

# Debiasing Conditional Stochastic Optimization

Lie He\*  
EPFL  
lie.he@epfl.ch

Shiva Prasad Kasiviswanathan  
Amazon  
kasivisw@gmail.com

## Abstract

In this paper, we study the conditional stochastic optimization (CSO) problem which covers a variety of applications including portfolio selection, reinforcement learning, robust learning, causal inference, etc. The sample-averaged gradient of the CSO objective is biased due to its nested structure, and therefore requires a high sample complexity to reach convergence. We introduce a general *stochastic extrapolation* technique that effectively reduces the bias. We show that for nonconvex smooth objectives, combining this extrapolation with variance reduction techniques can achieve a significantly better sample complexity than existing bounds. Additionally, we develop new algorithms for the finite-sum variant of the CSO problem that also significantly improve upon existing results. Finally, we believe that our debiasing technique has the potential to be a useful tool for addressing similar challenges in other stochastic optimization problems.

---

\*Work initiated during internship at Amazon.

# Contents

1	Introduction	3
2	Stochastic Extrapolation as a Tool for Bias Correction	5
3	Applying Stochastic Extrapolation in the CSO Problem	7
4	Applying Stochastic Extrapolation in the FCCO Problem	10
5	Applications	11
6	Concluding Remarks	12
A	Missing Details from Section 1	16
A.1	Other Related Work . . . . .	16
B	Missing Details from Section 2	16
C	Stationary Point Convergence Proofs from Section 3 (CSO)	20
C.1	Helpful Lemmas . . . . .	20
C.2	Convergence of BSGD . . . . .	22
C.3	Convergence of E-BSGD . . . . .	24
C.4	Convergence of BSpiderBoost . . . . .	26
C.5	Convergence of E-BSpiderBoost . . . . .	29
D	Stationary Point Convergence Proofs from Section 4 (FCCO)	31
D.1	E-BSpiderBoost for FCCO problem . . . . .	31
D.2	Convergence of NestedVR . . . . .	32
D.3	Convergence of E-NestedVR . . . . .	42
E	Missing Details from Section 5	46
E.1	Application of Instrumental Variable Regression . . . . .	46
E.2	Application of First-order MAML . . . . .	47
E.3	Application of Deep Average Precision Maximization . . . . .	47

# 1 Introduction

In this paper, we investigate the *conditional stochastic optimization* (CSO) problem as presented by Hu et al. [2020b], which is formulated as follows:

$$\min_{\mathbf{x} \in \mathbb{R}^d} F(\mathbf{x}) = \mathbb{E}_{\xi} [f_{\xi}(\mathbb{E}_{\eta|\xi} [g_{\eta}(\mathbf{x}; \xi)])], \quad (\text{CSO})$$

where  $\xi$  and  $\eta$  represent two random variables, with  $\eta$  conditioned on  $\xi$ . Additionally,  $f_{\xi} : \mathbb{R}^p \rightarrow \mathbb{R}$  and  $g_{\eta} : \mathbb{R}^d \rightarrow \mathbb{R}^p$  denote stochastic functions. The inner expectation is calculated with respect to the conditional distribution of  $\eta|\xi$ . Throughout this paper, in line with the established CSO framework [Hu et al., 2020b,a], we assume access to samples from the distribution  $\mathbb{P}(\xi)$  and the conditional distribution  $\mathbb{P}(\eta|\xi)$ .

Many machine learning tasks can be formulated as a CSO problem, e.g., policy evaluation and control in reinforcement learning [Dai et al., 2018; Nachum and Dai, 2020], and linearly-solvable Markov decision process [Dai et al., 2017]. Other examples of the CSO problem include instrumental variable regression [Muandet et al., 2020] and invariant learning [Hu et al., 2020b]. Moreover, the widely-used Model-Agnostic Meta-Learning (MAML) framework, which seeks to determine a meta-initialization parameter using metadata for related learning tasks that are trained through gradient-based algorithms, is another example of a CSO problem. In this context, tasks  $\xi$  are drawn randomly, followed by the drawing of samples  $\eta|\xi$  from the specified task [Finn et al., 2017]. It is noteworthy that the standard stochastic optimization problem  $\min_{\mathbf{x}} \mathbb{E}_{\xi} [f_{\xi}(\mathbf{x})]$  represents a degenerate case of the CSO problem, which is obtained by setting  $g_{\eta}$  as an identity function.

In numerous prevalent CSO problems, such as first-order MAML (FO-MAML) [Finn et al., 2017], the outer random variable  $\xi$  only takes value in a finite set (say in  $\{1, \dots, n\}$ ). These problems can be reformulated to have a finite-sum structure in the outer loop and referred to as *Finite-sum Coupled Compositional Optimization* (FCCO) problem in [Wang and Yang, 2022; Jiang et al., 2022]. In this paper, we also study this problem, formulated as:

$$\min_{\mathbf{x} \in \mathbb{R}^d} F_n(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n f_i(\mathbb{E}_{\eta|i} [g_{\eta}(\mathbf{x}; i)]). \quad (\text{FCCO})$$

The FCCO problem also has broad applications in machine learning for optimizing average precision, listwise ranking losses, neighborhood component analysis, deep survival analysis, deep latent variable models [Wang and Yang, 2022; Jiang et al., 2022].

Although the CSO and FCCO problems are widespread, they present challenges for optimization algorithms. Based on the special composition structure of CSO, using chain rule, under mild conditions, the full gradient of CSO is given by

$$\nabla F(\mathbf{x}) = \mathbb{E}_{\xi} \left[ \left( \mathbb{E}_{\eta|\xi} [\nabla g_{\eta}(\mathbf{x}; \xi)] \right)^{\top} \nabla f_{\xi}(\mathbb{E}_{\eta|\xi} [g_{\eta}(\mathbf{x}; \xi)]) \right].$$

Constructing an unbiased stochastic estimator for the gradient is generally computationally expensive (and even impossible). A straightforward estimation of  $\nabla F(\mathbf{x})$  is to estimate  $\mathbb{E}_{\xi}$  with 1 sample of  $\xi$ , estimate  $\mathbb{E}_{\eta|\xi} [g_{\eta}(\cdot)]$  with a set  $H_{\xi}$  of  $m$  independent and identically distributed (i.i.d.) samples drawn from the conditional distribution  $\mathbb{P}(\eta|\xi)$ , and  $\mathbb{E}_{\eta|\xi} [\nabla g_{\eta}(\cdot)]$  with a different set  $\tilde{H}_{\xi}$  of  $m$  i.i.d. samples drawn from the same conditional distribution, i.e.,

$$\nabla \hat{F}_m(\mathbf{x}) := \left( \frac{1}{m} \sum_{\eta \in H_{\xi}} \nabla g_{\eta}(\mathbf{x}; \xi) \right)^{\top} \nabla f_{\xi} \left( \frac{1}{m} \sum_{\eta \in \tilde{H}_{\xi}} g_{\eta}(\mathbf{x}; \xi) \right). \quad (1)$$

Note that  $\nabla \hat{F}_m(\mathbf{x})$  consists of two terms. The first term  $(1/m) \sum_{\eta \in H_{\xi}} \nabla g_{\eta}(\mathbf{x}; \xi)$  is an unbiased estimate of  $\mathbb{E}_{\eta|\xi} [\nabla g_{\eta}(\mathbf{x}; \xi)]$ . However, the second term is biased, i.e.,

$$\mathbb{E}_{\eta|\xi} [\nabla f_{\xi} \left( \frac{1}{m} \sum_{\eta \in \tilde{H}_{\xi}} g_{\eta}(\mathbf{x}; \xi) \right)] \neq \nabla f_{\xi}(\mathbb{E}_{\eta|\xi} [g_{\eta}(\mathbf{x}; \xi)]).$$

Consequently,  $\nabla \hat{F}_m(\mathbf{x})$  is a biased estimator of  $\nabla F(\mathbf{x})$ . To reach the  $\varepsilon$ -stationary point of  $F(\mathbf{x})$  (Definition 1), the bias has to be sufficiently small.

Optimization with biased gradients converges only to a neighborhood of the stationary point. While the bias diminishes with increasing batch size, it also introduces additional sample complexity. For nonconvex objectives, Biased Stochastic Gradient Descent (BSGD) requires a total sample complexity of  $O(\varepsilon^{-6})$  to reach an  $\varepsilon$ -stationary point [Hu et al., 2020b]. This contrasts with standard stochastic optimization, where

Problem	Old Bounds		Our Bounds	
	Algorithm	Bound	Algorithm	Bound
CSO	BSGD [Hu et al., 2020b]	$O(\varepsilon^{-6})$	E-BSGD	$O(\varepsilon^{-4.5})$
CSO	BSpiderBoost [Hu et al., 2020b] RT-MLMC [Hu et al., 2021]	$O(\varepsilon^{-5})$ $O(\varepsilon^{-4})$	E-BSpiderBoost	$O(\varepsilon^{-3.5})$
FCCO	MSVR-V2 [Jiang et al., 2022]	$O(n\varepsilon^{-3})$	E-NestedVR	$\begin{cases} O(n\varepsilon^{-3}) & \text{if } n \leq \varepsilon^{2/3} \\ O(\max\{f_{\frac{\rho}{\varepsilon^{2.5}}, \frac{\rho}{n}}, \frac{1}{n\varepsilon^4}g\}) & \text{if } n > \varepsilon^{2/3} \end{cases}$

Table 1: Sample complexities needed to reach  $\varepsilon$ -stationary point for FCCO and CSO problems with nonconvex smooth objectives. Assumptions are comparable, but our results require an additional mild regularity on  $f_\xi$  and  $g_\eta$ . For FCCO also see Footnote 1. Note that  $\Omega(\varepsilon^{-3})$  is a sample complexity lower bound for standard stochastic nonconvex optimization [Arjevani et al., 2022], and hence, also for the problems considered in this paper.

sample-averaged gradients are unbiased with a sample complexity of  $O(\varepsilon^{-4})$  [Ghadimi and Lan, 2013; Arjevani et al., 2022]. This discrepancy has spurred a multitude of proposals aimed at reducing the sample complexities of both CSO and FCCO problems. Hu et al. [2020b] introduced Biased SpiderBoost (BSpiderBoost), which, based on the variance reduction technique SpiderBoost from Wang et al. [2019], reduces the variance of  $\xi$  to achieve a sample complexity of  $O(\varepsilon^{-5})$  for the CSO problem. Hu et al. [2021] proposed multi-level Monte Carlo (MLMC) gradient methods V-MLMC and RT-MLMC to further enhance the sample complexity to  $O(\varepsilon^{-4})$ . The SOX [Wang and Yang, 2022] and MSVR-V2 [Jiang et al., 2022] algorithms concentrated on the FCCO problem and improved the sample complexity to  $O(n\varepsilon^{-4})$  and  $O(n\varepsilon^{-3})$ , respectively.

**Our Contributions.** In this paper, we improve the sample complexities for both the CSO and FCCO problems (see Table 1). To facilitate a clear and concise presentation, we will suppress the dependence on specific problem parameters throughout the ensuing discussion.

- (a) Our main technical tool in this paper is an extrapolation-based scheme to mitigate bias in gradient estimations. Considering a suitably differentiable function  $q(\cdot)$  and a random variable  $\delta \in D$ , we show that we can approximate the value of  $q(\mathbb{E}[\delta])$  via extrapolation from a limited number of evaluations of  $q(\delta)$ , while maintaining a minimal bias. Within the context of CSO and FCCO problems, this scheme is used in gradient estimation, where the function  $q$  corresponds to  $r f_\xi$  and the random variable  $\delta$  corresponds to  $g_\eta$ .
- (b) For the CSO problem, we present novel algorithms that integrate this above extrapolation-based scheme with BSGD and BSpiderBoost algorithms of Hu et al. [2020b]. Our algorithms, referred to as E-BSGD and E-BSpiderBoost, achieve a sample complexity of  $O(\varepsilon^{-4.5})$  and  $O(\varepsilon^{-3.5})$  respectively, in order to attain an  $\varepsilon$ -stationary point for nonconvex smooth objectives. Notably, the sample complexity of E-BSpiderBoost improves the best-known sample complexity of  $O(n\varepsilon^{-4})$  for the CSO problem from Hu et al. [2021].
- (c) For the FCCO problem<sup>1</sup> we propose a new algorithm that again combines the extrapolation-based scheme with a multi-level variance reduction applied to both inner and outer parts of the problem. Our algorithm, referred to as E-NestedVR, achieves a sample complexity of  $O(n\varepsilon^{-3})$  if  $n \leq \varepsilon^{2/3}$  and  $O(\max\{f_{\frac{\rho}{\varepsilon^{2.5}}, \frac{\rho}{n}}, \frac{1}{n\varepsilon^4}g\})$  if  $n > \varepsilon^{2/3}$  for nonconvex smooth objectives and second-order extrapolation scheme. Our bound is never worse than the  $O(n\varepsilon^{-3})$  bound of MSVR-V2 algorithm of Jiang et al. [2022] and is in fact better if  $n = \Omega(\varepsilon^{2/3})$ . As an illustration, when  $n = \Theta(\varepsilon^{1.5})$ , our bound of  $O(\varepsilon^{-3.25})$  is significantly better than the MSVR-V2 bound of  $O(\varepsilon^{-4.5})$ .

In terms of proof techniques, our approach diverges from conventional analyses for the CSO and FCCO problems in that we focus on explicitly bounding the bias and variance terms of the gradient estimator to establish the convergence guarantee. Compared to previous results, our improvements do require an additional mild regularity assumption on  $f_\xi$  and  $g_\eta$  mainly that  $r f_\xi$  is 3rd order differentiable. Firstly, as we discuss in Remark 1 most common instantiations of CSO/FCCO framework such as: 1) invariant logistic regression Hu et al. [2020b], 2) instrumental variable regression [Muandet et al., 2020], 3) first-order MAML for sine-wave

<sup>1</sup> For the FCCO problem we focus on  $n = O(\varepsilon^{-2})$  case, for  $n = \Omega(\varepsilon^{-2})$  we can just treat the FCCO problem as a CSO problem and get an  $O(\varepsilon^{-3.5})$  sample complexity bound via our E-BSpiderBoost algorithm.

few shot regression [Finn et al., 2017] and other problems, 4) deep average precision maximization [Qi et al., 2021a; Wang et al., 2022a], tend to satisfy this assumption. Secondly, we would like to highlight that the bounds derived from previous studies do not improve when incorporating this additional regularity assumption. Our results show that, this regularity assumption, which seems to be practically valid, can be exploited through a novel extrapolation-based bias reduction technique to provide substantial improvements in sample complexity.<sup>2</sup>

We defer some additional related work to Appendix A and conclude with some preliminaries. **Notation.** Vectors are denoted by boldface letters. For a vector  $\mathbf{x}$ ,  $\|\mathbf{x}\|_2$  denotes its  $\ell_2$ -norm. A function with  $k$  continuous derivatives is called a  $C^k$  function. We use  $a \asymp b$  to denote that  $a = Cb$  for some constant  $C > 0$ . We consider expectation over various randomness:  $\mathbb{E}_\xi[\cdot]$  denotes expectation over the random variable  $\xi$ ,  $\mathbb{E}_{\eta|\xi}[\cdot]$  denotes expectation over the conditional distribution of  $\eta|\xi$ . Unless otherwise specified, for a random variable  $X$ ,  $\mathbb{E}[X]$  denotes expectation over the randomness in  $X$ . We focus on nonconvex objectives in this paper and use the following standard convergence criterion for nonconvex optimization [Jain et al., 2017].

**Definition 1 ( $\varepsilon$ -stationary point)** For a differentiable function  $F(\cdot)$ , we say that  $\mathbf{x}$  is a first-order  $\varepsilon$ -stationary point if  $\|\nabla F(\mathbf{x})\|_2 \leq \varepsilon$ .

For notational convenience, in the rest of this paper, we omit the dependence on  $\xi$  (or  $i$  in the FCCO context) in the function  $g$  and use  $g_\eta(\mathbf{x})$  to represent  $g_\eta(\mathbf{x}; \xi)$ .

## 2 Stochastic Extrapolation as a Tool for Bias Correction

In this section, we present an approach for tackling the bias problem as appears in optimization procedures such as BSGD, BSpiderBoost, etc. Importantly, our approach addresses a general problem appearing in optimization settings and could be of independent interest. All missing details from this section are presented in Appendix B.

For ease of presentation, we start by considering the 1-dimensional case and assume a function  $q: \mathbb{R} \rightarrow \mathbb{R}$ , a constant  $s \geq \mathbb{R}$ . Let  $\delta$  be a random variable drawn from an arbitrary distribution  $D$  over  $\mathbb{R}$ . In Sections 3 and 4, we apply these ideas to the CSO and FCCO problems where the random variable  $\delta$  is played by  $g_\eta(\cdot)$  and function  $q$  is played by  $\nabla f_\xi$ . Informally stated, our goal in this section will be to

Efficiently approximate  $q(s + \mathbb{E}[\delta])$  with few evaluations of  $f(q(s + \delta)g_\delta \sim D$ .

An interesting case is when  $s = 0$ , where we are approximating  $q(\mathbb{E}[\delta])$  with evaluations of  $f(q(\delta)g_\delta \sim D$ . Now, if  $q$  is an affine function, then  $q(s + \mathbb{E}[\delta]) = \mathbb{E}[q(s + \delta)]$ . However, the equality does not hold true for general  $q$ , and there exists a bias, i.e.,  $\mathbb{E}[q(s + \mathbb{E}[\delta])] - \mathbb{E}[q(s + \delta)] > 0$ . In this section, we introduce a stochastic extrapolation-based method, where we use an affine combination of biased stochastic estimates, to achieve better approximation.

Suppose  $q \in C^{2k-1}$  is a continuous differentiable up to  $(2k-1)$ th derivative and let  $h = \mathbb{E}[\delta]$ . We expand  $q(s + \delta)$ , the most straightforward approximation of  $q(s + \mathbb{E}[\delta])$ , using Taylor series at  $s + h$ , and take expectation,

$$\begin{aligned} \mathbb{E}[q(s + \delta)] &= q(s + h) + q'(s + h) \mathbb{E}[\delta - h] + \frac{q''(s + h)}{2} \mathbb{E}[(\delta - h)^2] + \frac{q^{(3)}(s + h)}{6} \mathbb{E}[(\delta - h)^3] \\ &+ \dots + \frac{q^{(2k-2)}(s + h)}{(2k-2)!} \mathbb{E}[(\delta - h)^{(2k-2)}] + \frac{q^{(2k-1)}(\phi_\delta)}{(2k-1)!} \mathbb{E}[(\delta - h)^{2k-1}], \end{aligned} \quad (2)$$

where  $\phi_\delta$  between  $s$  and  $s + \delta$ . While  $\mathbb{E}[q(s + \delta)]$  matches  $q(s + h)$  in the first 2 terms, the third term is no longer zero. The approximation error (bias) is

$$j \mathbb{E}[q(s + \delta)] - q(s + h)j = j \frac{q''(s + h)}{2} \mathbb{E}[(\delta - h)^2] + \dots + \frac{q^{(2k-1)}(\phi_\delta)}{(2k-1)!} \mathbb{E}[(\delta - h)^{2k-1}]j.$$

In order to analyze the upper bound, we make the following assumption on  $D$  and  $q$ .

<sup>2</sup>Higher-order smoothness conditions have also been exploited in standard stochastic optimization for performance gains [Bubeck et al., 2019].

Assumption 1 (Bounded moments) For all  $\delta \in D$  has bounded higher-order moments:  $\sigma_l := j\mathbb{E}[(\delta - \mathbb{E}[\delta])^l]j < 1$  for  $l = 2, 3, \dots, 2k - 1$ .

Assumption 2 (Bounded derivatives) The  $q \in C^{2k-1}$  and has bounded derivatives, i.e.,  $a_l := \sup_{s \in \text{dom}(q)} j q^{(l)}(s) j < 1$  for  $l = 1, 2, \dots, 2k - 1$ .

In addition, we consider a sample averaged distribution  $D_m$  derived from  $D$  as follows.

Definition 2 Given a distribution  $D$  satisfying Assumption 1 and  $m \in \mathbb{N}^+$ , we define the distribution  $D_m$  that outputs  $\delta$  where  $\delta = \frac{1}{m} \sum_{i=1}^m \delta_i$  with  $\delta_i \stackrel{i.i.d.}{\sim} D$ .

The moments of such distribution  $D_m$  decrease with batch size  $m$  as  $k \geq 2$ ,  $j\mathbb{E}[(\delta - \mathbb{E}[\delta])^k]j = O(m^{-dk/2e})$  (see Lemma 1). Our desiderata would be to construct a scheme that uses some samples from the distribution  $D_m$  to construct an approximation of  $q(s + \mathbb{E}[\delta])$  that satisfies the following requirement.

Definition 3 ( $k$ th-order Extrapolation Operator) Given a function  $q : \mathbb{R} \rightarrow \mathbb{R}$  satisfying Assumption 2 and distribution  $D_m$  satisfying Assumption 1, we define a  $k$ th-order extrapolation operator  $T_{D_m}^{(k)}$  as an operator from  $C^{2k-1} \rightarrow C^{2k-1}$  that given  $N = N(k)$  i.i.d. samples  $\delta_1, \dots, \delta_N$  from  $D_m$  satisfies  $\forall s \in \mathbb{R}$ :  $j\mathbb{E}[T_m^{(k)}q(s)] - q(s + \mathbb{E}[\delta])j = O(m^{-k})$ .

We now propose a sequence of operators  $L_{D_m}^{(1)}, L_{D_m}^{(2)}, L_{D_m}^{(3)}, \dots$  that satisfy the above definition. The  $L_{D_m}^{(k)}q(s)$  is designed to ensure its Taylor expansion at  $s + h$  has a form of  $q(s + h) + O(\mathbb{E}[(\delta - h)^{2k-1}])$ . The remainder  $O(\mathbb{E}[(\delta - h)^{2k-1}])$  is bounded by  $O(m^{-k})$  due to Lemma 1.

A First-order Extrapolation Operator. We define the simplest operator

$$L_{D_m}^{(1)}q : s \mapsto [q(s + \delta)] \quad \text{where } \delta \stackrel{i.i.d.}{\sim} D_m.$$

In Proposition 2 (Appendix B), we show that  $L_{D_m}^{(1)}$  is a first-order extrapolation operator.<sup>3</sup>

A Second-order Extrapolation Operator. We define the following linear operator  $L_{D_m}^{(2)}$  which transforms  $q \in C^3$  into  $L_{D_m}^{(2)}q$  which has lesser bias (but similar variance, as shown later).

Definition 4 ( $L_{D_m}^{(2)}$  Operator) Given  $D_m$  and  $q$ , define the following operator,

$$L_{D_m}^{(2)}q : s \mapsto \left[ 2q\left(s + \frac{\delta_1 + \delta_2}{2}\right) - \frac{q(s + \delta_1) + q(s + \delta_2)}{2} \right] \quad \text{where } \delta_1, \delta_2 \stackrel{i.i.d.}{\sim} D_m.$$

Note that  $\frac{\delta_1 + \delta_2}{2}$  is same as sampling from  $D_{2m}$ . The absolute difference in the Taylor expansion of  $L_{D_m}^{(2)}q$  at  $s + h$  differs from  $q(s + h)$  as,


$$O\left(\left|\mathbb{E}\left[2\left(\frac{\delta_1 + \delta_2}{2} - h\right)^3 - \frac{1}{2}\left((\delta_1 - h)^3 + (\delta_2 - h)^3\right)\right]\right|\right) = O(j\mathbb{E}[(\delta - h)^3]j) \text{ for } \delta \stackrel{i.i.d.}{\sim} D_m. \quad (3)$$

The bias error of this scheme can be bounded through the following proposition.


Proposition 1 (Second-order Guarantee) Assume that distribution  $D_m$  and  $q(\cdot)$  satisfies Assumption 1 and 2 respectively with  $k = 2$ . Then, for all  $s \in \mathbb{R}$ ,  $\left|\mathbb{E}\left[L_{D_m}^{(2)}q(s)\right] - q(s + \mathbb{E}[\delta])\right| \leq a_3\sigma_3/m^2$ .

The above lemma shows that  $L_{D_m}^{(2)}$  is in fact a second-order extrapolation operator with  $k = 2$  under Definition 3. We will use this operator when we consider the CSO and FCCO problems later. Now, focusing on variance, we can relate the variance of  $L_{D_m}^{(2)}q(s)$  in terms of the variance of  $q(s + \delta)$ . In particular, a consequence of Lemma 2 is that

<sup>3</sup>Note that if the function  $q$  is only  $L_q$ -Lipschitz continuous, then  $j\mathbb{E}[q(s + \delta)] - q(s + \mathbb{E}[\delta])j \leq \sqrt{L_q^2 \mathbb{E}[\delta - \mathbb{E}[\delta]]^2} = \frac{L_q \sigma_2}{m^{1/2}}$ . Therefore, in this case,  $q(s + \delta)$  does not satisfy the first-order guarantee.



(a)  $q(s) = s^2/2$ ,  $\delta \sim \mathcal{N}(10, 100)$ ,  $m = 1$ .



(b)  $q(s) = s^4$ ,  $p(\delta) = \delta/2$  where  $\delta \in [0, 2]$ .

Figure 1: The Fig. 1a investigates the estimation errors of  $L^{(k)}q(s)$  with their number of observations. The shaded region represents the 95%-confidence interval. The Fig. 1b compares the biases of  $\mathbb{E}[L^{(k)}q(s)]$  with increasing inner batch size  $m$ .

$$\mathbb{E} \left[ (L_{D_m}^{(2)} q(s) - \mathbb{E}[L_{D_m}^{(2)} q(s)])^2 \right] = O(\mathbb{E}[(q(s + \delta) - \mathbb{E}[q(s + \delta)])^2]).$$

Extension of  $L_{D_m}^{(2)}$  to Higher-dimensional Case. If  $q: \mathbb{R}^p \rightarrow \mathbb{R}^\ell$  is a vector-valued function, then there is a straightforward extension of Definition 4. Now, for distribution  $D$  over  $\mathbb{R}^p$  and corresponding sampled averaged distribution  $D_m$ , and  $\mathbf{s} \in \mathbb{R}^p$

$$L_{D_m}^{(2)} q: \mathbf{s} \mapsto \left[ 2 q\left(\mathbf{s} + \frac{\delta_1 + \delta_2}{2}\right) - \frac{q(\mathbf{s} + \delta_1) + q(\mathbf{s} + \delta_2)}{2} \right] \quad \text{where } \delta_1, \delta_2 \stackrel{\text{i.i.d.}}{\sim} D_m. \quad (4)$$

Higher-order Extrapolation Operators. The idea behind the construction of  $L_{D_m}^{(2)}$  can be generalized to higher  $k$ 's. For example, in Proposition 3, we construct a third-order extrapolation operator  $L_{D_m}^{(3)}$  through higher degree Taylor series approximation

$$L_{D_m}^{(3)} q: \mathbf{s} \mapsto \left( \frac{1}{36} L_{D_m}^{(2)} + \frac{5}{9} L_{D_{2m}}^{(2)} - \frac{3}{4} L_{D_{3m}}^{(2)} + \frac{16}{9} L_{D_{4m}}^{(2)} - 3 L_{D_{6m}}^{(2)} \right) q(\mathbf{s}).$$

While this idea of expressing the  $k$ -th order operator as an affine combination of lower-order operators works for every  $k$ , explicit constructions soon become tedious.

In Fig. 1, we empirically demonstrate the effectiveness of extrapolation in stochastic estimation.<sup>4</sup> In Fig. 1a, we choose  $q(s) = s^2/2$ ,  $\delta \sim \mathcal{N}(10, 100)$ . For both  $L_{D_{6m}}^{(2)} q(s)$  and  $L_{D_m}^{(3)} q(s)$ , their estimation errors converge to 0 with increasing number of estimates. This coincides with Proposition 1 as  $a_3 = 0$  for quadratic  $q$ . In contrast, biased first order method only converge to a neighborhood. In Fig. 1b, we consider  $q(s) = s^4$  and  $p(\delta) = \delta/2$  where  $\delta \in [0, 2]$ . All three methods are biased and their biases decrease with  $m$ , i.e.  $O(m^{-k})$  for  $k$ th order method. Depending on the constants (e.g.  $a_i, \sigma_i$ ), a higher-order extrapolation method may need decently large  $m$  (burn-in phase) to outperform lower-order methods, i.e., start displaying the expected bias decrease.

### 3 Applying Stochastic Extrapolation in the CSO Problem

In this section, we apply the extrapolation-based scheme from the previous section to reduce the bias in the CSO problem. We focus on variants of BSGD and its accelerated version BSpiderBoost based on our second-order approximation operator (Definition 4). Let  $H_\xi$ ,  $\tilde{H}_\xi$ , and  $H_\xi^0$  indicate different sets, each of which contains  $m$  i.i.d. random variables/samples drawn from the conditional distribution  $\mathbb{P}(\eta|\xi)$ . Remember that, as mentioned earlier, we use  $g_\eta(\mathbf{x})$  to represent  $g_\eta(\mathbf{x}; \xi)$ .

<sup>4</sup>We use  $L_{D_{12m}}^{(1)}$ ,  $L_{D_{6m}}^{(2)}$ ,  $L_{D_m}^{(3)}$  to ensure that each estimate uses same amount of samples (12m).

---

**Algorithm 1 E-BSGD**


---

- 1: Input:  $\mathbf{x}^0 \in \mathbb{R}^d$ , step-size  $\gamma$ , batch sizes  $m$
  - 2: for  $t = 0, 1, \dots, T - 1$  do
  - 3:   Draw one sample  $\xi$  and compute extrapolated gradient  $G_{\text{E-BSGD}}^{t+1}$  from (7)
  - 4:    $\mathbf{x}^{t+1} = \mathbf{x}^t - \gamma G_{\text{E-BSGD}}^{t+1}$
  - 5: end for
  - 6: Output:  $\mathbf{x}^s$  picked uniformly at random from  $\{\mathbf{x}^t\}_{t=0}^{T-1}$
- 

Extrapolated BSGD. At time  $t$ , BSGD constructs a biased estimator of  $\nabla F(\mathbf{x}^t)$  using one sample  $\xi$  and  $2m$  i.i.d. samples from the conditional distribution as in (1)

$$G_{\text{BSGD}}^{t+1} = \left( \frac{1}{m} \sum_{\eta \in H_\xi} \nabla g_\eta(\mathbf{x}^t) \right)^\triangleright \nabla f_\xi \left( \frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t) \right). \quad (5)$$

To reduce this bias, we apply the second-order extrapolation operator from (4). At time  $t$ , we define  $D_{\mathbf{g}, \xi}^{t+1}$  to be the distribution of the random variable  $\frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t)$ . Then we apply  $L_{D_{\mathbf{g}, \xi}^{t+1}}^{(2)}$  by setting  $q$  to  $\nabla f_\xi$  and  $\mathbf{s} = \mathbf{0}$ , i.e.

$$L_{D_{\mathbf{g}, \xi}^{t+1}}^{(2)} \nabla f_\xi(0) := 2 \nabla f_\xi \left( \frac{1}{2m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t) + \frac{1}{2m} \sum_{\eta^0 \in H_\xi^0} g_{\eta^0}(\mathbf{x}^t) \right) \\ \frac{1}{2} \left( \nabla f_\xi \left( \frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t) \right) + \nabla f_\xi \left( \frac{1}{m} \sum_{\eta^0 \in H_\xi^0} g_{\eta^0}(\mathbf{x}^t) \right) \right), \quad (6)$$

where  $\frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t)$  and  $\frac{1}{m} \sum_{\eta^0 \in H_\xi^0} g_{\eta^0}(\mathbf{x}^t)$  are i.i.d. drawn from  $D_{\mathbf{g}, \xi}^{t+1}$ . In Algorithm 1, we present our extrapolated BSGD (E-BSGD) scheme, where we replace  $\nabla f_\xi \left( \frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t) \right)$  in (5) by  $L_{D_{\mathbf{g}, \xi}^{t+1}}^{(2)} \nabla f_\xi(0)$  resulting in this following gradient estimate:

$$G_{\text{E-BSGD}}^{t+1} = \left( \frac{1}{m} \sum_{\eta \in H_\xi} \nabla g_\eta(\mathbf{x}^t) \right)^\triangleright L_{D_{\mathbf{g}, \xi}^{t+1}}^{(2)} \nabla f_\xi(0). \quad (7)$$

---

**Algorithm 2 E-BSpiderBoost**


---

- 1: Input:  $\mathbf{x}^0 \in \mathbb{R}^d$ , step-size  $\gamma$ , batch sizes  $B_1, B_2$ , Probability  $p_{\text{out}}$
- 2: for  $t = 0, 1, \dots, T - 1$  do
- 3:   Draw  $\chi_{\text{out}}$  from Bernoulli( $p_{\text{out}}$ )
- 4:   if  $(t = 0)$  or  $(\chi_{\text{out}} = 1)$  then ▷ Large batch
- 5:     Draw  $B_1$  outer samples  $\tilde{f}\xi_1, \dots, \xi_{B_1} \mathcal{G}$
- 6:     Compute extrapolated gradient  $G_{\text{E-BSGD}}^{t+1}$  with (7)

$$G_{\text{E-BSB}}^{t+1} = \frac{1}{B_1} \sum_{\xi \in \mathcal{B}_1} G_{\text{E-BSGD}}^{t+1}$$

- 7:   else ▷ Small batch
- 8:     Draw  $B_2$  outer samples  $\tilde{f}\xi_1, \dots, \xi_{B_2} \mathcal{G}$
- 9:     Compute extrapolated gradient  $G_{\text{E-BSGD}}^{t+1}$  with (7)

$$G_{\text{E-BSB}}^{t+1} = G_{\text{E-BSB}}^t + \frac{1}{B_2} \sum_{\xi \in \mathcal{B}_2} (G_{\text{E-BSGD}}^{t+1} - G_{\text{E-BSGD}}^t)$$

- 10:   end if
  - 11:    $\mathbf{x}^{t+1} = \mathbf{x}^t - \gamma G_{\text{E-BSB}}^{t+1}$
  - 12: end for
  - 13: Output:  $\mathbf{x}^s$  picked uniformly at random from  $\{\mathbf{x}^t\}_{t=0}^{T-1}$
- 

Extrapolated BSpiderBoost. BSpiderBoost, proposed by Hu et al. [2020b], uses the variance reduction methods for nonconvex smooth stochastic optimization developed by Fang et al. [2018]; Wang et al. [2019].



BSpiderBoost builds upon BSGD and has two kinds of updates: a large batch and a small batch update. In each step, it decides which update to apply based on a random coin. With probability  $p_{\text{out}}$ , it selects a large batch update with  $B_1$  outer samples of  $\xi$ . With remaining probability  $1 - p_{\text{out}}$ , it selects a small batch update where the gradient estimator will be updated with gradient information in the current iteration generated with  $B_2$  outer samples of  $\xi$  and the information from the last iteration. Formally, it constructs a gradient estimate as follows,

$$G_{\text{BSB}}^{t+1} = \begin{cases} G_{\text{BSB}}^t + \frac{1}{B_2} \sum_{\xi \in \mathcal{B}_2, j \in \mathcal{B}_2} (G_{\text{BSGD}}^{t+1} - G_{\text{BSGD}}^t) & \text{with prob. } 1 - p_{\text{out}} \\ \frac{1}{B_1} \sum_{\xi \in \mathcal{B}_1, j \in \mathcal{B}_1} G_{\text{BSGD}}^{t+1} & \text{with prob. } p_{\text{out}}. \end{cases} \quad (8)$$

We propose our extrapolated BSpiderBoost scheme (formally defined in Algorithm 2) by replacing the BSGD gradient estimates in (8) with E-BSGD.

$$G_{\text{E-BSB}}^{t+1} = \begin{cases} G_{\text{E-BSB}}^t + \frac{1}{B_2} \sum_{\xi \in \mathcal{B}_2, j \in \mathcal{B}_2} (G_{\text{E-BSGD}}^{t+1} - G_{\text{E-BSGD}}^t) & \text{with prob. } 1 - p_{\text{out}} \\ \frac{1}{B_1} \sum_{\xi \in \mathcal{B}_1, j \in \mathcal{B}_1} G_{\text{E-BSGD}}^{t+1} & \text{with prob. } p_{\text{out}}. \end{cases} \quad (9)$$

Sample Complexity Analyses of E-BSGD and E-BSpiderBoost. We adopt the standard assumptions used in the literature [Qi et al., 2021b; Wang et al., 2022b; Wang and Yang, 2022; Zhang and Xiao, 2021]. All proofs are deferred to Appendix C.

Assumption 3 (Lower bound)  $F$  is lower bounded by  $F^*$ .

Assumption 4 (Bounded variance) Assume that  $g_\eta$  and  $r g_\eta$  have bounded variances, i.e., for all  $\xi$  in the support of  $\mathbb{P}(\xi)$  and  $\mathbf{x} \in \mathbb{R}^p$ ,  $\sigma_g^2 := \mathbb{E}_{\eta|\xi}[\|g_\eta(\mathbf{x}; \xi) - \mathbb{E}_{\eta|\xi}[g_\eta(\mathbf{x}; \xi)]\|_2^2] < 1$  and  $\zeta_g^2 := \mathbb{E}_{\eta|\xi}[\|r g_\eta(\mathbf{x}; \xi) - \mathbb{E}_{\eta|\xi}[r g_\eta(\mathbf{x}; \xi)]\|_2^2] < 1$ .

Assumption 5 (Lipschitz continuity/smoothness of  $f_\xi$  and  $g_\eta$ ) For all  $\xi$  in the support of  $\mathbb{P}(\xi)$ ,  $f_\xi(\cdot)$  is  $C_f$ -Lipschitz continuous (i.e.,  $\|f_\xi(\mathbf{x}) - f_\xi(\mathbf{x}^0)\|_2 \leq C_f \|\mathbf{x} - \mathbf{x}^0\|_2, \mathbf{x}, \mathbf{x}^0 \in \mathbb{R}^p$ ) and  $L_f$ -Lipschitz smooth (i.e.,  $\|r f_\xi(\mathbf{x}) - r f_\xi(\mathbf{x}^0)\|_2 \leq L_f \|\mathbf{x} - \mathbf{x}^0\|_2, \mathbf{x}, \mathbf{x}^0 \in \mathbb{R}^p$ ) for any  $\xi$ . Similarly, for all  $\xi$  in the support of  $\mathbb{P}(\xi)$  and  $\eta$  in the support of  $\mathbb{P}(\eta|\xi)$ ,  $g_\eta(\cdot; \xi)$  is  $C_g$ -Lipschitz continuous and  $L_g$ -Lipschitz smooth.

The smoothness of  $f_\xi$  and  $g_\eta$  naturally implies the smoothness of  $F$ . Zhang and Xiao [2021, Lemma 4.2] show that Assumption 5 ensures  $F$  is: 1)  $C_F$ -Lipschitz continuous with  $C_F = C_f C_g$ ; and 2)  $L_F$ -Lipschitz smooth with  $L_F = L_g C_f + C_g^2 L_f$ . We denote  $\tilde{L}_F = \zeta_g C_f + \sigma_g C_g L_f$ . Moreover, Assumption 5 also guarantees that  $f_\xi$  and  $g_\eta$  have bounded gradients. In addition,  $f_\xi$  and  $g_\eta$  are assumed to satisfy the following regularity condition in order to apply our extrapolation-based scheme from Section 2.

Assumption 6 (Regularity) For all  $\xi$  in the support of  $\mathbb{P}(\xi)$ ,  $r f_\xi$  is 3rd order differentiable with bounded derivatives (i.e.,  $a_l := \|r^{(l)} f_\xi(\mathbf{x})\|_2 < 1$  for  $l = 1, 2, 3, 4, \mathbf{x} \in \mathbb{R}^p$ ) and  $g_\eta$  has bounded third moment (i.e.,  $\sigma_3 := \mathbb{E}_{\eta|\xi}[\|g_\eta(\mathbf{x}; \xi) - \mathbb{E}_{\eta|\xi}[g_\eta(\mathbf{x}; \xi)]\|_2^3] < 1, \mathbf{x} \in \mathbb{R}^p$ ).

Remark 1 The core piece of Assumption 6 is the 3rd order differentiability of  $r f_\xi$  as other parts can be easily satisfied through appropriate boundedness assumptions. This condition though is satisfied by common instantiations of CSO/FCCO. We discuss some examples including invariant logistic regression, instrumental variable regression, first-order MAML for sine-wave few-shot regression task, deep average precision maximization in Section 5. Therefore, our improvements in sample complexity apply to all these problems.

Consider some time  $t > 0$ . Let  $G^{t+1}$  be a stochastic estimate of  $r F(\mathbf{x}^t)$  where  $\mathbf{x}^t$  is the current iterate. The next iterate  $\mathbf{x}^{t+1} := \mathbf{x}^t - \gamma G^t$ . Let  $\mathbb{E}[\cdot]$  denote the conditional expectation, where we condition on all the randomness until time  $t$ . We consider the bias and variance terms coming from our gradient estimate. Formally, we define the following two quantities

$$E_{\text{bias}}^{t+1} = \|r F(\mathbf{x}^t) - \mathbb{E}[G^{t+1}]\|_2^2, \quad E_{\text{var}}^{t+1} = \mathbb{E}[\|G^{t+1} - \mathbb{E}[G^{t+1}]\|_2^2].$$

Our idea of getting to an  $\varepsilon$ -stationary point (Definition 1) will be to ensure that  $E_{\text{bias}}^{t+1}$  and  $E_{\text{var}}^{t+1}$  are bounded. The main technical component of our analyses is in fact analyzing these bias and variance terms for the various gradient estimates considered. For this purpose, we first analyze the bias and variance terms for the (original) BSGD (Lemma 5) and BSpiderBoost (Lemma 7) algorithms, which are then used to get the corresponding bounds for our E-BSGD (Lemma 6) and E-BSpiderBoost (Lemma 8) algorithms. Through these bias and variance bounds, we establish the following main results of this section.

**Theorem 2 [E-BSGD Convergence]** Consider the (CSO) problem. Suppose Assumptions 3, 4, 5, 6 holds true and  $a_3, \sigma_3, L_F, C_F, \tilde{L}_F, C_g, F^*$  are the constants. Let step size  $\gamma = 1/(2L_F)$ . Then the output  $\mathbf{x}^s$  of E-BSGD (Algorithm 1) satisfies:  $\mathbb{E}[k_{\text{r}} F(\mathbf{x}^s) k_2^2] \leq \varepsilon^2$ , for nonconvex  $F$ , if the inner batch size  $m = \Omega(a_3 \sigma_3 C_g \varepsilon^{-1/2})$ , and the number of iterations

$$T = \Omega(L_F(F(\mathbf{x}^0) - F^*)(L_F^2/m + C_F^2)\varepsilon^{-4}).$$

The E-BSGD takes  $O(\varepsilon^{-4})$  iterations to converge and compute  $O(\varepsilon^{-0.5})$  gradients per iteration. Therefore, its resulting sample complexity is  $O(\varepsilon^{-4.5})$  which is more efficient than  $O(\varepsilon^{-6})$  of BSGD. Similar improvements can be observed for E-BSpiderBoost in Theorem 3.

**Theorem 3 [E-BSpiderBoost Convergence]** Consider the (CSO) problem under the same assumptions as Theorem 2. Let step size  $\gamma = 1/(13L_F)$ . Then the output  $\mathbf{x}^s$  of E-BSpiderBoost (Algorithm 2) satisfies:  $\mathbb{E}[k_{\text{r}} F(\mathbf{x}^s) k_2^2] \leq \varepsilon^2$ , for nonconvex  $F$ , if the inner batch size  $m = \Omega(a_3 \sigma_3 C_g \varepsilon^{-0.5})$ , the hyperparameters of the outer loop of E-BSpiderBoost  $B_1 = (\tilde{L}_F^2/m + C_F^2)\varepsilon^{-2}$ ,  $B_2 = \frac{1}{B_1}$ ,  $p_{\text{out}} = 1/B_2$ , and the number of iterations

$$T = \Omega(L_F(F(\mathbf{x}^0) - F^*)\varepsilon^{-2}).$$

The resulting sample complexity of E-BSpiderBoost is  $O(\varepsilon^{-3.5})$ , which improves  $O(\varepsilon^{-5})$  bound of BSpiderBoost [Hu et al., 2020b] and  $O(\varepsilon^{-4})$  bound of V-MLMC/RT-MLMC [Hu et al., 2021].

## 4 Applying Stochastic Extrapolation in the FCCO Problem

In this section, we apply the extrapolation-based scheme from Section 2 to the FCCO problem. Our focus will be on case where  $n = O(\varepsilon^{-2})$ . For larger  $n$ , we can treat the FCCO problem as a CSO problem and get an  $O(\varepsilon^{-3.5})$  bound from Theorem 3. All missing details are presented in Appendix D.

Now, a straightforward algorithm for FCCO is to use the finite-sum variant of SpiderBoost (or SPIDER) [Fang et al., 2018; Wang et al., 2019] in Algorithm 2. In this case, if we choose the outer batch sizes to be  $B_1 = n$ ,  $B_2 = \frac{1}{n}$  and the inner batch size to be  $m = \max\{\varepsilon^{-2}/n, \varepsilon^{-1/2}g\}$ . The resulting sample complexity of E-BSpiderBoost now becomes,  $O(\max\{\frac{1}{n}/\varepsilon^{2.5}, 1/\frac{1}{n}\varepsilon^4g\})$ , which recovers  $O(\varepsilon^{-3.5})$  bound as in Theorem 3 for  $n = \Theta(\varepsilon^{-2})$ . However, when  $n$  is small, such as  $n = O(1)$ , the sample complexity degenerates to  $O(\varepsilon^{-4})$  which is worse than the  $\Omega(\varepsilon^{-3})$  lower bound of stochastic optimization [Arjevani et al., 2022]. We leave the details to Theorem 7. We still use Assumptions 3, 4, 5, 6 for the analysis of FCCO problem, replacing the role of  $\xi$  with  $i$ .

**Extrapolated NestedVR.** We now introduce a nested variance reduction algorithm E-NestedVR which reaches low sample complexity for all choices of  $n$ . Missing proofs from this section are presented in Appendix D. For the stochasticities in the FCCO problem, our idea is to use two nested SpiderBoost variance reduction components: one for the outer random variable  $i$  and the other for the inner random variable  $\eta/j$ . In each outer (resp. inner) SpiderBoost step, we choose large batch  $B_1$  (resp.  $S_1$ ) with probability  $p_{\text{out}}$  (resp.  $p_{\text{in}}$ ); otherwise we choose small batch. Let  $H_i$  denote a set of  $m$  i.i.d. samples drawn from the conditional distribution  $\mathbb{P}(\eta/j)$ . Similarly, let  $\tilde{H}_i$  denote another set of  $m$  i.i.d. samples drawn from the same conditional distribution. For each given  $i$ , we approximate  $\mathbb{E}_{\eta/j}[g_\eta(\mathbf{x}^t)]$  with  $\mathbf{y}_i^{t+1}$  from distribution  $D_{\mathbf{y},i}^{t+1}$  where,

$$\mathbf{y}_i^{t+1} = \begin{cases} \frac{1}{S_1} \sum_{\eta \in 2H_i} g_\eta(\mathbf{x}^t) & \text{with prob. } p_{\text{in}} \text{ or } t = 0 \\ \mathbf{y}_i^t + \frac{1}{S_2} \sum_{\eta \in 2\tilde{H}_i} (g_\eta(\mathbf{x}^t) - g_\eta(\phi_i^t)) & \text{with prob. } 1 - p_{\text{in}}. \end{cases} \quad (10)$$

---

**Algorithm 3 E-NestedVR**


---

```

1: Input:  $\mathbf{x}^0 \in \mathbb{R}^d$ , step-size  $\gamma$ , batch sizes  $S_1, S_2, B_1, B_2$ , Probability  $p_{in}, p_{out}$ 
2: for  $t = 0, 1, \dots, T - 1$  do
3:   if  $(t = 0)$  or (with prob.  $p_{out}$ ) then ▷ Large outer batch
4:     for  $i \in B_1 \subseteq [n]$  with  $|B_1| = B_1$  do
5:       draw  $\mathbf{y}_i^{t+1}$  from distribution  $D_{\mathbf{y},i}^{t+1}$  defined in (10)
6:       compute  $\mathbf{z}_i^{t+1}$  using (11) and define  $\phi_i^t = \mathbf{x}^t$ 
7:     end for
8:      $G_{E-NVR}^{t+1} = \frac{1}{B_1} \sum_{i \in B_1} (\mathbf{z}_i^{t+1})^\top L_{D_{\mathbf{y},i}^{t+1}}^{(2)} r f_i(0)$ 
9:   else ▷ Small outer batch
10:    for  $i \in B_2$  with  $|B_2| = B_2$  do
11:      draw  $\mathbf{y}_i^{t+1}$  and  $\mathbf{y}_i^t$  from distribution  $D_{\mathbf{y},i}^{t+1}$  and  $D_{\mathbf{y},i}^t$  defined in (10)
12:      compute  $\mathbf{z}_i^{t+1}$  using (11) and define  $\phi_i^t = \mathbf{x}^t$ 
13:    end for
14:     $G_{E-NVR}^{t+1} = G_{E-NVR}^t + \frac{1}{B_2} \sum_{i \in B_2} (\mathbf{z}_i^{t+1})^\top (L_{D_{\mathbf{y},i}^{t+1}}^{(2)} r f_i(0) - L_{D_{\mathbf{y},i}^t}^{(2)} r f_i(0))$ 
15:  end if
16:   $\mathbf{x}^{t+1} = \mathbf{x}^t - \gamma G_{E-NVR}^{t+1}$ 
17: end for
18: Output:  $\mathbf{x}^s$  picked uniformly at random from  $\{\mathbf{x}^t\}_{t=0}^{T-1}$ 

```

---

Similarly, we approximate  $\mathbb{E}_{\eta^t} [r g_\eta(\mathbf{x}^t)]$  with  $\mathbf{z}_i^{t+1}$  defined as follows

$$\mathbf{z}_i^{t+1} = \begin{cases} \frac{1}{S_1} \sum_{\eta \in \mathcal{H}_i} r g_\eta(\mathbf{x}^t) & \text{with prob. } p_{in} \text{ or } t = 0 \\ \mathbf{z}_i^t + \frac{1}{S_2} \sum_{\eta \in \mathcal{H}_i} (r g_\eta(\mathbf{x}^t) - r g_\eta(\phi_i^t)) & \text{with prob. } 1 - p_{in}, \end{cases} \quad (11)$$

where  $\phi_i^t$  is the last time  $i$  is visited before time  $t$ . If  $i$  is not selected at time  $t$ , then  $\mathbf{y}_i^{t+1} = \mathbf{y}_i^t$  and  $\mathbf{z}_i^{t+1} = \mathbf{z}_i^t$ . Note that we use independent samples for  $\mathbf{y}_i^{t+1}$  and  $\mathbf{z}_i^{t+1}$ .

Finally, we present E-NestedVR in Algorithm 3 where second-order extrapolation operator  $L^{(2)}$  is applied to each occurrence of  $r f_i$ . We now analyze its convergence guarantee. Our analysis works by first looking at the effect of multi-level variance reduction without the extrapolation (that we refer to as NestedVR, Theorem 9, Appendix D.2), and then showing how extrapolation could further help to drive down the sample complexity.

**Theorem 4 [E-NestedVR Convergence]** *Consider the (FCCO) problem. Under the same assumptions as Theorem 2.*

*If  $n = O(\varepsilon^{-2/3})$ , then we choose the hyperparameters of E-NestedVR (Algorithm 3) as  $B_1 = B_2 = n, p_{out} = 1, S_1 = \tilde{L}_F \varepsilon^{-2}, S_2 = \tilde{L}_F \varepsilon^{-1}, p_{in} = \tilde{L}_F^{-1} \varepsilon, \gamma = O(\frac{1}{L_F})$ .*

*If  $n = \Omega(\varepsilon^{-2/3})$ , then we choose the hyperparameters of E-NestedVR as  $B_1 = n, B_2 = \frac{P_-}{n}, p_{out} = 1/P_-$ ,  $S_1 = S_2 = \max \left\{ a_3 \sigma_3 C_g \varepsilon^{-1/2}, \tilde{L}_F^2 / (n \varepsilon^2) \right\}$ ,  $p_{in} = 1, \gamma = O(\frac{1}{L_F})$ .*

*Then the output  $\mathbf{x}^s$  of E-NestedVR satisfies:  $\mathbb{E}[k r F(\mathbf{x}^s) k_2^2] \leq \varepsilon^2$ , for nonconvex  $F$ , if the number of iterations*

$$T = \Omega \left( L_F (F(\mathbf{x}^0) - F^*) \varepsilon^{-2} \right).$$

From Theorem 4, E-NestedVR has a sample complexity of  $O(n \varepsilon^{-3})$  in the small  $n$  regime ( $n = O(\varepsilon^{-2/3})$ ) and  $O(\max\{\tilde{L}_F^2 / n \varepsilon^{2.5}, 1 / n \varepsilon^4\})$  in the large  $n$  regime ( $n = \Omega(\varepsilon^{-2/3})$ ). Therefore, in the large  $n$  regime, this improves the  $O(n \varepsilon^{-3})$  sample complexity of MSVR-V2 [Jiang et al., 2022].

## 5 Applications

In this section, we demonstrate the numerical performance of our proposed algorithms. We focus on the application of invariant logistic regression here. In Appendix E, we discuss performance of our proposed




Figure 2: Performances of algorithms and their extrapolated versions on the invariant logistic regression task. Algorithms in each subplot use the same amount of inner batch size in each iteration. The shaded region represents the 95%-confidence interval.

algorithms on other common CSO/FCCO applications, including instrumental variable regression and first-order model-agnostic meta-learning.

Invariant Risk Minimization. Invariant learning has wide applications in machine learning and related areas [Mroueh et al., 2015; Anselmi et al., 2016]. Invariant logistic regression [Hu et al., 2020b] is formulated as follows:

$$\min_{\mathbf{x}} \mathbb{E}_{\xi=(\mathbf{a},b)} [\log(1 + \exp(-b\mathbb{E}_{\eta|\xi}[\eta]^{\top} \mathbf{x}))],$$

where  $\mathbf{a}$  and  $b$  represent a sample and its corresponding label, and  $\eta$  is a noisy observation of the sample  $\mathbf{a}$ . This first part can be considered as a CSO objective, with  $f_{\xi}(y) := \log(1 + \exp(-by))$  and  $g_{\eta}(\mathbf{x}; \xi) := \eta^{\top} \mathbf{x}$ . As the loss  $f_{\xi} \in \mathcal{C}^1$  is smooth, our results from Sections 3 and 4 are applicable.

An  $\ell_2$ -regularizer is added to ensure the existence of  $\cdot$ . Since the gradient of the penalization term is unbiased, we only have to consider the biasness of the data-dependent term. We generate a synthetic dataset with  $d = 10$  dimensions. The minimizer is drawn from Gaussian distribution  $\mathbf{x}^* \sim \mathcal{N}(0, 1) \subseteq \mathbb{R}^d$ . We draw invariant samples  $f(\mathbf{a}_i, b_i)g_i$  where  $\mathbf{a}_i \sim \mathcal{N}(0, 1) \subseteq \mathbb{R}^d$  and compute  $b_i = \text{sgn}(\mathbf{a}_i^{\top} \mathbf{x}^*)$ . Given each  $\xi = (\mathbf{a}_i, b_i)$ , we draw perturbed observations  $\eta \sim \mathcal{N}(\mathbf{a}_i, 100) \subseteq \mathbb{R}^d$ .

We consider drawing  $\xi$  from a large set ( $n = 50000$ ) and a small set ( $n = 50$ ) as CSO and FCCO problems respectively. As baselines, we implemented the BSGD and BSpiderBoost methods from [Hu et al., 2020b], V-MLMC approach from [Hu et al., 2021], and NestedVR approach from Appendix D.2 which achieves the same complexity as MSVR-V2 [Jiang et al., 2022] for the FCCO problem. The results are shown in Fig. 2. In the CSO setting, we compare biased gradient methods with their extrapolated variants (BSGD vs. E-BSGD, BSpiderBoost vs. E-BSpiderBoost, and NestedVR vs. E-NestedVR). The extrapolated versions of BSGD, BSpiderBoost, and NestedVR consistently reach lower error than their non-extrapolated counterparts, as is evident in Figure 2a. In this case, the performance of BSpiderBoost is similar to BSGD as also noted by the authors of these techniques [Hu et al., 2020b], and a drawback of BSpiderBoost seems to be that it is much harder to tune in practice. However, it is clear that E-BSGD outperforms BSGD, and E-BSpiderBoost outperforms BSpiderBoost, respectively. In the FCCO setting, we compare extrapolation based methods and MLMC based methods. Figure 2a, shows that E-NestedVR outperforms all other extrapolated algorithms, including the V-MLMC approach of [Hu et al., 2021], matching our theoretical findings.

## 6 Concluding Remarks

In this paper, we consider the conditional stochastic optimization CSO problem and its finite-sum variant FCCO. Due to the interplay between nested structure and stochasticity, most of the existing gradient estimates suffer from large biases and have large sample complexity of  $O(\varepsilon^{-5})$ . We propose stochastic extrapolation-based algorithms that tackle this bias problem and improve the sample complexities for both these problems. While we focus on nonconvex objectives, our proposed algorithms can also be beneficial when used with strongly convex, convex objectives. We also believe that similar ideas could also prove helpful for multi-level stochastic optimization problems [Zhang and Xiao, 2021] with nested dependency.

## Acknowledgements

We would like to thank Caner Turkmen, Sai Praneeth Karimireddy, and Martin Jaggi for helpful initial discussions surrounding this project.

## References

- Fabio Anselmi, Joel Z Leibo, Lorenzo Rosasco, Jim Mutch, Andrea Tacchetti, and Tomaso Poggio. Unsupervised learning of invariant representations. *Theoretical Computer Science*, 633:112–121, 2016.
- Yossi Arjevani, Yair Carmon, John C Duchi, Dylan J Foster, Nathan Srebro, and Blake Woodworth. Lower bounds for non-convex stochastic optimization. *Mathematical Programming*, pages 1–50, 2022.
- Sébastien Bubeck, Qijia Jiang, Yin Tat Lee, Yuanzhi Li, and Aaron Sidford. Near-optimal method for highly smooth convex optimization. In *Conference on Learning Theory*, pages 492–507. PMLR, 2019.
- Bo Dai, Niao He, Yunpeng Pan, Byron Boots, and Le Song. Learning from conditional distributions via dual embeddings. In *Artificial Intelligence and Statistics*, pages 1458–1467. PMLR, 2017.
- Bo Dai, Albert Shaw, Lihong Li, Lin Xiao, Niao He, Zhen Liu, Jianshu Chen, and Le Song. Sbeed: Convergent reinforcement learning with nonlinear function approximation. In *International Conference on Machine Learning*, pages 1125–1134. PMLR, 2018.
- Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. *Advances in neural information processing systems*, 27, 2014.
- Bradley Efron. *Bootstrap methods: another look at the jackknife*. Springer, 1992.
- Yuri M Ermoliev and Vladimir I Norkin. Sample average approximation method for compound stochastic optimization problems. *SIAM Journal on Optimization*, 23(4):2231–2263, 2013.
- Cong Fang, Chris Junchi Li, Zhouchen Lin, and Tong Zhang. Spider: Near-optimal non-convex optimization via stochastic path-integrated differential estimator. *Advances in Neural Information Processing Systems*, 31, 2018.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017.
- Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- Takashi Goda and Wataru Kitade. Constructing unbiased gradient estimators with finite variance for conditional stochastic optimization. *arXiv preprint arXiv:2206.01991*, 2022.
- YanJun Han, Jiantao Jiao, and Tsachy Weissman. Minimax estimation of divergences between discrete distributions. *IEEE Journal on Selected Areas in Information Theory*, 1(3):814–823, 2020.
- Yifan Hu, Xin Chen, and Niao He. Sample complexity of sample average approximation for conditional stochastic optimization. *SIAM Journal on Optimization*, 30(3):2103–2133, 2020a.
- Yifan Hu, Siqi Zhang, Xin Chen, and Niao He. Biased stochastic first-order methods for conditional stochastic optimization and applications in meta learning. *Advances in Neural Information Processing Systems*, 33: 2759–2770, 2020b.
- Yifan Hu, Xin Chen, and Niao He. On the bias-variance-cost tradeoff of stochastic optimization. *Advances in Neural Information Processing Systems*, 34:22119–22131, 2021.
- Prateek Jain, Purushottam Kar, et al. Non-convex optimization for machine learning. *Foundations and Trends® in Machine Learning*, 10(3-4):142–363, 2017.

- Wei Jiang, Gang Li, Yibo Wang, Lijun Zhang, and Tianbao Yang. Multi-block-single-probe variance reduced estimator for coupled compositional optimization. *arXiv preprint arXiv:2207.08540*, 2022.
- Jiantao Jiao and Yanjun Han. Bias correction with jackknife, bootstrap, and taylor series. *IEEE Transactions on Information Theory*, 66(7):4392–4418, 2020.
- Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in neural information processing systems*, 26, 2013.
- Youssef Mroueh, Stephen Voinea, and Tomaso A Poggio. Learning with group invariant features: A kernel perspective. *Advances in neural information processing systems*, 28, 2015.
- Krikamol Muandet, Arash Mehrjou, Si Kai Lee, and Anant Raj. Dual instrumental variable regression, 2020.
- Ofir Nachum and Bo Dai. Reinforcement learning via fenchel-rockafellar duality. *arXiv preprint arXiv:2001.01866*, 2020.
- Lam M Nguyen, Jie Liu, Katya Scheinberg, and Martin Takáč. Sarah: A novel method for machine learning problems using stochastic recursive gradient. In *International Conference on Machine Learning*, pages 2613–2621. PMLR, 2017.
- Qi Qi, Youzhi Luo, Zhao Xu, Shuiwang Ji, and Tianbao Yang. Stochastic optimization of areas under precision-recall curves with provable convergence. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 1752–1765, 2021a. URL <https://proceedings.neurips.cc/paper/2021/hash/0dd1bc593a91620daecf7723d2235624-Abstract.html>.
- Qi Qi, Youzhi Luo, Zhao Xu, Shuiwang Ji, and Tianbao Yang. Stochastic optimization of areas under precision-recall curves with provable convergence. *Advances in Neural Information Processing Systems*, 34: 1752–1765, 2021b.
- Sashank J Reddi, Ahmed Hefny, Suvrit Sra, Barnabas Poczos, and Alex Smola. Stochastic variance reduction for nonconvex optimization. In *International conference on machine learning*, pages 314–323. PMLR, 2016a.
- Sashank J Reddi, Suvrit Sra, Barnabás Póczos, and Alex Smola. Fast incremental method for nonconvex optimization. *arXiv preprint arXiv:1603.06159*, 2016b.
- Mark Schmidt, Nicolas Le Roux, and Francis Bach. Minimizing finite sums with the stochastic average gradient. *Mathematical Programming*, 162(1):83–112, 2017.
- Rahul Singh, Maneesh Sahani, and Arthur Gretton. Kernel instrumental variable regression. *Advances in Neural Information Processing Systems*, 32, 2019.
- John Tukey. Bias and confidence in not quite large samples. *Ann. Math. Statist.*, 29:614, 1958.
- Bokun Wang and Tianbao Yang. Finite-sum coupled compositional stochastic optimization: Theory and applications. In *International Conference on Machine Learning*, pages 23292–23317. PMLR, 2022.
- Guanghui Wang, Ming Yang, Lijun Zhang, and Tianbao Yang. Momentum accelerates the convergence of stochastic AUPRC maximization. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera, editors, *International Conference on Artificial Intelligence and Statistics, AISTATS 2022, 28-30 March 2022, Virtual Event*, volume 151 of *Proceedings of Machine Learning Research*, pages 3753–3771. PMLR, 2022a. URL <https://proceedings.mlr.press/v151/wang22b.html>.
- Guanghui Wang, Ming Yang, Lijun Zhang, and Tianbao Yang. Momentum accelerates the convergence of stochastic auprc maximization. In *International Conference on Artificial Intelligence and Statistics*, pages 3753–3771. PMLR, 2022b.

- Mengdi Wang, Ethan X Fang, and Han Liu. Stochastic compositional gradient descent: algorithms for minimizing compositions of expected-value functions. *Mathematical Programming*, 161:419–449, 2017.
- Zhe Wang, Kaiyi Ji, Yi Zhou, Yingbin Liang, and Vahid Tarokh. Spiderboost and momentum: Faster variance reduction algorithms. *Advances in Neural Information Processing Systems*, 32, 2019.
- Christopher Stroud Withers. Bias reduction by taylor series. *Communications in Statistics-Theory and Methods*, 16(8):2369–2383, 1987.
- Yu M Yermol'yev. A general stochastic programming problem. 1971.
- Junyu Zhang and Lin Xiao. Multilevel composite stochastic optimization via nested variance reduction. *SIAM Journal on Optimization*, 31(2):1131–1157, 2021.

## A Missing Details from Section 1

### A.1 Other Related Work

CSO. Dai et al. [2017] proposed a primal-dual stochastic approximation algorithm to solve a min-max reformulation of CSO, employing the kernel embedding techniques. However, this method requires convexity of  $f_\xi$  and linearity of  $g_\eta$ , which are not satisfied by general applications when neural networks are involved. Goda and Kitade [2022] showed that a special class of CSO problems can be unbiased, e.g., when  $f_\xi$  measures the squared error between some  $u(\xi)$  and  $\mathbb{E}_{\eta|\xi}[g_\eta(\mathbf{x}; \xi)]$ , giving rise to this objective function  $\mathbb{E}_\xi[(u(\xi) - \mathbb{E}_{\eta|\xi}[g_\eta(\mathbf{x}; \xi)])^2]$ . However, they did not show any improvement over the sample complexity of BSGD (i.e.,  $O(\varepsilon^{-6})$ ). Hu et al. [2020b] also analyzed lower bounds on the minimax error for the CSO problem and showed that for a specific class of biased gradients with  $O(\varepsilon)$  bias (same bias as BSGD) and variance  $O(1)$  the bound achieved by BSpiderBoost is tight. However, these lower bounds are not applicable in settings such as ours (and also to [Hu et al., 2021]) where the bias is smaller than the BSGD bias.

Variance Reduction. The reduction of variance in stochastic optimization is a crucial approach to decrease sample complexity, particularly when dealing with finite-sum formulations of the form  $\min_{\mathbf{x}} \frac{1}{n} \sum_{i=1}^n f_i(\mathbf{x})$ . Pioneering works such as Stochastic Average Gradient (SAG) [Schmidt et al., 2017], Stochastic Variance Reduced Gradient (SVRG) [Johnson and Zhang, 2013; Reddi et al., 2016a], and SAGA [Defazio et al., 2014; Reddi et al., 2016b] improved the iteration complexity from  $O(\varepsilon^{-4})$  in Stochastic Gradient Descent (SGD) to  $O(\varepsilon^{-2})$ . Subsequent research, including Stochastic Path-Integrated Differential Estimator (SPIDER) [Fang et al., 2018] and Stochastic Recursive Gradient Algorithm (SARAH) [Nguyen et al., 2017], expanded the application of these techniques to both finite-sum and online scenarios, where  $n$  is large or possibly infinite. These methods boast an improved sample complexity of  $\min(\frac{1}{n\varepsilon^2}, \varepsilon^{-3})$ . SpiderBoost [Wang et al., 2019], achieves the same near-optimal complexity performance as SPIDER, but allows a much larger step size and hence runs faster in practice than SPIDER. In this paper, we use a probabilistic variant of SpiderBoost as the variance reduction module for CSO and FCCO problems. We highlight that alternative techniques, such as SARAH, can also be applied and offer similar guarantees.

Bias Correction. One of the classic problems in statistics is to design procedures to reduce the bias of estimators. Well-established general bias correction techniques, such as the jackknife [Tukey, 1958], bootstrap [Efron, 1992], Taylor series [Withers, 1987; Han et al., 2020], have been extensively studied and applied in various contexts [Jiao and Han, 2020]. However, these methods are predominantly examined in relation to standard statistical distributions, with limited emphasis on their adaptability to optimization problems. Our proposed extrapolation-based approach is derived from sample-splitting methods [Han et al., 2020], specifically tailored and analyzed for optimization problems involving unknown distributions.

Stochastic Composition Optimization. Finally, a closely related class of problems, called stochastic composition optimization, has been extensively studied (e.g., [Yermol'yev, 1971; Ermoliev and Norkin, 2013; Wang et al., 2017]) in the literature where the goal is:

$$\min_{\mathbf{x} \in \mathbb{R}^d} \mathbb{E}_\xi [f_\xi(\mathbb{E}_\eta [g_\eta(\mathbf{x})])]. \quad (12)$$

Despite having nested expectations in their formulations (CSO) and (12) are fundamentally different: a) in stochastic composite optimization the inner randomness  $\eta$  is conditionally dependent on the outer randomness  $\xi$  and b) in CSO the inner random function  $g_\eta(\mathbf{x}, \xi)$  depends on both  $\xi$  and  $\eta$ . These differences lead to quite different sample complexity bounds for these problems, as explored in Hu et al. [2020a]. In fact, Zhang and Xiao [2021] presented a near optimal complexity of  $O(\min(\varepsilon^{-3}, \frac{1}{n\varepsilon^2}))$  for stochastic composite optimization problems using nested variance reduction.

## B Missing Details from Section 2

Lemma 1 (Moments of  $D_m$ ) *The moments of  $\delta \in D_m$  are bounded as follows*

$$\mathbb{E}[(\delta - \mathbb{E}[\delta])^2] = \frac{\sigma_2}{m}, \quad j \mathbb{E}[(\delta - \mathbb{E}[\delta])^3] = \frac{\sigma_3}{m^2}, \quad \mathbb{E}[(\delta - \mathbb{E}[\delta])^4] = \frac{\sigma_4}{m^3} + \frac{3(m-1)\sigma_2^2}{m^3}.$$

More generally, for  $k \geq 2$ ,  $j \mathbb{E}[(\delta - \mathbb{E}[\delta])^k] = O(m^{-dk/2e})$ .



Proof: Define  $\hat{\delta} = \delta - \mathbb{E}[\delta]$  as the centered random variable. Now

$$\mathbb{E}[(\delta - \mathbb{E}[\delta])^k] = \mathbb{E}[\hat{\delta}^k].$$

So we focus on  $\mathbb{E}[\hat{\delta}^k]$  in the remainder of the proof. For  $k = 2$ ,

$$j\mathbb{E}[\hat{\delta}^2]j = \frac{1}{m^2}j\mathbb{E}[\sum_{i=1}^m \hat{\delta}_i^2]j = \frac{1}{m^2} \left| \mathbb{E} \left[ \sum_i \hat{\delta}_i^2 + 2 \sum_{i<j} \hat{\delta}_i \hat{\delta}_j \right] \right| = \frac{\sigma_2}{m}.$$

For  $k = 3$ ,

$$\begin{aligned} j\mathbb{E}[\hat{\delta}^3]j &= \frac{1}{m^3}j\mathbb{E}[\sum_{i=1}^m \hat{\delta}_i^3]j \\ &= \frac{1}{m^3} \left| \mathbb{E} \left[ \sum_i \hat{\delta}_i^3 + 3 \sum_{i \neq j} \hat{\delta}_i^2 \hat{\delta}_j + 6 \sum_{i<j<k} \hat{\delta}_i \hat{\delta}_j \hat{\delta}_k \right] \right| \\ &= \frac{\sigma_3}{m^2}. \end{aligned}$$

For  $k = 4$ ,

$$\begin{aligned} j\mathbb{E}[\hat{\delta}^4]j &= \frac{1}{m^4}j\mathbb{E}[\sum_{i=1}^m \hat{\delta}_i^4]j \\ &= \frac{1}{m^4} \left| \mathbb{E} \left[ \sum_i \hat{\delta}_i^4 + 4 \sum_{i \neq j} \hat{\delta}_i^3 \hat{\delta}_j + 6 \sum_{i<j} \hat{\delta}_i^2 \hat{\delta}_j^2 + 24 \sum_{i<j<k<l} \hat{\delta}_i \hat{\delta}_j \hat{\delta}_k \hat{\delta}_l \right] \right| \\ &= \frac{1}{m^4} \left| m \mathbb{E}[\hat{\delta}_i^4] + 6 \frac{m(m-1)}{2} \mathbb{E}[\hat{\delta}_i^2] \mathbb{E}[\hat{\delta}_j^2] \right| \\ &= \frac{\sigma_4}{m^3} + \frac{3(m-1)\sigma_2^2}{m^3}. \end{aligned}$$

For  $k = 5$ ,

$$\begin{aligned} j\mathbb{E}[\hat{\delta}^5]j &= \frac{1}{m^5}j\mathbb{E}[\sum_{i=1}^m \hat{\delta}_i^5]j \\ &= \frac{1}{m^5} \left| \mathbb{E} \left[ \sum_i \hat{\delta}_i^5 + 10 \sum_{i \neq j} \hat{\delta}_i^3 \hat{\delta}_j^2 \right] \right| \\ &= \frac{1}{m^5} \left| m \mathbb{E}[\hat{\delta}_i^5] + 10m(m-1) \mathbb{E}[\hat{\delta}_i^3] \mathbb{E}[\hat{\delta}_j^2] \right| \\ &= \frac{\sigma_5}{m^4} + \frac{10(m-1)\sigma_3\sigma_2}{m^4}. \end{aligned}$$

For general  $k > 0$ , we expand the following term as a function of  $m$

$$j\mathbb{E}[\hat{\delta}^k]j = \frac{1}{m^k}j\mathbb{E}[\sum_{i=1}^m \hat{\delta}_i^k]j.$$

As  $\mathbb{E}[\hat{\delta}_i] = 0$  and  $\hat{\delta}_i$  and  $\hat{\delta}_j$  are independent for different  $i$  and  $j$ , the outcome has the following form

$$j\mathbb{E}[\hat{\delta}^k]j = \frac{1}{m^k} \mathcal{O} \left( \sum_{\substack{2a_2+3a_3+\dots+a_k=k \\ a_i \geq 0}} m^{\sum_{i=2}^k a_i} \sigma_2^{a_2} \sigma_3^{a_3} \dots \sigma_k^{a_k} \right) \quad (13)$$

where  $\sum_{i=2}^k a_i$  is the count of independent  $\hat{\delta}_i$  used in  $\sigma_2^{a_2} \sigma_3^{a_3} \dots \sigma_k^{a_k}$ . Among the terms in (13), the dominating one in terms of  $m$  is one with largest  $\sum_{i=2}^k a_i$ , i.e.

$$j\mathbb{E}[\hat{\delta}^k]j = \begin{cases} \frac{1}{m^k} \mathcal{O}(m^{k/2}) \sigma_2^{k/2} & \text{if } k \text{ even,} \\ \frac{1}{m^k} \mathcal{O}(m^{bk/2c}) \sigma_2^{bk/2c} \sigma_3 & \text{if } k \text{ odd.} \end{cases}$$

Then, we can simplify the upper right-hand side with

$$j\mathbb{E}[\hat{\delta}^k]j = \mathcal{O}(m^{-k+bk/2c}),$$

which gives all the desired results.

Proposition 2 (First-order Guarantee) *Assume that  $D_m$  and  $q(\cdot)$  satisfy Assumption 1 and 2 respectively with  $k = 1$ . Then,  $\forall s \in \mathbb{R}$ ,  $\left| \mathbb{E} \left[ L_{D_m}^{(1)} q(s) \right] - q(s + \mathbb{E}[\delta]) \right| \leq a_2 \sigma_2 / (2m)$ .*

Proof: Let  $h = \mathbb{E}[\delta]$ . If the function  $q \in \mathcal{C}^2$ , then the Taylor expansion at  $s + h$  with remainders leads to

$$\mathbb{E}[q(s + \delta)] = q(s + h) + q'(s + h) \mathbb{E}[\delta - h] + \frac{q''(\phi_1)}{2} \mathbb{E}[(\delta - h)^2]$$

where  $\phi_1$  between  $s$  and  $s + \delta$ . Then the error of extrapolation becomes

$$j \mathbb{E}[q(s + \delta)] - q(s + h)j = \left| \frac{q''(\phi_1)}{2} \mathbb{E}[(\delta - h)^2] \right| = \frac{j q''(\phi_1)}{2} \mathbb{E}[(\delta - h)^2].$$

By Assumption 2 and Lemma 1, we have that

$$j \mathbb{E}[q(s + \delta)] - q(s + h)j \leq \frac{a_2}{2} \mathbb{E}[(\delta - h)^2] = \frac{a_2}{2} \mathbb{E}[(\delta - h)^2] = \frac{a_2 \sigma_2}{2m}.$$

This completes the proof.

Proposition 1 (Second-order Guarantee) *Assume that distribution  $D_m$  and  $q(\cdot)$  satisfies Assumption 1 and 2 respectively with  $k = 2$ . Then, for all  $s \in \mathbb{R}$ ,  $\left| \mathbb{E} \left[ L_{D_m}^{(2)} q(s) \right] - q(s + \mathbb{E}[\delta]) \right| \leq a_3 \sigma_3 / m^2$ .*

Proof: Let  $h = \mathbb{E}[\delta]$ . If the function  $q \in \mathcal{C}^3$ , then the Taylor expansion at  $s + h$  with remainders leads to

$$\begin{aligned} \mathbb{E}[q(s + \delta_1)] &= q(s + h) + q'(s + h) \mathbb{E}[\delta_1 - h] + \frac{q''(s+h)}{2} \mathbb{E}[(\delta_1 - h)^2] + \frac{q^{(3)}(\phi_1)}{6} \mathbb{E}[(\delta_1 - h)^3] \\ \mathbb{E}[q(s + \delta_2)] &= q(s + h) + q'(s + h) \mathbb{E}[\delta_2 - h] + \frac{q''(s+h)}{2} \mathbb{E}[(\delta_2 - h)^2] + \frac{q^{(3)}(\phi_2)}{6} \mathbb{E}[(\delta_2 - h)^3] \\ \mathbb{E}[q(s + \frac{\delta_1 + \delta_2}{2})] &= q(s + h) + q'(s + h) \mathbb{E}[\frac{\delta_1 + \delta_2}{2} - h] + \frac{q''(s+h)}{2} \mathbb{E}[(\frac{\delta_1 + \delta_2}{2} - h)^2] + \frac{q^{(3)}(\phi_3)}{6} \mathbb{E}[(\frac{\delta_1 + \delta_2}{2} - h)^3] \end{aligned}$$

where  $\phi_1 \in [s, s + \delta_1]$ ,  $\phi_2 \in [s, s + \delta_2]$ ,  $\phi_3 \in [s, s + \frac{\delta_1 + \delta_2}{2}]$ .

As  $\mathbb{E}[\delta - h] = 0$ , the error of extrapolation becomes

$$\begin{aligned} j \mathbb{E}[L_{D_m}^{(2)} q(s)] - q(s + h)j &= \left| 2 \frac{q^{(3)}(\phi_3)}{6} \mathbb{E}[(\frac{\delta_1 + \delta_2}{2} - h)^3] - \frac{1}{2} \left( \frac{q^{(3)}(\phi_1)}{6} \mathbb{E}[(\delta_1 - h)^3] + \frac{q^{(3)}(\phi_2)}{6} \mathbb{E}[(\delta_2 - h)^3] \right) \right| \\ &= \frac{j q^{(3)}(\phi_3)}{3} \left| \mathbb{E}[(\frac{\delta_1 + \delta_2}{2} - h)^3] \right| + \frac{j q^{(3)}(\phi_1)}{12} \left| \mathbb{E}[(\delta_1 - h)^3] \right| + \frac{j q^{(3)}(\phi_2)}{12} \left| \mathbb{E}[(\delta_2 - h)^3] \right| \\ &= \frac{a_3}{3} \left( \left| \mathbb{E}[(\frac{\delta_1 + \delta_2}{2} - h)^3] \right| + \left| \mathbb{E}[(\delta_1 - h)^3] \right| + \left| \mathbb{E}[(\delta_2 - h)^3] \right| \right) \\ &= \frac{a_3 \sigma_3}{m^2}, \end{aligned}$$

where we used Assumption 2 in the second to last inequality and Lemma 1 in the last inequality.

Proposition 3 *Assume  $q \in \mathcal{C}^5$ . Then  $L_{D_m}^{(3)}$  as defined below is a third-order extrapolation operator.*

$$L_{D_m}^{(3)} q : s \mapsto \left( \frac{1}{36} L_{D_m}^{(2)} + \frac{5}{9} L_{D_{2m}}^{(2)} - \frac{3}{4} L_{D_{3m}}^{(2)} + \frac{16}{9} L_{D_{4m}}^{(2)} - 3 L_{D_{6m}}^{(2)} \right) q(s).$$

Proof: Let  $h = \mathbb{E}[\delta]$ . If  $q \in \mathcal{C}^{2k-1}$ , then  $q$  has the following Taylor expansion

$$\begin{aligned} \mathbb{E}[q(s + \delta)] &= \underbrace{q(s + h)}_{\text{zero order term}} + q'(s + h) \mathbb{E}[\delta - h] + \underbrace{\frac{q''(s+h)}{2} \mathbb{E}[(\delta - h)^2]}_{\text{second order term}} + \dots \\ &\quad + \frac{q^{(2k-1)}(s+h)}{(2k-1)!} \mathbb{E}[(\delta - h)^{2k-1}] + \frac{q^{(2k)}(\phi)}{2k!} \mathbb{E}[(\delta - h)^{2k}]. \end{aligned}$$

Eliminate the third order term in the Taylor expansion. Consider the following affine combination which

$$F_{D_m}^{(3)} q : s \mapsto \alpha_1 L_{D_m}^{(2)} q(s) + \alpha_2 L_{D_{2m}}^{(2)} q(s).$$

We determine  $\alpha_1$  and  $\alpha_2$  by expanding  $L_{D_m}^{(2)} q(s)$  and  $L_{D_{2m}}^{(2)} q(s)$  and analyze the coefficients of terms:

- (A ne). Taylor expansion of  $F_{D_m}^{(3)}q(s)$  at  $s+h$  should have zero order term  $q(s+h)$ , i.e.

$$\alpha_1 q(s+h) + \alpha_2 q(s+h) = q(s+h).$$

- (Eliminate third term). Taylor expansion of  $F_{D_m}^{(3)}q(s)$  at  $s+h$  should have third order term  $\mathbb{E}[(\delta-h)^3]$ . That is,

$$\alpha_1 \mathbb{E}[(\delta_1-h)^3] + \alpha_2 \mathbb{E}\left[\left(\frac{\delta_1+\delta_2}{2}-h\right)^3\right] = 0.$$

This is equivalent to

$$\alpha_1 \mathbb{E}[(\delta_1-h)^3] + \frac{\alpha_2}{4} \mathbb{E}\left[(\delta_1-h)^3\right] = 0.$$

Therefore,  $\alpha_1$  and  $\alpha_2$  can be determined through the following linear system

$$\begin{aligned} \alpha_1 + \alpha_2 &= 1 \\ \alpha_1 + \frac{1}{4}\alpha_2 &= 0. \end{aligned}$$

The solution is  $\alpha_1 = \frac{1}{3}$  and  $\alpha_2 = \frac{4}{3}$ .

For  $k=3$  order extrapolation, consider the following

$$L_{D_m}^{(3)}q : s \not\sim \alpha_1^\ell F_{D_m}^{(3)}q(s) + \alpha_2^\ell F_{D_{2m}}^{(3)}q(s) + \alpha_3^\ell F_{D_{3m}}^{(3)}q(s).$$

We determine  $\alpha_1^\ell$ ,  $\alpha_2^\ell$  and  $\alpha_3^\ell$  by satisfying the following two conditions

- (A ne). Taylor expansion of  $L_{D_m}^{(3)}q(s)$  at  $s+h$  should have zero order term  $q(x+h)$ , i.e.

$$(\alpha_1^\ell + \alpha_2^\ell + \alpha_3^\ell)q(x+h) = q(x+h).$$

- Taylor expansion of  $L_{D_m}^{(3)}q(s)$  at  $s+h$  should have 4th order term  $\mathbb{E}[(\delta-h)^4]$ . That is

$$\alpha_1^\ell \mathbb{E}[(\delta_1-h)^4] + \alpha_2^\ell \mathbb{E}\left[\left(\frac{\delta_1+\delta_2}{2}-h\right)^4\right] + \alpha_3^\ell \mathbb{E}\left[\left(\frac{\delta_1+\delta_2+\delta_3}{3}-h\right)^4\right] = 0.$$

This is equivalent to

$$\begin{aligned} \left(\alpha_1^\ell + \frac{\alpha_2^\ell}{8} + \frac{\alpha_3^\ell}{27}\right) \mathbb{E}[(\delta_1-h)^4] &= 0 \\ \left(\frac{3}{8}\alpha_2^\ell + \frac{2}{9}\alpha_3^\ell\right) (\mathbb{E}[(\delta_1-h)^2])^2 &= 0. \end{aligned}$$

Therefore,  $\alpha_1^\ell$ ,  $\alpha_2^\ell$  and  $\alpha_3^\ell$  can be determined through the following linear system

$$\begin{aligned} \alpha_1^\ell + \alpha_2^\ell + \alpha_3^\ell &= 1 \\ \alpha_1^\ell + \frac{1}{8}\alpha_2^\ell + \frac{1}{27}\alpha_3^\ell &= 0 \\ \alpha_1^\ell + \frac{3}{8}\alpha_2^\ell + \frac{2}{9}\alpha_3^\ell &= 0. \end{aligned}$$

The solution is  $\alpha_1^\ell = \frac{1}{12}$ ,  $\alpha_2^\ell = \frac{4}{3}$  and  $\alpha_3^\ell = \frac{9}{4}$ . Then consider the Taylor expansion of  $L_{D_m}^{(3)}q(s)$  at  $s+h$  with (2), we can

$$j \mathbb{E}[L_{D_m}^{(3)}q(s)] - q(s+h)j \cdot \left| \frac{q^{(5)}(\phi_\delta)}{5!} \mathbb{E}[(\delta-h)^5] \right| \cdot O(a_5 m^{-3})$$

where the first inequality uses the fact that  $L_{D_m}^{(3)}$  is an affine mapping and the last inequality uses Lemma 1.

Therefore,  $L_{D_m}^{(3)}$  is a 3rd-order extrapolation operator. We can expand it into

$$\begin{aligned} L_{D_m}^{(3)}q : s \not\sim & \frac{1}{12} \left( \frac{1}{3} L_{D_m}^{(2)}q(s) + \frac{4}{3} L_{D_{2m}}^{(2)}q(s) \right) - \frac{4}{3} \left( \frac{1}{3} L_{D_{2m}}^{(2)}q(s) + \frac{4}{3} L_{D_{4m}}^{(2)}q(s) \right) \\ & + \frac{9}{4} \left( \frac{1}{3} L_{D_{3m}}^{(2)}q(s) + \frac{4}{3} L_{D_{6m}}^{(2)}q(s) \right) \\ = & \left( \frac{1}{36} L_{D_m}^{(2)} + \frac{5}{9} L_{D_{2m}}^{(2)} - \frac{3}{4} L_{D_{3m}}^{(2)} - \frac{16}{9} L_{D_{4m}}^{(2)} + 3 L_{D_{6m}}^{(2)} \right) q(s). \end{aligned}$$

Lemma 2 (Variance Bound) Assume that  $q: \mathbb{R}^p \rightarrow \mathbb{R}^\ell$  is in  $C^3$  and  $D_m$  is the distribution in Assumption 1. Suppose that the variance of  $q(\mathbf{s} + \boldsymbol{\delta})$  is bounded as

$$\mathbb{E}[\|kq(\mathbf{s} + \boldsymbol{\delta}) - \mathbb{E}[q(\mathbf{s} + \boldsymbol{\delta})]\|_2^2] \leq \frac{V^2}{m} + C.$$

Then the variance of extrapolation  $L_{D_m}^{(2)}q(\mathbf{s})$  is upper bounded by

$$\mathbb{E} \left[ \left\| L_{D_m}^{(2)}q(\mathbf{s}) - \mathbb{E}[L_{D_m}^{(2)}q(\mathbf{s})] \right\|_2^2 \right] \leq 14 \left( \frac{V^2}{m} + C \right).$$

Proof: Let us use the definition of  $L_{D_m}^{(2)}q(\mathbf{s})$ :

$$\begin{aligned} & \mathbb{E} \left[ \left\| L_{D_m}^{(2)}q(\mathbf{s}) - \mathbb{E}[L_{D_m}^{(2)}q(\mathbf{s})] \right\|_2^2 \right] \\ &= \mathbb{E} \left[ \left\| 2q\left(\mathbf{s} + \frac{\boldsymbol{\delta}_1 + \boldsymbol{\delta}_2}{2}\right) - \frac{q(\mathbf{s} + \boldsymbol{\delta}_1) + q(\mathbf{s} + \boldsymbol{\delta}_2)}{2} - \left[ 2q\left(\mathbf{s} + \frac{\boldsymbol{\delta}_1 + \boldsymbol{\delta}_2}{2}\right) - \frac{q(\mathbf{s} + \boldsymbol{\delta}_1) + q(\mathbf{s} + \boldsymbol{\delta}_2)}{2} \right] \right\|_2^2 \right] \\ &= 3 \mathbb{E} \left[ \left\| 2q\left(\mathbf{s} + \frac{\boldsymbol{\delta}_1 + \boldsymbol{\delta}_2}{2}\right) - \mathbb{E}\left[2q\left(\mathbf{s} + \frac{\boldsymbol{\delta}_1 + \boldsymbol{\delta}_2}{2}\right)\right] \right\|_2^2 \right] + 3 \mathbb{E} \left[ \left\| \frac{q(\mathbf{s} + \boldsymbol{\delta}_1)}{2} - \mathbb{E}\left[\frac{q(\mathbf{s} + \boldsymbol{\delta}_1)}{2}\right] \right\|_2^2 \right] \\ &+ 3 \mathbb{E} \left[ \left\| \frac{q(\mathbf{s} + \boldsymbol{\delta}_2)}{2} - \mathbb{E}\left[\frac{q(\mathbf{s} + \boldsymbol{\delta}_2)}{2}\right] \right\|_2^2 \right] \\ &= 12 \left( \frac{V^2}{2m} + C \right) + \frac{3}{4} \left( \frac{V^2}{m} + C \right) + \frac{3}{4} \left( \frac{V^2}{m} + C \right) \\ &= \frac{15V^2}{2m} + \frac{27C}{2}. \end{aligned}$$

This completes the proof.

## C Stationary Point Convergence Proofs from Section 3 (CSO)

In this section, we provide the convergence proofs for the CSO problem. We start by establishing some helpful lemmas in Appendix C.1. In Appendix C.2, we reanalyze the BSGD algorithm to obtain explicit bias and variance bounds, which are then useful when we analyze E-BSGD in Appendix C.3. Similarly, we reanalyze BSpiderBoost in Appendix C.4 and use the resulting bias and variance bounds for the analysis of E-BSpiderBoost in Appendix C.5.

Note that throughout our analyses, we define  $\mathbb{E}^{t+1}[\cdot]$  as the expectation of randomness at time  $t+1$  conditioning on the randomness until time  $t$ . When there is no ambiguity, we use  $\mathbb{E}[\cdot]$  instead of  $\mathbb{E}^{t+1}[\cdot]$ .

### C.1 Helpful Lemmas

Lemma 3 (Sufficient Decrease) Suppose Assumption 5 holds true and  $\gamma \leq \frac{1}{2L_F}$  then

$$\|r F(\mathbf{x}^t)\|_2^2 \leq \frac{2(\mathbb{E}[F(\mathbf{x}^{t+1})] - F(\mathbf{x}^t))}{\gamma} + L_F \gamma E_{\text{var}}^{t+1} + E_{\text{bias}}^{t+1},$$

where  $\mathbb{E}[\cdot]$  denote conditional expectation over the randomness at time  $t$  conditioned on all of the past randomness until time  $t$ .

Proof: In this proof, we use  $\mathbb{E}[\cdot]$  to denote conditional expectation over the randomness at time  $t$  conditioned on all the past randomness until time  $t$ .

Let us expand  $F(\mathbf{x}^{t+1})$  and apply the  $L_F$ -smoothness of  $F$

$$\mathbb{E}[F(\mathbf{x}^{t+1})] - F(\mathbf{x}^t) \leq \gamma \mathbb{E}[\langle \nabla F(\mathbf{x}^t), G^{t+1} \rangle] + \frac{L_F \gamma^2}{2} \mathbb{E}[\|G^{t+1}\|_2^2].$$

Since  $\mathbb{E}[\|G^{t+1}\|_2^2] = \mathbb{E}[\|G^{t+1} - \mathbb{E}[G^{t+1}]\|_2^2] + \|\mathbb{E}[G^{t+1}]\|_2^2 = E_{\text{var}}^{t+1} + \|\mathbb{E}[G^{t+1}]\|_2^2$ , then

$$\mathbb{E}[F(\mathbf{x}^{t+1})] - F(\mathbf{x}^t) \leq \gamma \mathbb{E}[\langle \nabla F(\mathbf{x}^t), G^{t+1} \rangle] + \frac{L_F \gamma^2}{2} (E_{\text{var}}^{t+1} + \|\mathbb{E}[G^{t+1}]\|_2^2).$$

Expand the middle term with

$$\begin{aligned}\gamma \mathbb{E}[h r F(\mathbf{x}^t), G^{t+1} | j] &= \frac{\gamma}{2} \|r F(\mathbf{x}^t)\|_2^2 - \frac{\gamma}{2} \|\mathbb{E}[G^{t+1}]\|_2^2 + \frac{\gamma}{2} \|r F(\mathbf{x}^t) - \mathbb{E}[G^{t+1}]\|_2^2 \\ &= \frac{\gamma}{2} \|r F(\mathbf{x}^t)\|_2^2 - \frac{\gamma}{2} \|\mathbb{E}[G^{t+1}]\|_2^2 + \frac{\gamma}{2} E_{\text{bias}}^{t+1}.\end{aligned}$$

Combine with the inequality

$$\mathbb{E}[F(\mathbf{x}^{t+1})] - F(\mathbf{x}^t) \leq \frac{\gamma}{2} \|r F(\mathbf{x}^t)\|_2^2 - \frac{\gamma}{2} (1 - L_F \gamma) \|\mathbb{E}[G^{t+1}]\|_2^2 + \frac{\gamma}{2} E_{\text{bias}}^{t+1} + \frac{L_F \gamma^2}{2} E_{\text{var}}^{t+1}.$$

By taking  $\gamma = \frac{1}{2L_F}$ , we have that

$$\mathbb{E}[F(\mathbf{x}^{t+1})] - F(\mathbf{x}^t) \leq \frac{\gamma}{2} \|r F(\mathbf{x}^t)\|_2^2 - \frac{\gamma}{4} \|\mathbb{E}[G^{t+1}]\|_2^2 + \frac{\gamma}{2} E_{\text{bias}}^{t+1} + \frac{L_F \gamma^2}{2} E_{\text{var}}^{t+1}.$$

Re-arranging the terms we get the desired inequality.

A consequence of Lemma 3 is the following result.

Lemma 4 (Descent Lemma) *Suppose Assumption 5 holds true. By taking  $\gamma = \frac{1}{2L_F}$ , we have,*

$$\begin{aligned}\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ \|r F(\mathbf{x}^t)\|_2^2 \right] + \frac{1}{2T} \sum_{t=0}^{T-1} \mathbb{E} \left[ \|\mathbb{E}^t[G^{t+1} | j]\|_2^2 \right] \\ \leq \frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} + \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [E_{\text{bias}}^{t+1}] + \frac{L_F \gamma}{T} \sum_{t=0}^{T-1} \mathbb{E} [E_{\text{var}}^{t+1}]\end{aligned}$$

where the expectation is taken over all randomness from  $t = 0$  to  $T$ .

Proof: We denote the conditional expectation at time  $t$  in the descent lemma (Lemma 3) as  $\mathbb{E}^{t+1}[j | t]$  which conditions on all past randomness until time  $t$ . Then the descent lemma can be written as

$$\mathbb{E}^{t+1}[F(\mathbf{x}^{t+1}) | j | t] - F(\mathbf{x}^t) \leq \frac{\gamma}{2} \|r F(\mathbf{x}^t)\|_2^2 - \frac{\gamma}{4} \|\mathbb{E}^{t+1}[G^{t+1} | j | t]\|_2^2 + \frac{\gamma}{2} \mathbb{E}^{t+1}[E_{\text{bias}}^{t+1} | j | t] + \frac{L_F \gamma^2}{2} \mathbb{E}^{t+1}[E_{\text{var}}^{t+1} | j | t].$$

If we additionally consider the randomness at time  $t - 1$ , and apply  $\mathbb{E}^t[j | t - 1]$  to both sides

$$\begin{aligned}\mathbb{E}^t [\mathbb{E}^{t+1}[F(\mathbf{x}^{t+1}) | j | t] | j | t - 1] - \mathbb{E}^t [F(\mathbf{x}^t) | j | t - 1] &\leq \frac{\gamma}{2} \mathbb{E}^t [\|r F(\mathbf{x}^t)\|_2^2 | j | t - 1] \\ &\quad - \frac{\gamma}{4} \mathbb{E}^t [\|\mathbb{E}^{t+1}[G^{t+1} | j | t]\|_2^2 | j | t - 1] + \frac{\gamma}{2} \mathbb{E}^t [\mathbb{E}^{t+1}[E_{\text{bias}}^{t+1} | j | t] | j | t - 1] \\ &\quad + \frac{L_F \gamma^2}{2} \mathbb{E}^t [\mathbb{E}^{t+1}[E_{\text{var}}^{t+1} | j | t] | j | t - 1].\end{aligned}$$

By the law of iterative expectations, we have  $\mathbb{E}^t [\mathbb{E}^{t+1}[j | t] | j | t - 1] = \mathbb{E}^t [j | t - 1]$

$$\begin{aligned}\mathbb{E}^t [\mathbb{E}^{t+1} [F(\mathbf{x}^{t+1}) | j | t] | j | t - 1] - \mathbb{E}^t [F(\mathbf{x}^t) | j | t - 1] &\leq \frac{\gamma}{2} \mathbb{E}^t [\|r F(\mathbf{x}^t)\|_2^2 | j | t - 1] \\ &\quad - \frac{\gamma}{4} \mathbb{E}^t [\|\mathbb{E}^{t+1}[G^{t+1} | j | t]\|_2^2 | j | t - 1] + \frac{\gamma}{2} \mathbb{E}^t [\mathbb{E}^{t+1} [E_{\text{bias}}^{t+1} | j | t] | j | t - 1] \\ &\quad + \frac{L_F \gamma^2}{2} \mathbb{E}^t [\mathbb{E}^{t+1} [E_{\text{var}}^{t+1} | j | t] | j | t - 1].\end{aligned}$$

Similarly, we can apply  $\mathbb{E}^{t-1}[j | t - 2]$ ,  $\mathbb{E}^{t-2}[j | t - 3]$ , ...,  $\mathbb{E}^2[j | 1]$  and finally  $\mathbb{E}^1[ ]$

$$\begin{aligned}\mathbb{E}^1 \dots [\mathbb{E}^{t+1} [F(\mathbf{x}^{t+1})]] - \mathbb{E}^1 \dots [\mathbb{E}^t [F(\mathbf{x}^t)]] &\leq \frac{\gamma}{2} \mathbb{E}^1 \dots [\mathbb{E}^t [\|r F(\mathbf{x}^t)\|_2^2]] \\ \mathbb{E}^1 \dots [\mathbb{E}^t [\frac{\gamma}{4} \|\mathbb{E}^{t+1}[G^{t+1} | j | t]\|_2^2]] + \frac{\gamma}{2} \mathbb{E}^1 \dots [\mathbb{E}^{t+1} [E_{\text{bias}}^{t+1}]] \\ &\quad + \frac{L_F \gamma^2}{2} \mathbb{E}^1 \dots [\mathbb{E}^t [E_{\text{var}}^{t+1}]].\end{aligned}$$

Now that both sides of the inequality have no randomness, we can simplify the notation by applying  $\mathbb{E}^{t+1} \dots [\mathbb{E}^t [ ]]$  to both sides and by denoting

$$\mathbb{E}[ ] = \mathbb{E}^1 \dots [\mathbb{E}^{t+1} [ ]].$$

Then the descent lemma becomes

$$\mathbb{E}[F(\mathbf{x}^{t+1})] - \mathbb{E}[F(\mathbf{x}^t)] - \frac{\gamma}{2} \mathbb{E}[\|r F(\mathbf{x}^t)\|_2^2] - \frac{\gamma}{4} \mathbb{E}[\|\mathbb{E}^{t+1}[G^{t+1}j_t]\|_2^2] + \frac{\gamma}{2} \mathbb{E}[E_{\text{bias}}^{t+1}] + \frac{L_F \gamma^2}{2} \mathbb{E}[E_{\text{var}}^{t+1}].$$

Now we can sum the descent lemmas from  $t = 0$  to  $T - 1$

$$\begin{aligned} \sum_{t=0}^{T-1} \mathbb{E}[F(\mathbf{x}^{t+1})] - \sum_{t=0}^{T-1} \mathbb{E}[F(\mathbf{x}^t)] - \frac{\gamma}{2} \sum_{t=0}^{T-1} \mathbb{E} \left[ k r F(\mathbf{x}^t) k_2^2 \right] \\ - \frac{\gamma}{4} \sum_{t=0}^{T-1} \mathbb{E} \left[ \|\mathbb{E}^{t+1}[G^{t+1}j_t]\|_2^2 \right] + \frac{\gamma}{2} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{bias}}^{t+1}] + \frac{L_F \gamma^2}{2} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}]. \end{aligned}$$

After simplification and division by  $T$ , we get

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ k r F(\mathbf{x}^t) k_2^2 \right] + \frac{1}{2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ \|\mathbb{E}^{t+1}[G^{t+1}j_t]\|_2^2 \right] \\ \frac{2(\mathbb{E}[F(\mathbf{x}^T)] - \mathbb{E}[F(\mathbf{x}^0)])}{\gamma T} + \frac{\gamma}{2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{bias}}^{t+1}] + \frac{L_F \gamma^2}{2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}] \\ \frac{2(\mathbb{E}[F(\mathbf{x}^T)] - F^*)}{\gamma T} + \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{bias}}^{t+1}] + \frac{L_F \gamma}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}]. \end{aligned}$$

The following corollary is a consequence of Proposition 1.

Corollary 1 Assume  $r f_\xi$  in CSO satisfies

$$a_l := \sup_{\mathbf{x}} \sup_{\xi} \|r^{l+1} f_\xi(\mathbf{x})\|_2 < 1, \quad l = 1, 2, 3.$$

Let's further assume that the higher order moments of  $g_\eta(\cdot)$  are bounded,

$$\sigma_k = \sup_{\mathbf{x}} \sup_{\xi} \mathbb{E}_{\eta|\xi} \left[ \|g_\eta(\mathbf{x}) - \mathbb{E}_{\eta|\xi}[g_\eta(\mathbf{x})]\|_2^k \right] < 1, \quad k = 1 \dots 3.$$

Consider the  $L_{D_{g,\xi}^{(2)}}^{(2)} r f_\xi(0)$  defined in (6). Then  $\delta\xi$ ,

$$\left\| \mathbb{E} \left[ L_{D_{g,\xi}^{(2)}}^{(2)} r f_\xi(0) \right] - r f_\xi(\mathbb{E}_{\eta|\xi}[g_\eta(\mathbf{x}^t)]) \right\|_2^2 \leq \frac{C_e^2}{m^4},$$

where  $C_e^2(f; g) := a_3^2 \sigma_3^2$ .

## C.2 Convergence of BSGD

In this section, we reanalyze the BSGD algorithm of [Hu et al., 2020b] to obtain bounds on bias and variance of its gradient estimates. Theorem 5 shows that BSGD achieves an  $O(\varepsilon^{-6})$  sample complexity.

Lemma 5 (Bias and Variance of BSGD) *The bias and variance of BSGD are*

$$E_{\text{bias}}^{t+1} = \frac{\sigma_{\text{bias}}^2}{m}, \quad E_{\text{var}}^{t+1} = \frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2$$

where  $\sigma_{\text{in}}^2 = \zeta_g^2 C_f^2 + \sigma_g^2 C_g^2 L_f^2$ ,  $\sigma_{\text{out}}^2 = C_F^2$ , and  $\sigma_{\text{bias}}^2 = \sigma_g^2 C_g^2 L_f^2$ .

Proof: Denote  $G^{t+1} = G_{\text{BSGD}}^{t+1}$  (5) and denote  $\mathbb{E}[\cdot]$  as the conditional expectation  $\mathbb{E}^{t+1}[\cdot|j_t]$  which conditions on all past randomness until time  $t$ . Note that the  $r g_\eta$  can be estimated without bias, i.e.

$$\mathbb{E}_{\eta|\xi} \left[ \frac{1}{m} \sum_{\eta \sim \mathcal{H}_\xi} r g_\eta(\mathbf{x}) \right] = \mathbb{E}_{\eta|\xi} [r g_\eta(\mathbf{x})],$$

Then let's first look at the bias of BSGD

$$\begin{aligned}
E_{\text{bias}}^{t+1} &= \left\| r F(\mathbf{x}^{t+1}) - \mathbb{E}[G^{t+1}] \right\|_2^2 \\
&= \left\| \mathbb{E}_\xi \left[ \left( \mathbb{E}_{\eta|\xi} [r g_\eta(\mathbf{x}^t)] \right)^\triangleright \left( r f_\xi(\mathbb{E}_{\eta|\xi} [g_\eta(\mathbf{x}^t)]) - \mathbb{E}_{\eta|\xi} [r f_\xi(\frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t))] \right) \right] \right\|_2^2 \\
&\quad C_g^2 \mathbb{E}_\xi \left[ \left\| r f_\xi(\mathbb{E}_{\eta|\xi} [g_\eta(\mathbf{x}^t)]) - \mathbb{E}_{\eta|\xi} [r f_\xi(\frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t))] \right\|_2^2 \right] \\
&\quad C_g^2 L_f^2 \mathbb{E}_\xi \left[ \mathbb{E}_{\eta|\xi} \left[ \left\| \mathbb{E}_{\eta|\xi} [g_\eta(\mathbf{x}^t)] - \frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t) \right\|_2^2 \right] \right] \\
&\quad \frac{C_g^2 L_f^2}{m} \mathbb{E}_\xi \left[ \mathbb{E}_{\eta|\xi} \left[ \left\| \mathbb{E}_{\eta|\xi} [g_\eta(\mathbf{x}^t)] - g_\eta(\mathbf{x}^t) \right\|_2^2 \right] \right] \\
&= \frac{\sigma_g^2 C_g^2 L_f^2}{m} = \frac{\sigma_{\text{bias}}^2}{m}.
\end{aligned}$$

For the first inequality, we take the expectation outside the norm and bound  $r g_\eta$  with  $C_g$ .

On the other hand, the variance of BSGD can be decomposed into inner variance and outer variance

$$\begin{aligned}
E_{\text{var}}^{t+1} &= \mathbb{E}_\xi [\mathbb{E}_{\eta|\xi, \eta|\xi} [\|G^{t+1} - \mathbb{E}_\xi [\mathbb{E}_{\eta|\xi, \eta|\xi} [G^{t+1}]]\|_2^2]] \\
&= \mathbb{E}_\xi [\mathbb{E}_{\eta|\xi, \eta|\xi} [\| (G^{t+1} - \mathbb{E}_{\eta|\xi, \eta|\xi} [G^{t+1}]) + (\mathbb{E}_{\eta|\xi, \eta|\xi} [G^{t+1}] - \mathbb{E}_\xi [\mathbb{E}_{\eta|\xi, \eta|\xi} [G^{t+1}]] ) \|_2^2]] \\
&= \underbrace{\mathbb{E}_\xi [\mathbb{E}_{\eta|\xi, \eta|\xi} [\|G^{t+1} - \mathbb{E}_{\eta|\xi, \eta|\xi} [G^{t+1}]\|_2^2]]}_{\text{Inner variance}} + \underbrace{\mathbb{E}_\xi [\| \mathbb{E}_{\eta|\xi, \eta|\xi} [G^{t+1}] - \mathbb{E}_\xi [\mathbb{E}_{\eta|\xi, \eta|\xi} [G^{t+1}]] \|_2^2]}_{\text{Outer variance}}.
\end{aligned}$$

The inner variance is bounded as follows

$$\begin{aligned}
&\mathbb{E}_\xi [\mathbb{E}_{\eta|\xi, \eta|\xi} [\|G^{t+1} - \mathbb{E}_\xi [\mathbb{E}_{\eta|\xi, \eta|\xi} [G^{t+1}]]\|_2^2]] \\
&= \mathbb{E}_\xi \left[ \mathbb{E}_{\eta|\xi, \eta|\xi} \left[ \left\| \left( \frac{1}{m} \sum_{\eta \in H_\xi} r g_\eta(\mathbf{x}^t) - \mathbb{E}_{\eta|\xi} [r g_\eta(\mathbf{x}^t)] \right)^\triangleright r f_\xi(\frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t)) \right\|_2^2 \right] \right] \\
&\quad + \mathbb{E}_\xi \left[ \mathbb{E}_{\eta|\xi} \left[ \left\| \left( \mathbb{E}_{\eta|\xi} [r g_\eta(\mathbf{x}^t)] \right)^\triangleright \left( r f_\xi(\frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t)) - \mathbb{E}_{\eta|\xi} [r f_\xi(\frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t))] \right) \right\|_2^2 \right] \right] \\
&\quad C_f^2 \mathbb{E}_\xi \left[ \mathbb{E}_{\eta|\xi} \left[ \left\| \frac{1}{m} \sum_{\eta \in H_\xi} r g_\eta(\mathbf{x}^t) - \mathbb{E}_{\eta|\xi} [r g_\eta(\mathbf{x}^t)] \right\|_2^2 \right] \right] \\
&\quad + C_g^2 \mathbb{E}_\xi \left[ \mathbb{E}_{\eta|\xi} \left[ \left\| r f_\xi(\frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t)) - \mathbb{E}_{\eta|\xi} [r f_\xi(\frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t))] \right\|_2^2 \right] \right] \\
&= C_f^2 \mathbb{E}_\xi \left[ \mathbb{E}_{\eta|\xi} \left[ \left\| \frac{1}{m} \sum_{\eta \in H_\xi} r g_\eta(\mathbf{x}^t) - \mathbb{E}_{\eta|\xi} [r g_\eta(\mathbf{x}^t)] \right\|_2^2 \right] \right] \\
&\quad + C_g^2 \mathbb{E}_\xi \left[ \mathbb{E}_{\eta|\xi} \left[ \left\| r f_\xi(\frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t)) - r f_\xi(\mathbb{E}_\eta [g_\eta(\mathbf{x}^t)]) \right\|_2^2 \right] \right] \\
&\quad C_g^2 \mathbb{E}_\xi \left[ \left\| r f_\xi(\mathbb{E}_\eta [g_\eta(\mathbf{x}^t)]) - \mathbb{E}_{\eta|\xi} [r f_\xi(\frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t))] \right\|_2^2 \right] \\
&\quad \frac{\zeta_g^2 C_f^2}{m} + C_g^2 L_f \mathbb{E}_\xi \left[ \mathbb{E}_{\eta|\xi} \left[ \left\| \frac{1}{m} \sum_{\eta \in H_\xi} g_\eta(\mathbf{x}^t) - \mathbb{E}_{\eta|\xi} [g_\eta(\mathbf{x}^t)] \right\|_2^2 \right] \right] \\
&\quad \frac{C_f^2 \zeta_g^2}{m} + \frac{C_g^2 L_f}{m} \mathbb{E}_\xi \left[ \mathbb{E}_{\eta|\xi} \left[ \left\| g_\eta(\mathbf{x}^t) - \mathbb{E}_{\eta|\xi} [g_\eta(\mathbf{x}^t)] \right\|_2^2 \right] \right] \\
&\quad \frac{\zeta_g^2 C_f^2 + \sigma_g^2 C_g^2 L_f^2}{m} = \frac{\sigma_{\text{in}}^2}{m}.
\end{aligned}$$

The outer variance is independent of the inner batch size and can be bounded by

$$\mathbb{E}_\xi [\| \mathbb{E}_{\eta|\xi, \eta|\xi} [G^{t+1}] - \mathbb{E}_\xi [\mathbb{E}_{\eta|\xi, \eta|\xi} [G^{t+1}]] \|_2^2] = \mathbb{E}_\xi [\| \mathbb{E}_{\eta|\xi, \eta|\xi} [G^{t+1}] \|_2^2] - C_f^2 C_g^2 = C_F^2 = \sigma_{\text{out}}^2$$

Therefore, the variance is bounded as follows

$$E_{\text{var}}^{t+1} = \frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2.$$

This completes the proof.

Theorem 5 (BSGD Convergence) *Consider the (CSO) problem. Suppose Assumptions 3, 4, 5 holds true. Let step size  $\gamma = 1/(2L_F)$ . Then for BSGD,  $\mathbf{x}^s$  picked uniformly at random among  $\mathcal{F}\mathbf{x}^t \mathcal{G}_{t=0}^T$  satisfies:  $\mathbb{E}[k r F(\mathbf{x}^s) k_2^2] \leq \varepsilon^2$ , for nonconvex  $F$ , if the inner batch size  $m = \Omega(\sigma_{bias}^2 \varepsilon^{-2})$  and the number of iterations  $T = \Omega((F(\mathbf{x}^0) - F^*) L_F (\sigma_{in}^2/m + \sigma_{out}^2) \varepsilon^{-4})$ , where  $\sigma_{in}^2 = \zeta_g^2 C_f^2 + \sigma_g^2 C_g^2 L_f^2$ ,  $\sigma_{out}^2 = C_F^2$ , and  $\sigma_{bias}^2 = \sigma_g^2 C_g^2 L_f^2$ .*

Proof: Denote  $G^{t+1} = G_{\text{BSGD}}^{t+1}$  (1). Using descent lemma (Lemma 4) and bias-variance bounds of BSGD (Lemma 5)

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ k r F(\mathbf{x}^t) k_2^2 \right] + \frac{1}{2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ \left\| \mathbb{E}^{t+1}[G^{t+1} | t] \right\|_2^2 \right] \leq \frac{2(\mathbb{E}[F(\mathbf{x}^T)] - F^*)}{\gamma T} + L_F \gamma (\frac{\sigma_{in}^2}{m} + \sigma_{out}^2) + \frac{\sigma_{bias}^2}{m}$$

Then we can minimize the right-hand size by optimizing  $\gamma$  to

$$\gamma = \sqrt{\frac{2(F(\mathbf{x}^0) - F^*)}{L_F (\sigma_{in}^2/m + \sigma_{out}^2) T}}$$

which is smaller than the bound of step size  $\gamma = \frac{1}{2L_F}$  if  $T$  is greater than the following constant which does not rely on the target precision  $\varepsilon$

$$T \geq \frac{8L_F (F(\mathbf{x}^0) - F^*)}{\sigma_{in}^2/m + \sigma_{out}^2}.$$

Then the upper bound of gradient becomes

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ k r F(\mathbf{x}^t) k_2^2 \right] \leq \sqrt{\frac{2(F(\mathbf{x}^0) - F^*) L_F (\sigma_{in}^2/m + \sigma_{out}^2)}{T}} + \frac{\sigma_{bias}^2}{m}.$$

By taking inner batch size of at least

$$m \geq \frac{\sigma_{bias}^2}{\varepsilon^2},$$

and iteration  $T$  greater than

$$T \geq \frac{2(F(\mathbf{x}^0) - F^*) L_F (\sigma_{in}^2/m + \sigma_{out}^2)}{\varepsilon^4},$$

we have that

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ k r F(\mathbf{x}^t) k_2^2 \right] \leq 2\varepsilon^2.$$

By picking  $\mathbf{x}^s$  uniformly at random among  $\mathcal{F}\mathbf{x}^t \mathcal{G}_{t=0}^T$ , we get the desired guarantee.

The resulting sample complexity of BSGD to get to an  $\varepsilon$ -stationary point is  $O(\varepsilon^{-6})$ .

### C.3 Convergence of E-BSGD

In this section, we analyze the sample complexity of Algorithm 1 (E-BSGD) for the CSO problem.

Lemma 6 (Bias and Variance of E-BSGD) *The bias and variance of E-BSGD are*

$$E_{bias}^{t+1} = \frac{\sigma_{bias}^2}{m^4}, \quad E_{var}^{t+1} = 14(\frac{\sigma_{in}^2}{m} + \sigma_{out}^2)$$

where  $\sigma_{in}^2 = \zeta_g^2 C_f^2 + \sigma_g^2 C_g^2 L_f^2$ ,  $\sigma_{out}^2 = C_F^2$ , and  $\tilde{\sigma}_{bias}^2 = C_g^2 C_e^2$  with  $C_e^2$  defined in Corollary 1.

Proof: Denote  $G^{t+1} = G_{\text{E-BSGD}}^{t+1}$  (7). Like previously (Lemma 5), let  $\mathbb{E}[\cdot]$  denote the conditional expectation  $\mathbb{E}^{t+1}[\cdot | t]$  which conditions on all past randomness until time  $t$ . In E-BSGD, we apply extrapolation to  $r f_\xi(\cdot)$ . The bias can be estimated with the help of Corollary 1 as

$$\begin{aligned} E_{bias}^{t+1} &= \left\| r F(\mathbf{x}^{t+1}) - \mathbb{E}[G^{t+1}] \right\|_2^2 \\ &= C_g^2 \mathbb{E}_\xi \left[ \left\| r f_\xi(\mathbb{E}_{\eta | \xi}[g_\eta(\mathbf{x}^t)]) - \mathbb{E} \left[ L_{D_{\mathbf{g}, \xi}^{(2)}}^{(2)} r f_\xi(0) \right] \right\|_2^2 \right] \\ &\leq \frac{C_g^2 C_e^2}{m^4}. \end{aligned}$$



Since the variance of BSGD in Lemma 5 is upper bounded by  $\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2$ , then Lemma 2 gives

$$E_{\text{var}}^{t+1} \leq 14(\sigma_{\text{in}}^2/m + \sigma_{\text{out}}^2).$$

This proves the claimed bounds.

**Theorem 2 [E-BSGD Convergence]** Consider the (CSO) problem. Suppose Assumptions 3, 4, 5, 6 holds true and  $a_3, \sigma_3, L_F, C_F, \tilde{L}_F, C_g, F^*$  are the constants. Let step size  $\gamma = 1/(2L_F)$ . Then the output  $\mathbf{x}^s$  of E-BSGD (Algorithm 1) satisfies:  $\mathbb{E}[k\Gamma F(\mathbf{x}^s)k_2^2] \leq \varepsilon^2$ , for nonconvex  $F$ , if the inner batch size  $m = \Omega(a_3\sigma_3C_g\varepsilon^{-1/2})$ , and the number of iterations

$$T = \Omega(L_F(F(\mathbf{x}^0) - F^*)(\tilde{L}_F/m + C_F^2)\varepsilon^{-4}).$$

**Proof:** The proof is very similar to Theorem 5. Denote  $G^{t+1} = G_{\text{E-BSGD}}^{t+1}$  (7). Using descent lemma (Lemma 4) and bias-variance bounds of E-BSGD (Lemma 6)

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ k\Gamma F(\mathbf{x}^t)k_2^2 \right] + \frac{1}{2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ \|\mathbb{E}^{t+1}[G^{t+1}|t]\|_2^2 \right] \\ \leq \frac{2(\mathbb{E}[F(\mathbf{x}^T)] - F^*)}{\gamma T} + 14L_F\gamma(\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2) + \frac{C_g^2 C_\varepsilon^2}{m^4}. \end{aligned}$$

Then we optimize  $\gamma$  to

$$\gamma = \sqrt{\frac{(F(\mathbf{x}^0) - F^*)}{7L_F(\sigma_{\text{in}}^2/m + \sigma_{\text{out}}^2)T}}$$

which is smaller than the bound of step size  $\gamma = \frac{1}{2L_F}$  if  $T$  is greater than the following constant which does not rely on the target precision  $\varepsilon$

$$T \geq \frac{4L_F(F(\mathbf{x}^0) - F^*)}{7(\sigma_{\text{in}}^2/m + \sigma_{\text{out}}^2)}.$$

Then the gradient norm has the following upper bound.

$$\frac{1}{T} \sum_{t=0}^{T-1} k\Gamma F(\mathbf{x}^t)k_2^2 \leq 4\sqrt{\frac{7(F(\mathbf{x}^0) - F^*)L_F(\sigma_{\text{in}}^2 + \sigma_{\text{out}}^2)}{T}} + \frac{\sigma_{\text{bias}}^2}{m^4}.$$

In order to reach  $\varepsilon$ -stationary point, i.e.

$$\frac{1}{T} \sum_{t=0}^{T-1} k\Gamma F(\mathbf{x}^t)k_2^2 \leq \varepsilon^2,$$

we can enforce

$$4\sqrt{\frac{7(F(\mathbf{x}^0) - F^*)L_F(\sigma_{\text{in}}^2/m + \sigma_{\text{out}}^2)}{T}} \leq \varepsilon^2, \quad \frac{C_g^2 C_\varepsilon^2}{m^4} \leq \varepsilon^2.$$

By taking inner batch size of at least

$$m = \Omega(\tilde{\sigma}_{\text{bias}}^{1/2} \varepsilon^{-1/2}),$$

and iteration  $T$  greater than

$$T \geq \frac{112(F(\mathbf{x}^0) - F^*)L_F(\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2)}{\varepsilon^4},$$

we have that

$$\frac{1}{T} \sum_{t=0}^{T-1} k\Gamma F(\mathbf{x}^t)k_2^2 \leq 3\varepsilon^2.$$

By picking  $\mathbf{x}^s$  uniformly at random among  $\{\mathbf{x}^t\}_{t=0}^{T-1}$ , we get the desired guarantee.

## C.4 Convergence of BSpiderBoost

In this section, we reanalyze the BSpiderBoost algorithm of [Hu et al., 2020b] to obtain bounds on bias and variance of its gradient estimates. Theorem 5 shows that BSpiderBoost achieves an  $O(\varepsilon^{-5})$  sample complexity.

Let  $G_{\text{BSB}}^{t+1}$  as the BSpiderBoost gradient estimate

$$G_{\text{BSB}}^{t+1} = \begin{cases} G_{\text{BSB}}^t + \frac{1}{B_2} \sum_{\xi \in 2B_2} (G_{\text{BSGD}}^{t+1} - G_{\text{BSGD}}^t) & \text{with prob. } 1 - p_{\text{out}} \\ \frac{1}{B_1} \sum_{\xi \in 2B_1} G_{\text{BSGD}}^{t+1} & \text{with prob. } p_{\text{out}}. \end{cases}$$

Lemma 7 (Bias and Variance of BSpiderBoost) *If  $\gamma \leq \min\{f \frac{1}{2L_F}, \frac{\rho \overline{B_2}}{6L_F} g\}$ , then the bias and variance of BSpiderBoost are*

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{bias}}^{t+1}] &= \frac{2\sigma_{\text{bias}}^2}{m} + \frac{(1-p_{\text{out}})^3}{p_{\text{out}}B_2} \frac{56L_F^2\gamma^2}{T} \sum_{t=0}^{T-1} \mathbb{E}[\| \mathbb{E}^{t+1}[G^{t+1}jt] \|_2^2] + \left(\frac{1}{Tp_{\text{out}}} + 1\right) \frac{4(1-p_{\text{out}})^2}{B_1} (\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2) \\ \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}] &= \frac{28(1-p_{\text{out}})L_F^2\gamma^2}{B_2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\| \mathbb{E}^{t+1}[G^{t+1}jt] \|_2^2] + \left(\frac{1}{T} + p_{\text{out}}\right) \frac{2}{B_1} (\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2), \end{aligned}$$

where  $\sigma_{\text{in}}^2 = \zeta_g^2 C_f^2 + \sigma_g^2 C_g^2 L_f^2$ ,  $\sigma_{\text{out}} = C_F^2$ , and  $\sigma_{\text{bias}}^2 = \sigma_g^2 C_g^2 L_f^2$ .

Proof: Denote  $G^{t+1} = G_{\text{BSB}}^{t+1}$  (8). Like previously (Lemma 5), let  $\mathbb{E}[\cdot]$  denote the conditional expectation  $\mathbb{E}^{t+1}[\cdot]$  which conditions on all past randomness until time  $t$ . Denote  $G_L^{t+1}$  and  $G_S^{t+1}$  as the large batch and small batch in BSpiderBoost separately, i.e.,

$$\begin{cases} G_L^{t+1} = \frac{1}{B_1} \sum_{\xi \in 2B_1} G_{\text{BSGD}}^{t+1} & \text{with prob. } p_{\text{out}} \\ G_S^{t+1} = G^t + \frac{1}{B_2} \sum_{\xi \in 2B_2} (G_{\text{BSGD}}^{t+1} - G_{\text{BSGD}}^t) & \text{with prob. } 1 - p_{\text{out}}. \end{cases}$$

The bias of BSpiderBoost can be decomposed to its distance to BSGD and the distance from BSGD to the full gradient, i.e.,

$$\begin{aligned} E_{\text{bias}}^{t+1} &= \left\| r F(\mathbf{x}^{t+1}) - \mathbb{E}[G^{t+1}] \right\|_2^2 \\ &= 2 \left\| r F(\mathbf{x}^{t+1}) - \mathbb{E}[G_{\text{BSGD}}^{t+1}] \right\|_2^2 + 2 \left\| \mathbb{E}[G_{\text{BSGD}}^{t+1}] - \mathbb{E}[G^{t+1}] \right\|_2^2 \\ &= \frac{2\sigma_{\text{bias}}^2}{m} + 2 \left\| \mathbb{E}[G_{\text{BSGD}}^{t+1}] - \mathbb{E}[G^{t+1}] \right\|_2^2. \end{aligned} \tag{14}$$

where the last inequality uses the bias of BSGD from Lemma 5. Then the second term can be bounded as follows

$$\begin{aligned} \left\| \mathbb{E}[G_{\text{BSGD}}^{t+1}] - \mathbb{E}[G^{t+1}] \right\|_2^2 &= (1 - p_{\text{out}})^2 \left\| \mathbb{E}[G_{\text{BSGD}}^{t+1}] - \mathbb{E}[G_S^{t+1}] \right\|_2^2 \\ &= (1 - p_{\text{out}})^2 \left\| \mathbb{E}[G_{\text{BSGD}}^t] - G^t \right\|_2^2. \end{aligned}$$

By taking the expectation of randomness of  $G^t$

$$\begin{aligned} \left\| \mathbb{E}[G_{\text{BSGD}}^{t+1}] - \mathbb{E}[G^{t+1}] \right\|_2^2 &= (1 - p_{\text{out}})^2 \left( \left\| \mathbb{E}[G_{\text{BSGD}}^t] - \mathbb{E}[G^t] \right\|_2^2 + \mathbb{E} \left\| G^t - \mathbb{E}[G^t] \right\|_2^2 \right) \\ &= (1 - p_{\text{out}})^2 \left( \left\| \mathbb{E}[G_{\text{BSGD}}^t] - \mathbb{E}[G^t] \right\|_2^2 + E_{\text{var}}^t \right) \end{aligned}$$

Note that  $\left\| \mathbb{E}[G_{\text{BSGD}}^1] - \mathbb{E}[G^1] \right\|_2^2 = 0$  as the first iteration always chooses the large batch. Then as we always use large batch at  $t = 0$  we know that

$$\frac{1}{T} \sum_{t=0}^{T-1} \left\| \mathbb{E}[G_{\text{BSGD}}^{t+1}] - \mathbb{E}[G^{t+1}] \right\|_2^2 \leq \frac{(1-p_{\text{out}})^2}{p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1}. \tag{15}$$

Therefore combine (14) and (15) we can upper bound the bias

$$\frac{1}{T} \sum_{t=0}^{T-1} E_{\text{bias}}^{t+1} \leq \frac{2\sigma_{\text{bias}}^2}{m} + \frac{2(1-p_{\text{out}})^2}{p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1}. \tag{16}$$

Variance. Now we consider the variance,

$$\begin{aligned}
E_{\text{var}}^{t+1} &= \mathbb{E} \left[ \left\| G^{t+1} - \mathbb{E}[G^{t+1}] \right\|_2^2 \right] \\
&= (1 - p_{\text{out}}) \mathbb{E} \left[ \left\| G_S^{t+1} - \mathbb{E}[G_S^{t+1}] \right\|_2^2 \right] + p_{\text{out}} \mathbb{E} \left[ \left\| G_L^{t+1} - \mathbb{E}[G_L^{t+1}] \right\|_2^2 \right] \\
&= \frac{(1 - p_{\text{out}})}{B_2} \mathbb{E} \left[ \left\| G_{\text{BSGD}}^{t+1} - G_{\text{BSGD}}^t - \mathbb{E}[G_{\text{BSGD}}^{t+1} - G_{\text{BSGD}}^t] \right\|_2^2 \right] + \frac{p_{\text{out}}}{B_1} \mathbb{E} \left[ \left\| G_{\text{BSGD}}^{t+1} - \mathbb{E}[G_{\text{BSGD}}^{t+1}] \right\|_2^2 \right] \\
&= \frac{1 - p_{\text{out}}}{B_2} \mathbb{E} \left[ \left\| G_{\text{BSGD}}^{t+1} - G_{\text{BSGD}}^t - \mathbb{E}[G_{\text{BSGD}}^{t+1} - G_{\text{BSGD}}^t] \right\|_2^2 \right] + \frac{p_{\text{out}}}{B_1} (\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2)
\end{aligned} \tag{17}$$

where the last equality is because the large batch in BSpiderBoost is similar to BSGD.

$$E_{\text{var}}^1 = \mathbb{E} \left[ \left\| G^1 - \mathbb{E}[G^1] \right\|_2^2 \right] = \mathbb{E} \left[ \left\| G_L^1 - \mathbb{E}[G_L^1] \right\|_2^2 \right] = \frac{1}{B_1} \mathbb{E} \left[ \left\| G_{\text{BSGD}}^1 - \mathbb{E}[G_{\text{BSGD}}^1] \right\|_2^2 \right] = \frac{1}{B_1} (\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2). \tag{18}$$

Finally, we expand the variance at small batch size epoch

$$\begin{aligned}
&\mathbb{E} \left[ \left\| G_{\text{BSGD}}^{t+1} - G_{\text{BSGD}}^t - \mathbb{E}[G_{\text{BSGD}}^{t+1} - G_{\text{BSGD}}^t] \right\|_2^2 \right] \\
&= \underbrace{\mathbb{E} \left[ \left\| G_{\text{BSGD}}^{t+1} - G_{\text{BSGD}}^t - \mathbb{E}_{\eta, \xi, \eta, \xi} [G_{\text{BSGD}}^{t+1} - G_{\text{BSGD}}^t] \right\|_2^2 \right]}_{\text{Inner variance } T_{\text{in}}} \\
&\quad + \underbrace{\mathbb{E}_{\xi} \left[ \left\| \mathbb{E}_{\eta, \xi, \eta, \xi} [G_{\text{BSGD}}^{t+1} - G_{\text{BSGD}}^t] - \mathbb{E}[G_{\text{BSGD}}^{t+1} - G_{\text{BSGD}}^t] \right\|_2^2 \right]}_{\text{Outer variance } T_{\text{out}}}.
\end{aligned}$$

The outer variance  $T_{\text{out}}$  can be upper bounded as

$$\begin{aligned}
T_{\text{out}} &= \mathbb{E}_{\xi} \left[ \left\| \mathbb{E}_{\eta, \xi, \eta, \xi} [G_{\text{BSGD}}^{t+1} - G_{\text{BSGD}}^t] \right\|_2^2 \right] \\
&= \mathbb{E}_{\xi} \left[ \left\| \left( \mathbb{E}_{\eta, \xi} [r g_{\eta}(\mathbf{x}^t)] \right)^{\triangleright} \mathbb{E}_{\eta, \xi} [r f_{\xi}(\frac{1}{m} \sum_{\eta \in H_{\xi}} g_{\eta}(\mathbf{x}^t))] - \left( \mathbb{E}_{\eta, \xi} [r g_{\eta}(\mathbf{x}^{t-1})] \right)^{\triangleright} \mathbb{E}_{\eta, \xi} [r f_{\xi}(\frac{1}{m} \sum_{\eta \in H_{\xi}} g_{\eta}(\mathbf{x}^{t-1}))] \right\|_2^2 \right] \\
&= 2 \mathbb{E}_{\xi} \left[ \left\| \left( \mathbb{E}_{\eta, \xi} [r g_{\eta}(\mathbf{x}^t)] - \mathbb{E}_{\eta, \xi} [r g_{\eta}(\mathbf{x}^{t-1})] \right)^{\triangleright} \mathbb{E}_{\eta, \xi} [r f_{\xi}(\frac{1}{m} \sum_{\eta \in H_{\xi}} g_{\eta}(\mathbf{x}^t))] \right\|_2^2 \right] \\
&\quad + 2 \mathbb{E}_{\xi} \left[ \left\| \left( \mathbb{E}_{\eta, \xi} [r g_{\eta}(\mathbf{x}^{t-1})] \right)^{\triangleright} \mathbb{E}_{\eta, \xi} [r f_{\xi}(\frac{1}{m} \sum_{\eta \in H_{\xi}} g_{\eta}(\mathbf{x}^t)) - r f_{\xi}(\frac{1}{m} \sum_{\eta \in H_{\xi}} g_{\eta}(\mathbf{x}^{t-1}))] \right\|_2^2 \right] \\
&= 2L_g^2 C_f^2 \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 + 2C_g^4 L_f^2 \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 \\
&= 2L_F^2 \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 \\
&= 2L_F^2 \gamma^2 \|G^t\|_2^2.
\end{aligned}$$

The inner variance can be bounded by

$$\begin{aligned}
\tau_{\text{in}} & \leq 4 \mathbb{E} \left[ \left\| \left( \frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} (r g_\eta(\mathbf{x}^t) - r g_\eta(\mathbf{x}^{t-1})) - \mathbb{E}_{\eta|\xi} [r g_\eta(\mathbf{x}^t) - r g_\eta(\mathbf{x}^{t-1})] \right) \right\|_2^2 \right] \\
& + 4 \mathbb{E} \left[ \left\| \left( \mathbb{E}_{\eta|\xi} [r g_\eta(\mathbf{x}^t) - r g_\eta(\mathbf{x}^{t-1})] - (r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^t)) - \mathbb{E}_{\eta|\xi} [r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^t))]) \right) \right\|_2^2 \right] \\
& + 4 \mathbb{E} \left[ \left\| \left( \frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} (r g_\eta(\mathbf{x}^t) - r g_\eta(\mathbf{x}^{t-1})) - \left( r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^t)) - r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^{t-1})) \right) \right. \right. \\
& \quad \left. \left. - \mathbb{E}_{\eta|\xi} \left[ r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^t)) - r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^{t-1})) \right] \right) \right\|_2^2 \right] \\
& + 4 \mathbb{E} \left[ \left\| \left( \frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} (r g_\eta(\mathbf{x}^t) - r g_\eta(\mathbf{x}^{t-1})) - \mathbb{E}_{\eta|\xi} [r g_\eta(\mathbf{x}^t) - r g_\eta(\mathbf{x}^{t-1})] \right) \right. \right. \\
& \quad \left. \left. - \left( r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^t)) - r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^{t-1})) \right) \right\|_2^2 \right] \\
& \frac{4C_g^2}{m} \mathbb{E} \left[ \left\| r g_\eta(\mathbf{x}^t) - r g_\eta(\mathbf{x}^{t-1}) - \mathbb{E}_{\eta|\xi} [r g_\eta(\mathbf{x}^t) - r g_\eta(\mathbf{x}^{t-1})] \right\|_2^2 \right] + \frac{4L_g^2 C_f^2}{m} \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 \\
& + 4C_g^2 \mathbb{E} \left[ \left\| r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^t)) - r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^{t-1})) \right. \right. \\
& \quad \left. \left. - \mathbb{E}_{\eta|\xi} [r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^t)) - r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^{t-1}))] \right\|_2^2 \right] \\
& + \frac{4C_g^4 L_f^2}{m} \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 \\
& \frac{4L_g^2 C_f^2}{m} \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 + \frac{4L_g^2 C_f^2}{m} \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 \\
& + 4C_g^2 \mathbb{E} \left[ \left\| r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^t)) - r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^{t-1})) - (r f_\xi(\mathbb{E}_{\eta|\xi} [g_\eta(\mathbf{x}^t)]) - r f_\xi(\mathbb{E}_{\eta|\xi} [g_\eta(\mathbf{x}^{t-1})])) \right\|_2^2 \right] \\
& + 4C_g^2 \mathbb{E} \left[ \left\| (r f_\xi(\mathbb{E}_{\eta|\xi} [g_\eta(\mathbf{x}^t)]) - r f_\xi(\mathbb{E}_{\eta|\xi} [g_\eta(\mathbf{x}^{t-1})])) \right. \right. \\
& \quad \left. \left. - \mathbb{E}_{\eta|\xi} [r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^t)) - r f_\xi(\frac{1}{m} \sum_{\eta \in \mathcal{H}_\xi} g_\eta(\mathbf{x}^{t-1}))] \right\|_2^2 \right] \\
& + \frac{4C_g^4 L_f^2}{m} \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 \\
& \frac{8L_g^2 C_f^2}{m} \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 + \frac{8C_g^4 L_f^2}{m} \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 + \frac{4C_g^4 L_f^2}{m} \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 \\
& \frac{12L_F^2}{m} \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 \\
& = \frac{12L_F^2 \gamma^2}{m} \|G^t\|_2^2.
\end{aligned}$$

To sum up, the variance is bounded by

$$\begin{aligned}
\mathbb{E}_{\text{var}}^{t+1} & \leq \frac{2(1 - p_{\text{out}})L_F^2 \gamma^2}{B_2} (1 + \frac{6}{m}) \|G^t\|_2^2 + \frac{p_{\text{out}}}{B_1} (\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2) \\
& \leq \frac{14(1 - p_{\text{out}})L_F^2 \gamma^2}{B_2} \|G^t\|_2^2 + \frac{p_{\text{out}}}{B_1} (\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2) \\
& = \frac{14(1 - p_{\text{out}})L_F^2 \gamma^2}{B_2} \mathbb{E}^t [\|G^t\|_2^2] + \frac{p_{\text{out}}}{B_1} (\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2).
\end{aligned}$$

Then averaging over time and now we redefine  $\mathbb{E}[\cdot] = \mathbb{E}^T \dots [\mathbb{E}^0[\cdot]]$

$$\begin{aligned}
\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\mathbb{E}_{\text{var}}^{t+1}] & \leq \frac{14(1 - p_{\text{out}})L_F^2 \gamma^2}{B_2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] + \frac{\mathbb{E}_{\text{var}}^1}{T} + \frac{p_{\text{out}}}{B_1} (\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2) \\
& = \frac{14(1 - p_{\text{out}})L_F^2 \gamma^2}{B_2} \frac{1}{T} \sum_{t=0}^{T-1} (\mathbb{E}[\mathbb{E}_{\text{var}}^{t+1}] + \mathbb{E}[\|G^{t+1} j^t\|_2^2]) + \frac{\mathbb{E}_{\text{var}}^1}{T} + \frac{p_{\text{out}}}{B_1} (\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2).
\end{aligned}$$

If we take  $\gamma = \frac{p_{\text{out}} B_2}{6L_F}$ , then  $\frac{14(1 - p_{\text{out}})L_F^2 \gamma^2}{B_2} \leq \frac{1}{2}$ , therefore

$$\begin{aligned}
\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\mathbb{E}_{\text{var}}^{t+1}] & \leq \frac{28(1 - p_{\text{out}})L_F^2 \gamma^2}{B_2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1} j^t\|_2^2] + \frac{\mathbb{E}_{\text{var}}^1}{T} + \frac{2p_{\text{out}}}{B_1} (\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2) \\
& \stackrel{(18)}{\leq} \frac{28(1 - p_{\text{out}})L_F^2 \gamma^2}{B_2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1} j^t\|_2^2] + (\frac{1}{T} + p_{\text{out}}) \frac{2}{B_1} (\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2)
\end{aligned}$$

Then with (16), we can bound the bias by

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{bias}}^{t+1}] \leq \frac{2\sigma_{\text{bias}}^2}{m} + \frac{(1-p_{\text{out}})^3}{p_{\text{out}}B_2} \frac{56L_F^2\gamma^2}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbb{E}^t[G^{t+1}j_t]\|_2^2] + \left(\frac{1}{Tp_{\text{out}}} + 1\right) \frac{4(1-p_{\text{out}})^2}{B_1} (\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2)$$

Theorem 6 (BSpiderBoost Convergence) *Consider the (CSO) problem. Suppose Assumptions 3, 4, 5 holds true. Let step size  $\gamma = 1/(13L_F)$ . Then for BSpiderBoost,  $\mathbf{x}^s$  picked uniformly at random among  $f\mathbf{x}^t g_{t=0}^{T-1}$  satisfies:  $\mathbb{E}[k\mathcal{R} F(\mathbf{x}^s)k_2^2] \leq \varepsilon^2$ , for nonconvex  $F$ , if the inner batch size  $m = \Omega(\sigma_{\text{bias}}^2 \varepsilon^{-2})$ , the hyperparameters of the outer loop of BSpiderBoost are  $B_1 = (\sigma_{\text{in}}^2/m + \sigma_{\text{out}}^2)\varepsilon^{-2}$ ,  $B_2 = O(\varepsilon^{-1})$ ,  $p_{\text{out}} = 1/B_2$ , and the number of iterations  $T = \Omega(L_F(F(\mathbf{x}^0) - F^*)\varepsilon^{-2})$ , where  $\sigma_{\text{in}}^2 = \zeta_g^2 C_f^2 + \sigma_g^2 C_g^2 L_f^2$ ,  $\sigma_{\text{out}}^2 = C_F^2$ , and  $\sigma_{\text{bias}}^2 = \sigma_g^2 C_g^2 L_f^2$ .*

Proof: Denote  $G^{t+1} = G_{\text{BSB}}^{t+1}$  (8). Using descent lemma (Lemma 4) and bias-variance bounds of BSpiderBoost (Lemma 7)

$$\begin{aligned} & \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[k\mathcal{R} F(\mathbf{x}^t)k_2^2] + \frac{1}{2T} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbb{E}^t[G^{t+1}j_t]\|_2^2] \\ & \leq \frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} + L_F \gamma \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}] + \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{bias}}^{t+1}] \\ & \leq \frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} + L_F \gamma \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}] + \frac{2\sigma_{\text{bias}}^2}{m} + \frac{2}{p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}] \\ & \leq \frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} + \frac{2\sigma_{\text{bias}}^2}{m} + \frac{3}{p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}] \end{aligned}$$

where the last inequality use  $\gamma = \frac{1}{2L_F}$ . Use the variance estimation of  $G^{t+1}$  and choose  $B_2 p_{\text{out}} = 1$

$$\begin{aligned} & \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[k\mathcal{R} F(\mathbf{x}^t)k_2^2] + \frac{1}{2T} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbb{E}^t[G^{t+1}j_t]\|_2^2] \\ & \leq \frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} + \frac{2\sigma_{\text{bias}}^2}{m} + 84L_F^2\gamma^2 \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbb{E}^t[G^{t+1}j_t]\|_2^2] + \left(\frac{1}{Tp_{\text{out}}} + 1\right) \frac{6}{B_1} (\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2). \end{aligned}$$

Now we can let  $\gamma = \frac{1}{13L_F}$  such that  $84L_F^2\gamma^2 = \frac{1}{2}$

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[k\mathcal{R} F(\mathbf{x}^t)k_2^2] \leq \frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} + \frac{2\sigma_{\text{bias}}^2}{m} + \left(\frac{1}{Tp_{\text{out}}} + 1\right) \frac{6}{B_1} (\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2).$$

In order for the right-hand side to be  $\varepsilon^2$ , the inner batch size

$$m \geq \frac{2\sigma_{\text{bias}}^2}{\varepsilon^2},$$

and the outer batch size

$$B_1 = \frac{\sigma_{\text{in}}^2/m + \sigma_{\text{out}}^2}{\varepsilon^2}, \quad B_2 = \sqrt{B_1}, \quad p_{\text{out}} = \frac{1}{B_2}.$$

The step size  $\gamma$  is upper bounded by  $\min\{\frac{1}{2L_F}, \frac{1}{6L_F}, \frac{1}{13L_F}\}$ . As  $B_2 \geq 1$ , we can take  $\gamma = \frac{1}{13L_F}$ . So we need iteration  $T$  greater than

$$T \geq \frac{26L_F(F(\mathbf{x}^0) - F^*)}{\varepsilon^2}.$$

By picking  $\mathbf{x}^s$  uniformly at random among  $f\mathbf{x}^t g_{t=0}^{T-1}$ , we get the desired guarantee.

The resulting sample complexity of BSpiderBoost to get to an  $\varepsilon$ -stationary point is  $O(\varepsilon^{-5})$ .

## C.5 Convergence of E-BSpiderBoost

In this section, we analyze the sample complexity of Algorithm 2 (E-BSpiderBoost) for the CSO problem.

Lemma 8 (Bias and Variance of E-BSpiderBoost) *The bias and variance of E-BSpiderBoost are*

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}] & \leq \frac{28(1-p_{\text{out}})L_F^2\gamma^2}{B_2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbb{E}^t[G^{t+1}j_t]\|_2^2] + \left(\frac{1}{Tp_{\text{out}}} + 1\right) \frac{28p_{\text{out}}}{B_1} (\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2) \\ \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{bias}}^{t+1}] & \leq \frac{2\sigma_{\text{bias}}^2}{m^4} + \frac{2}{p_{\text{out}}} \frac{(1-p_{\text{out}})^2}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}], \end{aligned}$$

where  $\sigma_{\text{in}}^2 := \zeta_g^2 C_f^2 + \sigma_g^2 C_g^2 L_f^2$ ,  $\sigma_{\text{out}} = C_F^2$ , and  $\tilde{\sigma}_{\text{bias}}^2 = C_g^2 C_e^2$  with  $C_e^2$  defined in Corollary 1.

Proof: Denote  $G^{t+1} = G_{\text{E-BSB}}^{t+1}$  (9). Like previously (Lemma 5), let  $\mathbb{E}[\cdot]$  denote the conditional expectation  $\mathbb{E}^t[\cdot]$  which conditions on all past randomness until time  $t$ . Let  $G^{t+1} = G_{\text{E-BSGD}}^{t+1}$  be the E-BSGD update. We expand the bias as follows

$$\begin{aligned} E_{\text{bias}}^{t+1} &= \left\| r F(\mathbf{x}^{t+1}) - \mathbb{E}[G^{t+1}] \right\|_2^2 \\ &= 2 \left\| r F(\mathbf{x}^{t+1}) - \mathbb{E}[G_{\text{E-BSGD}}^{t+1}] \right\|_2^2 + 2 \left\| \mathbb{E}[G_{\text{E-BSGD}}^{t+1}] - \mathbb{E}[G^{t+1}] \right\|_2^2. \end{aligned}$$

From Lemma 6, we know that

$$\left\| r F(\mathbf{x}^{t+1}) - \mathbb{E}[G_{\text{E-BSGD}}^{t+1}] \right\|_2^2 \leq \frac{\sigma_{\text{bias}}^2}{m^4}.$$

The distance between  $\mathbb{E}[G_{\text{E-BSGD}}^{t+1}]$  and  $\mathbb{E}[G^{t+1}]$  can be bounded as follows.

$$\begin{aligned} \left\| \mathbb{E}[G_{\text{E-BSGD}}^{t+1}] - \mathbb{E}[G^{t+1}] \right\|_2^2 &= (1 - p_{\text{out}})^2 \left\| \mathbb{E}[G_{\text{E-BSGD}}^{t+1}] - (G^t + \mathbb{E}[G_{\text{E-BSGD}}^{t+1}] - G_{\text{E-BSGD}}^t) \right\|_2^2 \\ &= (1 - p_{\text{out}})^2 \left\| \mathbb{E}[G_{\text{E-BSGD}}^{t+1}] - G^t \right\|_2^2 \end{aligned}$$

Taking expectation with respect to  $G^t$

$$\left\| \mathbb{E}[G_{\text{E-BSGD}}^{t+1}] - \mathbb{E}[G^{t+1}] \right\|_2^2 \leq (1 - p_{\text{out}})^2 \left( \left\| \mathbb{E}[G_{\text{E-BSGD}}^t] - \mathbb{E}[G^t] \right\|_2^2 + \left\| G^t - \mathbb{E}[G^t] \right\|_2^2 \right).$$

where  $\left\| \mathbb{E}[G_{\text{E-BSGD}}^1] - \mathbb{E}[G^1] \right\|_2^2 = 0$ . By averaging over time we have

$$\frac{1}{T} \sum_{t=0}^{T-1} \left\| \mathbb{E}[G_{\text{E-BSGD}}^{t+1}] - \mathbb{E}[G^{t+1}] \right\|_2^2 \leq \frac{1 - p_{\text{out}}}{p_{\text{out}}} \frac{(1 - p_{\text{out}})^2}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}].$$

Then the bias is bounded by

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{bias}}^{t+1}] \leq \frac{2\sigma_{\text{bias}}^2}{m^4} + \frac{2}{p_{\text{out}}} \frac{(1 - p_{\text{out}})^2}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}].$$

Variance. Since the extrapolation only gives a constant overhead given Lemma 2

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ \left\| G_{\text{BSB}}^{t+1} - \mathbb{E}^t[G_{\text{BSB}}^{t+1} | \mathcal{I}^t] \right\|_2^2 \right] \leq \frac{28(1 - p_{\text{out}})L_F^2\gamma^2}{B_2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ \left\| \mathbb{E}^t[G_{\text{BSB}}^{t+1} | \mathcal{I}^t] \right\|_2^2 \right] + \left( \frac{1}{T} + p_{\text{out}} \right) \frac{28}{B_1} (\sigma_{\text{in}}^2 + \sigma_{\text{out}}^2).$$

Then the variance is bounded by

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}] \leq \frac{28(1 - p_{\text{out}})L_F^2\gamma^2}{B_2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ \left\| \mathbb{E}^t[G^{t+1} | \mathcal{I}^t] \right\|_2^2 \right] + \left( \frac{1}{T p_{\text{out}}} + 1 \right) \frac{28 p_{\text{out}}}{B_1} (\sigma_{\text{in}}^2 + \sigma_{\text{out}}^2).$$

**Theorem 3 [E-BSGD Convergence]** Consider the (CSO) problem under the same assumptions as Theorem 2. Let step size  $\gamma = 1/(13L_F)$ . Then the output  $\mathbf{x}^s$  of E-BSGD (Algorithm 2) satisfies:  $\mathbb{E}[\|r F(\mathbf{x}^s) - k_2^2\|] \leq \varepsilon^2$ , for nonconvex  $F$ , if the inner batch size  $m = \mathcal{O}(a_3 \sigma_3 C_g \varepsilon^{-0.5})$ , the hyperparameters of the outer loop of E-BSGD  $B_1 = (\tilde{L}_F^2/m + C_F^2)\varepsilon^{-2}$ ,  $B_2 = \frac{1}{B_1}$ ,  $p_{\text{out}} = 1/B_2$ , and the number of iterations

$$T = \Omega(L_F(F(\mathbf{x}^0) - F^*)\varepsilon^{-2}).$$

Proof: Denote  $G^{t+1} = G_{\text{E-BSB}}^{t+1}$  (9). Using descent lemma (Lemma 4) and bias-variance bounds of E-BSGD (Lemma 8)

$$\begin{aligned} & \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|r F(\mathbf{x}^t) - k_2^2\|] + \frac{1}{2T} \sum_{t=0}^{T-1} \mathbb{E}[\left\| \mathbb{E}^t[G^{t+1} | \mathcal{I}^t] \right\|_2^2] \\ & \leq \frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} + L_F \gamma \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}] + \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{bias}}^{t+1}] \\ & \leq \frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} + L_F \gamma \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}] + \frac{2\sigma_{\text{bias}}^2}{m^4} + \frac{2}{p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}] \\ & \leq \frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} + \frac{2\sigma_{\text{bias}}^2}{m^4} + \frac{3}{p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_{\text{var}}^{t+1}] \end{aligned}$$

where the last inequality use  $\gamma \leq \frac{1}{2L_F}$ . Use the variance estimation of  $G^{t+1}$  and choose  $B_2 p_{\text{out}} = 1$

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[k^T F(\mathbf{x}^t) k_2^2] + \frac{1}{2T} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbb{E}^t[G^{t+1} j_t]\|_2^2] \\ \leq \frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} + \frac{2\sigma_{\text{bias}}^2}{m^4} + 84L_F^2 \gamma^2 \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbb{E}^t[G^{t+1} j_t]\|_2^2] + \left(\frac{1}{Tp_{\text{out}}} + 1\right) \frac{84}{B_1} \left(\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2\right).$$

Now we can let  $\gamma \leq \frac{1}{13L_F}$  such that  $84L_F^2 \gamma^2 \leq \frac{1}{2}$

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[k^T F(\mathbf{x}^t) k_2^2] \leq \frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} + \frac{2\sigma_{\text{bias}}^2}{m^4} + \left(\frac{1}{Tp_{\text{out}}} + 1\right) \frac{84}{B_1} \left(\frac{\sigma_{\text{in}}^2}{m} + \sigma_{\text{out}}^2\right). \quad (19)$$

In order to make the right-hand side  $\varepsilon^2$ , the inner batch size

$$m = \Omega(\tilde{\sigma}_{\text{bias}}^2 \varepsilon^{-0.5}),$$

and the outer batch size

$$B_1 = \frac{(\sigma_{\text{in}}^2/m + \sigma_{\text{out}}^2)}{\varepsilon^2}, \quad B_2 = \sqrt{B_1}, \quad p_{\text{out}} = \frac{1}{B_2}.$$

The step size  $\gamma$  is upper bounded by  $\min\left\{\frac{1}{2L_F}, \frac{\rho_{B_2}}{6L_F}, \frac{1}{13L_F}g\right\}$ . As  $B_2 \geq 1$ , we can take  $\gamma = \frac{1}{13L_F}$ . So we need iteration  $T$  greater than

$$T \geq \frac{26L_F(F(\mathbf{x}^0) - F^*)}{\varepsilon^2}.$$

By picking  $\mathbf{x}^s$  uniformly at random among  $\{G^t \mathbf{x}^0\}_{t=0}^{T-1}$ , we get the desired guarantee.

## D Stationary Point Convergence Proofs from Section 4 (FCCO)

In this section, we provide the convergence proofs for the FCCO problem. We start by analyzing a variant of BSpiderBoost (Algorithm 2) for this case in Appendix D.1. In Appendix D.2, we present a multi-level variance reduction approach (called NestedVR) that applies variance reduction in both outer (over the random variable  $i$ ) and inner (over the random variable  $\eta_j$ ) loops. In Appendix D.3, we analyze E-NestedVR. As in the case of CSO analyses, our proofs go via bounds on bias and variance terms of these algorithms.

### D.1 E-BSpiderBoost for FCCO problem

*Theorem 7 Consider the (FCCO) problem. Suppose Assumptions 3, 4, 5, 6 holds true. Let step size  $\gamma = O(1/L_F)$ . Then the output of E-BSpiderBoost (Algorithm 2) satisfies:  $\mathbb{E}[k^T F(\mathbf{x}^s) k_2^2] \leq \varepsilon^2$ , for nonconvex  $F$ , if the inner batch size  $m = \Omega(\max\{\tilde{\sigma}_{\text{bias}}^{1/2} \varepsilon^{-1/2}, \sigma_{\text{in}}^2 n^{-1} \varepsilon^{-2}\})$ , the hyperparameters of the outer loop of E-BSpiderBoost  $B_1 = n, B_2 = \sqrt{n}, p_{\text{out}} = 1/B_2$ , and the number of iterations  $T = \Omega(L_F(F(\mathbf{x}^0) - F^*) \varepsilon^{-2})$ . The resulting sample complexity is*

$$O\left(L_F(F(\mathbf{x}^0) - F^*) \max\left\{\frac{\rho_{B_2}^{1/2}}{\varepsilon^{2.5}}, \frac{\sigma_{\text{in}}^2}{n\varepsilon^4}\right\}\right).$$

*Remark 8* • When  $n = O(1)$ , we have a complexity of  $O(\varepsilon^{-4})$ . This happens because we did not apply variance reduction for the inner loop.

- When  $n = \Theta(\varepsilon^{-2/3})$ , E-BSpiderBoost has same performance as MSVR-V2 [Jiang et al., 2022] of  $O(n\varepsilon^{-3}) = O(\varepsilon^{-11/3})$ .
- When  $n = \Theta(\varepsilon^{-1.5})$ , E-BSpiderBoost achieves a better sample complexity of  $O(\varepsilon^{-3.25})$  than  $O(\varepsilon^{-4.5})$  from MSVR-V2 [Jiang et al., 2022].
- When  $n = \Theta(\varepsilon^{-2})$ , we recover  $O(\varepsilon^{-3.5})$  sample complexity as in Theorem 3.

Proof: Denote  $G^{t+1} = G_{\text{E-BSB}}^{t+1}$  (9). As we are using the finite-sum variant of SpiderBoost for the outer loop of the CSO problem, we only need to change the (17) and (18) to reflect that the outer variance is 0 now instead of  $\frac{\sigma_{\text{out}}^2}{B_1}$  in the general CSO case. More concretely, we update (17) to

$$\begin{aligned} E_{\text{var}}^{t+1} &= \mathbb{E}[\|G^{t+1} - \mathbb{E}[G^{t+1}]\|_2^2] \\ &= (1 - p_{\text{out}}) \mathbb{E}[\|G_S^{t+1} - \mathbb{E}[G_S^{t+1}]\|_2^2] + p_{\text{out}} \mathbb{E}[\|G_L^{t+1} - \mathbb{E}[G_L^{t+1}]\|_2^2] \\ &= \frac{(1 - p_{\text{out}})}{B_2} \mathbb{E}[\|G_{\text{E-BSGD}}^{t+1} - G_{\text{E-BSGD}}^t - \mathbb{E}[G_{\text{E-BSGD}}^{t+1} - G_{\text{E-BSGD}}^t]\|_2^2] + \frac{p_{\text{out}}}{B_1} \mathbb{E}[\|G_{\text{E-BSGD}}^{t+1} - \mathbb{E}[G_{\text{E-BSGD}}^{t+1}]\|_2^2] \\ &= \frac{1 - p_{\text{out}}}{B_2} \mathbb{E}[\|G_{\text{E-BSGD}}^{t+1} - G_{\text{E-BSGD}}^t - \mathbb{E}[G_{\text{E-BSGD}}^{t+1} - G_{\text{E-BSGD}}^t]\|_2^2] + \frac{p_{\text{out}}}{B_1} \frac{\sigma_{\text{in}}^2}{m}. \end{aligned} \quad (20)$$

and change (18) to

$$E_{\text{var}}^1 = \mathbb{E}[\|G^1 - \mathbb{E}[G^1]\|_2^2] = \mathbb{E}[\|G_L^1 - \mathbb{E}[G_L^1]\|_2^2] = \frac{1}{B_1} \mathbb{E}[\|G_{\text{E-BSGD}}^1 - \mathbb{E}[G_{\text{E-BSGD}}^1]\|_2^2] = \frac{1}{B_1} \frac{\sigma_{\text{in}}^2}{m}. \quad (21)$$

Then our analysis only has to start from the updated version of (19)

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[k^T F(\mathbf{x}^t) k^2] = \frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} + \frac{2\sigma_{\text{bias}}^2}{m^4} + \left(\frac{1}{T p_{\text{out}}} + 1\right) \frac{84}{B_1} \frac{\sigma_{\text{in}}^2}{m}.$$

We would like all terms on the right-hand side to be bounded by  $\varepsilon^2$ . From  $\frac{2\sigma_{\text{bias}}^2}{m^4} \leq \varepsilon^2$  we know that

$$m = \Omega\left(\frac{\sigma_{\text{bias}}^{1/2}}{\varepsilon^{1/2}}\right).$$

From  $\left(\frac{1}{T p_{\text{out}}} + 1\right) \frac{84}{B_1} \frac{\sigma_{\text{in}}^2}{m} \leq \varepsilon^2$ , we know that

$$m = \Omega\left(\frac{\sigma_{\text{in}}^2}{n\varepsilon^2}\right).$$

From  $\frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} \leq \varepsilon^2$ , we can choose that

$$\gamma = O\left(\frac{1}{L_F}\right), \quad T = \Omega\left(\frac{L_F(F(\mathbf{x}^0) - F^*)}{\varepsilon^2}\right).$$

Now the total sample complexity for E-BSpiderBoost for the FCCO problem becomes

$$B_2 m T = O\left(L_F^2(F(\mathbf{x}^0) - F^*) \max\left\{\frac{\rho \bar{n} \sigma_{\text{bias}}^{1/2}}{\varepsilon^{2.5}}, \frac{\sigma_{\text{in}}^2}{n\varepsilon^4}\right\}\right).$$

By picking  $\mathbf{x}^s$  uniformly at random among  $\mathcal{F} \mathbf{x}^t g_{t=0}^T$ , we get the desired guarantee.

## D.2 Convergence of NestedVR

NestedVR Algorithm. We start by describing the NestedVR construction. We maintain states  $\mathbf{y}_i^{t+1}$  and  $\mathbf{z}_i^{t+1}$  to approximate

$$\mathbf{y}_i^{t+1} = \mathbb{E}_{\eta^i}[g_\eta(\mathbf{x}^t)], \quad \mathbf{z}_i^{t+1} = \mathbb{E}_{\eta^i}[r g_\eta(\mathbf{x}^t)].$$

In iteration  $t + 1$ , if  $i$  is selected, then the state  $\mathbf{y}_i^{t+1}$  is updated as follows

$$\mathbf{y}_i^{t+1} = \begin{cases} \frac{1}{S_1} \sum_{\eta \in H_i} g_\eta(\mathbf{x}^t) & \text{with prob. } p_{\text{in}} \\ \mathbf{y}_i^t + \frac{1}{S_2} \sum_{\eta \in H_i} (g_\eta(\mathbf{x}^t) - g_\eta(\phi_i^t)) & \text{with prob. } 1 - p_{\text{in}}, \end{cases}$$

where  $\phi_i^t$  is the last time node  $i$  is visited. If  $i$  is not selected, then

$$\mathbf{y}_i^{t+1} = \mathbf{y}_i^t.$$



In this case,  $\mathbf{y}_i^{t+1}$  was never used to compute  $r f_i(\mathbf{y}_i^{t+1})$  because  $i$  is not selected at the time  $t + 1$ . We use the following quantities

$$\hat{\mathbf{z}}_i^{t+1} = \mathbb{E}_{\eta_{ji}}[r g_\eta(\mathbf{x}^t)], \quad \mathbf{z}_i^{t+1} = \frac{1}{m} \sum_{\eta \in \mathcal{H}_i} r g_\eta(\mathbf{x}^t). \quad (22)$$

We use  $G_{\text{NVR}}^{t+1}$  as the actual updates,

$$G_{\text{NVR}}^{t+1} = \begin{cases} \frac{1}{B_1} \sum_{i \in \mathcal{B}_1} (\mathbf{z}_i^{t+1})^\top r f_i(\mathbf{y}_i^{t+1}) & \text{with prob. } p_{\text{out}} \\ G_{\text{NVR}}^t + \frac{1}{B_2} \sum_{i \in \mathcal{B}_2} ((\hat{\mathbf{z}}_i^{t+1})^\top r f_i(\mathbf{y}_i^{t+1}) - (\mathbf{z}_i^t)^\top r f_i(\tilde{\mathbf{y}}_i^t)) & \text{with prob. } 1 - p_{\text{out}}. \end{cases}$$

We can also use the following quantity  $\hat{G}_{\text{NVR}}^{t+1}$  as an auxiliary

$$\hat{G}_{\text{NVR}}^{t+1} = \begin{cases} \frac{1}{B_1} \sum_{i \in \mathcal{B}_1} (\hat{\mathbf{z}}_i^{t+1})^\top r f_i(\mathbf{y}_i^{t+1}) & \text{with prob. } p_{\text{out}} \\ \hat{G}_{\text{NVR}}^t + \frac{1}{B_2} \sum_{i \in \mathcal{B}_2} ((\hat{\mathbf{z}}_i^{t+1})^\top r f_i(\mathbf{y}_i^{t+1}) - (\hat{\mathbf{z}}_i^t)^\top r f_i(\tilde{\mathbf{y}}_i^t)) & \text{with prob. } 1 - p_{\text{out}}. \end{cases}$$

Here we use  $\tilde{\mathbf{y}}_i^t$  to represent an i.i.d. copy of  $\mathbf{y}_i^t$  where  $i$  is selected at time  $t$ .

The iterate  $\mathbf{x}^{t+1}$  is therefore updated

$$\mathbf{x}^{t+1} = \mathbf{x}^t - \gamma G_{\text{NVR}}^{t+1}.$$

Lemma 9 *The error between  $G_{\text{NVR}}^{t+1}$  and  $\hat{G}_{\text{NVR}}^{t+1}$  can be upper bounded as follows*

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ \left\| G_{\text{NVR}}^{t+1} - \hat{G}_{\text{NVR}}^{t+1} \right\|_2^2 \right] \leq \frac{1}{B_1} \frac{C_f^2 \sigma_g^2}{m} + \frac{4(1 - p_{\text{out}})}{B_2 m p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ \left\| G_i^{t+1} - G_i^t \right\|_2^2 \right].$$

Proof: In this proof, we ignore the subscript in  $G_{\text{NVR}}^{t+1}$  and  $\hat{G}_{\text{NVR}}^{t+1}$ , we bound the error between  $G^{t+1}$  and associated  $\hat{G}^{t+1}$  where

$$G_i^{t+1} = \left( \frac{1}{m} \sum_{\eta \in \mathcal{H}_i} r g_\eta(\mathbf{x}) \right)^\top r f_i(\mathbf{y}_i^{t+1}), \\ \hat{G}_i^{t+1} = \left( \mathbb{E}_{\eta_{ji}}[r g_\eta(\mathbf{x})] \right)^\top r f_i(\mathbf{y}_i^{t+1}).$$

Let's only consider the expectation over the randomness of  $r g_\eta$ ,

$$\begin{aligned} \mathbb{E}_{H_i} \left[ \left\| G_i^{t+1} - \hat{G}_i^{t+1} \right\|_2^2 \right] &= \mathbb{E}_{\eta_{ji}} \left[ \left\| \left( \frac{1}{m} \sum_{\eta \in \mathcal{H}_i} r g_\eta(\mathbf{x}) - \mathbb{E}_{\eta_{ji}}[r g_\eta(\mathbf{x})] \right) \right\|_2^2 \right] \mathbb{E} \left[ \left\| r f_i(\mathbf{y}_i^{t+1}) \right\|_2^2 \right] \\ &= \frac{C_f^2}{m} \mathbb{E}_{\eta_{ji}} \left[ \left\| r g_\eta(\mathbf{x}) - \mathbb{E}_{\eta_{ji}}[r g_\eta(\mathbf{x})] \right\|_2^2 \right] \\ &= \frac{C_f^2 \sigma_g^2}{m}. \end{aligned}$$

Then we can bound the error as follows

$$\begin{aligned} \mathbb{E} \left[ \mathbb{E}_{H_i} \left[ \left\| G^{t+1} - \hat{G}^{t+1} \right\|_2^2 \right] \right] &= \frac{p_{\text{out}}}{B_1} \mathbb{E} \left[ \mathbb{E}_{H_i} \left[ \left\| G_i^{t+1} - \hat{G}_i^{t+1} \right\|_2^2 \right] \right] \\ &\quad + (1 - p_{\text{out}}) \left( \left\| G^t - \hat{G}^t \right\|_2^2 + \frac{1}{B_2} \mathbb{E} \left[ \mathbb{E}_{H_i} \left[ \left\| G_i^{t+1} - G_i^t - \hat{G}_i^{t+1} + \hat{G}_i^t \right\|_2^2 \right] \right] \right) \\ &= \frac{p_{\text{out}}}{B_1} \frac{C_f^2 \sigma_g^2}{m} + (1 - p_{\text{out}}) \left\| G^t - \hat{G}^t \right\|_2^2 \\ &\quad + \frac{(1 - p_{\text{out}})}{B_2 m} \left( \mathbb{E} \left[ \left\| G_i^{t+1} - G_i^t \right\|_2^2 \right] \right) \\ &= \frac{p_{\text{out}}}{B_1} \frac{C_f^2 \sigma_g^2}{m} + (1 - p_{\text{out}}) \left\| G^t - \hat{G}^t \right\|_2^2 + \frac{(1 - p_{\text{out}})}{B_2 m} \left( \mathbb{E} \left[ \left\| G_i^{t+1} - G_i^t \right\|_2^2 \right] \right). \end{aligned}$$

Unroll the recursion gives

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ \left\| G^{t+1} - \hat{G}^{t+1} \right\|_2^2 \right] \leq \frac{1}{B_1} \frac{C_f^2 \sigma_g^2}{m} + \frac{4(1 - p_{\text{out}})}{B_2 m p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ \left\| G_i^{t+1} - G_i^t \right\|_2^2 \right].$$

Lemma 10 (Staleness) Define the staleness of iterates at time  $t$  as  $\Xi^t := \frac{1}{n} \sum_{j=1}^n k_{\mathbf{x}^t} \|\phi_j^t\|_2^2$  and let  $G^{t+1}$  be the gradient estimate, then

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\Xi^t] \leq \frac{6n^2}{B_2^2} \gamma^2 \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2]. \quad (23)$$

Proof: Like previously (Lemma 5), let  $\mathbb{E}[\cdot]$  denote the expectation conditioned on all previous randomness until  $t-1$ . It is clear that  $\Xi^0 = 0$ , so we only consider  $t > 0$ . We upper bound  $\mathbb{E}[\Xi^t]$  as follows,

$$\mathbb{E}[\Xi^t] = (1 - p_{\text{out}}) \underbrace{\frac{1}{n} \sum_{j=1}^n \mathbb{E}[\|\mathbf{x}^t - \phi_j^t\|_2^2]}_{\text{if time } t \text{ takes } B_2} + p_{\text{out}} \underbrace{\frac{1}{n} \sum_{j=1}^n \mathbb{E}[\|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2]}_{\text{if time } t \text{ takes } B_1(\phi_j^t = \mathbf{x}^{t-1})}.$$

Then we can expand  $\mathbb{E}[\Xi^t]$  as follows

$$\begin{aligned} \mathbb{E}[\Xi^t] &= \frac{1-p_{\text{out}}}{n} \sum_{j=1}^n \mathbb{E}[\|\mathbf{x}^t - \phi_j^t\|_2^2] + \frac{p_{\text{out}}}{n} \sum_{j=1}^n \mathbb{E}[\|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2] \\ &= \frac{1-p_{\text{out}}}{n} \sum_{j=1}^n \left( (1 + \frac{1}{\beta}) \mathbb{E}_i[\|\mathbf{x}^t - \phi_j^t\|_2^2] + (1 + \beta) \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 \right) + p_{\text{out}} \gamma^2 \mathbb{E}[\|G^t\|_2^2] \\ &= \frac{1-p_{\text{out}}}{n} \sum_{j=1}^n (1 + \frac{1}{\beta}) \mathbb{E}_i[\|\mathbf{x}^t - \phi_j^t\|_2^2] + (1 + \beta) \gamma^2 \mathbb{E}[\|G^t\|_2^2] \end{aligned}$$

where we use Cauchy-Schwarz inequality with coefficient  $\beta > 0$ . By the definition of  $\phi_j^t$ ,

$$\begin{aligned} \mathbb{E}[\Xi^t] &= \frac{1-p_{\text{out}}}{n} \sum_{j=1}^n (1 + \frac{1}{\beta}) \left( \frac{n}{n} \frac{B_2}{n} \|\mathbf{x}^t - \phi_j^t\|_2^2 + \frac{B_2}{n} \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 \right) + (1 + \beta) \gamma^2 \mathbb{E}[\|G^t\|_2^2] \\ &= (1 + \frac{1}{\beta}) (1 - \frac{B_2}{n}) \Xi^t + (1 + \beta) \gamma^2 \mathbb{E}[\|G^t\|_2^2]. \end{aligned}$$

By taking  $\beta = 2n/B_2$ , we have that  $(1 + \frac{1}{\beta})(1 - \frac{B_2}{n}) = 1 - \frac{B_2}{2n}$  and thus

$$\mathbb{E}[\Xi^t] \leq (1 - \frac{B_2}{2n}) \Xi^t + (1 + \frac{2n}{B_2}) \gamma^2 \mathbb{E}[\|G^t\|_2^2].$$

Note that  $\mathbb{E}[\Xi^0] = 0$ .

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\Xi^t] &\leq \frac{2n}{B_2} (1 + \frac{2n}{B_2}) \gamma^2 \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] \\ &= \frac{6n^2}{B_2^2} \gamma^2 \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2]. \end{aligned}$$

The following lemma describes how the inner variable changes inside the variance.

Lemma 11 Denote  $E_y^{t+1} := \mathbb{E}[\|\mathbf{y}_i^{t+1} - \mathbb{E}_{\eta^{ji}}[g_\eta(\mathbf{x}^t)]\|_2^2]$  to be the error from inner variance and  $p_{\text{out}} T = 1$ . Then

$$\frac{1}{T} \sum_{t=0}^{T-1} E_y^{t+1} \leq \frac{(1 - p_{\text{in}}) C_g^2}{p_{\text{in}} S_2} \frac{1}{T} \sum_{t=0}^{T-1} \Xi^t + \frac{2\sigma_g^2}{S_1}.$$

Meanwhile,  $E_y^1 = \mathbb{E}[\|\mathbf{y}_i^1 - \mathbb{E}_{\eta^{ji}}[g_\eta(\mathbf{x}^0)]\|_2^2] = \frac{\sigma_g^2}{S_1}$ .

Proof:

$$\begin{aligned} E_y^{t+1} &= p_{\text{in}} \frac{\sigma_g^2}{S_1} + (1 - p_{\text{in}}) \mathbb{E}_i[\mathbb{E}_{\eta^{ji}}[\|\mathbf{y}_i^t - \mathbb{E}_{\eta^{ji}}[g_\eta(\phi_i^t)]\|_2^2]] \\ &\quad + \frac{1 - p_{\text{in}}}{S_2} \mathbb{E}_i[\mathbb{E}_{\eta^{ji}}[\|g_\eta(\mathbf{x}^t) - g_\eta(\phi_i^t)\|_2^2]] \\ &= (1 - p_{\text{in}}) E_y^t + \frac{(1 - p_{\text{in}}) C_g^2}{S_2} \mathbb{E}_i[\mathbb{E}_{\eta^{ji}}[\|\mathbf{x}^t - \phi_i^t\|_2^2]] + p_{\text{in}} \frac{\sigma_g^2}{S_1}. \end{aligned}$$

As  $t = 0$  always uses the large batch,  $E_y^1 = \mathbb{E}[\|\mathbf{y}_i^1 - \mathbb{E}_{\eta^{ji}}[g_\eta(\mathbf{x}^0)]\|_2^2] = \frac{\sigma_g^2}{S_1}$ . Then

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} E_y^{t+1} &= \frac{(1-p_{\text{in}})C_g^2}{p_{\text{in}}S_2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_i[\mathbb{E}_{\eta^{ji}}[k\mathbf{x}^t - \phi_i^t k_2^2]] + \frac{\sigma_g^2}{S_1} + \frac{E_y^1}{p_{\text{in}}T} \\ &= \frac{(1-p_{\text{in}})C_g^2}{p_{\text{in}}S_2} \frac{1}{T} \sum_{t=0}^{T-1} \Xi^t + \frac{2\sigma_g^2}{S_1}. \end{aligned}$$

Lemma 12 *The error  $\mathbb{E}_i[\mathbb{E}_{p_{\text{in}}}[\mathbb{E}_{\eta^{ji}}[\|\mathbf{y}_i^{t+1} - \tilde{\mathbf{y}}_i^t\|_2^2]]]$  satisfies*

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^{T-1} \mathbb{E}_i[\mathbb{E}_{p_{\text{in}}}[\mathbb{E}_{\eta^{ji}}[\|\mathbf{y}_i^{t+1} - \tilde{\mathbf{y}}_i^t\|_2^2]]] &= \frac{4C_g^2\gamma^2}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] + \frac{4(1-p_{\text{in}})C_g^2}{S_2} \frac{1}{T} \sum_{t=0}^{T-1} \Xi^{t+1} \\ &\quad + \frac{6(1-p_{\text{in}})}{T} \sum_{t=0}^{T-1} E_y^{t+1}. \end{aligned}$$

Note that when  $p_{\text{in}} = 1$  and  $S_1 = S_2 = m$ , we recover the following

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^{T-1} \mathbb{E}_i[\mathbb{E}_{p_{\text{in}}}[\mathbb{E}_{\eta^{ji}}[\|\mathbf{y}_i^{t+1} - \tilde{\mathbf{y}}_i^t\|_2^2]]] &= \frac{1}{T} \sum_{t=1}^{T-1} \mathbb{E}_i[\mathbb{E}_{p_{\text{in}}}[\mathbb{E}_{\eta^{ji}}[\|\frac{1}{m} \sum_{\eta \sim H_\xi} g_\eta(\mathbf{x}^t) - g_\eta(\mathbf{x}^{t-1})\|_2^2]]] \\ &\quad + \frac{4C_g^2\gamma^2}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2]. \end{aligned}$$

Proof: For  $t \geq 2$ ,  $\mathbb{E}_i[\mathbb{E}_{p_{\text{in}}}[\mathbb{E}_{\eta^{ji}}[\|\mathbf{y}_i^{t+1} - \tilde{\mathbf{y}}_i^t\|_2^2]]]$  can be upper bounded as follows

$$\begin{aligned} \mathbb{E}_i[\mathbb{E}_{p_{\text{in}}}[\mathbb{E}_{\eta^{ji}}[\|\mathbf{y}_i^{t+1} - \tilde{\mathbf{y}}_i^t\|_2^2]]] &= p_{\text{in}} \mathbb{E}_i \left[ \mathbb{E}_{\eta^{ji}} \left[ \left\| \frac{1}{S_1} \sum_{\eta \sim H_i} (g_\eta(\mathbf{x}^t) - g_\eta(\mathbf{x}^{t-1})) \right\|_2^2 \right] \right] \\ &\quad + (1-p_{\text{in}}) \mathbb{E}_i \left[ \mathbb{E}_{\eta^{ji}} \left[ \left\| \mathbf{y}_i^t - \mathbf{y}_i^{t-1} + \frac{1}{S_2} \sum_{\eta \sim H_i} (g_\eta(\mathbf{x}^t) - g_\eta(\phi_i^t)) - (g_\eta(\mathbf{x}^{t-1}) - g_\eta(\phi_i^{t-1})) \right\|_2^2 \right] \right] \\ &\quad + p_{\text{in}} C_g^2 \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 + \frac{1-p_{\text{in}}}{S_2} \mathbb{E}_i \left[ \mathbb{E}_{\eta^{ji}} \left[ \left\| (g_\eta(\mathbf{x}^t) - g_\eta(\phi_i^t)) - (g_\eta(\mathbf{x}^{t-1}) - g_\eta(\phi_i^{t-1})) \right\|_2^2 \right] \right] \\ &\quad + 3(1-p_{\text{in}}) \left( \mathbb{E}_i \left[ \left\| \mathbf{y}_i^t - \mathbb{E}_{\eta^{ji}}[g_\eta(\phi_i^t)] \right\|_2^2 \right] + \mathbb{E}_i \left[ \left\| \mathbf{y}_i^{t-1} - \mathbb{E}_{\eta^{ji}}[g_\eta(\phi_i^{t-1})] \right\|_2^2 \right] + C_g^2 \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 \right) \\ &\quad + p_{\text{in}} C_g^2 \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 + \frac{2(1-p_{\text{in}})C_g^2}{S_2} (\Xi^t + \Xi^{t-1}) + 3(1-p_{\text{in}}) (E_y^t + E_y^{t-1} + C_g^2 \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2) \\ &\quad + (p_{\text{in}} + 3(1-p_{\text{in}})) C_g^2 \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 + \frac{2(1-p_{\text{in}})C_g^2}{S_2} (\Xi^t + \Xi^{t-1}) + 3(1-p_{\text{in}}) (E_y^t + E_y^{t-1}). \end{aligned}$$

For  $t = 1$ , we choose  $\tilde{\mathbf{y}}_i^1 = \mathbf{y}_i^1$

$$\begin{aligned} \mathbb{E}_i[\mathbb{E}_{p_{\text{in}}}[\mathbb{E}_{\eta^{ji}}[\|\mathbf{y}_i^2 - \tilde{\mathbf{y}}_i^1\|_2^2]]] &= p_{\text{in}} \mathbb{E}_i \left[ \mathbb{E}_{\eta^{ji}} \left[ \left\| \frac{1}{S_1} \sum_{\eta \sim H_i} (g_\eta(\mathbf{x}^1) - g_\eta(\mathbf{x}^0)) \right\|_2^2 \right] \right] \\ &\quad + (1-p_{\text{in}}) \mathbb{E}_i \left[ \mathbb{E}_{\eta^{ji}} \left[ \left\| \mathbf{y}_i^1 - \frac{1}{S_2} \sum_{\eta \sim H_i} (g_\eta(\mathbf{x}^1) - g_\eta(\mathbf{x}^0)) - \tilde{\mathbf{y}}_i^1 \right\|_2^2 \right] \right] \\ &\quad + C_g^2 \|\mathbf{x}^1 - \mathbf{x}^0\|_2^2. \end{aligned}$$

Then for summing up  $t = 1$  to  $T-1$

$$\begin{aligned} &\sum_{t=2}^{T-1} \mathbb{E}_i[\mathbb{E}_{p_{\text{in}}}[\mathbb{E}_{\eta^{ji}}[\|\mathbf{y}_i^{t+1} - \tilde{\mathbf{y}}_i^t\|_2^2]]] + \mathbb{E}_i[\mathbb{E}_{p_{\text{in}}}[\mathbb{E}_{\eta^{ji}}[\|\mathbf{y}_i^2 - \tilde{\mathbf{y}}_i^1\|_2^2]]] \\ &\quad + (p_{\text{in}} + 3(1-p_{\text{in}})) C_g^2 \sum_{t=2}^{T-1} \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 + \frac{2(1-p_{\text{in}})C_g^2}{S_2} \left( \sum_{t=2}^{T-1} \Xi^t + \sum_{t=2}^{T-1} \Xi^{t-1} \right) \\ &\quad + 3(1-p_{\text{in}}) \left( \sum_{t=2}^{T-1} E_y^t + \sum_{t=2}^{T-1} E_y^{t-1} \right) + C_g^2 \|\mathbf{x}^1 - \mathbf{x}^0\|_2^2 \\ &\quad + 4C_g^2 \sum_{t=1}^{T-1} \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 + \frac{4(1-p_{\text{in}})C_g^2}{S_2} \sum_{t=0}^{T-1} \Xi^{t+1} + 6(1-p_{\text{in}}) \sum_{t=0}^{T-1} E_y^{t+1}. \end{aligned}$$

Finally, the error has the following upper bound

$$\begin{aligned} &\frac{1}{T} \sum_{t=1}^{T-1} \mathbb{E}_i[\mathbb{E}_{p_{\text{in}}}[\mathbb{E}_{\eta^{ji}}[\|\mathbf{y}_i^{t+1} - \tilde{\mathbf{y}}_i^t\|_2^2]]] \\ &\quad + \frac{4C_g^2\gamma^2}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] + \frac{4(1-p_{\text{in}})C_g^2}{S_2} \frac{1}{T} \sum_{t=0}^{T-1} \Xi^{t+1} + \frac{6(1-p_{\text{in}})}{T} \sum_{t=0}^{T-1} E_y^{t+1}. \end{aligned}$$

Lemma 13 (Bias and Variance of NestedVR) *If the step size  $\gamma$  satisfies,*

$$\gamma^2 L_F^2 \max \left\{ \frac{(1-p_{in})}{p_{in} S_2} \frac{18}{B_2}, \frac{1-p_{out}}{B_2} \frac{(1-p_{in})}{p_{in} S_2} \frac{18n^2}{B_2^2}, \frac{(1-p_{out})}{B_2} \right\} \leq \frac{1}{16} \frac{1}{6}$$

*then the variance and bias of NestedVR are*

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} E_{var}^{t+1} &= 32 \left( \left( \frac{p_{out}}{B_1} + \frac{1-p_{out}}{B_2} \right) \frac{(1-p_{in})}{p_{in} S_2} \frac{18n^2}{B_2^2} + \frac{(1-p_{out})}{B_2} \right) \frac{\gamma^2 L_F^2}{T} \sum_{t=0}^{T-1} \|\mathbb{E}[G^{t+1}]\|_2^2 \\ &\quad + 96 \left( \frac{p_{out}}{B_1} + \frac{(1-p_{in})(1-p_{out})}{B_2} \right) \frac{L_F^2}{S_1} + \frac{(1-p_{out})}{T} \frac{8L_F^2}{B_1 S_1} \\ \frac{1}{T} \sum_{t=0}^{T-1} E_{bias}^{t+1} &= \frac{12(1-p_{in})}{p_{in} S_2} \frac{n^2}{B_2^2} \frac{L_F^2 \gamma^2}{T} \sum_{t=0}^{T-1} \|\mathbb{E}[G^{t+1}]\|_2^2 + \frac{4L_F^2}{S_1} \\ &\quad + \left( \frac{12(1-p_{in})}{p_{in} S_2} \frac{n^2}{B_2^2} L_F^2 \gamma^2 + \frac{2(1-p_{out})^2}{p_{out}} \right) \frac{1}{T} \sum_{t=0}^{T-1} E_{var}^{t+1}. \end{aligned}$$

Proof: Notations. Let us define the following terms,

$$G_i^{t+1} := (\mathbf{z}_i^{t+1})^{\triangleright} r f_i(\mathbf{y}_i^{t+1}), \quad \tilde{G}_i^t := (\mathbf{z}_i^t)^{\triangleright} r f_i(\tilde{\mathbf{y}}_i^t).$$

Note that the  $\tilde{G}^t$  computed at time  $t+1$  has same expectation as  $G^t$

$$\mathbb{E}^{t+1}[\tilde{G}^t j t] = \mathbb{E}^t[G^t j t - 1]. \quad (24)$$

Computing the bias. First consider the two cases in the outer loop

$$\begin{aligned} E_{bias}^{t+1} &= \left\| r F(\mathbf{x}^t) - \mathbb{E}^{t+1}[G^{t+1} j t] \right\|_2^2 \\ &\quad + 2 \underbrace{\left\| r F(\mathbf{x}^t) - \mathbb{E}^{t+1}[G_i^{t+1} j t] \right\|_2^2}_{A_1^{t+1}} + 2 \underbrace{\left\| \mathbb{E}^{t+1}[G_i^{t+1} j t] - \mathbb{E}^{t+1}[G^{t+1} j t] \right\|_2^2}_{A_2^{t+1}}. \end{aligned}$$

We expand  $A_2^{t+1}$  as follows

$$\begin{aligned} A_2^{t+1} &= \left\| \mathbb{E}^{t+1}[G_i^{t+1} j t] - \mathbb{E}^{t+1}[G^{t+1} j t] \right\|_2^2 \\ &= \left\| \mathbb{E}^{t+1}[G_i^{t+1} j t] - p_{out} \mathbb{E}^{t+1}[G_i^{t+1} j t] - (1-p_{out})(G^t + \mathbb{E}^{t+1}[G_i^{t+1} - \tilde{G}_i^t j t]) \right\|_2^2 \\ &= (1-p_{out})^2 \left\| G^t - \mathbb{E}^{t+1}[\tilde{G}_i^t j t] \right\|_2^2 \\ &= (1-p_{out})^2 \left\| G^t - \mathbb{E}^t[G_i^t j t - 1] \right\|_2^2 \end{aligned}$$

where we use (24) in the last equality. Now we take expectation with respect to randomness at  $t$  such that  $G^t$  is a random variable, then

$$\begin{aligned} A_2^{t+1} &= (1-p_{out})^2 \mathbb{E}^t \left[ \left\| G^t - \mathbb{E}^t[G_i^t j t - 1] \right\|_2^2 j t - 1 \right] \\ &= (1-p_{out})^2 \left( \left\| \mathbb{E}^t[G^t j t - 1] - \mathbb{E}^t[G_i^t j t - 1] \right\|_2^2 + E_{var}^t \right) \\ &= (1-p_{out})^2 (A_2^t + E_{var}^t) \end{aligned}$$

while at initialization we always use large batch

$$A_2^1 = \left\| \mathbb{E}^1[G_i^1] - \mathbb{E}^1[G^1] \right\|_2^2 = \left\| \mathbb{E}^1[G_i^1] - \mathbb{E}^1[G_i^1] \right\|_2^2 = 0.$$

Therefore, when we average over time  $t$

$$\frac{1}{T} \sum_{t=0}^{T-1} A_2^{t+1} = \frac{(1-p_{out})^2}{p_{out}} \frac{1}{T} \sum_{t=0}^{T-1} E_{var}^{t+1}. \quad (25)$$

On the other hand, let us consider the upper bound on  $A_1^{t+1}$

$$A_1^{t+1} = C_g^2 L_f^2 \mathbb{E} \left[ \left\| \mathbf{y}_i^{t+1} - \mathbb{E}_{\eta/j} [g_\eta(\mathbf{x}^t)] \right\|_2^2 \right] = C_g^2 L_f^2 E_y^{t+1}.$$

From Lemma 11 we know that

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} A_1^{t+1} &= C_g^2 L_f^2 \left( \frac{(1-p_{\text{in}})C_g^2}{p_{\text{in}}S_2} \frac{1}{T} \sum_{t=0}^{T-1} \Xi^t + \frac{2\sigma_g^2}{S_1} \right) \\ &= \frac{(1-p_{\text{in}})L_F^2}{p_{\text{in}}S_2} \frac{1}{T} \sum_{t=0}^{T-1} \Xi^t + \frac{2L_F^2}{S_1}. \end{aligned}$$

From Lemma 10 we know that

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} A_1^{t+1} &= \frac{(1-p_{\text{in}})L_F^2}{p_{\text{in}}S_2} \left( \frac{6n^2}{B_2^2} \gamma^2 \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] \right) + \frac{2L_F^2}{S_1} \\ &= \frac{6(1-p_{\text{in}})}{p_{\text{in}}S_2} \frac{n^2}{B_2^2} \frac{L_F^2 \gamma^2}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] + \frac{2L_F^2}{S_1} + \frac{6(1-p_{\text{in}})}{p_{\text{in}}S_2} \frac{n^2}{B_2^2} \frac{L_F^2 \gamma^2}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1}. \end{aligned}$$

Therefore, the bias has the following bound

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{bias}}^{t+1} &= \frac{12(1-p_{\text{in}})}{p_{\text{in}}S_2} \frac{n^2}{B_2^2} \frac{L_F^2 \gamma^2}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] + \frac{4L_F^2}{S_1} \\ &+ \left( \frac{12(1-p_{\text{in}})}{p_{\text{in}}S_2} \frac{n^2}{B_2^2} L_F^2 \gamma^2 + \frac{2(1-p_{\text{out}})^2}{p_{\text{out}}} \right) \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1}. \end{aligned} \quad (26)$$

Note that when  $p_{\text{in}} = 1$  and  $S_1 = S_2 = m$ , then this bias recovers BSpiderBoost in (16)

$$\frac{1}{T} \sum_{t=0}^{T-1} E_{\text{bias}}^{t+1} = \frac{4L_F^2}{m} + \frac{2(1-p_{\text{out}})^2}{p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1}.$$

Computing the variance. Let us decompose the variance into 3 parts:

$$\begin{aligned} E_{\text{var}}^{t+1} &= \mathbb{E} \left[ \left\| G^{t+1} - \mathbb{E}[G^{t+1}] \right\|_2^2 \right] \\ &= \mathbb{E} \left[ \left\| G^{t+1} - \hat{G}^{t+1} + \hat{G}^{t+1} - \mathbb{E}_{\eta_{ji}}[\hat{G}^{t+1}] + \mathbb{E}_{\eta_{ji}}[\hat{G}^{t+1}] - \mathbb{E}_i[\mathbb{E}_{\eta_{ji}}[\hat{G}^{t+1}]] \right\|_2^2 \right] \\ &= \underbrace{\mathbb{E} \left[ \left\| G^{t+1} - \hat{G}^{t+1} \right\|_2^2 \right]}_{E_{r,g}^{t+1}} + \underbrace{\mathbb{E}_i \left[ \left\| \mathbb{E}_{\eta_{ji}}[G^{t+1}] - \mathbb{E}_i[\mathbb{E}_{\eta_{ji}}[G^{t+1}]] \right\|_2^2 \right]}_{E_{\text{var,out}}^{t+1}} + \underbrace{\mathbb{E} \left[ \left\| G^{t+1} - \mathbb{E}_{\eta_{ji}}[G^{t+1}] \right\|_2^2 \right]}_{E_{\text{var,in}}^{t+1}} \end{aligned}$$

where  $E_{\text{var,out}}^{t+1}$  and  $E_{\text{var,in}}^{t+1}$  are the variance of outer loop and inner loop.

Inner Variance. For  $t \geq 1$ , we expand the inner variance

$$\begin{aligned} E_{\text{var,in}}^{t+1} &= p_{\text{out}} \mathbb{E} \left[ \left\| \frac{1}{B_1} \sum_i (\mathbb{E}_{\eta_{ji}}[r g_{\eta}(\mathbf{x}^t)])^{\triangleright} (r f_i(\mathbf{y}_i^{t+1}) - \mathbb{E}_{\eta_{ji}}[r f_i(\mathbf{y}_i^{t+1})]) \right\|_2^2 \right] \\ &+ (1-p_{\text{out}}) \mathbb{E} \left[ \left\| \frac{1}{B_2} \sum_i (G_i^{t+1} - \tilde{G}_i^t) - \mathbb{E}_{\eta_{ji}}[G_i^{t+1} - \tilde{G}_i^t] \right\|_2^2 \right] \\ &= \frac{p_{\text{out}}}{B_1} C_g^2 \mathbb{E} \left[ \left\| r f_i(\mathbf{y}_i^{t+1}) - \mathbb{E}_{\eta_{ji}}[r f_i(\mathbf{y}_i^{t+1})] \right\|_2^2 \right] + \frac{1-p_{\text{out}}}{B_2} \mathbb{E}_i \left[ \left\| G_i^{t+1} - \tilde{G}_i^t \right\|_2^2 \right] \\ &= \frac{p_{\text{out}}}{B_1} 4C_g^2 L_f^2 E_y^{t+1} + \frac{1-p_{\text{out}}}{B_2} \mathbb{E}_i \left[ \mathbb{E}_{p_{\text{in}}} \left[ \mathbb{E}_{\eta_{ji}} \left[ \left\| G_i^{t+1} - \tilde{G}_i^t \right\|_2^2 \right] \right] \right]. \end{aligned} \quad (27)$$

We bound the outer variance as

$$\begin{aligned} \mathbb{E}_i \left[ \mathbb{E}_{p_{\text{in}}} \left[ \mathbb{E}_{\eta_{ji}} \left[ \left\| G_i^{t+1} - \tilde{G}_i^t \right\|_2^2 \right] \right] \right] &= \mathbb{E}_i \left[ \mathbb{E}_{p_{\text{in}}} \left[ \mathbb{E}_{\eta_{ji}} \left[ \left\| G_i^{t+1} - (\mathbb{E}_{\eta_{ji}}[r g_{\eta}(\mathbf{x}^{t-1})])^{\triangleright} r f_i(\mathbf{y}_i^{t+1}) - \tilde{G}_i^t \right\|_2^2 \right] \right] \right] \\ &= 2 \mathbb{E}_i \left[ \mathbb{E}_{p_{\text{in}}} \left[ \mathbb{E}_{\eta_{ji}} \left[ \left\| G_i^{t+1} - (\mathbb{E}_{\eta_{ji}}[r g_{\eta}(\mathbf{x}^{t-1})])^{\triangleright} r f_i(\mathbf{y}_i^{t+1}) \right\|_2^2 \right] \right] \right] \\ &+ 2 \mathbb{E}_i \left[ \mathbb{E}_{p_{\text{in}}} \left[ \mathbb{E}_{\eta_{ji}} \left[ \left\| (\mathbb{E}_{\eta_{ji}}[r g_{\eta}(\mathbf{x}^{t-1})])^{\triangleright} r f_i(\mathbf{y}_i^{t+1}) - \tilde{G}_i^t \right\|_2^2 \right] \right] \right] \\ &= 2C_f^2 L_g^2 \|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2^2 + 2C_g^2 L_f^2 \mathbb{E}_i \left[ \mathbb{E}_{p_{\text{in}}} \left[ \mathbb{E}_{\eta_{ji}} \left[ \left\| \mathbf{y}_i^{t+1} - \tilde{\mathbf{y}}_i^t \right\|_2^2 \right] \right] \right]. \end{aligned} \quad (28)$$

For  $t = 0$ , as we only use large and small batch in the

$$\begin{aligned}
E_{\text{var},\text{in}}^1 &= \mathbb{E} \left[ \left\| \frac{1}{B_1} \sum_i (\mathbb{E}_{\eta^j i} [r g_{\eta}(\mathbf{x}^0)])^{\succ} (r f_i(\mathbf{y}_i^1) \quad \mathbb{E}_{\eta^j i} [r f_i(\mathbf{y}_i^1)]) \right\|_2^2 \right] \\
&\quad \frac{1}{B_1} C_g^2 \mathbb{E} \left[ \left\| r f_i(\mathbf{y}_i^1) \quad \mathbb{E}_{\eta^j i} [r f_i(\mathbf{y}_i^1)] \right\|_2^2 \right] \\
&\quad \frac{1}{B_1} 4C_g^2 L_f^2 E_y^1 \\
&\quad \frac{1}{B_1} 4C_g^2 L_f^2 \frac{\sigma_g^2}{S_1} \\
&\quad \frac{4L_F^2}{B_1 S_1}.
\end{aligned} \tag{29}$$

Therefore, average over time  $t = 0, \dots, T-1$  gives

$$\begin{aligned}
&\frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var},\text{in}}^{t+1} \quad \frac{p_{\text{out}}}{B_1} 4C_g^2 L_f^2 \frac{1}{T} \sum_{t=1}^T E_y^{t+1} + \frac{1-p_{\text{out}}}{B_2} \frac{1}{T} \sum_{t=1}^T \mathbb{E}_i [\mathbb{E}_{p_{\text{in}}} [\mathbb{E}_{\eta^j i} [\|G_i^{t+1} \quad \tilde{G}_i^t\|_2^2]]] + \frac{E_{\text{var},\text{in}}^1}{T} \\
&= \frac{p_{\text{out}}}{B_1} 4C_g^2 L_f^2 \frac{1}{T} \sum_{t=0}^{T-1} E_y^{t+1} + \frac{1-p_{\text{out}}}{B_2} \frac{1}{T} \sum_{t=1}^T \mathbb{E}_i \left[ \mathbb{E}_{p_{\text{in}}} \left[ \mathbb{E}_{\eta^j i} \left[ \|G_i^{t+1} \quad \tilde{G}_i^t\|_2^2 \right] \right] \right] + \frac{(1-p_{\text{out}})E_{\text{var},\text{in}}^1}{T} \\
&\quad \frac{p_{\text{out}}}{B_1} 4C_g^2 L_f^2 \frac{1}{T} \sum_{t=0}^{T-1} E_y^{t+1} + \frac{(1-p_{\text{out}})E_{\text{var},\text{in}}^1}{T} \\
&\quad + \frac{2(1-p_{\text{out}})}{B_2} \left( \frac{C_f^2 L_g^2}{T} \sum_{t=1}^T \|\mathbf{x}^t \quad \mathbf{x}^{t-1}\|_2^2 + C_g^2 L_f^2 \frac{1}{T} \sum_{t=1}^T \mathbb{E}_i [\mathbb{E}_{p_{\text{in}}} [\mathbb{E}_{\eta^j i} [\|\mathbf{y}_i^{t+1} \quad \tilde{\mathbf{y}}_i^t\|_2^2]]] \right) \\
&\quad \frac{p_{\text{out}}}{B_1} 4C_g^2 L_f^2 \frac{1}{T} \sum_{t=0}^{T-1} E_y^{t+1} + \frac{(1-p_{\text{out}})E_{\text{var},\text{in}}^1}{T} \\
&\quad + \frac{2(1-p_{\text{out}})}{B_2} \frac{C_f^2 L_g^2 \gamma^2}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|G^{t+1}\|_2^2] \\
&\quad + \frac{2(1-p_{\text{out}})}{B_2} C_g^2 L_f^2 \frac{1}{T} \sum_{t=1}^T \mathbb{E}_i [\mathbb{E}_{p_{\text{in}}} [\mathbb{E}_{\eta^j i} [\|\mathbf{y}_i^{t+1} \quad \tilde{\mathbf{y}}_i^t\|_2^2]]].
\end{aligned}$$

Let us first apply Lemma 12

$$\begin{aligned}
&\frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var},\text{in}}^{t+1} \quad \frac{p_{\text{out}}}{B_1} 4C_g^2 L_f^2 \frac{1}{T} \sum_{t=0}^{T-1} E_y^{t+1} + \frac{(1-p_{\text{out}})E_{\text{var},\text{in}}^1}{T} + \frac{2(1-p_{\text{out}})}{B_2} \frac{C_f^2 L_g^2 \gamma^2}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|G^{t+1}\|_2^2] \\
&\quad + \frac{2(1-p_{\text{out}})C_g^2 L_f^2}{B_2} \left( \frac{4C_g^2 \gamma^2}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|G^{t+1}\|_2^2] + \frac{4(1-p_{\text{in}})C_g^2}{S_2} \frac{1}{T} \sum_{t=0}^{T-1} \Xi^t + \frac{6(1-p_{\text{in}})}{T} \sum_{t=0}^{T-1} E_y^{t+1} \right) \\
&\quad \left( \frac{p_{\text{out}}}{B_1} + \frac{(1-p_{\text{in}})(1-p_{\text{out}})}{B_2} \right) \frac{12C_g^2 L_f^2}{T} \sum_{t=0}^{T-1} E_y^{t+1} + \frac{8(1-p_{\text{out}})L_F^2 \gamma^2}{B_2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|G^{t+1}\|_2^2] \\
&\quad + \frac{8(1-p_{\text{out}})(1-p_{\text{in}})C_g^4 L_f^2}{B_2 S_2} \frac{1}{T} \sum_{t=0}^{T-1} \Xi^t + \frac{(1-p_{\text{out}})E_{\text{var},\text{in}}^1}{T}
\end{aligned}$$

Then we apply Lemma 11 on the bound of  $\frac{1}{T} \sum_{t=0}^{T-1} E_y^{t+1}$

$$\begin{aligned}
&\frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var},\text{in}}^{t+1} \quad 24 \left( \frac{p_{\text{out}}}{B_1} + \frac{(1-p_{\text{in}})(1-p_{\text{out}})}{B_2} \right) \left( \frac{(1-p_{\text{in}})L_F^2}{p_{\text{in}}S_2} \frac{1}{T} \sum_{t=0}^{T-1} \Xi^t + \frac{L_F^2}{S_1} \right) \\
&\quad + \frac{8(1-p_{\text{out}})L_F^2 \gamma^2}{B_2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|G^{t+1}\|_2^2] \\
&\quad + \frac{8(1-p_{\text{out}})(1-p_{\text{in}})C_g^4 L_f^2}{B_2 S_2} \frac{1}{T} \sum_{t=0}^{T-1} \Xi^t + \frac{(1-p_{\text{out}})E_{\text{var},\text{in}}^1}{T} \\
&\quad 24 \left( \frac{p_{\text{out}}}{B_1} + \frac{(1-p_{\text{in}})(1-p_{\text{out}})}{B_2} + \frac{p_{\text{in}}(1-p_{\text{out}})}{B_2} \right) \frac{(1-p_{\text{in}})L_F^2}{p_{\text{in}}S_2} \frac{1}{T} \sum_{t=0}^{T-1} \Xi^t \\
&\quad + \frac{8(1-p_{\text{out}})L_F^2 \gamma^2}{B_2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|G^{t+1}\|_2^2] \\
&\quad + 24 \left( \frac{p_{\text{out}}}{B_1} + \frac{(1-p_{\text{in}})(1-p_{\text{out}})}{B_2} \right) \frac{L_F^2}{S_1} + \frac{(1-p_{\text{out}})E_{\text{var},\text{in}}^1}{T} \\
&\quad 24 \left( \frac{p_{\text{out}}}{B_1} + \frac{1-p_{\text{out}}}{B_2} \right) \frac{(1-p_{\text{in}})L_F^2}{p_{\text{in}}S_2} \frac{1}{T} \sum_{t=0}^{T-1} \Xi^t \\
&\quad + \frac{8(1-p_{\text{out}})L_F^2 \gamma^2}{B_2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|G^{t+1}\|_2^2] \\
&\quad + 24 \left( \frac{p_{\text{out}}}{B_1} + \frac{(1-p_{\text{in}})(1-p_{\text{out}})}{B_2} \right) \frac{L_F^2}{S_1} + \frac{(1-p_{\text{out}})E_{\text{var},\text{in}}^1}{T}.
\end{aligned}$$

From Lemma 10, we plug in the upper bound of  $\frac{1}{T} \sum_{t=0}^{T-1} \Xi^t$

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var},\text{in}}^{t+1} & \leq 24 \left( \frac{p_{\text{out}}}{B_1} + \frac{1-p_{\text{out}}}{B_2} \right) \frac{(1-p_{\text{in}})L_F^2}{p_{\text{in}}S_2} \left( \frac{6n^2}{B_2^2} \gamma^2 \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] \right) \\ & \quad + \frac{8(1-p_{\text{out}})L_F^2\gamma^2}{B_2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] \\ & \quad + 24 \left( \frac{p_{\text{out}}}{B_1} + \frac{(1-p_{\text{in}})(1-p_{\text{out}})}{B_2} \right) \frac{L_F^2}{S_1} + \frac{(1-p_{\text{out}})}{T} E_{\text{var},\text{in}}^1 \\ & \leq 8 \left( \left( \frac{p_{\text{out}}}{B_1} + \frac{1-p_{\text{out}}}{B_2} \right) \frac{(1-p_{\text{in}})}{p_{\text{in}}S_2} \frac{18n^2}{B_2^2} + \frac{(1-p_{\text{out}})}{B_2} \right) \frac{\gamma^2 L_F^2}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] \\ & \quad + 24 \left( \frac{p_{\text{out}}}{B_1} + \frac{(1-p_{\text{in}})(1-p_{\text{out}})}{B_2} \right) \frac{L_F^2}{S_1} + \frac{(1-p_{\text{out}})}{T} E_{\text{var},\text{in}}^1. \end{aligned}$$

Finally, we add the upper bound on with  $E_{\text{var},\text{in}}^1$  with (29)

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var},\text{in}}^{t+1} & \leq 8 \left( \left( \frac{p_{\text{out}}}{B_1} + \frac{1-p_{\text{out}}}{B_2} \right) \frac{(1-p_{\text{in}})}{p_{\text{in}}S_2} \frac{18n^2}{B_2^2} + \frac{(1-p_{\text{out}})}{B_2} \right) \frac{\gamma^2 L_F^2}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] \\ & \quad + 24 \left( \frac{p_{\text{out}}}{B_1} + \frac{(1-p_{\text{in}})(1-p_{\text{out}})}{B_2} \right) \frac{L_F^2}{S_1} + \frac{(1-p_{\text{out}})}{T} \frac{4L_F^2}{B_1 S_1}. \end{aligned} \quad (30)$$

Outer Variance. Now we consider the outer variance for  $t \geq 1$

$$\begin{aligned} E_{\text{var},\text{out}}^{t+1} & \leq \frac{(1-p_{\text{out}})^2}{B_2} \mathbb{E}_i \left[ \left\| \mathbb{E}_{\eta^j i} [G_i^{t+1}] - \mathbb{E}_{\eta^j i} [\tilde{G}_i^t] \right\|_2^2 \right] \\ & \leq \frac{(1-p_{\text{out}})^2}{B_2} \mathbb{E}_i \left[ \mathbb{E}_{\eta^j i} \left[ \left\| G_i^{t+1} - \tilde{G}_i^t \right\|_2^2 \right] \right]. \end{aligned}$$

Compared to (27) we know that the upper bound of is smaller than that of  $E_{\text{var},\text{in}}^{t+1}$ . Besides, whereas  $E_{\text{var},\text{out}}^1 = 0$  as we use large batch at  $t = 0$ . Therefore, the upper bound of  $E_{\text{var}}^{t+1}$  is upper bounded by  $2^*(30)$ . Variance of  $r g_\eta$ . From Lemma 9, we know that

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[E_r^{t+1}] & \leq \frac{1}{B_1} \frac{C_f^2 \sigma_g^2}{m} + \frac{4(1-p_{\text{out}})}{B_2 m p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ \left\| G_i^{t+1} - \tilde{G}_i^t \right\|_2^2 \right] \\ & \leq \frac{1}{B_1} \frac{C_f^2 \sigma_g^2}{m} + \frac{1}{m} E_{\text{var}}^{t+1} \end{aligned}$$

Finally, we use  $\mathbb{E}[\|G^{t+1}\|_2^2] = \|\mathbb{E}[G^{t+1}]\|_2^2 + E_{\text{var}}^{t+1}$ .

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1} & \leq 16 \left( \left( \frac{p_{\text{out}}}{B_1} + \frac{1-p_{\text{out}}}{B_2} \right) \frac{(1-p_{\text{in}})}{p_{\text{in}}S_2} \frac{18n^2}{B_2^2} + \frac{(1-p_{\text{out}})}{B_2} \right) \frac{\gamma^2 L_F^2}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] \\ & \quad + 16 \left( \left( \frac{p_{\text{out}}}{B_1} + \frac{1-p_{\text{out}}}{B_2} \right) \frac{(1-p_{\text{in}})}{p_{\text{in}}S_2} \frac{18n^2}{B_2^2} + \frac{(1-p_{\text{out}})}{B_2} \right) \frac{\gamma^2 L_F^2}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1} \\ & \quad + 48 \left( \frac{p_{\text{out}}}{B_1} + \frac{(1-p_{\text{in}})(1-p_{\text{out}})}{B_2} \right) \frac{L_F^2}{S_1} + \frac{(1-p_{\text{out}})}{T} \frac{8L_F^2}{B_1 S_1}. \end{aligned}$$

By taking step size  $\gamma$  to satisfy

$$\gamma^2 L_F^2 \max \left\{ \frac{p_{\text{out}}}{B_1} \frac{(1-p_{\text{in}})}{p_{\text{in}}S_2} \frac{18n^2}{B_2^2}, \frac{1-p_{\text{out}}}{B_2} \frac{(1-p_{\text{in}})}{p_{\text{in}}S_2} \frac{18n^2}{B_2^2}, \frac{(1-p_{\text{out}})}{B_2} \right\} \leq \frac{1}{16} \frac{1}{6}$$

which can be simplified to

$$\gamma^2 L_F^2 \max \left\{ \frac{(1-p_{\text{in}})}{p_{\text{in}}S_2} \frac{18}{B_2}, \frac{1-p_{\text{out}}}{B_2} \frac{(1-p_{\text{in}})}{p_{\text{in}}S_2} \frac{18n^2}{B_2^2}, \frac{(1-p_{\text{out}})}{B_2} \right\} \leq \frac{1}{16} \frac{1}{6}.$$

Then the coefficient of  $\frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1}$  is bounded by  $\frac{1}{2}$

$$16 \left( \left( \frac{p_{\text{out}}}{B_1} + \frac{1-p_{\text{out}}}{B_2} \right) \frac{(1-p_{\text{in}})}{p_{\text{in}}S_2} \frac{18n^2}{B_2^2} + \frac{(1-p_{\text{out}})}{B_2} \right) \gamma^2 L_F^2 \leq \frac{1}{2}.$$

The the variance has the following bound

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1} & \leq 32 \left( \left( \frac{p_{\text{out}}}{B_1} + \frac{1-p_{\text{out}}}{B_2} \right) \frac{(1-p_{\text{in}}) 18n^2}{p_{\text{in}} S_2 B_2^2} + \frac{(1-p_{\text{out}})}{B_2} \right) \frac{\gamma^2 L_F^2}{T} \sum_{t=0}^{T-1} \|\mathbb{E}[G^{t+1}]\|_2^2 \\ & \quad + 96 \left( \frac{p_{\text{out}}}{B_1} + \frac{(1-p_{\text{in}})(1-p_{\text{out}})}{B_2} \right) \frac{L_F^2}{S_1} + \frac{(1-p_{\text{out}})}{T} \frac{8L_F^2}{B_1 S_1}. \end{aligned}$$

Theorem 9 Consider the (FCCO) problem. Suppose Assumptions 3, 4, 5 holds true. Let step size  $\gamma = O(\frac{1}{nL_F})$ . Then for NestedVR,  $\mathbf{x}^s$  picked uniformly at random among  $\mathbf{f}\mathbf{x}^t g_{t=0}^{T-1}$  satisfies:  $\mathbb{E}[kr F(\mathbf{x}^s)k_2^2] \leq \varepsilon^2$ , for nonconvex  $F$ , if the hyperparameters of the inner loop  $S_1 = O(\tilde{L}_F^2 \varepsilon^{-2})$ ,  $S_2 = O(\tilde{L}_F \varepsilon^{-1})$ ,  $p_{\text{in}} = O(1/S_2)$ , the hyperparameters of the outer loop  $B_1 = n, B_2 = \frac{P}{n}, p_{\text{out}} = 1/B_2$ , and the number of iterations

$$T = \Omega \left( \frac{P n L_F (F(\mathbf{x}^0) - F^*)}{\varepsilon^2} \right).$$

The resulting sample complexity is

$$O \left( \frac{n L_F L_F (F(\mathbf{x}^0) - F^*)}{\varepsilon^3} \right).$$

In fact, it reaches this sample complexity for all  $\frac{P n p_{\text{out}}}{1-p_{\text{in}}} \leq \varepsilon$ .

Proof: Using descent lemma (Lemma 4) and bias-variance bounds of NestedVR (Lemma 13)

$$\begin{aligned} & \frac{1}{T} \sum_{t=0}^{T-1} kr F(\mathbf{x}^t)k_2^2 + \frac{1}{2T} \sum_{t=0}^{T-1} \|\mathbb{E}[G^{t+1}]\|_2^2 \\ & \quad \frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} + \frac{L_F \gamma}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1} + \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{bias}}^{t+1} \\ & \quad \underbrace{\frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T}}_{T_0} + \underbrace{\frac{4L_F^2}{S_1}}_{T_1} + \underbrace{\frac{12(1-p_{\text{in}})n^2 L_F^2 \gamma^2}{p_{\text{in}} S_2 B_2^2 T} \sum_{t=0}^{T-1} \|\mathbb{E}[G^{t+1}]\|_2^2}_{T_2} \\ & \quad + \underbrace{\left( \frac{12(1-p_{\text{in}})n^2 L_F^2 \gamma^2}{p_{\text{in}} S_2 B_2^2} + \frac{2(1-p_{\text{out}})^2}{p_{\text{out}}} + \gamma L_F \right) \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1}}_{T_3}. \end{aligned}$$

Compute  $T_0$ . In order to let  $T_0 \leq \varepsilon^2$ , we require that

$$\gamma T \leq \varepsilon^2. \quad (31)$$

Compute  $T_1$ . In order to let  $T_1$  to be smaller than  $\varepsilon^2$ , we need

$$S_1 = \frac{4L_F^2}{\varepsilon^2}.$$

Compute  $T_2$ . In order to let the coefficient of  $\frac{1}{T} \sum_{t=0}^{T-1} \|\mathbb{E}[G^{t+1}]\|_2^2$  in  $T_2$  to be less than  $\frac{1}{4}$ , i.e.

$$\frac{12(1-p_{\text{in}})n^2 L_F^2 \gamma^2}{p_{\text{in}} S_2 B_2^2} \leq \frac{1}{4}, \quad (32)$$

which requires  $\gamma$

$$\gamma \leq \frac{B_2 \sqrt{p_{\text{in}} S_2}}{7L_F n \sqrt{1-p_{\text{in}}}} = \frac{p_{\text{out}} \sqrt{p_{\text{in}} L_F}}{7\varepsilon L_F \sqrt{1-p_{\text{in}}}}. \quad (33)$$

Compute  $T_3$ . Let us now focus on  $T_3$  and notice that the middle term  $\frac{2(1-p_{\text{out}})^2}{p_{\text{out}}}$

$$\frac{2(1-p_{\text{out}})^2}{p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1}.$$

Using Lemma 13 we have that

$$\begin{aligned} & \frac{2(1-p_{\text{out}})^2}{p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1} \\ & \quad \underbrace{32 \frac{2(1-p_{\text{out}})^2}{p_{\text{out}}} \left( \left( \frac{p_{\text{out}}}{B_1} + \frac{1-p_{\text{out}}}{B_2} \right) \frac{(1-p_{\text{in}}) 18n^2}{p_{\text{in}} S_2 B_2^2} + \frac{(1-p_{\text{out}})}{B_2} \right) \gamma^2 L_F^2 \frac{1}{T} \sum_{t=0}^{T-1} \|\mathbb{E}[G^{t+1}]\|_2^2}_{T_{3,1}} \\ & \quad + \underbrace{96 \frac{2(1-p_{\text{out}})^2}{p_{\text{out}}} \left( \frac{p_{\text{out}}}{B_1} + \frac{(1-p_{\text{in}})(1-p_{\text{out}})}{B_2} \right) \frac{L_F^2}{S_1}}_{T_{3,2}} + \underbrace{\frac{2(1-p_{\text{out}})^2}{p_{\text{out}}} \frac{(1-p_{\text{out}})}{T} \frac{8L_F^2}{B_1 S_1}}_{T_{3,3}}. \end{aligned}$$



- Compute  $T_{3,3}$ : As we already know that  $S_1 = O(\varepsilon^{-2})$  and  $T = 1$  and  $B_1 p_{\text{out}} = 1$ . This imposes no more constraints, i.e.

$$S_1 = O\left(\frac{\tilde{L}_F^2}{\varepsilon^2}\right).$$

- Compute  $T_{3,2}$ : As  $S_1 = O(\varepsilon^{-2})$  and  $B_1 = n$  and  $B_2 = B_1 p_{\text{out}}$ , then it requires

$$\frac{(1 - p_{\text{in}})(1 - p_{\text{out}})^3}{p_{\text{out}}^2} = n.$$

- Compute  $T_{3,1}$ : In order to satisfy the following

$$32 \frac{2(1 - p_{\text{out}})^2}{p_{\text{out}}} \left( \left( \frac{p_{\text{out}}}{B_1} + \frac{1 - p_{\text{out}}}{B_2} \right) \frac{(1 - p_{\text{in}}) 18n^2}{p_{\text{in}} S_2 B_2^2} + \frac{(1 - p_{\text{out}})}{B_2} \right) \gamma^2 L_F^2 = \frac{1}{12}$$

we need to enforce

$$\gamma = \frac{p_{\text{in}} p_{\text{out}} \tilde{L}_F}{\varepsilon L_F (1 - p_{\text{in}})^{1/2} (1 - p_{\text{out}})^{3/2}}. \quad (34)$$

Now we go back to  $T_3$  and compare the other two coefficients

$$\frac{12(1 - p_{\text{in}}) n^2}{p_{\text{in}} S_2 B_2^2} L_F^2 \gamma^2 + \frac{2(1 - p_{\text{out}})^2}{p_{\text{out}}} + \gamma L_F.$$

As  $\gamma L_F = \frac{1}{2} \cdot \frac{2(1 - p_{\text{out}})^2}{p_{\text{out}}}$  we can safely ignore  $\gamma L_F$ . On the other hand, from (32) we know that the first term is also have

$$\frac{12(1 - p_{\text{in}}) n^2}{p_{\text{in}} S_2 B_2^2} L_F^2 \gamma^2 = \frac{1}{4} \cdot \frac{2(1 - p_{\text{out}})^2}{p_{\text{out}}}.$$

Constraints from the Bias-Variance Lemma (Lemma 13). By setting  $B_1 = n$  and  $S_1 = O(\frac{L_F^2}{\varepsilon^2})$ , this constraint translates to

$$\gamma^2 L_F^2 \max \left\{ \frac{(1 - p_{\text{in}}) \varepsilon^2}{p_{\text{in}}^2 B_2}, \frac{1 - p_{\text{out}}}{B_2} \frac{(1 - p_{\text{in}}) \varepsilon^2}{p_{\text{in}}^2} \frac{1}{p_{\text{out}}^2}, \frac{(1 - p_{\text{out}})}{B_2} \right\} = 1$$

which is weaker than (33).

Summary on the Limit on  $\gamma$ . Combine (33) and (34) and  $\gamma = \frac{1}{2L_F}$ , we have a final limit on step size  $\gamma$

$$\gamma = \min \left\{ \frac{p_{\text{out}} B_{\text{in}} L_F}{\varepsilon L_F^2 (1 - p_{\text{in}})}, \frac{1}{L_F} \right\} \quad (35)$$

Then the total sample complexity of NestedVR can be computed as

$$(\# \text{ of iters } T) = (\text{Avg. outer batch size } B_2 = B_1 p_{\text{out}}) \cdot (\text{Avg. inner batch size } S_2 = S_1 p_{\text{in}}).$$

This sample complexity has the following requirement

$$B_2 S_2 T = \frac{B_2 S_2 (T \gamma)}{\gamma} \stackrel{(31)}{=} \frac{B_2 S_2}{\varepsilon^2 \gamma} = \frac{n \varepsilon^{-2} p_{\text{in}} p_{\text{out}}}{\varepsilon^2 \gamma} \stackrel{(35)}{\geq} n \varepsilon^{-3}.$$

The lower bound  $n \varepsilon^{-3}$  is reached when in (35) we have

$$\frac{p_{\text{out}} B_{\text{in}} L_F}{\varepsilon L_F^2 (1 - p_{\text{in}})} = \frac{1}{L_F}.$$

That is,  $\frac{p_{\text{out}} B_{\text{in}}}{1 - p_{\text{in}}} = \varepsilon$ .

In particular, we can choose the following hyperparameters to reach  $O(n \varepsilon^{-3})$  sample complexity

$$B_1 = n, \quad B_2 = \frac{p_{\text{out}}}{n}, \quad p_{\text{out}} = \frac{1}{n}, \quad S_1 = O(\tilde{L}_F^2 \varepsilon^{-2}), \quad S_2 = O(\tilde{L}_F \varepsilon^{-1}), \quad p_{\text{in}} = O(\tilde{L}_F^{-1} \varepsilon)$$

The step size  $\gamma$  can be chosen as

$$\gamma \leq \frac{1}{\bar{n}L_F}.$$

and the iteration complexity

$$T = \Omega\left(\frac{D\bar{n}L_F(F(\mathbf{x}^0) - F^*)}{\varepsilon^2}\right).$$

Putting these together gives the claimed sample complexity bound. By picking  $\mathbf{x}^s$  uniformly at random among  $\{\mathbf{x}^t\}_{t=0}^{T-1}$ , we get the desired guarantee.

### D.3 Convergence of E-NestedVR

In this section, we analyze the sample complexity of Algorithm 3 (E-NestedVR) for the FCCO problem with

$$G_{\text{E-NVR}}^{t+1} = \begin{cases} \frac{1}{B_1} \sum_i (\mathbf{z}_i^{t+1})^{\triangleright} L_{D_{\mathbf{y},i}^{(2)}} r f_i(0) & \text{with prob. } p_{\text{out}} \\ G_{\text{E-NVR}}^t + \frac{1}{B_2} \sum_i \left( (\mathbf{z}_i^{t+1})^{\triangleright} L_{D_{\mathbf{y},i}^{(2)}} r f_i(0) - (\mathbf{z}_i^t)^{\triangleright} L_{D_{\mathbf{y},i}^{(2)}} r f_i(0) \right) & \text{with prob. } 1 - p_{\text{out}} \end{cases} \quad (36)$$

Lemma 14 (Bias and Variance of E-NestedVR) *If the step size  $\gamma$  satisfies*

$$\gamma^2 L_F^2 \max\left\{ \frac{(1-p_{\text{in}})18}{p_{\text{in}}S_2} \frac{1}{B_2}, \frac{1-p_{\text{out}}}{B_2} \frac{(1-p_{\text{in}})18n^2}{p_{\text{in}}S_2} \frac{1}{B_2^2}, \frac{(1-p_{\text{out}})}{B_2} \right\} \leq \frac{1}{16} \frac{1}{6}$$

*then the variance and bias of E-NestedVR are*

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1} &= 14 \cdot 32 \left( \left( \frac{p_{\text{out}}}{B_1} + \frac{1-p_{\text{out}}}{B_2} \right) \frac{(1-p_{\text{in}})18n^2}{p_{\text{in}}S_2} \frac{1}{B_2^2} + \frac{(1-p_{\text{out}})}{B_2} \right) \frac{\gamma^2 L_F^2}{T} \sum_{t=0}^{T-1} \|\mathbb{E}[G^{t+1}]\|_2^2 \\ &\quad + 14 \cdot 96 \left( \frac{p_{\text{out}}}{B_1} + \frac{(1-p_{\text{in}})(1-p_{\text{out}})}{B_2} \right) \frac{L_F^2}{S_1} + \frac{(1-p_{\text{out}})}{T} \frac{8L_F^2}{B_1S_1}. \\ \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{bias}}^{t+1} &= \frac{(1-p_{\text{in}})^3 L_F^2}{p_{\text{in}}S_2} \frac{6n^2}{B_2^2} \gamma^2 \frac{1}{T} \sum_{t=0}^{T-1} \|\mathbb{E}[G^{t+1}]\|_2^2 + \frac{2(1-p_{\text{in}})^2 L_F^2}{S_1} + \frac{C_g^2}{S_2^2} \\ &\quad + \left( \frac{(1-p_{\text{in}})^3 L_F^2}{p_{\text{in}}S_2} \frac{6n^2}{B_2^2} \gamma^2 + \frac{(1-p_{\text{out}})^2}{p_{\text{out}}} \right) \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1}. \end{aligned}$$

Proof: Note that this proof is very similar to NestedVR so we highlight the differences. Let  $G^{t+1} = G_{\text{E-NVR}}^{t+1}$  (36) be the E-NestedVR update and define

$$G_i^{t+1} := (\mathbf{z}_i^{t+1})^{\triangleright} L_{D_{\mathbf{y},i}^{(2)}} r f_i(0)$$

We expand the bias by inserting  $\mathbb{E}_{i,p_{\text{in}},\eta_j}[G_i^{t+1}]$

$$\begin{aligned} E_{\text{bias}}^{t+1} &= \|\mathbb{E}[r F(\mathbf{x}^{t+1}) - \mathbb{E}[G^{t+1}]]\|_2^2 \\ &= 2 \underbrace{\|\mathbb{E}[r F(\mathbf{x}^{t+1}) - \mathbb{E}_{i,p_{\text{in}},\eta_j}[G_i^{t+1}]]\|_2^2}_{A_1^{t+1}} + 2 \underbrace{\|\mathbb{E}_{i,p_{\text{in}},\eta_j}[G_i^{t+1}] - \mathbb{E}[G^{t+1}]\|_2^2}_{A_2^{t+1}}. \end{aligned}$$

Consider  $A_1^{t+1}$ . The term  $A_1^{t+1}$  captures the difference between full gradient and extrapolated gradient

$$\begin{aligned} A_1^{t+1} &= \left\| \mathbb{E}_i \left[ \left( \mathbb{E}_{\eta_j} [r g_{\eta}(\mathbf{x}^t)] \right)^{\triangleright} r f_i(\mathbb{E}[g_{\eta}(\mathbf{x}^t)]) - \mathbb{E}_{p_{\text{in}},\eta_j} \left[ \left( \mathbb{E}_{\eta_j} [r g_{\eta}(\mathbf{x}^t)] \right)^{\triangleright} L_{D_{\mathbf{y},i}^{(2)}} r f_i(0) \right] \right] \right\|_2^2 \\ &= C_g^2 \mathbb{E}_i \left[ \left\| r f_i(\mathbb{E}_{\eta_j}[g_{\eta}(\mathbf{x}^t)]) - \mathbb{E}_{p_{\text{in}},\eta_j} \left[ L_{D_{\mathbf{y},i}^{(2)}} r f_i(0) \right] \right\|_2^2 \right] \\ &\quad + 2C_g^2 \underbrace{\mathbb{E}_i \left[ \left\| r f_i(\mathbb{E}_{\eta_j}[g_{\eta}(\mathbf{x}^t)]) - \mathbb{E}_{p_{\text{in}}} [r f_i(\mathbb{E}_{\eta_j}[\mathbf{y}_i^{t+1}])] \right\|_2^2 \right]}_{=: A_{1,1}^{t+1}} \\ &\quad + 2C_g^2 \underbrace{\mathbb{E}_i \left[ \left\| \mathbb{E}_{p_{\text{in}}} [r f_i(\mathbb{E}_{\eta_j}[\mathbf{y}_i^{t+1}])] - \mathbb{E}_{p_{\text{in}},\eta_j} \left[ L_{D_{\mathbf{y},i}^{(2)}} r f_i(0) \right] \right\|_2^2 \right]}_{=: A_{1,2}^{t+1}}. \end{aligned}$$

The first term  $A_{1,1}^{t+1}$  can be upper bounded through smoothness of  $f_\xi$ , for  $t \geq 1$

$$\begin{aligned}
A_{1,1}^{t+1} &= \mathbb{E}_i \left[ \left\| r f_i(\mathbb{E}_{\eta_{ji}}[g_\eta(\mathbf{x}^t)]) - p_{\text{in}} r f_i(\mathbb{E}_{\eta_{ji}}[g_\eta(\mathbf{x}^t)]) - (1 - p_{\text{in}}) r f_i(\mathbf{y}_i^t + \mathbb{E}_{\eta_{ji}}[g_\eta(\mathbf{x}^t) - g_\eta(\phi_i^t)]) \right\|_2^2 \right] \\
&= (1 - p_{\text{in}})^2 \mathbb{E}_i \left[ \left\| r f_i(\mathbb{E}[g_\eta(\mathbf{x}^t)]) - r f_i(\mathbf{y}_i^t + \mathbb{E}_{\eta_{ji}}[g_\eta(\mathbf{x}^t) - g_\eta(\phi_i^t)]) \right\|_2^2 \right] \\
&\quad + (1 - p_{\text{in}})^2 L_f^2 \mathbb{E}_i \left[ \left\| \mathbb{E}_{\eta_{ji}}[g_\eta(\mathbf{x}^t)] - (\mathbf{y}_i^t + \mathbb{E}_{\eta_{ji}}[g_\eta(\mathbf{x}^t) - g_\eta(\phi_i^t)]) \right\|_2^2 \right] \\
&= (1 - p_{\text{in}})^2 L_f^2 \mathbb{E}_i \left[ \left\| \mathbf{y}_i^t - \mathbb{E}_{\eta_{ji}}[g_\eta(\phi_i^t)] \right\|_2^2 \right] \\
&= (1 - p_{\text{in}})^2 L_f^2 E_y^t.
\end{aligned}$$

For  $t = 0$ ,  $A_{1,1}^1 = 0$ , then

$$\frac{1}{T} \sum_{t=0}^{T-1} A_{1,1}^{t+1} \leq (1 - p_{\text{in}})^2 L_f^2 C_g^2 \frac{1}{T} \sum_{t=0}^{T-1} E_y^{t+1}. \quad (37)$$

On the other hand, with Lemma 6

$$\begin{aligned}
A_{1,2}^{t+1} &\leq p_{\text{in}} \mathbb{E}_i \left[ \left\| r f_i(\mathbb{E}_{\eta_{ji}}[g_\eta(\mathbf{x}^t)]) - \mathbb{E}_{\eta_{ji}} \left[ L_{D_{\mathbf{y}, S_1, i}}^{(2)} r f_i(0) \right] \right\|_2^2 \right] \\
&\quad + (1 - p_{\text{in}}) \mathbb{E}_i \left[ \left\| r f_i(\mathbf{y}_i^t + \mathbb{E}_{\eta_{ji}}[g_\eta(\mathbf{x}^t) - g_\eta(\phi_i^t)]) - \mathbb{E}_{\eta_{ji}} \left[ L_{D_{\mathbf{y}, S_2, i}}^{(2)} r f_i(0) \right] \right\|_2^2 \right] \\
&\leq \frac{p_{\text{in}} C_e^2}{S_1^4} + \frac{(1 - p_{\text{in}}) C_e^2}{S_2^4} \\
&\quad + \frac{C_e^2}{S_2^4}
\end{aligned}$$

where  $D_{\mathbf{y}, S_1, i}^{t+1}$  is the distribution of  $\frac{1}{S_1} \sum_{\eta \in S_1} g_\eta(\mathbf{x}^t)$  and  $D_{\mathbf{y}, S_2, i}^{t+1}$  is the distribution of

$$\mathbf{y}_i^t + \frac{1}{S_2} \sum_{\eta \in S_2} (g_\eta(\mathbf{x}^t) - \mathbb{E}[g_\eta(\phi_i^t)]).$$

Thus the  $A_1^{t+1}$  has the following upper bound

$$\frac{1}{T} \sum_{t=0}^{T-1} A_1^{t+1} \leq (1 - p_{\text{in}})^2 L_f^2 C_g^2 \frac{1}{T} \sum_{t=0}^{T-1} E_y^{t+1} + \frac{C_e^2}{S_2^4}. \quad (38)$$

Consider  $A_2^{t+1}$ . Let us expand  $A_2^{t+1}$  through recursion

$$\begin{aligned}
A_2^{t+1} &= \left\| \mathbb{E}_{i, p_{\text{in}}, \eta, \eta_{ji}}[G_i^{t+1}] - \mathbb{E}[G^{t+1}] \right\|_2^2 \\
&= (1 - p_{\text{out}})^2 \left\| G^t - \mathbb{E}_i[\mathbb{E}_{\eta, \eta_{ji}}[\tilde{G}_i^t]] \right\|_2^2 \\
&= (1 - p_{\text{out}})^2 \left( \left\| \mathbb{E}[G^t] - \mathbb{E}_i[\mathbb{E}_{\eta, \eta_{ji}}[\tilde{G}_i^t]] \right\|_2^2 + E_{\text{var}}^t \right) \\
&= (1 - p_{\text{out}})^2 (A_2^t + E_{\text{var}}^t).
\end{aligned}$$

For  $t = 0$ , we have that  $A_2^1 = 0$ , then average over time gives

$$\frac{1}{T} \sum_{t=0}^{T-1} A_2^{t+1} \leq \frac{(1 - p_{\text{out}})^2}{p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1}.$$

Therefore, the bias has the following bound

$$\frac{1}{T} \sum_{t=0}^{T-1} E_{\text{bias}}^{t+1} \leq (1 - p_{\text{in}})^2 L_f^2 C_g^2 \frac{1}{T} \sum_{t=0}^{T-1} E_y^{t+1} + \frac{C_e^2}{S_2^4} + \frac{(1 - p_{\text{out}})^2}{p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1}.$$

Using Lemma 11

$$\begin{aligned}
\frac{1}{T} \sum_{t=0}^{T-1} E_{\text{bias}}^{t+1} &\leq (1 - p_{\text{in}})^2 L_f^2 C_g^2 \left( \frac{(1 - p_{\text{in}}) C_e^2}{p_{\text{in}} S_2^4} \frac{1}{T} \sum_{t=0}^{T-1} \Xi^t + \frac{2\sigma_g^2}{S_1^2} \right) \\
&\quad + \frac{C_e^2}{S_2^4} + \frac{(1 - p_{\text{out}})^2}{p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1} \\
&\leq \frac{(1 - p_{\text{in}})^3 L_f^2}{p_{\text{in}} S_2^4} \frac{1}{T} \sum_{t=0}^{T-1} \Xi^t + \frac{2(1 - p_{\text{in}})^2 L_f^2 E_F}{S_1^2} + \frac{C_e^2}{S_2^4} + \frac{(1 - p_{\text{out}})^2}{p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1}.
\end{aligned}$$

Using Lemma 10 we have that

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{bias}}^{t+1} &= \frac{(1-p_{\text{in}})^3 L_F^2}{p_{\text{in}} S_2} \left( \frac{6n^2 \gamma^2}{B_2^2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] \right) + \frac{2(1-p_{\text{in}})^2 L_F^2}{S_1} + \frac{C_\varepsilon^2}{S_2^4} + \frac{(1-p_{\text{out}})^2}{p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1} \\ &= \frac{(1-p_{\text{in}})^3 L_F^2}{p_{\text{in}} S_2} \frac{6n^2 \gamma^2}{B_2^2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] + \frac{2(1-p_{\text{in}})^2 L_F^2}{S_1} + \frac{C_\varepsilon^2}{S_2^4} \\ &\quad + \left( \frac{(1-p_{\text{in}})^3 L_F^2}{p_{\text{in}} S_2} \frac{6n^2 \gamma^2}{B_2^2} + \frac{(1-p_{\text{out}})^2}{p_{\text{out}}} \right) \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1}. \end{aligned}$$

Variance. Combine the variance of NestedVR in Lemma 13 and Lemma 2 gives

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1} &= 14 \cdot 32 \left( \left( \frac{p_{\text{out}}}{B_1} + \frac{1-p_{\text{out}}}{B_2} \right) \frac{(1-p_{\text{in}}) 18n^2}{p_{\text{in}} S_2 B_2^2} + \frac{(1-p_{\text{out}})}{B_2} \right) \frac{\gamma^2 L_F^2}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] \\ &\quad + 14 \cdot 96 \left( \frac{p_{\text{out}}}{B_1} + \frac{(1-p_{\text{in}})(1-p_{\text{out}})}{B_2} \right) \frac{L_F^2}{S_1} + \frac{(1-p_{\text{out}})}{T} \frac{8L_F^2}{B_1 S_1}. \end{aligned}$$

Theorem 4 [*E-NestedVR Convergence*] Consider the (FCCO) problem. Under the same assumptions as Theorem 2.

If  $n = O(\varepsilon^{-2/3})$ , then we choose the hyperparameters of E-NestedVR (Algorithm 3) as  $B_1 = B_2 = n, p_{\text{out}} = 1, S_1 = \tilde{L}_F \varepsilon^{-2}, S_2 = \tilde{L}_F \varepsilon^{-1}, p_{\text{in}} = \tilde{L}_F \varepsilon, \gamma = O(\frac{1}{L_F})$ .

If  $n = \Omega(\varepsilon^{-2/3})$ , then we choose the hyperparameters of E-NestedVR as  $B_1 = n, B_2 = \frac{p_{\text{out}}}{n}, p_{\text{out}} = 1/p_{\text{out}}, S_1 = S_2 = \max \left\{ a_3 \sigma_3 C_g \varepsilon^{-1/2}, \tilde{L}_F / (n \varepsilon^2) \right\}, p_{\text{in}} = 1, \gamma = O(\frac{1}{L_F})$ .

Then the output  $\mathbf{x}^s$  of E-NestedVR satisfies:  $\mathbb{E}[k r F(\mathbf{x}^s) k_2^2] \leq \varepsilon^2$ , for nonconvex  $F$ , if the number of iterations

$$T = \Omega(L_F(F(\mathbf{x}^0) - F^*) \varepsilon^{-2}).$$

Proof: Denote.  $G^{t+1} = G_{\text{E-NVR}}^{t+1}$  (36). Using descent lemma (Lemma 3) and bias-variance of E-NestedVR (Lemma 14)

$$\begin{aligned} &\frac{1}{T} \sum_{t=0}^{T-1} k r F(\mathbf{x}^t) k_2^2 + \frac{1}{2T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] \\ &\quad \frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} + \frac{L_F \gamma}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1} + \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{bias}}^{t+1} \\ &\quad \frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T} + \frac{(1-p_{\text{in}})^3 L_F^2}{p_{\text{in}} S_2} \frac{6n^2 \gamma^2}{B_2^2} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] + \frac{2(1-p_{\text{in}})^2 L_F^2}{S_1} + \frac{C_\varepsilon^2}{S_2^4} \\ &\quad + \left( \frac{(1-p_{\text{in}})^3 L_F^2}{p_{\text{in}} S_2} \frac{6n^2 \gamma^2}{B_2^2} + \frac{(1-p_{\text{out}})^2}{p_{\text{out}}} + L_F \gamma \right) \frac{1}{T} \sum_{t=0}^{T-1} E_{\text{var}}^{t+1}. \end{aligned}$$

As we would like the right-hand side to be bounded by either  $\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2]$  or  $\varepsilon^2$ .

- Bound on  $\frac{2(F(\mathbf{x}^0) - F^*)}{\gamma T}$  with  $\varepsilon^2$ , i.e.

$$\gamma T \leq (F(\mathbf{x}^0) - F^*) \varepsilon^{-2} \quad (39)$$

- Coefficient of  $\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2]$  is bounded by  $\frac{1}{4}$ , i.e.

$$\frac{(1-p_{\text{in}})^3 L_F^2}{p_{\text{in}} S_2} \frac{6n^2 \gamma^2}{B_2^2} \leq \frac{1}{4}$$

which can be achieved by choosing the following step size

$$\gamma = \frac{p_{\text{out}} p_{\text{in}} S_1}{5 L_F (1-p_{\text{in}})^{3/2}}. \quad (40)$$

- Bound on  $\frac{2(1-p_{\text{in}})^2 L_F^2}{S_1}$  with  $\varepsilon^2$

$$\frac{2(1-p_{\text{in}})^2 L_F^2}{S_1} \leq \varepsilon^2. \quad (41)$$

- Bound  $\frac{C_e^2}{S_2^4}$  with  $\varepsilon^2$ . This leads to

$$S_2 \geq \sqrt{\frac{C_e}{\varepsilon}}. \quad (42)$$

- Bound on the variance. First notice from (40) and  $\gamma \leq \frac{1}{2L_F}$ ,

$$\begin{aligned} \frac{(1-p_{\text{in}})^3 E_F^2}{p_{\text{in}} S_2} \frac{6n^2}{B_2^2} \gamma^2 &\leq \frac{1}{4} \cdot \frac{(1-p_{\text{out}})^2}{p_{\text{out}}} \\ L_F \gamma &\leq \frac{1}{2} \cdot \frac{(1-p_{\text{out}})^2}{p_{\text{out}}}. \end{aligned}$$

Therefore, we only need to consider the upper bound on

$$\begin{aligned} &\frac{(1-p_{\text{out}})^2}{p_{\text{out}}} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_{\text{var}}^{t+1} \\ &14 \cdot 32 \frac{(1-p_{\text{out}})^2}{p_{\text{out}}} \left( \left( \frac{p_{\text{out}}}{B_1} + \frac{1-p_{\text{out}}}{B_2} \right) \frac{(1-p_{\text{in}}) 18n^2}{p_{\text{in}} S_2 B_2^2} + \frac{(1-p_{\text{out}})}{B_2} \right) \frac{\gamma^2 L_F^2}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|G^{t+1}\|_2^2] \\ &+ 14 \cdot 96 \frac{(1-p_{\text{out}})^2}{p_{\text{out}}} \left( \frac{p_{\text{out}}}{B_1} + \frac{(1-p_{\text{in}})(1-p_{\text{out}})}{B_2} \right) \frac{E_F^2}{S_1} + \frac{(1-p_{\text{out}})^3}{p_{\text{out}} T} \frac{8E_F^2}{B_1 S_1}. \end{aligned}$$

We impose the constraints for each term

$$\begin{aligned} \frac{(1-p_{\text{out}})^2}{p_{\text{out}}} \frac{p_{\text{out}}}{B_1} \frac{(1-p_{\text{in}}) 18n^2}{p_{\text{in}} S_2 B_2^2} L_F^2 \gamma^2 &\leq 1 \\ \frac{(1-p_{\text{out}})^2}{p_{\text{out}}} \frac{1-p_{\text{out}}}{B_2} \frac{(1-p_{\text{in}}) 18n^2}{p_{\text{in}} S_2 B_2^2} L_F^2 \gamma^2 &\leq 1 \\ \frac{(1-p_{\text{out}})^2}{p_{\text{out}}} \frac{1-p_{\text{out}}}{B_2} L_F^2 \gamma^2 &\leq 1 \\ \frac{(1-p_{\text{out}})^2}{p_{\text{out}}} \frac{p_{\text{out}}}{B_1} \frac{E_F^2}{S_1} &\leq \varepsilon^2 \\ \frac{(1-p_{\text{out}})^2}{p_{\text{out}}} \frac{(1-p_{\text{in}})(1-p_{\text{out}})}{B_2} \frac{E_F^2}{S_1} &\leq \varepsilon^2 \\ \frac{(1-p_{\text{out}})^3}{p_{\text{out}} T} \frac{8E_F^2}{B_1 S_1} &\leq \varepsilon^2. \end{aligned}$$

These can be simplified as

$$\gamma \cdot \frac{p_{\text{in}} p_{\text{out}}}{(1-p_{\text{out}})} \frac{1}{B_1} \frac{1}{p_{\text{in}} L_F} \leq 1 \quad (43)$$

$$\gamma \cdot \frac{p_{\text{in}} p_{\text{out}}^2}{(1-p_{\text{out}})^{3/2}} \frac{1}{B_1} \frac{1}{p_{\text{in}} L_F} \leq 1 \quad (44)$$

$$\gamma \cdot \frac{1}{(1-p_{\text{out}})^{3/2}} \frac{1}{L_F} \leq 1 \quad (45)$$

$$B_1 S_1 \leq \frac{(1-p_{\text{out}})^2 E_F^2}{\varepsilon^2} \quad (46)$$

$$B_1 S_1 \leq \frac{(1-p_{\text{out}})^3 (1-p_{\text{in}}) E_F^2}{\varepsilon^2 p_{\text{out}}^2} \quad (47)$$

$$B_1 S_1 \leq \frac{(1-p_{\text{out}})^3 E_F^2}{T \varepsilon^2 p_{\text{out}}} \quad (48)$$

- Constraints from Lemma 14

$$\gamma^2 L_F^2 \max \left\{ \frac{(1-p_{\text{in}}) 18}{p_{\text{in}} S_2 B_2}, \frac{1-p_{\text{out}}}{B_2} \frac{(1-p_{\text{in}}) 18n^2}{p_{\text{in}} S_2 B_2^2}, \frac{(1-p_{\text{out}})}{B_2} \right\} \leq \frac{1}{16} \leq \frac{1}{6}$$

which can be translated to

$$\gamma \cdot \frac{p_{\text{in}}}{L_F} \frac{1}{p_{\text{in}} L_F} \frac{1}{B_2} \leq 1 \quad (49)$$

$$\gamma \cdot \frac{p_{\text{in}} p_{\text{out}}}{L_F} \frac{1}{p_{\text{in}} L_F} \frac{1}{B_2} \leq 1 \quad (50)$$

$$\gamma \cdot \frac{1}{L_F} \frac{1}{p_{\text{out}}} \leq 1 \quad (51)$$

- Constraint from sufficient decrease lemma:

$$\gamma \leq \frac{1}{2L_F}. \quad (52)$$




Figure 3: Performances of algorithms and their extrapolated versions on the instrumental variable regression task. The shaded region represents the 95%-confidence interval.

We simplify the conditions noticing that 1) (48) is weaker than (46); 2) (45) and (51) are weaker than (52). Combine all the constraints on  $\gamma$ , i.e. (43), (44), (49), (50), (52)

$$\gamma \cdot \frac{1}{L_F} \min \left\{ \min \left\{ 1, \frac{p_{\text{out}}}{1-p_{\text{out}}} \right\} \frac{p_{\text{in}} p_{\text{out}}}{(1-p_{\text{out}})^2} \frac{p_{\text{in}}}{1-p_{\text{in}}} \frac{1}{L_F}, \min \left\{ 1, \frac{p_{\text{out}}}{1-p_{\text{out}}} \right\} \frac{p_{\text{in}} p_{\text{out}}}{1-p_{\text{in}}} \frac{1}{5L_F(1-p_{\text{in}})^{3/2}} \right\}.$$

This can be simplified as an upper bound

$$\gamma \cdot \frac{1}{L_F} \min \left\{ \frac{p_{\text{in}} p_{\text{out}}}{1-p_{\text{in}}}, \frac{p_{\text{in}} p_{\text{out}}}{1-p_{\text{out}}} \frac{p_{\text{in}}}{1-p_{\text{in}}}, \frac{p_{\text{in}} p_{\text{out}}}{1-p_{\text{in}}} \frac{1}{1-p_{\text{out}}}, 1 \right\}.$$

Now we consider two sets of hyperparameters depending on the size of  $n$  Case 1: For  $n = O(\varepsilon^{-2/3})$ , we choose the following set of hyperparameters

$$B_1 = B_2 = n, \quad p_{\text{out}} = 1, \quad S_1 = \tilde{L}_F^2 \varepsilon^{-2}, \quad S_2 = \tilde{L}_F \varepsilon^{-1}, \quad p_{\text{in}} = \tilde{L}_F^{-1} \varepsilon.$$

Then we have  $\gamma \cdot \frac{1}{L_F} \min \left\{ \frac{p_{\text{in}} p_{\text{out}}}{1-p_{\text{in}}}, 1 \right\} g = \frac{1}{L_F}$ , we have the total sample complexity of

$$B_2 S_2 T = \frac{B_2 S_2 T \gamma}{\gamma} \stackrel{(39)}{=} \frac{F(\mathbf{x}^0) F^* B_2 S_2}{\varepsilon^2 \gamma} = \frac{(F(\mathbf{x}^0) F^*) n L_F L_F}{\varepsilon^3}$$

Case 2: For  $n = \Omega(\varepsilon^{-2/3})$ , we choose the following set of hyperparameters

$$B_1 = n, \quad B_2 = \frac{p_{\text{out}}}{n}, \quad p_{\text{out}} = \frac{1}{n}.$$

In this case, (46) is stronger than (47) which requires  $S_1 \& \frac{L_F^2}{n \varepsilon^2}$

$$S_1 = S_2 = \max \left\{ \tilde{\sigma}_{\text{bias}}^{1/2} \varepsilon^{-1/2}, \frac{\sigma_{\text{in}}^2}{n \varepsilon^2} \right\}, \quad p_{\text{in}} = 1$$

Then we have  $\gamma \cdot \frac{1}{L_F} \min \left\{ \frac{p_{\text{in}} p_{\text{out}}}{1-p_{\text{out}}}, 1 \right\} g = \frac{1}{L_F}$ , we have the total sample complexity of

$$B_2 S_2 T = \frac{B_2 S_2 T \gamma}{\gamma} \stackrel{(39)}{=} \frac{F(\mathbf{x}^0) F^* B_2 S_2}{\varepsilon^2 \gamma} = (F(\mathbf{x}^0) F^*) \max \left\{ \frac{p_{\text{out}} \sigma_{\text{bias}}^{1/2}}{\varepsilon^{2.5}}, \frac{\sigma_{\text{in}}^2}{n \varepsilon^4} \right\}.$$

By picking  $\mathbf{x}^s$  uniformly at random among  $f \mathbf{x}^t g_{t=0}^T$ , we get the desired guarantee.

## E Missing Details from Section 5

### E.1 Application of Instrumental Variable Regression

Instrumental variable regression is a popular technique in econometrics that aims to estimate the causal effect of input  $X$  on a confounded outcome  $Y$  with an instrument variable  $Z$ , which, for a given  $X$ , is conditionally

independent of  $Y$ . As noted by [Muandet et al., 2020], the instrumental variable regression is a special case of the CSO problem. The instrumental variable regression problem is phrased as:

$$\min_w \mathbb{E}_{\xi=(Y,Z)} [\ell(Y, \mathbb{E}_{X|Z}[g_X(w)])]$$

where  $\xi = (Y, Z), \eta = X$ . This can be viewed in the CSO framework with  $f_\xi(\cdot) = \ell(Y, \cdot)$ . We choose  $\ell(y, \hat{y}) = \log \cosh(y - \hat{y})$  as regression loss function and  $g_X(w) = w^\top X$  to be a linear regressor. In this case,  $f_\xi \in \mathcal{C}^1$  with  $\Gamma f_\xi(\hat{y}) = \tanh(\hat{y} - Y)$ , and our results from Sections 3 and 4 apply. We generate the data similar to [Singh et al., 2019]

$$\begin{aligned} Z &\sim \text{Uniform}([ -3, 3]^2), \quad e \sim N(0, 1), \quad \delta \sim N(0, 0.1), \quad \gamma \sim \text{Exponential}(10) \\ X &= \frac{1}{2}z_1 + \frac{1}{2}e + \gamma, \quad Y = X + e + \delta \end{aligned}$$

where  $z_1$  is the first dimension of  $Z$ ,  $e$  is the confounding variable and  $\delta, \gamma$  are noises. In this experiment, we solve the instrumental variable regression using BSGD, BSpiderBoost, NestedVR and their extrapolated variants described in Sections 3 and 4. In each pair of experiments, the samples used per iteration are fixed same, i.e.: 1) BSGD uses  $m = 2$  and E-BSGD uses  $m = 1$ ; 2) For BSpiderBoost and E-BSpiderBoost, we use cycle length of 10, small batch and large batch in Spider to be 10 and 100 respectively, and we choose inner batch sizes  $m = 2$  for BSpiderBoost and  $m = 1$  for E-BSpiderBoost; 3) For NestedVR and E-NestedVR, we fix the outer batch size to 10 and 5 respectively, and choose fix the inner Spider Cycle to be 10 with large batch 100 and small batch 10. The results are presented in Figure 3. As is quite evident, the extrapolation variants achieve faster convergence in all 3 cases, confirming our theoretical findings.

## E.2 Application of First-order MAML

Over the past few years, the MAML framework [Finn et al., 2017] has become quite popular for few-shot supervised learning and meta reinforcement learning tasks. The first-order Model-Agnostic Meta-Learning (MAML) can be formulated mathematically as follows:

$$\min_{\mathbf{x}} \mathbb{E}_{i \sim p, D_{\text{query}}^i} \ell_i \left( \mathbb{E}_{D_{\text{supp}}^i} (\mathbf{x} - \alpha \Gamma \ell_i(\mathbf{x}, D_{\text{supp}}^i)), D_{\text{query}}^i \right)$$

where  $\alpha$  is the step size,  $D_{\text{supp}}^i$  and  $D_{\text{query}}^i$  are meta-training and meta-testing data respectively and  $\ell_i$  being the loss function of task  $i$ . Stated in the CSO framework,  $f_\xi(\mathbf{x}) := \ell_i(\mathbf{x}, D_{\text{query}}^i)$  and  $g_\eta(\mathbf{x}, \xi) := \mathbf{x} - \alpha \Gamma \ell_i(\mathbf{x}, D_{\text{supp}}^i)$  where  $\xi = (i, D_{\text{query}}^i)$  and  $\eta = D_{\text{supp}}^i$ .

In this context, lots of popular choices for  $f_\xi$  are smooth. For illustration purposes, we now discuss a widely used sine-wave few-shot regression task as appearing from the work of Finn et al. [2017], where the goal is to do a few-shot learning of a sine wave,  $A \sin(t - \phi)$ , using a neural network  $\Phi_{\mathbf{x}}(t)$  with smooth activations, where  $A$  and  $\phi$  represent the unknown amplitude and phase, and  $\mathbf{x}$  denotes the model weight. Each task  $i$  is characterized by  $(A^i, \phi^i, D_{\text{query}}^i)$ . In the first-order MAML training process, we randomly select a task  $i$ , and draw training data  $\eta = D_{\text{supp}}^i$ . Define the loss function for a given dataset  $D$  as  $\ell_i(\Phi_{\mathbf{x}}; D) = \frac{1}{2} \mathbb{E}_{t \sim D} \|A^i \sin(t - \phi^i) - \Phi_{\mathbf{x}}(t)\|_2^2$ . We then establish the outer function  $f_i(\mathbf{x}) = \ell_i(\Phi_{\mathbf{x}}; D_{\text{query}}^i)$  and inner function  $g_\eta(\mathbf{x}) = \mathbf{x} - \alpha \Gamma_{\mathbf{x}} \ell_i(\Phi_{\mathbf{x}}; D_{\text{supp}}^i)$ . As  $f_i$  is smooth, our results are applicable.

In Figure 4, we show the results of BSGD and E-BSGD applied to this problem. In this experiment, the amplitude  $A$  is drawn from a uniform distribution  $U(0.1, 5)$  and the phase  $\phi$  is drawn from  $U(0, \pi)$ . Both  $D_{\text{supp}}$  and  $D_{\text{query}}$  are independently drawn from  $U([ -5, 5])$ . The step size is set to  $\alpha = 0.01$ . The batch size is fixed to 10. The performances of BSGD and E-BSGD are very close. This is not surprising because finetuning step size  $\alpha$  is chosen to be small which significantly reduces the variance of  $g_\eta$ , making the bias of meta gradient to be very small ( $O(\alpha^2)$ ). Therefore, we observe similar performance of BSGD and E-BSGD. Similar trend also holds for BSpiderBoost and NestedVR compared to their extrapolated variants.

## E.3 Application of Deep Average Precision Maximization

The areas under precision-recall curve (AUPRC) has an unbiased point estimator that maximizes average precision (AP) [Qi et al., 2021a; Wang et al., 2022a]. Let  $S_+$  and  $S_-$  be the set of positive and negative




Figure 4: Performance of BSGD vs. E-BSGD on the few-shot sinusoid regression task.

samples and  $S = S \cup S_+$ . Let  $h_{\mathbf{w}}(\cdot)$  be a classifier parameterized with  $\mathbf{w}$  and  $\ell$  be a surrogate function, such as logistic or sigmoid. A smooth surrogate objective for maximizing average precision can be formulated as [Wang and Yang, 2022]:

$$F(\mathbf{w}) = \frac{1}{|S_+|} \sum_{\mathbf{x}_i \in S_+} \frac{\ell(h_{\mathbf{w}}(\mathbf{x}) - h_{\mathbf{w}}(\mathbf{x}_i))}{\sum_{\mathbf{x} \in S} \ell(h_{\mathbf{w}}(\mathbf{x}) - h_{\mathbf{w}}(\mathbf{x}_i))}$$

This problem can be seen as a conditional stochastic optimization problem with  $g_i(\mathbf{w}) = [\sum_{\mathbf{x} \in S_+} \ell(h_{\mathbf{w}}(\mathbf{x}) - h_{\mathbf{w}}(\mathbf{x}_i)), \sum_{\mathbf{x} \in S} \ell(h_{\mathbf{w}}(\mathbf{x}) - h_{\mathbf{w}}(\mathbf{x}_i))]$  and  $f_i: \mathbb{R}^2 \rightarrow \mathbb{R}$  is defined as  $f_i(\mathbf{y}) = \frac{[\mathbf{y}]_1}{[\mathbf{y}]_2}$  where  $[\mathbf{y}]_k$  denotes the  $k$ th coordinate of a vector  $\mathbf{y} \in \mathbb{R}^2$ . During the stochastic optimization of this objective, we draw uniformly at random  $\xi := \mathbf{x}_i$  (drawn from the set  $S_+$ ) as a positive sample and  $\eta/\xi = [F_{\mathbf{x}_1}, F_{\mathbf{x}_2}]$  where set  $\mathbf{x}_1$  is drawn uniformly at random from  $S_+$  and  $\mathbf{x}_2$  is drawn uniformly at random from  $S$  and functional  $F_{\mathbf{x}}(\mathbf{w}) := \ell(h_{\mathbf{w}}(\mathbf{x}) - h_{\mathbf{w}}(\mathbf{x}_i))$ . Note that  $f_i \in \mathcal{C}^1$  is smooth with gradient

$$\nabla f_i(\mathbf{y}) = \begin{bmatrix} \frac{1}{[\mathbf{y}]_2} \\ -\frac{[\mathbf{y}]_1}{([\mathbf{y}]_2)^2} \end{bmatrix}.$$

Therefore, our results from Sections 3 and 4 again apply.