DISTRIBUTIONALLY ROBUST BAYESIAN EXPERIMENTAL DESIGN VIA CONFIDENCE-BAND AMBIGUITY SETS

Anonymous authors
Paper under double-blind review

ABSTRACT

Model misspecification in Bayesian Experimental Design (BED) can lead to overconfident posterior inference and unreliable design selection. This limitation hinders the use of the approach in safety-critical applications, such as healthcare, where experimental outcomes must be accompanied by guarantees at a specified confidence level. To address this challenge, we propose a *Distributionally Robust BED* (DR-BED) framework that introduces a robust optimisation problem with a novel ambiguity set, constructed from a pair of lower-bound and upper-bound functions as confidence bands for the true Data Generating Process (DGP). We establish theoretical guarantees showing that, under standard regularity conditions on the experimental setting, namely Hölder continuity and compactness of the design space and bounded support of the experimental outcome, the constructed ambiguity set contains the true DGP with the experimenter-specified confidence level $1 - \alpha$. Notably, the confidence level can be updated dynamically throughout the sequential experimentation process, allowing experts to explore different robustness-efficiency trade-offs without repeating the experimental workflow. Building on the theoretical guarantees, we reformulate the DR-BED optimisation problem in its dual form and develop a stochastic subgradient algorithm for its practical implementation. Experiments on controlled and benchmark problems demonstrate that, under model misspecification, DR-BED consistently produces improved posterior beliefs and achieves lower distances to the true DGP compared to the baseline across multiple distributional metrics.

1 INTRODUCTION

Bayesian Experimental Design (BED) is a principled model-based framework for scientific discovery that enables efficient design selection (Chaloner & Verdinelli, 1995; Rainforth et al., 2024). Given a data-generating process (DGP) relating the design choices to the outcomes, standard BED recommends optimal designs that maximise Expected Information Gain (EIG) (Rainforth et al., 2024), or equivalently, the part of the outcome uncertainty that is explained by the model parameters (Sebastiani & Wynn, 2000). In practice, however, experts can rarely begin an experimental design with a complete characterisation of the model. Instead, they start with prior knowledge, perform experiments, analyse observations, and revise their beliefs, and then the model, and hence neither its parameters, necessarily fully capture the uncertainty..

To address this limitation, recent works move beyond reliance on a single probabilistic model, explicitly accounting for uncertainty over a collection of plausible data-generating distributions (Ouyang et al., 2020; Barlas et al., 2025). This perspective can be formalised within a *Distributionally Robust Experimental Design* framework, which is based on the notion of constructing an ambiguity set (Go & Isaac, 2022; Igea & Cicirello, 2025). This set is a confidence region that contains the true DGP with a user-specified confidence level. Experimental designs are then optimised against the worst-case distribution within this set, yielding designs that remain informative under model misspecification and preserve performance guarantees at the prescribed confidence level (Amortila et al., 2024; Catanach & Das, 2023; Kirschner et al., 2020).

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

While (Go & Isaac, 2022; Igea & Cicirello, 2025; Barlas et al., 2025) have proposed robust BED frameworks, we introduce an alternative formulation that provides answers to the following new question, and then compare it with the existing the solutions:

How can we optimise the ambiguity set to achieve informative designs under model misspecification while guaranteeing adherence to an experimenter-specified confidence level?

To answer this question, we introduce a *Distributionally Robust BED* framework in which the ambiguity set is constructed as a pair of functional confidence bands, which is novel in the context of BED. It allows the bands to evolve throughout the sequential design process, narrowing around the true DGP as statistical evidence accumulated from experimental observations. Such a framework not only explicitly accounts for uncertainty in the likelihood model but also enables the experimenter to progressively adjust the confidence level throughout the experiment by varying the width of the confidence bands, thereby reflecting their degree of conservativeness. Figure 1 illustrates the application of BED in early-stage drug discovery, where the goal is to understand the drug's effect on the body, and BED determines the optimal blood sampling times for patients.

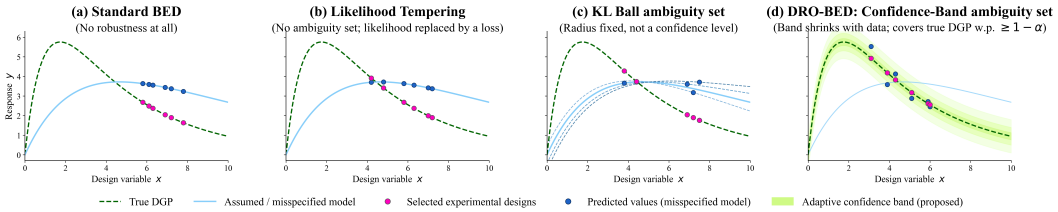


Figure 1: **Robustness to model misspecification in BED (Pharmacokinetic (PK) model) (a)** Standard BED is not robust to misspecification. It trusts the misspecified model and seeks designs that reduce posterior uncertainty under it. **(b)** Tempering the likelihood (or replacing it by a loss), does not necessarily suggest robust and informative designs. **(c)** Optimising against the worst-case distribution in a Kullback-Leibler (KL) ball (dashed blue lines) around the misspecified model, without guarantees that it contains the true DGP, results in ineffective designs. **(d)** Our method contains the true DGP with confidence level $1 - \alpha$ and produces designs that are both robust (optimal under worst-case) and informative (guiding posterior belief toward true DGP).

Our contributions are summarised as follows:

- **Adaptive confidence-band ambiguity sets for robust BED.** We introduce a distributionally robust BED framework in which the ambiguity set is defined by functional lower and upper bounds on the true data-generating process and is updated dynamically as data are collected
- **Theoretical guarantees.** We establish a confidence level guarantee for the proximity between the worst estimate and the true DGP under adaptive ambiguity set updates and prove bounded regret for robust sequential experiment selection.
- **Empirical validation.** We show on Pharmacokinetic and synthetic benchmarking simulations that our proposed framework recommends more informative designs under model misspecification than other robust BED frameworks, while respecting the specified confidence guarantee.

Related work Classical BED assumes a correctly specified model, yet real-world experiments are often affected by model misspecification (Chaloner & Verdinelli, 1995). Recent approaches improve robustness through alternative information measures, generalised Bayesian objectives, and robust active learning strategies (Catanach & Das, 2023; Barlas et al., 2025; Forster et al., 2025; Overstall et al., 2026; Tang et al., 2026). However, these methods typically modify the utility or inference procedure and do not provide explicit performance guarantees under a prescribed confidence level.

Distributionally robust optimisation (DRO) instead seeks designs that perform well against the worst-case distribution within an ambiguity set (Hu & Hong, 2013; Gao & Kleywegt, 2023; Caprio et al., 2024; Wen & Yang, 2026). Building on this perspective, we introduce a distributionally robust BED framework based on confidence-band ambiguity sets that are updated sequentially as data are

collected and that contain the true DGP with a user-specified confidence level $1-\alpha$. A detailed review of related work is provided in Appendix A.1.

2 Sequential Bayesian Experimental Design

Consider a supervised learning setting where the goal of an experimenter is to predict the outcome of an experiment, $y \in \mathscr{Y} := [a, b] \subset \mathbb{R}$, by observing the design vector $\boldsymbol{d} \in \mathscr{D} \subseteq \mathbb{R}^k$, where $\mathscr{D}$ denotes a compact design space. A limited experimental budget, $\mathfrak{B}$, is available for collecting experimental samples, $\mathscr{H} = \{(\boldsymbol{d}_i, y_i)\}_{i=1}^{\mathfrak{B}}$, and the experimental outcomes are assumed to be generated from an unknown conditional probability density $\tilde{q}(y \mid \boldsymbol{d})$, referred to as the true data-generating process (DGP). Since the true DGP is unknown, the experimenter specifies prior knowledge on DGP in the form of a parametric likelihood function $p(y \mid \boldsymbol{d}, \boldsymbol{\theta})$, where $\boldsymbol{\theta} \in \Theta$ denotes the model parameters. The corresponding hypothesis model class is defined as $\mathscr{P} = \{p(y \mid \boldsymbol{d}, \boldsymbol{\theta}) : \boldsymbol{\theta} \in \Theta\}$, and the learner's objective is to find the best model through determining the parameter $\boldsymbol{\theta}^*$ that maximises the expected decision-making utility. Following the Bayesian inference framework, the experimenter represents the uncertainty regarding the model parameters $\boldsymbol{\theta}$ through a prior distribution over the parameter space, denoted by $p(\boldsymbol{\theta})$. At each round t of the experiment, a new observation $(\boldsymbol{d}_t, y_t)$ is collected and used to obtain a distribution of posterior beliefs over the model parameters according to Bayes' theorem, $p(\boldsymbol{\theta} \mid \mathscr{H}_t) \propto p(\boldsymbol{\theta} \mid \mathscr{H}_{t-1})\, p(y_t \mid \boldsymbol{d}_t, \boldsymbol{\theta})$, where $p(y_t \mid \boldsymbol{d}_t, \boldsymbol{\theta})$ denotes the likelihood of observing the experimental outcome y_t given the design vector $\boldsymbol{d}_t$ and the model parameter $\boldsymbol{\theta}$. To conduct the experimental process efficiently, the experimenter specifies a utility function $U(\boldsymbol{d})$, and the optimal design at each round is selected by maximizing this utility, $\boldsymbol{d}^\star = \arg\max_{\boldsymbol{d} \in \mathscr{D}} U(\boldsymbol{d})$. When the utility function is the model log-likelihood, experiment selection reduces to maximising the Expected Information Gain (EIG), which measures the expected reduction in uncertainty about the model parameters $\boldsymbol{\theta}$ resulting from an experiment with design $\boldsymbol{d}$.

Since the exact likelihood function is often unknown in real-world applications, sequential experimental design is typically conducted under uncertainty regarding the assumed likelihood model. The following definition formalises model misspecification of the data generating process (Walker (2013), see Tang et al. (2026)).

Definition 2.1: Model Misspecification

A data-generating process is said to be *misspecified* if the defined hypothesis class $\mathscr{P}$ is incapable of representing the true DGP. In other words, misspecification occurs when there exists no parameter value $\boldsymbol{\theta} \in \Theta$ such that $p(y \mid \boldsymbol{d}, \boldsymbol{\theta}) = \tilde{q}(y \mid \boldsymbol{d})$. Equivalently, $\tilde{q}(y \mid \boldsymbol{d}) \notin \mathscr{P}$.

3 Proposed Distributionally Robust BED framework

In this section we present the DR-BED framework, where an experimenter explicitly accounts for the functional variation in the likelihood function. Instead of making design decisions by optimising expected utility under a single likelihood function, the experimenter sets a confidence level $1-\alpha$ which specifies a set of plausible likelihood functions, known as an ambiguity set $\mathscr{Q}$. We construct the ambiguity set as a pair of $(l_\alpha(y \mid \boldsymbol{d}), u_\alpha(y \mid \boldsymbol{d}))$, where $l_\alpha(y \mid \boldsymbol{d})$ is the lower bound and $u_\alpha(y \mid \boldsymbol{d})$ is the upper bound of a functional confidence band for the true DGP. Then, the objective in each round of experiments would be to choose the design that is maximally informative under the worst-case distribution within the specified ambiguity set as follows:

$$\boldsymbol{d}^\star = \arg\max_{\boldsymbol{d} \in \mathscr{D}} \inf_{q \in \mathcal{Q} = \{q : l_\alpha(y|\boldsymbol{d}) \le q(y|\boldsymbol{d}) \le u_\alpha(y|\boldsymbol{d}), \int q(y|\boldsymbol{d})\, dy = 1\}} U(\boldsymbol{d}, q), \tag{1}$$

where $\boldsymbol{d}^\star$ denotes the optimal design point. After constructing the confidence-band ambiguity set $\mathscr{Q} = \{q : l_\alpha(y|\boldsymbol{d}) \le q(y|\boldsymbol{d}) \le u_\alpha(y|\boldsymbol{d}), \int q(y|\boldsymbol{d})\, dy = 1\}$ in Section 3.1 and deriving a dual formulation of the robust BED objective in Section 3.2, we develop in Section 3.3 a novel subgradient descent algorithm to solve equation 1.

3.1 Ambiguity Set Construction

Suppose that at the start of each round $t \in \{1, 2, \ldots, \mathfrak{B}\}$ of experiments, there is a set of available historical observations $\mathscr{H}_t$ of the true DGP. Additionally, assume there is limited information avail-

able on the shape of the true DGP, namely the mode of the true DGP, $\mu(\boldsymbol{d})$, and an upper bound $\tilde{q}(y \mid \boldsymbol{d}) \leq \mathcal{U}(\boldsymbol{d})$. These types of assumptions have been made repeatedly in previous work on kernel density estimation (Jiang, 2017) and DRO (Chen et al., 2022). Furthermore, let $f(\boldsymbol{d}, y)$ represent the integrand of the utility function, $U(\boldsymbol{d}, q) = \mathbb{E}_q[f(\boldsymbol{d}, y)]$. With the functional confidence band ambiguity set $\mathscr{Q}_t = \{l_\alpha^t(y \mid \boldsymbol{d}),\, u_\alpha^t(y \mid \boldsymbol{d})\}$, we define the robust optimal design at round t to be $\boldsymbol{d}_t^\star \in \arg\max_{\boldsymbol{d}\in\mathscr{D}} \inf_{q\in\mathscr{Q}_t} \mathbb{E}_q[f(\boldsymbol{d}, y)]$.

Since the designs $\boldsymbol{d}_i$ that produce the experimental outcomes y_i are not necessarily the same, we use a localisation technique based on the *Nadaraya–Watson* weighting scheme (Cai, 2001) with a smooth kernel $\mathcal{K}$. This allows all observations to contribute according to their proximity-weighted relevance, enriching the available information on possible outcomes of each candidate design $\boldsymbol{d}$ while also helping provide the conditional coverage guarantee in Theorem 3.1.

Assumption 1 (Kernel regularity conditions (Chen et al., 2022))**.** *Assume that the kernel function $\mathcal{K}$ satisfies the following conditions:*

- *$\mathcal{K}$ is a probability density function, that is, $\int_{\mathbb{R}^k} \mathcal{K}(u)\, du = 1$.*
- *$\mathcal{K}$ is radially symmetric, so that $\mathcal{K}(u) = \kappa(\|u\|_2)$, where $\kappa : [0,\infty) \to [0,\infty)$ is a non-increasing function.*
- *$\mathcal{K}$ has an exponentially decaying tail, where there are constants $r > 0$, $C_r > 0$, and $\tau > 0$ such that $\kappa(t) \leq C_r \exp(-t^r) \quad \forall\, t > \tau$.*

Nadaraya–Watson Weights For any target design point $\boldsymbol{d} \in \mathscr{D}$, we localise the available observations around $\boldsymbol{d}$ using the Nadaraya-Watson weighting scheme, which assigns the following weight to the i-th historical observation $(\boldsymbol{d}_i, y_i)$ (Cai, 2001):

$$w_i(\boldsymbol{d}) = \frac{\mathcal{K}\big(\frac{\boldsymbol{d}-\boldsymbol{d}_i}{h}\big)}{\sum_{j=1}^N \mathcal{K}\Big(\frac{\boldsymbol{d}-\boldsymbol{d}_j}{h}\Big)}, \qquad i = 1, \ldots, N, \tag{2}$$

where N is the number of observations available at the start of round t of BED, and $h > 0$ is a bandwidth parameter controlling the degree of localisation. Increasing h allows information to be borrowed from a wider neighbourhood of the target design $\boldsymbol{d}$, while decreasing h concentrates the weights on nearby design points. Therefore, h determines the bias–variance trade-off of the kernel smoothing procedure.

By construction, the weights satisfy $\sum_{i=1}^N w_i(\boldsymbol{d}) = 1$, hence defining an empirical distribution centred around the target design point $\boldsymbol{d}$. Building on this fact, to estimate the cumulative distribution of the true DGP, we introduce the kernel-weighted empirical cumulative distribution function (CDF) of the outcome y for the target design point $\boldsymbol{d}$ as follows:

$$\widehat{F}_w(y \mid \boldsymbol{d}) := \sum_{i=1}^N w_i(\boldsymbol{d})\, \mathbb{1}_{\{y_i \leq y\}}, \tag{3}$$

where the outcomes are ordered as $y_{(1)} < \cdots < y_{(N)}$ based on the pooled observations. To construct the functional confidence bands of $\mathscr{Q}_t$, we partition the outcome space into $M < N$ bins with approximately equal kernel-weighted probability mass. Larger values of M provide a more accurate estimate of the true CDF and reduce discretisation bias. We will show that constructing functional confidence bands within these bins produces a lower and upper bounding functions for the true DGP (Appendix A.5).We define the index $k_m(\boldsymbol{d})$ of the last outcome in bin m, as follows:

$$k_m(\boldsymbol{d}) = \max\Big\{j :\, \widehat{F}_w\big(y_{(j)} \mid \boldsymbol{d}\big) \leq \frac{m}{M}\Big\} \quad m \in \{2, \cdots, M-2\}, \tag{4}$$

meaning that bin m consists of the outcomes $\{y_{k_{m-1}(\boldsymbol{d})+1}, \ldots, y_{k_m(\boldsymbol{d})}\}$, for $m = 1, \ldots, M$. We set by convention that $k_0(\boldsymbol{d}) = 0$ and $k_M(\boldsymbol{d}) = N$, implying that the first and last bins are $\{y_1, \ldots, y_{k_1(\boldsymbol{d})}\}$ and $\{y_{k_{M-1}(\boldsymbol{d})+1}, \ldots, y_N\}$, respectively. By construction, each bin captures approximately $1/M$ quantile of the weighted empirical probability mass. We define the corresponding true bin masses and their empirical estimates as

$$\Delta_m(\boldsymbol{d}) = \tilde{Q}\big(y_{(k_m(\boldsymbol{d}))} \mid \boldsymbol{d}\big) - \tilde{Q}\big(y_{(k_{m-1}(\boldsymbol{d}))} \mid \boldsymbol{d}\big) \quad m \in \{2, \cdots, M\}, \tag{5}$$

$$\widehat{\Delta}_m(\boldsymbol{d}) = \widehat{F}_w\big(y_{(k_m(\boldsymbol{d}))} \mid \boldsymbol{d}\big) - \widehat{F}_w\big(y_{(k_{m-1}(\boldsymbol{d}))} \mid \boldsymbol{d}\big) \quad m \in \{2, \cdots, M\} \tag{6}$$

respectively, where $\tilde{Q}(.)$ is the CDF of the true DGP $\tilde{q}(.)$. By convention $\Delta_1(\boldsymbol{d}) = \tilde{Q}\big(y_{(k_1(\boldsymbol{d}))} \mid \boldsymbol{d}\big)$ and $\Delta_{M+1}(\boldsymbol{d}) = 1 - \tilde{Q}\big(y_{(k_M(\boldsymbol{d}))} \mid \boldsymbol{d}\big)$. To calibrate the uncertainty associated with the weighted empirical process, we employ a weighted Bayesian bootstrap technique (Chen et al., 2022). Following this technique, let $X_1, \ldots, X_N$ denote the *i.i.d.* samples drawn from an exponential distribution with rate $\lambda = 1$. These samples induce bootstrap weights that approximate the finite-sample distribution of the binned empirical process:

$$\tilde{\Delta}_m(\boldsymbol{d}) := \frac{\sum_{j=k_{m-1}(\boldsymbol{d})+1}^{k_m(\boldsymbol{d})} w_j(\boldsymbol{d}) X_j}{\sum_{l=1}^{N} w_l(\boldsymbol{d}) X_l} \quad m \in \{2, \cdots, M\} \tag{7}$$

where the denominator serves as a normalising constant, ensuring that the masses across all bins sum to one. To gain intuition for this sampling scheme, note that, by the probability integral transform (Angus, 1994), the random variable $z = \tilde{Q}\big(y_{(k_m(\boldsymbol{d}))} \mid \boldsymbol{d}\big)$ is uniformly distributed. To sample from the uniform distribution, we use the inverse-CDF method (Chen et al., 2022), which directly generates ordered samples.

Definition 3.1: Weighted Bootstrap Calibration

For any confidence level $1 - \alpha_1$, let $c^-(\alpha_1)$ and $c^+(\alpha_1)$ denote the calibration constants for the weighted empirical bin mass. These bands are obtained by repeatedly simulating the bootstrap variables X_i, and using the kernel weights $\{w_i(\boldsymbol{d})\}_{i=1}^N$, such that the following probabilistic inequality holds:

$$\mathbb{P}\Big(c^-(\alpha_1) \leq \tilde{\Delta}_m(\boldsymbol{d}) \leq c^+(\alpha_1),\ \forall m = 1, \ldots, M\Big) \geq 1 - \alpha_1 \tag{8}$$

While the above calibration procedure provides uncertainty bands for the weighted empirical bin masses, connecting it to the true DGP requires probabilistic concentration bounds on the deviation between the true and empirical bin masses. To derive these bounds, we make the following smoothness assumption on the conditional density over the design space.

Assumption 2 (Hölder Continuity in the Design Space (Jiang, 2017))**.** *There exist constants $L > 0$ and $\gamma_D \in (0, 1]$ such that*

$$\sup_{y \in [a,b]} |\tilde{q}(y \mid \boldsymbol{d}) - \tilde{q}(y \mid \boldsymbol{d}')| \leq L \|\boldsymbol{d} - \boldsymbol{d}'\|^{\gamma_D} \qquad \forall\, \boldsymbol{d}, \boldsymbol{d}' \in \mathscr{D} \tag{9}$$

The above assumption ensures that local observations become progressively more representative of the true DGP at the target design $\boldsymbol{d}$. This locality property underpins the concentration result established in the following theorem.

Theorem 3.1: Conditional Coverage

Given the Nadaraya-Watson weighting scheme built according to Assumption 1, assume that the conditional density satisfies the Hölder continuity in the design space (Assumption 2). For sufficient number of samples, the following inequality holds,

$$\mathbb{P}(\tilde{q}(\cdot \mid \boldsymbol{d}) \in \mathscr{Q}_t(y \mid \boldsymbol{d}, \alpha)) \geq 1 - \alpha \tag{10}$$

where $\alpha > \alpha_1$ is a user specified significance level and α_1 is used for calibration of weighted empirical bin mass (see Definition 3.1). $\mathscr{Q}_t(y \mid \boldsymbol{d}, \alpha) = (l_\alpha^t(y \mid \boldsymbol{d}), u_\alpha^t(y \mid \boldsymbol{d}))$ is the ambiguity set at round t of BED.

Optimisation of confidence bands To optimise the confidence bands $\big(l_\alpha^t(y \mid \boldsymbol{d}), u_\alpha^t(y \mid \boldsymbol{d})\big)$, we first construct the polyhedral uncertainty set $\mathscr{H}_w(\boldsymbol{d}, \alpha)$ of admissible unimodal densities. Following Chen et al. (2022), we show in Appendix A.3 that the lower and upper confidence bands within each bin are step functions, implying that the functional bands of the ambiguity set are piecewise-constant densities. We therefore represent candidate densities by their values on a partition induced by the

outcome-space boundaries, the mode, the target outcome, and the bin endpoints. The resulting finite-dimensional feasible set is characterised by linear constraints enforcing unimodality, calibrated bin-mass bounds, normalisation, and the upper confidence bound on the true density. The complete linear programming formulation is provided in Appendix A.3.1. For any $(y, \boldsymbol{d})$, the lower and upper confidence bands are obtained by solving linear programs over $\mathscr{H}_w(\boldsymbol{d}, \alpha)$ that minimise and maximise the density value at y, respectively. Denoting the resulting extrema by $\widetilde{l}_{\alpha_1}(y \mid \boldsymbol{d})$ and $\widetilde{u}_{\alpha_1}(y \mid \boldsymbol{d})$, we define

$$l^t_\alpha(y \mid \boldsymbol{d}) = \max\Big\{0, \widetilde{l}_{\alpha_1}(y \mid \boldsymbol{d}) - Lh^{\gamma_D}\Big\}, \qquad u^t_\alpha(y \mid \boldsymbol{d}) = \widetilde{u}_{\alpha_1}(y \mid \boldsymbol{d}) + Lh^{\gamma_D}. \tag{11}$$

The resulting algorithmic procedure is summarised in Algorithm 2 in Appendix A.5.

3.2 Dual reformulation of an experimental design utility

Standard BED optimises *Expected Information Gain* at each round of experimentation, which is effective when the model is well-specified. Under model misspecification, however, each experimental round should take into account how closely the behaviour of the misspecified model matches that of the unknown true DGP. We achieve that in robust BED by recommending design points which are guaranteed to drive the misspecified model's behaviour as close as possible to the true DGP.

We propose a distributionally robust acquisition function, to be used in place of EIG of standard BED, which uses the midpoint of the constructed confidence band as a reference predictive density,

$$\hat{q}(y \mid \boldsymbol{d}) = \frac{1}{2}\left(l_\alpha(y \mid \boldsymbol{d}) + u_\alpha(y \mid \boldsymbol{d})\right) \tag{12}$$

Here, $\hat{q}(y \mid \boldsymbol{d})$ is constructed independently of the candidate density $q(y \mid \boldsymbol{d})$. We define the corresponding integrand as

$$f(\boldsymbol{d}, y) := \mathbb{E}_{\theta\sim p(\theta)}[\log p(y \mid \theta, \boldsymbol{d})] - \log \hat{q}(y \mid \boldsymbol{d}) \tag{13}$$

For any candidate $q(y \mid \boldsymbol{d})$ in the ambiguity set $\mathcal{Q}$, the utility of experiment $\boldsymbol{d}$ is

$$U(\boldsymbol{d}, q) := \mathbb{E}_{y\sim q(y|\boldsymbol{d})}[f(\boldsymbol{d}, y)] \tag{14}$$

which yields the *Distributionally Robust surrogate Expected Information Gain* (DR-EIG):

$$\text{DR-EIG}(\boldsymbol{d}) := \inf_{q\in\mathcal{Q}} U(\boldsymbol{d}, q) \tag{15}$$

The robust Bayesian experimental design problem is therefore

$$\boldsymbol{d}^\star \in \arg\max_{\boldsymbol{d}\in\mathcal{D}} \text{DR-EIG}(\boldsymbol{d}) \tag{16}$$

To solve this optimisation problem, we make the following regularity assumption:

Assumption 3. *For every $\boldsymbol{d} \in \mathscr{D}$,*

- *The utility is integrable: $\int_a^b |f(\boldsymbol{d}, y)|\, u_\alpha(y \mid \boldsymbol{d})\, dy < \infty$;*
- *The ambiguity set is non-empty for the given conservativeness level $1-\alpha$: $\mathscr{Q} \neq \emptyset$.*

Appendix A.4 shows that conducting the experiment with design $\boldsymbol{d}^\star$, rather than the design $\widetilde{\boldsymbol{d}}$ that is optimal under the true DGP, incurs regret bounded by the width of the ambiguity set. Appendix A.6 shows that through conic duality reformulation, the objective of the inner optimisation problem admits the following representation:

$$F(\boldsymbol{d}, \lambda) = \lambda - \int (\lambda + f(\boldsymbol{d}, y))^+ l_\alpha(y \mid \boldsymbol{d})\, dy + \int (-\lambda - f(\boldsymbol{d}, y))^+ u_\alpha(y \mid \boldsymbol{d})\, dy \tag{17}$$

where $(z)^+ = \max\{z, 0\}$. The corresponding optimization problem is

$$v^\star = -\inf_{\boldsymbol{d}\in\mathscr{D},\lambda} F(\boldsymbol{d}, \lambda), \tag{18}$$

Therefore, maximising the worst-case expected utility is equivalent to minimising the dual objective over the design variable $\boldsymbol{d}$ and the scalar dual parameter λ.

3.3 STOCHASTIC GRADIENT METHOD FOR DR-BED

Solving the dual optimisation problem in 18 using stochastic gradient method requires the subgradients $\partial_{\boldsymbol{d}} F(\cdot)$, and $\partial_\lambda F(\cdot)$. We first derive their analytical forms, and then, using the importance sampling identity method, estimate them through a Monte-Carlo method. Stochastic Gradient algorithm uses these estimate in each iteration of algorithm to update $\boldsymbol{d}$ and λ. Differentiating the dual objective with respect to $\boldsymbol{d}$, and λ are as follows:

$$\begin{aligned}\partial_{\boldsymbol{d}} F(\boldsymbol{d},\lambda) \ni &-\int_a^b l_\alpha(y \mid \boldsymbol{d}) f'(\boldsymbol{d},y) \mathbb{1}_{\{-f(\boldsymbol{d},y)\le\lambda\}}\, dy - \int_a^b u_\alpha(y \mid \boldsymbol{d}) f'(\boldsymbol{d},y) \mathbb{1}_{\{-f(\boldsymbol{d},y)\ge\lambda\}}\, dy \\ &- \int_a^b \partial_{\boldsymbol{d}} l_\alpha(y \mid \boldsymbol{d}) \big(\lambda + f(\boldsymbol{d},y)\big)^+ dy + \int_a^b \partial_{\boldsymbol{d}} u_\alpha(y \mid \boldsymbol{d}) \big(-\lambda - f(\boldsymbol{d},y)\big)^+ dy\end{aligned} \tag{19}$$

$$\partial_\lambda F(\boldsymbol{d},\lambda) \ni 1 - \int_a^b l_\alpha(y \mid \boldsymbol{d}) \mathbb{1}_{\{-f(\boldsymbol{d},y)\le\lambda\}}\, dy - \int_a^b u_\alpha(y \mid \boldsymbol{d}) \mathbb{1}_{\{-f(\boldsymbol{d},y)\ge\lambda\}}\, dy \tag{20}$$

$F(\boldsymbol{d},\lambda)$ in equation 17 depends on $\boldsymbol{d}$ in two ways, through the utility integrand function $f(\boldsymbol{d},y)$ and through the confidence bounds $l_\alpha(y \mid \boldsymbol{d})$ and $u_\alpha(y \mid \boldsymbol{d})$. Thus, the first two terms in the subgradient $\partial_{\boldsymbol{d}} F(\boldsymbol{d},\lambda)$ correspond to the variation in the objective function; the remaining terms are the bias correction terms induced by the movement of ambiguity set. In contrast, the subgradient $\partial_\lambda F(\boldsymbol{d},\lambda)$ does not depend on the partial derivatives of the confidence bands or the integrand function.

Under Assumption 3, let $y_1,\dots,y_B$ be uniformly distributed *i.i.d.* samples from a box $I \supseteq [a,b]$. Appendix A.6 shows that minibatch stochastic subgradients of F at $(\boldsymbol{d}_k,\lambda_k)$ are given as follows:

$$g_{\boldsymbol{d}} = \frac{|I|}{B}\sum_{i=1}^{B}\Big[\omega_i f'(\boldsymbol{d}_k,y_i) - \partial_{\boldsymbol{d}} l_\alpha(y_i \mid \boldsymbol{d}_k)\big(\lambda_k + f(\boldsymbol{d}_k,y_i)\big)^+ + \partial_{\boldsymbol{d}} u_\alpha(y_i \mid \boldsymbol{d}_k)\big(-\lambda_k - f(\boldsymbol{d}_k,y_i)\big)^+\Big], \tag{21}$$

$$g_\lambda = 1 + \frac{|I|}{B}\sum_{i=1}^{B}\omega_i, \tag{22}$$

where $\omega_i = \begin{cases} -l_\alpha(y_i \mid \boldsymbol{d}_k), & -f(\boldsymbol{d}_k,y_i) < \lambda_k \\ -u_\alpha(y_i \mid \boldsymbol{d}_k), & -f(\boldsymbol{d}_k,y_i) > \lambda_k \end{cases}$. The derivatives $\partial_{\boldsymbol{d}} l_\alpha(y_i \mid \boldsymbol{d}_k)$ and $\partial_{\boldsymbol{d}} u_\alpha(y_i \mid \boldsymbol{d}_k)$ are available through the dual multiplier of the LP problems in Equation 93, presented in more detail in Appendix A.6. The full procedure for the proposed stochastic optimisation algorithm is given in Appendix A.6, Algorithm 2.

We illustrate the behaviour of the algorithm on the source-location problem in Figure 2; experimental details are provided in Section 3.4. The left panel shows the normalised signal concentration over the design space. The subgradients guide the design toward regions of higher concentration, ultimately moving it closer to the peaks corresponding to the two sources. The right panel shows the observed concentrations (green dots) together with the predicted concentrations (blue dots). The red arrows denote the subgradients, which steer the midpoint towards a location closer to source 2.

3.4 EXPERIMENTS

In this section, we empirically validate Theorem 3.1 and compare the design efficiency of our proposed ***DR-BED*** algorithm with alternative methods. We choose **Bayesian Adaptive Design (BAD)** (Foster et al., 2019) and **Random** design selection strategies as the baselines for comparison purpose. We consider three well-known increasingly difficult experimental settings as our benchmark problems in BED, including (1) toy Polynomial Regression problem, (2) Pharmacokinetics (PK), a realistic dynamical system, and (3) Source Location Finding, a challenging nonlinear spatial system. We defer the details of the alternative methods, experimental settings, definitions of performance measures, and additional experimental results to Appendix A.7.

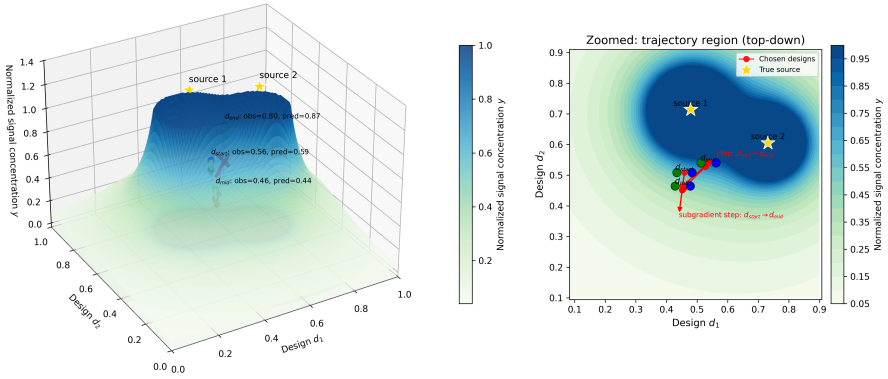


Figure 2: **SG-DRO algorithm steps for source location finding problem**: left panel shows the normalised signal concentration over the design space, and right panel shows the observed (greed dots) and predicted concentrations (blue dots).

Table 1: Comparison of methods after ten rounds of experiments.

	Method	Misspecified				Well-Specified			
		MMD	MSE	W1	NLL	MMD	MSE	W1	NLL
Poly	Rand	0.0983 (0.0229)	0.3001 (0.0902)	0.2995 (0.0457)	-1.1634 (0.1057)	0.0310 (0.0126)	0.2469 (0.0472)	0.1665 (0.0300)	-1.3870 (0.0498)
	BAD	0.0570 (0.0147)	0.3091 (0.0515)	0.2315 (0.0371)	-1.2964 (0.0550)	0.0242 (0.0055)	0.2448 (0.0260)	0.1560 (0.0176)	-1.4095 (0.0194)
	DR-BED	**0.0173 (0.0017)**	**0.0965 (0.0100)**	**0.0962 (0.0072)**	**-1.4808 (0.0080)**	**0.0153 (0.0014)**	**0.1157 (0.0098)**	**0.0865 (0.0060)**	**-1.4871 (0.0047)**
PK	Rand	0.0522 (0.0033)	5.0734 (2.0198)	0.7509 (0.2016)	-1.9916 (0.1832)	0.0321 (0.0135)	2.6380 (0.9239)	0.5572 (0.1264)	-1.6717 (0.1374)
	BAD	0.0358 (0.0037)	3.3916 (1.3225)	0.6929 (0.1552)	-2.0890 (0.1702)	0.0259 (0.0056)	1.4938 (0.6148)	0.4071 (0.0757)	-2.1121 (0.1435)
	DR-BED	**0.0229 (0.0012)**	**0.3711 (0.0422)**	**0.2993 (0.0068)**	**-2.2175 (0.0219)**	**0.0101 (0.0019)**	**0.3507 (0.0313)**	**0.1398 (0.0115)**	**-2.3672 (0.0258)**
SL	Rand	0.3554 (0.0954)	**0.0411 (0.0200)**	0.1371 (0.0310)	3.3525 (0.8238)	0.1317 (0.0462)	**0.0278 (0.0129)**	**0.0767 (0.0181)**	-0.4368 (0.6093)
	BAD	0.2647 (0.1253)	0.0647 (0.0392)	0.1321 (0.0484)	0.5653 (0.9418)	0.2878 (0.0786)	0.0543 (0.0174)	0.1304 (0.0252)	0.5706 (0.8608)
	DR-BED	**0.0827 (0.0267)**	0.0431 (0.0091)	**0.1234 (0.0221)**	**-1.4180 (0.0649)**	**0.0620 (0.0130)**	0.0418 (0.0087)	0.0923 (0.0056)	**-1.5597 (0.0449)**

3.5 Comparison with BAD

This section presents the results of the proposed *Distributionally Robust BED* (***DR-BED***) framework across benchmark problems under both well-specified and misspecified model scenarios. As shown in Table 1, in the well-specified case, DR-BED performs comparably to or better than the baselines. In the misspecified case, DR-BED achieves the lowest MMD while providing robustness guarantees at the user-specified confidence level of 90%, indicating that its selected designs yield a model whose behaviour more closely matches the true DGP despite the misspecified likelihood. Moreover, its lower MSE and NLL reflect more accurate point estimates together with better-calibrated predictive uncertainty, as NLL heavily penalises confident predictions that miss the true observations. Further details are provided in Appendix A.7.

Figure 3 shows coverage and band-width trends for several misspecified BED problems. The top row presents per-round coverage strip charts, where the round means (blue dots) typically exceed the target confidence level of 90%. The bottom row shows the corresponding band widths decreasing over time, demonstrating that the ambiguity set becomes progressively tighter as data are collected.

3.6 Comparison with the GBOED

To evaluate the effectiveness of the proposed algorithm, we compare DR-BED against GBOED (Barlas et al., 2025), one of the strong baselines for misspecified BED, see Appendix A.8 for details. Table 2 summarise the performance of the considered methods after 20 rounds of experiments on the source location finding problem under different dimensions.

In the four-dimensional setting, DR-BED improves on GBOED across all metrics, with lower MMD, W1, and MSE, and more than a twofold reduction in NLL. These results indicate both a closer match to the true DGP and substantially better predictive distributions under model misspecification. In the eight-dimensional setting, the two methods achieve similar MMD. GBOED attains lower MSE and W1, although the W1 difference is within the combined standard errors, while DR-BED achieves a markedly lower NLL (174.7 versus 303.9). This suggests that GBOED yields slightly more accurate

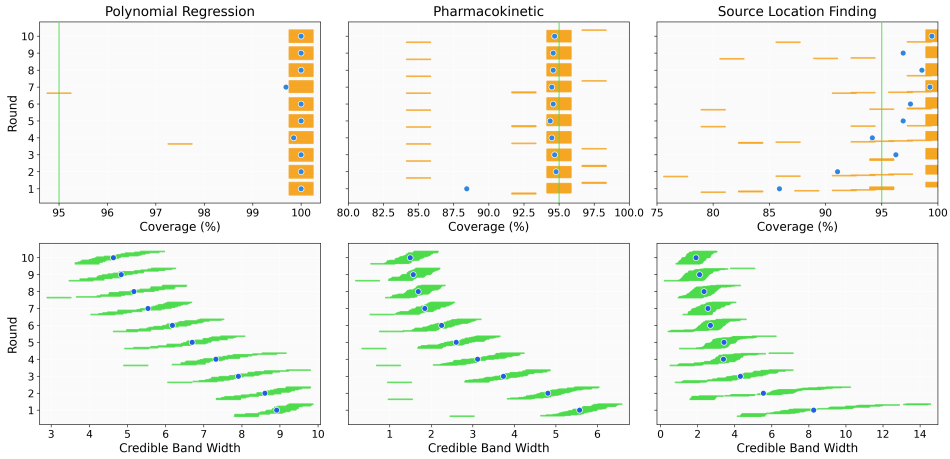


Figure 3: **Confidence level and Bandwidth trend**: Top row shows the coverage per round strip charts for different misspecified BED problems. Bottom row shows the corresponding band width shrinking as rounds progress, showing a tighter ambiguity set over the sequential design process.

Table 2: Results on the source location finding problem with misspecified likelihood models.

Dim.	Method	MMD	MSE	W1	NLL
$d = 4$	GBOED	1.033 (0.070)	1.762 (0.25)	1.157 (0.098)	199.1 (17)
	DR-BED	**0.9404 (0.035)**	**1.059 (0.14)**	**0.885 (0.070)**	**89.31 (12)**
$d = 8$	GBOED	0.8254 (0.022)	**0.5913 (0.048)**	**0.6199 (0.032)**	303.9 (9.5)
	DR-BED	**0.8197 (0.029)**	0.7774 (0.12)	0.7008 (0.062)	**174.7 (17)**

point predictions, whereas DR-BED produces better-calibrated predictive uncertainty, consistent with its distributionally robust formulation.

The computational time analysis in A.9 shows that the improved robustness comes at an additional computational cost. Since DR-BED optimises the worst case over an ambiguity set rather than a single nominal model, it adds an inner subgradient descent loop to each round of experimentation. However, this additional cost is reasonable in practical BED settings, especially when the likelihood function is misspecified.

4 CONCLUSION

We proposed a Distributionally Robust Bayesian Experimental Design (***DR-BED***) framework to address the model misspecification problem. To this end, ***DR-BED*** uses confidence-band ambiguity sets constructed from a pair of functional lower and upper bounds. We established a theoretical coverage guarantee showing that our proposed framework for ambiguity set construction contains the true DGP with a confidence level of $1 - \alpha$. To implement the theoretical insights and propose a practical method for robust design within the sequential BED, we develop a subgradient descent algorithm. Empirical results, aligned with related work, suggest that while the standard BED framework is fragile under model misspecification, our proposed ***DR-BED*** framework enables guided exploration of the design space, recommending designs that are supported by a theoretical guarantee at the confidence level. Such a guarantee is essential for safety-critical applications such as healthcare. In terms of limitations, this work tested exclusively myopic design selection. Future work can build on BED amortisation techniques, extending ***DR-BED*** to scalable design selection under non-myopic settings. Another challenge related to scalability is that the present study considers only low-dimensional outcomes. Adapting the methodology to high-dimensional experimental outcome spaces is therefore an important direction for future work.

AI USE STATEMENT

We used generative AI tools exclusively for editing the research paper to improve readability. The authors have reviewed all AI-assisted work, confirming that it is correct and it does not alter the meaning of text. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work targets generic problems of Bayesian experimental design and builds upon existing baselines and evaluation methodologies. Our studies do not involve human subjects beyond the research team, and we do not release any datasets. This study refers to the theoretical use of Bayesian experimental design, and precautions similarly to the other literature on this topic, will depend upon the applications to practical domains. We have no conflict of interests to declare. We do not have specific privacy issues to report.

REPRODUCIBILITY STATEMENT

We are committed to publicly releasing the source code upon eventual acceptance of the paper. Proofs of all theorems are included in Appendices A.3-A.4, the algorithm in Appendices A.5-A.6. We report the experimental details in Appendix A.7.

REFERENCES

Hany Abdulsamad, Sahel Iqbal, Christian A. Naesseth, Takuo Matsubara, and Adrien Corenflos. Maximin robust Bayesian experimental design. *CoRR*, abs/2603.14094, March 2026. URL https://doi.org/10.48550/arXiv.2603.14094.

Philip Amortila, Tongyi Cao, and Akshay Krishnamurthy. Mitigating covariate shift in misspecified regression with applications to reinforcement learning. In *The Thirty Seventh Annual Conference on Learning Theory*, pp. 130–160. PMLR, 2024.

John E Angus. The probability integral transform and related results. *SIAM review*, 36(4):652–654, 1994.

Yasir Zubayr Barlas, Sabina J Sloman, and Samuel Kaski. Robust experimental design via generalised bayesian inference. *arXiv preprint arXiv:2511.07671*, 2025.

Dimitris Bertsimas, Melvyn Sim, and Meilin Zhang. Adaptive distributionally robust optimization. *Management Science*, 65(2):604–618, 2019.

Zongwu Cai. Weighted Nadaraya-Watson regression estimation. *Statistics & probability letters*, 51 (3):307–318, 2001.

Michele Caprio, Maryam Sultana, Eleni G Elia, and Fabio Cuzzolin. Credal learning theory. *Advances in Neural Information Processing Systems*, 37:38665–38694, 2024.

Tommie A Catanach and Niladri Das. Metrics for bayesian optimal experiment design under model misspecification. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pp. 7707–7714. IEEE, 2023.

Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical science*, pp. 273–304, 1995.

Xi Chen, Qihang Lin, and Guanglin Xu. Distributionally robust optimization with confidence bands for probability density functions. *INFORMS Journal on Optimization*, 4(1):65–89, 2022.

A DasGupta and WJ Studden. Robust bayesian experimental designs in normal linear models. *The Annals of Statistics*, 19(3):1244–1256, 1991.

Chi Feng. *Optimal Bayesian experimental design in the presence of model error*. PhD thesis, Massachusetts Institute of Technology, 2015.

Alexander J Forster, Desi R Ivanova, and Tom Rainforth. Improving robustness to model misspecification in bayesian experimental design. In *Workshop at the 7th Symposium on Advances in Approximate Bayesian Inference (non-archival)*, volume 1, pp. 14, 2025.

Adam Foster, Martin Jankowiak, Elias Bingham, Paul Horsfall, Yee Whye Teh, Thomas Rainforth, and Noah Goodman. Variational Bayesian optimal experimental design. *Advances in neural information processing systems*, 32, 2019.

Rui Gao and Anton Kleywegt. Distributionally robust stochastic optimization with Wasserstein distance. *Mathematics of Operations Research*, 48(2):603–655, 2023.

Jinwoo Go and Tobin Isaac. Robust expected information gain for optimal bayesian experimental design using ambiguity sets. In *Uncertainty in Artificial Intelligence*, pp. 728–737. PMLR, 2022.

Joel Goh and Melvyn Sim. Distributionally robust optimization and its tractable approximations. *Operations research*, 58(4-part-1):902–917, 2010.

Zhaolin Hu and L Jeff Hong. Kullback-Leibler divergence constrained distributionally robust optimization. *Available at Optimization Online*, 1(2):9, 2013.

Daolang Huang, Ayush Bharti, Amauri Souza, Luigi Acerbi, and Samuel Kaski. Learning robust statistics for simulation-based inference under model misspecification. *Advances in neural information processing systems*, 36:7289–7310, 2023.

Felipe Igea and Alice Cicirello. An interacting Wasserstein gradient flow strategy to robust bayesian inference for application to decision-making in engineering. *Data-Centric Engineering*, 6:e19, 2025.

Yu Inatsu, Shogo Iwazaki, and Ichiro Takeuchi. Active learning for distributionally robust level-set estimation. In *International Conference on Machine Learning*, pp. 4574–4584. PMLR, 2021.

Bai Jiang, Qiang Sun, and Jianqing Fan. Bernstein's inequalities for general markov chains. *arXiv preprint arXiv:1805.10721*, 2018.

Heinrich Jiang. Uniform convergence rates for kernel density estimation. In *International Conference on Machine Learning*, pp. 1694–1703. PMLR, 2017.

Johannes Kirschner, Ilija Bogunovic, Stefanie Jegelka, and Andreas Krause. Distributionally robust bayesian optimization. In *International Conference on Artificial Intelligence and Statistics*, pp. 2174–2184. PMLR, 2020.

Daniel Levy, Yair Carmon, John C Duchi, and Aaron Sidford. Large-scale methods for distributionally robust optimization. *Advances in neural information processing systems*, 33:8847–8860, 2020.

Colin McDiarmid et al. On the method of bounded differences. *Surveys in combinatorics*, 141(1): 148–188, 1989.

Linhan Ouyang, Jianxiong Chen, Yizhong Ma, Chanseok Park, and Jionghua Jin. Bayesian closed-loop robust process design considering model uncertainty and data quality. *IISE Transactions*, 52 (3):288–300, 2020.

Antony Overstall and James McGree. Bayesian decision-theoretic design of experiments under an alternative model. *Bayesian Analysis*, 17(4):1021–1041, 2022.

Antony M Overstall, Jacinta Holloway-Brown, and James M McGree. Gibbs optimal design of experiments. *Statistical Science*, 41(3):582–597, 2026.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

Marek Petrik and Reazul Hasan Russel. Beyond confidence regions: Tight Bayesian ambiguity sets for robust MDPs. *Advances in neural information processing systems*, 32, 2019.

Hamed Rahimian and Sanjay Mehrotra. Frameworks and results in distributionally robust optimization. *Open Journal of Mathematical Optimization*, 3:1–85, 2022.

Tom Rainforth, Adam Foster, Desi R Ivanova, and Freddie Bickford Smith. Modern Bayesian experimental design. *Statistical Science*, 39(1):100–114, 2024.

Paola Sebastiani and Henry P Wynn. Maximum entropy sampling and optimal Bayesian experimental design. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(1): 145–157, 2000.

Aman Sinha, Hongseok Namkoong, Riccardo Volpi, and John Duchi. Certifying some distributional robustness with principled adversarial training. *arXiv preprint arXiv:1710.10571*, 2017.

Roubing Tang, Sabina J Sloman, and Samuel Kaski. Representative, informative, and de-amplifying: Requirements for robust Bayesian active learning under model misspecification. In *The 29th International Conference on Artificial Intelligence and Statistics*, 2026.

Blaza Toman and Joseph L Gastwirth. Robust bayesian experimental design and estimation for analysis of variance models using a class of normal mixtures. *Journal of statistical planning and inference*, 35(3):383–398, 1993.

Stephen G Walker. Bayesian inference with misspecified models. *Journal of statistical planning and inference*, 143(10):1621–1633, 2013.

Jiaqi Wen and Jianyi Yang. Distributionally robust optimization via generative ambiguity modeling. In *The Fourteenth International Conference on Learning Representations*, 2026. URL `https://openreview.net/forum?id=q67t0gFrdY`.

Junkang Wu, Jiawei Chen, Jiancan Wu, Wentao Shi, Xiang Wang, and Xiangnan He. Understanding contrastive learning via distributionally robust optimization. *Advances in Neural Information Processing Systems*, 36:23297–23320, 2023.

A Appendix

Contents

A.1 Related works 13
A.2 Sequential Bayesian Experimental Design 13
A.3 Proof for coverage theorem 14
A.3.1 Optimisation of confidence bands 23
A.4 Proof for Regret guarantee 23
A.5 Algorithm 1 (Conditional Kernel-Weighted Shape-Restricted Bands) 24
A.6 Algorithm 2 (Stochastic Gradient Method for Robust BED) 24
A.7 Experimental Details 30
A.7.1 Competing methods 30
A.7.2 Performance measures 31
A.7.3 Experimental setting 32
A.7.4 Additional Experimental results 34
A.8 DR-BED Performance in Higher-Dimensional Settings 47
A.9 Deployment Time Analysis 47

A.1 Related works

Model misspecification in Bayesian experimental design. Classical Bayesian experimental design (BED) seeks experiments that maximise an expected utility under a specified probabilistic model (Chaloner & Verdinelli, 1995). The quality of the resulting design therefore depends critically on the correctness of the assumed outcome model. In practice, however, we encounter model misspecification, motivating a growing literature on robust BED under model misspecification. Early work investigated optimal design under model uncertainty and robustness considerations (DasGupta & Studden, 1991; Toman & Gastwirth, 1993; Feng, 2015), while more recent approaches have developed Bayesian design criteria that remain effective across multiple candidate models or under distributional uncertainty (Overstall & McGree, 2022; Go & Isaac, 2022; Abdulsamad et al., 2026; Forster et al., 2025). Catanach & Das (2023) investigate alternative information metrics that are less sensitive to modelling assumptions. Barlas et al. (2025) propose a robust design objective by swapping the likelihood function with the Gibbs-expected information Gain. Overstall et al. (2026) develop Gibbs posterior approaches to experimental design, replacing exact likelihood assumptions with generalised Bayesian updating rules. While the approaches mentioned above solve the problem of robustness, none of them present the problem as a form of distributionally robust optimisation, which would allow to explicitly provide guarantees for a given significance level.

Closely related questions have also been studied in active learning settings. Huang et al. (2023) study learning and decision-making under model uncertainty; Tang et al. (2026) consider robust Bayesian active learning under model misspecification, proposing acquisition strategies that remain informative even when the learning model differs from the true DGP. Inatsu et al. (2021) consider the active learning setting in face of uncontrollable environment variables with unknown distribution, and employ robust active learning model for the worst-case probability threshold robustness measure.

Distributionally robust optimisation and experimental design. A related direction replaces uncertainty over models with uncertainty over data-generating distributions. Distributionally robust optimisation (DRO) (Goh & Sim, 2010; Rahimian & Mehrotra, 2022) seeks decisions that perform well against the worst-case distribution within a prescribed ambiguity set, providing robustness guarantees against distributional shift and modelling errors (Bertsimas et al., 2019; Wu et al., 2023; Levy et al., 2020; Kirschner et al., 2020). Petrik & Russel (2019) employ ambiguity sets for modelling robust Markov Decision processes. Recent work has begun to connect these ideas to Bayesian decision-making and experimental design. For example, Caprio et al. (2024) formulate Bayesian decision problems using credal sets to represent distributional ambiguity, while Wen & Yang (2026) develop distributionally robust BED frameworks that optimise acquisition policies against worst-case distributions. Similar robust formulations have been explored using Wasserstein-ball ambiguity sets (Gao & Kleywegt, 2023) and divergence-based uncertainty sets, such as those defined through Kullback-Leibler constraints (Hu & Hong, 2013). Collectively, these approaches relax the assumption that a single probabilistic model accurately describes reality and instead seek experimental strategies that remain effective across a family of plausible distributions. This work contributes to addressing the problem by employing distributionally robust optimisation via confidence-band ambiguity sets.

A.2 Sequential Bayesian Experimental Design

Sequential Bayesian Experimental Design (SBED) is a machine learning framework that proposes a method for the efficient sampling of experiments. Its goal is to find the best descriptive parametric model of the experimental outcomes under an affordable experimental budget (Rainforth et al., 2024). In the case of a well-specified likelihood model, where the experimental outcomes are assumed to be generated from an unknown true DGP $\tilde{q}(y \mid \boldsymbol{d})$, and a parametric surrogate of the form $p(y \mid \boldsymbol{d}, \boldsymbol{\theta})$ is available, SBED aims to choose the optimal informative design points in the design space of the assumed likelihood model. Following the Bayesian inference framework, as the true parameters $\boldsymbol{\theta}^*$ are unknown, the experimenter specifies prior knowledge in the form of a prior distribution over the parameter space $p(\boldsymbol{\theta})$. Then, at each round of the experiment, the learner seeks the design point $\boldsymbol{d}$ that maximises the expected reduction in uncertainty about the component of the experimental-outcome model that can be explained by the parameters.

At each round of the experiment, a new observation is collected and used to update the belief over the predictive model. To conduct the experimental process efficiently, the experimenter specifies a utility function $U(\boldsymbol{d})$, and the optimal design at each round is recommended by maximising it, $\boldsymbol{d}^\star = \arg\max_{\boldsymbol{d}\in\mathscr{D}} U(\boldsymbol{d})$. Traditionally, BED uses the *Expected Information Gain (EIG)*, as the utility function, which quantifies the expected uncertainty in the experimental outcome that can be explained by the model parameters $\boldsymbol{\theta}$, through performing an experiment with design vector $\boldsymbol{d}$ (Forster et al., 2025; Tang et al., 2026).

Definition A.1: Expected Information Gain

The *Expected Information Gain (EIG)* of conducting an experiment with design $\boldsymbol{d}$ is defined as the expected Kullback–Leibler (KL) divergence between the posterior and prior belief regarding the model parameters:

$$U(\boldsymbol{d}) = \mathbb{E}_{p(y|\boldsymbol{d})}\left[D_{\mathrm{KL}}\left(p(\boldsymbol{\theta} \mid y, \boldsymbol{d}) \,\|\, p(\boldsymbol{\theta})\right)\right] \tag{23}$$

The expectation is taken with respect to the predictive distribution $p(y \mid \boldsymbol{d}) = \int p(y \mid \boldsymbol{d}, \boldsymbol{\theta})\, p(\boldsymbol{\theta})\, d\boldsymbol{\theta}$. Furthermore, the *KL divergence* between a target distribution $p(\boldsymbol{\theta} \mid y, \boldsymbol{d})$ and a reference distribution $p(\boldsymbol{\theta})$ is defined as

$$D_{\mathrm{KL}}(p(\boldsymbol{\theta} \mid y, \boldsymbol{d}) \,\|\, p(\boldsymbol{\theta})) = \int p(\boldsymbol{\theta} \mid y, \boldsymbol{d}) \log \frac{p(\boldsymbol{\theta} \mid y, \boldsymbol{d})}{p(\boldsymbol{\theta})}\, d\boldsymbol{\theta} \tag{24}$$

where $p(\boldsymbol{\theta} \mid y, \boldsymbol{d})$ is absolutely continuous with respect to $p(\boldsymbol{\theta})$ (Barlas et al., 2025).

Illustrative Example of model misspecification. Consider a clinical trial scenario where the goal is to find a predictive model for a patient's experimental outcome $y \in \mathcal{Y}$ based on a selected treatment $\boldsymbol{d} \in \mathcal{D}$ and patient characteristics (e.g., age, medical history, and symptoms). Assume that the true treatment-outcome relationship is governed by an unknown DGP $\tilde{q}(y \mid \boldsymbol{d})$. However, due to limited clinical knowledge and incomplete understanding of patient-specific biological mechanisms, the clinical researcher relies on an assumed probabilistic model $p(y \mid \boldsymbol{d}, \boldsymbol{\theta})$ with uncertain parameters $\boldsymbol{\theta}$. Explicitly accounting for the model misspecification, it is desirable for the researcher to have a confidence level $1 - \alpha$ on the effectiveness of the adaptively selected treatment design decisions $\boldsymbol{d}$ at each round of experiments.

A.3 Proof for coverage theorem

Theorem A.1: Conditional Coverage

Suppose the Nadaraya-Watson weighting scheme, and assume that the conditional density satisfies the Hölder continuity in the design space. For sufficient number of samples, the following inequality holds,

$$\mathbb{P}(\tilde{q}(\cdot \mid \boldsymbol{d}) \in \mathscr{Q}_t(y \mid \boldsymbol{d}, \alpha)) \geq 1 - \alpha \tag{25}$$

where $\alpha > \alpha_1$ is user specified significance level and α_1 is used for calibration of weighted empirical group mass. $\mathscr{Q}_t(y \mid \boldsymbol{d}, \alpha) = (l^t_\alpha(y \mid \boldsymbol{d}), u^t_\alpha(y \mid \boldsymbol{d}))$ is the ambiguity set at round t of BED.

The proof builds upon four Lemmas including Lemma 1 on the bias of the weighted empirical CDF, Lemma 2 on concentration of the weighted empirical CDF, Lemma 3 on concentration of the kernel normalisation and upper bound of kernel weights, and Lemma 4 on reduction of the problem to a piecewise-constant optimisation problem.

In plain language, the proof proceeds in several steps. First, we derive an upper bound on the deviation between the expected weighted empirical distribution, $\mathbb{E}[\widehat{F}_w(\cdot)]$, and the CDF of the true data-generating process, $\tilde{Q}(\cdot)$. Next, using McDiarmid's concentration inequality (McDiarmid et al., 1989), we obtain an upper bound on the probability that the weighted empirical distribution deviates from its expectation. We then introduce two probabilistic events, E_1 and E_2, whose probabilities can be explicitly controlled. Conditioning on these events allows us to derive an upper bound on the

kernel weights. Finally, by collecting all intermediate bounds, applying the triangle inequality, and conditioning on E_1 and E_2, we obtain the left-hand side of the coverage inequality. The corresponding probability bound on the right-hand side then follows directly. The key idea of the proof lies in the construction of the events E_1 and E_2. These events are chosen so that they simultaneously enable the chain of inequalities needed to establish the left-hand side of Theorem 3.1 and admit tractable probability bounds. This makes it possible to derive the right-hand side of the desired coverage inequality and thus complete the proof.

We first introduce an additional regularity condition controlling the kernel normalisation term.

Assumption 4 (Regularity of the design density)**.** *The design points $\boldsymbol{d}_1, \ldots, \boldsymbol{d}_N$ are i.i.d. from a continuous distribution with density f_D on $\mathscr{D}$. There exist constants $0 < f_{\min} \leq f_{\max} < \infty$ and $h_0 > 0$ such that*

$$f_{\min} \leq f_D(u) \leq f_{\max} \tag{26}$$

for all u satisfying $\|u - \boldsymbol{d}\| \leq h_0$.

Lemma 1 (Bias of the weighted empirical CDF)**.** *Supposing Assumption 2 (Hölder-γ_D continuity), then for each $y \in [a, b]$,*

$$\left|\mathbb{E}[\widehat{F}_w(y \mid \boldsymbol{d}) \mid \boldsymbol{d}_1, \ldots, \boldsymbol{d}_N] - \tilde{Q}(y \mid \boldsymbol{d})\right| \leq L(b-a)h^{\gamma_D} + O(h^{\gamma_D}) \tag{27}$$

where $\tilde{Q}(y \mid \boldsymbol{d})$ is the CDF of true DGP $\tilde{q}(y \mid \boldsymbol{d})$, constants $L > 0$ and $\gamma_D \in (0, 1]$ come from continuity assumption, and the remainder term $O(h^{\gamma_D})$ is controlled by the exponential decay of the kernel function which is faster than polynomial order in h.

Proof. To gain intuition for Inequality 27, observe that the left-hand side quantifies the discrepancy between the weighted empirical distribution and the true CDF, together with the expected kernel mass assigned to a neighbourhood of radius h around the design point $\boldsymbol{d}$. Since the kernel concentrates its mass locally, most of the weight is assigned to points close to $\boldsymbol{d}$, while the contribution of distant points is negligible. Accordingly, the first term on the right-hand side corresponds to the local neighbourhood of $\boldsymbol{d}$, whereas the second term captures the residual contribution from points outside that neighbourhood. The derivation proceeds as follows:

Consider the definition of $\widehat{F}_w(y \mid \boldsymbol{d})$ in Equation 3 and take the expectation conditional on the design points used to compute the kernel weights:

$$\widehat{F}_w(y \mid \boldsymbol{d}) := \sum_{i=1}^{N} w_i(\boldsymbol{d})\, \mathbb{1}_{\{y_i \leq y\}} \tag{28}$$

$$\begin{aligned}
\mathbb{E}[\widehat{F}_w(y \mid \boldsymbol{d}) \mid \{\boldsymbol{d}_i\}] &= \mathbb{E}[\sum_{i=1}^{N} w_i(\boldsymbol{d})\, \mathbb{1}_{\{y_i \leq y\}}] \\
&= \sum_{i=1}^{N} \mathbb{E}[w_i(\boldsymbol{d})\, \mathbb{1}_{\{y_i \leq y\}}] \\
&= \sum_{i=1}^{N} w_i(\boldsymbol{d})\, \mathbb{E}[\mathbb{1}_{\{y_i \leq y\}}] \\
&= \sum_{i=1}^{N} w_i(\boldsymbol{d}) \tilde{Q}(y \mid \boldsymbol{d}_i)
\end{aligned} \tag{29}$$

Note that in the penultimate line:

$$\mathbb{E}\big[\mathbb{1}_{\{y_i \leq y\}}\big] = \int \mathbb{1}_{\{\theta \leq y\}} \tilde{q}(\theta \mid \boldsymbol{d}_i)\, d\theta = \widetilde{Q}(y \mid \boldsymbol{d}_i) \tag{30}$$

Since by the definition of the weights $\Sigma_i w_i(\boldsymbol{d}) = 1$, then:

$$\widetilde{Q}(y \mid \boldsymbol{d}_i) = 1 \times \widetilde{Q}(y \mid \boldsymbol{d}_i) = \left(\sum_i w_i(\boldsymbol{d})\right) \widetilde{Q}(y \mid \boldsymbol{d}_i) = \sum_i w_i(\boldsymbol{d}) \widetilde{Q}(y \mid \boldsymbol{d}_i) \tag{31}$$

Now to construct the left-hand side of the equation subtract $\widetilde{Q}(y \mid \boldsymbol{d})$ from both sides of inequality 27:

$$\begin{aligned}\mathbb{E}[\widehat{F}_w(y \mid \boldsymbol{d}) \mid \{\boldsymbol{d}_i\}] - \widetilde{Q}(y \mid \boldsymbol{d}) &= \sum_i w_i(\boldsymbol{d})\widetilde{Q}(y \mid \boldsymbol{d}_i) - \sum_i w_i(\boldsymbol{d})\widetilde{Q}(y \mid \boldsymbol{d}) \\ &= \sum_i w_i(\boldsymbol{d}) \left[\widetilde{Q}(y \mid \boldsymbol{d}_i) - \widetilde{Q}(y \mid \boldsymbol{d})\right]\end{aligned} \tag{32}$$

Taking the absolute value of both sides and the fact that $w_i(\boldsymbol{d}) > 0$, we obtain:

$$\left|\mathbb{E}[\widehat{F}_w(y \mid \boldsymbol{d}) \mid \{\boldsymbol{d}_i\}] - \tilde{Q}(y \mid \boldsymbol{d})\right| = \sum_{i=1}^{N} w_i(\boldsymbol{d}) \left|\tilde{Q}(y \mid \boldsymbol{d}_i) - \tilde{Q}(y \mid \boldsymbol{d})\right| \tag{33}$$

Now by using the continuity condition of Assumption 2, we construct an upper bound for the right-hand side of equation 33:

$$\begin{aligned}\left|\tilde{Q}(y \mid \boldsymbol{d}_i) - \tilde{Q}(y \mid \boldsymbol{d})\right| &= \left|\int_a^y \tilde{q}(\theta \mid \boldsymbol{d}_i) - \tilde{q}(\theta \mid \boldsymbol{d}) d\theta\right| \\ &\leq \int_a^y \left|\tilde{q}(\theta \mid \boldsymbol{d}_i) - \tilde{q}(\theta \mid \boldsymbol{d})\right| d\theta \\ &\leq \int_a^y L\|\boldsymbol{d} - \boldsymbol{d}_i\|^{\gamma_D} d\theta \\ &= L\|\boldsymbol{d} - \boldsymbol{d}_i\|^{\gamma_D} \int_a^y d\theta \\ &= L\|\boldsymbol{d} - \boldsymbol{d}_i\|^{\gamma_D} (y - a) \\ &\leq L(b-a)\|\boldsymbol{d} - \boldsymbol{d}_i\|^{\gamma_D}\end{aligned} \tag{34}$$

Now, plugging it into equation 33:

$$\begin{aligned}\left|\mathbb{E}[\widehat{F}_w(y \mid \boldsymbol{d}) \mid \{\boldsymbol{d}_i\}_{i=1}^N] - \widetilde{Q}(y \mid \boldsymbol{d})\right| &\leq \sum_{i=1}^{N} w_i(\boldsymbol{d}) L(b-a) \|\boldsymbol{d} - \boldsymbol{d}_i\|^{\gamma_D} \\ &= L(b-a) \sum_{i=1}^{N} w_i(\boldsymbol{d}) \|\boldsymbol{d} - \boldsymbol{d}_i\|^{\gamma_D} \\ &\leq L(b-a) \left\{ \sum_{i:\|\boldsymbol{d}-\boldsymbol{d}_i\| \leq h} w_i(\boldsymbol{d}) \|\boldsymbol{d} - \boldsymbol{d}_i\|^{\gamma_D} + \sum_{i:\|\boldsymbol{d}-\boldsymbol{d}_i\| > h} w_i(\boldsymbol{d}) \|\boldsymbol{d} - \boldsymbol{d}_i\|^{\gamma_D} \right\} \\ &\leq L(b-a) \left\{ \sum_{i:\|\boldsymbol{d}-\boldsymbol{d}_i\| \leq h} w_i(\boldsymbol{d}) h^{\gamma_D} + \sum_{i:\|\boldsymbol{d}-\boldsymbol{d}_i\| > h} w_i(\boldsymbol{d}) \|\boldsymbol{d} - \boldsymbol{d}_i\|^{\gamma_D} \right\} \\ &\leq L(b-a) \left\{ h^{\gamma_D} \sum_{i:\|\boldsymbol{d}-\boldsymbol{d}_i\| \leq h} w_i(\boldsymbol{d}) + \sum_{i:\|\boldsymbol{d}-\boldsymbol{d}_i\| > h} w_i(\boldsymbol{d}) \|\boldsymbol{d} - \boldsymbol{d}_i\|^{\gamma_D} \right\} \\ &\leq L(b-a) \left\{ h^{\gamma_D} + \sum_{i:\|\boldsymbol{d}-\boldsymbol{d}_i\| > h} w_i(\boldsymbol{d}) \|\boldsymbol{d} - \boldsymbol{d}_i\|^{\gamma_D} \right\} \\ &\leq L(b-a)h^{\gamma_D} + L(b-a) \sum_{i:\|\boldsymbol{d}-\boldsymbol{d}_i\| > h} w_i(\boldsymbol{d}) \|\boldsymbol{d} - \boldsymbol{d}_i\|^{\gamma_D} \\ &\leq L(b-a)h^{\gamma_D} + O(h^{\gamma_D})\end{aligned} \tag{35}$$

The last line follows from the definition of the Kernel, where the weights exponentially decay for points far from the target design $\boldsymbol{d}$. This leads us to the right-hand side of Lemma 1. □

Lemma 2 (Concentration of the weighted empirical CDF)**.** *Conditional on the design points* $\{\boldsymbol{d}_i\}_{i=1}^N$*, for every possible outcome* $y \in [a, b]$ *and any given* $\varepsilon > 0$*,*

$$\mathbb{P}\left\{\left|\widehat{F}_w(y \mid \boldsymbol{d}) - \mathbb{E}[\widehat{F}_w(y \mid \boldsymbol{d}) \mid \{\boldsymbol{d}_i\}]\right| \geq \varepsilon \;\middle|\; \{\boldsymbol{d}_i\}\right\} \leq 2\exp\left(-\frac{2\varepsilon^2}{\sum_{i=1}^N w_i(\boldsymbol{d})^2}\right) \tag{36}$$

Proof. The inequality 36, is a direct application from McDiarmid's concentration inequality (McDiarmid et al., 1989). In other words, McDiarmid's inequality states that if a function is bounded versus any changes to a single input at a time (holding other inputs unchanged), then its deviation from its expected value is also probabilistically bounded.

Here, according to the definition of $\widehat{F}_w(y \mid \boldsymbol{d})$, for a given observations $\{(\boldsymbol{d}_i, y_i)\}_{i=1}^N$:

$$\widehat{F}_w(y \mid \boldsymbol{d}) := \sum_{i=1}^N w_i(\boldsymbol{d})\, \mathbb{1}_{\{y_i \leq y\}} \tag{37}$$

Therefore, the maximum possible change in the function due to a change in only a single input would be as follows:

$$\begin{aligned}
&\left|\widehat{F}_w\left(y \mid \boldsymbol{d}; \{(\boldsymbol{d}_1, y_1), \ldots, (\boldsymbol{d}_i, y_i), \ldots, (\boldsymbol{d}_N, y_N)\}\right) - \widehat{F}_w\left(y \mid \boldsymbol{d}; \{(\boldsymbol{d}_1, y_1), \ldots, (\boldsymbol{d}_i, y_i'), \ldots, (\boldsymbol{d}_N, y_N)\}\right)\right| \\
&\quad = \left|\sum_{l\neq i, l=1}^N w_l(\boldsymbol{d})\, \mathbb{1}_{\{y_l \leq y\}} + w_i(\boldsymbol{d})\, \mathbb{1}_{\{y_i \leq y\}} - \sum_{l\neq i, l=1}^N w_l(\boldsymbol{d})\, \mathbb{1}_{\{y_l \leq y\}} - w_i(\boldsymbol{d})\, \mathbb{1}_{\{y_i' \leq y\}}\right| \\
&= |w_i(\boldsymbol{d})(\mathbb{1}_{\{y_i \leq y\}} - \mathbb{1}_{\{y_i' \leq y\}})| = w_i(\boldsymbol{d})|\mathbb{1}_{\{y_i \leq y\}} - \mathbb{1}_{\{y_i' \leq y\}}| \leq w_i(\boldsymbol{d}) \times 1 = w_i(\boldsymbol{d})
\end{aligned} \tag{38}$$

Therefore, changing an observation y_i changes $\widehat{F}_w(y \mid \boldsymbol{d})$ by at most $w_i(\boldsymbol{d})$. Applying McDiarmid's bounded-difference inequality completes the proof. □

Lemma 3 (Concentration of the kernel normalisation)**.** *Define the summation of the normalised kernels as follows:*

$$S(\boldsymbol{d}) := \sum_{j=1}^N \mathcal{K}_h(\boldsymbol{d}_j - \boldsymbol{d}) \tag{39}$$

where $\mathcal{K}_h(\boldsymbol{d}_j - \boldsymbol{d}) = h^{-k}\mathcal{K}((\boldsymbol{d}_j - \boldsymbol{d})/h)$*. Under Assumptions 3 and 4, there exists a constant* $c_1 > 0$*, depending on* $\kappa(0)$*, such that*

$$\mathbb{P}\left\{S(\boldsymbol{d}) \geq \frac{1}{2}Nh^k f_{\min}\right\} \geq 1 - \exp(-c_1 Nh^k f_{\min}) \tag{40}$$

Proof. This is a non trivial inequality that establishes a probabilistic lower bound on the summation of localised kernel. It follows from Bernstein's inequality (Jiang et al., 2018), for given set of independent random variables $X_1, .., X_n$ where $\mathbb{E}(X_i) = 0$, and $|X_i| \leq b$, we have:

$$\mathbb{P}\{\sum_{i=1}^n X_i \geq t\} \leq \exp(-\frac{t^2}{2(\sigma^2 + bt/3)}) \tag{41}$$

For which we first show that the requirements are met. By the definition of the Kernel function, each kernel contribution, its expectation and its variance are bounded as follows:

$$\mathcal{K}_h(\boldsymbol{d}_j - \boldsymbol{d}) = h^{-k}\mathcal{K}((\boldsymbol{d}_j - \boldsymbol{d})/h) \in (0, h^{-k}\kappa(0)] \tag{42}$$

$$\mathbb{E}[\mathcal{K}_h(\boldsymbol{d}_j - \boldsymbol{d})] = \int \mathcal{K}_h(u - \boldsymbol{d}) f_D(u)\, du \geq f_{\min} \int \mathcal{K}(v)\, dv = f_{\min} \tag{43}$$

$$\mathbb{E}[\mathcal{K}_h(\boldsymbol{d}_j - \boldsymbol{d})] = \int \mathcal{K}_h(u - \boldsymbol{d}) f_D(u)\, du \leq f_{\max} \int \mathcal{K}(v)\, dv = f_{\max} \tag{44}$$

Hence,

$$N f_{\min} \leq \mathbb{E}[S(\boldsymbol{d})] \leq N f_{\max} \tag{45}$$

Moreover, by the definition of variance, it turns out that:

$$\mathrm{Var}(\mathcal{K}_h(\boldsymbol{d}_j - \boldsymbol{d})) \leq \frac{\kappa(0)}{h^k}\mathbb{E}[\mathcal{K}_h(\boldsymbol{d}_j - \boldsymbol{d})] \tag{46}$$

Now, we apply Bernstein's inequality by defining X_i in inequality 41 as follows:

$$X_i = \mathcal{K}_h(\boldsymbol{d}_j - \boldsymbol{d}) - \mathbb{E}[\mathcal{K}_h(\boldsymbol{d}_j - \boldsymbol{d})] \tag{47}$$

Then:

$$\mathbb{E}[X_i] = 0 \tag{48}$$

$$Var(X_i) \leq \frac{\kappa(0)}{h^k}\mathbb{E}[\mathcal{K}_h(\boldsymbol{d}_j - \boldsymbol{d})] \tag{49}$$

The following derivation shows $b = \frac{\kappa(0)}{h^k}$:

$$\max |\mathcal{K}_h(\boldsymbol{d}_j - \boldsymbol{d}) - \mathbb{E}[\mathcal{K}_h(\boldsymbol{d}_j - \boldsymbol{d})]| \leq \max \mathcal{K}_h(\boldsymbol{d}_j - \boldsymbol{d}) = \frac{\kappa(0)}{h^k} \tag{50}$$

Constructing the left-hand side of the Bernstein's inequality leads to:

$$\sum_{i=1}^{N} X_i = \sum_{i=1}^{N} \left(\mathcal{K}_h(\boldsymbol{d}_i - \boldsymbol{d}) - \underbrace{\mathbb{E}[\mathcal{K}_h(\boldsymbol{d}_i - \boldsymbol{d})]}_{\mu} \right) = S(\boldsymbol{d}) - N\mu \tag{51}$$

Now, let $t = N(\mu - \frac{1}{2}h^k f_{min})$ and plug it into Bernstein's inequality:

$$\begin{aligned}
-\frac{t^2}{2\left(\sigma^2 + \frac{bt}{3}\right)} &= -\frac{t}{2\left(\frac{\sigma^2}{t} + \frac{b}{3}\right)} \\
&\leq -\frac{\underbrace{\frac{1}{2} N f_{\min}}_{\text{lower bound on } t}}{2\left(\underbrace{\frac{\kappa(0) f_{\max}}{h^k N \frac{1}{2} f_{\min}}}_{\text{upper bound on } \sigma^2/t} + \underbrace{\frac{\kappa(0)}{3h^k}}_{\text{upper bound on } b/3}\right)} \\
&= -\frac{N f_{\min}}{\frac{\kappa(0)}{h^k}\left(\frac{8 f_{\max}}{N f_{\min}} + \frac{4}{3}\right)} = -\frac{N h^k f_{\min}}{\kappa(0)\left(\frac{8 f_{\max}}{N f_{\min}} + \frac{4}{3}\right)} \\
&= -c_1 N h^k f_{\min}
\end{aligned} \tag{52}$$

Thus,

$$\mathbb{P}\left\{S(\boldsymbol{d}) \leq \frac{1}{2}Nh^k f_{\min}\right\} \leq \exp(-c_1 Nh^k f_{\min}) \tag{53}$$

$$\mathbb{P}\left\{S(\boldsymbol{d}) \geq \frac{1}{2}Nh^k f_{\min}\right\} \geq 1 - \exp(-c_1 Nh^k f_{\min}) \tag{54}$$

where $c_1 > 0$ depends on the kernel's upper bound. This completes the proof. □

Corollary 1 (Kernel weights upper bound)**.** *On the event of Lemma 2, with probability* $1 - \alpha^{KWC}$*, where* $\alpha^{KWC} = \exp(-c_1 Nh^k f_{\min})$*,*

$$\max_i w_i(\boldsymbol{d}) \leq \frac{2}{Nh^{2k} f_{\min}} \tag{55}$$

and

$$\sum_i w_i(\boldsymbol{d})^2 \leq \frac{2}{Nh^{2k} f_{\min}} \tag{56}$$

Proof.

$$w_i(\boldsymbol{d}) = \frac{\mathcal{K}\left(\frac{\boldsymbol{d}_i - \boldsymbol{d}}{h}\right)}{\sum_{j=1}^N \mathcal{K}\left(\frac{\boldsymbol{d}_j - \boldsymbol{d}}{h}\right)} \leq \frac{\kappa(0)}{\sum_{j=1}^N \mathcal{K}(\boldsymbol{d}_j - \boldsymbol{d})} = \frac{\kappa(0)}{h^k S(d)} \leq \frac{\kappa(0)}{h^k \frac{1}{2} Nh^k f_{min}} \leq \frac{2}{Nh^{2k} f_{min}} \tag{57}$$

Additionally,

$$\sum_i w_i(\boldsymbol{d})^2 \leq \sum_i \max_i w_i(\boldsymbol{d}) \times w_i(\boldsymbol{d}) \leq \max_i w_i(\boldsymbol{d}) \sum_i w_i(\boldsymbol{d}) \leq \max_i w_i(\boldsymbol{d}) \times 1 \leq \frac{2}{Nh^{2k} f_{\min}} \tag{58}$$

Note that according to the Kernel regularity conditions in assumption 1, $\kappa(0) \leq 1$. □

Lemma 4 (Reduction to piecewise-constant densities)**.** *Fix a design point* $\boldsymbol{d} \in \mathscr{D}$*, and consider the ordered breakpoints*

$$a = z_0 < z_1 < \cdots < z_{\widehat{N}-1} < z_{\widehat{N}} = b \tag{59}$$

and let $y = z_{\widehat{s}}$*. Consider* $\mathscr{E}$ *the class of densities* $q(y|\boldsymbol{d})$ *on* $[a, b]$ *that are unimodal with mode* $\mu(\boldsymbol{d})$*, bounded by* $\mathcal{U}(d)$ *and satisfying the group mass constraints as follows:*

$$c^-(\alpha_1) \leq \int_{y_{(k_{i-1}(\boldsymbol{d}))}}^{y_{(k_i(\boldsymbol{d}))}} q(t \mid \boldsymbol{d})\, dt \leq c^+(\alpha_1) \tag{60}$$

for all groups, where α_1 *is used for the calibration of weighted empirical distribution and* $k_m(\boldsymbol{d}) = \max\left\{j : \widehat{F}_w\big(y_{(j)} \mid \boldsymbol{d}\big) \leq \frac{m}{M}\right\}$*. The maximum and minimum possible value of* $q(y \mid \boldsymbol{d})$ *over this class can be attained by piecewise constant densities over all partitions* $\Omega = \{[z_{i-1}, z_i); j = 1, \cdots, M\}$*.*

Proof. First, we construct a piecewise-constant density function $\mathscr{F}(\cdot)$. We show that it belongs to $\mathscr{E}$ by verifying that it satisfies the density bounds of $\mathcal{U}(\boldsymbol{d})$, has mode $\mu(\boldsymbol{d})$, and satisfies the group-mass constraint. Second, we prove by construction that $\mathscr{F}(\cdot)$ can attain the minimum and maximum possible values of $q(y \mid \boldsymbol{d})$.

For each partition $\mathcal{I}_j = [z_{j-1}, z_j)$, define the length and total mass of the partition as follows:

$$\Delta_j = z_j - z_{j-1}, \qquad S_j = \int_{\mathcal{I}_j} q(y \mid \boldsymbol{d}) dy \tag{61}$$

Therefore, the average height of the density $q(y \mid \boldsymbol{d})$ on partition $\mathcal{I}_j$ and the related piecewise constant functions are as follows:

$$\bar{q}_j = \frac{S_j}{\Delta_j}, \qquad \mathscr{F}(y \mid \boldsymbol{d}) = \bar{q}_j \quad \forall y \in \mathcal{I}_j \tag{62}$$

By construction:

$$\int_{\mathcal{I}_j} \mathscr{F}(y \mid \boldsymbol{d})dy = \int_{\mathcal{I}_j} q(y \mid \boldsymbol{d})dy \tag{63}$$

$$c^-(\alpha_1) \leq \int_{y_{(k_{j-1}(\boldsymbol{d}))}}^{y_{(k_j(\boldsymbol{d}))}} q(t \mid \boldsymbol{d})\, dt = \int_{\mathcal{I}_j} \mathscr{F}(y \mid \boldsymbol{d})dy \leq c^+(\alpha_1) \tag{64}$$

Therefore, replacing $q(y \mid \boldsymbol{d})$ by $\mathscr{F}(y \mid \boldsymbol{d})$ preserves group total mass and group total math constraints. Additionally, it results in $\int_a^b \mathscr{F}(y \mid \boldsymbol{d})dy = 1$.

Moreover, by definition of $\mathscr{F}(y \mid \boldsymbol{d})$ it satisfies the boundedness condition $\mathscr{F}(y \mid \boldsymbol{d}) \leq \mathcal{U}(\boldsymbol{d})$. To show the unimodality of $\mathscr{F}(y \mid \boldsymbol{d})$, consider its properties relative to $q(y \mid \boldsymbol{d})$ in two folds of nondecreasing and nonincreasing regions.

For the nondecreasing part consider $y_h < y_l < \mathcal{U}(\boldsymbol{d})$. By unimodality of $q(y \mid \boldsymbol{d})$, we have $\forall y_h \in \mathcal{I}_j = [z_{j-1}, z_j)$ and $\forall y_l \in \mathcal{I}_{j+1} = [z_j, z_{j+1})$:

$$q(y_h \mid \boldsymbol{d}) \leq q(y_l \mid \boldsymbol{d}) \quad \text{and} \quad S_j = \int_{\mathcal{I}_j} q(y \mid \boldsymbol{d})dy \approx \int_{\mathcal{I}_{j+1}} q(y \mid \boldsymbol{d})dy = S_{j+1} \tag{65}$$

The partitions are chosen, in a way through the definition of $k_m(\boldsymbol{d})$, the selection of group size M and the kernel bandwidth h that the total mass in each partition is approximately equal to $1/M$. As in the nondecreasing part of any unimodal functions, the step lengths are decreasing:

$$\Delta_j \geq \Delta_{j+1} \quad \text{and} \quad S_j \approx S_{j+1} \quad \Rightarrow \quad \bar{q}_j \leq \bar{q}_{j+1} \tag{66}$$

Thus, $\mathscr{F}(y \mid \boldsymbol{d})$ satisfies the unimodality condition by meeting the constraint $\bar{q}_j \leq \bar{q}_{j+1}$.

For the nonincreasing part, analogous to the above part it can be shown that another set of constraint should be met, $\bar{q}_j \geq \bar{q}_{j+1}$.

Thus, by construction $\mathscr{F}$ belongs to $\mathscr{E}$, and the optimisation of $q(y \mid \boldsymbol{d})$ reduces to the the optimisation over the finite collection of partition heights $(\bar{q}_1, \cdots, \bar{q}_{\hat{N}})$ as follows:

$$\begin{gathered} 0 \leq \bar{q}_j \leq \mathcal{U}(\boldsymbol{d}) \\ \bar{q}_1 \leq \bar{q}_2 \leq \cdots \leq \bar{q}_{j_\mu} \quad \text{and} \quad \bar{q}_{j_\mu} \geq \bar{q}_{j_\mu+1} \geq \cdots \geq \bar{q}_{\hat{N}} \\ \sum_{j=1}^{\hat{N}} \Delta_j \bar{q}_j = 1 \\ c^-(\alpha_1) \leq \int_{\mathcal{I}_j} \mathscr{F}(y \mid \boldsymbol{d})\, dy \leq c^+(\alpha_1) \end{gathered} \tag{67}$$

Therefore, for the optimisation of $q(y \mid \boldsymbol{d})$ the maximum and minimum possible values for the density at the target design point $\boldsymbol{d}$ can be obtained by the above linear program over the non empty polyhedral feasible set at the exteremd points. Therefore, there exist an extremal piecewise constant density in $\mathscr{E}$ that gives the upper bound and lower bound for density $q(y \mid \boldsymbol{d})$. This completes the proof. □

Proof of Theorem: Define the following event,

$$E_1 := \left\{ S(\boldsymbol{d}) \geq \frac{1}{2} N h^k f_{\min} \right\} \tag{68}$$

where $\mathbb{P}\{E_1\} = 1 - \alpha^{KWC}$. According to Lemma 2:

$$\mathbb{P}\left\{\left|\widehat{F}_w(y \mid \boldsymbol{d}) - \mathbb{E}[\widehat{F}_w(y \mid \boldsymbol{d}) \mid \{\boldsymbol{d}_i\}]\right| \geq \varepsilon \;\middle|\; \{\boldsymbol{d}_i\}\right\} \leq 2\exp\left(-\frac{2\varepsilon^2}{\sum_{i=1}^{N} w_i(\boldsymbol{d})^2}\right) \tag{69}$$

On the occurrence of event E_1, by Corollary 1:

$$\mathbb{P}\left\{\left|\widehat{F}_w(y \mid \boldsymbol{d}) - \mathbb{E}[\widehat{F}_w(y \mid \boldsymbol{d}) \mid \{\boldsymbol{d}_i\}]\right| \geq \varepsilon \;\middle|\; \{\boldsymbol{d}_i\}, E_1\right\} \leq 2\exp\left(-Nh^{2k} f_{min}\varepsilon^2\right) \tag{70}$$

Choose $\varepsilon = \sqrt{\frac{\log(\frac{2(M+1)}{\alpha_1})}{Nh^{2k} f_{min}}}$, where α_1 is used for the weighted bootstrap calibration, therefore:

$$\mathbb{P}\left\{\left|\widehat{F}_w(y \mid \boldsymbol{d}) - \mathbb{E}[\widehat{F}_w(y \mid \boldsymbol{d}) \mid \{\boldsymbol{d}_i\}]\right| \geq \varepsilon \;\middle|\; \{\boldsymbol{d}_i\}, E_1\right\} \leq \frac{\alpha_1}{M+1} \tag{71}$$

Define the event A_i as follows:

$$A_i = \left\{|\widehat{F}_w(y_i \mid \boldsymbol{d}) - \mathbb{E}[\widehat{F}_w(y_i \mid \boldsymbol{d}) \mid \{\boldsymbol{d}_i\}]| > \varepsilon\right\} \qquad i = 1, \cdots M \tag{72}$$

Therefore, take the union of the above events, and use the inequality for the probability of a union of events:

$$\mathbb{P}\left\{\bigcup_{i=1}^{M} A_i \mid \{\boldsymbol{d}_i\}, E_1\right\} \leq \sum_{i=1}^{M} \mathbb{P}\{A_i \mid \{\boldsymbol{d}_i\}, E_1\} \leq M\,\frac{\alpha_1}{M+1} \leq \alpha_1 \tag{73}$$

Therefore,

$$\mathbb{P}\left\{\bigcap_{i=1}^{M} A_i^c \mid \{\boldsymbol{d}_i\}, E_1\right\} \geq 1 - \alpha_1 \tag{74}$$

The following union bound converts M individual probability bounds into a single useful probability bound.

$$\mathbb{P}\left\{\max_{0\leq i\leq M}\left|\widehat{F}_w(y \mid \boldsymbol{d}) - \mathbb{E}[\widehat{F}_w(y \mid \boldsymbol{d}) \mid \{\boldsymbol{d}_i\}]\right| \leq \varepsilon \;\middle|\; \{\boldsymbol{d}_i\}, E_1\right\} \geq 1 - \alpha_1 \tag{75}$$

From Lemma 1:

$$\left|\mathbb{E}[\widehat{F}_w(y \mid \boldsymbol{d}) \mid \boldsymbol{d}_1, \ldots, \boldsymbol{d}_N] - \tilde{Q}(y \mid \boldsymbol{d})\right| \leq L(b-a)h^{\gamma_D} + O(h^{\gamma_D}) \tag{76}$$

On the occurrence of the stochastic event in equation 75 and by using the triangle inequality, we have:

$$\mathbb{P}\left\{\left|\widehat{F}_w(y \mid \boldsymbol{d}) - \tilde{Q}(y \mid \boldsymbol{d})\right| \leq \underbrace{\varepsilon + L(b-a)h^{\gamma_D} + O(h^{\gamma_D})}_{\eta} \;\middle|\; \{\boldsymbol{d}_i\}, E_1\right\} \geq 1 - \alpha_1 \tag{77}$$

The above inequality, relates the weighted empirical distribution to the true CDF. By using the above probabilistic inequality we construct functional upper and lower bounds for the true DGP. Now, for the true CDF increments, the following inequality holds for any candidate experimental outcomes t_i and t_{i-1} :

$$\Delta_i^w(\boldsymbol{d}) := \tilde{Q}(t_i \mid \boldsymbol{d}) - \tilde{Q}(t_{i-1} \mid \boldsymbol{d}) \tag{78}$$

$$\widehat{\Delta}_i^w(\boldsymbol{d}) := \widehat{F}_w(t_i \mid \boldsymbol{d}) - \widehat{F}_w(t_{i-1} \mid \boldsymbol{d}) \tag{79}$$

$$\begin{aligned}\left|\Delta_i^w(\boldsymbol{d}) - \widehat{\Delta}_i^w(\boldsymbol{d})\right| &= \left|\tilde{Q}(t_i \mid \boldsymbol{d}) - \tilde{Q}(t_{i-1} \mid \boldsymbol{d}) - \widehat{F}_w(t_i \mid \boldsymbol{d}) + \widehat{F}_w(t_{i-1} \mid \boldsymbol{d})\right| \\ &\leq \left|\tilde{Q}(t_i \mid \boldsymbol{d}) - \widehat{F}_w(t_i \mid \boldsymbol{d})\right| + \left|\widehat{F}_w(t_{i-1} \mid \boldsymbol{d}) - \tilde{Q}(t_{i-1} \mid \boldsymbol{d})\right| \\ &\leq \underbrace{2\eta}_{\delta}\end{aligned} \tag{80}$$

Here, in the above equation, on the occurrence of event E_1 with probability $1 - \alpha^{KWC}$, and on the occurrence of event $E_2 = A_i|E_1$ with probability $1 - \alpha_1$, we can use the upper bound η. Otherwise, the probability of the failure of E_1 is α^{KWC}, and for E_2 is α_1. The probability of both of them occur successfully, would be as follows:

$$\mathbb{P}(E_1 \cap E_2) \geq 1 - \alpha^{KWC} - \alpha_1 \tag{81}$$

Thus, with the occurrence of those events:

$$\left|\Delta_i^w(\boldsymbol{d}) - \widehat{\Delta}_i^w\right| \leq \delta \tag{82}$$

where:

$$\begin{aligned}\delta &= 2\sqrt{\frac{\log\left(\frac{2(M+1)}{\alpha_1}\right)}{Nh^{2k}f_{\min}}} + 2L(b-a)h^{\gamma_D} + 2O\left(h^{\gamma_D}\right) \\ &\approx 2\sqrt{\frac{\log\left(\frac{2(M+1)}{\alpha_1}\right)}{Nh^{2k}f_{\min}}} + 2L(b-a)h^{\gamma_D}\end{aligned} \tag{83}$$

Therefore,

$$\widehat{\Delta}_i^w - \delta \leq \Delta_i^w(\boldsymbol{d}) \leq \widehat{\Delta}_i^w + \delta \tag{84}$$

$$\widehat{\Delta}_i^w - \delta \leq \tilde{Q}(t_i \mid \boldsymbol{d}) - \tilde{Q}(t_{i-1} \mid \boldsymbol{d}) \leq \widehat{\Delta}_i^w + \delta \tag{85}$$

$$\widehat{\Delta}_i^w - \delta \leq \int_{t_{i-1}}^{t_i} \tilde{q}(y \mid \boldsymbol{d}) \leq \widehat{\Delta}_i^w + \delta \tag{86}$$

Thus, by choosing $c^-(\alpha_1) = \widehat{\Delta}_i^w - \delta$ and $c^+(\alpha_1) = \widehat{\Delta}_i^w + \delta$ ensures that $\tilde{q}(\cdot \mid \boldsymbol{d})$ satisfies the group-mass constraints that define $\mathscr{H}_w(\boldsymbol{d}, \alpha)$. The remaining shape constraints and normalization hold by assumption. Hence $\tilde{q}(\cdot \mid \boldsymbol{d})$ belongs to the ambiguity set.

By Lemma 4, the optimal solutions of the associated linear programming model specify a pair of functional confidence bands $(l_\alpha(y \mid \boldsymbol{d}), u_\alpha(y \mid \boldsymbol{d}))$ which are respectively the minimum and maximum values of $q(y \mid \boldsymbol{d})$ over the feasible class $\mathscr{E}$. As the true density is one of the feasible distributions,

$$l_\alpha(y \mid \boldsymbol{d}) \leq \tilde{q}(y \mid \boldsymbol{d}) \leq u_\alpha(y \mid \boldsymbol{d}) \tag{87}$$

Combining the previously defined probabilistic events:

$$\mathbb{P}\left\{l_\alpha(y \mid \boldsymbol{d}) \leq \tilde{q}(y \mid \boldsymbol{d}) \leq u_\alpha(y \mid \boldsymbol{d}),\ \forall y \in [a, b]\right\} \geq 1 - \alpha^{KWC} - \alpha_1 \tag{88}$$

which proves the theorem. Therefore, for sufficient N, where the following holds,

$$1 - \alpha^{KWC} - \alpha_1 \geq 1 - \alpha \Rightarrow \alpha^{KWC} + \alpha_1 \leq \alpha \tag{89}$$

This completes the proof.

A.3.1 Optimisation of confidence bands

To optimise the confidence bands $(l^t_\alpha(y \mid \boldsymbol{d}), u^t_\alpha(y \mid \boldsymbol{d}))$, we first construct the polyhedral set $\mathscr{H}_w(\boldsymbol{d}, \alpha)$ of possible bands in each bin. We then obtain the functional bands of the ambiguity set by taking, within each bin, the minimum and maximum. To this end, fix a target pair $(y, \boldsymbol{d}) \in [a, b] \times \mathscr{D}$, and let $z_1 < \cdots < z_{\widehat{N}}$ denote the sorted distinct values in $\{a,\ b,\ \mu(\boldsymbol{d}),\ y,\ y_{(k_1(\boldsymbol{d}))}, \ldots, y_{(k_M(\boldsymbol{d}))}\}$ consisting of the outcome-space boundaries, the mode, the target outcome, and the bin endpoints. Here, $\widehat{N} \geq M$ is the number of distinct values in this collection. After ordering these values, we denote the mode and target outcome indexes as $\widehat{s}$ and $\widetilde{s}$ respectively, so that $z_{\widehat{s}} = y$, $z_{\widetilde{s}} = \mu(\boldsymbol{d})$.

Following Chen et al. (2022), we show in Appendix A.3 that the lower and upper confidence bands within each bin are step functions. Therefore, the functional bands of the ambiguity set are piecewise constant densities. Next, consider the intervals induced by the ordered values $z_1 < \cdots < z_{\widehat{N}}$. Let $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_{\widehat{N}-1})$ be the optimisation variables representing the values of the step function on the intervals $[z_j, z_{j+1}), \quad j = 1, \ldots, \widehat{N} - 1$. These values are determined by the following linear programming formulation:

$$\beta_1 \leq \beta_2 \leq \cdots \leq \beta_{\widetilde{s}-1}, \qquad \beta_{\widetilde{s}} \geq \beta_{\widetilde{s}+1} \geq \cdots \geq \beta_{\widehat{N}-1}, \tag{90}$$

$$c^-(\alpha_1) \leq \sum_{j:\, y_{k_{m-1}(\boldsymbol{d})+1} \leq z_j \leq y_{k_m(\boldsymbol{d})}} \beta_j (z_{j+1} - z_j) \leq c^+(\alpha_1), \quad m = 2, \ldots, M, \tag{91}$$

$$\sum_{j=1}^{\widehat{N}-1} \beta_j (z_{j+1} - z_j) = 1, \qquad 0 \leq \beta_j \leq \mathcal{U}(\boldsymbol{d}), \quad j = 1, \ldots, \widehat{N} - 1. \tag{92}$$

Note that Equation 90 enforces unimodality, requiring the step heights β_i to be nondecreasing before the mode $z_{\widetilde{s}}$ and nonincreasing afterwards. Equation 91 imposes the calibrated bin-mass bounds in each partition, while Equation 92 ensures that the step function integrates to one and satisfies the upper confidence bound on the true DGP. The lower and upper confidence bands at $(y, \boldsymbol{d})$ are obtained by solving the following optimisation problems:

$$\widetilde{l}^t_\alpha(y \mid \boldsymbol{d}) = \begin{cases} \min\limits_{\beta \in \mathscr{H}_w(\boldsymbol{d}, \alpha)} \beta_{\widehat{s}-1}, & y \leq \mu(\boldsymbol{d}) \\ \min\limits_{\beta \in \mathscr{H}_w(\boldsymbol{d}, \alpha)} \beta_{\widehat{s}}, & y > \mu(\boldsymbol{d}) \end{cases} \qquad \widetilde{u}^t_\alpha(y \mid \boldsymbol{d}) = \begin{cases} \max\limits_{\beta \in \mathscr{H}_w(\boldsymbol{d}, \alpha)} \beta_{\widehat{s}}, & y \leq \mu(\boldsymbol{d}) \\ \max\limits_{\beta \in \mathscr{H}_w(\boldsymbol{d}, \alpha)} \beta_{\widehat{s}-1}, & y > \mu(\boldsymbol{d}) \end{cases} \tag{93}$$

and computing $l^t_\alpha(y \mid \boldsymbol{d}) = \max\{0, \widetilde{l}_{\alpha_1}(y \mid \boldsymbol{d}) - Lh^{\gamma_D}\}$, $u^t_\alpha(y \mid \boldsymbol{d}) = \widetilde{u}_{\alpha_1}(y \mid \boldsymbol{d}) + Lh^{\gamma_D}$.The complete procedure is summarised in Algorithm 2 in Appendix A.5.

A.4 Proof for Regret guarantee

Proposition A.1: Bounded Regret Guarantee

Let $\boldsymbol{d}^\star$ denote the design that maximizes the DR-EIG objective ($U^{\mathrm{DR}}(\boldsymbol{d})$ in Equation 16), i.e. the utility evaluated under the worst-case density $q(\cdot \mid \boldsymbol{d})$ in the ambiguity set $\mathcal{Q}$, and let $\widetilde{\boldsymbol{d}}$ denote the design that maximizes the true utility $U^{\mathrm{true}}(\boldsymbol{d})$, evaluated under the true DGP $\tilde{q}(\cdot \mid \boldsymbol{d})$. Under the conditions of Theorem 3.1, with probability at least $1 - \alpha$,

$$\mathrm{Regret}(\boldsymbol{d}^\star) := U^{\mathrm{true}}(\widetilde{\boldsymbol{d}}) - U^{\mathrm{true}}(\boldsymbol{d}^\star) \ \leq \ \Delta(\alpha) \tag{94}$$

where $\Delta(\alpha)$ is the confidence-band width. That is, the design chosen by maximizing the robust surrogate U^{DR} suffers at most $\Delta(\alpha)$ loss in true utility relative to the design $\widetilde{\boldsymbol{d}}$.

Proof. We compare $\boldsymbol{d}^\star$ against $\widetilde{\boldsymbol{d}}$ in true utility by routing the comparison through the robust surrogate U^{DR}, which is the quantity $\boldsymbol{d}^\star$ actually optimizes. Adding and subtracting $U^{\mathrm{DR}}(\widetilde{\boldsymbol{d}})$ and $U^{\mathrm{DR}}(\boldsymbol{d}^\star)$ gives the decomposition

$$\mathrm{Regret}(\boldsymbol{d}^\star) = \underbrace{\left[U^{\mathrm{DR}}(\widetilde{\boldsymbol{d}}) - U^{\mathrm{DR}}(\boldsymbol{d}^\star)\right]}_{\text{(I)}} + \underbrace{\left[U^{\mathrm{DR}}(\boldsymbol{d}^\star) - U^{\mathrm{true}}(\boldsymbol{d}^\star)\right]}_{\text{(II)}} + \underbrace{\left[U^{\mathrm{true}}(\widetilde{\boldsymbol{d}}) - U^{\mathrm{DR}}(\widetilde{\boldsymbol{d}})\right]}_{\text{(III)}} \tag{95}$$

We bound each term as follows.

Term (I): Optimality of $\boldsymbol{d}^\star$. By construction, $\boldsymbol{d}^\star$ maximizes $U^{\mathrm{DR}}(\boldsymbol{d})$ over $\mathcal{D}$, so in particular $U^{\mathrm{DR}}(\widetilde{\boldsymbol{d}}) \le U^{\mathrm{DR}}(\boldsymbol{d}^\star)$, and therefore

$$U^{\mathrm{DR}}(\widetilde{\boldsymbol{d}}) - U^{\mathrm{DR}}(\boldsymbol{d}^\star) \le 0 \tag{96}$$

Term (II): The robust surrogate lower-bounds the truth. By Theorem 3.1, with probability at least $1-\alpha$ the true density lies in the ambiguity set at every design, $\tilde{q}(\cdot \mid \boldsymbol{d}) \in \mathcal{Q}(\boldsymbol{d}, \alpha)$ for all $\boldsymbol{d} \in \mathcal{D}$. On this event, since $U^{\mathrm{DR}}(\boldsymbol{d})$ is by definition an infimum over $\mathcal{Q}(\boldsymbol{d}, \alpha)$, and $\tilde{q}(\cdot \mid \boldsymbol{d})$ is one member of that set, $\forall \boldsymbol{d} \in \mathcal{D}$:

$$U^{\mathrm{DR}}(\boldsymbol{d}) = \inf_{q(\cdot|\boldsymbol{d}) \in \mathcal{Q}(\boldsymbol{d},\alpha)} \int_a^b f(\boldsymbol{d}, y)\, q(y \mid \boldsymbol{d})\, dy \;\le\; \int_a^b f(\boldsymbol{d}, y)\, \tilde{q}(y \mid \boldsymbol{d})\, dy = U^{\mathrm{true}}(\boldsymbol{d}) \tag{97}$$

Evaluating at $\boldsymbol{d} = \boldsymbol{d}^\star$ gives

$$U^{\mathrm{DR}}(\boldsymbol{d}^\star) - U^{\mathrm{true}}(\boldsymbol{d}^\star) \le 0 \tag{98}$$

Term (III): The confidence band controls the gap. On the same coverage event, the true density is sandwiched by the calibrated confidence band at $\widetilde{\boldsymbol{d}}$,

$$l_\alpha(y \mid \widetilde{\boldsymbol{d}}) \le \tilde{q}(y \mid \widetilde{\boldsymbol{d}}) \le u_\alpha(y \mid \widetilde{\boldsymbol{d}}), \qquad \forall y \in [a, b] \tag{99}$$

Since $U^{\mathrm{DR}}(\widetilde{\boldsymbol{d}})$ takes the infimum over all densities consistent with this band, the largest possible gap between $U^{\mathrm{true}}(\widetilde{\boldsymbol{d}})$ and $U^{\mathrm{DR}}(\widetilde{\boldsymbol{d}})$ occurs when the worst-case density is pushed to the least favorable endpoint of the band at every y, according to the sign of $f(\widetilde{\boldsymbol{d}}, y)$. This yields

$$U^{\mathrm{true}}(\widetilde{\boldsymbol{d}}) - U^{\mathrm{DR}}(\widetilde{\boldsymbol{d}}) \;\le\; \int_a^b \left|f(\widetilde{\boldsymbol{d}}, y)\right| \Big(u_\alpha(y \mid \widetilde{\boldsymbol{d}}) - l_\alpha(y \mid \widetilde{\boldsymbol{d}})\Big)\, dy \;\le\; \Delta(\alpha) \tag{100}$$

We define the confidence-band width term appearing in Term (III) as

$$\Delta(\alpha) \;:=\; \sup_{\boldsymbol{d} \in \mathcal{D}} \int_a^b \left|f(\boldsymbol{d}, y)\right| \Big(u_\alpha(y \mid \boldsymbol{d}) - l_\alpha(y \mid \boldsymbol{d})\Big)\, dy \tag{101}$$

It is called the confidence-band width because $u_\alpha(y \mid \boldsymbol{d}) - l_\alpha(y \mid \boldsymbol{d})$ is exactly the vertical gap between the upper and lower calibrated bands at y, and $\Delta(\alpha)$ is the $|f|$-weighted integral of that gap, maximized over all designs. This completes the proof. □

A.5 Algorithm 1 (Conditional Kernel-Weighted Shape-Restricted Bands)

This algorithm computes the conditional kernel-weighted shape-restricted confidence band at a point $(\boldsymbol{d}, y)$ as follows:

A.6 Algorithm 2 (Stochastic Gradient Method for Robust BED)

For a fixed experimental design $\boldsymbol{d}$, the expected information gain (EIG) can be written as

$$\mathrm{EIG}(\boldsymbol{d}) = \mathbb{E}_{y \sim p(y|\boldsymbol{d})}\left[\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta})}\left[\log p(y \mid \boldsymbol{\theta}, \boldsymbol{d})\right] - \log p(y \mid \boldsymbol{d})\right] \tag{102}$$

Algorithm 1 Conditional Kernel-Weighted Shape-Restricted Bands $\big(l^t_\alpha(y \mid \boldsymbol{d}), u^t_\alpha(y \mid \boldsymbol{d})\big)$

Require: Samples $\{(\boldsymbol{d}_1, y_1), \ldots, (\boldsymbol{d}_N, y_N)\}$, target $(\boldsymbol{d}, y) \in \mathscr{D} \times [a, b]$, kernel $\mathcal{K}$, bandwidth h, conditional mode $\mu(\cdot)$, upper bound $\mathcal{U}(\cdot)$, number of groups M, confidence levels α, and $\alpha_1 < \alpha$ and Hölder constants (L, γ_D).

1: Compute the kernel weights, $w_i(\boldsymbol{d})$ $i = 1, \ldots, N$, using Equation 2
2: Sort the pooled outcomes $y_{(1)} < \cdots < y_{(N)}$.
3: Compute estimated group probabilities using equation 3 and their index using equation 4
4: Draw $X_1, \ldots, X_N \overset{\text{i.i.d.}}{\sim} \text{Exp}(1)$.
5: Compute finite-sample distributions using equation 6; estimate $c^-(\alpha_1), c^+(\alpha_1)$ as per Definition 3.1.
6: Construct $z_1 < \cdots < z_{\widehat{N}}$ from $\{a, b, \mu(\boldsymbol{d}), y, y_{(k_1(\boldsymbol{d}))}, \ldots, y_{(k_M(\boldsymbol{d}))}\}$.
7: Determine $\widehat{s}$ and $\widetilde{s}$.
8: Construct $\mathscr{H}_w(\boldsymbol{d}, \alpha)$ using equation 90–equation 92.
9: Solve equation 93 to obtain $\widetilde{l}_{\alpha_1}(y \mid \boldsymbol{d})$ and $\widetilde{u}_{\alpha_1}(y \mid \boldsymbol{d})$.
10: Compute $l^t_\alpha(y \mid \boldsymbol{d}) = \max\{0, \widetilde{l}_{\alpha_1}(y \mid \boldsymbol{d}) - Lh^{\gamma_D}\}$, $u^t_\alpha(y \mid \boldsymbol{d}) = \widetilde{u}_{\alpha_1}(y \mid \boldsymbol{d}) + Lh^{\gamma_D}$.
11: **return** $\big(l^t_\alpha(y \mid \boldsymbol{d}), u^t_\alpha(y \mid \boldsymbol{d})\big)$.

where $p(y \mid \boldsymbol{d})$ is the marginal predictive density. In our setting, this predictive density is uncertain and represented by an ambiguity set $\mathcal{Q}$. We construct a fixed reference predictive density from the midpoint of the confidence band,

$$\hat{q}(y \mid \boldsymbol{d}) = \frac{1}{2}\left(l_\alpha(y \mid \boldsymbol{d}) + u_\alpha(y \mid \boldsymbol{d})\right) \tag{103}$$

which is constructed independently of the candidate predictive density $q(y \mid \boldsymbol{d}) \in \mathcal{Q}$. Replacing the predictive density in the reference term of the EIG integrand by $\hat{q}$ gives the surrogate information-gain integrand

$$f(\boldsymbol{d}, y) = \mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta})}\left[\log p(y \mid \boldsymbol{\theta}, \boldsymbol{d})\right] - \log \hat{q}(y \mid \boldsymbol{d}) \tag{104}$$

For any $q \in \mathcal{Q}$, we then define the corresponding utility as

$$U(\boldsymbol{d}, q) = \mathbb{E}_{y \sim q(y \mid \boldsymbol{d})}\left[f(\boldsymbol{d}, y)\right] \tag{105}$$

Thus, the reference density $\hat{q}$ remains fixed for a given $\boldsymbol{d}$, while q determines the distribution under which the surrogate integrand is averaged. To account for predictive distributional uncertainty, we adopt a distributionally robust criterion by taking the worst-case utility over $\mathcal{Q}$,

$$\text{DR-EIG}(\boldsymbol{d}) = \inf_{q \in \mathcal{Q}} U(\boldsymbol{d}, q) \tag{106}$$

and obtain the robust Bayesian experimental design problem

$$\boldsymbol{d}^\star \in \arg\max_{\boldsymbol{d} \in \mathcal{D}} \text{DR-EIG}(\boldsymbol{d}) \tag{107}$$

When the likelihood model is well-specified, the model hypothesis class $\mathscr{P} = \{p(y \mid \boldsymbol{d}, \boldsymbol{\theta}) : \boldsymbol{\theta} \in \Theta\}$ contains the true DGP, i.e., there exists $\tilde{\boldsymbol{\theta}} \in \Theta$ such that

$$p(y \mid \boldsymbol{d}, \tilde{\boldsymbol{\theta}}) = \tilde{q}(y \mid \boldsymbol{d}) \qquad \text{for all } \boldsymbol{d} \in \mathcal{D}. \tag{108}$$

Moreover, since each candidate predictive density $q(y \mid \boldsymbol{d}) \in \mathcal{Q}$ is itself constructed from the well-specified model class, i.e., for some $\boldsymbol{\theta} \in \Theta$, $q(y \mid \boldsymbol{d}) = p(y \mid \boldsymbol{d}, \boldsymbol{\theta})$, the ambiguity set $\mathcal{Q}$ is centred on $\mathscr{P}$ and contains the candidate corresponding to $\boldsymbol{\theta} = \tilde{\boldsymbol{\theta}}$. For this candidate,

$$q(y \mid \boldsymbol{d}) = p(y \mid \boldsymbol{d}, \tilde{\boldsymbol{\theta}}) = \tilde{q}(y \mid \boldsymbol{d}) \tag{109}$$

Under a well-specified prior belief $p(\boldsymbol{\theta})$ the prior-averaged marginal recovers the true DGP, $p(y \mid \boldsymbol{d}) = \mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta})}[p(y \mid \boldsymbol{d}, \boldsymbol{\theta})] = \tilde{q}(y \mid \boldsymbol{d})$. Thus, the reference predictive density and the candidate predictive density coincide with the true marginal,

$$\hat{q}(y \mid \boldsymbol{d}) = q(y \mid \boldsymbol{d}) = p(y \mid \boldsymbol{d}) \tag{110}$$

In this case, the utility reduces to the standard EIG,

$$U(\boldsymbol{d}, q) = \mathbb{E}_{y\sim p(y|\boldsymbol{d})}\left[\mathbb{E}_{\boldsymbol{\theta}\sim p(\boldsymbol{\theta})}\left[\log p(y \mid \boldsymbol{\theta}, \boldsymbol{d})\right] - \log p(y \mid \boldsymbol{d})\right] = \mathrm{EIG}(\boldsymbol{d}) \tag{111}$$

Hence, the proposed DR-EIG extends the standard EIG by replacing the predictive distribution used for averaging with a candidate distribution from $\mathcal{Q}$ and by using the confidence-band midpoint as a fixed reference density. In particular, $-\log \hat{q}(y \mid \boldsymbol{d})$ is the surprisal of y under the reference density, while $f(\boldsymbol{d}, y)$ is the resulting surrogate pointwise information-gain integrand. The infimum over $\mathcal{Q}$ then accounts for uncertainty in the predictive distribution.

The proposed Stochastic Gradient method optimises *DR-EIG* to recommend the next optimal design variable. To this end, it uses dual optimisation problem in equation 18, which requires the associated subgradients with respect to $\boldsymbol{d}$, and λ. In this section, we first derive the analytical form of $\partial_{\boldsymbol{d}} F(\cdot)$, and $\partial_{\lambda} F(\cdot)$, then, using the importance sampling identity method, we present their Monte-Carlo estimates to be used in each iteration of Stochastic Gradient algorithm.

Given the definition of the utility function in equation 14,

$$U(\boldsymbol{d}, q) := \mathbb{E}_{y\sim q(y|\boldsymbol{d})}\left[f(\boldsymbol{d}, y)\right] = \int_a^b f(\boldsymbol{d}, y) q(y \mid \boldsymbol{d})\, dy \tag{112}$$

where

$$\begin{aligned} f(\boldsymbol{d}, y) &:= \mathbb{E}_{\boldsymbol{\theta}\sim p(\boldsymbol{\theta})}\left[\log p(y \mid \boldsymbol{\theta}, \boldsymbol{d})\right] - \log \widehat{q}(y \mid \boldsymbol{d}) \\ &= \mathbb{E}_{\boldsymbol{\theta}\sim p(\boldsymbol{\theta})}\left[\log p(y \mid \boldsymbol{\theta}, \boldsymbol{d})\right] - \log \frac{(l_\alpha(y \mid \boldsymbol{d}) + u_\alpha(y \mid \boldsymbol{d}))}{2} \end{aligned} \tag{113}$$

the optimisation problem for recommending a design point at each round of BED is as follows:

$$\boldsymbol{d}^\star \in \arg\max_{\boldsymbol{d}\in\mathscr{D}} \inf_{q\in\mathscr{Q}=\{l_\alpha(y|\boldsymbol{d}),\, u_\alpha(y|\boldsymbol{d})\}} U(\boldsymbol{d}, q) \tag{114}$$

where the inner minimisation problem is:

$$\begin{aligned} v^\star := \inf_{q\in\mathscr{Q}=\{l_\alpha(y|\boldsymbol{d}),\, u_\alpha(y|\boldsymbol{d})\}} & \int_a^b f(\boldsymbol{d}, y) q(y \mid \boldsymbol{d})\, dy \\ & \int q(y \mid \boldsymbol{d})\, dy = 1 \\ & l_\alpha(y \mid \boldsymbol{d}) \le q(y \mid \boldsymbol{d}) \\ & q(y \mid \boldsymbol{d}) \le u_\alpha(y \mid \boldsymbol{d}) \end{aligned} \tag{115}$$

Then, the dual reformulation of inner minimisation problem in equation 115 is:

$$\begin{aligned} v^\star := - \inf_{\lambda\in\mathbb{R},\, \alpha(y|\boldsymbol{d})\ge 0,\, \beta(y|\boldsymbol{d})\ge 0} \sup_{q(y|\boldsymbol{d})} \Bigg\{ & \int -f(\boldsymbol{d}, y)\, q(y \mid \boldsymbol{d})\, dy + \lambda\left(1 - \int q(y \mid \boldsymbol{d})\, dy\right) \\ & + \int \alpha(y \mid \boldsymbol{d})\left(q(y \mid \boldsymbol{d}) - l_\alpha(y \mid \boldsymbol{d})\right)\, dy \\ & + \int \beta(y \mid \boldsymbol{d})\left(u_\alpha(y \mid \boldsymbol{d}) - q(y \mid \boldsymbol{d})\right)\, dy \Bigg\} \end{aligned} \tag{116}$$

where $\lambda, \alpha(y \mid \boldsymbol{d}), \beta(y \mid \boldsymbol{d})$ are dual multipliers.

The simplified form of problem 116, is as follows:

$$v^\star := -\inf_{\lambda,\alpha(y|\boldsymbol{d}),\,\beta(y|\boldsymbol{d})} \sup_{q(y|\boldsymbol{d})} \int \underbrace{\Big(-f(\boldsymbol{d},y)-\lambda+\alpha(y\mid\boldsymbol{d})-\beta(y\mid\boldsymbol{d})\Big)}_{\leq 0} q(y\mid\boldsymbol{d})\,dy + \lambda + \int \beta(y\mid\boldsymbol{d})u_\alpha(y\mid\boldsymbol{d})\,dy - \int \alpha(y\mid\boldsymbol{d})l_\alpha(y\mid\boldsymbol{d})\,dy \tag{117}$$

where, the integrand term must be non positive, otherwise the sup results in $+\infty$.

$$-f(\boldsymbol{d},y)-\lambda+\alpha(y\mid\boldsymbol{d})-\beta(y\mid\boldsymbol{d})\leq 0 \Rightarrow \beta(y\mid\boldsymbol{d})\geq -f(\boldsymbol{d},y)-\lambda+\alpha(y\mid\boldsymbol{d}) \tag{118}$$

Therefore,

$$\beta(y\mid\boldsymbol{d}) = \max\left\{-f(\boldsymbol{d},y)-\lambda+\alpha(y\mid\boldsymbol{d}),0\right\} \tag{119}$$

Now, replacing $\beta(y\mid\boldsymbol{d})$, equation 116 reduces to the following form:

$$v^\star := -\inf_{\lambda,\alpha(y|\boldsymbol{d}),\,\beta(y|\boldsymbol{d})} \left[\lambda + \int \max\left\{-f(\boldsymbol{d},y)-\lambda+\alpha(y\mid\boldsymbol{d}),\,0\right\}\left(u_\alpha(y\mid\boldsymbol{d})-l_\alpha(y\mid\boldsymbol{d})\right)\,dy + \int \left(\beta(y\mid\boldsymbol{d})-\alpha(y\mid\boldsymbol{d})\right)l_\alpha(y\mid\boldsymbol{d})\,dy\right] \tag{120}$$

If $\alpha(y\mid\boldsymbol{d})\leq f(\boldsymbol{d},y)+\lambda$, then $\beta(y\mid\boldsymbol{d})=0$, and therefore $\alpha(y\mid\boldsymbol{d})$ can increase without bound. Hence, $\alpha(y\mid\boldsymbol{d})\geq f(\boldsymbol{d},y)+\lambda$, and $\alpha(y\mid\boldsymbol{d})=\max\{f(\boldsymbol{d},y)+\lambda,0\}$. After replacing $\alpha(y\mid\boldsymbol{d})$ and $\beta(y\mid\boldsymbol{d})$ in the objective function at 120:

$$v^\star := -\inf_{\lambda}\left[\lambda - \int \left(\lambda+f(\boldsymbol{d},y)\right)^+ l_\alpha(y\mid\boldsymbol{d})\,dy + \int \left(-\lambda-f(\boldsymbol{d},y)\right)^+ u_\alpha(y\mid\boldsymbol{d})\,dy\right] \tag{121}$$

Therefore, the dual objective function is as follows:

$$F(\boldsymbol{d},\lambda) = \lambda - \int \left(\lambda+f(\boldsymbol{d},y)\right)^+ l_\alpha(y\mid\boldsymbol{d})\,dy + \int \left(-\lambda-f(\boldsymbol{d},y)\right)^+ u_\alpha(y\mid\boldsymbol{d})\,dy \tag{122}$$

Now, taking the partial derivative of the above function with respect to $\boldsymbol{d}$:

$$\partial_{\boldsymbol{d}}(\lambda+f(\boldsymbol{d},y))^+ = \begin{cases} 0, & -f(\boldsymbol{d},y)\geq\lambda \\ f'(\boldsymbol{d},y), & -f(\boldsymbol{d},y)<\lambda \end{cases} \tag{123}$$

$$\partial_{\boldsymbol{d}}(-\lambda-f(\boldsymbol{d},y))^+ = \begin{cases} 0, & -f(\boldsymbol{d},y)\leq\lambda \\ -f'(\boldsymbol{d},y), & -f(\boldsymbol{d},y)>\lambda \end{cases} \tag{124}$$

$$\partial_{\boldsymbol{d}}F(\boldsymbol{d},\lambda) = -\int l_\alpha(y\mid\boldsymbol{d})f'(\boldsymbol{d},y)\mathbb{1}_{\{-f(\boldsymbol{d},y)<\lambda\}}dy - \int u_\alpha(y\mid\boldsymbol{d})f'(\boldsymbol{d},y)\mathbb{1}_{\{-f(\boldsymbol{d},y)>\lambda\}}dy - \int \partial_{\boldsymbol{d}}l_\alpha(y\mid\boldsymbol{d})\left(\lambda+f(\boldsymbol{d},y)\right)^+dy + \int \partial_{\boldsymbol{d}}u_\alpha(y\mid\boldsymbol{d})\left(-\lambda-f(\boldsymbol{d},y)\right)^+dy \tag{125}$$

Note that, based on the definition of $f(.)$ in equation 113:

$$f'(\boldsymbol{d},y) = \int \partial_{\boldsymbol{d}}\log p(y\mid\boldsymbol{\theta},\boldsymbol{d})\,p(\boldsymbol{\theta})\,d\boldsymbol{\theta} - \frac{\partial_{\boldsymbol{d}}l_\alpha(y\mid\boldsymbol{d})+\partial_{\boldsymbol{d}}u_\alpha(y\mid\boldsymbol{d})}{l_\alpha(y\mid\boldsymbol{d})+u_\alpha(y\mid\boldsymbol{d})} \tag{126}$$

Let's define the following shorthands to get a better representation of the above subgradient:

$$w = \begin{cases} -l_\alpha(y \mid \boldsymbol{d}), & -f(\boldsymbol{d}, y) < \lambda \\ -u_\alpha(y \mid \boldsymbol{d}), & -f(\boldsymbol{d}, y) > \lambda \end{cases} \tag{127}$$

$$c^A = \frac{-w}{2\hat{q}} - (\lambda + f(\boldsymbol{d}, y))^+ \tag{128}$$

$$c^B = \frac{-w}{2\hat{q}} + (-\lambda - f(\boldsymbol{d}, y))^+ \tag{129}$$

Then,

$$\partial_{\boldsymbol{d}} F(\boldsymbol{d}, \lambda) = \iint w\, \partial_{\boldsymbol{d}} \log p(y \mid \boldsymbol{\theta}, \boldsymbol{d})\, p(\boldsymbol{\theta})\, d\boldsymbol{\theta}\, dy + \int \partial_{\boldsymbol{d}} l_\alpha(y \mid \boldsymbol{d})\, c^A\, dy + \int \partial_{\boldsymbol{d}} u_\alpha(y \mid \boldsymbol{d})\, c^B\, dy \tag{130}$$

where the subgradients of the functional lower and upper bounds are available from the dual multipliers in the LP optimisation problem in equations 90-92.

$$\partial_{\boldsymbol{d}} l_\alpha(y \mid \boldsymbol{d}) = \partial_{\boldsymbol{d}} \max\{0, \widetilde{l}_{\alpha_1}(y \mid \boldsymbol{d}) - Lh^{\gamma_D}\} = \partial_{\boldsymbol{d}} \widetilde{l}_{\alpha_1}(y \mid \boldsymbol{d}) \mathbb{1}_{\{\widetilde{l}_{\alpha_1}(y|\boldsymbol{d}) > Lh^{\gamma_D}\}} \tag{131}$$

$$\partial_{\boldsymbol{d}} u_\alpha(y \mid \boldsymbol{d}) = \partial_{\boldsymbol{d}} \widetilde{u}_{\alpha_1}(y \mid \boldsymbol{d}) \tag{132}$$

Since, in LP problems, the derivative of the optimal solution with respect to any right-hand-side (RHS) parameter equals the corresponding dual multiplier, these multipliers can be obtained directly by solving the LP, without incurring any additional computational cost. Consequently, no extra differentiation step is required, as they are already a byproduct of any LP solver. This motivates formulating the dual of problem 93, as follows.

$$\tilde{\mathscr{L}}(\boldsymbol{d}) = \max_{\substack{\lambda_{mono},\, \vartheta \\ \lambda^-,\, \lambda^+ \\ ,\, \mu^L,\, \mu^U}} \sum_{m=2}^{M} \lambda_m^-\, c^-(\alpha_1; \boldsymbol{d}) - \sum_{m=2}^{M} \lambda_m^+\, c^+(\alpha_1; \boldsymbol{d}) - \vartheta - \sum_{j=1}^{\widehat{N}-1} \mu_j^U\, \mathcal{U}(\boldsymbol{d}) \tag{133}$$

s.t.

$$\boldsymbol{e}_v + A_{mono}(\boldsymbol{d})^\top \lambda_{mono} - \sum_{m=2}^{M} \lambda_m^- A_m(\boldsymbol{d})^\top + \sum_{m=2}^{M} \lambda_m^+ A_m(\boldsymbol{d})^\top + \vartheta\, \boldsymbol{w}(\boldsymbol{d}) - \mu^L + \mu^U = \mathbf{0} \tag{134}$$

$$\lambda_{mono} \geq 0, \quad \lambda^- \geq 0, \quad \lambda^+ \geq 0, \quad \mu^L \geq 0, \quad \mu^U \geq 0, \quad \vartheta \text{ free} \tag{135}$$

Equation 133 gives the dual of the confidence-band LP 93, where the coefficient objects explicitly indexed by $\boldsymbol{d}$ including $c^-(\alpha_1; \boldsymbol{d})$, $c^+(\alpha_1; \boldsymbol{d})$, $\mathcal{U}(\boldsymbol{d})$, $A_{mono}(\boldsymbol{d})$, $A_m(\boldsymbol{d})$, and $\boldsymbol{w}(\boldsymbol{d})$ to make it explicit that the entire LP, including its constraint matrices, is in fact a family of LPs parameterized by the design $\boldsymbol{d}$. The optimal dual variables $\lambda_{mono}^\star, \lambda^{-\star}, \lambda^{+\star}, \vartheta^\star, \mu^{L\star}, \mu^{U\star}$ are obtained as a byproduct of solving the primal problem. Since these multipliers satisfy the KKT stationarity condition at the optimum for every $\boldsymbol{d}$, they let us differentiate the optimal value $\tilde{v}^\star(\boldsymbol{d})$ with respect to $\boldsymbol{d}$. The general subgradients form is as follows:

$$\begin{aligned} \partial_{\boldsymbol{d}} \tilde{\mathscr{L}}(\boldsymbol{d}) = {} & \lambda_{mono}^{\star\top} \big(\partial_{\boldsymbol{d}} A_{mono}(\boldsymbol{d})\big) \boldsymbol{\beta}^\star + \sum_{m=2}^{M} \lambda_m^{-\star} \Big(- \partial_{\boldsymbol{d}} A_m(\boldsymbol{d})\, \boldsymbol{\beta}^\star + \partial_{\boldsymbol{d}} c^-(\alpha_1; \boldsymbol{d}) \Big) \\ & + \sum_{m=2}^{M} \lambda_m^{+\star} \Big(\partial_{\boldsymbol{d}} A_m(\boldsymbol{d})\, \boldsymbol{\beta}^\star - \partial_{\boldsymbol{d}} c^+(\alpha_1; \boldsymbol{d}) \Big) + \vartheta^\star \big(\partial_{\boldsymbol{d}} \boldsymbol{w}(\boldsymbol{d})\big)^\top \boldsymbol{\beta}^\star - \sum_{j=1}^{\widehat{N}-1} \mu_j^{U\star}\, \partial_{\boldsymbol{d}} \mathcal{U}(\boldsymbol{d}) \end{aligned} \tag{136}$$

$$\partial_{\boldsymbol{d}}\widetilde{\ell}_{\alpha_1}(y \mid \boldsymbol{d}) = \begin{cases} \lambda_{mono}^{\star(1)\top}\partial_{\boldsymbol{d}}A_{mono}\boldsymbol{\beta}^{\star(1)} + \sum_{m=2}^{M}\lambda_m^{-\star(1)}\Big(-\partial_{\boldsymbol{d}}A_m\boldsymbol{\beta}^{\star(1)} + \partial_{\boldsymbol{d}}c^-(\alpha_1;\boldsymbol{d})\Big) & y \le \mu(\boldsymbol{d}), \\ \quad + \sum_{m=2}^{M}\lambda_m^{+\star(1)}\Big(\partial_{\boldsymbol{d}}A_m\boldsymbol{\beta}^{\star(1)} - \partial_{\boldsymbol{d}}c^+(\alpha_1;\boldsymbol{d})\Big) + \vartheta^{\star(1)}(\partial_{\boldsymbol{d}}\boldsymbol{w})^\top\boldsymbol{\beta}^{\star(1)} - \sum_{j=1}^{\widehat{N}-1}\mu_j^{U\star(1)}\partial_{\boldsymbol{d}}\mathcal{U}(\boldsymbol{d}), \\ \lambda_{mono}^{\star(2)\top}\partial_{\boldsymbol{d}}A_{mono}\boldsymbol{\beta}^{\star(2)} + \sum_{m=2}^{M}\lambda_m^{-\star(2)}\Big(-\partial_{\boldsymbol{d}}A_m\boldsymbol{\beta}^{\star(2)} + \partial_{\boldsymbol{d}}c^-(\alpha_1;\boldsymbol{d})\Big) & y > \mu(\boldsymbol{d}). \\ \quad + \sum_{m=2}^{M}\lambda_m^{+\star(2)}\Big(\partial_{\boldsymbol{d}}A_m\boldsymbol{\beta}^{\star(2)} - \partial_{\boldsymbol{d}}c^+(\alpha_1;\boldsymbol{d})\Big) + \vartheta^{\star(2)}(\partial_{\boldsymbol{d}}\boldsymbol{w})^\top\boldsymbol{\beta}^{\star(2)} - \sum_{j=1}^{\widehat{N}-1}\mu_j^{U\star(2)}\partial_{\boldsymbol{d}}\mathcal{U}(\boldsymbol{d}), \end{cases} \tag{137}$$

$$\partial_{\boldsymbol{d}}\widetilde{u}_{\alpha_1}(y \mid \boldsymbol{d}) = -\begin{cases} \lambda_{mono}^{\star(3)\top}\partial_{\boldsymbol{d}}A_{mono}\boldsymbol{\beta}^{\star(3)} + \sum_{m=2}^{M}\lambda_m^{-\star(3)}\Big(-\partial_{\boldsymbol{d}}A_m\boldsymbol{\beta}^{\star(3)} + \partial_{\boldsymbol{d}}c^-(\alpha_1;\boldsymbol{d})\Big) & y \le \mu(\boldsymbol{d}), \\ \quad + \sum_{m=2}^{M}\lambda_m^{+\star(3)}\Big(\partial_{\boldsymbol{d}}A_m\boldsymbol{\beta}^{\star(3)} - \partial_{\boldsymbol{d}}c^+(\alpha_1;\boldsymbol{d})\Big) + \vartheta^{\star(3)}(\partial_{\boldsymbol{d}}\boldsymbol{w})^\top\boldsymbol{\beta}^{\star(3)} - \sum_{j=1}^{\widehat{N}-1}\mu_j^{U\star(3)}\partial_{\boldsymbol{d}}\mathcal{U}(\boldsymbol{d}), \\ \lambda_{mono}^{\star(4)\top}\partial_{\boldsymbol{d}}A_{mono}\boldsymbol{\beta}^{\star(4)} + \sum_{m=2}^{M}\lambda_m^{-\star(4)}\Big(-\partial_{\boldsymbol{d}}A_m\boldsymbol{\beta}^{\star(4)} + \partial_{\boldsymbol{d}}c^-(\alpha_1;\boldsymbol{d})\Big) & y > \mu(\boldsymbol{d}). \\ \quad + \sum_{m=2}^{M}\lambda_m^{+\star(4)}\Big(\partial_{\boldsymbol{d}}A_m\boldsymbol{\beta}^{\star(4)} - \partial_{\boldsymbol{d}}c^+(\alpha_1;\boldsymbol{d})\Big) + \vartheta^{\star(4)}(\partial_{\boldsymbol{d}}\boldsymbol{w})^\top\boldsymbol{\beta}^{\star(4)} - \sum_{j=1}^{\widehat{N}-1}\mu_j^{U\star(4)}\partial_{\boldsymbol{d}}\mathcal{U}(\boldsymbol{d}), \end{cases} \tag{138}$$

where superscripts (1), (2), (3), and (4) specify which LP problem needs to be solved in Equation 14. Additionally, $\lambda_{mono}^{\star}$, $\lambda_m^{-\star}$, $\lambda_m^{+\star}$, $\vartheta^{\star}$, and $\mu_j^{U\star}$ represent the optimal dual values associated, respectively, with the monotonicity constraint on A_{mono} in Equation 12, the lower band $c^-(\alpha_1)$ in Equation 12, the upper band $c^+(\alpha_1)$ in Equation 12, the weight constraint on $\boldsymbol{w}^\top\boldsymbol{\beta}$ in Equation 12, and the constraint satisfying the upper bound $\mathcal{U}(\boldsymbol{d})$ of the true DGP in Equation 13. Furthermore, the subgradients $\partial_{\boldsymbol{d}}c^-(\alpha_1;\boldsymbol{d})$, $\partial_{\boldsymbol{d}}c^+(\alpha_1;\boldsymbol{d})$ are available either analytically, through the gradient of kernel $\mathcal{K}$, or numerically, using the same generated samples for finding the Bootstrap calibration in Section 3.1. Analogously, $\partial_{\boldsymbol{d}}\mathcal{U}(\boldsymbol{d})$ is available either through prior knowledge or a constructed surrogate model for $\mathcal{U}(\boldsymbol{d})$.

Moreover, the subgradiant of $F(\boldsymbol{d},\lambda)$ with respect to λ is:

$$\begin{aligned} \partial_\lambda F(\boldsymbol{d},\lambda) &= 1 - \int l_\alpha(y \mid \boldsymbol{d})\mathbb{1}_{-f(\boldsymbol{d},y)<\lambda}\,dy + \int -u_\alpha(y \mid \boldsymbol{d})\mathbb{1}_{-f(\boldsymbol{d},y)>\lambda}\,dy \\ &= 1 + \int w dy \end{aligned} \tag{139}$$

Therefore, using the importance sampling identity method, we turn the integrals in $\partial_\lambda F(.)$ and $\partial_{\boldsymbol{d}}F(.)$ into expectation, and apply Monte-Carlo method to approximate them as follows:

$$g_{\boldsymbol{d}} = \frac{|I|}{B}\sum_{i=1}^{B}\left\{\frac{w_i}{\mathcal{M}}\sum_{j}^{\mathcal{M}}\partial_{\boldsymbol{d}}\log p(y_i \mid \boldsymbol{d}_k,\theta_j) + \partial_{\boldsymbol{d}}l_\alpha(y_i \mid \boldsymbol{d}_k)c_i^A + \partial_{\boldsymbol{d}}u_\alpha(y_i \mid \boldsymbol{d}_k)c_i^B\right\} \tag{140}$$

$$g_\lambda = 1 + \frac{|I|}{B}\sum_{i=1}^{B}w_i \tag{141}$$

Returning back to the equation equation 114, we have the following outer problem:

$$\boldsymbol{d}^\star \in \underset{\boldsymbol{d}\in\mathcal{D}}{\arg\max}\; v^\star(\boldsymbol{d}) = \underset{\boldsymbol{d}\in\mathcal{D}}{\arg\max}\left[-\inf_{\lambda\in\mathbb{R}}F(\boldsymbol{d},\lambda)\right] \tag{142}$$

Since $-\inf_\lambda F(\boldsymbol{d},\lambda) = \sup_\lambda\big(-F(\boldsymbol{d},\lambda)\big)$, the following is an equivalent optimisation problem:

$$\boldsymbol{d}^\star \in \arg\max_{\boldsymbol{d}\in\mathcal{D}} \sup_{\lambda\in\mathbb{R}} [-F(\boldsymbol{d},\lambda)] \tag{143}$$

Folding both optimisations into a single problem that is usable for the stochastic gradient algorithm:

$$(\boldsymbol{d}^\star,\lambda^\star) \in \arg\max_{\boldsymbol{d}\in\mathcal{D},\ \lambda\in\mathbb{R}} [-F(\boldsymbol{d},\lambda)] \tag{144}$$

using the subgradients $\partial_{\boldsymbol{d}}F(\boldsymbol{d},\lambda)$ and $\partial_\lambda F(\boldsymbol{d},\lambda)$ to perform gradient ascent jointly in $(\boldsymbol{d},\lambda)$. Algorithm 2 presents the proposed stochastic gradient method for robust BED.

Algorithm 2 Stochastic Gradient Method for Robust BED

Require: Initial point $(\boldsymbol{d}_0,\lambda_0)\in\mathcal{D}\times\mathbb{R}$; sampling box $I\supseteq[a,b]$; step-size parameter $\eta_0>0$; batch size B; total iterations K

1: **for** $k=0,1,\ldots,K-1$ **do**

2: Set the step size $\eta_k \leftarrow \dfrac{\eta_0}{k+1}$.

3: Sample $y_1,\ldots,y_B$ independently and uniformly from I.

4: Run Algorithm 1 at $\boldsymbol{d}=\boldsymbol{d}_k$ to compute $l_\alpha(y_i \mid \boldsymbol{d}_k)$, $u_\alpha(y_i \mid \boldsymbol{d}_k)$, and $\partial_{\boldsymbol{d}} l_\alpha(y_i \mid \boldsymbol{d}_k)$, $\partial_{\boldsymbol{d}} u_\alpha(y_i \mid \boldsymbol{d}_k)$, for $i=1,\ldots,B$.

5: For each $i=1,\ldots,B$, set

$$\omega_i = \begin{cases} -l_\alpha(y_i \mid \boldsymbol{d}_k), & f(\boldsymbol{d}_k,y_i) > -\lambda_k, \\ -u_\alpha(y_i \mid \boldsymbol{d}_k), & f(\boldsymbol{d}_k,y_i) < -\lambda_k. \end{cases} \tag{145}$$

6: Construct the stochastic subgradients $g_{\boldsymbol{d}}$ and g_λ from equations equation 140 and equation 141.

7: Update the design variable:

$$\boldsymbol{d}_{k+1} \leftarrow \mathrm{Proj}_{\mathcal{D}}(\boldsymbol{d}_k - \eta_k g_{\boldsymbol{d}}) \tag{146}$$

8: Update the dual variable:

$$\lambda_{k+1} \leftarrow \lambda_k - \eta_k g_\lambda \tag{147}$$

9: Update the running averages:

$$\bar{\boldsymbol{d}}_{k+1} \leftarrow \frac{\sum_{i=0}^k \eta_i \boldsymbol{d}_i}{\sum_{i=0}^k \eta_i}, \qquad \bar{\lambda}_{k+1} \leftarrow \frac{\sum_{i=0}^k \eta_i \lambda_i}{\sum_{i=0}^k \eta_i} \tag{148}$$

10: **end for**

11: **return** $(\bar{\boldsymbol{d}}_K,\ \bar{\lambda}_K)$ and $v^\star \approx -F(\bar{\boldsymbol{d}}_K,\bar{\lambda}_K)$.

A.7 Experimental Details

This section provides additional details on the competing methods, the performance measures used for the comparison purposes, experimental settings of benchmarking problems and the obtained additional results. The experiments were conducted using relevant Python packages, such as PyTorch (Paszke et al., 2019). Additionally, interested readers are referred to relevant implementations in the literature, such as (Tang et al., 2026), (Barlas et al., 2025), and (Foster et al., 2019).

A.7.1 Competing methods

Random Design selection strategy In this algorithm at each round of BED the next design would be selected randomly from the design space. This baseline shows the setting that if the experimenter wants to assess the performance of suggested designs from an algorithm with the scenario if the designs are selected through a uniform random distribution over the design space without relying on any model-based utility for selecting designs. It allows to assess the value added by an algorithm that solves a stochastic optimisation problem to recommend designs.

Bayesian Adaptive Design (BAD) This algorithm suggests the design based on the traditional *Expected Information Gain* (EIG) as the utility function. It solves a stochastic optimisation problem that relies fully on the assumed statistical model. In other words, BAD assumes the experimenter is satisfied with the likelihood function and selects the design that maximises the expected information gain with respect to the assumed model. In our experiments, this baseline is useful to illustrate the differences in the performance between accounting for the model misspecification and ignoring the model unceratainty when selecting the optimal design.

A.7.2 PERFORMANCE MEASURES

This subsection describes the performance measures that we used and how they were computed. Considering D available designs d_i at each round of experiments, we sample N outcomes from the posterior predictive distribution resulting in $D \times N$ pairs of (d_i, y_i^m) available samples. Additionally, we draw M sample outcomes from the true DGP yielding in total $D \times M$ pair of $(d_i, y_i^{trueDGP})$ available samples from the true DGP. With these samples in hand, we compute the following measure at each round of experiments:

Maximum Mean Discrepancy (MMD): Maximum Mean Discrepancy (MMD) is a statistical measure used to compute the difference between two statistical distributions. Instead of directly computing the difference in the support of the distributions, it uses a kernel function to map the data into a Reproducing Kernel Hilbert Space (RKHS), in which the MMD then computes the difference in the mean embeddings of the two distributions. The MMD is calculated as follows (Tang et al., 2026):

$$\begin{aligned}\text{MMD}^2 = \frac{1}{D}\sum_{l=1}^{D}\Bigg[&\frac{1}{N^2}\sum_{i=1}^{N}\sum_{i'=1}^{N}k\big(y_i^m, y_{i'}^m\big)\\ &+\frac{1}{M^2}\sum_{j=1}^{M}\sum_{j'=1}^{M}k\big(y_j^{\text{true DGP}}, y_{j'}^{\text{true DGP}}\big)\\ &-\frac{2}{NM}\sum_{i=1}^{N}\sum_{j=1}^{M}k\big(y_i^m, y_j^{\text{true DGP}}\big)\Bigg]\end{aligned} \tag{149}$$

where $k(x, y)$ is the radial basis function (RBF) with a bandwidth h as follows:

$$k(x, y) = \exp\left(-\frac{\|x - y\|_2^2}{2h^2}\right) \tag{150}$$

A lower MMD, ideally close to 0, indicates that the two distributions are similar.

Wasserstein Distance (WD): It compares two candidate distributions in terms of how much transport cost is required to reshape one distribution to the other one. In the case of one dimensional, is shows on average how much shift is required to distribution to match the target distribution. The mathematical formula is as follows (Sinha et al., 2017):

$$\text{WD} = \frac{1}{D}\sum_{l=1}^{D}\frac{1}{N}\sum_{i=1}^{N}\left|y_i^m - y_i^{\text{true DGP}}\right| \tag{151}$$

where y_i^m and $y_i^{\text{true DGP}}$ denote the i-th sample from the predictive and true DGP distributions, respectively, for design d. This corresponds to the empirical 1-Wasserstein distance between the two one-dimensional distributions. A lower amount of WD, ideally close to 0, represents that a small amount of mass is required to push the distribution to turn it into the target distribution.

Root Mean Squared Error (RMSE): To compare the predictive performance of two distributions, we compute the deviation in the outcome for each sample and take the square of it, representing the estimation error for that sample. In this way, larger errors contribute more to the computation of the error, whereas the smaller errors contribute less. Taking the average over all available samples

and then the square root gives the RMSE. The mathematical formula is as follows (Barlas et al., 2025):

$$\text{RMSE} = \frac{1}{D}\sum_{l=1}^{D}\sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i^m - y_i^{\text{true DGP}})^2} \tag{152}$$

A lower RMSE means on average the predicted values are close to the true values.

Negative Log-Likelihood (NLL): It measures how close the predicted probability of the outcomes is to the real distribution of the true DGP. The mathematical formula is as follows (Barlas et al., 2025):

$$\text{NLL} = -\frac{1}{D}\sum_{l=1}^{D}\frac{1}{N}\sum_{i=1}^{N}\log\left(\frac{1}{M}\sum_{j=1}^{M}p\left(y_i^{\text{true DGP}} \mid \theta_j, d_l\right)\right) \tag{153}$$

where $y_i^{\text{true DGP}}$ are samples from the true data-generating process, θ_j are samples from posterior distribution of model parameters, and $p(\cdot \mid \theta, d_l)$ is the likelihood of the assumed statistical model computed at design point d_l. A lower NLL means the model is assigning probabilities to the outcomes that are close to those of the true DGP.

A.7.3 Experimental setting

Here we introduce the experimental settings of the benchmarking problems, for which the associated true DGP, misspecified model, and prior distribution of the model parameters are presented.

Polynomial Regression problem Polynomial regression is an experimental setting in which the outcome of the model is computed via a polynomial equation plus added noise. The goal is to specify the predictive model by finding the coefficients of the assumed model structure through collecting samples. A polynomial regression of degree two is as follows:

$$\begin{cases} \text{True DGP:} \qquad \mathcal{D} \to Y \\ y = \beta_0 + \beta_1 d + \beta_2 d^2 + \epsilon, \quad \epsilon \sim p(\epsilon) \end{cases} \tag{154}$$

In our experiments, we assume the following scenarios for the well-specified and misspecified models:

Well-specified likelihood: We assume the true DGP function is as follows:

$$y = 1 + 0.5d + 1.5d^2 + \epsilon, \quad \epsilon \sim \mathcal{N}(0, 1) \tag{155}$$

The experimenter has specified a model with the following structure:

$$y = \beta_0 + \beta_1 d + \beta_2 d^2 + \epsilon, \quad \epsilon \sim \mathcal{N}(0, 1) \tag{156}$$

The prior belief on the model parameters follow the distribution:

$$\boldsymbol{\beta} = \begin{bmatrix}\beta_0 \\ \beta_1 \\ \beta_2\end{bmatrix} \sim \mathcal{N}_3(\mathbf{0}, \tau^2 I^{(3)}) \tag{157}$$

where $\mathbf{0}$ is the zero mean vector and $I^{(3)}$ is the identity matrix of size 3.

Misspecified likelihood: We assume the same true DGP function as in the Equation (155). The experimenter has specified a model with the following structure:

$$y = \beta_0 + \beta_1 d + \epsilon, \quad \epsilon \sim \mathcal{N}(0, 1) \tag{158}$$

The prior belief on the model parameters follow the distribution:

$$\boldsymbol{\beta} = \begin{bmatrix}\beta_0 \\ \beta_1\end{bmatrix} \sim \mathcal{N}_2(\mathbf{0}, \tau^2 I^{(2)}) \tag{159}$$

where $\mathbf{0}$ is the zero mean vector and $I^{(2)}$ is the identity matrix of size 2.

Pharmacokinetics problem A pharmacokinetic model uses a mathematical equation to make predictions about the drug concentration in the blood or tissue of a patient after a given dose of that drug. It helps researchers answer questions such as the time required for a drug to be cleared from the body and the required amount of drug dosage for a patient. A one-compartment model with oral absorption uses the following formula to compute the concentration:

$$C(t) = \frac{F\, D_{\text{dose}}}{V_d} \cdot \frac{k_a}{k_a - k_e} \left(\exp(-k_e t) - \exp(-k_a t)\right) \tag{160}$$

where F represents the bioavailability, D_{dose} is the administered dose, and V_d is the volume of distribution. Additionally, k_a is the assumed absorption rate and k_e represents the elimination rate of the drug from the patient's body. Here, t is the elapsed time since the drug was consumed. Researchers make critical decisions regarding the patient's testing time in order to compute the concentration of the drug in the patient's blood or tissue.

Well-specified likelihood: We assume the drug concentration $C(.)$, measured at time t after consumption, is as follows:

$$c(t \mid \theta) = \frac{F\, D_{\text{dose}}}{V_d} \cdot \frac{k_a}{k_a - k_e} \left(\exp(-k_e t) - \exp(-k_a t)\right) \tag{161}$$

and the corresponding observations are collected as follows:

$$y(t \mid \theta) = c(t \mid \theta)(1 + \epsilon) + \eta \tag{162}$$

where $\epsilon \sim \mathcal{N}(0, 0.01)$ and $\eta \sim \mathcal{N}(0, 0.1)$ are multiplicative and additive observation noises, respectively. Additionally, we assume in the true DGP function $\theta = (k_a, k_e, V_d)$, $D_{dose} = 300$, and $F = 20$. The prior distribution for the model parameters θ is as follows:

$$\log \theta \sim \mathcal{N}\left(\begin{bmatrix} \log 1 \\ \log 0.1 \\ \log 18 \end{bmatrix}, \begin{bmatrix} 0.05 & 0 & 0 \\ 0 & 0.05 & 0 \\ 0 & 0 & 0.05 \end{bmatrix}\right) \tag{163}$$

Therefore, $y(t \mid \theta)$ has the following distribution:

$$y(t \mid \theta) \sim \mathcal{N}\left(c(t \mid \theta),\, 0.01\, c(t \mid \theta)^2 + 0.1\right) \tag{164}$$

Misspecified likelihood: We assume the same true DGP as in Equation 162. The experimenter has assumed a dual-absorption model with two absorption rates k_{a1} and k_{a2}. It is assumed that $k_{a1} = k_a$ and $k_{a2} = \rho k_{a1}$, where $\rho \in (0, 1)$. Assuming $f \in (0, 1)$, the mean concentration is as follows:

$$\begin{aligned} c_{\text{mis}}(t \mid \theta, \rho, f) = \frac{F\, D_{dose}}{V_d} \Bigg[& f \cdot \frac{k_{a1}}{k_{a1} - k_e} \left(\exp(-k_e t) - \exp(-k_{a1} t)\right) \\ & - (1 - f) \cdot \frac{k_{a2}}{k_{a2} - k_e} \left(\exp(-k_e t) - \exp(-k_{a2} t)\right) \Bigg] \end{aligned} \tag{165}$$

We set $\rho = 0.2$ and $f = 0.7$. The observation in the assumed model is as follows:

$$y_{\text{mis}}(t \mid \theta, \rho, f) \sim \mathcal{N}\left(c_{\text{mis}}(t \mid \theta, \rho, f),\, 0.02\, c_{\text{mis}}(t \mid \theta, \rho, f)^2 + 0.2\right) \tag{166}$$

Source location finding problem A source location finding problem assumes there are M unknown sources of signal radiation in the k-dimensional space. The experimenter inspects a chosen location and observes a noisy total signal intensity at the inspection point. The task is to find the unknown locations of the M sources $\theta = \{\theta_i\}_{i=1}^M$ of the signals based on the collected observations, while the total intensity of the signal at location d is as follows (Rainforth et al., 2024; Tang et al., 2026):

$$\mu(d, \theta) = b + \sum_{i=1}^{M} \frac{\alpha_i}{\kappa + ||d - \theta_i||^2} \tag{167}$$

where $\mu(\cdot)$ is the total concentration of signals, b represents the background density, α_i shows the strength of the i-th source, and κ is a small constant in the denominator.

Well-specified likelihood: We assume that there are 2 sources of signals, and the true DGP has the following formula:

$$\mu(d, \theta) = 10^{-1} + \frac{0.4}{10^{-4} + ||d - \theta_1||^2} + \frac{0.4}{10^{-4} + ||d - \theta_2||^2} \tag{168}$$

$$\log y \mid d, \theta \sim \mathcal{N}\left(\log \mu(d, \theta), 0.1\right) \tag{169}$$

The experimenter has specified a model with the following structure:

$$\mu(d, \theta) = 10^{-1} + \frac{0.4}{10^{-4} + ||d - \theta_1||^2} + \frac{0.4}{10^{-4} + ||d - \theta_2||^2} \tag{170}$$

$$\log y \mid d, \theta \sim \mathcal{N}\left(\log \mu(d, \theta), 0.1\right) \tag{171}$$

And the prior belief on the model parameters is a 2-dimensional gaussian distribution as follows:

$$\theta_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, I^{(2)}), \quad i = 1, 2 \tag{172}$$

where $\mathbf{0}$ is the zero vector and $I^{(2)}$ is the identity matrix of size 2.

Mis-specified likelihood: We assume the same true DGP as in equation 169. The experimenter has specified a model with the following structure:

$$\mu(d, \theta) = 2 \times 10^{-2} + \frac{0.5}{4 \times 10^{-4} + ||d - \theta_1||^2} + \frac{0.4}{4 \times 10^{-4} + ||d - \theta_2||^2} \tag{173}$$

$$\log y \mid d, \theta \sim \mathcal{N}\left(\log \mu(d, \theta), 0.1\right) \tag{174}$$

And the prior belief on the model parameters is a 2-dimensional gaussian distribution as follows:

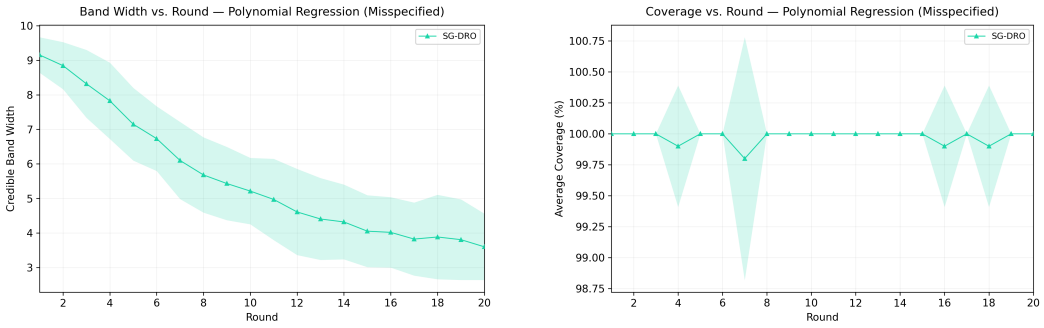
$$\theta_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, I^{(2)}), \quad i = 1, 2 \tag{175}$$

where $\mathbf{0}$ is the zero vector and $I^{(2)}$ is the identity matrix of size 2.

A.7.4 Additional Experimental results

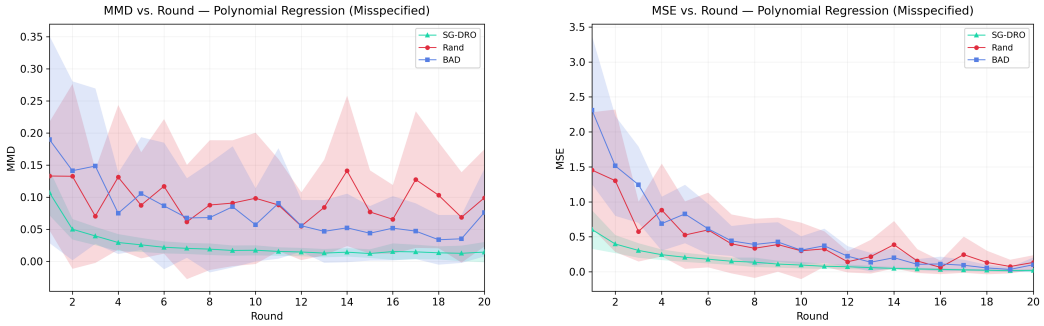
This section presents additional experimental results for the benchmarking problems. Experimental results show that, with the proposed framework, experts can efficiently work with misspecified models and choose design points that are guaranteed to have bounded regret aligned with the confidence level on the coverage.

Polynomial regression: Misspecified model Additional experimental results for the polynomial regression problem under the misspecified model are presented below.



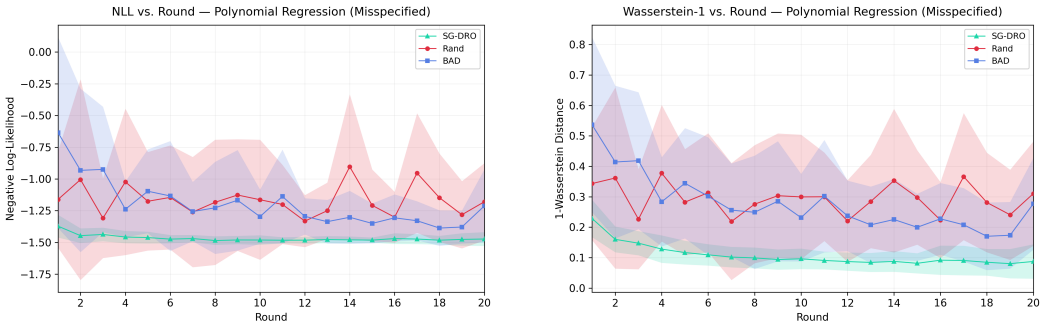
(a) Bandwidth

(b) Credible interval



(c) MMD

(d) MSE



(e) NLL

(f) Wasserstein distance

Figure 4: Results for the misspecified polynomial regression model.

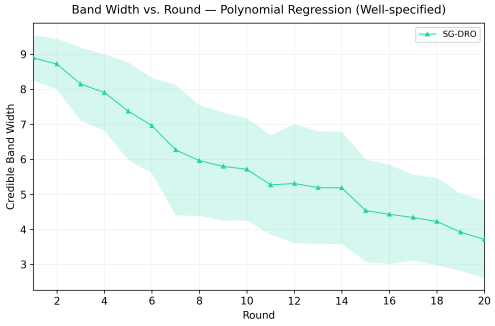
Figure 4a illustrates the decreasing trend in the average width of the constructed ambiguity set. It represents that the distance between the upper bound and lower bound decreases as the experiments progresses. The smaller width indicates that the true DGP in bounded by increasingly tight functional lower and upper bounds, which in turn means that the worst-case distribution in this bound is closer to the true DGP. Consistent with this, Figure 4b shows that the constructed ambiguity set covers the true DGP with a percentage exceeding the nominal confidence level of 90%, consistent with the coverage guarantee proposed in theorem 3.1 for true DGP.

In terms of predictive performance, a comparison of baseline methods in Figure 4c shows that the proposed SG-DRO algorithm has a superior performance with a decreasing trend in MMD. Analogous to the MMD results, Figure 4d shows that the proposed SG-DRO algorithm achieves a better performance in terms of MSE. It shows the predictive performance of misspecified likelihood model, updated using the design outcomes suggested by SG-DRO has a superior performance. This can be attributed to the principle that pushing the misspecified likelihood model towards the surro-

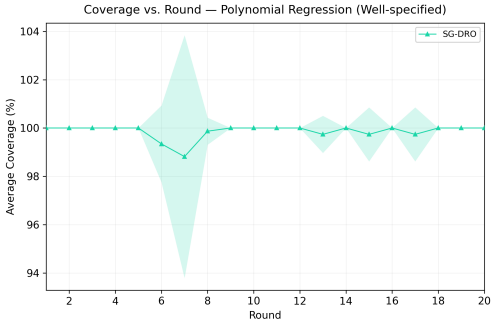
gate model at the midpoint of the lower and upper bounds has enabled DR-BED to outperform the other methods. Unlike the BAD method, which primarily focus on posterior uncertainty reduction, DR-BED provides a mechanism to guide the likelihood model towards the surrogate model using the information contained in the lower and upper bounds, which is useful in the case of misspecification. This guidance appears to be a key factor contributing to the improved predictive performance.

This improved generalisation is further reflected in Figure 4e, which shows a decreasing trend in the NLL suggesting improved generalisation performance of SG-DRO algorithm. Finally, the decreasing trend in the W1 distance in Figure 4f represents that the learned Misspecified likelihood model is converging as closely as possible to the true DGP. It is a critical result, as it shows that even under the misspecification, the learned model resembles the true DGP more closely than the other baseline methods, suggesting superior performance in real-world applications such as healthcare.

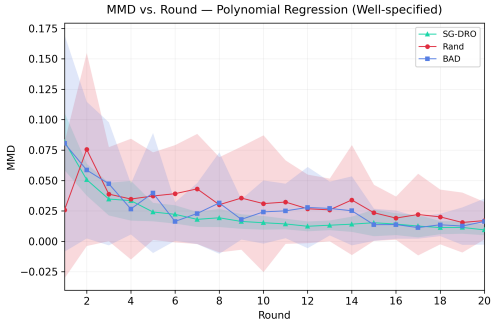
Polynomial regression: Well-specified model Figures for the polynomial regression problem under the well-specified model are presented below.



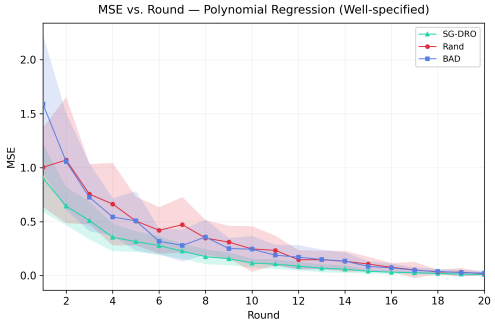
(a) Bandwidth



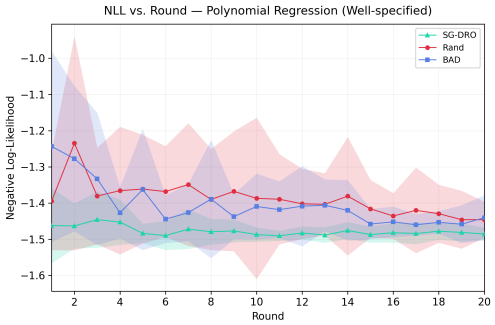
(b) Credible interval



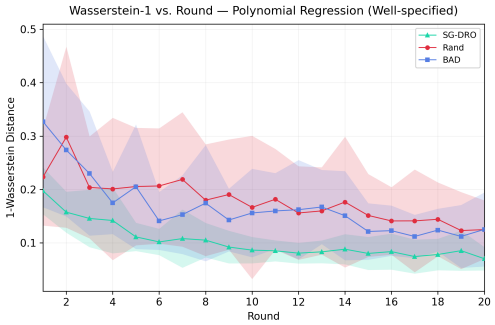
(c) MMD



(d) MSE



(e) NLL



(f) Wasserstein distance

Figure 5: Results for the well-specified polynomial regression model.

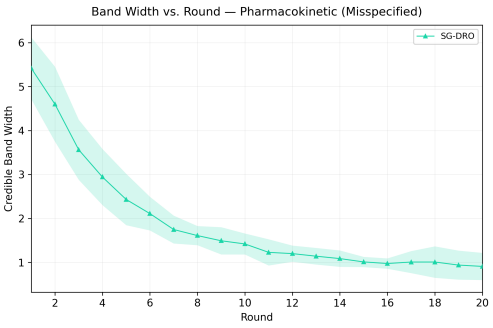
Figure 5a shows a decreasing trend in the width of the constructed ambiguity set under the well-specified likelihood model for polynomial regression. As the experiments progress and more evidence from the experiment become available, the encoded information in the functional bounds improves, yielding tighter bounds for the true DGP at the specified confidence level of 90%. Correspondingly, Figure 5b shows the convergence percentage of the constructed ambiguity set for the true DGP. It is aligned with the proposed theorem 3.1, as all the percentages are greater than the specified confidence level of 90%.

Turning to predictive performance, Figure 5c shows a decreasing trend in MMD measure across the baselines. As the model is well-specified, SG-DRO performs at least as well as the other baselines. There are some points in the early experiments where BAD and RAND show a better performance

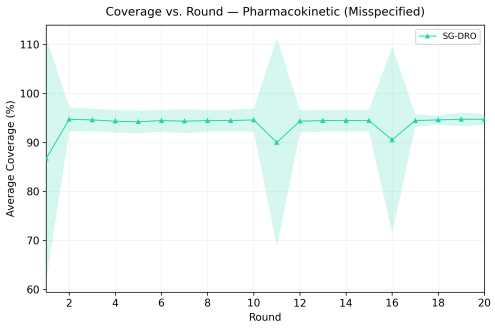
due to the learning process in the SG-DRO to learning the ambiguity set. As expected, Figure 5d shows that the MSE performance of the SG-DRO is superior to the competing baselines.

Although the other baselines also show a decreasing trend in the NLL in Figure 5e, reflecting their capacity for generalisation under well-specified likelihood model, the SG-DRO performs better across all the rounds of the experiments, which roots to the closeness of the targeted surrogate true DGP model at the midpoint of lower and upper bounds. Consistently, Figure 5f shows that, in the well-specified likelihood model, as expected, the SG-DRO leads to a distribution closer to the true DGP because of the constructed ambiguity set for the true DGP under the specified confidence level.

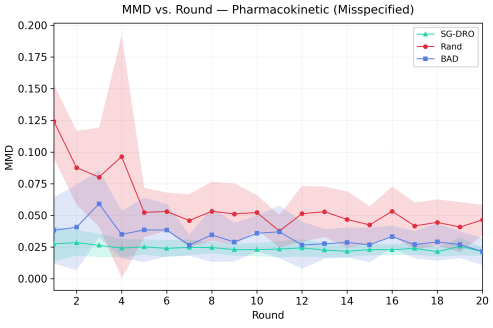
Pharmacokinetic problem Figures for the pharmacokinetic (PK) problem under the misspecified and well-specified models are presented below.



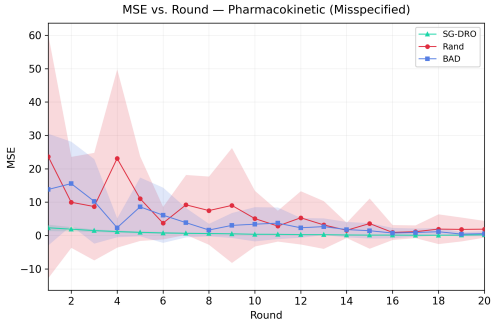
(a) Bandwidth



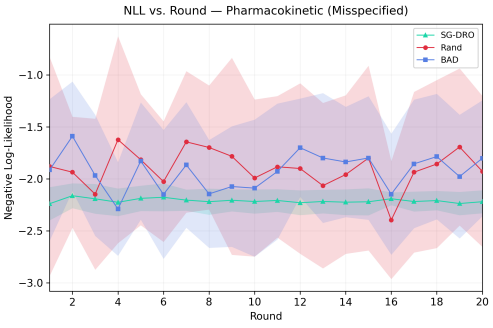
(b) Credible interval



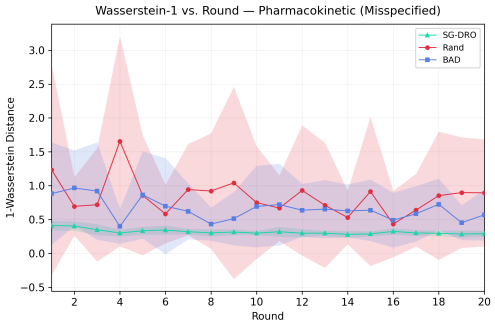
(c) MMD



(d) MSE



(e) NLL



(f) Wasserstein distance

Figure 6: Results for the misspecified pharmacokinetic (PK) model.

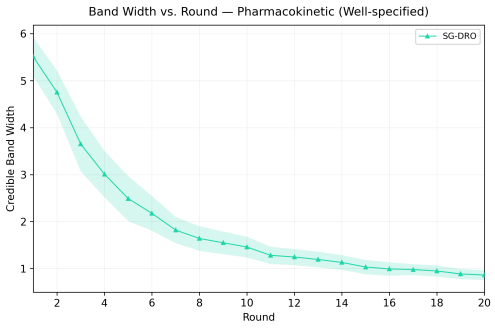
Figure 6a shows a decreasing trend in the distance between the constructed functional lower and upper bounds. This trend reflects an improvement in the information encoded in the functional bounds, and closeness to the true DGP. Robust design decisions made under these tight bounds would be closer to that of the optimal design under true DGP. It demonestrates the usefulness and applicability of proposed Distributionally Robust BED in the healthcare applications. In line with this, the average coverage percentage of the constructed ambiguity sets, shown in Figure 6b, exceeds the specified confidence level of 90% over the course of BED, except for the first design.

In terms of predictive performance, Figure 6c represents the lower values of MMD for the SG-DRO algorithm, which indicates that recommended designs under SG-DRO are more informative on the true DGP distribution. While the random design strategy selects designs randomly, and BAD

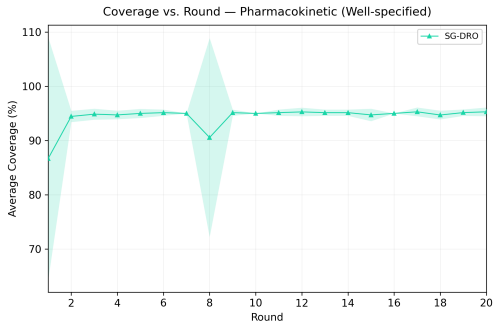
has targeted the uncertainty in the parameters of the misspecified model, the SG-DRO targets the surrogate of true DGP constructed at the midpoint of the constructed ambiguity sets. This leads SG-DRO to push the misspecified model toward parameter values that has closer to the true DGP. Consistent with this, Figure 6d shows that the predictive performance of the learned misspecified likelihood is superior compared to the other competing baseline methods, and the lower NLL in Figure 6e suggests the better generalisation performance of the SG-DRO algorithm in comparison to the baseline methods.

Finally, the lower W1 distance in Figure 6f shows that targeting the midpoint function between the lower and upper bound of the ambiguity set leads to learning such a misspecified likelihood model that is effectively close to the true DGP.

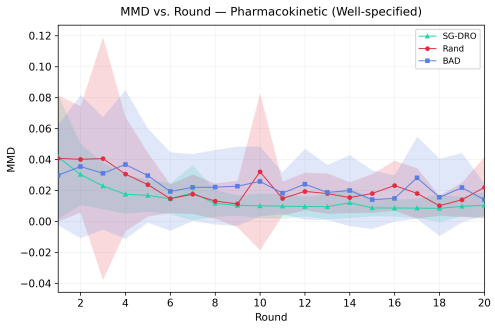
Pharmacokinetic problem: Well-specified model Figures for the pharmacokinetic (PK) problem under the well-specified model are presented below.



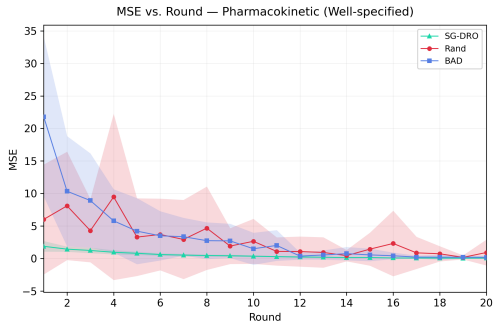
(a) Bandwidth



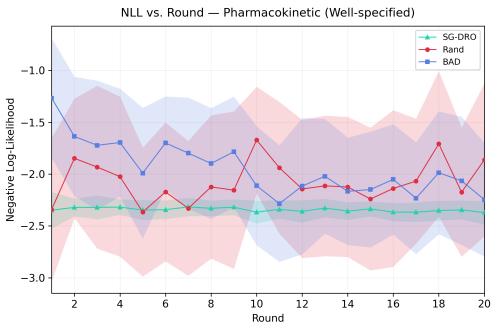
(b) Credible interval



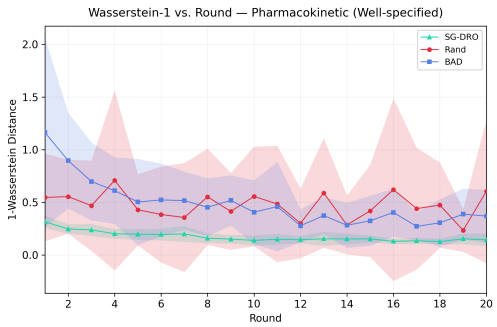
(c) MMD



(d) MSE



(e) NLL



(f) Wasserstein distance

Figure 7: Results for the well-specified pharmacokinetic (PK) model.

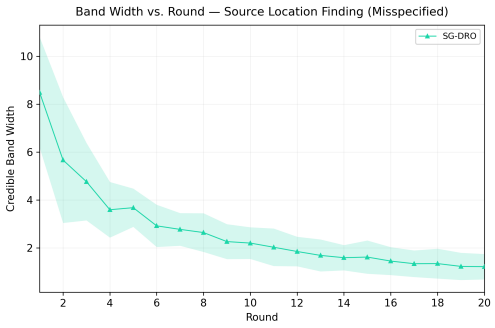
The decreasing trend in the distance between the functional bounds across the experiments, shown in Figure 7a, demonstrates the effectiveness of the proposed ambiguity set construction in terms of converging to the true DGP. In line with the proposed theorem 3.1, Figure 7b shows that the coverage percentage of the constructed ambiguity sets meets the specified confidence level of 90%.

This tighter ambiguity set translates into better predictive performance. The lower MMD of SG-DRO in Figure 7c indicates the lower discrepancy between the learned well-specified likelihood model and the true DGP, suggesting that the selected design points are effectively informative for learning the true DGP, rather than being solely concentrated on parameter informativeness. Correspondingly, the lower MSE in Figure 7d reflects the higher predictive performance of the learned likelihood model under the proposed SG-DRO algorithm, while the lower NLL over the experiments

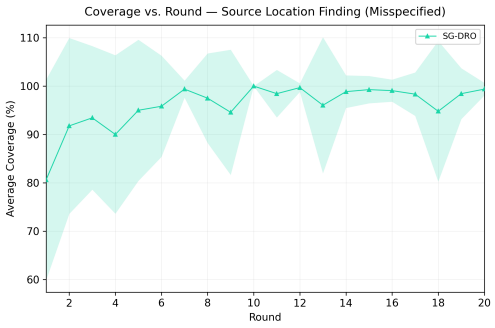
in Figure 7e shows the better generalisation performance of the learned model under well-specified likelihood model.

Finally, the lower W1 in Figure 7f indicates the closeness of the learned conditional likelihood model to the true DGP at each round of the experiment.

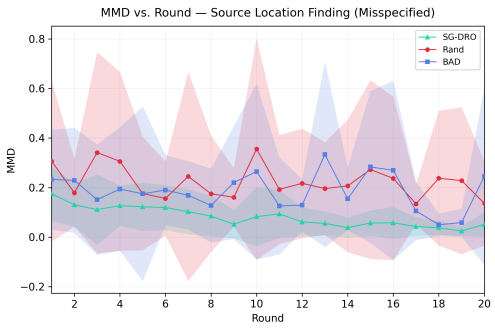
Source location finding: Misspecified model Figures for the source location finding problem under the misspecified model are presented below.



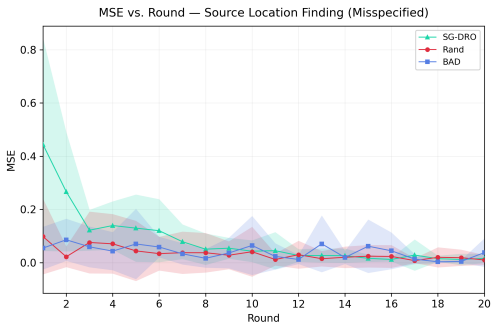
(a) Bandwidth



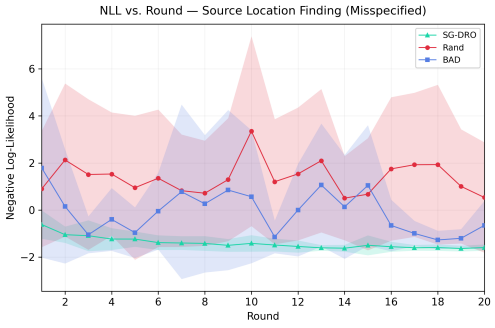
(b) Credible interval



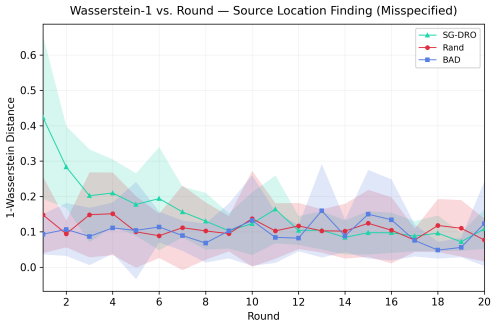
(c) MMD



(d) MSE



(e) NLL



(f) Wasserstein distance

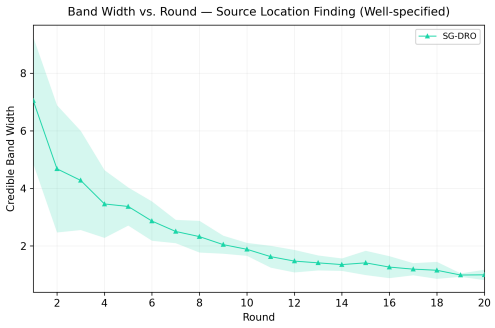
Figure 8: Results for the misspecified source location finding model.

Figure 8a shows a decreasing trend in the average distance between the constructed functional lower and upper bounds for the source location finding problem under the misspecified likelihood model. This trend reflects an improvement in the information encoded in the functional bounds and their increasing closeness to the true DGP. Therefore, robust design decisions made under these tight bounds are closer to the optimal design under the true DGP, demonstrating the usefulness and applicability of the proposed Distributionally Robust BED framework in complex experimental problems. In line with the proposed Theorem 3.1, Figure 8b shows that the average coverage percentage of the true DGP by the constructed ambiguity sets over the course of experiments exceeds the specified confidence level of 90%.

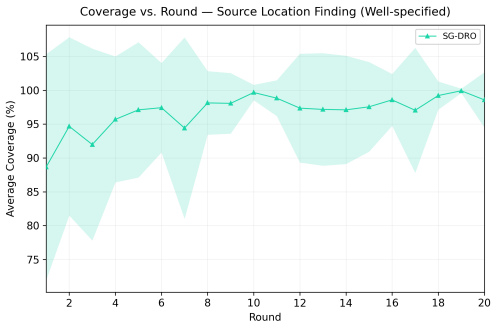
Turning to predictive performance, Figure 8c shows lower MMD values for the SG-DRO algorithm under a misspecified likelihood, indicating that the designs selected under SG-DRO are more informative about the true DGP distribution than those selected by the baseline methods. SG-DRO targets the surrogate of the true DGP constructed at the midpoint of the ambiguity sets. This leads SG-DRO to push the misspecified model toward parameter values that make it closer to the true DGP of the concentrated signal. Consistent with this, Figure 8d shows that the predictive performance of the learned misspecified likelihood model is acceptable compared to the other competing baseline methods in the source location finding problem, and the lower NLL in Figure 8e indicates better generalisation performance of SG-DRO in comparison to the baseline methods.

Figure 8f shows that the SG-DRO explores to learn the shape of the true DGP in the early stages of the experiments. Therefore, in the early stages, it requires more mass to push the learned conditional density to the true DGP. This exploration of the space helps to construct better ambiguity sets under the specified confidence level.

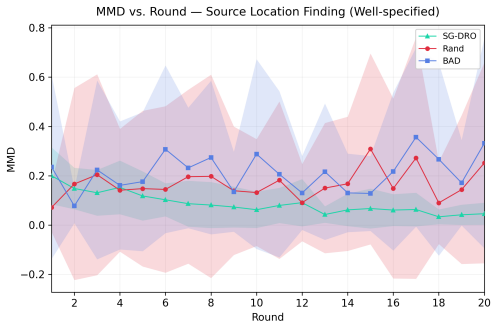
Source location finding: Well-specified model Figures for the source location finding problem under the well-specified model are presented below.



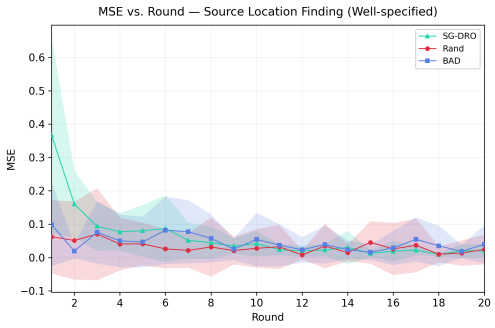
(a) Bandwidth



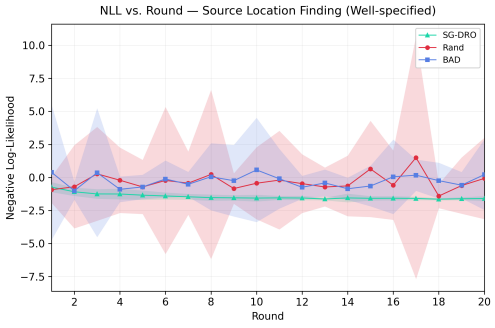
(b) Credible interval



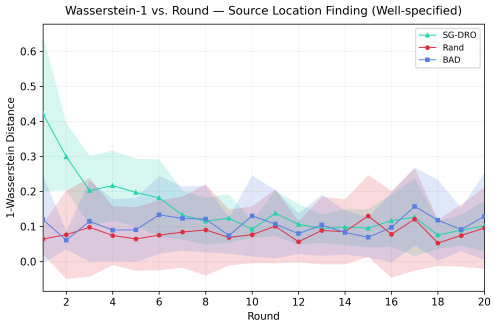
(c) MMD



(d) MSE



(e) NLL



(f) Wasserstein distance

Figure 9: Results for the well-specified source location finding model.

The decreasing trend in the average distance between the functional lower and upper bounds over the course of experiments, shown in Figure 9a, demonstrates the effectiveness of the proposed functional ambiguity set in converging to the true DGP. In line with the proposed theorem 3.1, Figure 9b shows that, on average, the coverage percentage of the true DGP by the constructed ambiguity sets exceeds the specified confidence level of 90% after the second round of experiments.

The lower MMD of SG-DRO in Figure 9c demonstrates the lower discrepancy between the learned model and the true DGP, showing that the selected design points over the experiments are effectively informative regarding the true DGP, rather than being restricted to the information that model parameters can provide about the true DGP. Although none of the methods fully dominate the others, the low value of the MSE in Figure 9d shows that the predictive performance of the learned likeli-

hood model under the proposed SG-DRO algorithm is appropriate and comparable to the baseline methods for the source location finding problem, and the lower NLL in Figure 9e indicates better generalisation performance of the learned model under the well-specified likelihood model.

The lower W1 in Figure 9f shows the closeness of the learned conditional likelihood model to the true DGP. This shows that the SG-DRO explores more to learn the shape of the distribution in the early stages of the experiments. Therefore, in the early stages, it requires more mass to push the learned conditional density toward the true DGP. This search does not have a severe effect on the distributional distortion, but also helps to construct better ambiguity sets under the specified confidence level.

A.8 DR-BED Performance in Higher-Dimensional Settings

To evaluate the effectiveness of the proposed algorithm, we compare DR-BED against GBOED, one of the strong baselines for misspecified BED.

Generalised Bayesian Optimal Experimental Design (GBOED). This method builds on Gibbs inference, which replaces the likelihood in the Bayesian update with a loss function, to make both design selection and inference robust. It introduces *pseudo* versions of information-theoretic quantities and uses them to derive a new acquisition function, the Gibbs expected information gain (Barlas et al., 2025).

Tables 3 and 4 summarise the performance of the considered methods after 20 rounds of experiments on the source location finding problem under different dimensions.

Table 3: 4-dimensional source location finding problem.

	Method	Misspecified likelihood model			
		MMD	**MSE**	**W1**	**NLL**
$d = 4$	GBOED	1.033 (0.070)	1.762 (0.25)	1.157 (0.098)	199.1 (17)
	DR-BED	**0.9404 (0.035)**	**1.059 (0.14)**	**0.885 (0.070)**	**89.31 (12)**

Best result per column are in bold and values in parentheses denote the standard error.

Table 4: 8-dimensional source location finding problem.

	Method	Misspecified likelihood model			
		MMD	**MSE**	**W1**	**NLL**
$d = 8$	GBOED	0.8254 (0.022)	**0.5913 (0.048)**	**0.6199 (0.032)**	303.9 (9.5)
	DR-BED	**0.8197 (0.029)**	0.7774 (0.12)	0.7008 (0.062)	**174.7 (17)**

Best result per column are in bold and values in parentheses denote the standard error.

In the four-dimensional setting, DR-BED outperforms GBOED across all four metrics. The reductions in MMD and W1 indicate that the designs selected by DR-BED lead to an inferred model that is closer to the true DGP, despite the misspecified likelihood. DR-BED also reduces the MSE by roughly 40%, showing more accurate point predictions. The largest improvement is observed in NLL, which DR-BED reduces by more than half, indicating a substantially better predictive distribution.

In the 8-dimensional setting, both methods achieve comparable MMD, with the difference lying well within one standard error. GBOED attains a lower MSE and W1, although the differences in W1 are within the combined standard errors of the two methods. Additionally, DR-BED achieves a lower NLL (174.7 versus 303.9). Since MSE evaluates only point predictions whereas NLL heavily penalises confident predictions that miss the true observations, this combination suggests that GBOED produces point estimates that are slightly more accurate on average but overconfident, whereas DR-BED yields better-calibrated predictive uncertainty. This behaviour is consistent with the distributionally robust formulation of DR-BED, which guards against the overconfidence that model misspecification typically induces.

A.9 Deployment Time Analysis

Table 5 summarises the computational cost of DR-BED relative to the considered baselines. We conduct the comparison under the misspecified scenario of the polynomial regression problem over 20 experimental rounds. The table shows the average time each method requires per round to select the corresponding design.

Table 5: Deployment time of different methods.

Method	Deployment Time(s)
Rand	0.01 (0.0001)
BAD	1.02 (0.0008)
GBOED	6.06 (0.0245)
DR-BED	23.89 (0.5671)

Values in parentheses denote the standard error.

The results show a trade-off between robustness and computational efficiency. The random design algorithm has negligible overhead because it requires no optimisation. BAD and GBOED take about 1 and 6 seconds per round, respectively. DR-BED is the most expensive, at roughly 24 seconds per round, because it optimizes the worst case over an ambiguity set rather than a single nominal model, which adds an inner subgradient descent loop to each design step. Even so, its per-round cost remains under one minute, which is reasonable in practical BED settings, where running a physical experiment typically takes far longer than selecting the design. Given the gains DR-BED achieves under model misspecification, this additional cost is an acceptable price for improved robustness. Developing an amortised version of DR-BED to further improve its computational efficiency is left for future research.