

# Budget-limited distribution learning in multifidelity problems

Yiming Xu<sup>\*1,2</sup> and Akil Narayan<sup>1,2</sup>

<sup>1</sup>Department of Mathematics, University of Utah, Salt Lake City, UT, USA, [yxu@math.utah.edu](mailto:yxu@math.utah.edu)

<sup>2</sup>Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT, USA,  
[akil@sci.utah.edu](mailto:akil@sci.utah.edu)

## Abstract

Multifidelity methods are widely used for estimating quantities of interest (QoI) in computational science by employing numerical simulations of differing costs and accuracies. Many methods approximate numerical-valued statistics that represent only limited information, e.g., scalar statistics, about the QoI. Further quantification of uncertainty, e.g., for risk assessment, failure probabilities, or confidence intervals, requires estimation of the full distributions. In this paper, we generalize the ideas in [Xu et al., *SIAM J. Sci. Comput.* 44.1 (2022), A150–A175] to develop a multifidelity method that approximates the full distribution of scalar-valued QoI. The main advantage of our approach compared to alternative methods is that we require no particular relationships among the high and lower-fidelity models (e.g. model hierarchy), and we do not assume any knowledge of model statistics including correlations and other cross-model statistics before the procedure starts. Under suitable assumptions in the framework above, we achieve provable 1-Wasserstein metric convergence of an algorithmically constructed distributional emulator via an exploration-exploitation strategy. We also prove that crucial policy actions taken by our algorithm are budget-asymptotically optimal. Numerical experiments are provided to support our theoretical analysis.

*Keywords:* multifidelity, Wasserstein distance, sequential decision-making, empirical measure, linear regression

## 1 Introduction

Estimation of output QoI from complex and large-scale simulations is an important task in many areas of computational science. A concrete example is a forward uncertainty quantification setup, where the QoI is an output of a physical system subject to modeled randomness/uncertainty, and the goal is to identify the typical behavior of the QoI by computing its expectation. A universal solution for this specific task is through Monte Carlo (MC) simulation [13], which in practice requires many repeated evaluations of an accurate forward model, and can be computationally infeasible for expensive models.

A modern collection of approaches that address this computational challenge is the suite of multifidelity methods [25]. Instead of operating on a single (high-fidelity) model alone, multifidelity methods combine several models of different accuracies and costs to accelerate computation. Low-fidelity models are less accurate but also inexpensive simulations, for example, resulting from numerical solvers of parametric partial differential equations (PDEs) using coarse discretizations. More generally, the lower fidelity models used in multifidelity methods arise from simplification or reduction of the high-fidelity model and thus are cheaper but less accurate. However, they often contain information that, if utilized properly, can contribute to characterizing QoI.

A prototypical example of a multifidelity method is the multilevel approach [9, 24, 12], which approximates the expectation of a scalar-valued QoI given by the high-fidelity model. Leveraging a telescoping

---

<sup>\*</sup>Mailing address: 155 1400 E, Salt Lake City, UT 84112

sum using hierarchical models, multilevel estimators make use of cross-model correlations to attain a smaller variance compared to a single-model MC estimator. Recent work has introduced a more general perspective for many existing methods within the multilevel framework, and provided a way to realize the optimal variance reduction among all linear unbiased estimators [29, 30]. Multilevel estimators are considered universal in the sense that they rely only on the covariance information of models, which is used as an input for the estimator construction. A similar approach has recently been developed in [34] that assumes only a linear model assumption but no *a priori* knowledge of covariance statistics.

Since multifidelity methods have been so successfully applied to the parametric estimation of QoI, it is natural to ask if it is possible to extend the same technique to also characterize their distributions or equivalent statistics such as characteristic functions. This question has been studied in recent works [10, 11, 19, 16]. The major application scenario under consideration is hierarchical models where special relationships between pairs of models of different fidelities are leveraged. For more general non-hierarchical multifidelity setups, paradigms for efficiently learning distributions of QoI are absent to our knowledge, which is a gap that this paper seeks to fill.

Learning a full distribution is possible when independent and identically distributed (i.i.d.) samples are available. Given enough samples, universal approaches using empirical cumulative distribution functions (CDFs) can be employed to estimate the true distribution. However, the general nonparametric nature of this approach is balanced by the slow convergence rates, which considerably limits the direct usage in applications where sampling is costly. In this paper, we employ such non-parametric estimators but ameliorate the cost using multifidelity strategies. Alternative distributional models in statistics are parametric, limiting the space of expressible distributions but enabling the use of classical tools such as the maximum likelihood estimation. However, realistic models are often so complex that it is difficult to identify an appropriately expressive parametric family, so we focus on a non-parametric strategy.

## 1.1 Contributions of this paper

Our approach is based on the Explore-Then-Commit (ETC) algorithm in bandit learning [18, 3], which partitions computational procedures into an *exploration* phase, where models are sampled to learn information about their relationships, and an *exploitation* phase, where the learned information guides the development and execution of a strategy that builds a predictor. Using such exploration/exploitation ideas to design adaptive algorithms in multifidelity estimation is not completely new; see [23] for a procedure that estimates the mean, and [7] for methods that build reduced models in the multifidelity context. These methods and others like them make substantially stronger assumptions than we do; in particular assumptions regarding hierarchical relationships between models and certain *a priori* knowledge of correlations and behavior of model costs<sup>1</sup> is common and ubiquitous. Our approach and setup are more relevant to [34], where no model relationships are assumed (in particular, no hierarchy need to be provided), no knowledge of correlations is provided, and no particular behavior of model costs is required. In the initial exploration phase, we learn about interactions between the high and lower fidelity models and construct a linear regression emulator for the high-fidelity output using a selection of inexpensive low-fidelity models. The exploration phase terminates after an adaptively-identified budget investment and is followed by the exploitation phase that expends the remaining budget to construct an empirical estimator for the regression emulator that approximates the full distribution of the high-fidelity model output.

Our approach, therefore, leverages models of different fidelities and costs as well as various statistical procedures to produce an efficient estimator for the unknown distribution of the high-fidelity QoI, which would be difficult and costly to estimate directly. Our procedure does not require any hierarchy or relationships between models and does not require *a priori* knowledge of any model or cross-model statistics.

From the statistical point of view, our approach could be viewed as “semi-parametric” (though this terminology should not be conflated with notions of semi-parametric regression analyses [28]). That is,

---

<sup>1</sup>For example, requiring that models with higher correlation relative to the high-fidelity model should also incur higher cost is one such model cost behavior.

while we do not prescribe a target family of parametric distributions for the high-fidelity model  $Y$  and we learn a non-parametric distribution for a regression residual discrepancy term, we do express the distribution of  $Y$  parametrically with respect to the (unknown) distributions of low-fidelity models  $X_i$ ,  $i \in \{1, \dots, n\}$ .

In summary, our contributions in this article are twofold:

- We introduce an adaptive sequential decision-making algorithm, “AETC-d”, Algorithm 1, that produces an estimator for the full distribution of a high-fidelity output given a prescribed budget. This estimator is built from low-fidelity models that have no prescribed hierarchy or relationships and the algorithm proceeds without initial knowledge of any statistics. (Cf. the approaches in [10, 11, 19, 16, 23] that require more restrictive assumptions.)
- We prove almost sure convergence of the algorithm-produced distributional estimator in the mean 1-Wasserstein metric for a large budget. (See Theorem 5.2 and Corollary 5.5.) We also show almost sure optimality guarantees for particular exploration decisions made by the adaptive algorithm. (Cf. the algorithm in [34] which is less efficient in exploration.)

We provide numerical examples that establish the efficacy of our approach. The strength of our theoretical guarantees (full distributional convergence with initial ignorance about model relationships or statistics) comes at the cost of certain technical assumptions. These are concretely described in Section 2.3, with the most stringent one being Assumption 2.1: that the conditional expectation of the high-fidelity model output on the low-fidelity model outputs can be written as a linear function of the latter. However, even in cases when this assumption is violated, we present a concrete strategy that empirically ameliorates this model misspecification; see section 6.1.3.

We emphasize again that our methodology does *not* require any particular knowledge of the models, i.e., hierarchical/nested structure or specific coupling assumptions, to guarantee convergence. The implementation requires only the identification of a trusted high-fidelity model, the ability to query the models themselves, and the cost of sampling each model relative to the cost of sampling the high-fidelity model. Our general framework is similar to [34], but the learning objectives (and hence also the ultimate algorithm and resulting theory) are quite different.

The rest of the paper is organized as follows. In Section 2, we set up the budget-limited distribution learning problem, and technically describe our assumptions and main results. In Section 3, we briefly review results concerning the convergence of empirical measures under Wasserstein metrics and provide a few technical results for later use. In Section 4, we propose an exploration-exploitation strategy for distribution learning and derive an asymptotically informative upper bound for the mean 1-Wasserstein error of the estimator. We then utilize this upper bound in Section 5.2 to devise an efficient adaptive algorithm, AETC-d, and establish a trajectory-wise optimality result for it. In Section 6, we provide a detailed numerical study of the AETC-d algorithm, investigating consistency, model misspecification, and optimality of exploration rates. In Section 7, we conclude by summarizing the main results of the paper.

## 2 Problem setup

### 2.1 Notation

Let  $Y, X_1, \dots, X_n \in \mathbb{R}$  be scalar-valued random outputs associated with the high-fidelity model and  $n$  low-fidelity surrogates, respectively. Let  $c_0$  and  $c_i$ ,  $i \in [n] := \{1, \dots, n\}$ , be the respective cost of sampling  $Y$  and  $X_i$ . The costs are assumed known and deterministic. No additional assumptions about the accuracy or costs of  $X_i$  relative to those of  $X_{i+1}$  are assumed in the following discussion. In particular, the index  $i$  does not represent an ordering based on cost, accuracy, or hierarchy.

Many recent advances in multifidelity methods center around the efficient estimation of  $\mathbb{E}[Y]$  [9, 24, 12, 29, 30, 34]. Under appropriate correlation and cost conditions and for a fixed budget, the estimators from these methods are significantly more accurate than the classical MC estimator using

i.i.d. samples of  $Y$ . For some applications, however, obtaining only precise estimates for parameter means is not sufficient. For instance, when estimating robust statistics such as the median, when building confidence intervals to quantify estimation uncertainty, or characterizing failure events and estimation of their corresponding probabilities, the full distribution of a QoI containing the complete information of modeled randomness is often required. Letting  $F_Y(y) = \mathbb{P}(Y \leq y)$  be the CDF of  $Y$ , we wish to find an efficient estimate for  $F_Y$  instead of only certain functionals of it. To make use of the low-fidelity models to this end, assumptions on cross-model correlations are not enough for this purpose. As opposed to imposing strong hierarchical assumptions on the models, we introduce an alternative parametric assumption on the relationship between  $X_i, i \in [n]$  and  $Y$  that allows efficient estimation of  $F_Y(y)$  through the  $X_i$ 's.

## 2.2 Linear regression

A simple yet useful assumption relating  $X_i$  and  $Y$  is through linear regression. For any  $S \subseteq [n]$ , we assume that

$$Y = X_S^\top \beta_S + \varepsilon_S \quad X_S = (1, (X_i)_{i \in [S]})^\top, \quad (1)$$

where  $X_S$  includes the intercept as the first component, and  $\varepsilon_S$  is a centered random variable with variance  $\sigma_S^2$  (i.e.  $\varepsilon_S \sim (0, \sigma_S^2)$ ) that is *independent* of  $X_S$ . Note that a similar linear model decomposition as (1) requiring  $\varepsilon_S$  only be uncorrelated with  $X_S$  always holds provided  $Y$  and  $X_S$  have bounded second moments, but independence requires more assumptions that we will describe further in Section 2.3.

The model (1) provides a way to simulate  $Y$  using only samples of  $X_S$  and a noise generator for  $\varepsilon_S$ . Under a fixed budget, sampling from  $X_S$  and  $\varepsilon_S$  may have a lower cost than sampling  $Y$  alone under appropriate cost assumptions, potentially leading to an estimator for  $F_Y$  with better accuracy. This is the core idea we propose and investigate.

## 2.3 Assumptions

This section codifies some particular assumptions that we make. None of these assumptions are required to employ our main algorithm, but our theoretical convergence guarantees do require all these assumptions. We begin by stating two assumptions that together fully characterize (1):

**Assumption 2.1.** *For every  $S \subseteq [n]$ , the  $X_S$ -conditional expectation of  $Y$ ,  $\mathbb{E}[Y|X_S]$ , is a linear function of  $X_S$ .*

**Assumption 2.2.** *For every  $S \subseteq [n]$ ,  $\varepsilon_S$  is independent of  $X_S$ .*

Although Assumption 2.1 initially appears strong, it is not overly restrictive since nonlinear functions of  $X_S$  can be added as additional regressors without changing the analysis. For example, linear regression of  $Y$  on the covariates  $X_S$  could be “extended” to a more general regression on the covariates  $X_S$  as well as quadratic covariates  $\text{vec}(X_S X_S^\top)$ . Since  $\mathbb{E}[Y|X_S]$  is a measurable function of  $X_S$ , it is possible to include a sufficient number of nonlinear terms of  $X_S$  to achieve a good approximation of  $\mathbb{E}[Y|X_S]$  under appropriate regularity assumptions. Nevertheless, identifying an optimal set of such regressors is often problem-dependent, and is not in the scope of this paper.

Assumption 2.2 is more difficult to inspect in practice. In the context of PDE simulations, if we take  $Y$  and  $X_S$  to be a collection of functions of discrete solutions associated with, e.g., spatial refinement of a mesh, then  $\varepsilon_S$  is unlikely to be independent of  $X_S$ . However, the magnitude of  $\varepsilon_S$  is often relatively small thus exerting little impact on the practical usage of the developed procedure in Section 5.2. In general, one can adopt a less restrictive relationship between  $\varepsilon_S$  and  $X_S$  (e.g.  $\varepsilon_S$  and  $X_S$  are uncorrelated) but that would make both simulating  $Y$  and the analysis intractable since correlation gives information only about second moments and not the full distribution. We provide a careful empirical examination of both Assumptions 2.1 and 2.2 in practice in the numerical results section.

The remaining assumptions we make are relatively mild, holding for a large class of problems.

**Assumption 2.3.** *The uncentered second-moment matrix  $\Lambda := \mathbb{E}[XX^\top]$  is invertible.*

Assumption 2.3 is equivalent to the rather mild requirements that the second moments of  $X$  exist, and that there is an event set of positive probability over which the components of  $X$  are linearly independent. Pairing this with the Cauchy interlacing theorem for eigenvalues results in the actual technical statement that we utilize:

Assumption 2.3  $\iff$  For every  $S \subseteq [n]$ ,  $\Lambda_S := \mathbb{E}[X_S X_S^\top]$  is invertible.

Our final two assumptions are more technical but are easily satisfied in practice, which we will motivate after presenting the assumptions. First, a random vector  $Z \in \mathbb{R}^k$  is a *jointly sub-exponential* random vector, if  $\sup_{\|a\|_2=1} \|a^\top Z\|_{\psi_1} < \infty$ , where  $\|\cdot\|_{\psi_1}$  is the 1-Orlicz norm [31]. Equivalently,  $Z$  is jointly sub-exponential if there exists  $C > 0$  such that

$$\sup_{\|a\|_2=1} \mathbb{P}(|a^\top Z| > z) \leq 2 \exp(-z/C) \quad \forall z \geq 0. \quad (2)$$

Our fourth assumption can now be articulated.

**Assumption 2.4.**  *$X = (X_1, \dots, X_n)$  is jointly sub-exponential in the sense of (2). Moreover, denoting by  $F_a(x)$  the CDF of  $a^\top X$  for  $a \in \mathbb{R}^n$ , we assume that  $C_{Lip} := \sup_{\|a\|_2=1} \|F_a\|_{Lip} < \infty$ .*

We will motivate the sub-exponential assumption later in this section. The  $a$ -uniform Lipschitz assumption precludes, for example, cases when  $F_X$ , the multivariate distribution function of  $X$ , has jump discontinuities in any direction  $a$ . If  $F_X$  is Lipschitz, then this implies the weaker Lipschitz condition above. Our final assumption is as follows:

**Assumption 2.5.** *There is a universal constant  $C$  such that, for any  $S \subseteq [n]$ , the noise satisfies  $|\varepsilon_S| \leq C$  almost surely.*

One specialized situation under which the sub-exponential part in Assumption 2.4 and Assumption 2.5 are satisfied is when  $Y$  and all the components of  $X$  are uniformly bounded by a constant almost surely. In this case, setting  $\beta_S$  in (1) as any finite coefficients (e.g. those corresponding to the best linear unbiased estimator) ensures that  $\varepsilon_S := Y - X_S^\top \beta_S$  is also bounded uniformly by a constant, hence satisfying Assumption 2.5. This boundedness condition immediately implies that  $X$  is jointly sub-exponential, hence satisfying part of Assumption 2.4.

The specialized situation above, i.e., the conditions that  $X$  and  $Y$  are bounded by a constant almost surely, is quite natural in the setting of well-posed parametric PDEs. For simplicity, we describe the argument for  $Y$  only. Let  $\mathcal{S} : \mathcal{K} \rightarrow V$  be the solution operator of a parametric PDE, mapping values on some compact subset  $\mathcal{K}$  of a Banach space to another Banach space  $V$ . In a prototypical elliptic PDE setup,  $\mathcal{K}$  is a set of diffusion coefficient functions (say uniformly bounded above and below), and  $V$  is the Sobolev space  $H_0^1(D)$  for some spatial domain  $D$ . We define the QoI  $Y : \mathcal{K} \rightarrow \mathbb{R}$  as  $Y(k) = \mathcal{L}(\mathcal{S}(k))$ ,  $k \in \mathcal{K}$ , where  $\mathcal{L} \in V^*$  is an element in the dual space of  $V$ . The randomness in  $Y$  arises from placing a probability measure on  $\mathcal{K}$ . A common assumption in these settings is the uniform boundedness assumption [6, Equation (2.20)] for elliptic problems stating that there is a constant  $C'$  such that  $\|\mathcal{S}(k)\|_V \leq C'$  for every  $k \in \mathcal{K}$ . Pairing this with the fact that  $\mathcal{L}$  is bounded on  $V$  implies the desired almost sure uniform bound  $|Y| \leq \|\mathcal{L}\|_{V^*} C' =: C$ . Similar arguments hold in the case when  $Y$  is computed from a well-posed discretization of  $\mathcal{S}$  and  $\mathcal{L}$ .

In summary, Assumptions 2.1 and 2.2 are a crucial necessity for our theoretical results, but we show in practice in the numerical results section how violation of either can be ameliorated. The remaining Assumptions 2.3, 2.4, and 2.5 hold in many reasonable practical settings as described above.

## 2.4 Algorithm overview

In this section, we summarize the high-level algorithmic idea of this paper. For clarity of presentation, first suppose that (1) holds and  $(\beta_S, \varepsilon_S)$  are known *a priori*. Then for some fixed  $S \subseteq [n]$ , one may

expend a given and fixed total budget  $B > 0$  to sample  $X_S$  and build an empirical CDF estimator for  $F_Y$ :

$$\frac{1}{N} \sum_{i \in [N]} \mathbf{1}_{Z_i \leq y}, \quad (3)$$

where  $Z_i$  are i.i.d. random variables sampled according to  $X_S^\top \beta_S + \varepsilon_S$ . That is, one first samples  $X_S$  to compute  $X_S^\top \beta_S$ , and then augments it with an independent noise  $\varepsilon_S$ . Of course, the number of affordable samples  $N$  under this model is,

$$N = \left\lfloor \frac{B}{c_{\text{ept}}(S)} \right\rfloor \quad c_{\text{ept}}(S) = \sum_{i \in S} c_i. \quad (4)$$

As opposed to the direct construction of the empirical CDF from i.i.d. samples of  $Y$ , the emulator (3) will admit a much larger sampling rate under a fixed budget if  $c_{\text{ept}}(S) \ll c_0$ , which can substantially accelerate convergence. Of course, our presentation above elides the real practical challenges of this approach: (i) oracle knowledge of  $\beta_S$  is unavailable, (ii) the distribution of  $\varepsilon_S$  is unknown, and (iii) the “optimal” model subset  $S$  is unknown. The algorithm we develop overcomes these challenges and accomplishes the construction described above via two phases:

(Exploration) A portion of the budget  $B$  is expended to explore  $X$  and  $Y$  in order to learn about the regression coefficients  $\beta_S$  and noise  $\varepsilon_S$ , and also to analyze cost-benefit tradeoffs for each model  $S$ . The budget expended in this phase is adaptively determined.

(Exploitation) The remaining budget is expended over a computed optimal model  $S^*$  identified at the end of Exploration. The distributional estimator for  $Y$  is computed as the empirical CDF of  $X_{S^*}^\top \beta_{S^*} + \varepsilon_{S^*}$  built over the exploitation samples. We simulate  $\varepsilon_{S^*}$  by bootstrapping the empirical residuals  $Y - X_{S^*}^\top \beta_{S^*}$  from exploration data; see (22) and (25) for a concrete description.

In philosophy, some of the challenges described above (ignorance of  $S$ ,  $\beta_S$ ) are addressed by the bandit-learning approach proposed in [34]. Indeed, our exploration-exploitation metamodel here is essentially identical to [34]. However, one main advance of this paper is the development of a new loss function to guarantee convergence in a probability metric, whereas direct usage of the algorithm in [34] guarantees only convergence of an estimator for the single-statistic mean.

## 2.5 Main results

We state our main results in more technical terms compared to the descriptions in Section 1.1:

**Algorithm 1:** We develop an exploration-exploitation algorithm based on ideas from bandit learning that produces an estimator for the full distribution of  $Y$ . The algorithm requires as input only the ability to query the models  $(X, Y)$ , knowledge of the model costs  $c_i$ ,  $i \in \{0\} \cup [n]$ , and an available computational budget  $B$ .

**Theorem 5.2:** Using Algorithm 1, we show that almost surely for a large budget  $B$ , the computationally chosen model  $S^*$  coincides with the optimal model that would be identified with access to oracle knowledge. We show similar asymptotic optimality of the resources expended during the exploration stage.

**Corollary 5.5:** We show that Algorithm 1 produces a computational emulator for the distribution of  $Y$  that budget-asymptotically converges almost surely to the true distribution in the 1-Wasserstein metric.

Finally, our numerical results showcase how Algorithm 1 outperforms the only competitor we are aware of in this general non-hierarchical multifidelity context: the empirical CDF estimator for  $Y$  by expending the full budget  $B$  on samples of  $Y$ .

### 3 Distribution metrics and related results

Before directly addressing the multifidelity problem, this section provides some necessary discussion regarding distribution metrics, empirical measures, and error bounds.

#### 3.1 Distance between distributions

The goal of distribution learning is to approximate a probability measure on  $\mathbb{R}$ . In this section, we discuss metrics to measure the discrepancy between  $F_Y$  and an estimated distribution. For two Borel probability measures  $\mu, \nu$  on  $\mathbb{R}$ , numerous metrics are available to measure discrepancy [8], such as the Kolmogorov distance, Wasserstein distances, Kullback–Leibler (KL) divergence, etc. In this article, we will focus on Wasserstein metrics, which are defined below.

**Definition 3.1** ( $p$ -th Wasserstein distance). *Let  $1 \leq p < \infty$ . The  $p$ -th Wasserstein distance between two Borel probability measures  $\mu, \nu$  on  $\mathbb{R}$  is defined as*

$$W_p(\mu, \nu) = \inf_{\pi} \mathbb{E}_{\pi} [|x - y|^p]^{1/p},$$

where infimum is taken over all Borel probability measures  $\pi$  on  $\mathbb{R}^2$  with marginals satisfying  $\pi_x \equiv \mu, \pi_y \equiv \nu$ .

Intuitively,  $W_p(\mu, \nu)$  corresponds to the minimal amount of work needed to transform  $\mu$  to  $\nu$ , with the cost function given by the  $p$ -th power of the moving distance. Hence it is frequently labeled as the optimal transport distance. A comprehensive study of the subject can be found in [32]. In the following discussion we are mostly concerned with  $p = 1$ .

Wasserstein distances are hard to compute in general. But when the metric space is the real line equipped with the Borel algebra, explicit formulas exist using inverse CDFs [4]:

$$W_p(\mu, \nu) = \int_0^1 |F_{\mu}^{-1}(t) - F_{\nu}^{-1}(t)|^p dt,$$

where  $F_{\mu}^{-1}(t) := \inf\{x \in \mathbb{R} : F_{\mu}(x) \geq t\}$  is the inverse CDF of  $\mu$ . When  $p = 1$ ,  $W_1(\mu, \nu)$  is the  $L^1(\mathbb{R})$ -norm of  $F_{\mu}(t) - F_{\nu}(t)$ :

$$W_1(\mu, \nu) = \int_0^1 |F_{\mu}^{-1}(t) - F_{\nu}^{-1}(t)| dt = \int_{\mathbb{R}} |F_{\mu}(t) - F_{\nu}(t)| dt. \quad (5)$$

A classical result in optimal transport is the Kantorovich–Rubinstein duality [32], which provides an alternative characterization for  $W_1(\mu, \nu)$  using Lipschitz test functions:

$$W_1(\mu, \nu) = \sup_{\|f\|_{\text{Lip}} \leq 1} \left| \int f d\mu - \int f d\nu \right|, \quad (6)$$

where  $\|\cdot\|_{\text{Lip}}$  is the Lipschitz constant. As a consequence,  $W_1(\mu, \nu)$  provides an upper bound for the difference between all linear functionals of  $\mu$  and  $\nu$  with uniformly bounded Lipschitz coefficients.

Moreover, assuming  $F_{\mu}$  has a uniformly bounded density  $f_{\mu}$ , one can bound the Kolmogorov distance between  $\mu$  and  $\nu$  by  $W_1(\mu, \nu)$  using (6) [5]:

$$d_K(\mu, \nu) := \sup_{x \in \mathbb{R}} |F_{\mu}(x) - F_{\nu}(x)| \leq 2\sqrt{\|f_{\mu}\|_{\infty} W_1(\mu, \nu)}. \quad (7)$$

In the situations of this article, at least one of the measures under comparison has a uniformly bounded density under Assumption 2.4. Thus, in what follows we will work with the Wasserstein metrics, i.e., the  $W_1$ -metric.

### 3.2 Convergence of empirical measures

Our analysis in Section 4 relies on sharp convergence rates of one-dimensional empirical measures under the mean  $W_p$  metric. Given i.i.d. samples of  $\mu : Z_1, \dots, Z_N$ , the associated empirical measure is

$$\mu_N = \frac{1}{N} \sum_{i \in [N]} \delta_{Z_i}, \quad (8)$$

where  $\delta_x$  denotes the Dirac mass at  $x$ . The question is to quantify the average convergence rate for  $\mathbb{E}[W_p^p(\mu_N, \mu)]^{1/p}$  at fixed  $N$ . A comprehensive analysis quantifying  $\mathbb{E}[W_p^p(\mu_N, \mu)]^{1/p}$  for fixed  $N$  can be found in [2]. Here we only collect relevant results to be used later.

**Lemma 3.2** ([2]). *Suppose  $\mu$  is a Borel probability measure on  $\mathbb{R}$  with finite  $(2+\delta)$ -th moment for some  $\delta > 0$ , i.e.,  $\int |x|^{2+\delta} d\mu(x) < \infty$ . Let  $\mu_N$  be the empirical measure defined in (8). Then for every  $N \geq 1$ ,*

$$\frac{J_0(\mu)}{\sqrt{2}} \frac{1}{\sqrt{N}} \leq \mathbb{E}[W_1(\mu_N, \mu)] \leq \frac{J_1(\mu)}{\sqrt{N}}, \quad (9)$$

where

$$J_0(\mu) := \int_{\mathbb{R}} F_{\mu}(x)(1 - F_{\mu}(x))dx \quad J_1(\mu) := \int_{\mathbb{R}} \sqrt{F_{\mu}(x)(1 - F_{\mu}(x))}dx. \quad (10)$$

The moment assumption on  $\mu$  ensures that  $J_1(\mu) < \infty$  so that the upper bound in (9) is nonvacuous. We prove next that (9) is tight when  $\mu$  is fast-decaying and has a bounded density, i.e.,  $J_0(\mu)$  and  $J_1(\mu)$  are of a similar order.

**Lemma 3.3.** *If there exist  $z \in \mathbb{R}$  and constants  $C > 1$  and  $r \geq 4$  such that*

$$\min\{F_{\mu}(x+z), 1 - F_{\mu}(x+z)\} \leq C|x|^{-r} \quad x \neq 0 \quad (11)$$

$$|F'_{\mu}(x)| \leq \sqrt{C} \quad (12)$$

then  $J_1(\mu) \leq 21CJ_0(\mu)$ .

*Proof.* Since  $J_0$  and  $J_1$  are invariant under shifts, without loss of generality we assume  $z = 0$ . We first upper bound  $J_1(F_{\mu})$  as follows:

$$J_1(F_{\mu}) = \int_{\mathbb{R}} \sqrt{F_{\mu}(x)(1 - F_{\mu}(x))}dx \stackrel{(11)}{\leq} \int_{\mathbb{R}} \sqrt{C} \min\left\{\frac{1}{2}, |x|^{-r/2}\right\} dx \leq 2\sqrt{C} \left(\frac{1}{2} + \frac{2}{r-2}\right) \leq 2\sqrt{C}.$$

On the other hand, for any  $0 < \delta < 1$ ,

$$\begin{aligned} J_0(F_{\mu}) &= \int_{\mathbb{R}} F_{\mu}(x)(1 - F_{\mu}(x))dx \geq \int_{\delta \leq F_{\mu}(x) \leq 1-\delta} F_{\mu}(x)(1 - F_{\mu}(x))dx \\ &\geq \int_{\delta \leq F_{\mu}(x) \leq 1-\delta} \delta(1-\delta)dx \stackrel{(12)}{\geq} \frac{\delta(1-\delta)(1-2\delta)}{\sqrt{C}}. \end{aligned}$$

Taking  $\delta = \frac{1}{2}(1 - \frac{1}{\sqrt{3}})$  yields  $J_0(F_{\mu}) > 0.096/\sqrt{C}$ . This combined with the upper bound on  $J_1(F_{\mu})$  proves the desired result.  $\square$

Similar nonasymptotic results for  $\mathbb{E}[W_p^p(\mu_N, \mu)]^{1/p}$  when  $p > 1$  can be found in [2, Corollary 5.10]. In those cases, the sufficient and necessary condition to achieve the optimal convergence rate has a more subtle dependence on the moments of a distribution [2, Corollary 6.14].

In later sections, we need to compare  $J_1$  of sums of independent random variables. The following notion is useful for the analysis:

**Definition 3.4.** Let  $\mathcal{F}(\mathbb{R}^d)$  be the space of uniformly bounded functions on  $\mathbb{R}^d$ . For every  $r \geq 0$ , the  $r$ -local maximum operator  $M_r : \mathcal{F}(\mathbb{R}^d) \rightarrow \mathcal{F}(\mathbb{R}^d)$  is defined as follows: For any  $f \in \mathcal{F}(\mathbb{R}^d)$ ,

$$M_r[f](x) = \sup_{z \in B(x, r)} f(z) \quad B(x, r) = \{z : \|x - z\|_2 \leq r\}.$$

**Lemma 3.5.** Let  $R_1 = Z + U$  and  $R_2 = Z + V$ , where  $Z$  and  $U/V$  are independent random variables with bounded  $(2 + \delta)$ -th moments for some  $\delta > 0$ . Suppose  $\text{supp}(Z) \subseteq [-r, r] \subset \mathbb{R}$  for some  $r > 0$ . For a random variable  $X$ , denote by  $F_X(x)$  the CDF of  $X$ . Then,

$$|J_1(F_{R_1}) - J_1(F_{R_2})| \leq \int_{\mathbb{R}} M_r[|F_U - F_V|]^{1/2}(x) dx,$$

where  $M_r$  is the  $r$ -local maximum operator in Definition 3.4.

*Proof.* By definition,

$$\begin{aligned} |J_1(F_{R_1}) - J_1(F_{R_2})| &\leq \int_{\mathbb{R}} \left| \sqrt{F_{R_1}(x)(1 - F_{R_1}(x))} - \sqrt{F_{R_2}(x)(1 - F_{R_2}(x))} \right| dx \\ &\leq \int_{\mathbb{R}} \sqrt{|F_{R_1}(x)(1 - F_{R_1}(x)) - F_{R_2}(x)(1 - F_{R_2}(x))|} dx \\ &= \int_{\mathbb{R}} \sqrt{|F_{R_1}(x) - F_{R_2}(x)| |1 - (F_{R_1}(x) + F_{R_2}(x))|} dx \\ &\leq \int_{\mathbb{R}} \sqrt{|F_{R_1}(x) - F_{R_2}(x)|} dx. \end{aligned} \tag{13}$$

Since  $Z$  and  $U/V$  are independent,

$$F_{R_1}(x) = \int_{\mathbb{R}} F_U(x - z) dF_X(z) = \int_{\mathbb{R}} F_X(x - z) dF_U(z) \tag{14}$$

$$F_{R_2}(x) = \int_{\mathbb{R}} F_V(x - z) dF_X(z) = \int_{\mathbb{R}} F_X(x - z) dF_V(z). \tag{15}$$

Substituting (14) and (15) into (13) yields the desired result:

$$\begin{aligned} \int_{\mathbb{R}} \sqrt{|F_{R_1}(x) - F_{R_2}(x)|} dx &= \int_{\mathbb{R}} \sqrt{\int_{\mathbb{R}} |F_U(x - z) - F_V(x - z)| dF_X(z)} dx \\ &\leq \int_{\mathbb{R}} M_r[|F_U - F_V|]^{1/2}(x) dx. \end{aligned}$$

□

### 3.3 A CDF estimate for marginals of a random vector

We have seen in (5) that if two one-dimensional random variables have “similar” CDFs, then they are close in the  $W_1$ -metric. In this section, we prove a type of converse of this statement. In particular, we show that two close marginals of a jointly sub-exponential vector have similar CDFs.

We recall that  $X \in \mathbb{R}^k$  is a jointly sub-exponential random variable if it satisfies (2). A useful property of sub-exponential random variables is that their maximum grows less than any polynomial rate:

**Lemma 3.6.** Let  $\{U_n\}_{n \in \mathbb{N}}$  be a sequence of sub-exponential random variables with uniformly bounded sub-exponential norm, i.e., there exists a constant  $C > 0$  such that  $\mathbb{P}(|U_n| > x) \leq 2 \exp(-x/C)$  for all  $n$ . Denote by  $U_n^* = \max_{i \in [n]} |U_i|$ . Then,  $U_n^* \lesssim \mathcal{O}(\log n)$  a.s., where the implicit constant is realization-dependent.

*Proof.* For every  $n$  and  $t_n > 0$ , it follows from a union bound that

$$\mathbb{P}(U_n^* \geq t_n) \leq \sum_{i \in [n]} \mathbb{P}(|U_i| \geq t) \leq 2n \exp\left(-\frac{t_n}{C}\right).$$

Setting  $t_n = 3C \log n$  and applying the Borel–Cantelli lemma finishes the proof.  $\square$

Now consider two marginals  $Z_1, Z_2$  of  $X$ :

$$Z_1 = a_1^\top X \quad Z_2 = a_2^\top X, \quad (16)$$

where  $a_1, a_2 \in \mathbb{R}^k$ . Denote by  $F_1(x)$  and  $F_2(x)$  the CDFs of  $Z_1$  and  $Z_2$ , respectively. If  $\|a_1 - a_2\|_2$  is small (which implies that  $W_1(F_1, F_2)$  is small by Cauchy-Schwarz inequality), then  $F_1(x) \approx F_2(x)$  for every fixed  $x \in \mathbb{R}$ . Since (2) ensures that the CDFs of both  $Z_1$  and  $Z_2$  quickly decay to zero as  $|x| \rightarrow \infty$ , the distance between  $F_1(x)$  and  $F_2(x)$  can be bounded in a uniform sense.

**Lemma 3.7.** *Let  $X \in \mathbb{R}^k$  be a sub-exponential random vector satisfying (2) for some constant  $C = C_1 > 0$ , and  $Z_1, Z_2$  be the marginals defined in (16). Denote by*

$$\delta := \|a_1 - a_2\|_2 \quad C_2 := \sqrt{\|a_1\|_2^2 + \|a_2\|_2^2} < \infty.$$

Suppose the CDFs of  $Z_1$  and  $Z_2$ ,  $F_1(x)$  and  $F_2(x)$ , are globally Lipschitz continuous; i.e.,

$$\max\{\|F_1\|_{Lip}, \|F_2\|_{Lip}\} \leq C_3$$

for some  $C_3 < \infty$ . Then, for any  $p \in (0, \infty)$ , there exists a constant  $K > 0$  that depends on  $C_1, C_2, C_3$  and  $p$  such that the following bounds hold true:

$$\|M_r[|F_1 - F_2|]\|_{L^p(\mathbb{R})}^p \leq \begin{cases} K(1+2r)(\delta + \delta^{5p/6}) \log\left(\frac{1}{\delta}\right) & 0 \leq \delta \leq 1/2 \\ 8\left(r + \frac{2^p C_1 C_2}{p}\right)\delta & \delta \geq 1/2 \end{cases} \quad (17)$$

where  $M_r$  is the  $r$ -local maximum operator in Definition 3.4.

*Proof.* Without loss of generality, we assume  $F_1(x) \geq F_2(x)$  for all  $x \in \mathbb{R}$ , otherwise, one can divide into two regimes and discuss them separately.

We first give a bound for  $M_r[|F_1 - F_2|](x)$  when  $|x| \geq 2r$ :

$$\begin{aligned} M_r[|F_1 - F_2|](x) &= \sup_{z \in B(x, r)} (F_1(z) - F_2(z)) \leq \sup_{z \in B(x, r)} \min\{F_1(z), 1 - F_2(z)\} \\ &\leq 2 \sup_{z \in B(x, r)} \exp\left(-\frac{|z|}{C_1 C_2}\right) \\ &\leq 2 \exp\left(-\frac{|x|}{2C_1 C_2}\right). \end{aligned} \quad (18)$$

When  $\delta \geq 1/2$ , since  $M_r[|F_1 - F_2|](x) \leq 1$ , we have

$$\begin{aligned} \|M_r[|F_1 - F_2|]\|_{L^p(\mathbb{R})}^p &\leq \int_{|x| < 2r} 1 \, dx + \int_{|x| \geq 2r} 2^p \exp\left(-\frac{p|x|}{2C_1 C_2}\right) \, dx \\ &\leq 4r + \frac{2^{p+2} C_1 C_2}{p} \\ &\leq 8\left(r + \frac{2^p C_1 C_2}{p}\right)\delta. \end{aligned}$$

To obtain the bound when  $\delta \leq 1/2$ , note that for  $t > 0$ ,  $\{Z_1 \leq x\} \subseteq \{Z_2 \leq x + t\} \cup \{Z_2 - Z_1 > t\}$ , which implies

$$M_r[|F_1 - F_2|](x) \leq \sup_{z \in B(x, r)} (F_2(z + t) - F_2(z) + \mathbb{P}(Z_2 - Z_1 > t)) \leq C_3 t + 2 \exp\left(-\frac{t}{C_1 \delta}\right), \quad (19)$$

where the first inequality follows from applying (2) to the normalized sub-exponential random variable  $(Z_1 - Z_2)/\delta$ . Setting  $t = C_1 \delta \log(1/\delta)$  yields

$$M_r[|F_1 - F_2|](x) \leq C_1 C_3 \delta \log\left(\frac{1}{\delta}\right) + 2\delta \leq (2 + C_1 C_3) \delta^{5/6},$$

where we used the facts  $\log(1/x) \leq x^{-1/6}$  and  $x < x^{5/6}$  when  $x \leq 1/2$ . Combining estimates (18) and (19), we can bound  $\|M_r[|F_1 - F_2|]\|_{L^p(\mathbb{R})}^p$  as follows: Fixing

$$T = \max\left\{2r, \frac{2C_1 C_2 \log(1/\delta)}{p}\right\},$$

we compute

$$\begin{aligned} \|M_r[|F_1 - F_2|]\|_{L^p(\mathbb{R})}^p &= \int_{|x| < T} (2 + C_1 C_3)^p \delta^{5p/6} dx + \int_{|x| \geq T} 2^p \exp\left(-\frac{p|x|}{2C_1 C_2}\right) dx \\ &\leq 2(2 + C_1 C_3)^p \delta^{5p/6} \left(2r + \frac{2C_1 C_2 \log(1/\delta)}{p}\right) + \frac{2^{p+2} C_1 C_2 \delta}{p} \\ &\lesssim (1 + 2r) \left(\delta + \delta^{5p/6}\right) \log\left(\frac{1}{\delta}\right). \end{aligned} \quad (20)$$

□

## 4 Empirical CDF estimator under linear regression

In this section, we develop a rigorous framework for estimation of  $F_Y(y)$  under (1). Our framework follows the overall idea presented in Section 2.4.

### 4.1 Exploration and exploitation

In this section, we introduce an exploration-exploitation strategy following the ideas in [34]. In (1), since neither the best parametric model  $S$  nor the corresponding coefficients  $\beta_S$  is known, it is necessary to expend some effort (budget) to decide on  $S$  before committing to (3). We split a given total budget  $B > 0$  into two parts, one for *exploration* and the other for *exploitation*. In the exploration stage, we collect  $m$  independent joint samples of all models  $(Y, X)$  to estimate  $\beta_S$  and  $F_{\varepsilon_S}$  (i.e. the CDF of  $\varepsilon_S$ ) for every  $S \subseteq [n]$ , based on which we decide the best model for exploitation. In the exploitation stage, we follow the decision made in the exploration phase and use (3) to construct an estimator for  $F_Y$  by plugging in the estimated coefficients. Denote

$$\begin{aligned} &\text{(exploration samples)} \quad X_{\text{epr}, \ell} := (1, X_{1, \ell}, \dots, X_{n, \ell}, Y_\ell)^\top \quad \ell \in [m] \\ &\text{(design matrix for } S) \quad Z_S := \begin{bmatrix} 1, (X_{i,1})_{i \in S} \\ \vdots \\ 1, (X_{i,m})_{i \in S} \end{bmatrix} \in \mathbb{R}^{m \times (s+1)} \\ &\text{(exploration responses)} \quad Y_{\text{epr}} := (Y_1, \dots, Y_m)^\top, \end{aligned}$$

where  $\ell$  is the sampling index, and  $N_S = \lfloor (B - c_{\text{epr}} m) / c_{\text{ept}}(S) \rfloor$  is the number of affordable samples for exploiting  $S$ , where  $c_{\text{epr}} = \sum_{i=0}^n c_i$  is the cost for an exploration sample. For  $m > |S| + 1$ ,  $\beta_S$  can be

estimated using standard least-squares. Assuming that  $Z_S$  has full column rank, the estimator for  $\beta_S$  is given by

$$\hat{\beta}_S = Z_S^\dagger Y_{\text{epr}} = (Z_S^\top Z_S)^{-1} Z_S^\top Y_{\text{epr}}, \quad (21)$$

The full distribution of  $\varepsilon_S$  can be estimated using the following plug-in empirical CDF estimator based on exploration residuals:

$$\hat{F}_{\varepsilon_S}(y) = \frac{1}{m} \sum_{\ell \in [m]} \mathbf{1}_{Y_\ell - X_{S,\ell}^\top \hat{\beta}_S \leq y}. \quad (22)$$

The resulting empirical estimator from model  $S$  is

$$\hat{F}_{Y,S}(y) := \frac{1}{N_S} \sum_{j \in [N_S]} \mathbf{1}_{A_j \leq y}, \quad (23)$$

where  $A_j$  are i.i.d. random variables sampled according to

$$Y' := X_S^\top \hat{\beta}_S + \hat{\varepsilon}_S, \quad (24)$$

that is, one first samples  $X_S$  to compute  $X_S^\top \hat{\beta}_S$  then adds it with an independently sampled noise from

$$\hat{\varepsilon}_S \sim \hat{F}_{\varepsilon_S}(y). \quad (25)$$

Note that samples of  $A_j$  use estimated statistics and hence do not rely on oracle information.

To measure the average quality of (23) as an estimator for  $F_Y$ , we use the following mean 1-Wasserstein distance as the *loss function*:

$$L_S(m) := \mathbb{E} [W_1(\hat{F}_{Y,S}, F_Y) | Z_S], \quad (26)$$

where the expectation averages out the *randomness in exploitation* as well as the *noise in exploration*. Note that the  $L_S(m)$  defined in (26) is still random due to the remaining randomness in  $Z_S$ , and one could alternatively define it by averaging all the randomness, but this would lead to the term  $\mathbb{E}[(Z_S^\top Z_S)^{-1}]$ , which is difficult to analyze. Thus we will pursue (26) in this article and consider the case when  $B \rightarrow \infty$ .

Obtaining exact asymptotically equivalent expressions for (26) is difficult. To find computably informative substitutes, we compute sharp estimates for an upper bound of (26) in the next section.

## 4.2 Upper bounds

We cannot directly compute  $L_S(m)$  in (26) in an algorithmic setting; our goal in this section is thus to derive a computable upper bound. We start by writing (26) in two parts using the triangle inequality:

$$L_S(m) \leq \mathbb{E} [W_1(F_Y, F_{Y'}) | Z_S] + \mathbb{E} [W_1(\hat{F}_{Y,S}, F_{Y'}) | Z_S], \quad (27)$$

where  $F_{Y'}$  is the CDF of  $Y'$ ; see (24) for the definition of  $Y'$ .

We provide some intuition for the bound (27). The first term measures the mean  $W_1$  distance between  $X_S^\top \beta_S + \varepsilon_S$  and  $X_S^\top \hat{\beta}_S + \hat{\varepsilon}_S$ , which depends on the accuracy of  $\hat{\beta}_S$  and  $\hat{\varepsilon}_S$ , hence on the exploration rate  $m$ . The second term measures the mean convergence rate of empirical measures, which depends on the exploitation rate  $N_S$ . A good exploration-exploitation strategy (i.e. determination of  $m$ ) will balance these quantities. We will now produce a computable asymptotic upper bound for (27) that can be used to find such an  $m$ .

**Lemma 4.1.** *Under Assumptions 2.1-2.5, and given any  $\delta > 0$ , it holds almost surely that for sufficiently large  $m$  and every  $B > c_{\text{epr}}m$ ,*

$$\mathbb{E}[W_1(F_Y, F_{Y'})|Z_S] \leq \frac{(2+\delta)\sqrt{s+1}\sigma_S + J_1(F_{\varepsilon_S})}{\sqrt{m}}, \quad (28a)$$

$$\mathbb{E}\left[W_1(\hat{F}_{Y,S}, F_{Y'})|Z_S\right] \leq (1+\delta)\frac{J_1(F_Y)}{\sqrt{N_S}}. \quad (28b)$$

The proof of Lemma 4.1 is quite technical and can be found in Appendix A. An immediate consequence of Lemma 4.1 is that  $L_S$  can be estimated by a more computable quantity that serves as an asymptotic upper bound. This fact will be used for algorithm design in Section 5.2.

**Theorem 4.2.** *Assume  $B > c_{\text{epr}}m$ . Under Assumptions 2.1-2.5, with probability 1,*

$$\limsup_{B,m \uparrow \infty} \frac{L_S(m)}{G_S(m)} \leq 1 \quad G_S(m) := \sqrt{\frac{k_1(S)}{m}} + \sqrt{\frac{k_2(S)}{B - c_{\text{epr}}m}}, \quad (29)$$

where

$$k_1(S) = (2\sqrt{s+1}\sigma_S + J_1(F_{\varepsilon_S}))^2 \quad k_2(S) = c_{\text{ept}}(S)J_1^2(F_Y).$$

*Proof of Theorem 4.2.* Combining (27) and Lemma 4.1 yields that for any  $\delta > 0$ ,

$$\limsup_{B,m \uparrow \infty} \frac{L_S(m)}{G_S(m)} \leq 1 + \delta \quad \text{a.s.}$$

The proof is finished by sending  $\delta \downarrow 0$ .  $\square$

**Remark 4.3.** *Theorem 4.2 does not require  $B$  and  $m$  diverge at a certain rate as long as  $B > c_{\text{epr}}m$  is satisfied to ensure  $G_S(m)$  is well-defined. This will be made clear from the proof of Lemma 4.1, where we shall see that consistency of the parameters only depends on the exploration stage.*

## 5 Algorithm

In this section, we first use the asymptotic upper bound  $G_S$  from (29) to analyze the optimal exploration rate for each exploitation choice  $S \subseteq [n]$  in an exploration-exploitation policy, which allows us to find a deterministic strategy that explores and exploits optimally. Then we propose an adaptive procedure that, along each trajectory, resembles the best exploration-exploitation policy. Integer rounding effects defining  $N_S$  are subsequently ignored to simplify analysis.

### 5.1 Optimal exploration based on $G_S$

The quantity  $G_S(m)$  defined in Theorem 4.2 provides a computable criterion to evaluate the model  $S$  as a simulator for  $Y$ . For fixed  $S$ ,  $G_S(m)$  is a strictly convex function of  $m$  in the domain, attaining its minimum at

$$m^*(S) = \frac{B}{c_{\text{epr}} + \left(\frac{c_{\text{epr}}^2 k_2(S)}{k_1(S)}\right)^{1/3}}, \quad (30)$$

with optimum value

$$G_S^* := G_S(m^*(S)) = \frac{[(c_{\text{epr}} k_1(S))^{1/3} + k_2(S)^{1/3}]^{3/2}}{\sqrt{B}} \propto \left[(c_{\text{epr}} k_1(S))^{1/3} + k_2(S)^{1/3}\right]^{3/2}. \quad (31)$$

(31) is the (asymptotic) minimum loss of exploiting model  $S$  with optimal exploration rate (30). The model with the smallest value  $G_S^*$  is considered the optimal model under the “upper bound criterion”, which is our terminology for using  $G_S$  as a criterion. We assume the optimal model is unique and denoted by  $S_{\text{opt}}$ , i.e.,

$$S_{\text{opt}} = \arg \min_{S \subseteq [n]} \left[ (c_{\text{epr}} k_1(S))^{1/3} + k_2(S)^{1/3} \right]^{3/2}. \quad (32)$$

We call the policy that spends  $m^*(S_{\text{opt}})$  rounds on exploration and then selects model  $S_{\text{opt}}$  for exploitation a *perfect exploration-exploitation policy*. This is similar to the perfect uniform exploration policies in [34] where exact asymptotics of the loss function is used for estimation of first-order statistics of QoI.

To investigate the performance of a perfect exploration-exploitation policy, we compare it to the empirical CDF estimator for  $Y$  using samples of  $Y$  only. Let  $\widehat{F}_{Y,\text{dir}}(B)$  denote the empirical CDF of  $Y$  with the whole budget devoted to sampling  $Y$ . The number of admissible samples for  $\widehat{F}_{Y,\text{dir}}(B)$  is  $B/c_0$ , which, combined with the lower bound in (9), implies the following lower bound for the mean  $W_1$  distance between  $\widehat{F}_{Y,\text{dir}}(B)$  and  $F_Y$ : For every  $B > c_0$ ,

$$\mathbb{E} \left[ W_1 \left( \widehat{F}_{Y,\text{dir}}(B), F_Y \right) \right] \geq \sqrt{\frac{c_0 J_0^2(F_Y)}{2B}}. \quad (33)$$

Under the same budget, the average  $W_1$  distance between  $F_Y$  and the estimator given by the perfect exploration-exploitation policy is asymptotically bounded by  $G_{S_{\text{opt}}}^*$ , as guaranteed by Theorem 4.2. This motivates us to introduce the ratio between the lower bound in (33) and  $G_{S_{\text{opt}}}^*$  as a measure for the efficiency of perfect exploration-exploitation policies relative to  $\widehat{F}_{Y,\text{dir}}(B)$ :

$$\begin{aligned} \frac{\sqrt{\frac{c_0 J_0^2(F_Y)}{2B}}}{G_{S_{\text{opt}}}^*} &= \sqrt{\frac{c_0 J_0^2(F_Y)}{2 \left[ (c_{\text{epr}} k_1(S_{\text{opt}}))^{1/3} + k_2(S_{\text{opt}})^{1/3} \right]^3}} \stackrel{\text{(Jensen)}}{\geq} \sqrt{\frac{c_0 J_0^2(F_Y)}{8(c_{\text{epr}} k_1(S_{\text{opt}}) + k_2(S_{\text{opt}}))}} \\ &\geq \frac{1}{\sqrt{8(\kappa_0^2 + \kappa_1^2)}} \geq \frac{1}{4 \max\{\kappa_0, \kappa_1\}}, \end{aligned} \quad (34)$$

where

$$\kappa_0 = \frac{2\sigma_{S_{\text{opt}}} \sqrt{n+1} + J_1(F_{\varepsilon_{S_{\text{opt}}}})}{J_0(F_Y)} \quad \kappa_1 = \sqrt{\frac{c_{\text{ept}}(S_{\text{opt}})}{c_0} \frac{J_1(F_Y)}{J_0(F_Y)}}.$$

The efficiency ratio (34) is large if  $\max\{\kappa_0, \kappa_1\}$  is small. Lemma 3.3 tells us that, under appropriate assumptions on  $Y$ ,  $J_1(F_Y)/J_0(F_Y)$  is of constant order. In this case,  $\max\{\kappa_0, \kappa_1\} \ll 1$  if  $\sigma_{\text{opt}}^2 \ll 1$  and  $\sqrt{c_{\text{ept}}(S_{\text{opt}})/c_0} \ll 1$ , which corresponds to the scenario where the exploration is efficient and exploitation sampling rate is large. Under such circumstances, the perfect exploration-exploitation policy is expected to demonstrate a superior performance over the empirical CDF estimator based only on the samples of  $Y$ .

## 5.2 An adaptive algorithm

Finding perfect exploration-exploitation policies requires evaluation of (30) and (31), which uses oracle information of model statistics such as  $\sigma_S^2$ ,  $J_1(F_{\varepsilon_S})$  and  $J_1(F_Y)$ . These statistics are not available in practice but can be approximated in an online fashion using exploration data. At each step  $t \geq n+2$  in exploration, we define

$$\begin{aligned} \widehat{\sigma}_S^2(t) &= \frac{1}{t-|S|-1} \left\| Y_{\text{epr}}(t) - Z_S(t) \widehat{\beta}_S(t) \right\|_2^2 & \widehat{J}_{Y,1}(t) &= J_1(F_{Y_{\text{epr}}(t)}(y)) \\ \widehat{J}_{\varepsilon_S,1}(t) &= J_1(F_{\widehat{\varepsilon}_S(t)}(y)), \end{aligned} \quad (35)$$

where the parameter  $t$  indicates that estimates/data are based on the first  $t$  rounds of exploration, and  $F_{Y_{\text{epr}}(t)}(y)$  is the empirical CDF of  $Y$  based on the exploration samples  $Y_{\text{epr}}(t)$ . Under Assumptions 2.1-2.5, the consistency of  $\hat{\sigma}_S^2(t)$  is well known in the linear model literature; see [17] for instance. The following lemma shows that both  $\hat{J}_{Y,1}(t)$  and  $\hat{J}_{\varepsilon_S,1}(t)$  are also consistent.

**Lemma 5.1.** *Under Assumptions 2.1-2.5,  $\hat{J}_{Y,1}(t) \rightarrow J_1(F_Y)$  and  $\hat{J}_{\varepsilon_S,1}(t) \rightarrow J_1(F_{\varepsilon_S})$  as  $t \rightarrow \infty$ .*

*Proof.* We only provide the proof for  $\hat{J}_{Y,1}(t) \rightarrow J_1(F_Y)$ ; the consistency of  $\hat{J}_{\varepsilon_S,1}(t)$  can be shown similarly together with the second part of the proof of Lemma 4.1.

Let  $I_t$  be the support of the  $t$  exploration samples of  $Y$  that are used to produce  $F_{Y_{\text{epr}}(t)}(y)$ . Assumptions 2.4-2.5 ensure that  $Y$  is sub-exponential. By Lemma 3.6,

$$|I_t| \lesssim \log t \quad a.s., \quad (36)$$

where the implicit constant is realization-dependent.

On the other hand, according to the Dvoretzky-Kiefer-Wolfowitz inequality [20], for every  $t \in \mathbb{N}$ , with probability at least  $1 - 2t^{-2}$ ,  $d_K(F_Y(y), F_{Y_{\text{epr}}(t)}(y)) \leq \sqrt{\log t/t}$ , where  $d_K$  is the Kolmogorov distance defined in (7). This combined with a Borel–Cantelli argument yields that

$$d_K(F_Y(y), F_{Y_{\text{epr}}(t)}) \lesssim \sqrt{\frac{\log t}{t}} \quad a.s. \quad (37)$$

We now consider a trajectory along which both (36) and (37) hold. Define  $I'_t = I_t \cup [-\log t, \log t]$ . In this case, for all large  $t$ ,

$$\begin{aligned} \left| \hat{J}_{Y,1}(t) - J_1(F_Y) \right| &\stackrel{(13)}{\leq} \int_{\mathbb{R}} \sqrt{|F_Y(y) - F_{Y_{\text{epr}}(t)}(y)|} dy \\ &\leq \int_{I'_t} \sqrt{|F_Y(y) - F_{Y_{\text{epr}}(t)}(y)|} dy + \int_{(I'_t)^c} \sqrt{|F_Y(y) - F_{Y_{\text{epr}}(t)}(y)|} dy \\ &\lesssim \log t \cdot d_K(F_Y(y), F_{Y_{\text{epr}}(t)}) + \int_{I'_t} \sqrt{|F_Y(y) - F_{Y_{\text{epr}}(t)}(y)|} dy. \end{aligned} \quad (38)$$

As  $t \rightarrow \infty$ , the first term in (38) diminishes to 0 because of (36) and (37), and the second term diminishes to 0 as a result of dominated convergence.  $\square$

Plugging (35) into  $G_S$  and (30) allows us to estimate the optimal loss for each model  $S$  at exploration step  $t$ :

$$\rho_S = \hat{G}_S(\hat{m}^*(S) \vee t; t), \quad (39)$$

where

$$\hat{k}_1(S; t) = \left( 2\sqrt{s+2}\hat{\sigma}_S(t) + \hat{J}_{\varepsilon_S,1}(t) \right)^2, \quad \hat{k}_2(S; t) = c_{\text{ept}}(S)\hat{J}_{Y,1}^2(t) \quad (40)$$

$$\hat{G}_S(m; t) = \sqrt{\frac{\hat{k}_1(S; t)}{m}} + \sqrt{\frac{\hat{k}_2(S; t)}{B - c_{\text{epr}}m}} \quad \hat{m}^*(S; t) = \frac{B}{c_{\text{epr}} + \left( \frac{c_{\text{epr}}^2 \hat{k}_2(S; t)}{\hat{k}_1(S; t)} \right)^{1/3}}. \quad (41)$$

The second argument in  $\hat{G}_S(\cdot; \cdot)$  denotes the number of samples used to estimate the parameters in (40) and (41), and the first argument denotes the function variable. (39) can be used to decide which model is optimal to exploit at time  $t$ , and the corresponding estimated optimal stopping time  $\hat{m}^*$  will indicate

if more exploration is needed. The details are given in Algorithm 1, the adaptive Explore-Then-Commit algorithm for multifidelity distribution learning (AETC-d).

---

**Algorithm 1:** AETC algorithm for multifidelity distribution learning (AETC-d)

---

**Input:** samplers for high-fidelity model  $Y$  and  $n$  low-fidelity models  $X_i$ ;  $B$ : total budget;  $c_i$ : cost parameters (i.e. the exploration cost  $c_{\text{epr}} = \sum_{i=0}^n c_i$  and the exploitation cost  $c_S = \sum_{i \in S} c_i$  for  $S \subseteq [n]$ )

**Output:** An estimate for  $F_Y$

- 1: compute the maximum exploration round  $M = \lfloor B/c_{\text{epr}} \rfloor$
- 2: **if**  $M \leq n + 2$  **then**
- 3:   report “budget is too small”
- 4: **else**
- 5:   collect  $t = (n + 2)$  samples for exploration
- 6:   **while**  $n + 2 \leq t \leq M$  **do**
- 7:     **for**  $S \subseteq [n]$  **do**
- 8:       compute  $\hat{k}_1(S; t)$  and  $\hat{k}_2(S; t)$  as (40) using  $t$  exploration samples
- 9:       compute  $\hat{m}^*(S; t)$  and  $\rho_S$  as in (41) and (39)
- 10:     **end for**
- 11:     find the estimated optimal model  $S^* = \arg \min_{S \subseteq [n]} \rho_S$
- 12:     **if**  $\hat{m}^*(S^*; t) > 2t$  **then**
- 13:       take  $t$  new exploration samples, and set  $t \leftarrow 2t$
- 14:     **else if**  $t < \hat{m}^*(S^*; t) \leq 2t$  **then**
- 15:       take  $\lceil (\hat{m}^*(S^*; t) - t)/2 \rceil$  new exploration samples, and set  $t \leftarrow \lceil \frac{t + \hat{m}^*(S^*; t)}{2} \rceil$
- 16:     **else**
- 17:       exhaust the remaining budget  $(B - c_{\text{epr}}t)$  to draw  $N_{S^*} := \lfloor (B - c_{\text{epr}}t)/c_{S^*} \rfloor$  i.i.d. joint samples of  $X_{S^*}$ ,  $\{X_{\text{ept}, S^*, j}\}_{j \in [N_{S^*}]}$
- 18:       transform  $X_{\text{ept}, S^*, j}$  to the surrogate emulator samples  $A_j := X_{\text{ept}, S^*, j}^\top \hat{\beta}_{S^*} + \hat{\varepsilon}_{S^*, j}$ , where  $X_{\text{ept}, S^*, j}$  includes the intercept term 1 in its first component,  $\hat{\beta}_{S^*}$  is the estimated least-squares coefficients based on  $t$  exploration samples as in (21), and  $\hat{\varepsilon}_{S^*, j}$  are i.i.d. random variables (independent of  $X_{\text{ept}, S^*, j}$ ) generated from the empirical noise simulator  $\hat{F}_{\varepsilon_{S^*}}$  that is learned using  $t$  exploration samples as in (22)
- 19:       compute the exploitation estimator for  $F_Y$  as in (23):  $\hat{F}_{Y, S^*}(y) = \frac{1}{N_{S^*}} \sum_{j \in [N_{S^*}]} \mathbf{1}_{A_j \leq y}$ , and set  $t \leftarrow M + 1$
- 20:     **end if**
- 21:   **end while**
- 22: **end if**

---

When  $M = \lfloor B/c_{\text{epr}} \rfloor$  is too small, Algorithm 1 terminates immediately due to a deficient budget for initial sampling in exploration. This regime is of lesser interest since the full correlation information between the high- and low-fidelity models cannot even be estimated. Note that compared to the AETC algorithm in [34], Algorithm 1 does not use regularization and takes a more aggressive exploration to accelerate computation. (The regularization parameters could be added to stabilize the algorithm when the budget is small, but this would result in choosing additional input parameters.) Our main theoretical results, asymptotic analysis of Algorithm 1, are given in the next section.

### 5.3 Asymptotic performance of the AETC-d: optimality and consistency

**Theorem 5.2.** *Let  $m(B)$  and  $S(B)$  denote the exploration rate and selected model in Algorithm 1, respectively. Under Assumptions 2.1-2.5, with probability 1,*

$$\lim_{B \uparrow \infty} \frac{m(B)}{m^*(S_{\text{opt}})} = 1 \quad \lim_{B \uparrow \infty} S(B) = S_{\text{opt}}, \quad (42)$$

where  $S_{\text{opt}}$  is defined in (32).

*Proof.* Without loss of generality we assume  $M = \lfloor B/c_{\text{epr}} \rfloor > n + 2$ . The exploration rate  $t$  grows nonlinearly with respect to the index that counts the iterations of the WHILE loop in step 6 of Algorithm 1. For convenience, 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

$$n + 2 \leq t_q \leq t_{q+1} \leq 2t_q \quad 1 \leq q < q(B). \quad (43)$$

Since the desired results in Theorem 5.2 only concern decisions made during exploration, the probability space we will be working on is the product space of all exploration samples. Under Assumptions 2.3–2.5 and using Lemma 5.1,  $k_1(S; t) > 0$  a.s. for all  $t \geq n + 2$  ( $\widehat{k}_1(S; t) = 0$  implies that the residuals in the exploration are all the same), and both estimators  $\widehat{k}_1(S)$  and  $\widehat{k}_2(S)$  are consistent for  $k_1(S; t)$  and  $k_2(S; t)$ , respectively, for all  $S \subseteq [n]$ :

$$\widehat{k}_1(S; t) = \left( 2\sqrt{s+2}\widehat{\sigma}_S(t) + \widehat{J}_{\varepsilon_S, 1}(t) \right)^2 \xrightarrow{t \rightarrow \infty} k_1(S) = \left( 2\sqrt{s+2}\sigma_S + J_1(F_{\varepsilon_S}) \right)^2 \quad (44a)$$

$$\widehat{k}_2(S; t) = c_{\text{ept}}(S)\widehat{J}_{Y, 1}^2(t) \xrightarrow{t \rightarrow \infty} k_2(S) = c_{\text{ept}}(S)J_1^2(F_Y). \quad (44b)$$

Thus,

$$\underline{\widehat{k}}_1(S) := \inf_{t \geq n+2} \widehat{k}_1(S; t) > 0 \quad \forall S \subseteq [n] \quad \text{a.s.}, \quad (45)$$

where  $\underline{\widehat{k}}_1(S)$  is realization-dependent.

We claim that, with probability 1, the adaptive exploration rate  $m(B)$  diverges, i.e.  $m(B) \rightarrow \infty$  as  $B \rightarrow \infty$ . The proof of this fact is similar to [34, Theorem 5.2], but we give the details for convenience. If  $m(B)$  did not diverge, there would exist a stochastic realization  $\omega$  along which (44a) and (44b) hold but  $m(B, \omega) < \infty$  as  $B \rightarrow \infty$ , i.e., there exists a  $M(\omega) \in \mathbb{N}$  such that  $m(B, \omega) \leq M(\omega) < \infty$  for all  $B > 0$ . Since (44b) holds for  $\omega$ , there exists  $C(\omega) < \infty$  such that

$$\sup_t \max_{S \subseteq [n]} k_2(S; t, \omega) \leq C(\omega). \quad (46)$$

The stopping rule in Algorithm 1 implies that

$$M(\omega) \geq m(B; \omega) \geq \frac{B}{c_{\text{epr}} + \left( \frac{c_{\text{epr}}^2 \widehat{k}_2(S^*; t, \omega)}{\underline{\widehat{k}}_1(S^*; t, \omega)} \right)^{1/3}} \stackrel{(45), (46)}{\geq} \frac{B}{c_{\text{epr}} + \left( \frac{c_{\text{epr}}^2 C(\omega)}{\underline{\widehat{k}}_1(S; \omega)} \right)^{1/3}} \xrightarrow{B \uparrow \infty} \infty, \quad (47)$$

where  $S^*$  in (47) is the estimated optimal model (which is random) in the  $q$ -th loop iteration. This is a contradiction, which occurs for all  $\omega$  such that (44) holds, i.e., the contradiction occurs with probability 1. Thus,  $m(B; \omega)$  diverges, proving that  $m(B) \uparrow \infty$  with probability 1.

We now consider any trajectory  $\omega$  along which (44a) and (44b) hold and  $m(B; \omega)$  diverges. We will prove that both results in (42) hold for such an  $\omega$ . Fix  $\delta < 1/2$  sufficiently small. Since  $S_{\text{opt}}$  is assumed unique, (44a) and (44b) together with a continuity argument implies that, there exists a sufficiently large  $T(\omega)$ , such that for all  $t \geq T(\omega)$ ,

$$\max_{(1-\delta)m^*(S_{\text{opt}}) \leq m \leq (1+\delta)m^*(S_{\text{opt}})} \widehat{G}_{S_{\text{opt}}}(m; t) < \min_{S \subseteq [n], S \neq S_{\text{opt}}} \min_{0 < m < B/c_{\text{epr}}} \widehat{G}_S(m; t). \quad (48)$$

$$1 - \delta \leq \frac{\widehat{m}^*(S; t)}{m^*(S; t)} \leq 1 + \delta \quad \forall S \subseteq [n]. \quad (49)$$

To see why (48) can be achieved for sufficiently small  $\delta$ , one may without loss of generality assume  $B = 1$  (which appears as a multiplicative scaling factor). Since  $S_{\text{opt}}$  is unique, the maximum of  $G_{S_{\text{opt}}}(m)$  in a small neighborhood of its minimizer,  $m^*(S_{\text{opt}})$ , is strictly smaller than  $\min_{S \subseteq [n], S \neq S_{\text{opt}}} \min_{0 < m < B/c_{\text{epr}}} G_S(m)$ .

On the other hand,  $\widehat{G}_{S_{\text{opt}}}(m; t)$  converges uniformly to  $G_{S_{\text{opt}}}(m)$  in any sufficiently small compact neighborhood around  $m^*(S_{\text{opt}})$ , and  $\min_{S \subseteq [n], S \neq S_{\text{opt}}} \min_{0 < m < B/c_{\text{epr}}} \widehat{G}_S(m; t) \rightarrow \min_{S \subseteq [n], S \neq S_{\text{opt}}} \min_{0 < m < B/c_{\text{epr}}} G_S(m)$  as a result of strong consistency. This justifies the existence of a small  $\delta$  for which (48) is achievable.

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

$$\min_{S \subseteq [n]} m^*(S) > 4T(\omega) \quad (50)$$

$$t_{q(B)} = m(B; \omega) > 4T(\omega). \quad (51)$$

Consider  $q' < q(B)$  that satisfies  $t_{q'-1} < T(\omega) \leq t_{q'}$ . Such a  $q'$  always exists due to (51), and satisfies

$$t_{q'} \stackrel{(43)}{\leq} 2t_{q'-1} < 2T(\omega) \stackrel{(50)}{\leq} \frac{1}{2} \min_{S \subseteq [n]} m^*(S) \stackrel{(49), \delta < 1/2}{\leq} \widehat{m}^*(S; t_{q'}).$$

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,

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

This, along with (48) and (49), tells us that  $S_{\text{opt}}$  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  $2t_{q'} \leq \widehat{m}^*(S_{\text{opt}}; t_{q'})$ , then

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

which implies

$$(1 - \delta)m^*(S_{\text{opt}}) \stackrel{(49)}{\leq} t_{q'+1} \vee \widehat{m}^*(S_{\text{opt}}; t_{q'+1}) \leq (1 + \delta)m^*(S_{\text{opt}}). \quad (52)$$

If  $t_{q'} \leq \widehat{m}^*(S_{\text{opt}}; t_{q'}) < 2t_{q'}$ , then

$$t_{q'+1} = \left\lceil \frac{t_{q'} + \widehat{m}^*(S_{\text{opt}}; t_{q'})}{2} \right\rceil \leq \widehat{m}^*(S_{\text{opt}}; t_{q'}) \leq (1 + \delta)m^*(S_{\text{opt}}),$$

which also implies (52). But (52) combined with (48) and (49) implies that  $S_{\text{opt}}$  is again the estimated optimal model in the  $(q'+1)$ -th loop iteration. Applying the above argument inductively proves  $S(B) = S_{\text{opt}}$ , i.e., the second statement in (42). Note (52) holds true until the algorithm terminates, which combined with the termination criteria  $t_{q(B)} \geq \widehat{m}^*(S_{\text{opt}}; t_{q(B)}) \geq (1 - \delta)m^*(S_{\text{opt}})$  implies

$$1 - \delta \leq \frac{m(B)}{m^*(S_{\text{opt}})} = \frac{t_{q(B)}}{m^*(S_{\text{opt}})} \leq 1 + \delta.$$

The first part of (42) now follows by noting that  $\delta$  can be set arbitrarily small.  $\square$

**Remark 5.3.** *As the budget goes to infinity, when looking at the exploration rate as well as the selected model for exploitation, one finds that almost every policy realization of Algorithm 1 is an asymptotically perfect exploration-exploitation policy. This establishes a trajectory-wise optimality result for the AETC-d under the upper bound criterion. Nevertheless, since Algorithm 1 determines the exploration rate  $m(B)$  in an adaptive fashion, we cannot directly conclude the optimality for the average loss as defined in (26).*

**Remark 5.4.** *By the design of Algorithm 1, for large  $B$ , the total number of loop iteration steps is  $\mathcal{O}(\log B)$  rather than  $\mathcal{O}(B)$  when increasing  $t$  step by step, which considerably increases computational efficiency when dealing with a large budget. The same idea can also be used in the AETC algorithm in [34].*

As a consequence, we have the following consistency result for the CDF estimator produced by the AETC-d algorithm:

**Corollary 5.5.** *Denote the CDF estimator produced by Algorithm 1 as  $\widehat{F}_{Y,aetcd}(B)$ . Under Assumptions 2.1-2.5, then with probability 1,*

$$\lim_{B \uparrow \infty} W_1 \left( \widehat{F}_{Y,aetcd}(B), F_Y \right) = 0.$$

*Proof.* For fixed  $B$ , let  $N(B)$  be the exploitation sampling rate in Algorithm 1, and denote the parametric model used for exploitation as  $Y'(B)$ , i.e.,

$$Y'(B) = X_{S(B)}^\top \widehat{\beta}_{S(B)} + \widehat{\varepsilon}_{S(B)}, \quad (53)$$

where the two terms on the right-hand side of (53) are independent. By Theorem 5.2, with probability 1,  $S(B) = S_{\text{opt}}$  for all sufficiently large  $B$ , and  $m(B)/m^*(S_{\text{opt}}) \rightarrow 1$ , i.e.,  $m(B) \rightarrow \infty$  and  $N(B) \rightarrow \infty$  as  $B \rightarrow \infty$ . Under Assumption 2.4, we apply (17) with  $p = 1$  and  $r = 0$  together with the strong consistency of  $\widehat{\beta}_{S_{\text{opt}}}$  [17, Theorem 1] to conclude that  $W_1(X_{S(B)}^\top \widehat{\beta}_{S(B)}, X_{S_{\text{opt}}}^\top \beta_{S_{\text{opt}}}) \rightarrow 0$  as  $B \rightarrow \infty$  almost surely. Adjusting the argument in (64), one can also show that  $W_1(\widehat{\varepsilon}_{S(B)}, \varepsilon_{S_{\text{opt}}}) \rightarrow 0$ . By the additivity of the  $W_1$  metric, we conclude  $W_1(Y'(B), Y) \rightarrow 0$ . (The randomness here only depends on exploration.)

Now fix a trajectory along which  $N(B) \rightarrow \infty$  and  $W_1(Y'(B), Y) \rightarrow 0$  as  $B \rightarrow \infty$ . Let  $\{B_k\}$  be an arbitrary sequence such that  $B_k \uparrow \infty$  as  $k \uparrow \infty$ . Since convergence in  $W_1$  implies convergence in distribution,  $\{F_{Y'(B_k)}\}$  is  $\delta$ -tight<sup>2</sup> [33, Section 17.5]. This observation, combined with the fact that  $\widehat{F}_{Y,aetcd}(B_k)$  is an empirical measure of  $Y'(B_k)$  consisting of  $N(B_k)$  samples, implies that  $\widehat{F}_{Y,aetcd}(B_k)$  converges to  $F_Y$  in distribution almost surely [1, Theorem 1]. To lift the convergence to  $W_1$ , it only remains to show  $\int |x| d\widehat{F}_{Y,aetcd}(B_k) \rightarrow \int |x| dF_Y$  as  $k \rightarrow \infty$ , which can be verified using (53) and the strong law of large numbers. The proof is finished by noting that  $\{B_k\}$  is arbitrary.  $\square$

## 6 Numerical experiments

In this section, we demonstrate the performance of the AETC-d (Algorithm 1) for multifidelity estimation of univariate distributions. We will focus on consistency, optimality of exploration rates, and model misspecification. Four methods will be considered for estimating  $F_Y(y)$ :

- (ECDF-Y): Empirical CDF estimator for  $F_Y$  based on the samples of  $Y$  only.
- (AETC-d): Algorithm 1.
- (AETC-d-no): A modification of Algorithm 1, where the noise in (23) is omitted, i.e., we set  $\widehat{\varepsilon}_S$  to 0.
- (AETC-d-q): A modification of Algorithm 1 using quantile regression; see Appendix B for a detailed description of the algorithm.

The modifications AETC-d-no and AETC-d-q are introduced to empirically investigate plausible alternatives that one may consider (which do not enjoy our theoretical guarantees).

To evaluate results, we compute and report an empirical mean  $W_1$  distance (error) between  $F_Y$  and the estimated CDF given by the algorithms over 200 samples. Since AETC-d produces random estimators due to the exploration phase, the experiment is repeated 100 times with the 5%-50%-95%-quantiles recorded to measure this extra uncertainty.

We emphasize that direct comparison with most alternative multilevel/multifidelity procedures is not possible here, as our theory and the associated empirical errors reported here are in terms of the  $W_1$  error on the full distribution of the output.

<sup>2</sup>A sequence of probability measures  $\{P_k\}$  defined on a metric space is called  $\delta$ -tight if for every  $\varepsilon > 0$ , there exist a compact measurable set  $K$  and a sequence  $\delta_k \downarrow 0$  such that  $P_k(K^{\delta_k}) > 1 - \varepsilon$  for every  $k$ , where  $K^{\delta_k} := \{x : \text{dist}(x, K) < \delta_k\}$ .

## 6.1 Ishigami function

In this example, we investigate the performance of the proposed AETC-d algorithm (Algorithm 1) on a multifidelity algebraic system consisting of Ishigami functions [14], which have been widely used as a test model for uncertainty quantification. We consider an extended version of the setup in [26]. The high-fidelity model output corresponds to the following random variable:

$$Y = \sin Z_1 + a \sin^2 Z_2 + b Z_3^4 \sin Z_1 + c \sin^3 Z_4 + d \sin^4 Z_5, \quad (54)$$

where  $a, b, c, d$  are deterministic constants, and  $Z_i, i \in [5]$  are independent random variables,

$$Z_{1,2,3,4,5} \stackrel{\text{iid}}{\sim} \text{Unif}(-\pi, \pi).$$

In the following experiments, we set  $a = 5$  and  $b = 0.1$ . For the low-fidelity models, we will consider two different scenarios with different model assumptions.

### 6.1.1 Perfect model assumptions

Let  $c = 1, d = 0.1$ . We first consider a synthetic dataset consisting of two low-fidelity models:

$$\begin{aligned} X_1 &= \sin Z_1 + a \sin^2 Z_2 + b Z_3^4 \sin Z_1 + c \sin^3 Z_4 \\ X_2 &= \sin Z_1 + a \sin^2 Z_2 + b Z_3^4 \sin Z_1. \end{aligned} \quad (55)$$

In this case, the linear model assumption (1) is satisfied, i.e.,  $\mathbb{E}[Y|X_2] = X_2 + c\mathbb{E}[\sin^3 Z_4] + d\mathbb{E}[\sin^4 Z_5]$ ,  $\mathbb{E}[Y|X_1] = \mathbb{E}[Y|X_1, X_2] = X_1 + d\mathbb{E}[\sin^4 Z_5]$ . The correlation between  $Y$  and  $X_1, X_2$  are approximately 0.999 and 0.986, respectively. The cost of sampling  $Y, X_1$  and  $X_2$  are assumed hierarchical, assigned as  $(c_0, c_1, c_2) = (1, 0.05, 0.001)$ .

We first compare the performance of AETC-d with ECDF-Y. The total budget  $B$  in the experiment ranges from 10 to  $10^5$ . The ground truth is taken as an empirical CDF of  $Y$  computed using  $10^7$  independent samples. Accuracy results for ECDF-Y and AETC-d are reported in Figure 1(a). In this example, AETC-d consistently outperforms ECDF-Y by a factor of around 4. The average estimation error decays to 0 as the budget goes to infinity, verifying the consistency result in Corollary 5.5.

To inspect the optimality of exploration rates selected by AETC-d, we fix the total budget at  $B = 10^3$  and consider a deterministic version of Algorithm 1 with fixed exploration rate  $m$ . The new algorithm collects  $m$  exploration samples to estimate the parametric coefficients and then selects  $S_{\text{opt}}$  for exploitation. In this example,  $S_{\text{opt}} = \{1\}$ , and the maximum exploration rate is  $M = \lfloor B/1.051 \rfloor = 951$ . We apply the deterministic algorithm and compute the average  $W_1$  distance between the resulting estimator and  $F_Y$  over 100 independent experiments for  $m = (10, 30, 50, 100, 150, 200, 300, 400, 500, 600)$ . According to Theorem 4.2, for fixed  $m$ , the average error of the estimator given by the deterministic algorithm is asymptotically bounded by a function of the form

$$f(m; \alpha_1, \alpha_2) = \frac{\alpha_1}{\sqrt{m}} + \frac{\alpha_2}{\sqrt{\frac{B}{c_{\text{epr}}} - m}} \quad 0 < m < \frac{B}{c_{\text{epr}}}, \quad (56)$$

for some constants  $\alpha_1, \alpha_2 > 0$ . Assuming the analysis in Theorem 4.2 is tight, we would expect the mean error of the deterministic algorithm, as a function of  $m$ , to approximately fit a curve of the form (56). (Other parametric forms may be used if one obtains a different upper bound for the loss.) We use a least-squares procedure to obtain such a curve by optimizing over  $\alpha_1, \alpha_2$ . We also run AETC-d 100 times and record the 5%-50%-95% quantiles of the adaptively determined exploration rates. The results are illustrated in Figure 1(b).

By comparing the mean errors committed by the deterministic algorithm at various fixed exploration rates, we find that the 5%-50%-95% quantiles of the exploration rate chosen by the AETC-d algorithm, 178-191-203, are larger than the optimal trade-off point, which is near  $m = 41$ . Although the difference between the AETC-d and optimal trade-offs is large, it is modest relative to the total budget  $B$ . In

particular, the relative over-exploration rate by AETC-d in this example is  $(191 - 41)/951 \approx 15.8\%$ . According to (44) in Theorem 5.2, the over-exploration rate will tend to stabilize when the budget becomes large. Such a discrepancy is not unexpected since AETC-d is constructed and optimized with respect to  $G_S$  rather than the true loss. Nevertheless, the loss *value* of the AETC-d algorithm is close to the minimum. This empirically suggests that the upper bound criterion is useful in terms of balancing the exploration and exploitation errors.

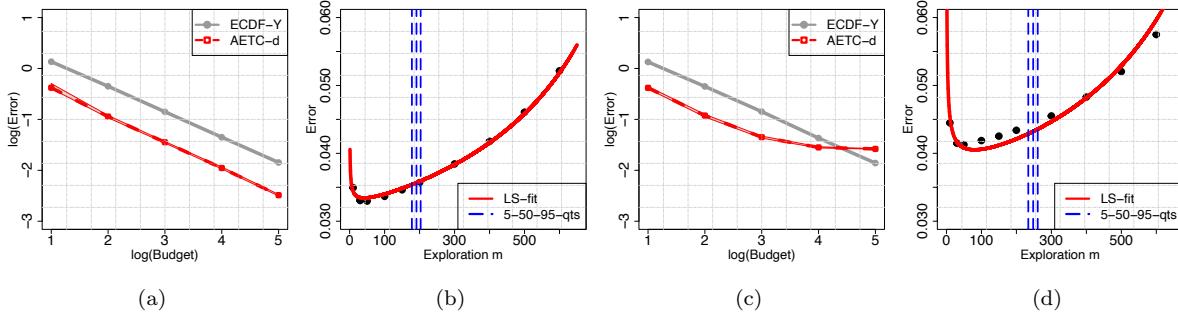


Figure 1:  $(\log_{10})$  mean  $W_1$  distance between  $F_Y$  and the estimated CDFs given by ECDF-Y and AETC-d in model (55) (a) and model (57) (c), with the 5%-50%-95% quantiles plotted for AETC-d to measure its uncertainty in the exploration phase. Mean  $W_1$  error of the deterministic algorithm at different exploration rates  $m$ , and the quantiles of the exploration rate from AETC-d in model (57) (b) and model (57) (d).

### 6.1.2 Approximate linear assumptions

Let  $c = d = 0$ . We now consider the same low-fidelity models which were used in [26] for sensitivity analysis:

$$\begin{aligned} X_1 &= \sin Z_1 + 0.95a \sin^2 Z_2 + bZ_3^4 \sin Z_1 \\ X_2 &= \sin Z_1 + 0.6a \sin^2 Z_2 + 9bZ_3^2 \sin Z_1. \end{aligned} \quad (57)$$

The correlations between  $Y$  and  $X_1, X_2$  are approximately 0.999 and 0.950, respectively. Under the current setup, the linear model assumption (1) is not satisfied. Nevertheless, the correlations between  $Y$  and  $X_1, X_2$  suggest that the relationship between the high-fidelity and low-fidelity models is approximately linear. The cost of sampling  $Y, X_1$  and  $X_2$  and the budget range is the same as in the previous case. We repeat the same experiments (both the accuracy and the optimality of exploration rates) and report the corresponding results in Figure 1(c,d).

Despite model misspecification, AETC-d still demonstrates reasonable performance until the total budget exceeds  $10^3$ . When the budget is sufficiently large, both exploration and exploitation errors are so small that model misspecification errors start to dominate. A further investigation of the model misspecification effects in this example will be carried out in Section 6.1.3. The optimal model selected by the AETC-d is  $S_{\text{opt}} = \{1, 2\}$ . By comparing the two plots in the right panel in Figure 6.1.3, we see that AETC-d again overexplores and has a close-to-optimal loss value.

### 6.1.3 Model misspecification

In this section, we identify the misspecified assumptions that prevent the experiment in Section 6.1.2 from converging. For every  $S \subseteq [n]$ , the following decomposition holds:

$$Y = \mathbb{E}[Y|X_S] + (Y - \mathbb{E}[Y|X_S]),$$

where  $\mathbb{E}[Y|X_S]$  is a measurable function of  $X_S$ , and  $Y - \mathbb{E}[Y|X_S]$  is a centered random variable that is uncorrelated with  $\mathbb{E}[Y|X_S]$ . The linear model assumption (1) (or equivalently, Assumptions 2.1-2.2) assumes the following:

- $\mathbb{E}[Y|X_S]$  a linear function of  $X_S$ ;
- $Y - \mathbb{E}[Y|X_S]$  is independent of  $\mathbb{E}[Y|X_S]$ .

Violation of either may cause the inconsistent behavior of AETC-d in this example.

To inspect the linearity assumption, we expand the feature space by incorporating nonlinear terms of the existing regressors to fit a larger linear model. To test the noise independence condition, we can either drop the noise in (23) to reduce the erroneous noisy effect<sup>3</sup>, or use a different method such as quantile regression to model the potential heteroscedasticity of the noise. A brief description of using quantile regression combined with AETC-d in our setup is given in Appendix B.

In the following experiment, we expand the original model by including higher-order terms of the existing features. We include additional regressors  $\{X_1^2, X_1^3, X_2^2, X_2^3, X_1 X_2, X_1^2 X_2, X_1 X_2^2\}$ , and call this enlarged model L. Note that in L, an exploitation model is no longer defined by the subset of regressors, but instead by the admissible  $\sigma$ -field generated by the regressors. In fact, since both  $X_1$  and  $X_1^2$  generate the same  $\sigma$ -field,  $\mathbb{E}[Y|X_1] = \mathbb{E}[Y|X_1^2]$ . As a result, the linear model assumption (1) cannot hold simultaneously for both regressors as two distinct models unless  $X_1$  is a constant a.s. For convenience, for each  $S \subseteq [n]$ , we define the model associated with  $S$  as the intersection of the linear space spanned by the existing features and the  $\sigma$ -field generated by  $X_S$ , with exploitation cost given by the total cost of the low-fidelity models in  $S$ . For instance, if  $S = \{1\}$ , then the corresponding model is  $\text{span}\{1, X_1, X_1^2, X_1^3\}$ , with exploitation cost  $c_1$  per sample. In this case, any model generating the same  $\sigma$ -field can be viewed as a sub-model under our definition. Also, in the exploitation stage of AETC-d, we consider two alternative methods to estimate  $F_Y$ : AETC-no and AETC-q. The performance of these modified estimators in both the original and the expanded models is compared in Figure 2(a-c). For ease of illustration, we only plot the median of errors instead of the 5%-50%-95% quantiles region in the 100 experiments.

Figure 2 shows that the overall performance of all the methods under comparison improves in the enlarged model L when the budget exceeds  $10^3$ . This implies that adding nonlinear terms of the existing regressors can help reduce the model misspecification effects. The improvement for AETC-d in model L diminishes when the budget reaches a new threshold near  $10^5$ , after which more regressors are needed to further increase model expressivity. In this example, noise has a less dominant role in affecting model convergence. Nevertheless, using quantile regression for reconstruction is beneficial for mitigating the noise misspecification effects.

We end this section by providing an instance of the pointwise absolute error between  $F_Y$  and the estimated CDFs given by the ECDF-Y, the AETC-d, the AETC-d (L), the AETC-d-no (L), and the AETC-d-q (L) when  $B = 10^4$ . For better visualization, we plot the distribution of these errors (i.e. CDF) rather than the pointwise curves. The result is given in Figure 2(d).

#### 6.1.4 Cost assumptions

Relative information between cross-model correlations and model costs may also impact the performance of the AETC-d algorithm. To understand the relationship between them, we fix the models and vary the (relative) cost assigned to each model. In particular, we test the following four cost conditions on both model (55) (perfect model assumptions) and model (57) (approximate model assumptions):

- $(c_0, c_1, c_2) = (1, 0.5, 0.001)$ ;
- $(c_0, c_1, c_2) = (1, 0.05, 0.001)$ ; (original choice in sections 6.1.1 and 6.1.2)

---

<sup>3</sup>When the variance ratio between  $\varepsilon_S$  and  $Y$  is small,  $Y \approx X_S^\top \beta_S$ , so that adding the noise emulator has little impact on the accuracy of the resulting estimator. When the ratio is moderate, adding  $\varepsilon_S$  as an independent component will degrade the quality of the estimator if the independence assumption is violated.

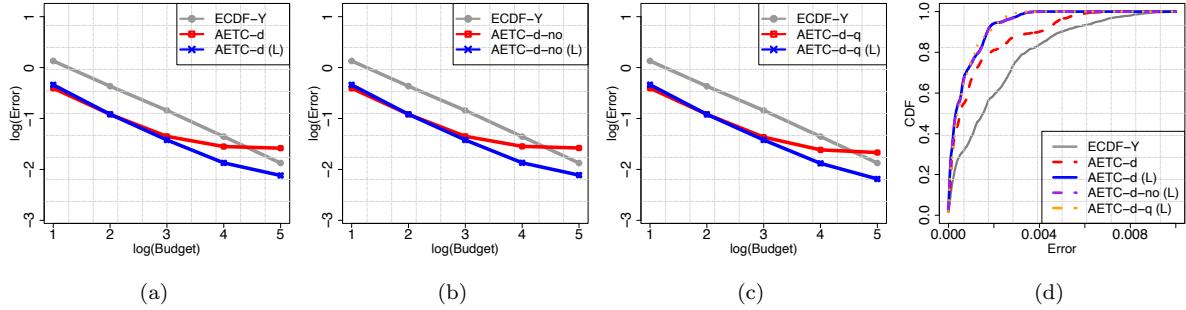


Figure 2: Median estimation errors of the estimated CDFs given by ECDF-Y and AETC-d with different exploitation strategies in both the original and the expanded model L (a)-(c). An instance of distribution of pointwise absolute errors of the CDFs given by ECDF-Y, AETC-d, AETC-d (L), AETC-d-no (L), and AETC-d-q (L) when  $B = 10^4$  (d).

- $(c_0, c_1, c_2) = (1, 0.005, 0.001)$ ;
- $(c_0, c_1, c_2) = (1, 0.001, 0.001)$ ,

For model (57), we also test the enlarged model L that includes the higher order terms; see Section 6.1.3. The results are reported in Figure 3. For ease of illustration, we only plot the error curves given by ECDF-Y and AETC-d.

In Figure 3, the limiting models selected by the AETC-d algorithm in scenarios (a)-(d) are  $\{2\}$ ,  $\{1\}$ ,  $\{1\}$ , and  $\{1\}$ , respectively. In particular, when model  $\{1\}$  becomes almost as expensive as the high-fidelity model, the AETC-d algorithm turns to use the less expensive low-fidelity model to better leverage cost and accuracy. When both low-fidelity models become equally cheap, the AETC-d algorithm selects  $\{1\}$  only. In none of the above scenarios did AETC-d choose  $\{1, 2\}$  for exploitation when  $B$  is large, which is consistent with the observation  $\mathbb{E}[Y|X_1, X_2] = \mathbb{E}[X_1]$ , i.e., adding  $X_2$  to  $X_1$  does not provide additional information in terms of understanding  $Y$ . This suggests that the AETC-d algorithm can filter redundant information among models in model selection. Finally, note that the performance of AETC-d in (a) is even slightly better than in (b) despite a larger cost for the chosen low-fidelity models. One would expect that model  $\{2\}$  is also slightly better than model  $\{1\}$  when  $(c_0, c_1, c_2) = (1, 0.05, 0.001)$ . (We used the word ‘‘slightly’’ because model  $\{2\}$  should have a similar efficiency in (a)-(d) since only the cost of  $\{1\}$  is varied and the high-fidelity cost dominates.) This shows that the  $G_S$  criteria may select a suboptimal model when the efficiency gap between different models is small. AETC-d has a better chance of finding the optimal model when the gap is large (i.e. in (a) and (d)).

For scenarios (e)-(g), the asymptotically selected models in both the original and enlarged model L are  $\{1, 2\}$ . As a result, they have the same error threshold below which the model misspecification effects start to dominate. The major difference lies in the relative efficiency of AETC-d over ECDF-Y. For instance, in (g), the cost ratio between the high- and low-fidelity models is large, and the computational gain of AETC-d is more prominent over ECDF-Y compared to (e) and (f) before the error threshold, which, as expected, arrives under a smaller budget. For scenario (h), the AETC-d algorithm chooses model  $\{1\}$  only. For model  $\{1\}$ , model misspecification occurred at a similar threshold as in  $\{1, 2\}$ . In this case, adding only quadratic and cubic terms of  $X_1$  is not sufficient to mitigate the effects (as opposed to adding seven higher order terms for  $\{1, 2\}$ ). Unfortunately, the AETC-d algorithm makes decisions assuming all the models are correct and finds the most efficient one among them.

## 6.2 Parametrized PDEs

In the last experiment, we consider a multifidelity model given by a parametric elliptic equation. The setup is taken from [34]. We consider an elliptic PDE over a square spatial domain  $D = [0, 1]^2$  that

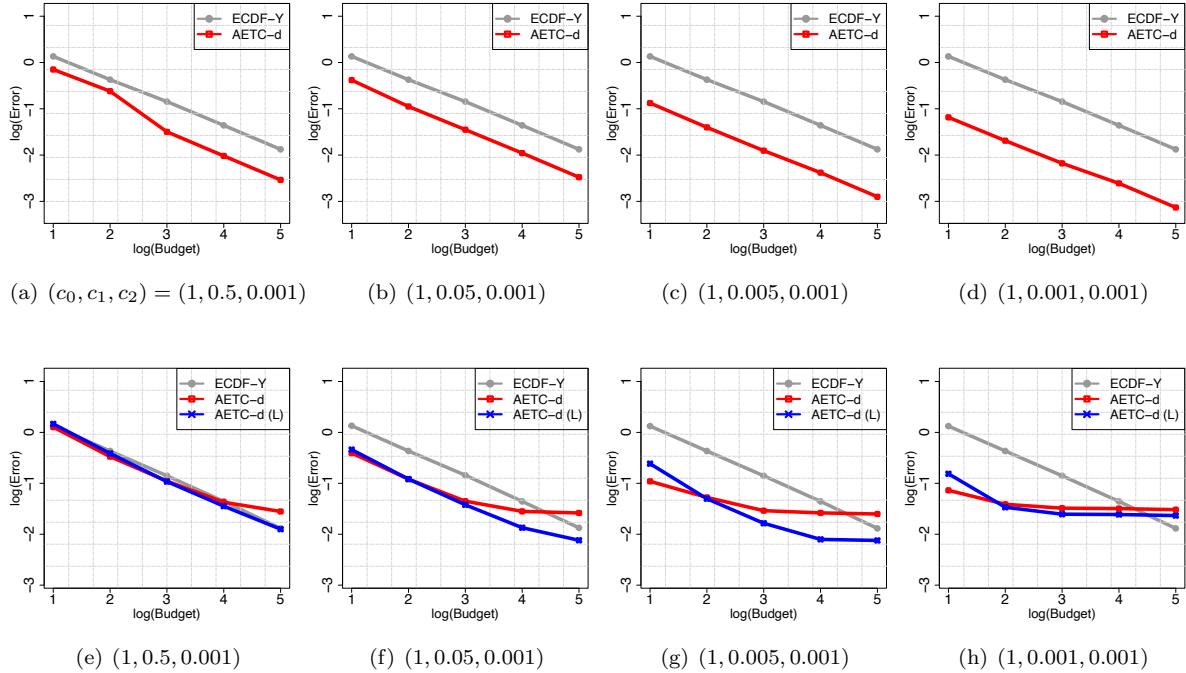


Figure 3: Median estimation errors of the estimated CDFs given by ECDF-Y and AETC-d in both model (55) (a)-(d) and model (57) (e)-(h) under different cost assumptions.

governs displacement in linear elasticity. The geometry and boundary conditions are shown in Figure 4.

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-Loëve 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 an exponential covariance kernel on  $D$ , i.e.,

$$\text{corr}(\kappa(\mathbf{p}, \mathbf{x}), \kappa(\mathbf{p}, \mathbf{y})) = \exp(-\|\mathbf{x} - \mathbf{y}\|_1/a),$$

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:

$$E := \int_D (\mathbf{u} \cdot \mathbf{F}) dx. \quad (58)$$

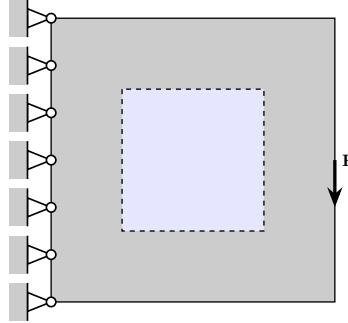


Figure 4: Geometry and boundary conditions for the linear elastic structure with the square domain.

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.

In this example, we form a multifidelity hierarchy through mesh coarsening through mesh parameter  $h$ . The model solved with mesh size  $h = 2^{-7}$  is the high-fidelity model. We create three low-fidelity models based on more economical discretizations:  $h = 2^{-3}, 2^{-2}, 2^{-1}$ . The outputs of these models are the energy QoI computed from the respective approximate solutions.

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.,  $c_0 = 4096, c_1 = 16, c_2 = 4, c_3 = 1$ . The correlations between the outputs of  $Y$  and  $X_1, X_2, X_3$  are 0.940, 0.841, -0.146, respectively. The total budget  $B$  ranges from  $10^5$  to  $10^7$ .

To mitigate potential model misspecification effects, we include second-order interactions between  $X_i, i \in [3]$  as additional regressors, i.e.,  $X_i X_j, i, j \in [3]$ . (We do observe model misspecification effects when  $B$  is near  $10^7$  without the interaction terms.) In this case, we have seven different models, namely,

$$\begin{aligned} \{i\} &\sim \text{span}\{1, X_i, X_i^2\} & i \in [3] \\ \{i, j\} &\sim \text{span}\{1, X_i, X_j, X_i^2, X_j^2, X_i X_j\} & i, j \in [3], i \neq j \\ \{1, 2, 3\} &\sim \text{span}\{1, X_1, X_2, X_3, X_1^2, X_2^2, X_3^2, X_1 X_2, X_1 X_3, X_2 X_3\}, \end{aligned}$$

with exploitation cost defined as the sum of the cost of the low-fidelity models used to build the regressors. The oracle  $F_Y$  is taken as an empirical CDF constructed from two million independent samples of the high-fidelity model. Accuracy results for ECDF-Y and AETC-d are reported in the first plot in Figure 5. For visualization, we also provide an instance of the estimated densities (`density` function in R [27] with the default bandwidth (data-dependent), and the Gaussian kernel for smoothing) and some statistics (mean, variance, skewness, and kurtosis) given by ECDF-Y and AETC-d using the same training dataset when  $B = 10^6$ .

Figure 5 shows that within the budget range of this experiment, AETC-d is asymptotically consistent and outperforms ECDF-Y by a substantial margin. The statistics computed by AETC-d are accurate up to the fourth order and significantly more accurate than the ones given by simple MC. The estimated density given by AETC-d at  $B = 10^6$  almost matches the density obtained from two million independent samples of  $Y$ , which approximately costs  $10^{10}$  budget units (and required around a week to generate on our hardware).

We notice that model misspecification seems to have little impact in this experiment. A possible explanation is that the selected regressors (which coincide with the regressors of the optimal model  $S_{\text{opt}} = \{1, 2, 3\}$ ) are sufficiently expressible for  $Y$ , and the magnitude of the corresponding model

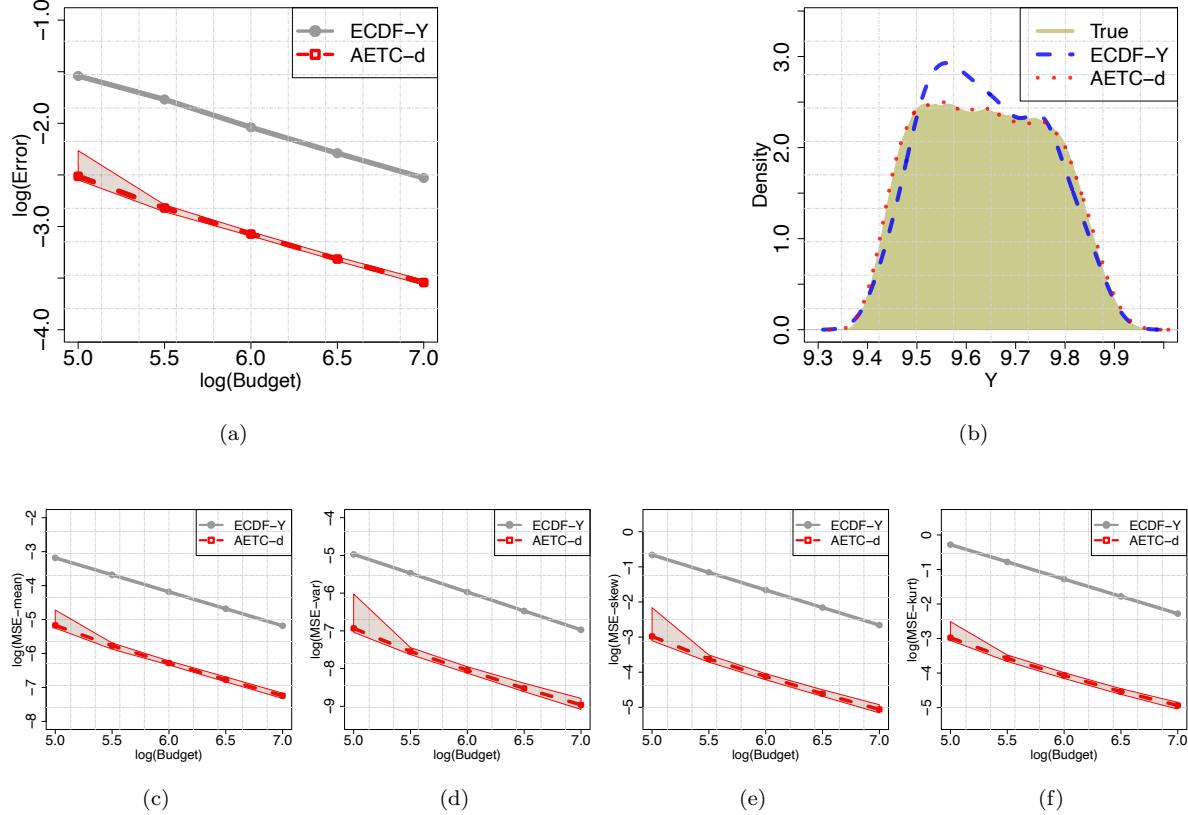


Figure 5:  $(\log_{10})$  mean  $W_1$  distance between  $F_Y$  and the estimated CDFs given by ECDF-Y and AETC-d. The 5%-50%-95% quantiles are plotted for the AETC-d algorithm to measure its uncertainty in the exploration phase (a). An instance of the densities given by ECDF-Y and AETC-d when  $B = 10^6$  (b). Comparison of the  $(\log_{10})$  mean-squared errors (MSEs) of various estimated statistics by ECDF-Y and AETC-d (bottom); statistics from bottom left to right are the mean (c), variance (d), skewness (e), and kurtosis (f), respectively.

variance is much smaller than the variance of  $Y$ , making the potential noise misspecification effect negligible under the budget range of this example. We substantiate this hypothesis by fitting a linear model using the complete training data, and we approximately compute the variance ratio between the model noise and  $Y$ , which is  $3 \times 10^{-5} \ll 1$ .

## 7 Conclusions

In this paper, we introduce an efficient strategy for learning the distribution of a scalar-valued QoI in the multifidelity setup. Under a linear model assumption, we propose a semi-parametric approach for approximating the distribution by leveraging samples of models of different resolutions/costs. The main novelty in our analysis is to provide an asymptotically informative and computationally estimable upper bound for the average 1-Wasserstein distance between the estimator and the true distribution and use it to devise an adaptive algorithm, AETC-d, for efficient budget allocation. We show that, for a large budget, the AETC-d is consistent, and explores and exploits optimally under a proposed upper bound criterion. Our setup and algorithm require neither a model hierarchy nor an *a priori* estimate of cross-model correlations. We also discuss several approaches to mitigating the model misspecification impact when a linear regression assumption is violated.

Distribution learning is a much harder problem compared to parameter/statistic estimation in general. Our method takes an initial step towards addressing this problem under a linear model assumption and provides a potentially effective way to fully quantify the uncertainty of the QoI associated with a high-fidelity model. Our algorithm is robust when the model noise is relatively small. However, practical guidance to ameliorate model misspecification effects in practice is still lacking when the budget is sufficiently large. In future work, we will investigate alternative multifidelity CDF emulators that are less sensitive to the linear model assumptions and possess better robustness properties when the model is misspecified.

## Acknowledgement

We would like to thank the referees for their time and helpful comments which significantly improved the presentation of the manuscript. Y. Xu and A. Narayan are partially supported by National Science Foundation DMS-1848508. A. Narayan is partially supported by the Air Force Office of Scientific Research award FA9550-20-1-0338. Y. Xu would like to thank Dr. Xiaou Pan for clarifying a uniform consistency result in quantile regression. We also thank Dr. Ruijian Han for a careful reading of an early draft, and for providing several comments that improved the presentation of the manuscript.

## 8 Appendices

### A Proof of Lemma 4.1

Let  $|S| = s$ . We first prove (28a). Recall the definition of  $Y$  and  $Y'$ :

$$Y = X_S^\top \beta_S + \varepsilon_S \quad Y' = X_S^\top \hat{\beta}_S + \hat{\varepsilon}_S$$

where  $\varepsilon_S, \hat{\varepsilon}_S$  are independent of  $X_S$  by Assumption 2.2.

Define the true empirical distribution of  $\varepsilon_S$  using exploration samples as

$$\tilde{\varepsilon}_S \sim \frac{1}{m} \sum_{\ell \in [m]} \delta_{Y_\ell - X_{S,\ell}^\top \beta_S}.$$

By the additive property of  $W_1$  metric under independence [22] and the triangle inequality, we can upper bound the average  $W_1$  distance between  $F_Y$  and  $F_{Y'}$  by conditioning on the exploration data:

$$\begin{aligned} & \mathbb{E}[W_1(F_Y, F_{Y'})|Z_S, Y_{\text{epr}}] \\ & \leq \mathbb{E}[W_1(F_{X_S^\top \beta_S}, F_{X_S^\top \hat{\beta}_S})|Z_S, Y_{\text{epr}}] + \mathbb{E}[W_1(F_{\varepsilon_S}, F_{\hat{\varepsilon}_S})|Z_S, Y_{\text{epr}}] \\ & \leq \mathbb{E}[W_1(F_{X_S^\top \beta_S}, F_{X_S^\top \hat{\beta}_S})|Z_S, Y_{\text{epr}}] + \mathbb{E}[W_1(F_{\varepsilon_S}, F_{\tilde{\varepsilon}_S})|Z_S, Y_{\text{epr}}] + \mathbb{E}[W_1(F_{\tilde{\varepsilon}_S}, F_{\hat{\varepsilon}_S})|Z_S, Y_{\text{epr}}]. \end{aligned} \quad (59)$$

For the first term in (59), note  $\hat{\beta}_S - \beta_S$  satisfies

$$\hat{\beta}_S - \beta_S \sim (0, \sigma_S^2 (Z_S^\top Z_S)^{-1}). \quad (60)$$

Averaging out the randomness of exploration noise, we have that, almost surely,

$$\begin{aligned} \mathbb{E}[W_1(F_{X_S^\top \beta_S}, F_{X_S^\top \hat{\beta}_S})|Z_S] & \leq \left( \mathbb{E}[W_2^2(F_{X_S^\top \beta_S}, F_{X_S^\top \hat{\beta}_S})|Z_S] \right)^{1/2} \\ & \leq \left( \mathbb{E}[|X_S^\top (\hat{\beta}_S - \beta_S)|^2|Z_S] \right)^{1/2} \\ & = \left( \text{tr} \left( \mathbb{E}[X_S X_S^\top] \mathbb{E}[(\hat{\beta}_S - \beta_S)(\hat{\beta}_S - \beta_S)^\top|Z_S] \right) \right)^{1/2} \\ & = \left( \frac{\sigma_S^2}{m} \text{tr}(\Lambda_S (m^{-1} Z_S^\top Z_S)^{-1}) \right)^{1/2} \\ & \simeq \sqrt{\frac{s+1}{m}} \sigma_S, \end{aligned} \quad (61)$$

where the first step uses Jensen's inequality, and the last step follows from the law of large numbers and Assumption 2.3.

For the second term in (59), note that  $\tilde{\varepsilon}_S$  is the empirical distribution of  $\varepsilon_S$  based on  $m$  exploration samples, which does not depend on  $Z_S$  under Assumption 2.2. Applying the nonasymptotic estimates on the convergence rate of empirical measures in Lemma 3.2 obtains

$$\mathbb{E}[W_1(F_{\varepsilon_S}, F_{\tilde{\varepsilon}_S})|Z_S] = \mathbb{E}[W_1(F_{\varepsilon_S}, F_{\tilde{\varepsilon}_S})] \leq \frac{J_1(F_{\varepsilon_S})}{\sqrt{m}}, \quad (62)$$

where  $J_1$  is defined in (10).

For the third term in (59), consider the natural coupling between  $\hat{\varepsilon}_S$  and  $\tilde{\varepsilon}_S$ :  $\tilde{\varepsilon}_S^\leftarrow(\tilde{\tau}_\ell) = \hat{\varepsilon}_S^\leftarrow(\hat{\tau}_\ell)$ , where  $\leftarrow$  denotes the preimage of a map and

$$\tilde{\tau}_\ell = Y_\ell - X_{S,\ell}^\top \beta_S \quad \hat{\tau}_\ell = Y_\ell - X_{S,\ell}^\top \hat{\beta}_S. \quad (63)$$

In this case,

$$\begin{aligned} \mathbb{E}[W_1(F_{\tilde{\varepsilon}_S}, F_{\hat{\varepsilon}_S})|Z_S] & \leq (\mathbb{E}[W_2^2(F_{\tilde{\varepsilon}_S}, F_{\hat{\varepsilon}_S})|Z_S])^{1/2} \leq (\mathbb{E}[|\tilde{\varepsilon}_S - \hat{\varepsilon}_S|^2|Z_S])^{1/2} \\ & = \left( \frac{1}{m} \sum_{\ell \in [m]} \mathbb{E}[(X_{S,\ell}^\top (\hat{\beta}_S - \beta_S))^2|Z_S] \right)^{1/2} \\ & = \sqrt{\frac{s+1}{m}} \sigma_S. \end{aligned} \quad (64)$$

Putting (61), (62), (64) together finishes the proof of (28a).

We next prove (28b). Conditioned on  $Z_S$  and  $Y_{\text{epr}}$ ,  $Y'$  is a random variable with bounded  $r$ -th moments for all  $r > 2$ . Appealing to Lemma 3.2 and averaging over the exploration noise, we have

$$\mathbb{E} \left[ W_1 \left( \hat{F}_{Y,S}, F_{Y'} \right) | Z_S, Y_{\text{epr}} \right] \leq \frac{\mathbb{E}[J_1(F_{Y'})|Z_S, Y_{\text{epr}}]}{\sqrt{N_S}} \implies \mathbb{E} \left[ W_1 \left( \hat{F}_{Y,S}, F_{Y'} \right) | Z_S \right] \leq \frac{\mathbb{E}[J_1(F_{Y'})|Z_S]}{\sqrt{N_S}}, \quad (65)$$

where  $J_1$  is defined in (10). The desired result would follow if we can show that  $\mathbb{E}[J_1(F_{Y'})|Z_S]$  converges to  $J_1(F_Y)$  a.s. as  $m \rightarrow \infty$ . To this end, we introduce the following intermediate random variables:

$$Y'' = X_S^\top \beta_S + \tilde{\varepsilon}_S \quad Y''' = X_S^\top \beta_S + \hat{\varepsilon}_S.$$

We will prove the desired result by verifying the following convergence statements respectively:

- (a).  $|\mathbb{E}[J_1(F_Y)|Z_S] - \mathbb{E}[J_1(F_{Y''})|Z_S]| \rightarrow 0$  a.s.;
- (b).  $|\mathbb{E}[J_1(F_{Y''})|Z_S] - \mathbb{E}[J_1(F_{Y'''})|Z_S]| \rightarrow 0$  a.s.;
- (c).  $|\mathbb{E}[J_1(F_{Y'''})|Z_S] - \mathbb{E}[J_1(F_{Y'})|Z_S]| \rightarrow 0$  a.s.

Without loss of generality, we assume  $\text{supp}(\varepsilon_S) \subseteq [-1, 1]$  and  $\|\beta_S\|_2 = 1$ ; the general case can be considered similarly by taking appropriate scaling involving a constant  $C$  in Assumption 2.5.

We introduce the following quantity for our analysis:

$$K_m^* = \max_{\ell \in [m]} \|X_{S,\ell}\|_2. \quad (66)$$

It is clear that  $K_m^*$  depends only on  $Z_S$ . Under Assumption 2.4,  $X_{S,\ell}$ 's are i.i.d. sub-exponential random variables with uniformly bounded sub-exponential norm. By Lemma 3.6,

$$K_m^* \lesssim \log m \quad a.s., \quad (67)$$

where the implicit constant is realization-dependent.

To prove (a), we condition on  $Z_S$  and  $Y_{\text{epr}}$ . Using (13)-(15) in Lemma 3.5,

$$\begin{aligned} & |\mathbb{E}[J_1(F_Y)|Z_S, Y_{\text{epr}}] - \mathbb{E}[J_1(F_{Y''})|Z_S, Y_{\text{epr}}]| \\ & \leq \mathbb{E} \left[ \int_{\mathbb{R}} \sqrt{|F_Y(y) - F_{Y''}(y)|} dy | Z_S, Y_{\text{epr}} \right] \\ & = \mathbb{E} \left[ \int_{\mathbb{R}} \sqrt{\left| \int_{\mathbb{R}} F_{X_S^\top \beta_S}(y-z) dF_{\varepsilon_S}(z) - \int_{\mathbb{R}} F_{X_S^\top \beta_S}(y-z) dF_{\tilde{\varepsilon}_S}(z) \right|} dy | Z_S, Y_{\text{epr}} \right]. \end{aligned} \quad (68)$$

Under Assumption 2.4, for every  $y$ , since  $\|\beta_S\|_2 = 1$ ,  $F_{X_S^\top \beta_S}(y-z)$  as a function of  $z$ , is  $C_{\text{Lip}}$ -Lipschitz. By the Kantorovich-Rubinstein duality (6),

$$\left| \int_{\mathbb{R}} F_{X_S^\top \beta_S}(y-z) dF_{\varepsilon_S}(z) - \int_{\mathbb{R}} F_{X_S^\top \beta_S}(y-z) dF_{\tilde{\varepsilon}_S}(z) \right| \leq C_{\text{Lip}} W_1(\varepsilon_S, \tilde{\varepsilon}_S). \quad (69)$$

Meanwhile, note  $\text{supp}(\varepsilon_S) = \text{supp}(\tilde{\varepsilon}_S) \subseteq [-1, 1]$ . This combined with the fact that  $X_S^\top \beta_S$  is sub-exponential implies that

$$\begin{aligned} & \left| \int_{\mathbb{R}} F_{X_S^\top \beta_S}(y-z) dF_{\varepsilon_S}(z) - \int_{\mathbb{R}} F_{X_S^\top \beta_S}(y-z) dF_{\tilde{\varepsilon}_S}(z) \right| \\ & = \left| \int_{\mathbb{R}} 1 - F_{X_S^\top \beta_S}(y-z) dF_{\varepsilon_S}(z) - \int_{\mathbb{R}} 1 - F_{X_S^\top \beta_S}(y-z) dF_{\tilde{\varepsilon}_S}(z) \right| \\ & \leq \frac{1}{2} \left( M_1[F_{X_S^\top \beta_S}](y) + M_1[1 - F_{X_S^\top \beta_S}](y) \right) \\ & \leq \exp \left( -\frac{\max\{|y-1|, |y+1|\}}{C} \right) \\ & \leq \exp \left( -\frac{|y|}{2C} \right) \quad |y| \geq 2, \end{aligned} \quad (70)$$

where  $C$  is an absolute constant depending only on the sub-exponential norm of  $\|X_S\|_2$ , and  $M_1$  is the 1-local maximum operator in Definition 3.4. Substituting (69) and (70) into (68) and applying a truncated estimate,

$$\begin{aligned} & |\mathbb{E}[J_1(F_Y)|Z_S, Y_{\text{epr}}] - \mathbb{E}[J_1(F_{Y''})|Z_S, Y_{\text{epr}}]| \\ & \leq \mathbb{E} \left[ \int_{|y|<\max\{2,4C\log m\}} \sqrt{W_1(\varepsilon_S, \tilde{\varepsilon}_S)} dy + \int_{|y|\geq\max\{2,4C\log m\}} \exp\left(-\frac{|y|}{4C}\right) dy |Z_S, Y_{\text{epr}} \right] \\ & \leq \mathbb{E} \left[ (4 + 8C\log m) \sqrt{W_1(\varepsilon_S, \tilde{\varepsilon}_S)} + \frac{2}{m} |Z_S, Y_{\text{epr}} \right]. \end{aligned}$$

Since  $\sqrt{W_1(\varepsilon_S, \tilde{\varepsilon}_S)}$  is independent of  $Z_S$ , taking expectation over the exploration noise together with Jensen's inequality and Lemma 3.2 yields

$$|\mathbb{E}[J_1(F_Y)|Z_S] - \mathbb{E}[J_1(F_{Y''})|Z_S]| \leq (4 + 8C\log m) \mathbb{E}[W_1(\varepsilon_S, \tilde{\varepsilon}_S)]^{1/2} + \frac{2}{m} \lesssim \frac{\log m}{\sqrt{m}} \rightarrow 0. \quad (71)$$

To prove (b), note that conditioning on  $Z_S$  and  $Y_{\text{epr}}$ , the difference between  $\tilde{\tau}_\ell$  and  $\hat{\tau}_\ell$  is bounded as follows:

$$|\tilde{\tau}_\ell - \hat{\tau}_\ell| = |X_{S,\ell}^\top(\hat{\beta}_S - \beta_S)| \leq \|X_{S,\ell}\|_2 \|\hat{\beta}_S - \beta_S\|_2 \stackrel{(66)}{\leq} K_m^* \delta \quad \delta = \|\hat{\beta}_S - \beta_S\|_2, \quad (72)$$

where  $\hat{\tau}_\ell, \tilde{\tau}_\ell$  are defined in (63). Moreover, since  $\text{supp}(\varepsilon_S) \subseteq [-1, 1]$ ,  $|\tilde{\tau}_\ell| \leq 1$ . This combined with (72) implies

$$\text{supp}(\tilde{\varepsilon}_S) \cup \text{supp}(\hat{\varepsilon}_S) \subseteq [-r, r] \quad r = 1 + K_m^* \delta. \quad (73)$$

The rest is similar to the proof of statement (a),

$$\begin{aligned} & |\mathbb{E}[J_1(F_{Y''})|Z_S, Y_{\text{epr}}] - \mathbb{E}[J_1(F_{Y''''})|Z_S, Y_{\text{epr}}]| \\ & \leq \mathbb{E} \left[ \int_{\mathbb{R}} \sqrt{|F_{Y''}(y) - F_{Y''''}(y)|} dy |Z_S, Y_{\text{epr}} \right] \\ & \leq \mathbb{E} \left[ \int_{\mathbb{R}} \sqrt{\frac{1}{m} \sum_{\ell \in [m]} |F_{X_S^\top \beta_S}(y - \hat{\tau}_\ell) - F_{X_S^\top \beta_S}(y - \tilde{\tau}_\ell)|} dy |Z_S, Y_{\text{epr}} \right]. \end{aligned}$$

It is easy to verify using the Lipschitz assumption and the tail bound of  $X_S^\top \beta_S$  that

$$\begin{aligned} & \frac{1}{m} \sum_{\ell \in [m]} |F_{X_S^\top \beta_S}(y - \hat{\tau}_\ell) - F_{X_S^\top \beta_S}(y - \tilde{\tau}_\ell)| \leq C_{\text{Lip}} K_m^* \delta \\ & \frac{1}{m} \sum_{\ell \in [m]} |F_{X_S^\top \beta_S}(y - \hat{\tau}_\ell) - F_{X_S^\top \beta_S}(y - \tilde{\tau}_\ell)| \leq \exp\left(-\frac{|y|}{2C}\right) \quad |y| \geq 2r. \end{aligned}$$

Thus,

$$\begin{aligned} & |\mathbb{E}[J_1(F_{Y''})|Z_S, Y_{\text{epr}}] - \mathbb{E}[J_1(F_{Y''''})|Z_S, Y_{\text{epr}}]| \\ & \leq \mathbb{E} \left[ \int_{|y|<\max\{2r,4C\log m\}} \sqrt{C_{\text{Lip}} K_m^* \delta} dy + \int_{|y|\geq\max\{2r,4C\log m\}} \exp\left(-\frac{|y|}{4C}\right) dy |Z_S, Y_{\text{epr}} \right] \\ & \leq \mathbb{E} \left[ (4r + 8C\log m) \sqrt{C_{\text{Lip}} K_m^* \delta} + \frac{2}{m} |Z_S, Y_{\text{epr}} \right]. \end{aligned}$$

Averaging out exploration noise and applying Jensen's inequality,

$$\begin{aligned}
|\mathbb{E}[J_1(F_{Y''})|Z_S] - \mathbb{E}[J_1(F_{Y'''})|Z_S]| &\lesssim \mathbb{E} \left[ (K_m^* \delta)^{3/2} + \log m \sqrt{K_m^* \delta} + \frac{2}{m} |Z_S| \right] \\
&\lesssim (K_m^*)^{3/2} \mathbb{E}[\delta^2 |Z_S|^{3/4} + \log m \sqrt{K_m^*} \mathbb{E}[\delta^2 |Z_S|^{1/4} + \frac{1}{m} \\
&\stackrel{(60),(67)}{\lesssim} \frac{(\log m)^{3/2}}{m^{1/4}} \rightarrow 0 \quad a.s.
\end{aligned}$$

To prove (c), recall from (73) that conditioning on  $Z_S$  and  $Y_{\text{epr}}$ ,  $\text{supp}(\hat{\varepsilon}_S) \subseteq [-r, r]$ . Applying Lemma 3.5,

$$|\mathbb{E}[J_1(F_{Y'''})|Z_S, Y_{\text{epr}}] - \mathbb{E}[J_1(F_{Y'})|Z_S, Y_{\text{epr}}]| \leq \mathbb{E} \left[ \left\| M_r [F_{X_S^\top \hat{\beta}_S} - F_{X_S^\top \beta_S}] \right\|_{L_{\mathbb{R}}^{1/2}}^{1/2} |Z_S, Y_{\text{epr}} \right]. \quad (74)$$

If  $\delta < 1/2$ , then  $1/2 < \|\hat{\beta}_S\|_2 < 3/2$ . In this case, the  $C_1, C_2, C_3$  in Lemma 3.7 are absolute constants. According to Lemma 3.7 with  $p = 1/2$ ,

$$\left\| M_r [F_{X_S^\top \hat{\beta}_S} - F_{X_S^\top \beta_S}] \right\|_{L_{\mathbb{R}}^{1/2}}^{1/2} \lesssim (r+1) \delta^{5/12} \log(1/\delta) \leq (r+1) \delta^{1/4}, \quad (75)$$

where we used  $\log(1/\delta) < \delta^{-1/6}$  when  $\delta \leq 1/2$ .

If  $\delta \geq 1/2$ , the same result in Lemma 3.7 implies

$$\left\| M_r [F_{X_S^\top \hat{\beta}_S} - F_{X_S^\top \beta_S}] \right\|_{L_{\mathbb{R}}^{1/2}}^{1/2} \lesssim (r+1+\delta) \delta. \quad (76)$$

Substituting (75) and (76) into (74) yields that

$$|\mathbb{E}[J_1(F_{Y'''})|Z_S, Y_{\text{epr}}] - \mathbb{E}[J_1(F_{Y'})|Z_S, Y_{\text{epr}}]| \lesssim \mathbb{E}[(r+1) \delta^{1/4} + (r+1+\delta) \delta |Z_S, Y_{\text{epr}}].$$

Taking expectation over the exploration noise and applying Jensen's inequality,

$$\begin{aligned}
&|\mathbb{E}[J_1(F_{Y'''})|Z_S] - \mathbb{E}[J_1(F_{Y'})|Z_S]| \\
&\lesssim \mathbb{E}[\delta^{1/4} |Z_S|] + \mathbb{E}[\delta |Z_S|] + K_m^* \mathbb{E}[\delta^{5/4} |Z_S|] + (K_m)^2 \mathbb{E}[\delta^2 |Z_S|] \\
&\lesssim \mathbb{E}[\delta^2 |Z_S|^{1/8} + \mathbb{E}[\delta^2 |Z_S|^{1/2} + K_m^* \mathbb{E}[\delta^2 |Z_S|^{5/8} + (K_m)^2 \mathbb{E}[\delta^2 |Z_S|] \\
&\stackrel{(60),(66)}{\lesssim} \frac{1}{m^{1/8}} \rightarrow 0 \quad a.s.
\end{aligned}$$

(28b) is proved by combining statements (a), (b), (c).

## B A quantile regression framework

Quantile regression offers an alternative approach to simulating  $Y$  through a random coefficient interpretation [15]. For any  $S \subseteq [n]$  and  $\tau \in (0, 1)$ , we assume the conditional  $\tau$ -th quantile of  $Y$  on  $X_S$  satisfies

$$F_{Y|X_S}^{-1}(\tau) = X_S^\top \beta_S(\tau), \quad (77)$$

where  $\beta_S(\tau)$  the  $\tau$ -th coefficient vector. (77) is a standard quantile regression formulation, and can be used to model heteroscedastic noise effects.

$$\hat{\beta}_S(\tau) = \arg \min_{\beta \in \mathbb{R}^{s+1}} \frac{1}{m} \sum_{\ell \in [m]} \rho_\tau(Y_\ell - X_{\text{epr}, \ell}^\top \beta) \quad \rho_\tau(x) = x(\tau - \mathbf{1}_{x < 0}).$$

Thus, (77) approximately equals

$$\widehat{F}_{Y|X_S}^{-1}(\tau) = X_S^\top \widehat{\beta}_S(\tau). \quad (78)$$

As opposed to (23), (78) provides a way to simulate  $Y$  based on  $X_S$  via inverse transform sampling:

$$Y \approx X_S^\top \widehat{\beta}_S(U). \quad U \sim \text{Unif}(0, 1) \perp\!\!\!\perp X_S. \quad (79)$$

In our case,  $X_{\text{epr},\ell}, \ell \in [m]$  are i.i.d. samples so (77) fits into a random design quantile regression framework as analyzed in [21], where the authors established a strong consistency result for  $\widehat{\beta}_S(\tau)$  under suitable conditions. The consistency result can be further proven to hold uniformly for all  $\tau \in [\delta, 1 - \delta]$  for any fixed  $\delta > 0$ , which justifies the asymptotic behavior of the procedure in (78) as  $m, N_S \rightarrow \infty$ .

In the quantile regression framework, obtaining the optimal choices for  $m$  and  $S$  is much harder than in the linear regression setup. The AETC-d-q algorithm in Section 6 implements (79) with  $m$  set as the adaptive exploration rate given by the AETC-d,  $S$  as the corresponding model output for exploitation, and  $U$  approximated via  $\frac{1}{K} \sum_{j \in [K]} \delta_{\frac{j}{K+1}}$  with  $K = 100$ .

## References

- [1] R. Beran, L Le Cam, and P. Millar. “Convergence of stochastic empirical measures”. *Journal of multivariate analysis* 23.1 (1987), pp. 159–168.
- [2] S. Bobkov and M. Ledoux. *One-dimensional empirical measures, order statistics, and Kantorovich transport distances*. Vol. 261. 1259. American Mathematical Society, 2019.
- [3] S. Bubeck, N. Cesa-Bianchi, et al. “Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems”. *Foundations and Trends® in Machine Learning* 5.1 (2012), pp. 1–122.
- [4] S. Cambanis, G. Simons, and W. Stout. “Inequalities for  $E k(x, y)$  when the marginals are fixed”. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* 36.4 (1976), pp. 285–294.
- [5] S. Chatterjee. “Lecture notes on Stein’s method”. *Stanford lecture notes* (2007).
- [6] A. Cohen and R. DeVore. “Approximation of high-dimensional parametric PDEs”. *Acta Numerica* 24 (2015), pp. 1–159. ISSN: 1474-0508. doi: [10.1017/S0962492915000033](https://doi.org/10.1017/S0962492915000033).
- [7] I.-G. Farcas. “Context-aware model hierarchies for higher-dimensional uncertainty quantification”. PhD thesis. Technische Universität München, 2020.
- [8] A. L. Gibbs and F. E. Su. “On Choosing and Bounding Probability Metrics”. *International Statistical Review* 70.3 (2002), pp. 419–435. doi: [10.1111/j.1751-5823.2002.tb00178.x](https://doi.org/10.1111/j.1751-5823.2002.tb00178.x).
- [9] M. B. Giles. “Multilevel monte carlo path simulation”. *Operations research* 56.3 (2008), pp. 607–617.
- [10] 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. doi: [10.1137/140960086](https://doi.org/10.1137/140960086).
- [11] M. B. Giles, T. Nagapetyan, and K. Ritter. “Adaptive multilevel Monte Carlo approximation of distribution functions”. *arXiv preprint arXiv:1706.06869* (2017).
- [12] 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. doi: [10.1016/j.jcp.2020.109257](https://doi.org/10.1016/j.jcp.2020.109257).
- [13] J. M. Hammersley and D. C. Handscomb. *Monte Carlo Methods*. Methuen, London, 1964.
- [14] T. Ishigami and T. Homma. “An importance quantification technique in uncertainty analysis for computer models”. In: *[1990] Proceedings. First International Symposium on Uncertainty Modeling and Analysis*. IEEE Comput. Soc. Press, 1990. doi: [10.1109/isuma.1990.151285](https://doi.org/10.1109/isuma.1990.151285).
- [15] R. Koenker. “Fundamentals of Quantile Regression”. In: *Quantile Regression*. Cambridge University Press, pp. 26–67. doi: [10.1017/cbo9780511845734.002](https://doi.org/10.1017/cbo9780511845734.002).

[16] S. Krumscheid and F. Nobile. “Multilevel Monte Carlo Approximation of Functions”. *SIAM/ASA Journal on Uncertainty Quantification* 6.3 (2018), pp. 1256–1293. DOI: [10.1137/17m1135566](https://doi.org/10.1137/17m1135566).

[17] T. L. Lai and C. Z. Wei. “Least Squares Estimates in Stochastic Regression Models with Applications to Identification and Control of Dynamic Systems”. *The Annals of Statistics* 10.1 (1982), pp. 154–166. DOI: [10.1214/aos/1176345697](https://doi.org/10.1214/aos/1176345697).

[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. DOI: [10.1002/2016wr019475](https://doi.org/10.1002/2016wr019475).

[20] P. Massart. “The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality”. *The Annals of Probability* (1990), pp. 1269–1283.

[21] X. Pan and W.-X. Zhou. “Multiplier bootstrap for quantile regression: non-asymptotic theory under random design”. *Information and Inference: A Journal of the IMA* (2020). DOI: [10.1093/imaiiai/iaaa006](https://doi.org/10.1093/imaiiai/iaaa006).

[22] V. M. Panaretos and Y. Zemel. “Statistical aspects of Wasserstein distances”. *Annual review of statistics and its application* 6 (2019), pp. 405–431.

[23] B. Peherstorfer. “Multifidelity Monte Carlo estimation with adaptive low-fidelity models”. *SIAM/ASA Journal on Uncertainty Quantification* 7.2 (2019), pp. 579–603.

[24] 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. DOI: [10.1137/15m1046472](https://doi.org/10.1137/15m1046472).

[25] B. Peherstorfer, K. Willcox, and M. Gunzburger. “Survey of Multifidelity Methods in Uncertainty Propagation, Inference, and Optimization”. *SIAM Review* 60.3 (2018), A550–A591.

[26] E. Qian, B. Peherstorfer, D. O’Malley, V. V. Vesselinov, and K. Willcox. “Multifidelity Monte Carlo Estimation of Variance and Sensitivity Indices”. *SIAM/ASA Journal on Uncertainty Quantification* 6.2 (2018), pp. 683–706. DOI: [10.1137/17m1151006](https://doi.org/10.1137/17m1151006).

[27] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria, 2020.

[28] P. M. Robinson. “Root-N-Consistent Semiparametric Regression”. *Econometrica* 56.4 (1988), p. 931. DOI: [10.2307/1912705](https://doi.org/10.2307/1912705).

[29] D. Schaden and E. Ullmann. “On Multilevel Best Linear Unbiased Estimators”. *SIAM/ASA Journal on Uncertainty Quantification* 8.2 (2020), pp. 601–635. DOI: [10.1137/19m1263534](https://doi.org/10.1137/19m1263534).

[30] 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.

[31] R. Vershynin. *High-dimensional probability: An introduction with applications in data science*. Vol. 47. Cambridge university press, 2018.

[32] C. Villani. “The metric side of optimal transportation”. In: *Graduate Studies in Mathematics*. American Mathematical Society, 2003, pp. 205–235. DOI: [10.1090/gsm/058/08](https://doi.org/10.1090/gsm/058/08).

[33] D. Williams. *Probability with martingales*. Cambridge university press, 1991.

[34] Y. Xu, V. Keshavarzzadeh, R. M. Kirby, and A. Narayan. “A bandit-learning approach to multifidelity approximation”. *SIAM Journal on Scientific Computing* 44.1 (2022), A150–A175.