Verlag, 2004. [MR2045385](#)
- Y.-C. Chen. Generalized cluster trees and singular measures. *arXiv preprint arXiv:1611.02762*, 2016. [MR3953448](#)

- Y.-C. Chen, C. R. Genovese, S. Ho, and L. Wasserman. Optimal ridge detection using coverage risk. In *Advances in Neural Information Processing Systems*, pages 316–324, 2015a.
- Y.-C. Chen, C. R. Genovese, and L. Wasserman. Asymptotic theory for density ridges. *The Annals of Statistics*, 43(5):1896–1928, 2015b. [MR3375871](#)
- Y.-C. Chen, C. R. Genovese, R. J. Tibshirani, and L. Wasserman. Non-parametric modal regression. *The Annals of Statistics*, 44(2):489–514, 2016. [MR3476607](#)
- Y.-C. Chen, C. R. Genovese, and L. Wasserman. Density level sets: Asymptotics, inference, and visualization. *Journal of the American Statistical Association*, pages 1–13, 2017. [MR3750891](#)
- Y. Choe. Information criterion for minimum cross-entropy model selection. *arXiv preprint arXiv:1704.04315*, 2017.
- Y. Choe, E. Byon, and N. Chen. Importance sampling for reliability evaluation with stochastic simulation models. *Technometrics*, 57(3):351–361, 2015. [MR3384950](#)
- Y. Choe, Q. Pan, and E. Byon. Computationally efficient uncertainty minimization in wind turbine extreme load assessments. *Journal of Solar Energy Engineering*, 138(4):041012–041012–8, 2016.
- Y. Choe, H. Lam, and E. Byon. Uncertainty quantification of stochastic simulation for black-box computer experiments. *Methodology and Computing in Applied Probability*, Oct 2017. ISSN 1573-7713. [MR3873620](#)
- J.-M. Cornuet, J.-M. Marin, A. Mira, and C. P. Robert. Adaptive multiple importance sampling. *Scandinavian Journal of Statistics*, 39(4):798–812, 2012. [MR3000850](#)
- B. Efron. *The jackknife, the bootstrap and other resampling plans*. SIAM, 1982. [MR0659849](#)
- B. Efron. Bootstrap methods: another look at the jackknife. In *Breakthroughs in Statistics*, pages 569–593. Springer, 1992. [MR0659849](#)
- U. Einmahl and D. M. Mason. Uniform in bandwidth consistency of kernel-type function estimators. *The Annals of Statistics*, 33(3):1380–1403, 2005. [MR2195639](#)
- V. Elvira, L. Martino, and C. P. Robert. Rethinking the effective sample size. *arXiv preprint arXiv:1809.04129*, 2018.
- M. Evans and T. Swartz. Methods for approximating integrals in statistics with special emphasis on Bayesian integration problems. *Statistical Science*, 10(3):254–272, 1995. [MR1390518](#)
- C. R. Genovese, M. Perone-Pacifco, I. Verdinelli, and L. Wasserman. Non-parametric ridge estimation. *The Annals of Statistics*, 42(4):1511–1545, 2014. [MR3262459](#)
- D. T. Gillespie. Approximate accelerated stochastic simulation of chemically reacting systems. *The Journal of Chemical Physics*, 115(4):1716–1733, 2001.
- E. Giné and A. Guillou. Rates of strong uniform consistency for multivariate kernel density estimators. In *Annales de l'Institut Henri Poincaré (B) Probability and Statistics*, volume 38, pages 907–921. Elsevier, 2002. [MR1955344](#)
- P. Glasserman and J. Li. Importance sampling for portfolio credit risk. *Man-*

- agement Science*, 51(11):1643–1656, 2005.
- P. W. Glynn and D. L. Iglehart. Importance sampling for stochastic simulations. *Management Science*, 35(11):1367–1392, 1989. [MR1024494](#)
- P. Graf, R. R. Damiani, K. Dykes, and J. M. Jonkman. Advances in the assessment of wind turbine operating extreme loads via more efficient calculation approaches. In *35th Wind Energy Symposium*, pages 1–19, Grapevine, TX, 2017. AIAA SciTech Forum.
- P. A. Graf, G. Stewart, M. Lackner, K. Dykes, and P. Veers. High-throughput computation and the applicability of Monte Carlo integration in fatigue load estimation of floating offshore wind turbines. *Wind Energy*, 19(5):861–872, 2016.
- R. B. Gramacy and H. K. Lee. Bayesian treed Gaussian process models with an application to computer modeling. *Journal of the American Statistical Association*, 103(483):1119–1130, 2012. [MR2528830](#)
- L. Györfi, M. Kohler, A. Krzyzak, and H. Walk. *A distribution-free theory of nonparametric regression*. Springer Science & Business Media, 2006. [MR1920390](#)
- P. Heidelberger. Fast simulation of rare events in queueing and reliability models. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 5(1):43–85, 1995.
- D. A. Henderson, R. J. Boys, K. J. Krishnan, C. Lawless, and D. J. Wilkinson. Bayesian emulation and calibration of a stochastic computer model of mitochondrial dna deletions in substantia nigra neurons. *Journal of the American Statistical Association*, 104(485):76–87, 2012. [MR2663034](#)
- International Electrotechnical Commission. IEC/TC88, 61400-1 ed. 3, Wind Turbines - Part 1: Design Requirements., 2005.
- B. J. Jonkman. TurbSim user’s guide: version 1.50. Technical Report NREL/TP-500-46198, National Renewable Energy Laboratory, Golden, Colorado, 2009.
- J. M. Jonkman and M. L. Buhl Jr. FAST User’s Guide. Technical Report NREL/EL-500-38230, National Renewable Energy Laboratory, Golden, Colorado, 2005.
- J. M. Jonkman, S. Butterfield, W. Musial, and G. Scott. Definition of a 5-MW reference wind turbine for offshore system development. Technical Report NREL/TP-500-38060, National Renewable Energy Laboratory, Golden, Colorado, 2009.
- H. Kahn and A. W. Marshall. Methods of reducing sample size in Monte Carlo computations. *Journal of the Operations Research Society of America*, 1(5):263–278, 1953. [MR0056266](#)
- A. Kong. A note on importance sampling using standardized weights. *University of Chicago, Dept. of Statistics, Tech. Rep.*, 348, 1992.
- D. P. Kroese, T. Taimre, and Z. I. Botev. *Handbook of Monte Carlo methods*. New York: John Wiley and Sons., 2011.
- E. Lawrence, S. V. Wiel, and R. Bent. Model bank state estimation for power grids using importance sampling. *Technometrics*, 55(4):426–435, 2013. [MR3176548](#)
- W. Li, Z. Tan, and R. Chen. Two-stage importance sampling with mixture

- proposals. *Journal of the American Statistical Association*, 108(504):1350–1365, 2013. [MR3174713](#)
- L. Manuel, H. H. Nguyen, and M. F. Barone. On the use of a large database of simulated wind turbine loads to aid in assessing design standard provisions. In *Proceedings of the 51st AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, Grapevine, Texas, Jan. 2013.
- P. Moriarty. Database for validation of design load extrapolation techniques. *Wind Energy*, 11(6):559–576, 2008.
- E. A. Nadaraya. On estimating regression. *Theory of Probability & Its Applications*, 9(1):141–142, 1964.
- J. C. Neddermeyer. Computationally efficient nonparametric importance sampling. *Journal of the American Statistical Association*, 104(486):788–802, 2009. [MR2541595](#)
- A. Owen and Y. Zhou. Safe and effective importance sampling. *Journal of the American Statistical Association*, 95(449):135–143, 2000. [MR1803146](#)
- A. B. Owen. Importance sampling. In *Monte Carlo Theory, Methods and Examples*, chapter 9. 2018.
- A. B. Owen, Y. Maximov, and M. Chertkov. Importance sampling the union of rare events with an application to power systems analysis. *Electronic Journal of Statistics*, 13(1):231–254, 2019. [MR3905126](#)
- V. Picheny, D. Ginsbourger, Y. Richet, and G. Caplin. Quantile-based optimization of noisy computer experiments with tunable precision. *Technometrics*, 55(1):2–13, 2013. [MR3038476](#)
- M. Plumlee and R. Tuo. Building accurate emulators for stochastic simulations via quantile kriging. *Technometrics*, 56(4):466–473, 2014. [MR3290616](#)
- Z. Qian, C. C. Seepersad, V. R. Joseph, J. K. Allen, and C. J. Wu. Building surrogate models based on detailed and approximate simulations. *Journal of Mechanical Design*, 128(4):668–677, 2006.
- R. A. Rigby and D. M. Stasinopoulos. Generalized additive models for location, scale and shape. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 54(3):507–554, 2005. [MR2137253](#)
- D. W. Scott. *Multivariate density estimation: theory, practice, and visualization*. John Wiley & Sons, 2015. [MR3329609](#)
- J. Staum. *Monte Carlo computation in finance*, pages 19–42. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. [MR2743886](#)
- Y. Sun, D. W. Apley, and J. Staum. Efficient nested simulation for estimating the variance of a conditional expectation. *Operations Research*, 59(4):998–1007, 2011. [MR2844419](#)
- A. B. Tsybakov. Introduction to nonparametric estimation. Revised and extended from the 2004 french original. Translated by Vladimir Zaiats, 2009. [MR2724359](#)
- A. W. van der Vaart. *Asymptotic statistics*, volume 3. Cambridge University Press, 2000. [MR1652247](#)
- A. W. van der Vaart and J. A. Wellner. Weak convergence. In *Weak Convergence and Empirical Processes*, pages 16–28. Springer, 1996. [MR1385671](#)
- G. G. Wang and S. Shan. Review of metamodeling techniques in support of

- engineering design optimization. *Journal of Mechanical design*, 129(4):370–380, 2007.
- L. Wasserman. *All of nonparametric statistics*. Springer Science & Business Media, 2006. [MR2172729](#)
- G. S. Watson. Smooth regression analysis. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 359–372, 1964. [MR0185765](#)
- C. F. J. Wu. Post-Fisherian experimentation: From physical to virtual. *Journal of the American Statistical Association*, 110(510):612–620, 2015. [MR3367251](#)
- M. You, E. Byon, J. J. Jin, and G. Lee. When wind travels through turbines: A new statistical approach for characterizing heterogeneous wake effects in multi-turbine wind farms. *IISE Transactions*, 49(1):84–95, 2017.
- P. Zhang. Nonparametric importance sampling. *Journal of the American Statistical Association*, 91(435):1245–1253, 1996. [MR1424622](#)