

# An approximate control variates approach to multifidelity distribution estimation

Ruijian Han<sup>1</sup>, Akil Narayan<sup>2</sup>, and Yiming Xu<sup>3</sup>

<sup>1</sup>Department of Applied Mathematics, The Hong Kong Polytechnic University

<sup>2</sup>Scientific Computing and Imaging Institute, and Department of Mathematics, University of Utah

<sup>3</sup>Corporate Model Risk, Wells Fargo

## Abstract

Forward simulation-based uncertainty quantification that studies the output distribution of quantities of interest (QoI) is a crucial component for computationally robust statistics and engineering. There is a large body of literature devoted to accurately assessing statistics of QoI, and in particular, multilevel or multifidelity approaches are known to be effective, leveraging cost-accuracy tradeoffs between a given ensemble of models. However, effective algorithms that can estimate the full distribution of outputs are still under active development. In this paper, we introduce a general multifidelity framework for estimating the cumulative distribution functions (CDFs) of vector-valued QoI associated with a high-fidelity model under a budget constraint. Given a family of appropriate control variates obtained from lower fidelity surrogates, our framework involves identifying the most cost-effective model subset and then using it to build an approximate control variates estimator for the target CDF. We instantiate the framework by constructing a family of control variates using intermediate linear approximators and rigorously analyze the corresponding algorithm. Our analysis reveals that the resulting CDF estimator is uniformly consistent and budget-asymptotically optimal, with only mild moment and regularity assumptions. The approach provides a robust multifidelity CDF estimator that is adaptive to the available budget, does not require *a priori* knowledge of cross-model statistics or model hierarchy, and is applicable to general output dimensions. We demonstrate the efficiency and robustness of the approach using several test examples.

## 1 Introduction

Physical systems are often modeled with computational simulations or emulators, and as such, understanding the error in these constructed approximations is of utmost importance. One particular source of “error” is due to the input uncertainty in these models, either through uncertainty in input parameters (which can be finite- or infinite-dimensional) or through modeled stochasticity in the system, e.g., systems driven with white noise processes. To make the resulting models trustworthy, it is crucial to quantify the resulting uncertainty in QoI; that is, to estimate the QoI’s distribution or some statistical summary of it. One popular approach for achieving this is through Monte Carlo (MC) simulation, which is easy to implement and provides robust results but has a slow convergence rate. A typical MC procedure requires drawing a large number of samples or running repeated experiments, which is expensive given the increasing complexity of modern problems.

To address this issue, methods based on multilevel/multifidelity modeling have been developed to estimate the statistics of QoI associated with the (high-fidelity) model [8, 23, 22, 14, 26, 6, 29, 27, 13, 15, 5, 7]; see [9, 24] for a detailed survey on the related methods. The core idea behind multilevel/multifidelity methods lies in leveraging models of different accuracies and costs to improve computational efficiency. However, a major limitation of the existing literature is that it predominantly focuses on the estimation

of the statistical mean of the QoI (or perhaps some other scalar-valued descriptive statistics), providing only partial insight into the uncertainty of the model output. A more comprehensive understanding would require assessing, for example, higher order statistics of the QoI, or even the entire distribution.

Existing methods to estimate CDFs in the multilevel/multifidelity setup have seen notable success [9, 19, 10, 17, 2, 30]. In [9], the authors proposed a multilevel approach to computing the CDFs of univariate random variables arising from stochastic differential equations (SDEs) and derived an upper bound for the cost in terms of the error. The methodology in [9] was further developed and applied in a number of subsequent works [19, 17, 10, 2]. In particular, [19] designed an *a posteriori* optimization strategy to calibrate the smoothing function and showed its superiority over MC in oil reservoir simulations; [17] generalized the ideas in [9] to approximate more general parametric expectations such as characteristic functions; [10] applied an adaptive approach for parameter selection that yields an improved cost bound; [2] provides a novel computable error estimator to enhance algorithm tuning. Despite the substantive contributions of these approaches, nearly all of them make relatively restrictive assumptions regarding model hierarchy (e.g., the model cost versus accuracy tradeoffs), and do not immediately extend to the general non-hierarchical multifidelity setup. For this more general multifidelity estimation of CDFs, the only work we are aware of is the adaptive explore-then-commit algorithm for distribution learning (AETC-d) [30]. However, the large-budget performance of AETC-d is restricted by its own set of statistical assumptions that are often too stringent to satisfy in practice. Moreover, the output of the QoI in all the above references is assumed to be a scalar.

## 1.1 Contributions of this article

The main goal of this article is to provide novel solutions that mitigate the deficiencies described above through the development of a workable and efficient algorithm for estimating the CDF in a general non-hierarchical multifidelity approximation setting under computational budget constraints. Our aim is to design a method that satisfies the following criteria: 1) it requires as input neither cross-model statistics nor model hierarchy; 2) it can provide distributional estimates for vector-valued QoI outputs, and 3) it is empirically robust and enjoys theoretical guarantees. Although our approach uses a similar meta-algorithm as in [29, 30] (all borrowing ideas from the explore-then-commit algorithm in bandit learning [18]), it contains a substantial amount of new ingredients that allow for both significant generalizations of learning scope and improvement of robustness. In more technical language, our contributions are twofold:

- We propose a control variates-based exploration-exploitation strategy for multifidelity CDF estimation under a budget constraint. The *exploration* step leverages statistical estimation to select a subset of low-fidelity models for the control variates construction, followed by the *exploitation* step that utilizes the learned information to build an approximate control variates estimator for the target CDF. Our estimator for the CDF applies to both scalar-valued and vector-valued QoI, which differentiates it from existing methods that apply only to scalar-valued QoI.
- Through examination of the average weighted- $L^2$  loss that balances errors in exploration and exploitation, we design a new meta-algorithm, the control variates multifidelity distribution learning algorithm (cvMDL, Algorithm 1), that accomplishes model (subset) selection and CDF estimation. Using control variates constructed from linear approximators, we establish both uniform consistency and budget-asymptotic optimality of the estimator produced by cvMDL. Our analysis illustrates that the proposed procedure is more robust to the restrictive model assumptions from [30].

A verbatim usage of our approaches produces an estimator of a CDF that itself is not a distribution function: The estimator is not necessarily monotonic (yet it does enjoy the previously-mentioned theoretical guarantees). To mitigate this artifact, we utilize an empirical algorithmic correction that restores monotonicity of the estimated CDFs, and additionally makes its manipulation more computationally convenient (e.g. for extraction of quantiles and conditional expectations), and we observe that in some cases this empirical correction further reduces errors.

An outline of the paper is as follows. Section 2 gives a brief overview of the control variates method. Section 3 introduces a multifidelity CDF estimation framework based on approximate control variates estimators. An explicit formula for the mean weighted- $L^2$  loss that balances exploration and exploitation is derived and used to obtain a lower bound for the error reduction. Section 4 provides a computational construction for the control variates through linear approximators. Section 5 develops our new meta-algorithm (cvMDL) that accomplishes autonomous model selection together with an algorithmic correction to preserve the monotonicity of the resulting CDF estimators. The meta-algorithm cvMDL itself does not specify how to compute the control variates: A specialization to using the linear approximations from Section 4 yields a computationally explicit algorithm that we study in detail, establishing both uniform consistency and budget-asymptotic optimality. Section 6 contains a detailed simulation study and Section 7 includes the concluding remarks.

## 1.2 Notation

For  $n \in \mathbb{N}$ , let  $[n] := \{1, \dots, n\}$ . We use bold upper-class and lower-class letters to denote matrices and vectors, respectively. The Euclidean ( $\ell^2$ ) norm on a vector  $\mathbf{v}$  is denoted  $\|\mathbf{v}\|_2$ . For a matrix  $\mathbf{A}$ ,  $\mathbf{A}^\top$  is the transpose and  $\mathbf{A}^\dagger$  is the pseudoinverse;  $\mathbf{A}^\dagger$  coincides with the regular inverse  $\mathbf{A}^{-1}$  when  $\mathbf{A}$  is invertible. The  $i$ -th column of  $\mathbf{A}$  is denoted by  $\mathbf{A}^{(i)}$ . The Frobenius norm of  $\mathbf{A}$  is denoted by  $\|\mathbf{A}\|_F = (\sum_i \|\mathbf{A}^{(i)}\|_2^2)^{1/2}$ . For a set  $T \subseteq \mathbb{R}^d$ , we denote its interior as  $T^\circ$ , and  $\mathbf{1}_T(\mathbf{x}) = \mathbf{1}_{\mathbf{x} \in T}$  as the indicator function on  $T$ . For two vectors  $\mathbf{x} = (\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(d)})$  and  $\mathbf{y} = (\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(d)})$ , we use  $\vee$  and  $\wedge$  to denote the componentwise *max* and *min* operators, respectively, i.e.,

$$\begin{aligned}\mathbf{x} \vee \mathbf{y} &= \left( \max\{\mathbf{x}^{(1)}, \mathbf{y}^{(1)}\}, \dots, \max\{\mathbf{x}^{(d)}, \mathbf{y}^{(d)}\} \right) \\ \mathbf{x} \wedge \mathbf{y} &= \left( \min\{\mathbf{x}^{(1)}, \mathbf{y}^{(1)}\}, \dots, \min\{\mathbf{x}^{(d)}, \mathbf{y}^{(d)}\} \right).\end{aligned}$$

Moreover, we say  $\mathbf{x} \leq \mathbf{y}$  if  $\mathbf{x}^{(i)} \leq \mathbf{y}^{(i)}$  for all  $i \in [d]$ . For a random vector  $X \in \mathbb{R}^d$ , we let  $F_X(\mathbf{x}) = \mathbb{P}(X^{(1)} \leq \mathbf{x}^{(1)}, \dots, X^{(d)} \leq \mathbf{x}^{(d)})$  denote its CDF. For two sequences of random variables  $\{a_m(\omega)\}$  and  $\{b_m(\omega)\}$  where  $\omega$  is a probabilistic event, we write  $a_m(\omega) \lesssim b_m(\omega)$  if almost surely (a.s.),  $a_m(\omega) \leq C(\omega)b_m(\omega)$  for all  $m \in \mathbb{N}$ , where the constant  $C(\omega)$  is independent of  $m$ .

We consider the QoI that are output from computational models as random variables that jointly lie in some common probability space  $(\Omega, \mathcal{F}, \mathbb{P})$ . For convenience, we let

$$Y = (Y^{(1)}, \dots, Y^{(d)})^\top \in \mathbb{R}^d \quad X_i = (X_i^{(1)}, \dots, X_i^{(d_i)})^\top \in \mathbb{R}^{d_i} \quad i \in [n]$$

denote the high-fidelity and the  $i$ -th low-fidelity model output, respectively, where  $d, d_i \in \mathbb{N}$  are the corresponding output dimensions of  $Y$  and  $X_i$ , and there are  $n$  low-fidelity models in total.  $\mathbb{E}[\cdot]$ ,  $\text{Var}[\cdot]/\text{Cov}[\cdot]$ , and  $\text{Corr}[\cdot]$  are the expectation, variance/covariance, and correlation operators respectively, and  $\perp\!\!\!\perp$  stands for probabilistic independence.

We assume the sampling costs for  $Y$  and  $X_1, \dots, X_n$ , denoted by positive numbers  $c_0$  and  $c_1, \dots, c_n$ , are deterministic and known. For  $S \subseteq [n]$ , let  $c_S = \sum_{i \in S} c_i$ , corresponding to the cost of sampling all (low-fidelity) models from subset  $S$ . We let  $B > 0$  be the total budget, which is deterministic and known. Moreover, for  $S \subseteq [n]$ , we let  $X_S = (X_i^\top)_{i \in S}^\top \in \mathbb{R}^{|d_S|}$ , where  $d_S = \sum_{i \in S} d_i$ , and  $X_{S^+} = (1, X_S^\top)^\top$ , where the latter will be used when considering a linear model approximation with the intercept/bias term.

The central goal in the rest of the article is to develop a multifidelity estimator for  $F_Y(\mathbf{x})$  through drawing samples of  $(Y, X_{[n]})$  and of  $X_S$  for some adaptively-determined  $S \subseteq [n]$ , subject to the sampling cost not exceeding the total budget constraint  $B > 0$ . No other high-level assumptions are made. In other words, we assume only that  $Y$  is a known high-fidelity model, we do not assume any ordering/hierarchy in the models  $X_{[n]}$ , and we do not assume known statistics (e.g. correlations) between any models.

### 1.3 cvMDL summary

The notation we have introduced is enough to present the overall cvMDL algorithm, although actual computations that make the algorithm practical require more technical details which will be provided in Section 2 through Section 5.

The cvMDL (meta-)algorithm is shown in Algorithm 1. The idea is as follows: We first gather  $m$  full joint samples of  $(Y, X_{[n]})$  through *exploration* to identify (i) how models are related, (ii) which model subset  $S$  optimally balances cost versus accuracy, and (iii) whether more samples  $m$  are needed to certify a robust exploration or whether the choice of  $S$  is statistically robust enough to proceed with *exploitation*. All these decisions are made in lines 1 through 13. In the literature on statistical decision-making, the choice of taking exploration samples from the full set of models  $(Y, X_{[n]})$  is called a *uniform exploration policy*. Exploration is followed by *exploitation*, where we exhaust the remaining computational budget to sample the optimal model subset  $X_S$ . The exploitation in line 14 is an application of a particular approximate control variates estimator for  $F_Y$ . More detailed descriptions of the steps in Algorithm 1 are as follows:

- *Collect a minimal amount of data:* Line 1 ensures that the number of exploration samples  $m$  is set large enough so that non-degenerate statistics can be computed.
- *Analyze model subsets  $S$ :* Lines 3 through 6 identify for each model subset  $S$  both an estimated number of optimal exploration samples  $\hat{m}_S^*$  along with the corresponding loss function  $\hat{L}(m \vee \hat{m}_S^*; m)$ . The value  $\hat{L}(z; m)$  is an estimator with the currently-available  $m$  exploration samples and measures the estimated loss if we eventually use  $z$  exploration samples. We require the input  $z \leftarrow m \vee \hat{m}_S^*$  since if  $m > \hat{m}_S^*$  then the number of exploration samples should be  $m$ , and not  $\hat{m}_S^*$  (we cannot take fewer exploration samples than already committed and we assume  $-\hat{L}(z; m)$  is unimodal). The definitions of  $\hat{L}$ ,  $\hat{m}_S^*$ , and  $\hat{S}^*$  are given in (5.7) and (5.8).
- *Decide whether to continue exploring:* Lines 7 through 13 decide whether to continue exploring or to switch to exploitation. First, the estimated optimal model subset  $\hat{S}^*$  is computed by choosing the minimal loss from the previous bullet. If the current number of exploration samples ( $m$ ) meets or exceeds the estimated optimal number of samples  $\hat{m}_{\hat{S}^*}^*$  required for the optimal subset  $\hat{S}^*$ , then exploration stops and we proceed to exploitation. Otherwise, we continue exploration, with the precise number of additional exploration samples determined by the function  $Q(\cdot, \cdot)$  that is defined in (5.13).
- *Construct an exploitation estimator:* After exploration terminates and an “optimal” model subset  $\hat{S}^*$  has been identified, line 14 expends the remaining computational budget on sampling  $X_{\hat{S}^*}$  to construct a CDF estimator  $\tilde{F}_{\hat{S}^*}$  for  $F_Y$ , which is defined in (5.11).

The precise details of how the loss function is computed and the CDF estimator is constructed is the topic of Section 5, with Sections 2 to 4 serving to make requisite mathematical and statistical definitions.

A more detailed version of the algorithm is given in Algorithm 2 in Section A, which lists more explicit computational steps that must be taken. The coming sections are devoted to the theoretical construction of quantities in Algorithm 1; in particular Sections 2 and 3 provide a construction of a loss function that is the integral part of the exploration decision-making.

## 2 Control variates

We first introduce the control variates method, which is a standard approach for variance reduction in MC simulation. For a random variable  $X$  with bounded variance  $\sigma_X^2 > 0$ , the size- $m$  MC estimator for  $\mathbb{E}[X]$  based on i.i.d. data  $X_\ell$ ,  $\hat{x} = \sum_{\ell \in [m]} X_\ell / m$ , is unbiased and has variance  $\sigma_X^2 / m$ . Given a random vector  $Z = (Z^{(1)}, \dots, Z^{(d)})^\top \in \mathbb{R}^d$  that lives in the same probability space as  $X$ , one may use joint i.i.d samples of  $(X, Z)$ , i.e.,  $(X_\ell, Z_\ell^\top) = (X_\ell, Z_\ell^{(1)}, \dots, Z_\ell^{(d)})$ ,  $\ell \in [m]$ , to construct a control

---

**Algorithm 1:** cvMDL: See Section 1.3 for accompanying discussion.

---

**Input:**  $B$ : total budget, exploration = TRUE

**Output:** an estimator for  $F_Y(\mathbf{x})$

- 1: Initialize exploration sample size  $m = \sum_{i \in [n]} d_i + 2$ , generate  $m$  samples of  $(Y, X_{[n]})$ .
  - 2: **while** exploration = TRUE **do**
  - 3:   **for**  $S \subseteq [n]$  **do**
  - 4:     Compute  $\hat{m}_S^*$  and  $\hat{L}_S(\cdot; m)$ ;
  - 5:     Compute the minimal expected loss  $\hat{L}_S(m \vee \hat{m}_S^*; m)$ ;
  - 6:   **end for**
  - 7:   Choose  $\hat{S}^* = \arg \min_{S \subseteq [n]} \hat{L}_S(m \vee \hat{m}_S^*; m)$ ;
  - 8:   **if**  $m < \hat{m}_{\hat{S}^*}^*$  **then**
  - 9:     Increase  $m$ :  $m \leftarrow Q(m, \hat{m}_{\hat{S}^*}^*)$ , generate additional samples of  $(Y, X_{[n]})$ .
  - 10:   **else**
  - 11:     exploration = FALSE
  - 12:   **end if**
  - 13: **end while**
  - 14: Use remaining budget  $B - (c_0 + c_{[n]})m$  to generate  $\hat{S}^*$  exploitation estimator  $\tilde{F}_{\hat{S}^*}$ .
- 

variates estimator  $\hat{x}_c$  for  $\mathbb{E}[X]$ :

$$\hat{x}_c = \frac{1}{m} \sum_{\ell \in [m]} X_\ell - \frac{1}{m} \sum_{\ell \in [m]} (Z_\ell^\top \beta - \mathbb{E}[Z]^\top \beta),$$

where  $\beta \in \mathbb{R}^d$  is some appropriately chosen vector and  $\mathbb{E}[Z]$  is assumed known. The estimator  $\hat{x}_c$  is also unbiased and has variance

$$\sigma_c^2 = \text{Var}[\hat{x}_c] = \frac{\text{Var}[X - Z^\top \beta]}{m},$$

which is minimized when  $\beta$  is the least squares coefficient for the centered linear regression (i.e.  $X - \mathbb{E}[X] \sim Z - \mathbb{E}[Z]$ ):

$$\beta = \text{Cov}[Z]^{-1} \text{Cov}[Z, X], \quad (2.1)$$

and the minimal variance value is,

$$\sigma_*^2 = \min_{\beta \in \mathbb{R}^d} \sigma_c^2 = \frac{(1 - \rho^2) \sigma_X^2}{m} \quad \rho = \text{Corr}(X, Z^\top \text{Cov}[Z]^{-1} \text{Cov}[Z, X]).$$

Hence, when  $|\rho| \approx 1$ , the variance reduction is significant, in which case  $Z^\top \beta$  accounts for most of the variance of  $X$ .

In our multifidelity setup, we will use  $X_S$  as our control variate  $Z$ , and so in practice,  $\mathbb{E}[Z]$  may be unknown. In this case, one may consider the following *approximate control variates* estimator that uses an independent size- $N$  MC estimator in place of  $\mathbb{E}[Z]$  using samples  $\tilde{Z}_j$ :

$$\hat{x}_{ac} = \frac{1}{m} \sum_{\ell \in [m]} X_\ell - \frac{1}{m} \sum_{\ell \in [m]} \left( Z_\ell^\top \beta - \frac{1}{N} \sum_{j \in [N]} \tilde{Z}_j^\top \beta \right) \quad Z_\ell \perp \tilde{Z}_j, (\ell, j) \in [m] \times [N], \quad (2.2)$$

and this has total variance

$$\sigma_*^2 = \sigma_c^2 + \frac{\text{Var}[Z^\top \beta]}{N} = \frac{(1 - \rho^2) \sigma_X^2}{m} + \frac{\rho^2 \sigma_X^2}{N}. \quad (2.3)$$

Construction of such approximate control variates estimators has been recently studied in the multifidelity estimation of first-order statistics [14, 29, 13].

**Remark 2.1.** When  $\text{Cov}[Z]^{-1}$  and  $\text{Cov}[Z, X]$  are unknown, we can use empirical estimators instead, which will result in higher order trajectory-wise statistical errors in  $m$  (which are different from the mean-squared error that we subsequently consider); see [11, 21] for a detailed discussion.

### 3 Variance reduction for CDF estimation

Control variates are more general than described in Section 2, and can be applied to CDF estimation of nonlinear functions of random variables [12, 16]. For example, in risk management applications [12], the authors considered using the *delta-gamma* approximation<sup>1</sup> (i.e. the second-order Taylor expansion) of a loss function  $L$  at a given position  $\mathbf{x}$  along random market move direction  $\boldsymbol{\eta}$  as control variates to compute the value-at-risk (VaR) at a given confidence level. More precisely, note that

$$-(L(\mathbf{x} + \boldsymbol{\eta}) - L(\mathbf{x})) =: \ell(\mathbf{x}) \approx \widehat{\ell}(\mathbf{x}) := -\nabla L(\mathbf{x})^\top \boldsymbol{\eta} - \frac{1}{2} \boldsymbol{\eta}^\top \nabla^2 L(\mathbf{x}) \boldsymbol{\eta}.$$

Fixing  $C$ ,  $\mathbf{1}_{\widehat{\ell}(\mathbf{x}) \leq C}$  can be used as a control variate for  $\mathbf{1}_{\ell(\mathbf{x}) \leq C}$  to compute the latter's expectation, which in particular provides a way to compute CDF's (since CDF values are expectations of indicator functions). More advanced approximation techniques have been introduced in [16] to construct other control variates in VaR computation.

We apply a similar idea in our multifidelity setup here, but an appropriate replacement for the delta-gamma approximation must be constructed, and our estimation procedure is subject to a total budget constraint. In our setup, a specific functional form may be computationally difficult to produce, and Taylor-like approximations can be inaccurate outside local regions. Our alternative strategy in this setting is to employ a global emulator for  $Y$  based on information of  $X_{[n]}$ , e.g. linear combinations of  $X_{[n]}$ , which can be effective when the correlation between these quantities is large – in practice for parametric PDE's, this situation is commonplace. Nonlinear approximations may also be considered based on appropriate relevant information between the high- and low-fidelity models, although we do not explore this possibility here.

In the rest of the section, we introduce a general multifidelity approach to estimating  $F_Y(x)$  subject to a budget constraint.

#### 3.1 Control variates for multifidelity CDF estimation

The following simple observation is crucial:

$$F_Y(\mathbf{x}) = \mathbb{E}[\mathbf{1}_{Y \leq \mathbf{x}}] \quad \mathbf{x} \in \mathbb{R}^d.$$

If we fix  $S \subseteq [n]$ , the control variate based on the random variables  $X_S$  that maximizes variance reduction (and hence is optimal) is  $\mathbb{E}[\mathbf{1}_{Y \leq \mathbf{x}} | X_S]$  [25]. Of course, this quantity requires the orthogonal projection of  $Y$  onto the sigma-field generated by  $X_S$ , which is computationally intractable without special assumptions (e.g. joint normality). Thus as a matter of practice, we will resort to approximations of  $\mathbb{E}[\mathbf{1}_{Y \leq \mathbf{x}} | X_S]$ . For the moment, we use  $h(X_S; \mathbf{x})$  to denote a general  $X_S$ -measurable function that will serve as the control variates for  $\mathbf{1}_{Y \leq \mathbf{x}}$ . We make a particular choice for  $h$  in Section 4.

Analogous to (2.2), we construct an approximate control variates estimator for  $F_Y(\mathbf{x})$ , where the  $m$  and  $N$  in (2.2) will be related by the budget constraint. Since different  $S$  will be considered simultaneously, we take a *uniform exploration policy* that first collects  $m$  i.i.d joint samples of the *full* models

$$\text{Exploration samples: } \{(X_{\text{ep},\ell,1}^\top, \dots, X_{\text{ep},\ell,n}^\top, Y_{\text{ep},\ell}^\top)^\top\}_{\ell \in [m]} \subset \mathbb{R}^{d+\sum_{i=1}^n d_i} \quad (3.1)$$

<sup>1</sup>Here we refer to the “full” delta-gamma approximation. The more commonly used delta-gamma approximation in practice does not consider the second-order cross terms.

for variance reduction and then commits the remainder of the budget to collect  $N_S$  samples of a selected model subset (say  $S$ ) of low-fidelity models

$$\text{Exploitation samples: } \{X_{\text{ept},j,S}\}_{j \in [N_S]} \quad (3.2)$$

to compute the mean of control variates, where  $m$  and  $N_S$  are related by the budget constraint:

$$N_S = \frac{B - c_{\text{epr}}m}{c_S} \quad c_{\text{epr}} = \sum_{i=0}^n c_i, \quad c_S = \sum_{i \in S} c_i, \quad (3.3)$$

where we ignore integer rounding effects to simplify the discussion. The control variates estimator for  $F_Y(\mathbf{x})$  based on  $h(X_S; \mathbf{x})$  is,

$$\widehat{F}_S(\mathbf{x}) = \frac{1}{m} \sum_{\ell \in [m]} \mathbf{1}_{Y_{\text{epr},\ell} \leq \mathbf{x}} - \frac{1}{m} \sum_{\ell \in [m]} \left( \alpha(\mathbf{x}) h(X_{\text{epr},\ell,S}; \mathbf{x}) - \frac{1}{N_S} \sum_{j \in [N_S]} \alpha(\mathbf{x}) h(X_{\text{ept},j,S}; \mathbf{x}) \right), \quad (3.4)$$

where  $\alpha(\mathbf{x})$  is the optimal scaling coefficient as (2.1):

$$\alpha(\mathbf{x}) = \text{Cov}[h(X_S; \mathbf{x})]^{-1} \text{Cov}[\mathbf{1}_{Y \leq \mathbf{x}}, h(X_S; \mathbf{x})]. \quad (3.5)$$

Note  $\alpha(\mathbf{x})$  is undefined if  $\text{Cov}[h(X_S; \mathbf{x})] = 0$ . In this case, the value of the estimator  $\widehat{F}_S(\mathbf{x})$  does not depend on  $\alpha(\mathbf{x})$ , and we set  $\alpha(\mathbf{x})$  to 0 for convenience. The quantity  $\widehat{F}_S(\mathbf{x})$  is an unbiased estimator for  $F_Y(\mathbf{x})$ , with variance

$$\text{Var}[\widehat{F}_S(\mathbf{x})] = \frac{(1 - \rho_S^2(\mathbf{x}))F_Y(\mathbf{x})(1 - F_Y(\mathbf{x}))}{m} + \frac{\rho_S^2(\mathbf{x})F_Y(\mathbf{x})(1 - F_Y(\mathbf{x}))}{N_S},$$

where

$$\rho_S(\mathbf{x}) = \text{Corr}[\mathbf{1}_{Y \leq \mathbf{x}}, h(X_S; \mathbf{x})]. \quad (3.6)$$

### 3.2 A control variates loss function

To measure the overall accuracy of  $\widehat{F}_S(\mathbf{x})$ , we introduce the loss  $L_S$  defined by the average  $\omega(\mathbf{x})$ -weighted  $L^2$ -norm square of  $\widehat{F}_S(\mathbf{x}) - F_Y(\mathbf{x})$ :

$$L_S := \mathbb{E} \left[ \int_{\mathbb{R}^d} \omega(\mathbf{x}) |\widehat{F}_S(\mathbf{x}) - F_Y(\mathbf{x})|^2 d\mathbf{x} \right], \quad (3.7)$$

where  $\omega(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$  is a weight function. It is worth noting that the  $w(\mathbf{x})$ -weighted  $L^2$ -norm square is related to other more widely used metrics on distributions. For example, it reduces to the Cramér-von Mises distance if taking  $\omega(\mathbf{x})d\mathbf{x} = dF_Y(\mathbf{x})$ .

By Tonelli's theorem,

$$L_S = \int_{\mathbb{R}^d} \omega(\mathbf{x}) \text{Var}[\widehat{F}_S(\mathbf{x})] d\mathbf{x} = \frac{k_1(S)}{m} + \frac{k_2(S)}{B - c_{\text{epr}}m}, \quad (3.8)$$

where

$$\begin{aligned} k_1(S) &= \int_{\mathbb{R}^d} \omega(\mathbf{x}) (1 - \rho_S^2(\mathbf{x})) F_Y(\mathbf{x}) (1 - F_Y(\mathbf{x})) d\mathbf{x} \\ k_2(S) &= c_S \int_{\mathbb{R}^d} \omega(\mathbf{x}) \rho_S^2(\mathbf{x}) F_Y(\mathbf{x}) (1 - F_Y(\mathbf{x})) d\mathbf{x}. \end{aligned} \quad (3.9)$$

Since  $k_1(S)$  and  $k_2(S)$  are nonnegative, a sufficient and necessary condition for  $k_1(S)$  and  $k_2(S)$  being well-defined (i.e. finite) is

$$k_1(S) + c_S^{-1}k_2(S) = \int_{\mathbb{R}^d} \omega(\mathbf{x})F_Y(\mathbf{x})(1 - F_Y(\mathbf{x}))d\mathbf{x} < \infty, \quad (3.10)$$

which may not be true for arbitrary choice of  $\omega$ . For instance, when  $\omega(\mathbf{x}) \equiv 1$ , (3.10) is true when  $d = 1$  if  $\mathbb{E}[|Y|^{1+\delta}] < \infty$  for some  $\delta > 0$ . However, when  $d \geq 2$ , (3.10) is generally not true when the support for the distribution of  $Y$  is unbounded since  $F_Y^{-1}([\varepsilon, 1 - \varepsilon])$  may have infinite Lebesgue measure in  $\mathbb{R}^d$  for some  $\varepsilon > 0$ . For such scenarios, requiring that  $\omega(\mathbf{x})$  is integrable ensures (3.10), i.e.,

$$\int_{\mathbb{R}^d} \omega(\mathbf{x})F_Y(\mathbf{x})(1 - F_Y(\mathbf{x}))d\mathbf{x} < \int_{\mathbb{R}^d} \omega(\mathbf{x})d\mathbf{x} < \infty. \quad (3.11)$$

Some typical choices for integrable  $\omega(\mathbf{x})$  include  $\omega(\mathbf{x}) = \mathbf{1}_T$  where  $T \subset \mathbb{R}^d$  is a bounded domain of interest or  $\omega(\mathbf{x})$  with reasonably fast decaying tails as  $\|\mathbf{x}\|_2 \rightarrow \infty$ . In the following discussion, we assume (3.10) holds true. We make different choices for  $\omega$  in our numerical results of Section 6.

### 3.3 Exploration-exploitation trade-off

Equation (3.8) is a similar exploration-exploitation loss trade-off that was originally formulated in [29], where  $k_1$  and  $k_2$  measure the errors committed by the exploration and the exploitation, respectively.  $L_S$  is a strictly convex function for valid  $m$ , i.e., for  $0 < m < B/c_{\text{epr}}$ , and achieves its unique minimum at,

$$m_S^* = \frac{B}{c_{\text{epr}} + \sqrt{\frac{c_{\text{epr}}k_2(S)}{k_1(S)}}}, \quad (3.12)$$

with the corresponding minimum loss

$$L_S^* := \min_{0 < m < \frac{B - c_{\text{epr}}m}{c_S}} L_S(m_S^*) = \frac{(\sqrt{c_{\text{epr}}k_1(S)} + \sqrt{k_2(S)})^2}{B} =: \frac{\gamma_S}{B} \quad (3.13)$$

This allow us to identify an optimal subset  $S$  as one that minimizes the optimal loss,

$$S^* = \arg \min_{S \subseteq [n]} \gamma_S. \quad (3.14)$$

Our uniform exploration policy can be called *optimal* if it collects  $m_{S^*}^*$  joint samples for exploration and uses model  $S^*$  for exploitation. This is, in effect, a model selection procedure, as an optimal exploration policy selects the model (subset) that yields the smallest error via optimally balancing the trade-off between exploration and exploitation. In the following discussion, we assume  $S^*$  is unique.

With oracle information, the minimum error achieved by a uniform exploration policy is given by  $L_{S^*}^* = \gamma_{S^*}/B$ . On the other hand, the error achieved by an empirical CDF (ECDF) estimator for  $F_Y$  with the same budget is  $c_{\text{epr}} \int_{\mathbb{R}^d} \omega(\mathbf{x})F_Y(\mathbf{x})(1 - F_Y(\mathbf{x}))d\mathbf{x}/B$ . This ECDF procedure amounts to devoting the full budget to sampling the high-fidelity model  $Y$ , and ignoring the lower fidelity models. Taking the quotient of the two errors yields a lower bound for the relative efficiency of a uniform exploration policy compared to the ECDF estimator:

$$\begin{aligned} \frac{c_{\text{epr}} \int_{\mathbb{R}^d} \omega(\mathbf{x})F_Y(\mathbf{x})(1 - F_Y(\mathbf{x}))d\mathbf{x}/B}{\gamma_{S^*}/B} &= \frac{c_{\text{epr}} \int_{\mathbb{R}^d} \omega(\mathbf{x})F_Y(\mathbf{x})(1 - F_Y(\mathbf{x}))d\mathbf{x}}{\left(\sqrt{c_{\text{epr}}k_1(S^*)} + \sqrt{k_2(S^*)}\right)^2} \\ &\geq \frac{c_{\text{epr}} \int_{\mathbb{R}^d} \omega(\mathbf{x})F_Y(\mathbf{x})(1 - F_Y(\mathbf{x}))d\mathbf{x}}{2(c_{\text{epr}}k_1(S^*) + k_2(S^*))} \\ &\geq \frac{1}{2\left(\frac{c_S}{c_{\text{epr}}} + \mathbb{E}_{\mathbf{x}}[(1 - \rho_{S^*}^2(\mathbf{x}))]\right)} \geq \frac{1}{4}, \end{aligned}$$



where the expectation  $\mathbb{E}_{\mathbf{x}}[\cdot]$  is taken with respect to

$$\mathbf{x} \sim \frac{\omega(\mathbf{x})F_Y(\mathbf{x})(1 - F_Y(\mathbf{x}))d\mathbf{x}}{\int_{\mathbb{R}^d} \omega(\mathbf{z})F_Y(\mathbf{z})(1 - F_Y(\mathbf{z}))d\mathbf{z}}.$$

Hence, the relative efficiency of a uniform exploration policy compared to the ECDF estimator is unconditionally bounded below by  $1/4$ , and hence the uniform exploration policy can at worst realize a loss value of 4 times a naive ECDF procedure. On the other hand, the relative efficiency is  $\gg 1$  if both  $c_S/c_{\text{epr}}$  and  $\mathbb{E}_{\mathbf{x}}[(1 - \rho_{S^*}^2(\mathbf{x}))]$  are small. This happens, for instance, if  $X_{S^*}$  has a much smaller sampling cost than  $Y$  and  $h(X_S; \mathbf{x})$  are “good” control variates for  $\mathbf{1}_{Y \leq x}$  uniformly for all  $\mathbf{x} \in \mathbb{R}^d$ , both of which are realistic occurrences in multifidelity applications.

## 4 Choosing control variates from linear approximations

We now discuss our procedure for selecting the control variate  $h$ , which boils down to constructing approximations of  $\mathbb{E}[\mathbf{1}_{Y \leq x} | X_S]$  that both retain high correlation with  $Y$  and are budget-friendly. While one may generate special forms for approximations in particular cases, our goal is a simple and generic choice that is useful for many practical applications.

Recall that  $X_{S^+} = (1, X_i^\top)_{i \in S}^\top \in \mathbb{R}^{d_S+1}$ . For  $i \in [d]$ , let  $\beta_{S^+}^{(i)}$  be the optimal linear projection coefficients for estimating the  $i$ -th component of  $Y$  using  $X_{S^+}$ :

$$\beta_{S^+}^{(i)} = \text{Cov}[X_{S^+}]^{-1} \text{Cov}[X_{S^+}, Y^{(i)}] \in \mathbb{R}^{d_S+1}. \quad (4.1)$$

Letting

$$\mathbf{B}_{S^+} = [\beta_{S^+}^{(1)}, \dots, \beta_{S^+}^{(d)}] \in \mathbb{R}^{(d_S+1) \times d}, \quad (4.2)$$

the least squares approximation of  $Y$  using linear combinations of  $X_S$  and 1 is given by

$$H_S(X_S) = (X_{S^+}^\top \mathbf{B}_{S^+})^\top = \begin{bmatrix} X_{S^+}^\top \beta_{S^+}^{(1)} \\ \vdots \\ X_{S^+}^\top \beta_{S^+}^{(d)} \end{bmatrix} \in \mathbb{R}^d.$$

When all quantities are scalars, i.e.,  $d = d_1 = \dots = d_n = 1$ , one can directly manipulate  $H_S$  to estimate the statistics of  $Y$  [29, 30]. Such an approach is easy to implement and enjoys certain robustness for first-order statistics [13], but is more prone to model misspecification effects (e.g. expressibility of the linear model, noise assumption, etc.) when the whole distribution of  $Y$  is to be learned due to the limitation of linear approximation [30].

To address the issue, we take an additional nonlinear step beyond  $H_S$ . In particular, we consider the following family of control variates that slice the estimator  $H_S$ :

$$h(X_S; \mathbf{x}) = \mathbf{1}_{H_S \leq x}. \quad (4.3)$$

Intuitively, we may expect  $\mathbf{1}_{Y \leq x}$  and  $h(X_S; \mathbf{x})$  to be correlated if  $\mathbb{E}[\|Y - H_S\|_2^2]$  is small. However, this may not be true for  $\mathbf{x}$  approaching the tails of  $Y$ . For instance, assuming  $d = 1$  and a standard joint Gaussian random vector  $(X, Y)$  with correlation  $\rho$ , it follows from straightforward computation that

$$\lim_{x \rightarrow -\infty} \text{Corr}(\mathbf{1}_{Y \leq x}, \mathbf{1}_{X \leq x}) = \lim_{x \rightarrow 0} \frac{C_{X,Y}(x, x)}{x} = \begin{cases} 1, & |\rho| = 1, \\ 0, & |\rho| < 1 \end{cases}$$

where  $C_{X,Y}(x, y) = \mathbb{P}(\Phi^{-1}(X) \leq x, \Phi^{-1}(Y) \leq y)$  is the Gaussian copula and  $\Phi^{-1}$  is the quantile of a standard normal distribution; see [20]. Hence,  $\mathbf{1}_{X \leq x}$  is not a good control variate for  $\mathbf{1}_{Y \leq x}$  unless  $|\rho| = 1$ , i.e., only if  $X \propto Y$ . Nevertheless, our experiments in Section 6 show that in practice  $h(X_S; \mathbf{x})$

provides a reasonable control variates choice for many scenarios in multifidelity simulations, and thus suggests that situations described above are less common for the applications of our interest.

Choosing  $h$  as in (4.3), the coefficient  $\alpha(\mathbf{x})$  in (3.5) can be explicitly computed as

$$\alpha(\mathbf{x}) = \begin{cases} \frac{F_{Y \vee H_S}(\mathbf{x}) - F_Y(\mathbf{x})F_{H_S}(\mathbf{x})}{F_{H_S}(\mathbf{x})(1 - F_{H_S}(\mathbf{x}))} & \mathbf{x} \in \text{supp}(F_{H_S}(\mathbf{x}))^\circ \\ 0 & \text{otherwise} \end{cases} \quad (4.4)$$

One useful technical result is that  $\alpha(\mathbf{x})$  is bounded.

**Lemma 4.1.** *Let  $\alpha(\mathbf{x})$  be given as in (4.4). Then,  $|\alpha(\mathbf{x})| \leq 1$ .*

*Proof.* It suffices to check that for  $\mathbf{x} \in \text{supp}(F_{H_S}(\mathbf{x}))^\circ$ ,  $\alpha(\mathbf{x}) \leq 1$  and  $-\alpha(\mathbf{x}) \leq 1$  hold simultaneously:

$$\begin{aligned} \frac{F_{Y \vee H_S}(\mathbf{x}) - F_Y(\mathbf{x})F_{H_S}(\mathbf{x})}{F_{H_S}(\mathbf{x})(1 - F_{H_S}(\mathbf{x}))} &\leq \frac{F_Y(\mathbf{x}) \wedge F_{H_S}(\mathbf{x}) - F_Y(\mathbf{x})F_{H_S}(\mathbf{x})}{F_{H_S}(\mathbf{x})(1 - F_{H_S}(\mathbf{x}))} \\ &= \frac{F_Y(\mathbf{x})}{F_{H_S}(\mathbf{x})} \wedge \frac{1 - F_Y(\mathbf{x})}{1 - F_{H_S}(\mathbf{x})} \leq 1, \end{aligned}$$

and

$$\begin{aligned} \frac{F_Y(\mathbf{x})F_{H_S}(\mathbf{x}) - F_{Y \vee H_S}(\mathbf{x})}{F_{H_S}(\mathbf{x})(1 - F_{H_S}(\mathbf{x}))} &\leq \frac{F_Y(\mathbf{x})F_{H_S}(\mathbf{x}) - F_Y(\mathbf{x}) + 1 - F_{H_S}(\mathbf{x})}{F_{H_S}(\mathbf{x})(1 - F_{H_S}(\mathbf{x}))} \wedge \frac{F_Y(\mathbf{x})F_{H_S}(\mathbf{x})}{F_{H_S}(\mathbf{x})(1 - F_{H_S}(\mathbf{x}))} \\ &\leq \frac{1 - F_Y(\mathbf{x})}{F_{H_S}(\mathbf{x})} \wedge \frac{F_Y(\mathbf{x})}{1 - F_{H_S}(\mathbf{x})} \leq 1. \end{aligned}$$

□

We will discuss computational aspects of using (4.3) as control variates in Section 5.1.

**Remark 4.2.** As an alternative to our choice (4.3), one may use generalized linear models (GLMs) to approximate  $\mathbb{E}[\mathbf{1}_{Y \leq \mathbf{x}} | X_S]$ . With  $d = 1$  and for binary outputs, a frequently used GLM is logistic regression:

$$h(X_S; x) = \frac{1}{1 + \exp(X_{S+}^\top \boldsymbol{\theta}_{S+})} \quad x \in \mathbb{R},$$

where  $\boldsymbol{\theta}_{S+}$  is some parameter that may depend on  $x$  and can be learned in exploration. Theoretically, different  $\boldsymbol{\theta}_{S+}$  can be assigned to different  $x$ , resulting in formally infinitely many parameters necessitating estimation, which is challenging from finite exploration data with an appropriate parametrization. Alternatively, one may first partition  $\mathbb{R}$  into intervals

$$\mathbb{R} = \bigcup_{i \in [k]} [y_{i-1}, y_i) \quad -\infty = y_0 < y_1 < \dots < y_{k-1} < y_k = +\infty,$$

and for every  $x \in [y_{i-1}, y_i)$ , one uses an appropriately chosen  $\boldsymbol{\theta}_{S+}$  associated with  $y_{i-1}$ . Of course, the question of how to choose  $y_i$  requires attention, and good choices are generally problem-dependent. For simplicity, we do not pursue this approach in this paper.

## 5 Algorithms

It is useful at this point to revisit the cvMDL algorithm in Algorithm 1: the loss function  $L_S$  in (3.8) is the desired loss function to optimize over but requires oracle statistics. Thus, we replace it with an approximation  $\widehat{L}_S$  that we describe in this section. Additionally, the computations in lines 3 through 6 are now more transparent: The oracle computations are given by (3.14) and (3.12). In a practical algorithmic setting, we replace these with approximate computations, which is the topic of this section.

When using approximate quantities to compute  $L_S$ , the explicit exploration-exploitation loss decomposition in (3.8) may no longer be true. Nevertheless, if the quantities we estimate are sufficiently accurate, then such a decomposition is expected to be approximately valid. Thus, in devising practical algorithms, we will use the oracle loss form (3.8) (with estimated coefficients) instead of (3.7) as the criteria for model selection. We will present in the numerical section some empirical evidence that such a replacement has little impact on model selection.

Since our estimators change when new exploration samples are collected, the dependence on this number of exploration samples must be made explicit. For  $S \subseteq [n]$ , we let  $\widehat{L}_S(m; t)$  denote an estimator for the loss function  $L_S$  after having collected  $t$  exploration samples. We then let  $\widehat{m}_S^*$  be the corresponding estimator for the optimal exploration sample size  $m_S^*$ . Summarizing this: the intuition behind the cvMDL algorithm is that we use currently collected exploration data ( $t$  samples) to find the estimated optimal model ( $\widehat{S}^*$ ) and the corresponding exploration rate ( $\widehat{m}_{\widehat{S}^*}^*$ ). Based on the value of  $\widehat{m}_{\widehat{S}^*}^*$  relative to  $t$ , we decide whether to continue to explore or to switch to exploitation. Although cvMDL is similar to the AETC and AETC-d algorithms in [29, 30, 13], the process of constructing both the model selection parameters,  $k_1(S)$  and  $k_2(S)$ , as well as the exploitation estimator, is distinct. These distinctions, and in particular their theoretical and practical consequences, will be further explained in the following sections.

## 5.1 Estimators for oracle quantities

In this section, we discuss how to estimate  $L_S$ ,  $m_S^*$ , and  $\alpha(\mathbf{x})$  from exploration data when instantiating cvMDL using the linear approximators as introduced in Section 4. The control variates  $h(X_S; \mathbf{x}) = \mathbf{1}_{H_S \leq \mathbf{x}}$  belong to a parametric family characterized by  $\beta_{S^+}^{(i)}$ ,  $i \in [d]$  from (4.1), which can be estimated from exploration data.

Recall from (3.1) in Section 3.1 that the  $\ell$ -th exploration sample of all low-fidelity models in  $S$  is denoted by  $X_{\text{opr}, \ell, S}$ . Similarly, we define  $X_{\text{opr}, \ell, S^+} = (1, X_{\text{opr}, \ell, S}^\top)^\top$ . To estimate  $\beta_{S^+}^{(i)}$ , we use the least-squares estimator:

$$\widehat{\beta}_{S^+}^{(i)} = \mathbf{Z}_S^\dagger \mathbf{Y}^{(i)} \quad \mathbf{Z}_S = \begin{bmatrix} X_{\text{opr}, 1, S^+}^\top \\ \vdots \\ X_{\text{opr}, m, S^+}^\top \end{bmatrix} \in \mathbb{R}^{m \times (d_S + 1)} \quad \mathbf{Y}^{(i)} = \begin{bmatrix} Y_{\text{opr}, 1}^{(i)} \\ \vdots \\ Y_{\text{opr}, m}^{(i)} \end{bmatrix} \in \mathbb{R}^m, \quad (5.1)$$

where  $(X_{\text{opr}, \ell, 1}^\top, \dots, Y_{\text{opr}, \ell}^\top)_{\ell \in [m]}^\top$  are joint exploration samples, and the design matrix  $\mathbf{Z}_S^\top \mathbf{Z}_S$  is assumed to have full column rank<sup>2</sup>.

For  $\mathbf{x} \in \mathbb{R}^d$ ,  $h(X_S; \mathbf{x})$  can in turn be estimated as

$$\widehat{h}(X_S; \mathbf{x}) = \mathbf{1}_{\widehat{H}_S(X_S) \leq \mathbf{x}} \quad \widehat{H}_S(X_S) = \widehat{\mathbf{B}}_{S^+}^\top X_{S^+} = \begin{bmatrix} X_{S^+}^\top \widehat{\beta}_{S^+}^{(1)} \\ \vdots \\ X_{S^+}^\top \widehat{\beta}_{S^+}^{(d)} \end{bmatrix}, \quad (5.2)$$

where  $X_S$  is a general notation for the low-fidelity model output, and

$$\widehat{\mathbf{B}}_{S^+} = [\widehat{\beta}_{S^+}^{(1)}, \dots, \widehat{\beta}_{S^+}^{(d)}]. \quad (5.3)$$

For ease of notation, we write  $\widehat{H}_S(X_S)$  and  $\widehat{h}_S(X_S; \mathbf{x})$  as  $\widehat{H}_S$  and  $\widehat{h}_S$  when  $X_{S^+}$  is generic and has nothing to do with exploration and exploitation data. We introduce some additional notation for

<sup>2</sup>This explains the minimal exploration size condition in Algorithm 1, which is a necessary condition for full rank here.

quantities involving both estimated coefficients and empirical CDFs using exploration data:

$$\begin{aligned}\widehat{F}_Y(\mathbf{x}) &= \frac{1}{m} \sum_{\ell \in [m]} \mathbf{1}_{Y_{\text{epr},\ell} \leq \mathbf{x}} \\ \widehat{F}_{\widehat{H}_S}(\mathbf{x}) &= \frac{1}{m} \sum_{\ell \in [m]} \widehat{h}(X_{\text{epr},\ell,S}; \mathbf{x}) = \frac{1}{m} \sum_{\ell \in [m]} \mathbf{1}_{\widehat{H}_S(X_{\text{epr},\ell,S}) \leq \mathbf{x}} \\ \widehat{F}_{Y \vee \widehat{H}_S}(\mathbf{x}) &= \frac{1}{m} \sum_{\ell \in [m]} \mathbf{1}_{Y_{\text{epr},\ell} \vee \widehat{H}_S(X_{\text{epr},\ell,S}) \leq \mathbf{x}}.\end{aligned}$$

To compute the loss function approximation, we build approximations to  $k_1$  and  $k_2$  in (3.9), which in turn requires us to compute  $\rho_S^2$  in (3.6). For this purpose, observe that

$$(1 - \rho_S^2(\mathbf{x}))F_Y(\mathbf{x})(1 - F_Y(\mathbf{x})) = \mathbb{E}[(\mathbf{1}_{Y \leq \mathbf{x}} - F_Y(\mathbf{x})) - \alpha(\mathbf{x})(\mathbf{1}_{H_S \leq \mathbf{x}} - F_{H_S}(\mathbf{x}))]^2,$$

where  $\alpha(\mathbf{x})$  is defined in (4.4). The quantity in the expectation is the mean squared regression residual between two (centered) Bernoulli random variables  $\mathbf{1}_{Y \leq \mathbf{x}}$  and  $\mathbf{1}_{H_S \leq \mathbf{x}}$ . Thus, a natural estimator for  $(1 - \rho_S^2(\mathbf{x}))F_Y(\mathbf{x})(1 - F_Y(\mathbf{x}))$  is to compute an empirical mean-squared difference between  $\mathbf{1}_{Y \leq \mathbf{x}}$  and a regressor with covariates  $\mathbf{1}_{\widehat{H}_S \leq \mathbf{x}}$ , which requires data for  $Y$ . Since we have (uncentered) data for  $Y$  on the exploration samples  $Y_{\text{epr},j}$ ,  $j \in [m]$ , then we need only evaluate a regressor for  $Y$  with covariates  $\widehat{H}_S$  and 1 on the exploration data sites. This results in the following estimator  $\mathcal{K}_1$  for  $(1 - \rho_S^2(\mathbf{x}))F_Y(\mathbf{x})(1 - F_Y(\mathbf{x}))$

$$\mathcal{K}_1(\mathbf{x}) = \frac{1}{m} \sum_{\ell \in [m]} (\mathbf{1}_{Y_{\text{epr},\ell} \leq \mathbf{x}} - r_\ell(\mathbf{x}))^2 \quad \begin{bmatrix} r_1(\mathbf{x}) \\ \vdots \\ r_m(\mathbf{x}) \end{bmatrix} = \mathbf{W}_S \mathbf{W}_S^\dagger \begin{bmatrix} \mathbf{1}_{Y_{\text{epr},1} \leq \mathbf{x}} \\ \vdots \\ \mathbf{1}_{Y_{\text{epr},m} \leq \mathbf{x}} \end{bmatrix}, \quad (5.4)$$

where

$$\mathbf{W}_S = \begin{bmatrix} 1 & \widehat{h}_S(X_{\text{epr},1,S}; \mathbf{x}) \\ \vdots & \vdots \\ 1 & \widehat{h}_S(X_{\text{epr},m,S}; \mathbf{x}) \end{bmatrix}.$$

This in turn allows us to estimate  $\rho_S^2(\mathbf{x})F_Y(\mathbf{x})(1 - F_Y(\mathbf{x}))$  as

$$\mathcal{K}_2(\mathbf{x}) = \widehat{F}_Y(\mathbf{x})(1 - \widehat{F}_Y(\mathbf{x})) - \mathcal{K}_1(\mathbf{x}). \quad (5.5)$$

Consequently, we can estimate  $k_1(S)$  and  $k_2(S)$  as

$$\widehat{k}_1(S) = \int_{\mathbb{R}^d} \omega(\mathbf{x}) \mathcal{K}_1(\mathbf{x}) d\mathbf{x} \quad \widehat{k}_2(S) = c_S \int_{\mathbb{R}^d} \omega(\mathbf{x}) \mathcal{K}_2(\mathbf{x}) d\mathbf{x}. \quad (5.6)$$

The above estimators for  $k_1(S)$  and  $k_2(S)$  are positive and actually coincide with empirical estimators for these quantities whenever defined (see Section B.3), which is a crucial realization for our consistency results later. Plugging the above estimates into (3.8) and (3.12) yields estimates for  $L_S$  and  $m_S^*$ :

$$\widehat{L}_S(z; m) = \frac{\widehat{k}_1(S)}{z} + \frac{\widehat{k}_2(S)}{B - c_{\text{epr}}z} \quad \widehat{m}_S^* = \frac{B}{c_{\text{epr}} + \sqrt{\frac{c_{\text{epr}} \widehat{k}_2(S)}{\widehat{k}_1(S)}}}. \quad (5.7)$$

Note  $\widehat{L}_S(z; m)$  has two parameters, where the second indicates the number of exploration samples used to compute  $\widehat{k}_1(S)$  and  $\widehat{k}_2(S)$ , and the first is the variable of  $\widehat{L}_S$ . We define  $\widehat{S}^*$  as the optimal model selected by this estimator,

$$\widehat{S}^* = \arg \min_{S \subseteq [n]} \widehat{L}_S(\widehat{m}_S^*; m), \quad (5.8)$$

which parallels the oracle choice (3.14). We have now described all quantities needed to complete the exploration phase of Algorithm 1, i.e., lines 1 - 13. What remains is to describe how the CDF estimator  $\tilde{F}_{\hat{S}^*}$  in line 14 is generated.

Our exploitation goal is to generate an estimator for (3.4), and so we also need to estimate  $\alpha(\mathbf{x})$ :

$$\hat{\alpha}(\mathbf{x}) = \frac{\hat{F}_{Y \vee \hat{H}_S}(\mathbf{x}) - \hat{F}_Y(\mathbf{x})\hat{F}_{\hat{H}_S}(\mathbf{x})}{\hat{F}_{\hat{H}_S}(\mathbf{x})(1 - \hat{F}_{\hat{H}_S}(\mathbf{x}))} \quad \mathbf{x} \in \text{supp}(\hat{F}_{\hat{H}_S}(\mathbf{x}))^\circ, \quad (5.9)$$

and 0 otherwise. By similar reasoning as in Lemma 4.1, one has

$$|\hat{\alpha}(\mathbf{x})| \leq 1. \quad (5.10)$$

Finally, the *exploitation* estimator  $\tilde{F}_S(\mathbf{x})$  for  $F_Y(\mathbf{x})$  based on estimated parameters, utilizes  $N_S$  exploitation samples (i.e., exhausts the remaining budget  $B$ ) and is given by,

$$\tilde{F}_S(\mathbf{x}) := \hat{F}_Y(\mathbf{x}) - \frac{1}{m} \sum_{\ell \in [m]} \left( \hat{\alpha}(\mathbf{x}) \hat{h}_S(X_{\text{epr}, \ell, S}; \mathbf{x}) - \frac{1}{N_S} \sum_{j \in [N_S]} \hat{\alpha}(\mathbf{x}) \hat{h}_S(X_{\text{ept}, \ell, S}; \mathbf{x}) \right), \quad (5.11)$$

where  $S = \hat{S}^*$  is the selected model based on  $\hat{k}_1(S)$  and  $\hat{k}_2(S)$ . By inspection, we observe that  $\tilde{F}_S(\mathbf{x})$  is a piecewise affine correction of  $\hat{F}_Y$ , where the correction is based on the control variates  $\hat{h}_S$ .

**Remark 5.1.** The estimator  $\hat{\alpha}(\mathbf{x})$  is undefined and manually set to 0 for  $\mathbf{x}$  outside the support of  $\hat{F}_{\hat{H}_S}$ , as in that case the denominator vanishes. To circumvent this, one can assign values of  $\hat{\alpha}(\mathbf{x})$  for  $\mathbf{x}$  outside the support of  $\hat{F}_{\hat{H}_S}$  to values slightly within the support. To be more precise, consider  $d = 1$ , and write  $x$  in place of  $\mathbf{x}$ . Assuming  $\alpha(x)$  is a continuous function of  $x$  in  $\text{supp}(F_{H_S})$  and  $\text{supp}(\hat{F}_{\hat{H}_S}(x)) = [x_{\min}, x_{\max}]$ , for  $x \in \text{supp}(\hat{F}_{\hat{H}_S}(x))^c$ , we may estimate  $\alpha(x)$  outside  $[x_{\min}, x_{\max}]$  as

$$\hat{\alpha}(x) = \begin{cases} \frac{\hat{F}_{Y \vee \hat{H}_S}(x(\tau)) - \hat{F}_Y(x(\tau))\tau}{\tau(1-\tau)} & x \leq x_{\min} \\ \frac{\hat{F}_{Y \vee \hat{H}_S}(x(1-\tau)) - \hat{F}_Y(x(1-\tau))(1-\tau)}{\tau(1-\tau)} & x \geq x_{\max}, \end{cases} \quad (5.12)$$

where  $x(\tau)$  and  $x(1-\tau)$  are the  $\tau$  and  $1-\tau$  quantiles of  $\hat{F}_{\hat{H}_S}$  for some small  $\tau \in (0, 1)$ :

$$x(\tau) = \hat{F}_{\hat{H}_S}^{-1}(\tau) \quad x(1-\tau) = \hat{F}_{\hat{H}_S}^{-1}(1-\tau).$$

This allows us to get nontrivial and more accurate estimates of  $F_Y$  outside  $[x_{\min}, x_{\max}]$  (i.e. in the tail regime). When  $d \geq 2$ , one may generalize the ideas above by projecting the points in the tail regime to some bounded set in  $\mathbb{R}^d$  that contains most of the measure in the domain, which we do not pursue further here.

This completes a comprehensive description of how all quantities in Algorithm 1 are computed.

## 5.2 Monotonicity of the exploitation CDF estimator

By construction,  $\tilde{F}_S(\mathbf{x})$  is a piecewise constant function on some  $d$ -dimensional rectangular partition of  $\mathbb{R}^d$ , but not necessarily a monotone nondecreasing function in each direction due to the fluctuations of estimators used in the construction. To address this issue, we introduce a dimension-wise recursive-sorting post-processing procedure on values in the range of  $\tilde{F}_S$  to recover the desired monotonicity. Represent  $\tilde{F}_S(\mathbf{x})$  as a  $d$ -dimensional tensor,  $\mathbf{T} \in \mathbb{R}^{\otimes_{i \in [d]} M_i}$ , i.e., with index set  $\mathbf{I} = \otimes_{i \in [d]} (z_{i,1}, \dots, z_{i, M_i})$ , where  $-\infty = z_{i,1} \leq \dots \leq z_{i, M_i} = +\infty$ , such that

$$\tilde{F}_S(\mathbf{x}) = \mathbf{T}_{z_{1,s_1}, \dots, z_{d,s_d}} \quad \mathbf{x} \in \prod_{i \in [d]} [z_{i,s_i}, z_{i,s_i+1}).$$

The desired monotonicity in each dimension can be recovered by alternately sorting values in  $\mathbf{T}$  in each dimension until converging. An example when  $d = 2$  is given below:

$$\begin{array}{ccc} \begin{bmatrix} 0.7 & 0.4 & 0 \\ 0.3 & 0.5 & 0.2 \\ 1 & 0.8 & 0.6 \end{bmatrix} & \xrightarrow{\text{sort rows}} & \begin{bmatrix} 0 & 0.4 & 0.7 \\ 0.2 & 0.3 & 0.5 \\ 0.6 & 0.8 & 1 \end{bmatrix} & \xrightarrow{\text{sort columns}} & \begin{bmatrix} 0 & 0.3 & 0.5 \\ 0.2 & 0.4 & 0.7 \\ 0.6 & 0.8 & 1 \end{bmatrix} \\ \begin{bmatrix} 0.7 & 0.4 & 0 \\ 0.3 & 0.5 & 0.2 \\ 1 & 0.8 & 0.6 \end{bmatrix} & \xrightarrow{\text{sort columns}} & \begin{bmatrix} 0.3 & 0.4 & 0 \\ 0.7 & 0.5 & 0.2 \\ 1 & 0.8 & 0.6 \end{bmatrix} & \xrightarrow{\text{sort rows}} & \begin{bmatrix} 0 & 0.3 & 0.4 \\ 0.2 & 0.5 & 0.7 \\ 0.6 & 0.8 & 1 \end{bmatrix} \end{array}$$

As shown above, sorting will finally end up in some stationary point with desired monotonicity (see Theorem 5.2), but different orders of sorting may lead to different sorted CDF representations when  $d \geq 2$ . However, in our case,  $\tilde{F}_S(\mathbf{x})$  is itself a perturbation of the CDF of  $Y$ , so the sorting procedure is often beneficial for stabilizing the algorithm. A more detailed empirical study on this will be given in Section 6. The sorting procedure described converges (i.e., achieves monotonicity in the values of  $\mathbf{T}$ ) in a finite number of iterations:

**Theorem 5.2.** *Assume that all the entries in  $\mathbf{T}$  are distinct. Fixing the orders of sorting, the alternating sorting algorithm described above will converge to a stationary point with desired monotonicity within a finite number of iterations.*

*Proof.* Sort the entries in increasing order:  $z^{(1)} < \dots < z^{(M)}$ , where  $M = M_1 \dots M_d$ . It is easy to see that at the beginning of the algorithm, the index of  $z^{(1)}$  is strictly decreasing in each direction. As a result,  $z^{(1)}$  will arrive at the entry of  $\mathbf{T}$  with index  $(1, \dots, 1)$  after finite steps of iteration, and after that, it remains unchanged in the subsequent iteration. In fact, for every  $s < M$ , assuming  $z^{(1)}, \dots, z^{(s)}$  have reached their final positions after which no change will occur, the index of  $z^{(s+1)}$  will be decreasing in each direction if the algorithm has not converged yet. The result follows by noting that  $M$  is finite, and a stationary point must possess desired monotonicity.  $\square$

### 5.3 Exploration sampling

One final technical detail we must describe is the precise action taken when we decide to continue exploring. I.e., we must define the function  $Q(m, \hat{m}_{\hat{S}^*}^*)$  in line 9 of Algorithm 1. When the current number  $m$  of exploration samples is smaller than the estimated optimal number of samples  $\hat{m}_{\hat{S}^*}^*$ , the function  $Q$  determines how to increase  $m$ .

A natural choice for  $Q$  is  $Q(m, \hat{m}_{\hat{S}^*}^*) = m + 1$ , i.e., simply increase by a single additional exploration sample. In practice, we observe that this behavior can be overly conservative and time-consuming when  $B$  is large. As an alternative, one could use a more aggressive strategy, say  $Q(m, \hat{m}_{\hat{S}^*}^*) = \frac{1}{2} \left( m + \hat{m}_{\hat{S}^*}^* \right)$ , which more quickly closes the gap between  $m$  and  $\hat{m}_{\hat{S}^*}^*$ . However, there are situations when this is too aggressive. For example, if  $m$  is very small (such as at initialization) then estimated quantities can be poor approximations, and in some cases  $\hat{m}_{\hat{S}^*}^*$  is significantly overestimated, and thus increasing  $m$  to  $\frac{1}{2} \left( m + \hat{m}_{\hat{S}^*}^* \right)$  can actually result in substantially overshooting the oracle value of  $m_{\hat{S}^*}^*$ . The probability of such an event is often positive and does not vanish as  $B$  increases.

As a compromise between these conservative and aggressive behaviors, we choose the following form:

$$Q(m, \hat{m}_{\hat{S}^*}^*) = \begin{cases} 2m, & m < \frac{\hat{m}_{\hat{S}^*}^*}{2} \\ \frac{1}{2} \left( m + \hat{m}_{\hat{S}^*}^* \right), & \frac{\hat{m}_{\hat{S}^*}^*}{2} \leq m < \hat{m}_{\hat{S}^*}^* \end{cases} \quad (5.13)$$

Since  $\hat{m}_{\hat{S}^*}^*$  is proportional to  $B$ , the above choice ensures that there is a sufficient amount of time for the algorithm to take exponential exploration whose growth manner is independent of the value of  $\hat{m}_{\hat{S}^*}^*$ , which ensures both efficiency and accuracy of the algorithm. We note that our particular choice of  $Q$  is not theoretically special, and serves only as a convenient practical choice that we have identified

through empirical testing. Our theoretical conclusions will assume that  $Q$  has the form above. However, since our theory is budget-asymptotic, reasonable alternative choices to the expression above would not change the theoretical conclusions.

We have completed all technical descriptions of Algorithm 1. A more fleshed-out pseudocode version is given in Section A that details every step that must be taken. Next, we establish that Algorithm 1 enjoys optimality guarantees relative to model selection and budget allocation strategies produced by an oracle.

## 5.4 Model consistency and optimality

We now provide theoretical guarantees for Algorithm 1. In summary, we will show that as the budget  $B$  tends to infinity, the model subset  $\widehat{S}^*$  chosen along with the number of exploration samples  $m$  taken by Algorithm 1, both converge to the oracle optimal model  $S^*$  and the optimal number of exploration samples  $m_{S^*}^*$ , respectively.

We need some technical assumptions in order to proceed with our results. Since we estimate quadratic moments, we require quadratic moments to exist, and we also require that there are no linear combinations of low-fidelity model outputs that are perfectly correlated with any other low-fidelity model output. These are codified in the following two assumptions.

**Assumption 5.3.** The models  $X_{[n]}$  and  $Y$  have bounded second moments:

$$\mathbb{E}[\|X_{[n]}\|_2^2 + \mathbb{E}\|Y\|_2^2] < \infty. \quad (5.14)$$

**Assumption 5.4.** The uncentered second moment matrix  $\mathbb{E}[X_{[n]+}X_{[n]+}^\top]$  is invertible, where  $X_{[n]+} = (1, X_{[n]}^\top)^\top$ .

Assumption 5.3 is the minimal moment condition on model outputs that we require to make oracle quantities well-defined. In practice, random variables that violate Assumption 5.4 exhibit perfect multicollinearity and are relatively rare. Assumption 5.4 being violated does not cause any conceptual breakdown of our procedure; the only consequence is that all the linear regression procedures suffer from a lack of identifiability of optimal covariates. While there are numerous standard procedures to remedy multicollinearity, such as covariate removal or regularization, violation of this assumption did not surface in our experiments, so we do not utilize any of these remedies.

The model selection procedure requires estimating the average  $\omega$ -weighted  $L^2$  norm. This requires us to make certain assumptions about  $\omega$ .

**Assumption 5.5.** The weight  $\omega(\mathbf{x})$  is chosen so that *either* of the following conditions is true:

- (a)  $\|\omega\|_{L^\infty(\mathbb{R}^d)} < \infty$  (e.g.  $\omega(\mathbf{x}) \equiv 1$ ) and  $d = 1$ ; or
- (b)  $\|\omega\|_{L^1(\mathbb{R}^d)} < \infty$ .

The final more technical assumption we require involves some regularity on distribution functions. In particular, we will show  $\mathbf{x}$ -pointwise convergence of the estimator  $\widehat{\alpha}(\mathbf{x})$  to the oracle parameter  $\alpha(\mathbf{x})$ , and to accomplish this we require bounds on the local variations of  $F_{H_S}$  and  $F_{H_S \vee Y}$  constructed in the model selection procedure. More technically, a sufficient assumption is a bounded local variations condition involving CDFs of certain  $d$ -dimensional sketches of  $X_{[n]}$  and  $Y$ .

**Assumption 5.6.** Define

$$V(\mathbf{A}) = (X_{S+}^\top \mathbf{A})^\top \in \mathbb{R}^d \quad \mathbf{A} = [\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(d)}] \in \mathbb{R}^{d_S \times d},$$

and recall the optimal coefficient matrix  $\mathbf{B}_{S+}$  in (4.2). We assume the CDFs of  $V(\mathbf{A})$  and  $V(\mathbf{A}) \vee Y$ , denoted by  $F_{V(\mathbf{A})}$  and  $F_{V(\mathbf{A}) \vee Y}$ , are globally Lipschitz near  $\mathbf{B}_{S+}$  for all  $S$ . In particular, there exists  $\varepsilon > 0$  such that

$$\max_{S \subseteq [n]} \sup_{\mathbf{A}: \|\mathbf{A} - \mathbf{B}_{S+}\|_F \leq \varepsilon} \{\|F_{V(\mathbf{A})}\|_{\text{Lip}} + \|F_{V(\mathbf{A}) \vee Y}\|_{\text{Lip}}\} = C < \infty,$$

where  $\|\cdot\|_{\text{Lip}}$  is the Lipschitz constant defined as

$$\|f\|_{\text{Lip}} = \sup_{\mathbf{x} \neq \mathbf{x}'} \frac{|f(\mathbf{x}) - f(\mathbf{x}')|}{\|\mathbf{x} - \mathbf{x}'\|_2} \quad f : \mathbb{R}^d \rightarrow \mathbb{R}.$$

This final assumption is less transparent than our previous ones but is true in every experiment we have run. For example, a (n unnecessarily strong) sufficient condition to ensure that Assumption 5.6 holds would be to assume that both  $Y$  and all linear combinations of  $X_{[n]}$  have bounded densities, and that every high-fidelity covariate  $Y^{(i)}$  is correlated with every low-fidelity covariate  $X_j^{(r)}$ , i.e.,  $\min_{i,j,r} |\text{Corr}(Y^{(i)}, X_j^{(r)})| > 0$ . Alternatively, one could only assume that the same bounded density condition, and the rather reasonable condition that the oracle regression coefficients  $\mathbf{B}_{S^+}$  select at least one non-deterministic covariate for every  $S$ .

We can now present our main results regarding applying the cvMDL algorithm with  $h(X_S; \mathbf{x})$  constructed using linear approximations, with the corresponding loss function parameters estimated from (5.6) and (5.7). In particular, we have that the adaptive exploration rate  $m(B)$  asymptotically matches the optimal (oracle) exploration rate  $m_{S^*}^*$  defined in Section 3.3, and the selected model  $S(B)$  converges to the optimal (oracle) model  $S^*$  as  $B \rightarrow \infty$ :

**Theorem 5.7** (Uniform consistency and asymptotic optimality of Algorithm 1). *Let  $h(X_S; \mathbf{x})$  be defined in (4.3), i.e., we use the linear approximation estimators from Section 4, and assume the model parameters are estimated via (5.6) and (5.7). Then consider Algorithm 1 with an input budget  $B$ , and let*

- $m(B) = \widehat{m}_{\widehat{S}^*}$  be the number of exploration samples chosen by Algorithm 1,
- $S(B) = \widehat{S}^*$  be the model selected for exploitation,
- $\widetilde{F}(\mathbf{x}; B) = \widetilde{F}_{\widehat{S}^*}(\mathbf{x})$  be the output CDF estimator for  $F_Y$ .

Under Assumptions 5.3, 5.4, 5.5, and 5.6, then with probability 1,

$$\lim_{B \rightarrow \infty} \frac{m(B)}{m_{S^*}^*} = 1, \quad (5.15a)$$

$$\lim_{B \rightarrow \infty} S(B) = S^*, \quad (5.15b)$$

$$\lim_{B \rightarrow \infty} \sup_{\mathbf{x} \in \mathbb{R}^d} |\widetilde{F}(\mathbf{x}; B) - F_Y(\mathbf{x})| = 0. \quad (5.15c)$$

where  $S^*$  and  $m_{S^*}^*$  are the unique optimal (oracle) model choice and exploration sample size defined in Section 3.3.

The proof is given in Section D. The result (5.15c) should not come as a surprise since uniform consistency is generally true for empirical CDF estimators. Therefore, while (5.15a) and (5.15b) show that Algorithm 1 exhibits optimality (relative to an oracle) for the choice of exploration samples and sample allocation across models, (5.15c) is not evidence that the multifidelity estimator  $\widetilde{F}(\mathbf{x}; B)$  is superior to the empirical CDF estimator that uses only the high-fidelity samples, although it confirms that  $\widetilde{F}(\mathbf{x}; B)$  behaves as expected. The major difference that distinguishes  $\widetilde{F}(\mathbf{x}; B)$  from a standard empirical CDF estimator is a constant term resulting from the mean  $\omega$ -weighted  $L^2$  convergence rate; see the discussion near the end of Section 3.3.

The statements in Theorem 5.7 and [30, Theorem 5.2] are similar, but in the former the requisite assumptions are *much* weaker and the guarantees are stronger. In fact, for [30, Theorem 5.2] to hold, one must assume that  $\mathbb{E}[Y|X_S]$  is a linear function of  $X_S$  and  $(Y - \mathbb{E}[Y|X_S]) \perp\!\!\!\perp X_S$  for all  $S \subseteq [n]$ , which can be challenging to verify in practice and in some practical cases does not hold. However, none of these assumptions is needed in Theorem 5.7. Additionally, Theorem 5.7 ensures convergence for a *multivariate* distribution function instead of the univariate convergence statements in [30, Theorem 5.2].



## 5.5 A brief view into proving Theorem 5.7

While we leave the technical parts of proving Theorem 5.7 to the appendix, we can summarize the crucial intermediate results that allow the proof to succeed. The major results we need revolve around consistency of various estimators as  $m$  and/or  $N_S$  approach infinity. The following two sets of results leverage the assumptions to conclude consistency of intermediate computations in the algorithm.

The first collection of results shows that the finite-sample estimators for quantities computed in the exploration phase are consistent as the number of exploration samples  $m$  tends to infinity.

**Lemma 5.8** (Asymptotic consistency of exploration estimators). *We have the following technical estimates and consistency results for all  $S \subseteq [n]$ :*

(i) Under Assumptions 5.3 and 5.6, then with probability 1,

$$\sup_{\|\mathbf{A} - \mathbf{B}_{S^+}\|_F < \varepsilon} \sup_{\mathbf{x} \in \mathbb{R}^d} |F_{V(\mathbf{A})}(\mathbf{x}) - F_{V(\mathbf{B}_{S^+})}(\mathbf{x})| \lesssim \|\mathbf{A} - \mathbf{B}_{S^+}\|_F^{2/3} \quad (5.16a)$$

$$\sup_{\|\mathbf{A} - \mathbf{B}_{S^+}\|_F < \varepsilon} \sup_{\mathbf{x} \in \mathbb{R}^d} |F_{V(\mathbf{A}) \vee Y}(\mathbf{x}) - F_{V(\mathbf{B}_{S^+}) \vee Y}(\mathbf{x})| \lesssim \|\mathbf{A} - \mathbf{B}_{S^+}\|_F^{2/3} \quad (5.16b)$$

where the implied constant in  $\lesssim$  is absolute.

(ii) Under Assumptions 5.3 and 5.4, then with probability 1,

$$\lim_{m \rightarrow \infty} \widehat{\mathbf{B}}_{S^+} = \mathbf{B}_{S^+} \quad (5.17)$$

(iii) Under Assumptions 5.3, 5.4, and 5.6, then with probability 1,

$$\lim_{m \rightarrow \infty} \sup_{\mathbf{x} \in \mathbb{R}^d} |\widehat{F}_{\widehat{H}_S}(\mathbf{x}) - F_{H_S}(\mathbf{x})| = 0 \quad \lim_{m \rightarrow \infty} \sup_{\mathbf{x} \in \mathbb{R}^d} |\widehat{F}_{Y \vee \widehat{H}_S}(\mathbf{x}) - F_{Y \vee H_S}(\mathbf{x})| = 0.$$

(iv) Under Assumptions 5.3, 5.4, and 5.6, then almost surely as  $m \rightarrow \infty$  we have that,

$$\mathcal{K}_1(\mathbf{x}) \rightarrow (1 - \rho_S^2(\mathbf{x}))F_Y(\mathbf{x})(1 - F_Y(\mathbf{x})) \quad \mathcal{K}_2(\mathbf{x}) \rightarrow \rho_S^2(\mathbf{x})F_Y(\mathbf{x})(1 - F_Y(\mathbf{x})) \quad (5.18)$$

for all  $\mathbf{x} \in \mathbb{R}^d$ .

(v) Under Assumptions 5.3, 5.4, 5.5, and 5.6, then with probability 1,  $\lim_{m \rightarrow \infty} \widehat{k}_1(S) = k_1(S)$  and  $\lim_{m \rightarrow \infty} \widehat{k}_2(S) = k_2(S)$ .

(vi) Under Assumptions 5.3, 5.4, and 5.6, for  $\mathbf{x} \in (\text{supp}(F_{H_S}))^\circ$ ,  $\widehat{\alpha}(\mathbf{x})$  is a consistent estimator of  $\alpha(\mathbf{x})$  almost surely, i.e.,  $\lim_{m \rightarrow \infty} \widehat{\alpha}(\mathbf{x}) = \alpha(\mathbf{x})$  for every  $\mathbf{x} \in (\text{supp}(F_{H_S}))^\circ$ .

The proof is given in Section B.

**Remark 5.9.** Note  $\widehat{\alpha}(\mathbf{x})$  may not be consistent outside  $(\text{supp}(F_{H_S}))^\circ$ , where the value of  $\alpha(\mathbf{x})$  is set to be zero in the definition for convenience; see (4.4). However, this has no impact on the accuracy of the exploitation estimator as  $\mathbf{1}_{H_S \leq \mathbf{x}}$  is constant.

Our second intermediate result shows then that the exploitation estimator for the CDF of  $Y$  is consistent asymptotically in both the exploration sample count  $m$  and the exploitation sample count  $N_S$ .

**Lemma 5.10** (Uniform asymptotic consistency of the exploitation CDF estimator). *Under Assumptions 5.3, 5.4, and 5.6, then with probability 1,*

$$\sup_{\mathbf{x} \in \mathbb{R}^d} |\widetilde{F}_S(\mathbf{x}) - F_Y(\mathbf{x})| \rightarrow 0 \quad m, N_S \rightarrow \infty. \quad (5.19)$$

See Section C for the proof.

Again, the proof of our main result, Theorem 5.7, is in Section D, which leverages the results in Lemma 5.8 and Lemma 5.10. One additional high-level step needed to prove Theorem 5.7 is to show that Algorithm 1 for asymptotically large budget  $B$  results in both  $m$  and  $N_S$  going to infinity. This is the first part of the proof presented in Section D.

## 6 Numerical simulations

In this section, we compare cvMDL and its variants with other algorithms including ECDF and AETC-d on several forward uncertainty quantification scenarios, including a single-output parametric PDE problem and a multi-output SDE problem. We label algorithms under consideration as follows:

- (ECDF) The empirical CDF estimator for  $F_Y$  using the high-fidelity samples only;
- (AETC-d) The AETC-d algorithm from [30];
- (cvMDL) Algorithm 1 with  $h(X_S; \mathbf{x})$  constructed using the linear approximations of Section 4;
- (cvMDL-sorted) cvMDL with the exploitation CDF monotonicity fix in Section 5.2;
- (cvMDL\*) cvMDL that estimates  $\hat{\alpha}(\mathbf{x})$  in the tail regime using (5.12) with  $\tau = 0.05$  when  $d = 1$ ;
- (cvMDL\*-sorted) cvMDL\* with the CDF monotonicity fix.

For the weight function in the cvMDL algorithm and its variants, we choose  $\omega(x) \equiv 1$  for all  $x \in \mathbb{R}$  when  $Y$  is scalar-valued, but in a case-dependent manner when  $Y$  is vector-valued. Since the estimators produced by the cvMDL-type and AETC-d algorithms are random (depending on the exploration data), for every budget value  $B$ , we repeat the experiment 100 times and report both the average of the mean  $\omega$ -weighted  $L^2$  error and the corresponding 5%-95% quantiles to measure the uncertainty.

### 6.1 Parametric PDEs

We consider a multifidelity model associated with a parametric elliptic equation, where lower fidelity models are identified through mesh coarsening. The setup is taken from [29, Section 7.1]. The elliptic PDE is over a square spatial domain  $D = [0, 1]^2$  that governs displacement in linear elasticity; see Figure 1. The parametric version of this problem equation seeks the displacement field  $\mathbf{u} = (u, v)^\top$  that is the solution to the PDE system,

$$-\nabla \cdot (\kappa(\mathbf{p}, \mathbf{x}) \boldsymbol{\sigma}(\mathbf{x}, \mathbf{p})) = \mathbf{F}(\mathbf{x}), \quad \forall (\mathbf{p}, \mathbf{x}) \in \mathcal{P} \times D$$

$$\boldsymbol{\sigma} = \begin{bmatrix} \sigma_1 & \sigma_{12} \\ \sigma_{12} & \sigma_2 \end{bmatrix}, \quad \begin{bmatrix} \sigma_1 \\ \sigma_2 \\ \sigma_{12} \end{bmatrix} = \frac{1}{1 - \nu^2} \begin{bmatrix} \frac{\partial u}{\partial x_1} + \frac{\partial v}{\partial x_2} \\ \frac{\partial v}{\partial x_2} + \nu \frac{\partial u}{\partial x_1} \\ \frac{1-\nu}{2} \left( \frac{\partial u}{\partial x_1} - \frac{\partial v}{\partial x_2} \right) \end{bmatrix}$$

where  $\mathbf{p} \in \mathbb{R}^4$  is a random vector with independent components uniformly distributed on  $[-1, 1]$ . We have fixed-displacement boundary conditions on the left wall, with the forcing  $\mathbf{F}$  being nonzero only on the right edge of the structure and equal to the constant 1. We set the Poisson ratio to  $\nu = 0.3$ , and  $\kappa(\mathbf{p}, \mathbf{x})$  is a scalar modeled as a truncated Karhunen-Loève expansion, given by


$$\kappa(\mathbf{p}, \mathbf{x}) = 1 + 0.5 \sum_{i=1}^4 \sqrt{\lambda_i} \phi_i(\mathbf{x}) p_i,$$

where  $(\lambda_i, \phi_i)$  are ordered eigenpairs of the exponential covariance kernel  $K$  on  $D$ , i.e.,

$$K(\mathbf{x}, \mathbf{y}) = \exp(-\|\mathbf{x} - \mathbf{y}\|_1/a),$$

where  $\|\cdot\|_1$  is the  $\ell^1$ -norm on vectors, and we choose  $a = 0.7$ . The displacement  $\mathbf{u}$  is used to compute a scalar QoI, the structural *compliance* or energy norm of the solution, which is the measure of elastic energy absorbed in the structure as a result of loading:

$$E := \int_D (\mathbf{u} \cdot \mathbf{F}) d\mathbf{x}. \quad (6.1)$$



Model $S$	{1}	{2}	{3}	{4}	{1, 2}	{1, 3}	{1, 4}	{2, 3}
$\gamma_S$	123.8	149.3	203.9	304.8	25.2	48.6	93.7	62.2
$m_S^*$ when $B = 10^7$	1998	2231	2337	2390	1253	1657	1909	2054


Model $S$	{2, 4}	{3, 4}	{1, 2, 3}	{1, 2, 4}	{1, 3, 4}	{2, 3, 4}	<b>{1, 2, 3, 4}</b>
$\gamma_S$	107.6	129.7	11.8	11.7	14.3	11.9	<b>11.5</b>
$m_S^*$ when $B = 10^7$	2175	2292	669	734	976	1540	<b>638</b>

**Figure 1:** Left: Geometry and boundary conditions for the linear elastic structure. Right: Oracle scaled loss  $\gamma_S$  (3.13) and optimal exploration sample count  $m_S^*$  (3.12) for different choices of  $S$ . The optimal model  $S$  is typed in boldface. Oracle statistics are computed with 50000 samples.

We solve the above system for each fixed  $\mathbf{p}$  via the finite element method with standard bilinear square isotropic finite elements on a rectangular mesh [1].

In this example, we form a multifidelity hierarchy through mesh coarsening via the mesh parameter  $h$ . The elastic energy  $E$  computed with mesh size  $h = 2^{-7}$  is the high-fidelity model (i.e., a scalar-valued  $Y$ ). We create four low-fidelity models  $X_1, \dots, X_4$  based on more economical discretizations:  $h = 2^{-4}, 2^{-3}, 2^{-2}, 2^{-1}$ . The outputs of these models are the energy  $E$  computed from the respective approximate solutions, all scalars. Hence, in this example  $d = d_i = 1$  for  $i \in [4]$ .

The cost for each model is the computational time, which we take to be inversely proportional to the mesh size squared, i.e.,  $h^2$ . This corresponds to using a linear solver of optimal linear complexity. We normalize cost so that the model with the lowest fidelity has unit cost, i.e.,  $(c_0, c_1, c_2, c_3, c_4) = (4096, 64, 16, 4, 1)$ . The correlations between the outputs of  $Y$  and  $X_1, X_2, X_3, X_4$  are 0.976, 0.940, 0.841,  $-0.146$ , respectively. The total budget  $B$  is taken on the interval  $[10^5, 10^7]$ .



**Figure 2:** Mean  $\omega(x)$ -weighted  $L^2$  error between  $F_Y$  and the estimated CDFs given by ECDF, AETC-d, cvMDL, and its variants, with the 5%-95% quantiles (for ease of visualization, we only plot the quantiles for AETC-d and cvMDL\*-sorted) to measure the uncertainty (a). Frequency of different models selected by cvMDL (b). Scatter plot of the estimated  $\alpha(x)$  and  $\rho(x)$  when  $S = \{1, 2, 3, 4\}$  using 50000 i.i.d. samples in the 1%-99%-quantile regime of  $Y$ . Gaussian kernel smoothing is applied to both data Gaussian kernel with bandwidth  $h = 0.0358$  chosen using 5-fold cross validation (c).


### 6.1.1 Results for estimating the distribution

In Figure 2(a), we see that AETC-d has the smallest error for smaller budgets but its asymptotic convergence is constrained by the model misspecification effects (associated with theoretical assumptions on the applicability of AETC-d), i.e., the error curve starts to plateau when  $B$  exceeds  $10^6$ . Although this can be mitigated by including additional nonlinear (e.g. polynomial) terms as additional covariates, trustworthy practical guidance is still lacking for this approach. On the other hand, both cvMDL and its variants demonstrate superior performance over ECDF, with cvMDL\*-sorted achieving a result competitive AETC-d without the plateau effect.

In Figure 2(b), we note that as the budget increases, the model  $\hat{S}^*$  selected by cvMDL converges to  $\{1, 2, 3, 4\}$ , which is the same as the optimal model computed under oracle statistics in Figure 1 (right). We note that the suboptimal model  $S = \{2, 3, 4\}$  is selected often by cvMDL, but not other models whose  $\gamma_S$  is close to that of  $\{1, 2, 3, 4\}$  (i.e. models  $\{1, 2, 3\}, \{1, 2, 4\}, \{1, 3, 4\}$ ). We believe this occurrence is due to the aggressive exploration steps taken by Algorithm 1, in particular when we double exploration sample counts ( $m \leftarrow 2m$ ) and then suboptimal models  $S$  with large values of  $m_S^*$  (e.g.,  $S = \{2, 3, 4\}$ ) become the preferred computational model.

The significance error reduction accomplished by cvMDL is indicated in Figure 2(c) by the values of  $\rho_S(x) = \text{Corr}[\mathbf{1}_{Y \leq x}, h(X_S; x)]$  where  $S = \{1, 2, 3, 4\}$ , which are close to 1. For cvMDL variants, either estimating  $\alpha(x)$  in the tail regime through continuity extension (5.9) (cvMDL\*) or sorting CDF values to ensure monotonicity (cvMDL/cvMDL\*-sorted) can help further reduce the errors. The former is particularly helpful in the small-budget regime where exploration data are not sufficient to estimate the full support of the output.

The weight function  $\omega(x)$  in this scenario is constant on  $\mathbb{R}$  thus the estimators produced by cvMDL-type estimators are expected to capture the global structure of  $F_Y$  (e.g. bulk and tails). To inspect this, we fix  $B = 10^7$  and plot the estimated CDFs by ECDF, AETC-d, and cvMDL\*-sorted in the tail and bulk regimes separately. The CDFs of the pointwise errors (at 1000 discretization points) in the three regimes are shown in Figure 3. It can be seen that cvMDL\*-sorted has the smallest errors in all three regimes.



**Figure 3:** One realization of how CDF errors given by cvMDL\*-sorted, AETC-d, and ECDF are distributed when  $B = 10^7$ . We plot CDF errors in three different regimes: (a) the lower tail of  $Y$  defined by the 0 – 0.05 quantile region, (b) the bulk defined by the 0.05 – 0.95 quantile region, (c) the upper tail defined by the 0.95 – 1.00 quantile region.


### 6.1.2 Risk metrics

To further compare the estimated distributions in the tail regime, we compare some risk metrics of the estimated CDFs. For example, one frequently used metric is the conditional VaR (CVaR), also called the expected shortfall, which is defined as the conditional expectation of  $Y$  in a tail regime (here,

$Y$  being large):

$$C(a) = \mathbb{E}[Y | F_Y(Y) \geq a] = \frac{1}{1-a} \int_a^1 F_Y^{-1}(x) dx \quad 0 < a < 1.$$

Assuming  $F_Y$  is known,  $C(a)$  can be numerically computed using root-finding algorithms. Fixing  $B = 10^7$  as before, we use the estimated CDFs by ECDF, AETC-d, and cvMDL\*-sorted to compute the CVaR of  $Y$  for  $a = 0.95$  and  $0.99$ , respectively. The experiment is repeated 50 times, and the corresponding statistics are summarized using boxplots in Figure 4 (a)-(b). For both choices of  $a$ , cvMDL\*-sorted outperforms the other methods by a noticeable margin. It is worth noting that although AETC-d and cvMDL\*-sorted have similar errors under the tested budget globally (Figure 2 (a)), the model misspecification effects result in the former estimates being systematically biased upward. The cvMDL\*-sorted estimates, on the other hand, remain unbiased.



**Figure 4:** (a): Boxplots of the CVaR (0.95) computed using the estimated CDFs given by ECDF, AETC-d, and cvMDL\*-sorted when  $B = 10^7$  with 50 experiments. (b): Same but for CVaR (0.99). (c): Inspection of how well the estimated loss  $\widehat{L}_S$  mimics and oracle loss curve as a function of  $m$ . The discrete data are fitted using a function  $f(m; a, b)$  of the form  $\frac{a}{m} + \frac{b}{B - c_{ep}m}$ , with fitted values for  $a$  and  $b$  given by  $2.64 \times 10^{-4}$  and  $5.57$ , respectively.

### 6.1.3 Oracle versus estimated loss

Finally, we investigate the model selection criteria used in cvMDL. Note for each  $S \subseteq [n]$ , there is a discrepancy between the exact loss function versus the estimator  $\widehat{L}_S$  constructed with empirical data. We now inspect if this approximation is reasonable. To numerically determine if the exploration-exploitation trade-off is optimal, we fix  $B = 10^7$  and  $S = \{1, 2, 3, 4\}$ . For a given value of  $m$ , we first take  $m$  exploration samples to estimate the control variates parameters and then use them to build an estimator  $\widetilde{F}_S$  for  $F_Y$  as in (5.11). We then compute the (exact) mean weighted  $L^2$  loss associated with this value of  $m$ . We repeat the experiment 10 times and compute the average loss value. We compile results of the above for  $m$  in the range from 200 to 1800. The results are reported in Figure 4 (c). It can be seen that the optimal exploration rate under oracle loss  $L_S$ , 638 (see Figure 1, right), almost matches the empirically identified optimal exploration rate, which is around 736. The small gap can be attributed to the underestimation of exploration error committed due to the finite-sample estimation of parameters.

## 6.2 Extrema of Geometric Brownian Motion (GBM)

GBM is a continuous-time stochastic process that is widely used in financial modeling. In a simple setting, GBM  $S_t$  is a random process with initial state  $s_0 > 0$  whose evolution is described using the

stochastic differential equation (SDE),

$$\begin{aligned} dS_t &= \mu S_t dt + \sigma S_t dW_t & t \geq 0 \\ S_0 &= s_0, \end{aligned}$$

where both  $\mu$  and  $\sigma > 0$  are constants, and  $W_t$  is a standard Brownian Motion process. A unique explicit solution for  $S_t$  exists and can be written as


$$S_t = s_0 \exp\left(\left(\mu - \frac{\sigma^2}{2}\right)t + \sigma W_t\right).$$

Set  $\mu = 0.05, \sigma = 0.2, s_0 = 1$ . We are interested in estimating the joint distribution of the extreme values of  $S_t$  over the time interval  $[0, 1]$ :

$$(S_{\min}, S_{\max}) \in \mathbb{R}^2 \quad S_{\min} := \min_{0 \leq t \leq 1} S_t, \quad S_{\max} := \max_{0 \leq t \leq 1} S_t.$$

I.e., our output quantity of interest is the vector  $(S_{\min}, S_{\max})$ . We evaluate these quantities by discretizing the SDE in time with time step  $\Delta t$  and computing the discrete extrema. The computational complexity (cost) of the corresponding procedure is proportional to the number of grid points used for discretization.

We construct a multifidelity model for this problem based on time discretization. In particular, we consider four different time scales  $\Delta t \in \{2^{-4}, 2^{-6}, 2^{-8}, 2^{-14}\}$ , where the pair of extreme values computed using the smallest mesh size ( $2^{-14}$ ) is treated as the high-fidelity model, and the three low-fidelity models  $X_1, X_2, X_3$  are the pairs computed under mesh size  $2^{-8}, 2^{-6}, 2^{-4}$ , respectively, where for joint samples the randomness of  $W_t$  is simulated from the same realization of GBM used in the high-fidelity model. The oracle CDF of the high-fidelity model is computed using MC with  $10^5$  samples and its shape is visualized in Figure 5 (left, middle). We also provide oracle model loss and exploration sample count in Figure 5 (right), and oracle correlations between the outputs of the high- and low-fidelity models in Table 1, and




**Figure 5:** Left and Middle: Oracle CDF of  $(S_{\min}, S_{\max})$  in the high-fidelity model computed using  $10^5$  Monte Carlo samples. Right: Oracle scaled loss  $\gamma_S$  (3.13) and  $m_S^*$  (3.12) for budget  $B = 10^6$ , computed using 50000 samples.

Model outputs	$S_{\min}(2^{-8})$	$S_{\max}(2^{-8})$	$S_{\min}(2^{-6})$	$S_{\max}(2^{-6})$	$S_{\min}(2^{-4})$	$S_{\max}(2^{-4})$
$S_{\min}(2^{-14})$	0.999	0.682	0.997	0.682	0.984	0.680
$S_{\max}(2^{-14})$	0.681	0.999	0.681	0.998	0.674	0.988


**Table 1:** Oracle correlations between the high-fidelity and low-fidelity model outputs computed using 50000 samples.

We define the cost for each model as their theoretical computational complexity which is inversely proportional to the time scale. With the cost of the cheapest model normalized to 1, we have  $(c_0, c_1, c_2, c_3) =$

(1024, 16, 4, 1). The total budget  $B$  takes values in  $[10^4, 10^6]$ . In this example, all models are two-dimensional random vectors so AETC-d cannot be directly applied. For cvMDL and its variants, setting  $\omega(\mathbf{x}) \equiv 1$  violates Assumption 5.5. Instead, since  $S_{\min} \leq s_0 = 1 \leq S_{\max}$ , we choose  $\omega(\mathbf{x}) = \mathbf{1}_T$  as an indicator function on a two-dimensional bounded region  $T \subset \mathbb{R}^2$  where the most likely outcomes reside. For instance, here we take  $T = [0.5, 1] \times [1, 3]$ . The statistics of the estimation errors and the selected models by cvMDL are reported in Figure 6(a),(b). Panel (c) shows that the correlation coefficient  $\rho_S(\mathbf{x})$ , is close to unity over the entire domain, suggesting that our chosen control variate (3.6) is a good choice. For better comparison, we also provide in Figure 7 an instance of a heatmap of the absolute estimation errors of ECDF, cvMDL, and cvMDL-sorted when  $B = 10^6$ .



**Figure 6:** (a): Mean  $\omega(\mathbf{x})$ -weighted  $L^2$  error between  $F_Y$  and the estimated CDFs given by ECDF, cvMDL, and cvMDL-sorted, with the 5%-50%-95% quantiles to measure the uncertainty. (b): Frequency of different models selected by cvMDL. (c) Estimated  $\rho_S(\mathbf{x})$  from (3.6) when  $S = \{1\}$  using 50000 i.i.d. samples for  $\mathbf{x} \in T$ .



**Figure 7:** An instance of absolute pointwise estimation errors of (a) ECDF, (b) cvMDL, and (c) cvMDL-sorted for budget  $B = 10^6$ .

Figures 6 and 7 show that cvMDL is consistent on the region  $T$  and substantially outperforms ECDF on average. As the budget goes to infinity, the model selected by cvMDL converges to the single low-fidelity model  $\{1\}$ , which coincides with the optimal model computed using oracle statistics in Figure 5 (right). With additional sorting to stabilize the algorithm, cvMDL-sorted further reduces the errors of cvMDL, which is consistent with the observations in the 1d case. In the pre-asymptotic regime when the budget is small, the models selected by cvMDL have relatively large fluctuations, but these stabilize for larger budgets.

## 7 Conclusions

We developed a versatile framework for efficiently estimating the CDFs of QoI subject to a budget constraint. To implement this framework, we constructed a set of binary control variables based on linear surrogates and integrated them into an adaptive meta-algorithm (cvMDL) that estimates the CDFs. We were able to establish both uniform consistency and trade-off optimality for the corresponding algorithm as the budget tends to infinity.

Although our framework is built upon an existing bandit-learning paradigm, our treatment of exploration and exploitation distinguishes itself from the previous works. In particular, the new approach employed in our framework leads to innovative estimators that dramatically lessen the reliance on underlying model assumptions. Furthermore, the approach allows for the treatment of different types of model outputs. To the best of our knowledge, our framework provides the first robust multifidelity CDF estimator under a budget constraint that can deal with both heterogeneous model sets and multi-valued outputs at the same time, meanwhile requiring no a priori cross-model statistics.

## Acknowledgment

RH is partially supported by the Hong Kong Research Grants Council (grant no. 14301821) and Start-up Fund for New Recruits, The Hong Kong Polytechnic University. AN is partially supported by NSF DMS-1848508 and AFOSR FA9550-20-1-0338. The views expressed in this article do not necessarily represent the views of Wells Fargo.

## References

- [1] E. Andreassen, A. Clausen, M. Schevenels, B. S. Lazarov, and O. Sigmund. “Efficient topology optimization in MATLAB using 88 lines of code”. *Structural and Multidisciplinary Optimization* 43.1 (2011), pp. 1–16.
- [2] Q. Ayoul-Guilmard, S. Ganesh, S. Krumscheid, and F. Nobile. “Quantifying uncertain system outputs via the multi-level Monte Carlo method—distribution and robustness measures”. *arXiv preprint arXiv:2208.07252* (2022).
- [3] S. Bobkov and M. Ledoux. “One-dimensional empirical measures, order statistics, and Kantorovich transport distances”. *Memoirs of the American Mathematical Society* 261.1259 (2019).
- [4] P. Bühlmann and S. Van De Geer. *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media, 2011.
- [5] M. Croci, K. Willcox, and S. Wright. “Multi-output multilevel best linear unbiased estimators via semidefinite programming”. *arXiv preprint arXiv:2301.07831* (2023).
- [6] I.-G. Farcas. “Context-aware model hierarchies for higher-dimensional uncertainty quantification”. PhD thesis. Technische Universität München, 2020.
- [7] I.-G. Farcas, B. Peherstorfer, T. Neckel, F. Jenko, and H.-J. Bungartz. “Context-aware learning of hierarchies of low-fidelity models for multi-fidelity uncertainty quantification”. *Computer Methods in Applied Mechanics and Engineering* 406 (2023), p. 115908.
- [8] M. B. Giles. “Multilevel monte carlo path simulation”. *Operations research* 56.3 (2008), pp. 607–617.
- [9] M. B. Giles, T. Nagapetyan, and K. Ritter. “Multilevel Monte Carlo approximation of distribution functions and densities”. *SIAM/ASA journal on Uncertainty Quantification* 3.1 (2015), pp. 267–295.
- [10] M. B. Giles, T. Nagapetyan, and K. Ritter. “Adaptive multilevel monte carlo approximation of distribution functions”. *arXiv preprint arXiv:1706.06869* (2017).
- [11] P. Glasserman. *Monte Carlo methods in financial engineering*. Vol. 53. Springer, 2004.
- [12] P. Glasserman, P. Heidelberger, and P. Shahabuddin. “Efficient Monte Carlo methods for value-at-risk” (2000).
- [13] A. Gorodetsky, J. Jakeman, A. Narayan, E. Ullman, and Y. Xu. “Generalized bandit-learning multifidelity approximation”. *Preprint* (2023).



- [14] A. A. Gorodetsky, G. Geraci, M. S. Eldred, and J. D. Jakeman. “A generalized approximate control variate framework for multifidelity uncertainty quantification”. *Journal of Computational Physics* 408 (2020), p. 109257.
- [15] A. Gruber, M. Gunzburger, L. Ju, and Z. Wang. “A Multifidelity Monte Carlo Method for Realistic Computational Budgets”. *arXiv preprint arXiv:2206.07572* (2022).
- [16] T. C. Hesterberg and B. L. Nelson. “Control variates for probability and quantile estimation”. *Management Science* 44.9 (1998), pp. 1295–1312.
- [17] S. Krumscheid and F. Nobile. “Multilevel Monte Carlo approximation of functions”. *SIAM/ASA Journal on Uncertainty Quantification* 6.3 (2018), pp. 1256–1293.
- [18] T. Lattimore and C. Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- [19] D. Lu, G. Zhang, C. Webster, and C. Barbier. “An improved multilevel Monte Carlo method for estimating probability distribution functions in stochastic oil reservoir simulations”. *Water resources research* 52.12 (2016), pp. 9642–9660.
- [20] A. J. McNeil, R. Frey, and P. Embrechts. *Quantitative risk management: concepts, techniques and tools-revised edition*. Princeton university press, 2015.
- [21] A. B. Owen. “Monte Carlo theory, methods and examples” (2013).
- [22] B. Peherstorfer. “Multifidelity Monte Carlo estimation with adaptive low-fidelity models”. *SIAM/ASA Journal on Uncertainty Quantification* 7.2 (2019), pp. 579–603.
- [23] B. Peherstorfer, K. Willcox, and M. Gunzburger. “Optimal Model Management for Multifidelity Monte Carlo Estimation”. *SIAM Journal on Scientific Computing* 38.5 (2016), A3163–A3194.
- [24] B. Peherstorfer, K. Willcox, and M. Gunzburger. “Survey of Multifidelity Methods in Uncertainty Propagation, Inference, and Optimization”. *SIAM Review* 60.3 (2018), A550–A591.
- [25] C. R. Rao, C. R. Rao, M. Statistiker, C. R. Rao, and C. R. Rao. *Linear statistical inference and its applications*. Vol. 2. Wiley New York, 1973.
- [26] D. Schaden and E. Ullmann. “On Multilevel Best Linear Unbiased Estimators”. *SIAM/ASA Journal on Uncertainty Quantification* 8.2 (2020), pp. 601–635.
- [27] D. Schaden and E. Ullmann. “Asymptotic analysis of multilevel best linear unbiased estimators”. *SIAM/ASA Journal on Uncertainty Quantification* 9.3 (2021), pp. 953–978.
- [28] R. Vershynin. *High-dimensional probability: An introduction with applications in data science*. Vol. 47. Cambridge university press, 2018.
- [29] Y. Xu, V. Keshavarzadeh, R. M. Kirby, and A. Narayan. “A bandit-learning approach to multifidelity approximation”. *SIAM Journal on Scientific Computing* 44.1 (2022), A150–A175.
- [30] Y. Xu and A. Narayan. “Budget-limited distribution learning in multifidelity problems”. *Numerische Mathematik* (2022), pp. 1–42.

## A Detailed cvMDL pseudocode

We provide here a more detailed version of the cvMDL summary routine in Algorithm 1. The main additions here are that we explicitly point to equations that must be utilized to implement the algorithm.

---

**Algorithm 2:** The detailed cvMDL algorithm.

---

**Input:**  $B$ : total budget, model costs  $c_0, c_1, \dots, c_n$   
**Output:** an estimator for  $F_Y(\mathbf{x})$

- 1: Initialize exploration = TRUE
- 2: Initialize  $m = \sum_{i \in [n]} d_i + 2$
- 3: Generate  $m$  exploration samples of  $(Y, X_{[n]})$
- 4: **while** exploration = TRUE **do**
- 5:   **for**  $S \subseteq [n]$  **do**
- 6:     Compute regression coefficients  $\hat{\beta}_{S^+}^{(i)}$ ,  $i \in [d]$  from (5.1)
- 7:     Construct  $\hat{H}_S(X_S)$  and  $\hat{h}(X_S; \mathbf{x})$  from (5.2)
- 8:     Compute regression coefficients  $r_j(\mathbf{x})$ ,  $j \in [m]$  from (5.4)
- 9:     Construct  $\mathcal{K}_1$  and  $\mathcal{K}_2$  from (5.4) and (5.5), respectively
- 10:     Evaluate  $\hat{k}_1(S)$  and  $\hat{k}_2(S)$  using (5.6) and a quadrature rule on  $\mathbb{R}^d$
- 11:     Compute  $\hat{m}_S^*$  and  $\hat{L}_S(\cdot; m)$  from (5.7)
- 12:     Compute the minimal expected loss  $\hat{L}_S(m \vee \hat{m}_S^*; m)$  from (5.7)
- 13:   **end for**
- 14:   Choose  $\hat{S}^* = \arg \min_{S \subseteq [n]} \hat{L}_S(m \vee \hat{m}_S^*; m)$ ;
- 15:   **if**  $m < \hat{m}_{\hat{S}^*}^*$  **then**
- 16:     Generate  $Q(m, \hat{m}_{\hat{S}^*}^*) - m$  additional samples of  $(Y, X_{[n]})$ , where  $Q$  is given in (5.13)
- 17:     Increase  $m$ :  $m \leftarrow Q(m, \hat{m}_{\hat{S}^*}^*)$
- 18:   **else**
- 19:     exploration = FALSE
- 20:   **end if**
- 21: **end while**
- 22: Generate  $N_{\hat{S}^*}$  samples of  $X_{\hat{S}^*}$ , with  $N_S$  given in (3.3)
- 23: Construct  $\hat{\alpha}(\mathbf{x})$  for  $S \leftarrow \hat{S}^*$  using (5.9)
- 24: Generate  $\hat{S}^*$  exploitation estimator  $\tilde{F}_{\hat{S}^*}$  using (5.11).

---

## B Proof of Lemma 5.8

This section contains the proofs of statements (i) through (vi) in Lemma 5.8. The proof of statement (ii), the asymptotic consistency of  $\hat{B}_{S^+}$ , is a direct result of the strong law of large numbers (SLLN), so we omit this proof.

### B.1 Proof of Statement (i)

We only prove (5.16b) as the proof for (5.16a) is similar. Denote by  $\mathbf{e}_i$  the  $i$ -th unit vector in  $\mathbb{R}^d$ , i.e.,  $\mathbf{e}_d^{(j)} = \delta_{ij}$ ,  $j \in [d]$ , where  $\delta$  is the Kronecker notation, and  $\mathbf{e} = \sum_{i \in [d]} \mathbf{e}_i$  is the all-one vector in  $\mathbb{R}^d$ . For fixed  $\mathbf{x} \in \mathbb{R}^d$ , without loss of generality, we assume  $F_{V(\mathbf{A}) \vee Y}(\mathbf{x}) \leq F_{V(\mathbf{B}_{S^+}) \vee Y}(\mathbf{x})$ , as the other case is similar by changing the roles of  $F_{V(\mathbf{A}) \vee Y}(\mathbf{x})$  and  $F_{V(\mathbf{B}_{S^+}) \vee Y}(\mathbf{x})$ . In this case, note

$$V(\mathbf{B}_{S^+}) \vee Y \leq \mathbf{x} \Rightarrow V(\mathbf{A}) \vee Y \leq \mathbf{x} + \Delta \mathbf{x} \Rightarrow V(\mathbf{A}) \vee Y - \|\Delta \mathbf{x}\|_\infty \mathbf{e} \leq \mathbf{x},$$

where

$$\Delta \mathbf{x} = \sum_{i \in [d]} \|X_{S^+}^\top(\mathbf{A}^{(i)} - \mathbf{B}_{S^+}^{(i)})\|_2 \mathbf{e}_i.$$

Hence, under Assumptions 5.3 and 5.6, for  $t > 0$ ,

$$\begin{aligned} |F_{V(\mathbf{A}) \vee Y}(\mathbf{x}) - F_{V(\mathbf{B}_{S^+}) \vee Y}(\mathbf{x})| &= F_{V(\mathbf{B}_{S^+}) \vee Y}(\mathbf{x}) - F_{V(\mathbf{A}) \vee Y}(\mathbf{x}) \\ &\leq F_{V(\mathbf{A}) \vee Y - \|\Delta \mathbf{x}\|_\infty}(\mathbf{x}) - F_{V(\mathbf{A}) \vee Y}(\mathbf{x}) \\ &\leq F_{V(\mathbf{A}) \vee Y - t\mathbf{e}}(\mathbf{x}) - F_{V(\mathbf{A}) \vee Y}(\mathbf{x}) + \mathbb{P}(\|\Delta \mathbf{x}\|_\infty \geq t) \\ &= F_{V(\mathbf{A}) \vee Y}(\mathbf{x} + t\mathbf{e}) - F_{V(\mathbf{A}) \vee Y}(\mathbf{x}) + \mathbb{P}(\|\Delta \mathbf{x}\|_\infty \geq t) \\ &\lesssim C\sqrt{dt} + \sum_{i \in [d]} \mathbb{P}\left(\|X_{S^+}^\top(\mathbf{A}^{(i)} - \mathbf{B}_{S^+}^{(i)})\|_2 \geq t\right) \\ &\stackrel{(5.14)}{\lesssim} C\sqrt{dt} + \frac{1}{t^2} \sum_{i \in [d]} \|\mathbf{A}^{(i)} - \mathbf{B}_{S^+}^{(i)}\|_2^2, \end{aligned}$$

where the penultimate inequality follows from the Lipschitz condition on  $F_{V(\mathbf{A}) \vee Y}$  and a union bound, and the last inequality follows from the Markov inequality. Taking  $t = \|\mathbf{A} - \mathbf{B}_{S^+}\|_F^{2/3}$  yields the desired result.

## B.2 Proof of Statement (iii)

We only prove the first statement; the second can be proved similarly. Recall that

$$\widehat{F}_{\widehat{H}_S}(\mathbf{x}) = G(\widehat{\mathbf{B}}_{S^+}) \quad G(\mathbf{A}) := \frac{1}{m} \sum_{\ell \in [m]} \mathbf{1}_{(X_{\text{epr}, \ell, S^+}^\top \mathbf{A})^\top \leq \mathbf{x}} \quad \mathbf{A} \in \mathbb{R}^{(d_S+1) \times d},$$

where  $X_{\text{epr}, \ell, S^+}$  denotes the  $\ell$ -th exploration sample of  $X_S$  with intercept. It follows from the direct computation that

$$\begin{aligned} &\sup_{\mathbf{x} \in \mathbb{R}^d} |\widehat{F}_{\widehat{H}_S}(\mathbf{x}) - F_{H_S}(\mathbf{x})| \\ &\leq \sup_{\mathbf{x} \in \mathbb{R}^d} \left| G(\widehat{\mathbf{B}}_{S^+}) - \mathbb{P}((X_{S^+}^\top \widehat{\mathbf{B}}_{S^+})^\top \leq \mathbf{x}) \right| + \sup_{\mathbf{x} \in \mathbb{R}^d} \left| \mathbb{P}((X_{S^+}^\top \widehat{\mathbf{B}}_{S^+})^\top \leq \mathbf{x}) - \mathbb{P}((X_{S^+}^\top \mathbf{B}_{S^+})^\top \leq \mathbf{x}) \right| \\ &\leq \sup_{\mathbf{A} \in \mathbb{R}^{(d_S+1) \times d}} \sup_{\mathbf{x} \in \mathbb{R}^d} \left| G(\mathbf{A}) - \mathbb{P}((X_{S^+}^\top \mathbf{A})^\top \leq \mathbf{x}) \right| + \sup_{\mathbf{x} \in \mathbb{R}^d} \left| \mathbb{P}((X_{S^+}^\top \widehat{\mathbf{B}}_{S^+})^\top \leq \mathbf{x}) - \mathbb{P}((X_{S^+}^\top \mathbf{B}_{S^+})^\top \leq \mathbf{x}) \right| \\ &= \underbrace{\sup_{\mathbf{A} \in \mathbb{R}^{(d_S+1) \times d}} \left| G(\mathbf{A}) - \mathbb{E}[G(\mathbf{A})] \right|}_{\Lambda_{m,1}} + \underbrace{\sup_{\mathbf{x} \in \mathbb{R}^d} \left| \mathbb{P}((X_{S^+}^\top \widehat{\mathbf{B}}_{S^+})^\top \leq \mathbf{x}) - \mathbb{P}((X_{S^+}^\top \mathbf{B}_{S^+})^\top \leq \mathbf{x}) \right|}_{\Lambda_{m,2}} \end{aligned}$$

where  $X_{S^+}$  a general notation that is independent of  $\widehat{\mathbf{B}}_{S^+}$ , and  $\Lambda_{m,1}$  has no supremum over  $\mathbf{x}$  (i.e.  $\mathbf{x}$  is treated as fixed) since one is able to alter the intercept coefficients in  $\mathbf{A}$  to yield different values of  $\mathbf{x} \in \mathbb{R}^d$  without changing the coefficients of  $X_S$ . In what follows, we show that both  $\Lambda_{m,1}$  and  $\Lambda_{m,2}$  converge to 0 a.s.

To bound  $\Lambda_{m,1}$ , we appeal to the empirical process theory. For any  $\mathbf{A} \in \mathbb{R}^{(d_S+1) \times d}$ , the indicator function  $\mathbf{1}_{(X_{\ell, S^+}^\top \mathbf{A})^\top \leq \mathbf{x}} \leq 1$ . According to Massart concentration inequality [4, Theorem 14.2], we have for any  $t > 0$  such that

$$\mathbb{P}(\Lambda_{m,1} > \mathbb{E}[\Lambda_{m,1}] + t) \leq \exp(-mt^2/8).$$

We take  $t = 4\sqrt{\log(m)/m}$ , which leads to

$$\mathbb{P}\left(\Lambda_{m,1} > \mathbb{E}[\Lambda_{m,1}] + 4\sqrt{\frac{\log m}{m}}\right) \leq m^{-2}. \quad (\text{B.1})$$

Since  $\sum_{m=1}^{\infty} m^{-2} < \infty$ , by the Borel-Cantelli lemma, we conclude that

$$\Lambda_{m,1} \leq \mathbb{E}[\Lambda_{m,1}] + 4\sqrt{\frac{\log m}{m}} \quad (\text{B.2})$$

for all sufficiently large  $m$  a.s. To bound  $\mathbb{E}[\Lambda_{m,1}]$ , note that the supremum in  $\mathbb{E}[\Lambda_{m,1}]$  is taken over all indicator functions defined on half-planes in  $\mathbb{R}^{d_S \times d}$  (the constant dimension is only a shift), which has Vapnik–Chervonenkis (VC) dimension  $(d_S d + 1)$ ; see [28, Section 8.3]. According to [28, Theorem 8.3.23], there exists a universal constant  $C'$  such that  $\mathbb{E}[\Lambda_{m,1}] \leq C' \sqrt{(d_S d + 1)/m}$ . This combined with (B.2) shows that  $\Lambda_{m,1} \rightarrow 0$  a.s.

To bound  $\Lambda_{m,2}$ , note that by Statement (ii) in Lemma 5.8, a.s. for all sufficiently large  $m$ ,  $\|\widehat{\mathbf{B}}_{S^+} - \mathbf{B}_{S^+}\|_F < \varepsilon$  where  $\varepsilon$  is the same as in Assumption 5.6. Since  $X_{S^+}$  is independent of  $\widehat{\mathbf{B}}_{S^+}$ , conditioning on  $\|\widehat{\mathbf{B}}_{S^+} - \mathbf{B}_{S^+}\|_F < \varepsilon$  and applying Statement (i) of Lemma 5.8,  $\Lambda_{m,2} \lesssim \|\widehat{\mathbf{B}}_{S^+} - \mathbf{B}_{S^+}\|_F^{2/3}$ . Now taking  $m \rightarrow \infty$  shows that  $\Lambda_{m,2} \rightarrow 0$  a.s.

### B.3 Proof of Statement (iv)

Note  $\mathcal{K}_1(\mathbf{x}) + \mathcal{K}_2(\mathbf{x}) = \widehat{F}_Y(\mathbf{x})(1 - \widehat{F}_Y(\mathbf{x}))$ , which is a consistent estimator for  $F_Y(\mathbf{x})(1 - F_Y(\mathbf{x}))$  for all  $\mathbf{x} \in \mathbb{R}^d$  a.s. as a result of SLLN. Therefore, it suffices to prove the consistency for  $\mathcal{K}_2(\mathbf{x})$  only.

Note  $\mathcal{K}_2(\mathbf{x})$  in (5.4) can be rewritten as

$$\mathcal{K}_2(\mathbf{x}) = \widehat{\rho}_S^2(\mathbf{x}) \widehat{F}_Y(\mathbf{x})(1 - \widehat{F}_Y(\mathbf{x})) = \begin{cases} \frac{(\widehat{F}_{Y \vee \widehat{H}_S}(\mathbf{x}) - \widehat{F}_Y(\mathbf{x}) \widehat{F}_{\widehat{H}_S}(\mathbf{x}))^2}{\widehat{F}_{\widehat{H}_S}(\mathbf{x})(1 - \widehat{F}_{\widehat{H}_S}(\mathbf{x}))} & \mathbf{x} \in (\text{supp}(\widehat{F}_{\widehat{H}_S}))^\circ \\ 0 & \text{otherwise} \end{cases} \quad (\text{B.3})$$

where

$$\widehat{\rho}_S^2(\mathbf{x}) = \begin{cases} \frac{(\widehat{F}_{Y \vee \widehat{H}_S}(\mathbf{x}) - \widehat{F}_Y(\mathbf{x}) \widehat{F}_{\widehat{H}_S}(\mathbf{x}))^2}{\widehat{F}_{\widehat{H}_S}(\mathbf{x})(1 - \widehat{F}_{\widehat{H}_S}(\mathbf{x})) \widehat{F}_Y(\mathbf{x})(1 - \widehat{F}_Y(\mathbf{x}))} & \mathbf{x} \in (\text{supp}(\widehat{F}_Y))^\circ \cap (\text{supp}(\widehat{F}_{\widehat{H}_S}))^\circ \\ 0 & \text{otherwise} \end{cases}$$

is the empirical estimator for  $\rho_S^2(\mathbf{x})$ . On the other hand,

$$\rho_S^2(\mathbf{x}) F_Y(\mathbf{x})(1 - F_Y(\mathbf{x})) = \begin{cases} \frac{(F_{Y \vee H_S}(\mathbf{x}) - F_Y(\mathbf{x}) F_{H_S}(\mathbf{x}))^2}{F_{H_S}(\mathbf{x})(1 - F_{H_S}(\mathbf{x}))} & \mathbf{x} \in (\text{supp}(F_{H_S}))^\circ \\ 0 & \text{otherwise} \end{cases} \quad (\text{B.4})$$

Comparing (B.3) and (B.4), the desired result follows from statement (iii) in Lemma 5.8.

### B.4 Proof of statement (v)

We prove the consistency of  $\widehat{k}_2(S)$ ; the consistency of  $\widehat{k}_1(S)$  can be proved similarly. By statement (iv) in Lemma 5.8,  $\mathcal{K}_2(\mathbf{x})$  converges to  $\rho_S^2(\mathbf{x}) F_Y(\mathbf{x})(1 - F_Y(\mathbf{x}))$  for all  $\mathbf{x} \in \mathbb{R}^d$  as  $m \rightarrow \infty$  a.s.

We first prove the first case where  $d = 1$  and  $\|\omega\|_{L^\infty(\mathbb{R})} = C < \infty$ , and we change the notation  $\mathbf{x}$  to the lowercase  $x$ . Under the moment condition in Assumption 5.3, according to [3, Theorem 2.13],

$$W_1(F_Y, \widehat{F}_Y) = \int_{\mathbb{R}} |\widehat{F}_Y(x) - F_Y(x)| dx \rightarrow 0 \quad m \rightarrow \infty,$$

where  $W_1$  is the Wasserstein-1 metric. Fix an arbitrary trajectory in the sample space such that  $\mathcal{K}_2(x) \rightarrow \rho_S^2(x)F_Y(x)(1 - F_Y(x))$  and  $\int_{\mathbb{R}} |\widehat{F}_Y(x) - F_Y(x)|dx \rightarrow 0$ . In the following, we shall treat  $\mathcal{K}_2(x)$  as a deterministic sequence.

To show the consistency of  $\widehat{k}_2(S)$ , it remains to justify the change of order of taking limit and integration:

$$\begin{aligned} \lim_{m \rightarrow \infty} \widehat{k}_2(S) &= \lim_{m \rightarrow \infty} c_S \int_{\mathbb{R}} \omega(x) \mathcal{K}_2(x) dx = c_S \int_{\mathbb{R}} \lim_{m \rightarrow \infty} \omega(x) \mathcal{K}_2(x) dx \\ &= c_S \int_{\mathbb{R}} \omega(x) \rho_S^2(x) F_Y(x) (1 - F_Y(x)) dx = k_2(S), \end{aligned}$$

for which we appeal to the Vitali convergence theorem. To apply the Vitali convergence theorem, we need to verify that the sequence  $\omega(x)\mathcal{K}_2(x)$  is uniformly integrable and has absolutely continuous integrals. To this end, recall the representation  $\mathcal{K}_2(x)$  in (B.3). Since the square of the empirical correlation estimator is bounded by 1, a.s.,

$$\omega(x)\mathcal{K}_2(x) \leq \omega(x)\widehat{F}_Y(x)(1 - \widehat{F}_Y(x)) \leq \frac{C}{4} < C.$$

The absolutely continuous integrals part follows immediately from the uniform boundedness. For uniform integrability, we first observe for every measurable  $I \subseteq \mathbb{R}$ ,

$$\begin{aligned} \int_I |\omega(x)F_Y(x)(1 - F_Y(x)) - \omega(x)\widehat{F}_Y(x)(1 - \widehat{F}_Y(x))| dx &\leq \int_I \omega(x) |\widehat{F}_Y(x) - F_Y(x)| dx \\ &\leq C \int_I |\widehat{F}_Y(x) - F_Y(x)| dx \rightarrow 0. \end{aligned}$$

Thus,

$$\begin{aligned} \int_{|x|>M} \omega(x)\mathcal{K}_2(x) dx &\leq \int_{|x|>M} \omega(x)\widehat{F}_Y(x)(1 - \widehat{F}_Y(x)) dx \\ &\leq \int_{|x|>M} \omega(x)F_Y(x)(1 - F_Y(x)) dx + C \int_{|x|>M} |\widehat{F}_Y(x) - F_Y(x)| dx \\ &\lesssim \int_{|x|>M} \frac{C}{x^2} dx + C \int_{\mathbb{R}} |\widehat{F}_Y(x) - F_Y(x)| dx, \end{aligned}$$

where the last step follows from Assumption 5.3 and the Chebyshev inequality. For every  $\varepsilon > 0$ , we can choose  $m$  and  $M$  sufficiently large so the right-hand side is less than  $\varepsilon$ . The uniform integrability follows by enlarging  $M$  to accommodate the first  $m$  terms.

The proof for (b) is similar. It suffices to verify the change of order for the sequence  $\omega(\mathbf{x})\mathcal{K}_2(\mathbf{x})$ . Since  $\omega(\mathbf{x})\mathcal{K}_2(\mathbf{x}) \leq \omega(\mathbf{x})$  and the latter is integrable and independent of  $m$ , the dominated convergence does the rest.

## B.5 Proof of statement (vi)

For  $\mathbf{x} \in (\text{supp}(F_{H_S}))^\circ$ , it is easy to show via a contradiction argument that  $\mathbf{x} \in \text{supp}(\widehat{F}_{\widehat{H}_S})$  for all sufficiently large  $m$  a.s. By statement (iii) in Lemma 5.8,  $\widehat{F}_{Y \vee \widehat{H}_S}$  and  $\widehat{F}_{\widehat{H}_S}$  are consistent estimators. Meanwhile,  $\widehat{F}_Y(\mathbf{x})$  is a consistent estimator for  $F_Y(\mathbf{x})$  due to SLLN. Therefore, we obtain

$$\widehat{\alpha}(\mathbf{x}) = \frac{\widehat{F}_{Y \vee \widehat{H}_S}(\mathbf{x}) - \widehat{F}_Y(\mathbf{x})\widehat{F}_{\widehat{H}_S}(\mathbf{x})}{\widehat{F}_{\widehat{H}_S}(\mathbf{x})(1 - \widehat{F}_{\widehat{H}_S}(\mathbf{x}))} \rightarrow \alpha(\mathbf{x}) = \frac{F_{Y \vee H_S}(\mathbf{x}) - F_Y(\mathbf{x})F_{H_S}(\mathbf{x})}{F_{H_S}(\mathbf{x})(1 - F_{H_S}(\mathbf{x}))}$$

as  $m \rightarrow \infty$  almost surely.

## C Proof of Lemma 5.10

Recall in (5.11) that

$$\begin{aligned}\tilde{F}_S(\mathbf{x}) &= \hat{F}_Y(\mathbf{x}) - \frac{1}{m} \sum_{\ell \in [m]} \left( \hat{\alpha}(\mathbf{x}) \hat{h}_S(X_{\text{epr}, \ell, S}; \mathbf{x}) - \frac{1}{N_S} \sum_{j \in [N_S]} \hat{\alpha}(\mathbf{x}) \hat{h}_S(X_{\text{ept}, \ell, S}; \mathbf{x}) \right) \\ &= \hat{F}_Y(\mathbf{x}) - \hat{\alpha}(\mathbf{x}) \hat{F}_{\hat{H}_S}(\mathbf{x}) + \hat{\alpha}(\mathbf{x}) \left( \frac{1}{N_S} \sum_{j \in [N_S]} \mathbf{1}_{(X_{\text{ept}, j, S}^\top + \hat{\mathbf{B}}_{S^+})^\top \leq \mathbf{x}} \right).\end{aligned}$$

Thus,

$$\begin{aligned}\sup_{\mathbf{x} \in \mathbb{R}^d} |\tilde{F}_S(\mathbf{x}) - F_Y(\mathbf{x})| &\leq \sup_{\mathbf{x} \in \mathbb{R}^d} |\hat{F}_Y(\mathbf{x}) - F_Y(\mathbf{x})| + \sup_{\mathbf{x} \in \mathbb{R}^d} |\hat{\alpha}(\mathbf{x}) (\hat{F}_{\hat{H}_S}(\mathbf{x}) - F_{H_S}(\mathbf{x}))| \\ &\quad + \sup_{\mathbf{x} \in \mathbb{R}^d} \left| \hat{\alpha}(\mathbf{x}) \left( \frac{1}{N_S} \sum_{j \in [N_S]} \mathbf{1}_{(X_{\text{ept}, j, S}^\top + \hat{\mathbf{B}}_{S^+})^\top \leq \mathbf{x}} - F_{H_S}(\mathbf{x}) \right) \right| \\ &\stackrel{(5.10)}{\leq} \underbrace{\sup_{\mathbf{x} \in \mathbb{R}^d} |\hat{F}_Y(\mathbf{x}) - F_Y(\mathbf{x})|}_{(i)} + \underbrace{\sup_{\mathbf{x} \in \mathbb{R}^d} |\hat{F}_{\hat{H}_S}(\mathbf{x}) - F_{H_S}(\mathbf{x})|}_{(ii)} \\ &\quad + \underbrace{\sup_{\mathbf{x} \in \mathbb{R}^d} \left| \left( \frac{1}{N_S} \sum_{j \in [N_S]} \mathbf{1}_{(X_{\text{ept}, j, S}^\top + \hat{\mathbf{B}}_{S^+})^\top \leq \mathbf{x}} - F_{H_S}(\mathbf{x}) \right) \right|}_{(iii)}.\end{aligned}$$

Note (i) converges to 0 as  $m \rightarrow \infty$  due to the multivariate Glivenko-Cantelli theorem. (ii) converges to 0 as  $m \rightarrow \infty$  due to statement (iii) in Lemma 5.8. A similar argument as in the proof of statement (iii) of Lemma 5.8 can be used to prove that (iii) converges to 0 as  $N_S \rightarrow \infty$ , which we do not repeat here.

## D Proof of Theorem 5.7

To reduce notational confusion with  $m$ , we use  $t$  to denote the number of exploration samples. The exploration rate grows nonlinearly with respect to an index that counts the iterations of the while loop in Algorithm 1. We let  $q$  denote the loop iteration index, and  $t_q$  the corresponding exploration rate, i.e.,  $t_1 = n + 2$ . Let  $q(B)$  be the total iteration steps in Algorithm 1, which is random. It follows from the definition that  $t_{q(B)} = m(B)$  and

$$n + 1 \leq t_q \leq t_{q+1} \leq 2t_q \quad 1 \leq q < q(B). \quad (\text{D.1})$$

We first show that  $m(B)$  diverges as  $B \rightarrow \infty$  a.s. According to statement (v) in Lemma 5.8,  $\hat{k}_1(S) \rightarrow k_1(S)$ ,  $\hat{k}_2(S) \rightarrow k_2(S)$  for  $S \subseteq [n]$  a.s. As a result, for almost every realization  $\omega \in \Omega$  (where  $\Omega$  denotes the product space of exploration samples), there exists an  $0 < L(\omega), L'(\omega) < \infty$ ,

$$\sup_{t > n+1} \max_{S \subseteq [n]} \hat{k}_1(S; \omega) < L(\omega) < \infty \quad \inf_{t > n+1} \min_{S \subseteq [n]} \hat{k}_2(S; \omega) > L'(\omega) > 0 \quad (\text{D.2})$$

The exploration stopping criterion of Algorithm 1 requires that

$$m(B; \omega) \geq \hat{m}_{S(B; \omega)}^* \geq \frac{B}{c_{\text{epr}} + \sqrt{\frac{c_{\text{epr}} L(\omega)}{L'(\omega)}}} \rightarrow \infty \quad B \rightarrow \infty. \quad (\text{D.3})$$

Thus,

$$\lim_{B \rightarrow \infty} m(B; \omega) = \infty. \quad (\text{D.4})$$

We now work with a fixed realization  $\omega$  along which  $m(B; \omega) \rightarrow \infty$  as  $B \rightarrow \infty$ , and  $\widehat{k}_1(S), \widehat{k}_2(S)$  converge to the true parameters as  $t \rightarrow \infty$ . We will prove that both (5.15a) and (5.15b) hold for such an  $\omega$ . Fix  $\delta < 1/2$  sufficiently small. Since  $S^*$  is assumed unique, a continuity argument implies that, there exists a sufficiently large  $T(\delta; \omega)$ , such that for all  $t \geq T(\delta; \omega)$ ,

$$\max_{(1-\delta)m_{S^*}^* \leq m \leq (1+\delta)m_{S^*}^*} \widehat{L}_{S^*}(m; t) < \min_{S \subseteq [n], S \neq S^*} \widehat{L}_S^*(t). \quad (\text{D.5})$$

$$1 - \delta \leq \frac{\widehat{m}_S^*(t; \omega)}{m_S^*} \leq 1 + \delta \quad \forall S \subseteq [n], \quad (\text{D.6})$$

where  $\widehat{L}_{S^*}(\cdot; t)$  is the estimated loss function for  $S^*$  using  $t$  exploration samples, and  $\widehat{L}_S^*(t)$  is the estimated  $L_S^*$  in (3.13) using  $t$  exploration samples.

Since  $m_S^*$  scales linearly in  $B$  and  $m(B; \omega)$  diverges as  $B \rightarrow \infty$ , there exists a sufficiently large  $B(\delta; \omega)$  such that for  $B > B(\delta; \omega)$ ,

$$\min_{S \subseteq [n]} m_S^* > 4T(\delta; \omega) \quad (\text{D.7})$$

$$t_{q(B)} = m(B; \omega) > 4T(\delta; \omega). \quad (\text{D.8})$$

Consider  $q' < q(B)$  that satisfies  $t_{q'-1} < T(\delta; \omega) \leq t_{q'}$ . Such a  $q'$  always exists due to (D.8), and satisfies

$$t_{q'} \stackrel{(\text{D.1})}{\leq} 2t_{q'-1} < 2T(\delta; \omega) \stackrel{(\text{D.7})}{\leq} \frac{1}{2} \min_{S \subseteq [n]} m_S^* \stackrel{(\text{D.6}), \delta < 1/2}{\leq} \widehat{m}_S^*(t_{q'}; \omega).$$

This inequality tells us that in the  $q'$ -th loop iteration, for all  $S \subseteq [n]$ , the corresponding estimated optimal exploration rate is larger than the current exploration rate. In this case,

$$\widehat{L}_S(t_{q'} \vee \widehat{m}_S^*(t_{q'}; \omega); t_{q'}) = \widehat{L}_S(\widehat{m}_S^*(t_{q'}; \omega); t_{q'}) = \widehat{L}_S^*(t_{q'}) \quad S \subseteq [n].$$

This, along with (D.5) and (D.6), tells us that  $S^*$  is the estimated optimal model in the current step, and more exploration is needed.

To see what  $t_{q'+1}$  should be, we consider two separate cases. If  $2t_{q'} \leq \widehat{m}_{S^*}^*(t_{q'}; \omega)$ , then

$$T(\delta; \omega) < t_{q'+1} = 2t_{q'} \leq \widehat{m}_{S^*}^*(t_{q'}; \omega) \leq (1 + \delta)m_{S^*}^*,$$

which implies

$$(1 - \delta)m_{S^*}^* \stackrel{(\text{D.6})}{\leq} t_{q'+1} \vee \widehat{m}_{S^*}^*(t_{q'+1}; \omega) \leq (1 + \delta)m_{S^*}^*. \quad (\text{D.9})$$

If  $t_{q'} \leq \widehat{m}_{S^*}^*(t_{q'}; \omega) < 2t_{q'}$ , then

$$t_{q'+1} = \left\lceil \frac{t_{q'} + \widehat{m}_{S^*}^*(t_{q'}; \omega)}{2} \right\rceil \leq \widehat{m}_{S^*}^*(t_{q'}; \omega) \stackrel{(\text{D.6})}{\leq} (1 + \delta)m_{S^*}^*,$$

which also implies (D.9). But (D.9) combined with (D.5) and (D.6) implies that  $S^*$  is again the estimated optimal model in the  $(q' + 1)$ -th loop iteration. Applying the above argument inductively proves  $S(B) = S^*$ , i.e. (5.15b). Note (D.9) holds true until the algorithm terminates, which combined with the termination criteria  $t_{q(B)} \geq \widehat{m}_{S^*}^*(t_{q(B)}; \omega) \geq (1 - \delta)m_{S^*}^*$  implies

$$1 - \delta \leq \frac{m(B; \omega)}{m_{S^*}^*} = \frac{t_{q(B)}}{m_{S^*}^*} \leq 1 + \delta.$$

(5.15a) now follows by noting that  $\delta$  can be set arbitrarily small.

Finally, let  $\tilde{F}'(\mathbf{x}; B)$  be chosen as in (5.11) with  $S = S^*$ ,  $m = m_{S^*}^*$  and  $N_S = (B - c_{\text{epr}} m_{S^*}^*) / c_{S^*}$ . Note both  $m, N_S$  are deterministic and diverge as  $B \rightarrow \infty$ . By the triangle inequality,

$$\sup_{\mathbf{x} \in \mathbb{R}^d} |\tilde{F}(\mathbf{x}; B) - F_Y(\mathbf{x})| \leq \sup_{\mathbf{x} \in \mathbb{R}^d} |\tilde{F}(\mathbf{x}; B) - \tilde{F}'(\mathbf{x}; B)| + \sup_{\mathbf{x} \in \mathbb{R}^d} |\tilde{F}'(\mathbf{x}; B) - F_Y(\mathbf{x})|.$$

As  $B \rightarrow \infty$ , the first term on the right-hand side converges to 0 due to (5.15a) and (5.15b) in Theorem 5.7, and the second term on the right-hand side converges to 0 due to Theorem 5.10. This proves (5.15c).