Proceedings of the 2024 Winter Simulation Conference
H. Lam, E. Azar, D. Batur, S. Gao, W. Xie, S. R. Hunter, and M. D. Rossetti, eds.

AN EFFICIENT FINITE-DIFFERENCE APPROXIMATION

Guo Liang[1], Guangwu Liu[2], and Kun Zhang[1]

[1]Institute of Statistics and Big Data, Renmin University of China, Beijing, CHINA
[2]Dept. of Mgmt. Sciences, City University of Hong Kong, Hong Kong SAR, CHINA

ABSTRACT

Estimating stochastic gradients is pivotal in fields like service systems in operations research. The classical method for this estimation is the finite-difference approximation, which entails generating samples at perturbed inputs. Nonetheless, practical challenges persist in determining the perturbation and obtaining an optimal finite-difference estimator with the smallest mean squared error (MSE). To tackle this problem, we propose a double sample-recycling approach in this paper. Firstly, pilot samples are recycled to estimate the optimal perturbation. Secondly, recycling these pilot samples and generating new samples at the estimated perturbation lead to an efficient finite-difference estimator. In numerical experiments, we apply the estimator in two examples, and numerical results demonstrate its robustness, as well as coincidence with the theory presented, especially in the case of small sample sizes.

1 INTRODUCTION

Stochastic gradient estimation involves estimating a function's derivative. In realistic scenarios, the exact function expression is often unknown, and one can only approximate it through stochastic simulation. Such setting is often regarded as black-box or zeroth-order (Ghadimi and Lan 2013; Nesterov and Spokoiny 2015). Stochastic gradient estimation finds significant applications in stochastic optimization, machine learning, financial engineering and operations research. In this paper, we investigate the finite-difference method, which is straightforward and effective for estimating the gradient (Asmussen and Glynn 2007; Fu 2006; Glasserman 2013; L'Ecuyer 1991).

Finite-difference methods include forward, backward, and central finite-difference (CFD) methods. In the finite-difference mechanism, the choice of perturbation size controls the balance between bias and variance. Typically, if the perturbation is too small, the variance of the estimated quantity will be too large, and if the perturbation is too large, the bias will be unacceptable. Using a metric of minimizing the MSE of the estimator, Fox and Glynn (1989) and Zazanis and Suri (1993) studied this trade-off and demonstrated that without using common random numbers (CRNs), the optimal order of perturbation size is $n^{-1/4}$ for forward and backward finite-difference methods and $n^{-1/6}$ for the CFD method, where n is the total number of samples. Glasserman (2013) suggested the extrapolation approach to further reduce the bias. A linear combination of the CFD output values under two different perturbations is employed to produce an estimator with decreased bias. When CRNs are not employed, the method's optimal perturbation is of order $n^{-1/10}$.

Although the order of the optimal perturbation is known for the finite-difference methods, the constants preceding the order are especially essential when the sample size is small (Li and Lam 2020; Lam et al. 2022). However, obtaining information about the constants before the order is challenging, even more difficult than estimating gradients because they contain higher-order derivatives and noises. Li and Lam (2020) proposed a two-stage approach to address this issue. In the first stage, they utilized regression to estimate the higher-order derivative and sample variance to estimate the function's noise. However, the noise estimation is somewhat coarse. Lam et al. (2022) enhanced the standard estimator regarding the choice of perturbation subject to the ambiguity of the model characteristics.

In this paper, we focus on enhancing the performance of the finite-difference method. We propose a double sample-recycling (DSR) approach to the finite-difference estimator for stochastic gradient estimation to more accurately estimate the unknown constants and effectively utilize the samples. Our work also proposes a procedure in two stages and focuses on sample-recycling. Specifically, we apply two sample-recycling techniques in the two stages to obtain an accurate estimator of perturbation and a finite-difference estimator. We support our theory with empirical results.

The rest of this paper is organized as follows. In Section 2, we give some background on how to use the CFD method for gradient estimation. In Section 3, we introduce the bootstrap method, apply it to gradient estimation and provide theoretical guarantees for it. Section 4 guides the process of reusing samples and provides theoretical evidence for the effectiveness of this operation. In Section 5, we give some experiments to verify the theoretical results, followed by conclusions in Section 6.

2 BACKGROUND

In this section, we introduce the background of the finite-difference method for stochastic gradient estimation. More details can be found in Asmussen and Glynn (2007), Fu (2006), Glasserman (2013). Consider a model that depends on a single parameter θ, where θ varies within some range $\Theta \subset \mathbb{R}$. Denote that $\alpha(\cdot)$ is a performance measure of interest and assume that we can only estimate it by simulation. Within the simulation for any chosen $\theta \in \Theta$, each trail gives an unbiased but noisy estimate of $\alpha(\theta)$, denoted by $Y(\theta)$, i.e., $\alpha(\theta) = \mathbb{E}[Y(\theta)]$. Suppose that we do not use the CRNs, i.e., for $\theta_1 \neq \theta_2 \in \Theta$, $Y(\theta_1)$ and $Y(\theta_2)$ are independent. We would like to estimate the first-order derivative $\alpha'(\theta_0)$, where $\theta_0 \in \Theta$ is the point of interest. In this paper, we consider the CFD method which yields an improvement in the convergence rate of the bias compared to the forward (and backward) finite-difference method.

The CFD scheme is similar to the definition of the derivative and utilizes the information at the neighboring points on both sides of θ_0. That is, $\alpha'(\theta_0)$ is approximated by

$$\widetilde{\alpha}'_h(\theta_0) = \frac{\alpha(\theta_0+h) - \alpha(\theta_0-h)}{2h}, \tag{1}$$

where h is a perturbation parameter. As h tends to 0, $\widetilde{\alpha}'_h(\theta_0)$ tends to $\alpha'(\theta_0)$. Denote

$$\Delta(h) = \frac{Y(\theta_0+h) - Y(\theta_0-h)}{2h}. \tag{2}$$

Evidently, $\Delta(h)$ is a noisy estimate of $\alpha'(\theta_0)$ and this estimate is the output of the CFD method.

Specifically, the CFD scheme sets the perturbation parameter $h > 0$. At $\theta_0 + h$, the simulation is repeated independently, and n independently and identically distributed (i.i.d.) observations $\{Y_1(\theta_0+h), \ldots, Y_n(\theta_0+h)\}$ are obtained. Then, n i.i.d. observations $\{Y_1(\theta_0-h), \ldots, Y_n(\theta_0-h)\}$ are simulated at $\theta_0 - h$. Using the sample average, we construct a CFD estimator of $\alpha'(\theta_0)$, denoted by

$$\widehat{\Delta}_{n,h} = \frac{1}{n}\sum_{i=1}^{n} \Delta_i(h) = \frac{1}{n}\sum_{i=1}^{n} \frac{Y_i(\theta_0+h) - Y_i(\theta_0-h)}{2h}. \tag{3}$$

For any $i = 1, \ldots, n$, $Y_i(\theta_0+h)$ and $Y_i(\theta_0-h)$ are the unbiased estimators of $\alpha(\theta_0+h)$ and $\alpha(\theta_0-h)$, respectively, so $\Delta_i(h)$ and $\widehat{\Delta}_{n,h}$ are both the unbiased estimators of $\widetilde{\alpha}'_h(\theta_0)$. Let $Z_i(h)$ denote the zero-mean error associated with $\Delta_i(h)$, and finally let $\bar{Z}_{n,h} = \sum_{i=1}^{n} Z_i(h)/n$ be the zero-mean error associated with $\widehat{\Delta}_{n,h}$. Then, $\widehat{\Delta}_{n,h} = \widetilde{\alpha}'_h(\theta_0) + \bar{Z}_{n,h}$.

The choice of h is crucial because it controls the quality of $\widehat{\Delta}_{n,h}$. If h is too large, the expectation of $\widehat{\Delta}_{n,h}$, i.e., $\widetilde{\alpha}'_h(\theta_0)$ will be too far from $\alpha'(\theta_0)$. In other words, the bias of $\widehat{\Delta}_{n,h}$ will be unacceptable. If h is too small, the variance of $\widehat{\Delta}_{n,h}$ will explode due to the location of h in the denominator. To choose a proper h, we should perform a trade-off between the bias and variance of $\widehat{\Delta}_{n,h}$. To calculate the bias and variance elaborately, we make the following assumptions.

Assumption 2.1 $\alpha(\theta)$ is thrice continuously differentiable in a neighborhood of θ_0 and $\alpha^{(3)}(\theta_0) \neq 0$.

Assumption 2.2 The standard deviation of $Y(\theta)$, denoted by $\sigma(\theta)$, is continuous at θ_0 and $\sigma(\theta) > 0$.

Assumptions 2.1 and 2.2 can be found in many other works, such as Glasserman (2013). Under Assumption 2.1, according to the Taylor expansion, we have

$$\alpha(\theta_0 + h) = \alpha(\theta_0) + \alpha'(\theta_0)h + \frac{\alpha^{(2)}(\theta_0)}{2}h^2 + \frac{\alpha^{(3)}(\theta_0)}{6}h^3 + o(h^3),$$
$$\alpha(\theta_0 - h) = \alpha(\theta_0) - \alpha'(\theta_0)h + \frac{\alpha^{(2)}(\theta_0)}{2}h^2 - \frac{\alpha^{(3)}(\theta_0)}{6}h^3 + o(h^3).$$

Subtraction eliminates $\alpha^{(2)}(\theta_0)$, leaving

$$\widetilde{\alpha}'_h(\theta_0) = \alpha'(\theta_0) + Bh^2 + o(h^2),$$

where $B = \alpha^{(3)}(\theta_0)/6$. Therefore, the bias of $\widehat{\Delta}_{n,h}$ is

$$\text{Bias}\left[\widehat{\Delta}_{n,h}\right] = \mathbb{E}\widehat{\Delta}_{n,h} - \alpha'(\theta_0) = \widetilde{\alpha}'_h(\theta_0) - \alpha'(\theta_0) = Bh^2 + o(h^2). \quad (4)$$

From (3), the variance of $\widehat{\Delta}_{n,h}$ is

$$\text{Var}\left[\widehat{\Delta}_{n,h}\right] = \frac{1}{n}\text{Var}\left[\frac{Y(\theta_0+h)+Y(\theta_0-h)}{2h}\right] = \frac{\sigma^2(\theta_0+h)+\sigma^2(\theta_0-h)}{4nh^2}.$$

Under Assumption 2.2, we have $\sigma^2(\theta_0+h) = \sigma^2(\theta_0)+o(1)$ and $\sigma^2(\theta_0-h) = \sigma^2(\theta_0)+o(1)$. Therefore,

$$\text{Var}\left[\widehat{\Delta}_{n,h}\right] = \frac{\sigma^2(\theta_0)+o(1)}{2nh^2}. \quad (5)$$

Then the MSE of $\widehat{\Delta}_{n,h}$ is

$$\text{MSE}\left[\widehat{\Delta}_{n,h}\right] = \text{Bias}^2\left[\widehat{\Delta}_{n,h}\right] + \text{Var}\left[\widehat{\Delta}_{n,h}\right] = (B+o(1))^2 h^4 + \frac{\sigma^2(\theta_0)+o(1)}{2nh^2}. \quad (6)$$

It follows from (6) that h controls a trade-off between the bias and variance, and thus we should choose an appropriate h to minimize the MSE. According to the relationship between the arithmetic average and the geometric average, we get

$$B^2h^4 + \frac{\sigma^2(\theta_0)}{2nh^2} \geq 3\left(\frac{B^2\sigma^4(\theta_0)}{16n^2}\right)^{1/3},$$

where the equality holds if and only if $h^* = \left(\sigma^2(\theta_0)/(4nB^2)\right)^{1/6}$. The optimal bias, variance and MSE are thus

$$\begin{cases} \text{Bias}^* = \left(\frac{B\sigma^2(\theta_0)}{4n}\right)^{1/3} + o\left(n^{-1/3}\right), \\ \text{Var}^* = \left(\frac{B^2\sigma^4(\theta_0)}{2n^2}\right)^{1/3} + o\left(n^{-2/3}\right), \\ \text{MSE}^* = 3\left(\frac{B^2\sigma^4(\theta_0)}{16n^2}\right)^{1/3} + o\left(n^{-2/3}\right). \end{cases} \quad (7)$$

To obtain an optimal $\widehat{\Delta}_{n,h}$, we need to determine the value of h. Although we know that $h^* = Cn^{-1/6}$, where $C = \left(\sigma^2(\theta_0)/(4B^2)\right)^{1/6}$, the constant C is typically unknown. The definition of C indicates that it requires the estimation of some values depending on the model information, such as $\alpha^{(3)}(\theta_0)$ and $\sigma^2(\theta_0)$. However, estimating these values is more challenging than estimating the gradient because it involves the higher-order derivatives of $\alpha(\theta)$. Furthermore, experimental findings indicate that assigning a perturbation with a significant deviation from the true one leads to substantial discrepancies in the MSE of the CFD estimator. These motivate us to explore an approach to address this problem and achieve the optimal MSE.

3 SAMPLE-RECYCLING FOR PERTURBATION SELECTION

In this section, we propose a resampling procedure to estimate C, specifically B and $\sigma^2(\theta_0)$. Our inspiration comes from equations (4) and (5), which demonstrate that $\mathbb{E}\widehat{\Delta}_{n,h}$ and $\mathrm{Var}\left[\widehat{\Delta}_{n,h}\right]$ are linear with respect to (w.r.t.) h^2 and $1/h^2$, respectively. However, $\mathbb{E}\widehat{\Delta}_{n,h}$ and $\mathrm{Var}\left[\widehat{\Delta}_{n,h}\right]$ are typically unknown and we turn to seek their surrogates. A straightforward approach is to use the sample mean and variance, respectively. Nevertheless, for any perturbation h, we need many sample pairs (sometimes we also use "samples") to obtain these surrogates, which is unavailable due to the limitation of sample size. To address this limitation, we introduce a resampling technique, the bootstrap, to estimate the sample mean and variance. On the other hand, to improve the accuracy of estimating B and $\sigma^2(\theta_0)$, we select K different values of h, expanding (4) to encompass K linear equations correspondingly. Then least-squares regression is applied to derive the slope of linear equations, i.e., the constant B. Similar procedure is also used to obtain the estimate of $\sigma^2(\theta_0)$.

3.1 Bootstrap Sample Mean and Variance

The bootstrap technique is described as follows. For a fixed h, assume that n_b sample pairs have been generated, denoted by $\{(Y_i(\theta_0+h), Y_i(\theta_0-h)), i = 1,2,\ldots,n_b\}$. Correspondingly, as defined in (2), we have n_b outputs $\mathbf{\Delta}(h) \triangleq \{\Delta_i(h), i = 1,2,\ldots,n_b\}$. We pick with replacement n_b times from $\mathbf{\Delta}(h)$ independently and randomly and then get a group of bootstrap samples $\mathbf{\Delta}^*(h) = \{\Delta_i^*(h), i = 1,2,\ldots,n_b\}$. Using these bootstrap samples, the estimator of $\alpha'(\theta_0)$ is expressed as

$$\widehat{\Delta}^b_{n_b,h} = \frac{1}{n_b}\sum_{i=1}^{n_b}\Delta_i^*(h).$$

It is natural to expect that the bootstrap mean and variance of $\widehat{\Delta}^b_{n_b,h}$ will exhibit similar asymptotic properties as described in (4) and (5), respectively, when n_b is sufficiently large. For these, let $\mathbb{E}_*$ and Var_* denote the expectation and variance under the bootstrap probability measure $\mathbb{P}_*(\cdot) \triangleq \mathbb{P}(\cdot|\mathbf{\Delta}(h))$. Then the asymptotic properties of the bootstrap mean and variance of $\widehat{\Delta}^b_{n_b,h}$ are summarized in Theorem 1.

Theorem 1 Denote $\nu_4(h) = \mathbb{E}\left[Z_1(h)\right]^4$ and $\nu_4 = \lim_{h\to 0}\mathbb{E}\left[hZ_1(h)\right]^4$. If $\nu_4 < \infty$, $\alpha(\theta)$ is five times continuously differentiable in a neighborhood of θ_0 and $\alpha^{(5)}(\theta_0) \neq 0$, then under Assumptions 2.2, as $n_b \to \infty$,

$$\mathbb{E}_*\widehat{\Delta}^b_{n_b,h} = \alpha'(\theta_0) + Bh^2 + Dh^4 + o(h^4) + \bar{Z}_{n_b,h}, \tag{8}$$

$$\mathrm{Var}_*\left[\widehat{\Delta}^b_{n_b,h}\right] = \frac{(n_b-1)(\sigma^2(\theta_0)+o(1))}{2n_b^2h^2} + \phi(h), \tag{9}$$

where $D = \alpha^{(5)}(\theta_0)/120$ and $\phi(h)$ is a zero-mean error with

$$\mathrm{Var}[\phi(h)] = \frac{(n_b-1)^2}{n_b^4h^4}\left(\frac{\nu_4+o(1)}{n_b} - \frac{n_b-3}{n_b(n_b-1)}\frac{\sigma^4(\theta_0)+o(1)}{4}\right). \tag{10}$$

Theorem 1 demonstrates that $\mathbb{E}_*\widehat{\Delta}^b_{n_b,h}$ and $\mathrm{Var}_*\left[\widehat{\Delta}^b_{n_b,h}\right]$ inherit the asymptotic properties of $\mathbb{E}\widehat{\Delta}_{n,h}$ and $\mathrm{Var}\left[\widehat{\Delta}_{n,h}\right]$, respectively. That is, they are linear w.r.t. h^2 and $1/h^2$, respectively. As a result, we can estimate B and $\sigma^2(\theta_0)$ through these linear relationships, which we will present in the next subsection. Another optional approach is to bootstrap $\{(Y_i(\theta_0+h), i=1,2,...,n_b\}$ and $\{(Y_i(\theta_0-h), i=1,2,...,n_b\}$ separately. This decoupling operation yields conclusions similar to those in Theorem 1. For convenience, we directly bootstrap $\boldsymbol{\Delta}(h)$ in this paper.

3.2 Regression for Perturbation Selection

From Theorem 1, we can use the least-squares regression to derive the estimates of B and $\sigma^2(\theta_0)$. Specifically, we consider

$$\boldsymbol{Y}_e = \boldsymbol{X}_e\boldsymbol{\beta}_e + \mathscr{E}_e, \quad \boldsymbol{Y}_v = \boldsymbol{X}_v\boldsymbol{\beta}_v + \mathscr{E}_v, \tag{11}$$

where

$$\begin{gathered}\boldsymbol{Y}_e = \left[\mathbb{E}_*\widehat{\Delta}^b_{n_b,h_1},...,\mathbb{E}_*\widehat{\Delta}^b_{n_b,h_K}\right]^\top, \boldsymbol{X}_e = \begin{bmatrix} 1 & \dots & 1 \\ h_1^2 & \dots & h_K^2 \end{bmatrix}^\top, \\ \boldsymbol{\beta}_e = \left[\alpha'(\theta_0), B\right]^\top, \mathscr{E}_e = \left[Dh_1^4 + o(h_1^4) + \bar{Z}_{n_b,h_1},...,Dh_K^4 + o(h_K^4) + \bar{Z}_{n_b,h_K}\right]^\top,\end{gathered} \tag{12}$$

and

$$\begin{gathered}\boldsymbol{Y}_v = \left[\mathrm{Var}_*\left[\widehat{\Delta}^b_{n_b,h_1}\right],...,\mathrm{Var}_*\left[\widehat{\Delta}^b_{n_b,h_K}\right]\right]^\top, \ \boldsymbol{X}_v = \left[\frac{n_b-1}{2n_b^2h_1^2},...,\frac{n_b-1}{2n_b^2h_K^2}\right]^\top, \\ \boldsymbol{\beta}_v = \sigma^2(\theta_0), \ \mathscr{E}_v = \left[\frac{(n_b-1)o(1)}{2n_b^2h_1^2} + \phi(h_1),...,\frac{(n_b-1)o(1)}{2n_b^2h_K^2} + \phi(h_K)\right]^\top.\end{gathered} \tag{13}$$

Using the least-squares method, we derive the estimators of $\boldsymbol{\beta}_e$ and $\boldsymbol{\beta}_v$, denoted by $\widehat{\boldsymbol{\beta}}_e$ and $\widehat{\boldsymbol{\beta}}_v$, respectively, where

$$\widehat{\boldsymbol{\beta}}_e = \left(\boldsymbol{X}_e^\top\boldsymbol{X}_e\right)^{-1}\boldsymbol{X}_e^\top\boldsymbol{Y}_e, \quad \widehat{\boldsymbol{\beta}}_v = \left(\boldsymbol{X}_v^\top\boldsymbol{X}_v\right)^{-1}\boldsymbol{X}_v^\top\boldsymbol{Y}_v.$$

Let the first term of $\widehat{\boldsymbol{\beta}}_e$ be $\widehat{\alpha}'(\theta_0)$ and the second term be $\widehat{B}$. In the following, we show that $\widehat{\boldsymbol{\beta}}_e$ and $\widehat{\boldsymbol{\beta}}_v$ are consistent estimators of $\boldsymbol{\beta}_e$ and $\boldsymbol{\beta}_v$, respectively.

3.2.1 Consistency of $\widehat{B}$

The following theorem presents the consistency of $\widehat{B}$. The conditions about $h_k(k=1,...,K)$ in Theorem 2 ensure that the regression error vanishes fast enough, which is necessary for the consistency of regression estimators.

Theorem 2 Suppose that $\alpha(\theta)$ is five times continuously differentiable in a neighborhood of θ_0 and $\alpha^{(5)}(\theta_0) \neq 0$. For any $k=1,...,K(K\geq 2)$, denote $h_k = c_k n_b^\gamma (c_k \neq 0, \gamma < 0)$ and for any $j \neq k$, let $c_j \neq c_k$. Then, under Assumption 2.1,

$$\mathbb{E}\left[\widehat{B}\right] - B = H_K n_b^{2\gamma} + o\left(n_b^{2\gamma}\right), \quad \mathrm{Var}\left[\widehat{B}\right] = V_K\frac{\sigma^2(\theta_0)+o(1)}{2n_b^{1+6\gamma}}, \tag{14}$$

where H_K and V_K are constants depending on $c_1,...,c_K$. Specifically,

$$H_K = \frac{K\sum_{k=1}^K Dc_k^6 - \sum_{k=1}^K c_k^2 \sum_{k=1}^K Dc_k^4}{K\sum_{k=1}^K c_k^4 - \left(\sum_{k=1}^K c_k^2\right)^2}, \quad V_K = \frac{-K^2\sum_{k=1}^K c_k^2 + \left(\sum_{k=1}^K c_k^2\right)^2\sum_{k=1}^K 1/c_k^2}{\left(K\sum_{k=1}^K c_k^4 - \left(\sum_{k=1}^K c_k^2\right)^2\right)^2}.$$

In addition, if $-1/6<\gamma<0$, then $\widehat{B}\xrightarrow{p}B$, where $\xrightarrow{p}$ means convergence in probability.

Theorem 2 indicates that for all $h_k=c_kn_b^{\gamma}$, where $k=1,...,K$ and $c_k\neq 0$, as long as $-1/6<\gamma<0$, $\widehat{B}$ is a consistent estimator of B. Furthermore, from (14), the convergence rate of the bias of $\widehat{B}$ is $O\left(n_b^{2\gamma}\right)$, and that of the variance of $\widehat{B}$ is $O\left(n_b^{-1-6\gamma}\right)$.

Except for fixed values of $c_1,...,c_K$, they can be generated from a proper distribution $\mathscr{P}_0$. In this case, H_K and V_K are both random. Since $c_1,...,c_K$ are coefficients in the perturbations, they must not equal 0, and are typically bounded. For these purposes, there are many optional distributions, such as the truncated normal distribution, denoted by $\psi\left(\mu_0,\sigma_0^2,L,U\right)$. Specifically, for any $k=1,...,K$, let $c_k\overset{d}{=}X\mathbb{1}_{\{L\leq X\leq U\}}$, where $\overset{d}{=}$ represents the same distribution, $X\sim\mathscr{N}\left(\mu_0,\sigma_0^2\right)$, and $L>0$ and $U>0$ denote the lower bound and upper bound, respectively. Because $0<L\leq c_k\leq U$, H_K and V_K are bounded.

When we generate $c_1,...,c_K$ from $\mathscr{P}_0$, the bias of $\widehat{B}$ is

$$\text{Bias}\left[\widehat{B}\right]=\mathbb{E}\left[\mathbb{E}\left[\widehat{B}-B\middle|\mathscr{P}_0\right]\right]=\mathbb{E}[H_K]n_b^{2\gamma}+o\left(n_b^{2\gamma}\right), \tag{15}$$

where the last equality is because $c_1,...,c_K$ are independent with the pilot samples and the bias result in Theorem 2. If $c_k\sim\psi\left(\mu_0,\sigma_0^2,L,U\right)$ for $k=1,...,K$, then c_k's are bounded. Therefore, H_K and the constants in $o\left(n_b^{2\gamma}\right)$ are all bounded.

Likewise, the variance of $\widehat{B}$ is

$$\begin{aligned}\text{Var}\left[\widehat{B}\right]&=\text{Var}\left[\mathbb{E}\left[\widehat{B}\middle|\mathscr{P}_0\right]\right]+\mathbb{E}\left[\text{Var}\left[\widehat{B}\middle|\mathscr{P}_0\right]\right]\\&=\text{Var}\left[B+H_Kn_b^{2\gamma}+o\left(n_b^{2\gamma}\right)\right]+\mathbb{E}\left[V_K\frac{\sigma^2(\theta_0)+o(1)}{2n_b^{1+6\gamma}}\right]\\&=\text{Var}[H_K]n_b^{4\gamma}+o\left(n_b^{4\gamma}\right)+\mathbb{E}[V_K]\frac{\sigma^2(\theta_0)+o(1)}{2n_b^{1+6\gamma}}.\end{aligned} \tag{16}$$

Comparing (14) with (15) and (16), it becomes evident that when c_k's are generated from $\mathscr{P}_0$, the bias of $\widehat{B}$ still converges at a rate of $O\left(n_b^{2\gamma}\right)$, while the variance of $\widehat{B}$ is increased by $O\left(n_b^{4\gamma}\right)$. However, it is important to note that the convergence rate of the MSE of $\widehat{B}$ remains unchanged. Therefore, during the pilot stage, the h_k's selected based on the tradeoff between bias and variance in (14) maintain the same order as those selected using the tradeoff between bias and variance in (15) and (16).

Similarly to Theorem 2, we can also obtain the asymptotic property of $\widehat{\alpha}'(\theta_0)$, which is useful for the recycling of the pilot samples. Due to space limitations, we omit the description about it.

3.2.2 Consistency of $\widehat{\beta}_v$

Now we provide Theorem 3 which demonstrates the convergence rate of the bias and variance of $\widehat{\boldsymbol{\beta}}_v$.

Theorem 3 Suppose that $\sigma(\theta)$ is continuously differentiable at θ_0 and $\sigma'(\theta_0)$ is bounded. For any $k=1,...,K(K\geq 1)$, denote $h_k=c_kn_b^{\gamma}(c_k\neq 0,\gamma<0)$ and for any $j\neq k$, let $c_j\neq c_k$. Then, under Assumption 2.2,

$$\mathbb{E}\left[\widehat{\boldsymbol{\beta}}_v\right]-\boldsymbol{\beta}_v=H_K^{\dagger}n_b^{2\gamma}+o\left(n_b^{2\gamma}\right),\quad\text{Var}\left[\widehat{\boldsymbol{\beta}}_v\right]=V_K^{\dagger}\frac{4\nu_4(n_b-1)-\sigma^4(\theta_0)(n_b-3)}{n_b(n_b-1)}+o\left(\frac{1}{n_b}\right),$$

where $H_K^\dagger$ and $V_K^\dagger$ are constants depending on $c_1,...,c_K$. Specifically,

$$H_K^\dagger = \sigma'(\theta_0)^2 \left(\sum_{k=1}^K \frac{1}{c_k^4}\right)^{-1} \sum_{k=1}^K \frac{1}{c_k^2}, \quad V_K^\dagger = \left(\sum_{k=1}^K \frac{1}{c_k^4}\right)^{-2} \sum_{k=1}^K \frac{1}{c_k^8}.$$

Consequently, we have $\widehat{\boldsymbol{\beta}}_v \xrightarrow{p} \boldsymbol{\beta}_v$.

Theorem 3 indicates that for all $h_k = c_k n_b^\gamma$, where $k = 1,...,K$ and $c_k \neq 0$, as long as $\gamma < 0$, $\widehat{\boldsymbol{\beta}}_v$ is a consistent estimator of $\boldsymbol{\beta}_v$. The convergence rate of the variance of $\widehat{\boldsymbol{\beta}}_v$ is $O\left(n_b^{-1}\right)$, which is independent of γ and faster than that of $\widehat{B}$. Additionally, the convergence rate of the bias of $\widehat{\boldsymbol{\beta}}_v$ is $O\left(n_b^{2\gamma}\right)$, which is equal to that of $\widehat{B}$.

In light of $\mathrm{Var}\left[\widehat{\boldsymbol{\beta}}_v\right]$ unrelated to γ, we focus on (14) and choose a suitable γ to balance the bias and variance of $\widehat{B}$. Specifically, according to Theorem 2, we consider the following optimization problem:

$$\begin{aligned} &\underset{\gamma, c_1,...,c_K}{\text{minimize}} \quad \mathrm{MSE}\left[\widehat{B}\right] = H_K^2 n_b^{4\gamma} + V_K \frac{\sigma^2(\theta_0)}{2n_b^{1+6\gamma}}, \\ &\text{subject to} \quad -1/6 < \gamma < 0. \end{aligned} \tag{17}$$

Solving (17) gives $\gamma = -1/10$. On the other hand, if $c_1,...,c_K$ are generated from $\mathscr{P}_0$ randomly, the optimization problem comes from the combination of (15) and (16):

$$\begin{aligned} &\underset{\gamma, \mathscr{P}_0}{\text{minimize}} \quad \mathbb{E}[H_K^2] n_b^{4\gamma} + \mathrm{Var}[H_K] n_b^{4\gamma} + \mathbb{E}[V_K] \frac{\sigma^2(\theta_0)}{2n_b^{1+6\gamma}}, \\ &\text{subject to} \quad -1/6 < \gamma < 0. \end{aligned} \tag{18}$$

Solving (18) also gives $\gamma = -1/10$. That is, whether $c_1,...,c_K$ are fixed or not, the optimal selection of h_k's is of order $O\left(n_b^\gamma\right)$ with $\gamma = -1/10$.

A more exact setting about $\{c_1,...,c_K\}$ or $\mathscr{P}_0$ necessitates knowledge of D, which is more challenging than the estimation of $\alpha'(\theta_0)$. In this paper, we treat $\{c_1,...,c_K\}$ and $\mathscr{P}_0$ as the hyperparameters whose selection is beyond the scope of our discussion.

4 SAMPLE-RECYCLING FOR FINITE-DIFFERENCE OUTPUT

In Section 3, we have provided the estimators $\widehat{B}$ and $\widehat{\boldsymbol{\beta}}_v$ and shown that both of them are consistent under mild conditions. In this section, we construct the double sample-recycling CFD (DSR-CFD) estimator through the remaining $n_2 = n - Kn_b$ sample pairs as well as the pilot sample pairs. We have to consider two questions: (1) how to generate the n_2 sample pairs; and (2) how to recycle the pilot sample pairs.

To answer the first question, note that we hope to utilize the whole n sample pairs, including the pilot sample pairs and the n_2 sample pairs. Therefore, the desired CFD estimator should be constructed based on the perturbation $\widehat{h}_n = \left(\widehat{\boldsymbol{\beta}}_v/(4n\widehat{B}^2)\right)^{1/6}$.

To answer the second question, recall that our objective is to estimate $\alpha'(\theta_0)$ under the criteria of minimizing the MSE of its estimator. To reuse the pilot samples, we transform $\Delta(h_k)$ which are constructed based on the pilot samples, by adjusting their location and scale, so that after the transformation, they have the same expectation and variance as $\Delta\left(\widehat{h}_n\right)$. Note that when no confusion arises, the subscript i of $\Delta_i(\cdot)$ is omitted to denote a generic index $i = 1,...,n$.

Specifically, for any $j = 1,...,n_b$ and $k = 1,...,K$, we transform $\Delta_j(h_k)$ to

$$\left(\frac{\mathrm{Var}\left[\Delta\left(\widehat{h}_n\right)\right]}{\mathrm{Var}[\Delta(h_k)]}\right)^{1/2}\left[\Delta_j(h_k)-\mathbb{E}\Delta(h_k)\right]+\mathbb{E}\Delta\left(\widehat{h}_n\right). \tag{19}$$

Evidently, $[\Delta_j(h_k)-\mathbb{E}\Delta(h_k)]/\left(\mathrm{Var}[\Delta(h_k)]\right)^{1/2}$ in (19) represents a standardization with mean 0 and variance 1. Then (19) possesses the same expectation and variance as $\Delta\left(\widehat{h}_n\right)$. Therefore, a desired estimator for $\alpha'(\theta_0)$ can be formulated as:

$$\frac{n_2}{n}\frac{1}{n_2}\sum_{i=1}^{n_2}\Delta_i\left(\widehat{h}_n\right)+\frac{Kn_b}{n}\frac{1}{Kn_b}\sum_{k=1}^{K}\sum_{j=1}^{n_b}(19). \tag{20}$$

This is a weighted average of an estimator at $\widehat{h}_n$ constructed through the n_2 samples, and the one by the transformation from the pilot samples. It is reasonable to expect that under certain mild conditions, such as ensuring that the covariance of two terms in (20) converges faster than the variance of either term, the bias, variance and MSE of (20) all achieve their optimal values as in (7).

For any $j = 1,...,n_b$ and $k = 1,...,K$, we substitute the regression estimators into (19), and (19) is approximated by

$$\Delta_j^t(h_k)=\frac{|h_k|}{\left|\widehat{h}_n\right|}\left(\Delta_j(h_k)-\left(1,h_k^2\right)\widehat{\boldsymbol{\beta}}_e\right)+\left(1,\widehat{h}_n^2\right)\widehat{\boldsymbol{\beta}}_e. \tag{21}$$

Then, the *DSR-CFD estimator*, i.e., the estimated version of (20), is established as follows:

$$\widehat{\Delta}_{n,\widehat{h}_n}=\frac{1}{n}\left(\sum_{i=1}^{n_2}\Delta_i\left(\widehat{h}_n\right)+\sum_{k=1}^{K}\sum_{j=1}^{n_b}\Delta_j^t(h_k)\right). \tag{22}$$

Next, we analyze the theoretical properties of $\widehat{\Delta}_{n,\widehat{h}_n}$. The results are summarized in the following theorem. For the sake of convenience, we consider the non-random hyperparameters $c_1,...,c_K$ such that when n_b is given, $h_1,...,h_K$ are also given. The result is similar when considering $\mathscr{P}_0$ as the hyperparameter, i.e., $c_1,...,c_K$ are random.

Theorem 4 Suppose that $\alpha(\theta)$ is five times continuously differentiable in a neighborhood of θ_0 and $\alpha^{(5)}(\theta_0)\neq 0$. For any $k = 1,...,K(K\geq 2)$, let $h_k = c_k n_b^{-1/10}(c_k\neq 0)$ and for any $j\neq k$, $c_j\neq c_k$. If $n_b,n\to\infty$, then under Assumptions 2.1 and 2.2,

$$\mathbb{E}\left[\widehat{\Delta}_{n,\widehat{h}_n}\right]=\alpha'(\theta_0)+\left(\frac{B\sigma^2(\theta_0)}{4n}\right)^{1/3}+\left(\frac{4B^2}{\sigma^2(\theta_0)}\right)^{1/6}\sqrt{\frac{n_b}{n}}D\boldsymbol{c}^\top\boldsymbol{P}\boldsymbol{c}^{\mathbf{4}}n^{-1/3}+o\left(n^{-1/3}\right), \tag{23}$$

$$\mathrm{Var}\left[\widehat{\Delta}_{n,\widehat{h}_n}\right]=\left(\frac{B^2\sigma^4(\theta_0)}{2n^2}\right)^{1/3}+\left(\frac{B^2\sigma^4(\theta_0)}{2}\right)^{1/3}\frac{n_b}{n}\left[\|\,\mathrm{Diag}(\boldsymbol{c}^{-1})\boldsymbol{P}\boldsymbol{c}\,\|_2^2-K\right]n^{-2/3}+o\left(n^{-2/3}\right), \tag{24}$$

where $\boldsymbol{c}=[|c_1|,...,|c_K|]^\top$, $\boldsymbol{c}^4=[c_1^4,...,c_K^4]^\top$, $\boldsymbol{P}=\boldsymbol{I}-\boldsymbol{X}_e(\boldsymbol{X}_e^\top\boldsymbol{X}_e)^{-1}\boldsymbol{X}_e^\top$, $\boldsymbol{X}_e$ is defined as in (12), $\mathrm{Diag}(\boldsymbol{c}^{-1})=\mathrm{Diag}\left(\frac{1}{|c_1|},...,\frac{1}{|c_K|}\right)$, sign$(B)$ is the sign of B, and $\|\,\boldsymbol{v}\,\|_2=\left(\sum_{k=1}^K v_k^2\right)^{1/2}$ for any $\boldsymbol{v}\in\mathbb{R}^K$.

Comparing (23) and (24) with (7), the changes of bias, variance, and MSE in the DSR-CFD method compared to the optimal CFD method are obvious, which depend on $\boldsymbol{c}^\top\boldsymbol{P}\boldsymbol{c}^{\mathbf{4}}$ and $\|\,\mathrm{Diag}(\boldsymbol{c}^{-1})\boldsymbol{P}\boldsymbol{c}\,\|_2^2$ that

Algorithm 1: DSR-CFD method.

Input: The number of total sample pairs n, the number of perturbation parameters K, the number of pilot sample pairs at each perturbation n_b, the number of bootstrap I and initial perturbation generator $\mathscr{P}_0$.

Step 1. *Estimate the unknown constants B, $\sigma^2(\theta_0)$ and the optimal perturbation.*

1. Generate $c_1, c_2, \ldots, c_K \overset{i.i.d.}{\sim} \mathscr{P}_0$ and let $h_k = c_k n_b^{-1/10}$ for any $k = 1,2,\ldots,K$. Generate pilot samples and calculate $\Delta_j(h_k)$, where $j = 1,2,\ldots,n_b$ and $k = 1,2,\ldots,K$.
2. Perform bootstrap resampling I times for each set of samples at $h_1, h_2, \ldots, h_K$ and calculate

$$\bar{\Delta}^b_{n_b,h_k} = \frac{1}{I}\sum_{q=1}^{I}\widehat{\Delta}^b_{n_b,h_k}(q), \quad s^2_{n_b,h_k} = \frac{1}{I}\sum_{q=1}^{I}\left[\widehat{\Delta}^b_{n_b,h_k}(q) - \bar{\Delta}^b_{n_b,h_k}\right]^2.$$

3. For $k = 1,\ldots,K$, replace $\mathbb{E}_*\widehat{\Delta}^b_{n_b,h_k}$ and $\mathrm{Var}_*\left[\widehat{\Delta}^b_{n_b,h_k}\right]$ with $\bar{\Delta}^b_{n_b,h_k}$ and $s^2_{n_b,h_k}$, respectively. Perform (weighted) regression to obtain $\widehat{B}$ and $\widehat{\boldsymbol{\beta}}_v$. Set the perturbation as $\widehat{h}_n = \left(\widehat{\boldsymbol{\beta}}_v/(4n\widehat{B}^2)\right)^{1/6}$.

Step 2. *Construct the DSR-CFD estimator of $\alpha'(\theta_0)$.*

1. Reuse $\Delta_j(h_k)$, where $j = 1,2,\ldots,n_b$ and $k = 1,2,\ldots,K$. That is, calculate $\Delta^t_j(h_k) = \frac{|h_k|}{\left|\widehat{h}_n\right|}\left(\Delta_j(h_k) - \left(1,h_k^2\right)\widehat{\boldsymbol{\beta}}_e\right) + \left(1,\widehat{h}_n^2\right)\widehat{\boldsymbol{\beta}}_e$ for $j = 1,2,\ldots,n_b$, and $k = 1,2,\ldots,K$.
2. Denote the number of remaining sample pairs as $n_2 = n - Kn_b$. Generate n_2 sample pairs at $\widehat{h}_n$ and calculate $\left\{\Delta_1(\widehat{h}_n),\ldots,\Delta_{n_2}(\widehat{h}_n)\right\}$.
3. The DSR-CFD estimator

$$\widehat{\Delta}_{n,\widehat{h}_n} = \frac{1}{n}\left(\sum_{i=1}^{n_2}\Delta_i\left(\widehat{h}_n\right) + \sum_{k=1}^{K}\sum_{j=1}^{n_b}\Delta^t_j(h_k)\right).$$

Output: $\widehat{\Delta}_{n,\widehat{h}_n}$.

are typically manageable. Even more attractive is the fact that if B and $\boldsymbol{D}\boldsymbol{c}^\top\boldsymbol{P}\boldsymbol{c}^{\boldsymbol{4}}$ are of opposite signs, the squared bias of $\widehat{\Delta}_{n,\widehat{h}_n}$ may be smaller than that of the optimal CFD estimator (see (7)). In addition, when $c_1,\ldots,c_K$ are chosen such that $\|\,\mathrm{Diag}(\boldsymbol{c}^{-1})\boldsymbol{P}\boldsymbol{c}\,\|_2^2 \leq K$, the variance will be reduced.

The full implementation of DSR-CFD method is shown in Algorithm 1. In Section 5, we will present two numerical examples to validate the aforementioned properties of our proposed DSR-CFD estimator.

5 NUMERICAL EXPERIMENTS

In this section, we use two examples to check the performance of our method. Examples 5.1 is from Li and Lam (2020). Example 5.2 is the generic $M/M/1$ queueing system from Lam et al. (2022). All numerical results are based on 1000 replications.

We will compare the performance of our proposed method and the Estimation-Minimization CFD (EM-CFD) method proposed by Li and Lam (2020). In the first and second examples, we let the initial perturbation generators $\mathscr{P}_0 = \psi(0,0.1,0.01,\infty)$ and $\mathscr{P}_0 = \psi(0,1,0.1,\infty)$, respectively. The number of bootstrap $I = 1000$. During the first phase, we set the number of regression points as $K = 10$ or 20.

Example 5.1 We consider estimating the first order derivatives of the polynomial function $\alpha(\theta) = 1 - 6\theta + 36\theta^2 - 53\theta^3 + 22\theta^5$ at different θ_0. Assume that the observed variables at θ obey a normal distribution with mean $\alpha(\theta)$ and variance 0.05, i.e., $Y(\theta) \sim \mathcal{N}(\alpha(\theta), 0.05)$. From a simple calculation we obtain

$$\alpha'(\theta_0) = -6 + 72\theta_0 - 159\theta_0^2 + 110\theta_0^4, \quad B = -53 + 220\theta_0^2, \quad \sigma^2(\theta_0) = 0.05,$$

as the real parameters to measure the performance of our proposed method. During the implementation of each algorithm, we assume that the parameters $\alpha'(\theta_0)$, B and $\sigma^2(\theta_0)$ are unknown.

Table 1 presents a comparison of bias, variance, and MSE between the DSR-CFD estimator (denoted by Bias-DSR, Var-DSR and MSE-DSR, respectively) and the optimal CFD (Opt-CFD) estimator (denoted by Bias-Opt, Var-Opt and MSE-Opt, respectively) across various sample sizes and estimation points (corresponding to different problems) in Example 5.1. In Table 1, B and D are respectively the true values and Opt-CFD is meant to set B and $\sigma^2(\theta_0)$ as their true values, i.e., $B = -53 + 220\theta_0^2$ and $\sigma^2(\theta_0) = 0.05$. We take 50% of the total sample to estimate the unknown constants and set $K = 10$ for DSR-CFD method.

Table 1: Comparison of the bias, variance and MSE of the DSR-CFD and Opt-CFD estimators in Example 5.1.

Budget	x	B	D	Bias-DSR	Bias-Opt	Var-DSR	Var-Opt	MSE-DSR	MSE-Opt
10^3	0	−53	22	−0.1335	−0.0829	0.0117	0.0154	0.0295	0.0222
	0.2	−44.2	22	−0.1356	−0.0789	0.0115	0.0149	0.0299	0.0211
	0.4	−17.5	22	−0.1264	−0.0584	0.0105	0.0067	0.0265	0.0101
	0.6	26.2	22	0.0170	0.0693	0.0102	0.0101	0.0105	0.0149
	0.8	87.8	22	0.0658	0.1016	0.0163	0.0212	0.0206	0.0315
	1	167	22	0.0991	0.1294	0.0217	0.0323	0.0315	0.0490
10^4	0	−53	22	−0.0583	−0.0415	0.0021	0.0031	0.0055	0.0049
	0.2	−44.2	22	−0.0558	−0.0339	0.0021	0.0031	0.0052	0.0042
	0.4	−17.5	22	−0.0544	−0.0279	0.0017	0.0015	0.0047	0.0023
	0.6	26.2	22	0.0099	0.0343	0.0018	0.0021	0.0019	0.0033
	0.8	87.8	22	0.0340	0.0484	0.0033	0.0044	0.0044	0.0068
	1	167	22	0.0549	0.0596	0.0055	0.0069	0.0085	0.0104
10^5	0	−53	22	−0.0245	−0.0196	4.5424e-04	6.7200e-04	0.0011	0.0011
	0.2	−44.2	22	−0.0222	−0.0179	4.2904e-04	6.3269e-04	9.2243e-04	9.5076e-04
	0.4	−17.5	22	−0.0208	−0.0126	2.8572e-04	3.3534e-04	7.1882e-04	4.9412e-04
	0.6	26.2	22	0.0088	0.0150	3.4222e-04	4.2576e-04	4.1963e-04	6.4935e-04
	0.8	87.8	22	0.0200	0.0220	7.4610e-04	9.9544e-04	0.0011	0.0015
	1	167	22	0.0270	0.0277	9.8655e-04	0.0015	0.0017	0.0023

From Table 1, we have the following observations:

- Bias: The bias is influenced by the signs of B and D. If B and D have the same (opposite) signs, the absolute bias of the DSR-CFD method is smaller (larger) than that of the Opt-CFD method. This aligns with Theorem 4 because $\boldsymbol{c}^\top \boldsymbol{P}\boldsymbol{c}^{\mathbf{4}}$ is negative in our experiment. When B and D share the same (opposite) signs, B and $D\boldsymbol{c}^\top \boldsymbol{P}\boldsymbol{c}^{\mathbf{4}}$ exhibit opposite (same) signs.
- Variance: It is observed that as the sample size increases, the variance of the DSR-CFD estimator consistently outperforms that of the Opt-CFD estimator, aligning with Theorem 4. For instance, at $x = 1$ with a sample budget of 10^5, the variances for the DSR-CFD and Opt-CFD estimators are 9.8655×10^{-4} and 0.0015, respectively.
- MSE: It is important to note that the bias and variance of the DSR-CFD estimator consistently approximate and are smaller than those of the Opt-CFD estimator, respectively. Consequently, when employing MSE as the error criterion, the DSR-CFD estimator consistently performs nearly as well as (or even outperforms) the Opt-CFD estimator.

Example 5.2 We consider a generic $M/M/1$ queueing system which is empty initially. Both of the arrival distribution and the service distribution are the exponential distribution and the arrival rate and service rate are denoted by λ and μ, respectively. Observations are the averaged system time of the first N customers and we are interested in the gradient of this quantity w.r.t. the service rate. For this example, we consider two cases.

- *Case 1:* We consider a critically loaded system, where we set $\lambda = \mu = 4$ and $N = 10$. In this case, the true derivative is -0.2501, which is calculated using the likelihood ratio/score function method (Glynn 1990) with 10^6 simulation repetitions (Lam et al. 2022).
- *Case 2:* We consider a non-critically loaded queue, where we set $\lambda = 3$, $\mu = 5$ and $N = 10$. In this case, the true derivative is -0.1136.

Figure 1 illustrates the numerical results for different methods in Example 5.2 when the sample-pair size changes from $n = 60$ to $n = 1000$. We compare the following three methods: The conventional CFD (Con-CFD) method, EM-CFD method and DSR-CFD method. For the Con-CFD method, we do not know any information about the model in advance and set $B = 5$ and $\sigma^2(\theta_0) = 1$. When considering the EM-CFD method, we set $r = 0.1$ and generate $h_1, \ldots, h_{rn} \overset{i.i.d.}{\sim} \mathcal{N}\left(0, 1 \times (rn)^{-1/5}\right)$. When applying the DSR-CFD method, we set $K = 20$ in the pilot phase and $r = 1$.

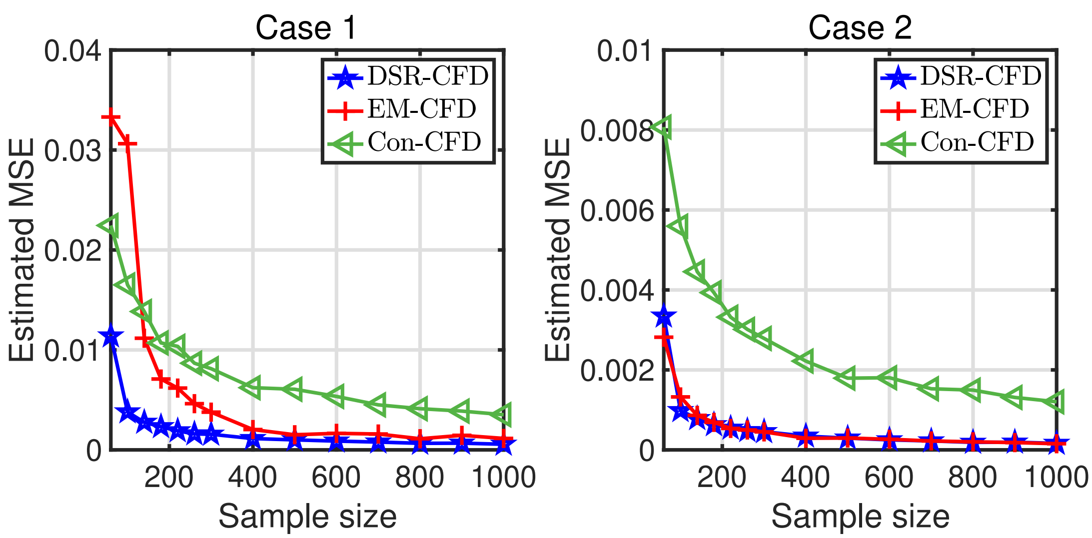

Figure 1: Comparison of the MSE among different estimators for cases 1 and 2 in Example 5.2.

Figure 1 illustrates that the performance of the Con-CFD estimator is inferior to that of the other two estimators. This observation emphasizes the importance of estimating unknown constants within finite-difference methods. In case 1 of Example 5.2, the DSR-CFD estimator consistently yields the most favorable results. For instance, with a sample-pair size of 60, the MSEs for the DSR-CFD and EM-CFD estimators are 0.011 and 0.033, respectively. In case 2, the performances of these two estimators are comparable. In conclusion, the DSR-CFD estimator performs well across various problems.

6 CONCLUSIONS

The choice of the perturbation significantly influences the accuracy of the finite-difference approximation in stochastic gradient estimation. Although the order of the optimal perturbation is known, determining the unknown constants remains a challenge. To address this issue and construct an efficient finite-difference approximation, in this paper, we develop a double sample-recycling approach, which utilizes samples efficiently. We provide complete theoretical analyses, particular algorithms, and numerical experiments. The theories and the findings of the numerical experiments are consistent. In particular, theoretical analyses reveal that the proposed estimator has a reduced variance compared to the optimal finite-difference estimator, and in some cases, a decrease in bias.

ACKNOWLEDGMENTS

The research of the first and third authors was supported by National Natural Science Foundation of China (NNSFC) grants 72101260. The research of the second author was supported partially by the NNSFC and the Research Grants Council of Hong Kong (RGC-HK), under the RGC-HK General Research Fund Project 11508620, and NSFC/RGC-HK Joint Research Scheme under project N_CityU 105/21.

REFERENCES

Asmussen, S. and P. W. Glynn. 2007. *Stochastic Simulation: Algorithms and Analysis*. Springer.

Fox, B. L. and P. W. Glynn. 1989. "Replication Schemes for Limiting Expectations". *Probability in the Engineering and Informational Sciences* 3(3):299–318.

Fu, M. C. 2006. "Gradient Estimation". *Handbooks in Operations Research and Management Science*.

Ghadimi, S. and G. Lan. 2013. "Stochastic First- and Zeroth-Order Methods for Nonconvex Stochastic Programming". *SIAM Journal on Optimization* 23(4):2341–2368.

Glasserman, P. 2013. *Monte Carlo Methods in Financial Engineering*. Springer Science & Business Media.

Glynn, P. W. 1990. "Likelihood Ratio Gradient Estimation for Stochastic Systems". *Communications of the ACM* 33(10):75–84.

Lam, H., X. Zhang, and X. Zhang. 2022. "Enhanced Balancing of Bias-Variance Tradeoff in Stochastic Estimation: A Minimax Perspective". *Operations Research*.

L'Ecuyer, P. 1991. "An Overview of Derivative Estimation". In *1991 Winter Simulation Conference Proceedings.* IEEE.

Li, H. and H. Lam. 2020. "Optimally Tuning Finite-Difference Estimators". In *2020 Winter Simulation Conference (WSC)*. IEEE.

Nesterov, Y. and V. Spokoiny. 2015. "Random Gradient-Free Minimization of Convex Functions". *Foundations of Computational Mathematics* 17(2):527–566.

Zazanis, M. A. and R. Suri. 1993. "Convergence Rates of Finite-Difference Sensitivity Estimates for Stochastic Systems". *Operations Research* 41(4):694–703.

AUTHOR BIOGRAPHIES

GUO LIANG is a Ph.D. Candidate in the Institute of Statistics and Big Data at Renmin University of China. His research interests include stochastic simulation, stochastic gradient estimation, financial engineering and risk management. His email address is liangguo000221@ruc.edu.cn.

GUANGWU LIU is a Professor in College of Business at City University of Hong Kong. His research interests include stochastic simulation, financial engineering, and risk management. His e-mail address is msgw.liu@cityu.edu.hk and his website is https://www.cb.cityu.edu.hk/staff/guanliu/.

KUN ZHANG is an Assistant Professor in the Institute of Statistics and Big Data at Renmin University of China. He holds a PhD in management science from City University of Hong Kong. His research interests include simulation optimization, machine learning, business analytics, financial engineering and risk management. His email address is kunzhang@ruc.edu.cn.