entropy

MDPI

Article

Decision-Making Under Model Misspecification: DRO with Robust Bayesian Ambiguity Sets

Charita Dellaporta [1,*,†,‡], Patrick O'Hara [2,‡] and Theodoros Damoulas [2,3]

[1] Department of Statistical Science, University College London, London WC1E 6BT, UK
[2] Department of Computer Science, University of Warwick, Coventry CV4 7AL, UK; t.damoulas@warwick.ac.uk (T.D.)
[3] Department of Statistics, University of Warwick, Coventry CV4 7AL, UK
* Correspondence: h.dellaporta@ucl.ac.uk
† Part of this work was undertaken while CD was affiliated with 3.
‡ These authors contributed equally to this work.

Abstract

Distributionally Robust Optimisation (DRO) protects risk-averse decision-makers by considering the worst-case risk within an ambiguity set of distributions based on the empirical distribution or a model. To further guard against finite, noisy data, model-based approaches admit Bayesian formulations that propagate uncertainty from the posterior to the decision-making problem. However, when the model is misspecified, the decision-maker must stretch the ambiguity set to contain the data-generating process (DGP), leading to overly conservative decisions. We address this challenge by introducing DRO with Robust ayesian Ambiguity Sets (DRO-RoBAS) to model misspecification. These are Maximum Mean Discrepancy ambiguity sets centred at a robust posterior predictive distribution that incorporates beliefs about the DGP. We show that the resulting optimisation problem obtains a dual formulation in the Reproducing Kernel Hilbert Space and we give probabilistic guarantees on the tolerance level of the ambiguity set. Our method outperforms other Bayesian and empirical DRO approaches in out-of-sample performance on the Newsvendor and Portfolio problems with various cases of model misspecification.

Keywords: robustness; Bayesian inference; stochastic optimisation; misspecification; divergence-based inference

Check for updates

Academic Editor: Michel Broniatowski

Received: 24 February 2026
Revised: 31 March 2026
Accepted: 9 April 2026
Published: 11 April 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Decision-makers frequently encounter the challenge of optimising under uncertainty since the data-generating process (DGP) is not fully known. As a result, they must rely on available data and model families to estimate the DGP within the optimisation objective. However, the data may be noisy, the data distribution might change with time, or the model might be misspecified or poorly fitted, leading to distributional uncertainty. Decision-making in this setting is a critical challenge in various applications such as inventory planning [1], portfolio optimisation [2] and distribution shifts in machine learning applications [3].

A risk-averse decision-maker might choose to hedge against distributional uncertainty by finding the decision that minimises the *worst-case* risk over a set of distributions. This worst-case protection is at the heart of *Distributionally Robust Optimisation* (*DRO*) that defines an ambiguity set of distributions with respect to an estimator of the DGP. This can be fully

data-driven, using the empirical measure of the observations (see, e.g., [4–7]), or model-based when expert knowledge is available to fit a model to available data, resulting in a model-based estimator for the DGP [8–10]. Both approaches are sensitive to the choice of DGP estimator and additional estimation error might exist in model-based DRO due to a poor fit or model uncertainty.

To overcome this, recently developed Bayesian formulations of DRO use posterior beliefs to inform the optimisation problem [11] or the ambiguity set itself [12]. However, these methods inherit the sensitivity of Bayesian posteriors to model misspecification (see, e.g., [13,14]). A key goal in DRO methodology is to choose the size of the ambiguity set such that the DGP falls within it with high probability, as illustrated in Figure 1 (left) for the two formulations of Bayesian Ambiguity Sets (BAS) [12] and our proposed Robust BAS (RoBAS). If the estimate is not accurate—for example, when the model is misspecified—then a much larger size will be required to contain the DGP as illustrated in Figure 1 (right). The price to pay for this large size is the inclusion of many probability distributions that are unlikely to occur, and which could be very pessimistic with respect to the objective function, leading to an overly conservative decision. If the decision-maker wrongly assumes the model is well specified and incorrectly chooses an overly optimistic ambiguity set size, then the DGP may not lie in the set, and the decision could be overly optimistic compared to the out-of-sample outcome, often referred to as the optimiser's curse [4].

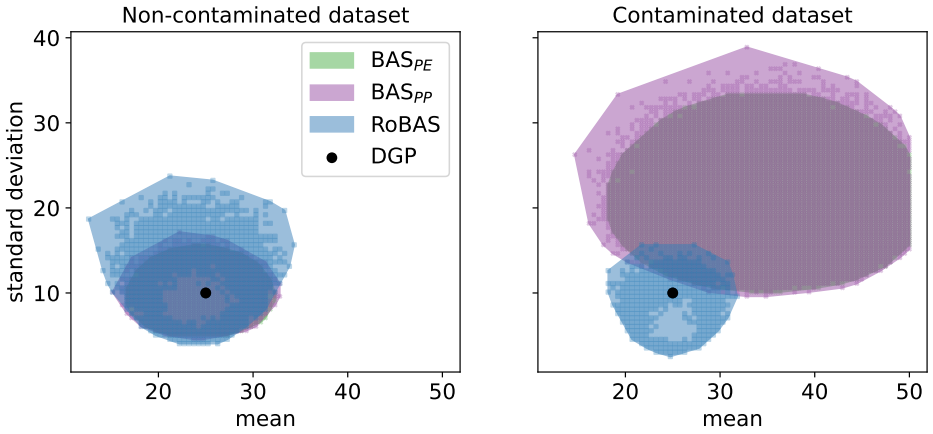

Figure 1. Illustration of (approximated) BAS_{PE}, BAS_{PP} and RoBAS (ours) with $\mathbb{P}_\theta = \mathcal{N}(\mu, \sigma^2)$ over a grid of (μ, σ) pairs for a fixed ϵ. In the well-specified case (**left**), all ambiguity sets include the DGP while RoBAS covers a slightly bigger area than BAS_{PE} and BAS_{PP}. For a contaminated dataset (**right**), RoBAS continues to contain the DGP and maintains a similar area, whereas the BAS formulations exclude it and cover a much larger area of distributions further away from the DGP.

Uncertainty over the DGP is extensively studied outside stochastic optimisation. Consider a parametric model $\mathbb{P}_\theta$, indexed by the parameter of interest θ. In the Bayesian framework, uncertainty about the parameter is typically expressed directly through prior beliefs. However, recent work in robust Bayesian inference by Lyddon et al. [15] takes a different approach: uncertainty about the parameter is now induced by uncertainty, in the form of prior beliefs, in the DGP. This concept lies at the core of the Bayesian Nonparametric Learning (NPL) framework [15,16] which relaxes the well-specified model assumption imposed by standard Bayesian inference. In this spirit, we approach the challenge of DRO under model misspecification by extending the recently proposed DRO-BAS [12] framework to tackle model misspecification through a robust NPL posterior coupled with the Maximum Mean Discrepancy (MMD) inside the ambiguity set, thus introducing DRO with Robust Bayesian Ambiguity Sets (DRO-RoBAS). While DRO-BAS targets distributional

uncertainty with respect to the DGP, the robustness offered is not sufficient under model misspecification, see Figure 1.

2. Background

Let $x \in \mathcal{X} \subseteq \mathbb{R}^d$ be a decision-making variable for the cost function $f : \mathcal{X} \times \Xi \to \mathbb{R}$ with data space $\Xi \subseteq \mathbb{R}^D$. Let $\{\xi_i\}_{i=1}^n \overset{\text{iid}}{\sim} \mathbb{P}^\star$ be observations from the DGP $\mathbb{P}^\star \in \mathcal{P}(\Xi)$, where $\mathcal{P}(\Xi)$ denotes the space of Borel probability measures on Ξ. Furthermore, consider a parametric model family $\mathcal{P}_\Theta := \{\mathbb{P}_\theta : \theta \in \Theta\} \subset \mathcal{P}(\Xi)$ indexed by parameter of interest $\theta \in \Theta \subseteq \mathbb{R}^k$. We say the model is *misspecified* if $\mathbb{P}^\star \notin \mathcal{P}_\Theta$. DRO methods construct an ambiguity set $\mathcal{A}$, based on an estimator of $\mathbb{P}^\star$, called the *nominal distribution*, and minimise the worst-case expected cost over $\mathcal{A}$.

The DRO literature typically categorises ambiguity sets into two classes [17]: *moment-based* and *discrepancy-based*. The former contain distributions that satisfy constraints on the moments of $\mathbb{P}^\star$, without necessarily considering an estimator. In contrast, discrepancy-based ambiguity sets consist of distributions close to the nominal according to a specified discrepancy measure. Examples include Integral Probability Metrics (IPMs) [18], such as the Wasserstein distance [4] and the MMD [5], as well as ϕ-divergences like the Kullback–Leibler (KL) divergence [7]. Regardless of the choice of ambiguity set, the resulting minimax problem can be seen as a game between the decision-maker who chooses x and an adversary who chooses the worst-case distribution in $\mathcal{A}$:

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{A}} \mathbb{E}_{\xi \sim \mathbb{P}}[f_x(\xi)] \tag{1}$$

where $f_x(\xi) := f(x, \xi)$. Although most discrepancy-based DRO methods are fully empirical, i.e., the estimator is obtained only via $\xi_{1:n}$, sometimes, such as in regression settings, the decision-maker needs to model the variables' relationship via a model family $\mathcal{P}_\Theta$, describing the DGP. Model-based DRO methods (e.g., [8–10]) use the observations to obtain an estimator $\mathbb{P}_{\hat{\theta}} \in \mathcal{P}_\Theta$ and use this as the nominal distribution. Thus, a poorly chosen $\mathbb{P}_{\hat{\theta}}$ far from $\mathbb{P}^\star$ (in some distance sense) requires a large $\mathcal{A}$, leading to overly pessimistic decisions. This has led to Bayesian formulations of DRO which propagate uncertainty about θ in the optimisation problem.

2.1. Bayesian Formulations of DRO

Shapiro et al. [11] introduced Bayesian DRO (BDRO) which defines an *expected worst-case* risk objective:

$$\min_{x \in \mathcal{X}} \mathbb{E}_{\Pi(\theta|\xi_{1:n})} \left[\sup_{\mathbb{P} \,:\, d_{\mathrm{KL}}(\mathbb{P}||\mathbb{P}_\theta) \leq \epsilon} \mathbb{E}_{\xi \sim \mathbb{P}}[f_x(\xi)] \right] \tag{2}$$

where $\Pi(\theta \mid \xi_{1:n})$ denotes the parameter posterior distribution for model family $\mathcal{P}_\Theta$. However, risk-averse decision-makers are interested in worst-case risk formulations. For this reason, Dellaporta et al. [12] proposed two formulations of the DRO with Bayesian Ambiguity Sets (DRO-BAS) that correspond to a *worst-case* optimisation problem with ambiguity sets informed by the standard Bayesian posterior. In particular, they defined DRO-BAS$_{\mathrm{PP}}$:

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{P} \,:\, d_{\mathrm{KL}}\left(\mathbb{P}||\mathbb{E}_{\Pi(\theta|\xi_{1:n})}[\mathbb{P}_\theta]\right) \leq \epsilon} \mathbb{E}_{\xi \sim \mathbb{P}}[f_x(\xi)], \tag{3}$$

based on a KL-based ambiguity set with nominal distribution the posterior predictive, and DRO-BAS$_{\text{PE}}$:

$$\min_{x\in\mathcal{X}} \sup_{\mathbb{P}\,:\,\mathbb{E}_{\Pi(\theta|\xi_{1:n})}[d_{\text{KL}}(\mathbb{P}||\mathbb{P}_\theta)]\leq\epsilon} \mathbb{E}_{\xi\sim\mathbb{P}}[f_x(\xi)]. \tag{4}$$

which considers the expected KL under the posterior distribution. The authors showcased improved out-of-sample robustness compared to BDRO in a number of Exponential family models. Although DRO-BAS in the standard Bayesian setting offers an intuitive, posterior-informed ambiguity set, it can be severely affected by model misspecification. Indeed, BAS$_{\text{PE}}$ only considers probability measures $\mathbb{P}$ that are absolutely continuous with respect to $\mathbb{P}_\theta$ (denoted by $\mathbb{P} \ll \mathbb{P}_\theta$) and also admit an expected KL divergence close enough to $\mathbb{P}_\theta$. Since the expectation is informed by the posterior, a non-robust posterior will likely lie far away from the DGP. Similarly, BAS$_{\text{PP}}$ considers probability measures $\mathbb{P}$ that admit small KL-divergence with respect to $\mathbb{P}_n^{\text{pred}}$, where $\mathbb{P}_n^{\text{pred}} := \mathbb{E}_{\theta\sim\Pi(\theta|\xi_{1:n})}[\mathbb{P}_\theta]$ and $\mathbb{P} \ll \mathbb{P}_n^{\text{pred}}$. Hence, the sensitivity of the Bayesian posterior will propagate to the posterior predictive and the resulting ambiguity set. A similar observation was made by Shapiro et al. [11] for the BDRO method in the misspecified case.

To remedy this, we exploit the flexibility of the DRO-BAS framework which allows us to choose a different posterior distribution and discrepancy measure, suitable for model misspecification. The notion of targeting a different discrepancy measure, other than the KL divergence, to induce robustness in the Bayesian posterior has been well established in the Bayesian inference literature. The NPL posterior [15,16] was introduced to, among others, remedy the sensitivity of Bayesian inference to model misspecification by removing the assumption that the model is correct. This is done by setting uncertainty, via nonparametric prior beliefs, directly on the DGP rather than on the parameter of interest. Incorporating DGP uncertainty in decision-making has also been considered by Wang et al. [19], who explored a nonparametric Dirichlet Process (DP) model for the DGP. Unlike the current paper, this work is not suited for parametric models and considers a weighted objective, with only one counterpart corresponding to a worst-case risk. We focus on decision-making under *parametric* models, which are especially useful for interpretability in decision-making, while also incorporating *nonparametric* prior beliefs about the DGP.

To achieve this, we leverage the work of Dellaporta et al. [20] who extended the NPL posterior to discrepancy-based loss functions and showed robustness guarantees when the Maximum Mean Discrepancy (MMD) is used. Using an MMD-based loss allows us to also employ the MMD to construct the ambiguity set. NPL is a natural choice for the DRO-BAS framework as distributional uncertainty in DRO stems directly from uncertainty in the DGP. Before we introduce the NPL posterior and the DRO-RoBAS framework, we first give a brief overview of robust Bayesian inference methodologies based on divergencies.

2.2. Robust Bayesian Inference via Divergences

To mitigate the lack of robustness in standard Bayesian inference, a class of approaches known as Generalised Bayesian Inference (GBI) [21–23] has been introduced. In this framework, one replaces the log-likelihood with a general empirical loss function $l_n : \Xi^n \times \Theta \to \mathbb{R}$, together with a learning rate $\beta > 0$. The resulting posterior distribution for prior $\pi(\theta)$ is defined through the density

$$\Pi_{\text{GBI}}(\theta|\xi_{1:n}) \propto \exp(-\beta l_n(\xi_{1:n},\theta))\pi(\theta).$$

In most cases, the loss function l_n is chosen in relation to the model family $\mathcal{P}_\Theta$. In particular, the choice of $l_n(\xi_{1:n},\theta) := -\log p(\xi_{1:n},\theta)$ and $\beta = 1$ recovers the standard Bayesian

posterior. This observation highlights that the sensitivity of the standard Bayesian posterior to model misspecification stems from this specific choice of loss. To improve robustness, alternative generalised posteriors have been proposed by selecting discrepancies D with desirable properties and select l_n induced by an approximation of $D(\mathbb{P}^\star, \mathbb{P}_\theta)$ using observed data. Examples of such works based on robust divergences or discrepancies include [24–27]. Beyond GBI approaches, Bayesian NPL provides an alternative generalised formulation of Bayesian inference, which in some cases is also grounded in discrepancy-based loss functions. Unlike GBI methods, NPL places uncertainty directly on the DGP, thereby enhancing robustness to model misspecification. As distributional robustness similarly seeks to account for uncertainty in the DGP, we adopt this framework in the present work and introduce it below.

Robust NPL Posterior

In this work, we propose an alternative formulation of DRO-BAS, based on the robust NPL posterior introduced by Lyddon et al. [15] and Fong et al. [16]. We introduce a DP prior $\mathbb{Q} \sim \mathrm{DP}(\alpha, \mathbb{F})$ on the DGP $\mathbb{P}^\star$ where $\alpha > 0$ and $\mathbb{F} \in \mathcal{P}(\Xi)$. Here, $\mathbb{F}$ represents our prior beliefs about the DGP and the concentration parameter α dictates the strength of the beliefs with $\alpha = 0$ representing a non-informative prior. To see this, note that given data $\xi_{1:n}$, the posterior is

$$\begin{aligned} \mathbb{Q} \mid \xi_{1:n} &\sim \mathrm{DP}(\alpha', \mathbb{F}'), \\ \alpha' := \alpha + n, \quad \mathbb{F}' &:= \frac{\alpha}{\alpha + n}\mathbb{F} + \frac{n}{\alpha + n}\mathbb{P}_n \end{aligned} \tag{5}$$

where $\mathbb{P}_n := \frac{1}{n}\sum_{i=1}^n \delta_{\xi_i}$ is the empirical measure and δ_ξ denotes the Dirac measure at $\xi \in \Xi$. For $\alpha = 0$, the DP posterior is centred directly on $\mathbb{P}_n$. If $\mathbb{P}^\star$ was known, we could directly compute:

$$\theta_L(\mathbb{P}^\star) := \arg\min_{\theta \in \Theta} \mathbb{E}_{\xi \sim \mathbb{P}^\star}[L(\xi; \theta)] \tag{6}$$

where $L : \Xi \times \Theta \to \mathbb{R}$ denotes any loss function. Note that this NPL objective does not assume that the model is well specified but simply looks for the most likely value of θ under the expectation of the DGP or, equivalently, the parameter value that best describes the data under the candidate model. Since $\mathbb{P}^\star$ is unknown and we instead have a nonparametric posterior over it, we can propagate our posterior beliefs to the parameter of interest through the push-forward measure $(\theta_L)_\#(\mathrm{DP}(\alpha', \mathbb{F}'))$ to give a posterior Π_{NPL} on Θ. Sampling from this posterior can be done through the Posterior Bootstrap [16]: For B Posterior Bootstrap iterations, at iteration $j \in [B]$:

1. Sample $\mathbb{Q}^{(j)}$ from the posterior $\mathrm{DP}(\alpha', \mathbb{F}')$.
2. Compute $\theta^{(j)} = \theta_L(\mathbb{Q}^{(j)})$ where $\theta_L(\cdot)$ as in (6).

Dellaporta et al. [20] suggested using a discrepancy-based loss function in (6) which we introduce below.

2.3. Maximum Mean Discrepancy

The MMD belongs to the family of IPMs [28]. Let $\mathcal{H}_k$ be a Reproducing Kernel Hilbert Space (RKHS), for kernel $k : \Xi \times \Xi \to \mathbb{R}$ and norm $\|\cdot\|_k$. For $\mathcal{P}_k(\Xi) := \{\mathbb{P} \in \mathcal{P}(\Xi) : \int_\Xi \sqrt{k(\xi, \xi)}\mathbb{P}(d\xi) < \infty\}$, the MMD between $\mathbb{P}, \mathbb{Q} \in \mathcal{P}_k(\Xi)$ is defined as:

$$\mathbb{D}_k(\mathbb{P}, \mathbb{Q}) := \sup_{f \in \mathcal{H}_k, \|f\|_k \leq 1} \left|\mathbb{E}_{\mathbb{P}}[f(\xi)] - \mathbb{E}_{\mathbb{Q}}[f(\xi)]\right|. \tag{7}$$

Dellaporta et al. [20] define the NPL target in (6) as:

$$\theta_k(\mathbb{P}^\star) := \arg\min_{\theta \in \mathbb{R}^k} \mathbb{D}_k(\mathbb{P}^\star, \mathbb{P}_\theta). \tag{8}$$

One of the attractive properties of the MMD is that the supremum in (7) can be obtained in closed form as

$$\begin{aligned} \mathbb{D}_k^2(\mathbb{P}, \mathbb{Q}) = \mathbb{E}_{\xi,\xi' \sim \mathbb{P}}[k(\xi, \xi')] - 2\mathbb{E}_{\xi \sim \mathbb{P}, \xi' \sim \mathbb{Q}}[k(\xi, \xi')] \\ + \mathbb{E}_{\xi,\xi' \sim \mathbb{Q}}[k(\xi, \xi')] \end{aligned} \tag{9}$$

and can be approximated via sampling (see, e.g., [29]). The resulting NPL posterior with the MMD is called an NPL-MMD posterior.

We explore the gains of using the MMD both in the robust NPL posterior and in the Bayesian ambiguity set. Since the NPL-MMD posterior will target the point in the model family closest (w.r.t. the MMD) to $\mathbb{P}^\star$, the MMD-based ambiguity set may require a smaller radius to include $\mathbb{P}^\star$, resulting in less conservative decisions.

2.4. DRO with the Maximum Mean Discrepancy

The MMD has previously been used as a distance metric in the DRO context by Staib and Jegelka [5] who considered the ambiguity set

$$\mathcal{B}_\epsilon^k(\mathbb{P}_n) := \{\mathbb{P} \in \mathcal{P}_k(\Xi) : \mathbb{D}_k(\mathbb{Q}, \mathbb{P}_n) \leq \epsilon\}$$

where $\epsilon > 0$ is the radius of the MMD ball, k is the kernel and $\mathbb{P}_n := \frac{1}{n}\sum_{i=1}^n \delta_{\xi_i}$ is the empirical measure of the observations $\xi_{1:n}$. As highlighted by the authors, this ambiguity set has several advantages. First, MMD-DRO can be readily applied to complex data structures such as images or graphs by choosing an appropriate kernel defined on the corresponding data space. This is in contrast to DRO with other choices of distance metrics like the popular Wasserstein distance, in which case, a lot of theoretical and optimisation results rely on specific choices of ground metric, which limit their applicability to more complex data. Moreover, existing finite sample conctration results (see, e.g., [29,30]) can be used for radius selection. Importantly, these results do not suffer from the curse of dimensionality and do not require any assumptions on the DGP, contrary to other distance metrics like the Wasserstein [4]. Although several works have provided remedies to this curse [31], results usually require assumptions on the DGP or the loss function and are hence not always applicable. However, these advantages come at a cost: Ref. [5] explored the optimisation of the MMD-DRO objective by deriving an upper bound which, however, requires for the loss function to be a member of the RKHS, which is often hard to verify. To remedy this, Zhu et al. [6] studied the same ambiguity set along with an extended family of kernel-based ambiguity sets, collectively called Kernel DRO, and provided strong duality results of the optimisation problem which do not require the loss function to be a member of the corresponding RKHS. In this work, we adopt this dual formulation and give a detailed explanation in Section 3.1. Chen et al. [32] generalised this to conditional Kernel DRO, leveraging conditional distributions and Romao et al. [33] explored this framework for dynamic programming.

3. DRO with Robust Bayesian Ambiguity Sets

We propose a robust version of DRO-BAS$_{\text{PP}}$ (3) via the MMD and the NPL-MMD posterior predictive defined as:

$$\mathbb{P}_n^{\text{pred(NPL)}} := \mathbb{E}_{\mathbb{Q} \sim \text{DP}(c', \mathbb{F}')}[\mathbb{P}_{\theta_k(\mathbb{Q})}]. \tag{10}$$

We assume that for $\mathbb{Q} \in \mathcal{P}(\Xi)$, the map $\mathbb{Q} \mapsto \theta_k(\mathbb{Q})$ is measurable such that $\theta_k(\mathbb{Q}) \in \arg\min_{\theta\in\Theta} D_k(\mathbb{Q}, \mathbb{P}_\theta)$, so that the expectation in (10) is well defined. Throughout, we also assume that $P_n^{\text{pred(NPL)}} \in \mathcal{P}_k(\Xi)$, which holds, for example, when the kernel k is bounded. A sufficient condition ensuring the existence of the minimiser in (8) is given in Section 3.2. Notice that contrary to the standard Bayesian posterior predictive, $\mathbb{P}_n^{\text{pred(NPL)}}$ is defined through marginalisation over the nonparametric posterior over the DGP and is defined for any choice of nonparametric prior $\text{DP}(\alpha, \mathbb{F})$ as defined in Section 2.1. Since the MMD can be approximated only via samples (Section 2.3), a closed-form density for the predictive in (10) is not required. We define the following Robust Bayesian Ambiguity Set (RoBAS) with the NPL posterior predictive:

$$\mathcal{B}_\epsilon^k(\mathbb{P}_n^{\text{pred(NPL)}}) := \{\mathbb{P} \in \mathcal{P}_k(\Xi) : \mathbb{D}_k(\mathbb{P}, \mathbb{P}_n^{\text{pred(NPL)}}) \leq \epsilon\}.$$

Note that $\mathcal{B}_\epsilon^k(\mathbb{P}_n^{\text{pred(NPL)}})$ forms a ball when k is a characteristic kernel, as this makes the MMD a probability metric. This property is desirable as it guarantees that the MMD will be zero if and only if $\mathbb{P} \equiv \mathbb{P}_n^{\text{pred(NPL)}}$. We hence obtain the following DRO-RoBAS worst-case risk problem:

$$\min_{x\in\mathcal{X}} \sup_{\mathbb{P} \in \mathcal{B}_\epsilon^k(\mathbb{P}_n^{\text{pred(NPL)}})} \mathbb{E}_{\xi\sim\mathbb{P}}[f_x(\xi)] \tag{11}$$

Similarly to DRO-BAS, this optimisation problem corresponds to a worst-case risk over a set of probability measures informed by posterior beliefs. Posterior beliefs about θ are obtained *via posterior beliefs about the DGP* and the map θ_k in (8). Since the goal of DRO is to target uncertainty about the DGP, the NPL posterior is a natural choice to inform the ambiguity set as it takes into account any prior beliefs about the DGP. Moreover, by targeting the MMD, rather than the KL divergence as in the BAS case, RoBAS is *not* restricted to probability measures that are absolutely continuous with respect to $\mathbb{P}_\theta$.

Intuitively, RoBAS is expected to be a better-informed ambiguity set than BAS when the model is misspecified since it is informed by a robust posterior predictive and a robust discrepancy measure. This is better understood through a toy example. Figure 2 shows a Gaussian location model in the presence of outliers. The top panel shows the DGP of the training data contaminated with 20% of outliers along with a pathological model $\mathbb{P}_{\text{pathological}}$ case with a mean larger than that of the DGP. In the BAS_{PE} case, the expected KL from the DGP to the model is significantly larger than that from the pathological model $\mathbb{P}_{\text{pathological}}$ due to the sensitivity of the Bayesian posterior to outliers. A similar result holds for the KL divergence between the posterior predictive and the DGP and pathological model in the BAS_{PP} case. In contrast, the MMD from the NPL posterior predictive to the DGP is much smaller compared to that of the pathological model. Alternative DRO-RoBAS formulations are provided in Appendix C. The results of this section are in Appendix A.

3.1. Duality of the DRO-RoBAS Problem

We first formulate our optimisation problem as a kernel DRO problem [6]. This allows us to obtain a dual formulation of (11) in the RKHS which can be optimised using kernel methods. Let $\phi : \Xi \to \mathcal{H}_k$ denote the feature map associated with kernel k and let $\mu_\mathbb{P} \in \mathcal{H}_k$ denote the kernel mean embedding of the probability measure $\mathbb{P} \in \mathcal{P}_k(\Xi)$, i.e., $\mu_\mathbb{P} := \mathbb{E}_{\xi\sim\mathbb{P}}[\phi(\xi)]$. Then the MMD is equivalently defined as $\mathbb{D}_k(\mathbb{P}, \mathbb{Q}) = \|\mu_\mathbb{P} - \mu_\mathbb{Q}\|_k$ ([29], Lemma 4). Throughout we assume that $\mathbb{P}_n^{\text{pred(NPL)}} \in \mathcal{P}_k(\Xi)$. Note that this condition is trivially satisfied for a bounded kernel k. Consider the following set satisfying the conditions of RoBAS:

$$\mathcal{C}^\star := \{\mu \in \mathcal{H}_k : \|\mu - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_k \leq \epsilon\}. \tag{12}$$

The associated ambiguity set induced by $\mathcal{C}^\star$ is

$$\mathcal{K}_{\mathcal{C}^\star} := \{\mathbb{P} \in \mathcal{P}_k(\Xi), \mu_{\mathbb{P}} \in \mathcal{C}^\star\} \equiv \mathcal{B}_\epsilon^k(\mathbb{P}_n^{\text{pred(NPL)}}) \tag{13}$$

and our DRO-RoBAS problem in (11) is equivalent to:

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{P}, \mu_{\mathbb{P}}} \mathbb{E}_{\xi \sim \mathbb{P}}[f_x(\xi)] \text{s.t.} \mathbb{P} \in \mathcal{P}_k(\Xi), \mu_{\mathbb{P}} \in \mathcal{C}^\star. \tag{14}$$

The equivalence can be seen through the $\mathcal{C}^\star$-induced ambiguity set for distributions $\mathbb{P}$ which can be written as $\mathcal{K}_{\mathcal{C}^\star}$.

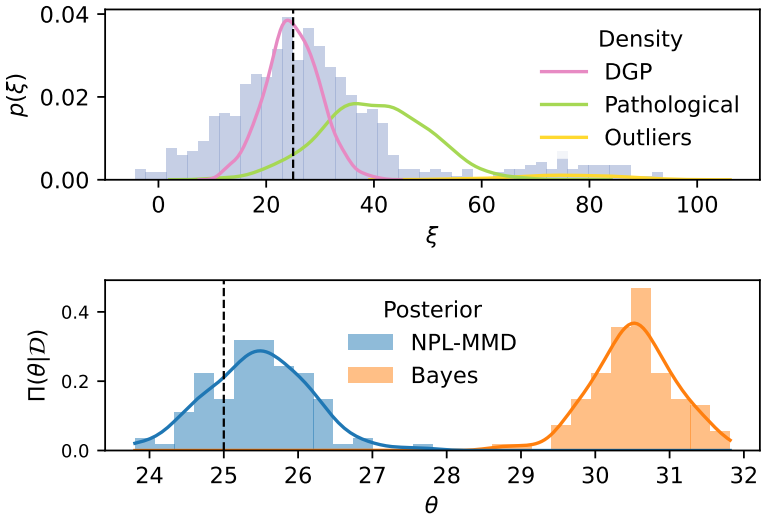

Figure 2. Contaminated Gaussian location example. (**Top**): Histogram of observed data along with the true (DGP), outlier and pathological densities. (**Bottom**): Posterior marginal distributions for NPL-MMD (blue) and standard Bayes (orange). The true mean is indicated with a black dotted line. For BAS_{PE} it holds that $\mathbb{E}_{\Pi_{\text{Bayes}}}[d_{\text{KL}}(\mathbb{P}_{\text{pathological}}, \mathbb{P}_\theta)] \approx 0.17 < 0.42 \approx \mathbb{E}_{\Pi_{\text{Bayes}}}[d_{\text{KL}}(\mathbb{P}^\star, \mathbb{P}_\theta)]$ and similarly for BAS_{PP} it holds that $d_{\text{KL}}(\mathbb{P}_{\text{pathological}}, \mathbb{P}_n^{\text{pred}}) \approx 0.18 < 0.38 \approx d_{\text{KL}}(\mathbb{P}^\star, \mathbb{P}_n^{\text{pred}})$. In contrast for RoBAS we have $\mathbb{D}_k(\mathbb{P}_{\text{pathological}}, \mathbb{P}_n^{\text{pred(NPL)}}) \approx 0.65 > 0.55 \approx \mathbb{D}_k(\mathbb{P}^\star, \mathbb{P}_n^{\text{pred(NPL)}})$. This example is inspired by Figure 1 of Gao and Kleywegt [34].

Before we proceed to our dual formulation, we introduce the main result from Zhu et al. [6] which gives a dual formulation of Kernel DRO problems for general sets $\mathcal{C}$ satisfying certain assumptions.

Theorem 1 (Zhu et al. [6], Theorem 3.1)**.** *Assume $\mathcal{C} \subset \mathcal{H}_k$ is closed convex, $f_x(\cdot)$ is proper, upper semi-continuous, and $\text{ri}(\mathcal{K}_{\mathcal{C}}) \neq \emptyset$, where $\text{ri}(\mathcal{K}_{\mathcal{C}})$ denotes the relative interior of $\mathcal{K}_{\mathcal{C}}$. Then the primal problem:*

$$\min_x \sup_{\mathbb{P}, \mu} \mathbb{E}_{\xi \sim \mathbb{P}}[f_x(\xi)] \quad s.t. \quad \mathbb{P} \in \mathcal{P}, \mu_{\mathbb{P}} = \mu, \ \mu \in \mathcal{C}$$

is equivalent to:

$$\min_{x, g_0 \in \mathbb{R}, g \in \mathcal{H}_k} g_0 + \delta_{\mathcal{C}}^\star(g) \quad s.t. \quad f_x(\xi) \leq g_0 + g(\xi), \forall \xi \in \Xi$$

where $\delta^\star_{\mathcal{C}}(g) := \sup_{\mu\in\mathcal{C}}\langle g,\mu\rangle_{\mathcal{H}_k}$ *the support function of* $\mathcal{C}$.

This theorem gives an effective way to transition from the primal to the dual formulation by using the support function of the set $\mathcal{C}$. Importantly, in contrast to other kernel-based dual formulations (e.g., [5]), this theorem does *not* require the objective function f to be a member of the RKHS $\mathcal{H}_k$. We first derive the support function of $\mathcal{C}^\star$ in the DRO-RoBAS case. We denote $\mathbb{E}_{\mathbb{Q}\sim \mathrm{DP}(\alpha',\mathbb{F}')}$ and $\mathbb{E}_{\xi\sim\mathbb{P}_{\theta_k(\mathbb{Q})}}$ by $\mathbb{E}_{\mathrm{DP}_{\xi_{1:n}}}$ and $\mathbb{E}_{\mathbb{P}_{\theta_k(\mathbb{Q})}}$ respectively.

Proposition 1. *Let* $\mathcal{C}^\star$ *be defined by (12). Then we have* $\delta^\star_{\mathcal{C}^\star}(g) = \mathbb{E}_{\mathrm{DP}_{\xi_{1:n}}}\left[\mathbb{E}_{\mathbb{P}_{\theta_k(\mathbb{Q})}}[g(\xi)]\right] + \epsilon\|g\|_k$.

We can now apply Theorem 1 to our problem.

Corollary 1. *Let* $\mathcal{C}^\star$ *as in (12) and* $f_x(\cdot)$ *proper, upper semi-continuous. Then problem (14) is equivalent to:*

$$\begin{aligned}\min_{x,g_0\in\mathbb{R},g\in\mathcal{H}_k}\quad & g_0 + \mathbb{E}_{\mathrm{DP}_{\xi_{1:n}}}\left[\mathbb{E}_{\mathbb{P}_{\theta_k(\mathbb{Q})}}[g(\xi)]\right] + \epsilon\|g\|_k \\ \text{subject to}\quad & f_x(\xi)\le g_0 + g(\xi), \forall\xi\in\Xi.\end{aligned} \tag{15}$$

Computation of (15):

The problem in (15) can be solved by the batch approach with discretization of a semi-infinite programme (SIP) [35] suggested in Zhu et al. [6], in addition to a Sample Average Approximation (SAA). Let $\{\hat{\xi}_i\}_{i=1}^N$ be samples from the nested expectation in (15) and $\{\zeta_j\}_{j=1}^m \subseteq \Xi^m$ be a set of discretisation points. Then the problem can be approximated by:

$$\begin{aligned}\min_{x,g_0\in\mathbb{R},g\in\mathcal{H}_k}\quad & g_0 + \frac{1}{N}\sum_{i=1}^N g(\hat{\xi}_i) + \epsilon\|g\|_k \\ \text{s.t.}\quad & f_x(\zeta_j)\le g_0 + g(\zeta_j), \forall j\in[m].\end{aligned} \tag{16}$$

We can now apply the distributional robust version of the Representer theorem ([6], Lemma B.1) which states that it is sufficient to parametrise g by $g(\cdot) = \sum_{i=1}^N \alpha_i k(\hat{\xi}_i,\cdot) + \sum_{j=1}^m \alpha_{N+j}k(\zeta_j,\cdot)$ for some $\alpha_k\in\mathbb{R}$, for all $k = 1,\ldots,N+m$.

3.2. Tolerance Level Guarantees

We start by using the generalisation error results for the NLP-MMD posterior to obtain a result in probability that the DGP lies within our ambiguity set. First, we give a concentration type bound for $\mathbb{E}_{\mathbb{Q}\sim\mathrm{DP}(\alpha',\mathbb{F}')}[\mathbb{D}_k(\mathbb{P}^\star,\mathbb{Q})]$. In practice, exact sampling from a DP is not possible, so we consider the approximation of the DP suggested in the NPL literature [15,16,20] to sample during the MMD Posterior Bootstrap. In particular, denote by $\widehat{\mathrm{DP}}_{\xi_{1:n}}$ the probability measure on $\mathcal{P}(\Xi)$ induced by the following sampling process for $(w_{1:n},\tilde{w}_{1:\tau}) \sim \mathrm{Dir}\left(1,\ldots,1,\frac{\alpha}{\tau},\ldots,\frac{\alpha}{\tau}\right)$ and $\tilde{\xi}_{1:\tau}\overset{\mathrm{iid}}{\sim}\mathbb{F}$:

$$\mathbb{Q} := \sum_{i=1}^n w_i\delta_{\xi_i} + \sum_{k=1}^\tau \tilde{w}_k\delta_{\tilde{\xi}_k} \sim \widehat{\mathrm{DP}}_{\xi_{1:n}}. \tag{17}$$

The associated approximate posterior predictive is $\hat{\mathbb{P}}_n^{\mathrm{pred(NPL)}} := \mathbb{E}_{\mathbb{Q}\sim\widehat{\mathrm{DP}}_{\xi_{1:n}}}[\mathbb{P}_{\theta_k(\mathbb{Q})}]$. We provide a concentration result for the MMD between $\mathbb{P}^\star$ and $\hat{\mathbb{P}}_n^{\mathrm{pred(NPL)}}$ as the approximated predictive is used in practice. However, similar results can be derived for the exact case with $\mathbb{P}_n^{\mathrm{pred(NPL)}}$ following the arguments in [20]. Additionally, all theoretical results

regarding the duality of the DRO-RoBAS framework from Section 3.1 hold exactly the same for the approximated DP as they are proven for a general posterior. We make the following assumptions:

Assumption 1. *For every* $\mathbb{Q} \in \mathcal{P}(\Xi)$ *there exists* $c > 0$ *such that the set* $\{\theta \in \Theta : \mathbb{D}_k(\mathbb{Q}, \mathbb{P}_\theta) \leq \inf_{\theta \in \Theta} \mathbb{D}_k(\mathbb{Q}, \mathbb{P}_\theta) + c\}$ *is bounded.*

Assumption 2. *The kernel* k *is such that* $|k(\xi, \xi')| \leq M$, $M < \infty$, *for any* $\xi, \xi' \in \Xi$.

Assumption 1 ensures that a minimiser in (8) exists and is a common assumption made in MMD estimator methods (see [30]). Assumption 2 is needed to obtain a concentration inequality for the NPL posterior and it is often made in methods using MMD estimators (see, e.g., [20,30,36,37]), as it ensures robustness guarantees (see, e.g., [20,30,36,37]). Intuitively, bounded kernels control the contribution of extreme or unlikely observations, preventing them from having an arbitrarily large effect on the distance and thereby yielding a robust measure of discrepancy. Many commonly used kernels are bounded, such as the Gaussian, Matern and Exponential kernels (see, e.g., [38]).

Theorem 2. *Suppose Assumptions 1 and 2 hold. Then with probability at least* $1 - \delta$*:*

$$\mathbb{D}_k(\mathbb{P}^\star, \hat{\mathbb{P}}_n^{pred(NPL)}) \leq \inf_{\theta \in \Theta} \mathbb{D}_k(\mathbb{P}_\theta, \mathbb{P}^\star) + C_{n,M,\alpha} \tag{18}$$

where $C_{n,M,\alpha}$ *is a constant depending on the number of samples* n*, the upper bound of the kernel* M *and the concentration parameter on the DP prior* α*.*

Remark 1. $C_{n,M,\alpha}$ *has an overall rate of* $1/\sqrt{n}$ *consistent with existing results for minimum MMD estimators [30,36]. Moreover, given an upper bound of the kernel* M*, the constant is fully known. Hence, if* $\inf_{\theta \in \Theta} \mathbb{D}_k(\mathbb{P}_\theta, \mathbb{P}^\star)$ *can be reasonably approximated, this result can be used to select the radius ensuring RoBAS includes* $\mathbb{P}^\star$ *with high probability. However, in practice, this theoretical radius can be over-conservative and a suitable value can be chosen via cross-validation or related bootstrapping procedures following standard practices in the DRO literature (see, e.g., [12,39,40]).*

We can now obtain an upper bound for the target optimisation problem for large enough ϵ.

Corollary 2. *Suppose Assumptions 1 and 2 hold and let* $C_{n,M,\alpha}$ *as in Theorem 2. Then, for* $\epsilon \geq C_{n,M,\alpha} + \inf_{\theta \in \Theta} \mathbb{D}_k(\mathbb{P}_\theta, \mathbb{P}^\star)$*, with probability at least* $1 - \delta$*:*

$$\mathbb{E}_{\xi \sim \mathbb{P}^\star}[f_x(\xi)] \leq \sup_{\mathcal{B}_\epsilon^k(\hat{\mathbb{P}}_n^{pred(NPL)})} \mathbb{E}_{\xi \sim \mathbb{P}}[f_x(\xi)].$$

In the special case of Huber's contamination model Huber [41], we can obtain a guarantee similar to Theorem 2 which depends on the contamination level.

Corollary 3 (Huber's cont. model)**.** *Suppose* $\mathbb{P}^\star = (1 - \eta)\mathbb{P}_{\theta_0} + \eta\mathbb{Q}$ *for some* $\theta_0 \in \Theta$, $\mathbb{Q} \in \mathcal{P}(\Xi)$ *and* $\eta \in [0, 1]$*. Suppose Assumptions 1 and 2 hold and let* $C_{n,M,\alpha}$ *as in Theorem 2. Then with probability at least* $1 - \delta$*:* $\mathbb{D}_k(\mathbb{P}_{\theta_0}, \hat{\mathbb{P}}_n^{pred(NPL)}) \leq 4\eta + 2C_{n,M,\alpha}$.

4. Experiments

We evaluate our method on several different DGPs, model families and misspecification settings for two decision-making problems: the Newsvendor and the Portfolio.

We compare our method to existing Bayesian formulations of DRO—DRO-BAS [12] and Bayesian DRO (BDRO) [11]—both based on the KL divergence and standard Bayesian posterior. To assess how much robustness in our framework is gained through the robust posterior compared to the choice of the MMD in the ambiguity set, we further compare against the empirical method which uses an MMD ball around the empirical measure (denoted by Empirical MMD). This was presented in Staib and Jegelka [5] and also forms a special case of Kernel DRO [6]. Implementation details are provided in Appendix B.

We explore two types of misspecification.

1. *Model misspecification* which occurs when the DGP $\mathbb{P}^\star$ does not belong to the model family $\mathcal{P}_\Theta$, e.g., if $\mathbb{P}^\star$ is multimodal while $\mathcal{P}_\Theta$ assumes unimodality. This affects Bayesian DRO methods (BDRO, DRO-BAS, DRO-RoBAS) but not empirical approaches, as the latter do not rely on a model.
2. *Huber's contamination model* [41] which is a specific type of model misspecification (see Figure 2) wherein the training DGP is $\mathbb{P}^\star = (1-\eta)\tilde{\mathbb{P}} + \eta\mathbb{Q}$ for some $\eta \in [0,1]$ and $\tilde{\mathbb{P}}, \mathbb{Q} \in \mathcal{P}(\Xi)$. Contamination, limited to the training set, impacts both Bayesian and empirical DRO methods since the test distribution is assumed to be $\tilde{\mathbb{P}}$. Huber contamination relates to concepts like distribution shift and out-of-distribution robustness (e.g., [42]) and its importance in DRO has attracted increasing attention in recent work [43,44].

4.1. The Newsvendor Problem

We start with the commonly explored Newsvendor problem (e.g., [45]). The goal is to choose the optimal amount of products to buy based on consumers' demand. The cost is defined as: $f(x,\xi) := h\max(x-\xi,0) + b\max(\xi-x,0)$ where $x \in \mathbb{R}^D_{\geq 0}$ is the number of product units ordered, $\xi \in \mathbb{R}^D$ is the consumers' demand, $0 \in \mathbb{R}^D$ is the zero vector and b and h denote the backorder and holding cost per unit respectively. In all examples we follow the implementation of Shapiro et al. [11], Dellaporta et al. [12] and set $b = 8$ and $h = 3$. We run each experiment $J = 100$ times for $n = 20$ observations and compute the out-of-sample mean and variance of the cost incurred.

We consider two models and several DGP cases as follows. First, we assume the demand $\xi \in \Xi$ follows a Gaussian distribution with known variance, i.e., $\mathbb{P}_\theta := \mathcal{N}(\theta, \sigma^2\mathcal{I}_{D\times D})$ while the DGP is a bimodal Gaussian distribution (case 1 above): $\mathbb{P}^\star := 0.5\mathcal{N}(\theta_1^\star, \sigma^2\mathcal{I}_{D\times D}) + 0.5\mathcal{N}(\theta_2^\star, \sigma^2\mathcal{I}_{D\times D})$ for $D = 1$ and $D = 5$. Furthermore, we explore the same Gaussian model with a contaminated Gaussian training DGP (case 2 above): $\mathbb{P}^\star_{\text{train}} := (1-\eta)\mathcal{N}(\theta^\star, \sigma^2) + \eta\mathcal{N}(\theta', \sigma^2)$ and an Exponential model $\mathbb{P}_\theta := \text{Exp}(\theta)$ with a contaminated Exponential DGP for the training DGP (case 2 above): $\mathbb{P}^\star_{\text{train}} := (1-\eta)\text{Exp}(\theta^\star) + \eta\mathcal{N}(\mu, \sigma)$ for $\eta \in \{0.0, 0.1, 0.2\}$.

Figure 3 presents the out-of-sample mean and variance of the methods for the bimodal univariate and multivariate Gaussian DGPs.

The effect of model misspecification is notably more pronounced for the DRO-BAS and BDRO instantiations, which are based on the standard Bayesian posterior and the KL divergence. In the DRO-RoBAS case, this robustness is likely due to the fact that the obtained NPL-MMD posterior is bimodal, even though the model itself is unimodal. In contrast, the standard Bayesian posterior is highly sensitive to misspecification, resulting in a unimodal posterior concentrated between the two modes. Consequently, DRO-BAS and BDRO require very large values of ϵ to capture the true DGP, leading to conservative decisions that incur high out-of-sample costs. This is evident as they achieve lower mean-variance as ϵ increases.

We further observe (Figure 3) that in the univariate case, empirical MMD achieves a lower (difference of $\leq$1) out-of-sample mean for most values of $\epsilon < 1$. However, DRO-

RoBAS consistently shows a lower out-of-sample variance. In the multivariate case, the performances of the two methods are similar, though empirical MMD outperforms DRO-RoBAS in both mean and variance. Notably, this comparison pits a Bayesian method under model misspecification against a completely empirical method unaffected by this misspecification.

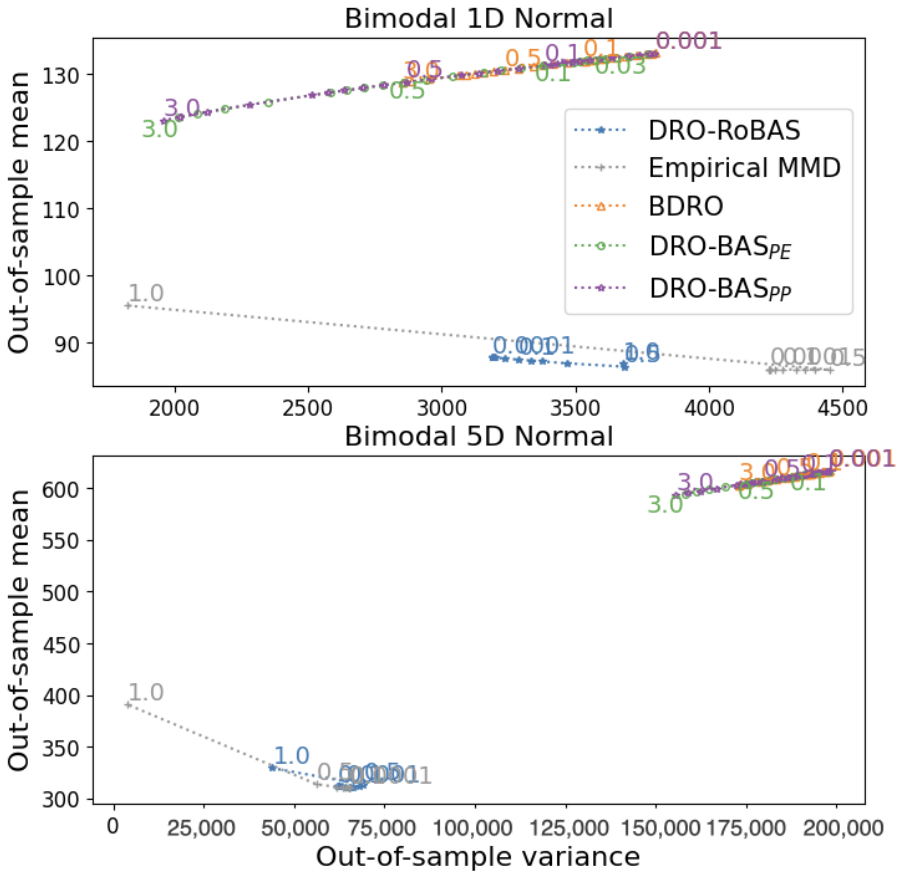

Figure 3. The out-of-sample mean and variance for the Newsvendor problem with a *misspecified* Gaussian location model and a bimodal Gaussian DGP. Results are shown for the univariate ($D = 1$, (**top**)) and the multivariate ($D = 5$, (**bottom**)) cases, with markers representing ϵ values. For illustration purposes, the bottom-left area of the multivariate case is shown in a zoomed-in view.

However, it is promising that DRO-RoBAS remains highly competitive against this baseline. The next example, based on contamination models, illustrates a scenario where robustness is crucial for both model-based and empirical methods.

In the second simulation, we consider the Huber contamination models [41] where the training set is contaminated, whereas the test set is not. Figure 4 demonstrates that both DRO-BAS formulations outperform DRO-RoBAS and the empirical MMD method in the well-specified case, where there is no contamination in the training set, and the training and test distributions are identical. In misspecified cases, where $\eta > 0$, DRO-RoBAS shows greater robustness compared to the other methods in terms of the out-of-sample mean-variance trade-off. This suggests that the robust posterior and robust distance measures in RoBAS contribute to a better-informed ambiguity set concerning the *test set generating process*.

This example further illustrates that, while the motivation for a robust ambiguity set stemmed from concerns about *model misspecification*, even entirely empirical methods, like the empirical MMD, can be sensitive to misspecifications arising from discrepancies between the training and test distributions.

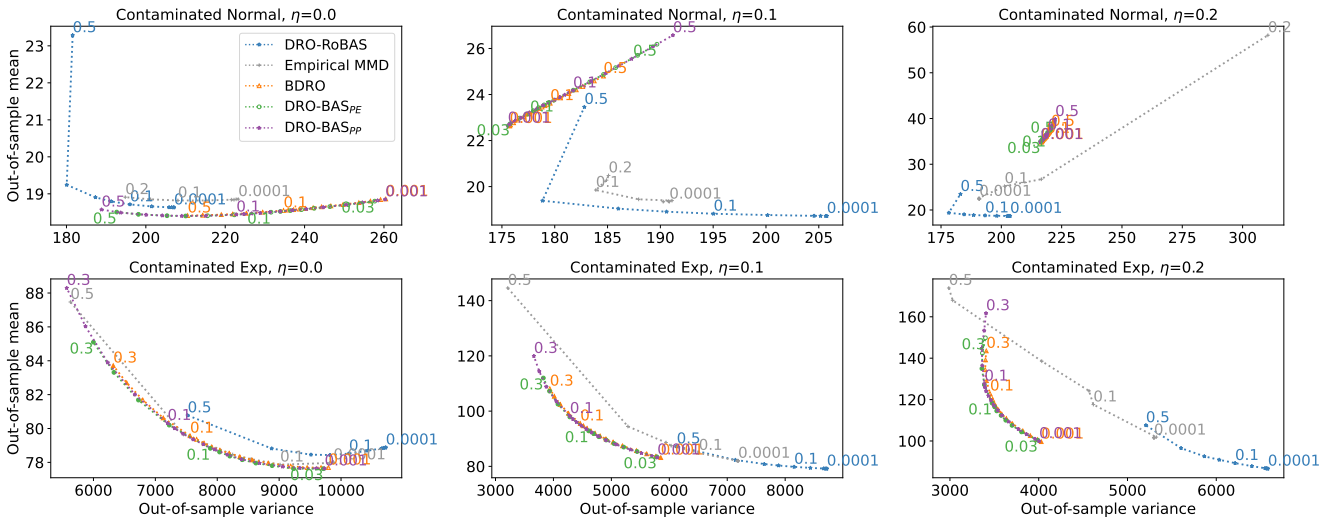

Figure 4. Out-of-sample mean-variance trade-off in the Newsvendor problem for a Gaussian location model (**top**) and an Exponential model (**bottom**) with a contaminated training dataset. Results are shown for contamination levels $\eta = 0.0$ (**left**), $\eta = 0.1$ (**middle**), and $\eta = 0.2$ (**right**). Each marker represents a specific ϵ value, with some labelled for reference.

4.2. The Portfolio Optimisation Problem

We continue with the multi-dimensional Portfolio problem, also considered by Shapiro et al. [11], which chooses stock weightings ($x \in \mathbb{R}^D$) to maximise returns. The objective function is $f_x(\xi) = -\xi^\top x$ which corresponds to maximising the return and the optimisation is subject to the constraints $x_i \geq 0$ for all $i = 1, \ldots, D$ and $\sum_{i=1}^{D} x_i = 1$. We generate $n = 100$ observations from a $5D$ Gaussian DGP with contamination on three-out-of-five dimensions: $\mathbb{P}^\star = (1-\eta)\mathcal{N}(\mu^\star, \Sigma^\star) + \eta\mathcal{N}(\mu', \Sigma^\star)$. We use a multivariate Gaussian model with unknown mean and variance.

Figure 5 shows DRO-RoBAS is unaffected by the contamination, whilst empirical MMD, DRO-BAS, and BDRO are negatively affected. Consider $\eta = 0.1$: for $\epsilon < 0.2$, empirical DRO performance quickly degrades compared to DRO-RoBAS; but, for $\epsilon \geq 0.2$, empirical MMD performs similarly due to the MMD robustness. This effect is magnified for $\eta = 0.2$, demonstrating that the empirical nominal distribution is unreliable due to outliers, whilst DRO-RoBAS benefits from a robust nominal—the NPL-MMD posterior predictive—thus performs better for small ϵ.

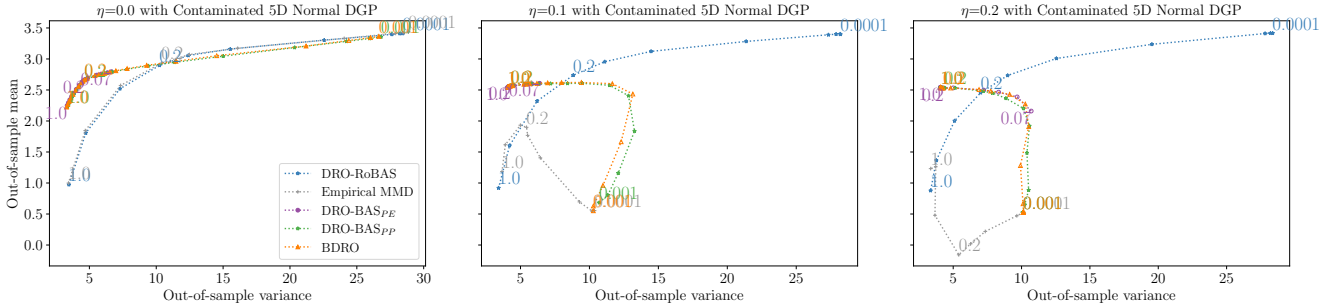

Figure 5. Out-of-sample mean-variance trade-off in the Portfolio problem for a 5D contaminated Gaussian DGP with $\eta = 0.0, 0.1, 0.2$. Note that the goal is to maximise returns, so *larger* out-of-sample mean is better.

4.3. Computational Time

The increased robustness of DRO-RoBAS comes at the cost of higher computational demands (see Tables A1 and A2 of Appendix B). This cost arises from the complex optimi-

sation problem in the RKHS and the longer sampling time required for the NPL posterior. However, as previously demonstrated, this cost is justified by improved out-of-sample performance across various cases of model misspecification. Moreover, by leveraging the NPL-MMD and the MMD, DRO-RoBAS can be used for any choice of model family, even likelihood-free models. Notice that DRO-BAS$_{\text{PE}}$ is limited to Exponential family models whereas the computational cost of BDRO and DRO-BAS$_{\text{PP}}$ increases considerably if the posterior is not available in closed form, as methods like Markov Chain Monte Carlo are needed for posterior sampling. This highlights the flexibility and robustness of DRO-RoBAS despite its computational demands. Nevertheless, possible scalability improvements are discussed in Section 5.

5. Conclusions

Bayesian formulations of DRO for decision-making problems can suffer from model misspecification as the ambiguity set heavily relies on the non-robust Bayesian posterior. We addressed this challenge by using a robust NPL posterior to inform the ambiguity set and leveraging the MMD to construct *both* the posterior and the ambiguity set itself. We show that DRO-RoBAS admits a dual formulation in the RKHS and we provide probabilistic guarantees for the tolerance level such that the resulting optimisation problem upper bounds the true objective with high probability.

Scalability improvements for DRO-RoBAS can be achieved through existing tools from kernel methods such as Fourier features [46] and low-rank kernel matrix approximations [47] that are left for future work. Our empirical evidence suggests that if the model is well specified, or the level of misspecification is low, then existing Bayesian formulations like DRO-BAS can have better performance and scalability. At the same time, when model misspecification is moderate or high then DRO-RoBAS achieves significantly improved out-of-sample performance and robustness. Note that any prior knowledge on the misspecification level can be naturally incorporated into our framework to further boost performance.

Finally, the construction of Robust Bayesian Ambiguity Sets can extend beyond the choices of NPL and MMD. For example, another instantiation of our framework arises if we employ Generalised Bayesian Inference (GBI) [21]. One of the motivations of GBI is to induce robustness with respect to model misspecification [22–26,48,49] by targeting a different divergence than the KL. While these methods do not directly impose uncertainty on the DGP like NPL does, they can produce robust GBI posteriors, making it worthwhile to integrate into DRO-RoBAS. Notably, the duality results in Section 3.1 are based on the MMD choice; however, they hold under a general posterior and are *not* dependent on the NPL framework.

Author Contributions: Conceptualization, C.D., P.O. and T.D.; methodology, C.D., P.O. and T.D.; software, C.D. and P.O.; validation, C.D., P.O. and T.D.; formal analysis, C.D., P.O. and T.D.; investigation, C.D., P.O. and T.D.; writing—original draft preparation, C.D. and P.O.; writing—review and editing, C.D., P.O. and T.D. All authors have read and agreed to the published version of the manuscript.

Funding: CD acknowledges support from EPSRC grant [EP/T51794X/1] as part of the Warwick CDT in Mathematics and Statistics and EPSRC grant [EP/Y022300/1]. PO and TD acknowledge support from a UKRI Turing AI acceleration Fellowship [EP/V02678X/1] and a Turing Impact Award from the Alan Turing Institute. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC-BY) licence to any Author Accepted Manuscript version arising from this submission.

Data Availability Statement: The code to reproduce the simulated experiments are openly available at https://github.com/PatrickOHara/mis-dro-code (accessed on 8 April 2026).

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Proofs of Theoretical Results

In this section, we provide detailed proofs of the theoretical results concerning the duality of the DRO-RoBAS problem and the tolerance level guarantees.

Appendix A.1. Proof of Proposition 1

First, we derive the support function of the effective ambiguity set in terms of kernel mean embeddings defined in Section 3.1:

$$\mathcal{C}^\star := \{\mu \in \mathcal{H}_k : \|\mu - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_k \leq \epsilon\} \tag{A1}$$

where we denote by $\mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}[\cdot]$ the expectation under $\mathbb{E}_{\mathbb{Q}\sim\text{DP}(c',\mathbb{F}')}[\cdot]$ and we have defined the NPL-MMD posterior predictive as $\mathbb{P}_n^{\text{pred(NPL)}} := \mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}[\mathbb{P}_{\theta_k(\mathbb{Q})}]$. We start with the following Lemma.

Lemma A1. *For any $\epsilon > 0$ we have:*

$$\sup_{\|\mu-\mu_{\mathbb{P}_n^{pred(NPL)}}\|_k\leq\epsilon} \left\langle g, \mu - \mu_{\mathbb{P}_n^{pred(NPL)}} \right\rangle_{\mathcal{H}_k} = \epsilon\|g\|_k.$$

Proof. The proof follows the same logic as the proof for the ambiguity set corresponding to the MMD ball around the empirical measure, provided in Appendix A.2 of Zhu et al. [6]. To prove the equality statement we will prove both sides of the inequality. We first prove that the left-hand side is less than or equal to the right-hand side. Applying the Cauchy–Schwarz inequality, we obtain:

$$\sup_{\|\mu-\mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_k\leq\epsilon} \left\langle g, \mu - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle_{\mathcal{H}_k} \leq \sup_{\|\mu-\mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_k\leq\epsilon} \|g\|_k\|\mu - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_k = \epsilon\|g\|_k. \tag{A2}$$

For the opposite direction, let $\mu' := \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} + \epsilon\frac{g}{\|g\|_k}$. Then

$$\begin{aligned}
&\|\mu' - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|^2_{\mathcal{H}_k} \\
&= \left\langle \mu' - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}, \mu' - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle_{\mathcal{H}_k} \\
&= \langle \mu', \mu' \rangle_{\mathcal{H}_k} - 2\left\langle \mu', \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle_{\mathcal{H}_k} + \left\langle \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}, \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle_{\mathcal{H}_k} \\
&= \left\langle \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} + \epsilon\frac{g}{\|g\|_k}, \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} + \epsilon\frac{g}{\|g\|_k} \right\rangle_{\mathcal{H}_k} + \left\langle \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}, \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle_{\mathcal{H}_k} \\
&\quad - 2\left\langle \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} + \epsilon\frac{g}{\|g\|_k}, \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle_{\mathcal{H}_k} \\
&= \left\langle \epsilon\frac{g}{\|g\|_k}, \epsilon\frac{g}{\|g\|_k} \right\rangle_{\mathcal{H}_k} \\
&= \epsilon^2\frac{\langle g, g\rangle_{\mathcal{H}_k}}{\|g\|_k^2} \\
&= \epsilon^2
\end{aligned}$$

and hence

$$\|\mu' - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_k = \epsilon$$

which proves that $\mu' \in \mathcal{C}^\star$. To complete the proof it suffices to show that

$$\left\langle g, \mu' - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle = \epsilon \|g\|_k. \tag{A3}$$

By definition of μ' we have:

$$\begin{aligned}\left\langle g, \mu' - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle_{\mathcal{H}_k} &= \left\langle g, \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} + \epsilon \frac{g}{\|g\|_k} - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle \\ &= \left\langle g, \epsilon \frac{g}{\|g\|_k} \right\rangle \\ &= \epsilon \|g\|_k.\end{aligned}$$

Since $\mu' \in \mathcal{C}^\star$ we have

$$\sup_{\|\mu - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_k \leq \epsilon} \left\langle g, \mu - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle_{\mathcal{H}_k} \geq \left\langle g, \mu' - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle_{\mathcal{H}_k} = \epsilon \|g\|_k$$

which completes the proof. □

We are now ready to prove Proposition 1.

Proof of Proposition 1. The result follows from the definition of the support function by applying Lemma A1 and the reproducing property as follows:

$$\begin{aligned}\delta^\star_{\mathcal{C}^\star}(g) &= \sup_{\mu \in \mathcal{C}^\star} \langle g, \mu \rangle_{\mathcal{H}_k} \\ &= \left\langle g, \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle_{\mathcal{H}_k} + \sup_{\mu \in \mathcal{C}^\star} \langle g, \mu \rangle_{\mathcal{H}_k} - \left\langle g, \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle_{\mathcal{H}_k} \\ &= \left\langle g, \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle_{\mathcal{H}_k} + \sup_{\mu \in \mathcal{C}^\star} \left\langle g, \mu - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle_{\mathcal{H}_k} \\ &\overset{(1)}{=} \mathbb{E}_{\text{DP}_{\xi_{1:n}}} \left[\mathbb{E}_{\mathbb{P}_{\theta_k}(\mathbb{Q})} [g(x)] \right] + \sup_{\|\mu - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_k \leq \epsilon} \left\langle g, \mu - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}} \right\rangle_{\mathcal{H}_k} \\ &\overset{(2)}{=} \mathbb{E}_{\text{DP}_{\xi_{1:n}}} \left[\mathbb{E}_{\mathbb{P}_{\theta_k}(\mathbb{Q})} [g(x)] \right] + \epsilon \|g\|_k.\end{aligned}$$

Equality (1) follows from the reproducing property and the definition of a kernel mean embedding by noting that:

$$\mu_{\mathbb{P}_n^{\text{pred(NPL)}}} := \mathbb{E}_{\xi \sim \mathbb{P}_n^{\text{pred(NPL)}}} [\phi(\xi)] = \mathbb{E}_{Q \sim \text{DP}_{\xi_{1:n}}} [\mathbb{E}_{\xi \sim \mathbb{P}_{\theta_k}(\mathbb{Q})} [\phi(\xi)]] = \mathbb{E}_{\text{DP}_{\xi_{1:n}}} [\mu_{\mathbb{P}_{\theta_k}(\mathbb{Q})}] \tag{A4}$$

where ϕ denotes the feature map associated with kernel k. Equality (2) follows from Lemma A1. □

Appendix A.2. Proof of Corollary 1

To prove the Corollary it suffices to verify the assumptions of Theorem 1 for our set $\mathcal{C}^\star$.

Lemma A2. *$\mathcal{C}^\star$, as defined in (A1), is closed and convex. Furthermore, ambiguity set $\mathcal{K}_{\mathcal{C}^\star}$ satisfies* $\text{ri}(\mathcal{K}_{\mathcal{C}^\star}) \neq \emptyset$, *where*

$$\mathcal{K}_{\mathcal{C}^\star} := \{\mathbb{P} \in \mathcal{P}_k(\Xi), \mu_{\mathbb{P}} \in \mathcal{C}^\star\} = \{\mathbb{P} : \mathbb{D}_k(\mathbb{P}, \mathbb{P}_n^{pred(NPL)}) \leq \epsilon\}. \tag{A5}$$

Proof. *Convexity:* Convexity follows trivially by the triangle inequality. Let $\mu_1, \mu_2 \in \mathcal{C}^\star$ and $\lambda \in [0,1]$. Then

$$\begin{aligned}\left\|\lambda\mu_1 + (1-\lambda)\mu_2 - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\right\|_k &\leq \lambda\left\|\mu_1 - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\right\|_k + (1-\lambda)\left\|\mu_2 - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\right\|_k \\ &\leq \lambda\epsilon + (1-\lambda)\epsilon = \epsilon\end{aligned}$$

where the first inequality follows from the triangle inequality applied to the RKHS norm. Hence, $\lambda\mu_1 + (1-\lambda)\mu_2 \in \mathcal{C}^\star$ and $\mathcal{C}^\star$ is convex.

Closeness of $\mathcal{C}^\star$: By definition, a closed set is a set whose complement is open, so it suffices to show the set $\mathcal{H}_k \backslash \mathcal{C}^\star$ is open. Again by definition, set $\mathcal{H}_k \backslash \mathcal{C}^\star$ is open if, for all $y \in \mathcal{H}_k \backslash \mathcal{C}^\star$, there exists $r > 0$ such that any point $z \in \mathcal{H}_k$ satisfying $\|y - z\|_{\mathcal{H}_k} < r$ also belongs to $\mathcal{H}_k \backslash \mathcal{C}^\star$. Given any $y \in \mathcal{H}_k \backslash \mathcal{C}^\star$, let $r = \|y - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_{\mathcal{H}_k} - \epsilon$. Observe that $r > 0$ because $\|y - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_{\mathcal{H}_k} > \epsilon$. Now, for all $z \in \mathcal{H}_k$ such that $\|y - z\|_{\mathcal{H}_k} < r$, we can apply the triangle inequality to obtain

$$\|y - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_{\mathcal{H}_k} \leq \|z - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_{\mathcal{H}_k} + \|z - y\|_{\mathcal{H}_k}.$$

After rearranging, we have

$$\begin{aligned}\|z - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_{\mathcal{H}_k} &\geq \|y - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_{\mathcal{H}_k} - \|z - y\|_{\mathcal{H}_k} \\ &> \|y - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_{\mathcal{H}_k} - r = \epsilon.\end{aligned}$$

Therefore, $\|z - \mu_{\mathbb{P}_n^{\text{pred(NPL)}}}\|_{\mathcal{H}_k} > \epsilon$, so z lies in set $\mathcal{H}_k \backslash \mathcal{C}^\star$, which concludes our claim that $\mathcal{H}_k \backslash \mathcal{C}^\star$ is open and proves that $\mathcal{C}^\star$ is closed.

Non-empty relative interior of $\mathcal{K}_{\mathcal{C}^\star}$: It suffices to prove that $\mathcal{K}_{\mathcal{C}^\star}$ is non-empty and convex. Convexity follows from an analogous proof of convexity of $\mathcal{C}^\star$ since MMD satisfies the triangle inequality. $\mathcal{K}_{\mathcal{C}^\star}$ is non-empty since, for any $\epsilon > 0$, $\mathbb{P}_n^{\text{pred(NPL)}} \in \mathcal{K}_{\mathcal{C}^\star}$. □

Appendix A.3. Proof of Theorem 2

We first provide a Lemma adapted from the results of Dellaporta et al. [20] which bounds the expected MMD (under a finite DGP sample) between the DGP and a sample from the approximate DP posterior. This is necessary to quantify the generalisation error between the DGP and the obtained model $\mathbb{P}_{\theta_k(\mathbb{Q})}$ for a DP sample $\mathbb{Q}$.

Lemma A3. *Let $\mathbb{Q} \sim \widehat{\text{DP}}_{\xi_{1:n}}$ be defined as in (17). Assume that $\xi_{1:n} \overset{\text{iid}}{\sim} \mathbb{P}^\star$ and that the kernel k is such that $|k(\xi, \xi')| \leq K$, $K < \infty$, for any $\xi, \xi' \in \Xi$. Then*

$$\mathbb{E}_{\xi_{1:n} \overset{\text{iid}}{\sim} \mathbb{P}^\star}\left[\mathbb{E}_{\widehat{\text{DP}}_{\xi_{1:n}}}[\mathbb{D}_k(\mathbb{P}^\star, \mathbb{Q})]\right] \leq \sqrt{\frac{K}{n}} + \sqrt{\frac{2K(n-1) + c(c+1)}{(c+n)(c+n+1)}} + \sqrt{\frac{Kc(c+1)}{(c+n)(c+n+1)}}. \tag{A6}$$

Proof. The result follows directly from Lemmas 6 and 11 of Dellaporta et al. [20] with some small modifications. First notice that by the triangle inequality

$$\mathbb{D}_k(\mathbb{P}^\star, \mathbb{Q}) \leq \mathbb{D}_k(\mathbb{P}^\star, \hat{\mathbb{P}}_n) + \mathbb{D}_k(\hat{\mathbb{P}}_n, \mathbb{Q}) \tag{A7}$$

where $\hat{\mathbb{P}}_n := \frac{1}{n}\sum_{i=1}^n \delta_{\xi_{1:n}}$. By Dellaporta et al. [20], Lemma 6, we have that

$$\mathbb{E}_{\xi_{1:n} \overset{\text{iid}}{\sim} \mathbb{P}^\star}\left[\mathbb{D}_k(\mathbb{P}^\star, \hat{\mathbb{P}}_n)\right] \leq \sqrt{\frac{K}{n}}. \tag{A8}$$

Moreover, by Dellaporta et al. [20], Lemma 11, we have that:

$$\mathbb{E}_{\xi_{1:n} \overset{\text{iid}}{\sim} \mathbb{P}^\star}\left[\mathbb{E}_{\mathbb{Q}\sim\widehat{\mathrm{DP}}_{\xi_{1:n}}}\left[\mathbb{D}_k(\hat{\mathbb{P}}_n, \mathbb{Q})\right]\right] \leq \sqrt{\frac{2K(n-1)+c(c+1)}{(c+n)(c+n+1)}} + \sqrt{\frac{Kc(c+1)}{(c+n)(c+n+1)}}. \tag{A9}$$

Taking expectations and using Equations (A8) and (A9) in (A7) we obtain the required result:

$$\begin{aligned}\mathbb{E}_{\xi_{1:n} \overset{\text{iid}}{\sim} \mathbb{P}^\star}\left[\mathbb{E}_{\mathbb{Q}\sim\widehat{\mathrm{DP}}_{\xi_{1:n}}}[\mathbb{D}_k(\mathbb{P}^\star, \mathbb{Q})]\right] &\leq \mathbb{E}_{\xi_{1:n} \overset{\text{iid}}{\sim} \mathbb{P}^\star}\left[\mathbb{D}_k(\mathbb{P}^\star, \hat{\mathbb{P}}_n)\right] + \mathbb{E}_{\xi_{1:n} \overset{\text{iid}}{\sim} \mathbb{P}^\star}\left[\mathbb{E}_{\mathbb{Q}\sim\widehat{\mathrm{DP}}_{\xi_{1:n}}}\left[\mathbb{D}_k(\hat{\mathbb{P}}_n, \mathbb{Q})\right]\right] \\ &\leq \sqrt{\frac{K}{n}} + \sqrt{\frac{2K(n-1)+c(c+1)}{(c+n)(c+n+1)}} + \sqrt{\frac{Kc(c+1)}{(c+n)(c+n+1)}}.\end{aligned} \tag{A10}$$

□

Based on this bound we can now provide a result in probability.

Lemma A4. *Assume that* $\xi_{1:n} \overset{\text{iid}}{\sim} \mathbb{P}^\star$ *and that kernel* k *is such that* $|k(\xi, \xi')| \leq K$, $K < \infty$, *for any* $\xi, \xi' \in \Xi$. *Then with probability* $1-\delta$ *we have:*

$$\begin{aligned}\mathbb{E}_{\mathbb{Q}\sim\widehat{\mathrm{DP}}_{\xi_{1:n}}}[\mathbb{D}_k(\mathbb{P}^\star, \mathbb{Q})] \leq \sqrt{\frac{K}{n}} + \sqrt{\frac{2K(n-1)+c(c+1)}{(c+n)(c+n+1)}} \\ + \sqrt{\frac{Kc(c+1)}{(c+n)(c+n+1)}} + \frac{\sqrt{2nK\left(\log\frac{1}{\delta}\right)}}{c+n}.\end{aligned} \tag{A11}$$

Proof. We follow the proof technique of Lemma 1 of Briol et al. [30] based on McDiarmid's inequality [50]. First, notice that by definition of $\widehat{\mathrm{DP}}_{\xi_{1:n}}$ in (17) we can re-write the objective as

$$\mathbb{E}_{\mathbb{Q}\sim\widehat{\mathrm{DP}}_{\xi_{1:n}}}[\mathbb{D}_k(\mathbb{P}^\star, \mathbb{Q})] = \mathbb{E}_{w\sim\mathrm{Dir}}\left[\mathbb{E}_{\tilde{\xi}_{1:T} \overset{\text{iid}}{\sim} \mathbb{F}}\left[\mathbb{D}_k(\mathbb{P}^\star, \widehat{\mathbb{Q}}_{\xi_{1:n}})\right]\right] \tag{A12}$$

where $\widehat{\mathbb{Q}}_{\xi_{1:n}} = \sum_{i=1}^n w_i \delta_{\xi_i} + \sum_{k=1}^\tau \tilde{w}_k \delta_{\tilde{\xi}_k} \sim \widehat{\mathrm{DP}}_{\xi_{1:n}}$ and $\mathbb{E}_{w\sim\mathrm{Dir}}$ denotes the expectation under $\mathrm{Dir}\left(1, \ldots 1, \frac{c}{\tau}, \ldots, \frac{c}{\tau}\right)$. Let $h(\xi_1, \ldots, \xi_n) := \mathbb{E}_{w\sim\mathrm{Dir}}\left[\mathbb{E}_{\tilde{\xi}_{1:T} \overset{\text{iid}}{\sim} \mathbb{F}}\left[\mathbb{D}_k(\mathbb{P}^\star, \widehat{\mathbb{Q}}_{\xi_{1:n}})\right]\right]$. Then for all $\{\xi_i\}_{i=1}^n, \xi_i' \in \Xi$ and $\xi'_{1:n} := \{\xi_1, \ldots, \xi_{i-1}, \xi_i', \xi_{i+1}, \ldots, \xi_n\}$ we have

$$\begin{aligned}&\left|h(\xi_1, \ldots, \xi_{i-1}, \xi_i, \xi_{i+1}, \ldots, \xi_n) - h(\xi_1, \ldots, \xi_{i-1}, \xi_i', \xi_{i+1}, \ldots, \xi_n)\right| \\ &= \left|\mathbb{E}_{w\sim\mathrm{Dir}}\left[\mathbb{E}_{\tilde{\xi}_{1:T} \overset{\text{iid}}{\sim} \mathbb{F}}\left[\mathbb{D}_k(\mathbb{P}^\star, \widehat{\mathbb{Q}}_{\xi_{1:n}})\right]\right] - \mathbb{E}_{w\sim\mathrm{Dir}}\left[\mathbb{E}_{\tilde{\xi}_{1:T} \overset{\text{iid}}{\sim} \mathbb{F}}\left[\mathbb{D}_k(\mathbb{P}^\star, \widehat{\mathbb{Q}}_{\xi'_{1:n}})\right]\right]\right| \\ &= \left|\mathbb{E}_{w\sim\mathrm{Dir}}\left[\mathbb{E}_{\tilde{\xi}_{1:T} \overset{\text{iid}}{\sim} \mathbb{F}}\left[\mathbb{D}_k(\mathbb{P}^\star, \widehat{\mathbb{Q}}_{\xi_{1:n}}) - \mathbb{D}_k(\mathbb{P}^\star, \widehat{\mathbb{Q}}_{\xi'_{1:n}})\right]\right]\right| \\ &\leq \mathbb{E}_{w\sim\mathrm{Dir}}\left[\mathbb{E}_{\tilde{\xi}_{1:T} \overset{\text{iid}}{\sim} \mathbb{F}}\left[\left|\mathbb{D}_k(\mathbb{P}^\star, \widehat{\mathbb{Q}}_{\xi_{1:n}}) - \mathbb{D}_k(\mathbb{P}^\star, \widehat{\mathbb{Q}}_{\xi'_{1:n}})\right|\right]\right] \\ &= \mathbb{E}_{w\sim\mathrm{Dir}}\left[\mathbb{E}_{\tilde{\xi}_{1:T} \overset{\text{iid}}{\sim} \mathbb{F}}\left[\left|\|\mu^\star_{\mathbb{P}} - \mu_{\widehat{\mathbb{Q}}_{x_{1:n}}}\|_k - \|\mu^\star_{\mathbb{P}} - \mu_{\widehat{\mathbb{Q}}_{\xi'_{1:n}}}\|_k\right|\right]\right] \\ &\leq \mathbb{E}_{w\sim\mathrm{Dir}}\left[\mathbb{E}_{\tilde{\xi}_{1:T} \overset{\text{iid}}{\sim} \mathbb{F}}\left[\|\mu^\star_{\mathbb{P}} - \mu_{\widehat{\mathbb{Q}}_{\xi_{1:n}}} - \mu^\star_{\mathbb{P}} + \mu_{\widehat{\mathbb{Q}}_{\xi'_{1:n}}}\|_k\right]\right] \\ &= \mathbb{E}_{w\sim\mathrm{Dir}}\left[\mathbb{E}_{\tilde{\xi}_{1:T} \overset{\text{iid}}{\sim} \mathbb{F}}\left[\|w_i(k(\xi_i, \cdot) - k(\xi_i', \cdot))\|_k\right]\right] \\ &= \mathbb{E}_{w\sim\mathrm{Dir}}\left[w_i\|k(\xi_i, \cdot) - k(\xi_i', \cdot)\|_k\right] \\ &\leq \mathbb{E}_{w\sim\mathrm{Dir}}[w_i] 2\sqrt{K} \\ &= \frac{2\sqrt{K}}{c+n}\end{aligned} \tag{A13}$$

where in the first inequality we used Jensen's inequality and in the second we used the reverse triangle inequality. Then by McDiarmid's inequality [50] we have that for any $\epsilon > 0$:

$$\begin{aligned}
&\mathbb{P}\Bigg(\mathbb{E}_{w\sim\mathrm{Dir}}\left[\mathbb{E}_{\tilde{\xi}_{1:T}\overset{\mathrm{iid}}{\sim}\mathbb{F}}\left[\mathbb{D}_k(\mathbb{P}^\star,\hat{\mathbb{Q}}_{\xi_{1:n}})\right]\right] \\
&\qquad -\mathbb{E}_{\xi_{1:n}\overset{\mathrm{iid}}{\sim}\mathbb{P}^\star}\left[\mathbb{E}_{w\sim\mathrm{Dir}}\left[\mathbb{E}_{\tilde{\xi}_{1:T}\overset{\mathrm{iid}}{\sim}\mathbb{F}}\left[\mathbb{D}_k(\mathbb{P}^\star,\hat{\mathbb{Q}}_{\xi_{1:n}})\right]\right]\right] \geq \epsilon\Bigg) \\
&\leq \exp\left(\frac{-2\epsilon^2}{\frac{4nK}{(n+c)^2}}\right) \\
&= \exp\left(\frac{-\epsilon^2(n+c)^2}{2nK}\right)
\end{aligned} \tag{A14}$$

Let $\delta := \exp\left(\frac{-\epsilon^2(n+c)^2}{2nK}\right)$ then it follows that with probability $1-\delta$:

$$\begin{aligned}
&\mathbb{E}_{w\sim\mathrm{Dir}}\left[\mathbb{E}_{\tilde{\xi}_{1:T}\overset{\mathrm{iid}}{\sim}\mathbb{F}}\left[\mathbb{D}_k(\mathbb{P}^\star,\hat{\mathbb{Q}}_{\xi_{1:n}})\right]\right] \\
&\leq \mathbb{E}_{\xi_{1:n}\overset{\mathrm{iid}}{\sim}\mathbb{P}^\star}\left[\mathbb{E}_{w\sim\mathrm{Dir}}\left[\mathbb{E}_{\tilde{\xi}_{1:T}\overset{\mathrm{iid}}{\sim}\mathbb{F}}\left[\mathbb{D}_k(\mathbb{P}^\star,\hat{\mathbb{Q}}_{\xi_{1:n}})\right]\right]\right] + \epsilon \\
&= \mathbb{E}_{\xi_{1:n}\overset{\mathrm{iid}}{\sim}\mathbb{P}^\star}\left[\mathbb{E}_{w\sim\mathrm{Dir}}\left[\mathbb{E}_{\tilde{\xi}_{1:T}\overset{\mathrm{iid}}{\sim}\mathbb{F}}\left[\mathbb{D}_k(\mathbb{P}^\star,\hat{\mathbb{Q}}_{\xi_{1:n}})\right]\right]\right] + \frac{\sqrt{2nK\left(\log\frac{1}{\delta}\right)}}{n+c} \\
&\leq \sqrt{\frac{K}{n}} + \sqrt{\frac{2K(n-1)+c(c+1)}{(c+n)(c+n+1)}} + \sqrt{\frac{Kc(c+1)}{(c+n)(c+n+1)}} \\
&\quad + \frac{\sqrt{2nK\left(\log\frac{1}{\delta}\right)}}{c+n}
\end{aligned} \tag{A15}$$

where the last inequality follows from Lemma A3. □

We are now ready to prove the main result.

Proof of Theorem 2. First note that

$$\begin{aligned}
\mathbb{D}_k(\mathbb{P}^\star,\mathbb{P}_{\theta_k(\mathbb{Q})}) - \inf_{\theta\in\Theta}\mathbb{D}_k(\mathbb{P}_\theta,\mathbb{P}^\star) &\leq \mathbb{D}_k(\mathbb{P}^\star,\mathbb{Q}) + \mathbb{D}_k(\mathbb{Q},\mathbb{P}_{\theta_k(\mathbb{Q})}) - \inf_{\theta\in\Theta}\mathbb{D}_k(\mathbb{P}_\theta,\mathbb{P}^\star) \\
&= \mathbb{D}_k(\mathbb{P}^\star,\mathbb{Q}) + \inf_{\theta\in\Theta}\mathbb{D}_k(\mathbb{P}_\theta,\mathbb{Q}) - \inf_{\theta\in\Theta}\mathbb{D}_k(\mathbb{P}_\theta,\mathbb{P}^\star) \\
&\leq \mathbb{D}_k(\mathbb{P}^\star,\mathbb{Q}) + \left|\inf_{\theta\in\Theta}\mathbb{D}_k(\mathbb{P}_\theta,\mathbb{Q}) - \inf_{\theta\in\Theta}\mathbb{D}_k(\mathbb{P}_\theta,\mathbb{P}^\star)\right| \\
&\leq \mathbb{D}_k(\mathbb{P}^\star,\mathbb{Q}) + \sup_{\theta\in\Theta}|\mathbb{D}_k(\mathbb{P}_\theta,\mathbb{Q}) - \mathbb{D}_k(\mathbb{P}_\theta,\mathbb{P}^\star)| \\
&\leq \mathbb{D}_k(\mathbb{P}^\star,\mathbb{Q}) + \sup_{\theta\in\Theta}\mathbb{D}_k(\mathbb{Q},\mathbb{P}^\star) \\
&= 2\mathbb{D}_k(\mathbb{P}^\star,\mathbb{Q})
\end{aligned} \tag{A16}$$

where we used the triangle inequality in the first inequality, the definition of $\theta_k(\cdot)$ in the first equality and the reversed triangle inequality in the last inequality. The third inequality follows from the fact (mentioned in Briol et al. [30], Theorem 1) that since k is bounded, the family $\inf_{\theta\in\Theta}\mathbb{D}_k(\mathbb{P}_\theta,\cdot)$ is uniformly bounded and for bounded functions f, g we have $\|\inf_\theta f(\theta) - \inf_\theta g(\theta)\| \leq \sup_\theta |f(\theta) - g(\theta)|$. Hence, by Jensen's inequality, since the MMD is convex, we have

$$\mathbb{D}_k(\mathbb{P}^\star, \hat{\mathbb{P}}_n^{\text{pred(NPL)}}) = \mathbb{D}_k(\mathbb{P}^\star, \mathbb{E}_{\widehat{\text{DP}}_{\xi_{1:n}}}[\mathbb{P}_{\theta_k(\mathbb{Q})}]) \leq \mathbb{E}_{\widehat{\text{DP}}_{\xi_{1:n}}}[\mathbb{D}_k(\mathbb{P}^\star, \mathbb{P}_{\theta_k(\mathbb{Q})})] \tag{A17}$$

and therefore

$$\begin{aligned}\mathbb{D}_k(\mathbb{P}^\star, \hat{\mathbb{P}}_n^{\text{pred(NPL)}}) - \inf_{\theta\in\Theta} \mathbb{D}_k(\mathbb{P}_\theta, \mathbb{P}^\star) &\leq \mathbb{E}_{\mathbb{Q}\sim\widehat{\text{DP}}_{\xi_{1:n}}}[\mathbb{D}_k(\mathbb{P}^\star, \mathbb{P}_{\theta_k(\mathbb{Q})})] - \inf_{\theta\in\Theta} \mathbb{D}_k(\mathbb{P}_\theta, \mathbb{P}^\star) \\ &\leq 2\mathbb{E}_{\mathbb{Q}\sim\widehat{\text{DP}}_{\xi_{1:n}}}[\mathbb{D}_k(\mathbb{P}^\star, \mathbb{Q})].\end{aligned} \tag{A18}$$

By Lemma A4 it follows that with probability $1-\delta$ the advertised result holds. □

Appendix A.4. Proof of Corollary 3

Recall that $\mathbb{P}^\star = (1-\eta)\mathbb{P}_{\theta_0} + \eta\mathbb{Q}$ for some $\theta_0 \in \Theta$, $\mathbb{Q} \in \mathcal{P}(\Xi)$ and $\eta \in [0,1]$. Corollary 3 is an immediate consequence of Lemma 3.3 of Chérief-Abdellatif and Alquier [36] which states that for any $\theta \in \Theta$:

$$|\mathbb{D}_k(\mathbb{P}_\theta, \mathbb{P}^\star) - \mathbb{D}_k(\mathbb{P}_\theta, \mathbb{P}_{\theta_0})| \leq 2\eta. \tag{A19}$$

Using this and Theorem 2 we can prove the advertised result as follows:

Proof of Corollary 3. It follows by (A19) that:

$$\mathbb{D}_k(\mathbb{P}_{\theta_k(\mathbb{Q})}, \mathbb{P}_{\theta_0}) \leq 2\eta + \mathbb{D}_k(\mathbb{P}_{\theta_k(\mathbb{Q})}, \mathbb{P}^\star). \tag{A20}$$

and hence

$$\mathbb{E}_{\widehat{\text{DP}}_{\xi_{1:n}}}[\mathbb{D}_k(\mathbb{P}_{\theta_k(\mathbb{Q})}, \mathbb{P}_{\theta_0})] \leq 2\eta + \mathbb{E}_{\widehat{\text{DP}}_{\xi_{1:n}}}[\mathbb{D}_k(\mathbb{P}_{\theta_k(\mathbb{Q})}, \mathbb{P}^\star)]. \tag{A21}$$

We then have:

$$\begin{aligned}\mathbb{D}_k(\mathbb{P}_{\theta_0}, \hat{\mathbb{P}}_n^{\text{pred(NPL)}}) &\leq \mathbb{E}_{\widehat{\text{DP}}_{\xi_{1:n}}}[\mathbb{D}_k(\mathbb{P}_{\theta_0}, \mathbb{P}_{\theta_k})] \\ &\leq 2\mathbb{E}_{\widehat{\text{DP}}_{\xi_{1:n}}}[\mathbb{D}_k(\mathbb{P}_{\theta_0}, \mathbb{Q})] + \inf_\theta \mathbb{D}_k(\mathbb{P}_{\theta_0}, \mathbb{P}_\theta) \\ &= 2\mathbb{E}_{\widehat{\text{DP}}_{\xi_{1:n}}}[\mathbb{D}_k(\mathbb{P}_{\theta_0}, \mathbb{Q})] \\ &\leq 4\eta + 2\mathbb{E}_{\widehat{\text{DP}}_{\xi_{1:n}}}[\mathbb{D}_k(\mathbb{P}^\star, \mathbb{Q})]\end{aligned}$$

and the result follows from Lemma A4. The first inequality follows from Jensen's inequality as in (A17), the second inequality follows from (A16), the equality holds from the fact that $\mathbb{P}_{\theta_0} \in \mathcal{P}_\Theta$ and the last inequality follows form (A21). □

Appendix B. Additional Experimental Details

We implemented the dual problems for DRO-RoBAS (Corollary 1) in Python version 3.11 using CVXPY version 1.5.2 and the MOSEK solver version 10.1.28 following the implementation of Kernel DRO algorithms by Zhu et al. [6]. We used 200 discretisation points for the constraint, as described in Computation of (15). For the implementation of the NPL-MMD posterior, we follow the median heuristic suggested by Gretton et al. [29] and the ADAM optimiser [51] with learning rate $h = 0.1$. For the DP approximation we used $\alpha = 0$ and truncation limit $\tau = 100$. For all Bayesian methods we use $N = 900$ total Monte Carlo samples to approximate the expectations in the optimisation objectives.

The out-of-sample mean and variance were calculated as:

$$m(\epsilon) = \frac{1}{JT}\sum_{j=1}^{J}\sum_{t=1}^{T} f(x_\epsilon^{(j)}, \xi_{n+t}),$$
$$v(\epsilon) = \frac{1}{JT-1}\sum_{j=1}^{J}\sum_{t=1}^{T}\Big(f(x_\epsilon^{(j)}, \xi_{n+t}) - m(\epsilon)\Big)^2, \tag{A22}$$

where $x_\epsilon^{(j)}$ is the obtained solution on training dataset $\xi_{1:n}^{(j)}$ and we set $T = 50$ for the number of test observations.

Detailed results on the average solve and sampling times and associated standard deviations are provided in Tables A1 and A2.

Data-Generating Process Settings

For the Newsvendor problem, we first considered a Gaussian distribution with known variance, i.e., $\mathbb{P}_\theta := \mathcal{N}(\theta, \sigma^2\mathcal{I}_{D\times D})$ while the DGP is a bimodal Gaussian distribution: $\mathbb{P}^\star := 0.5\mathcal{N}(\theta_1^\star, \sigma^2\mathcal{I}_{D\times D}) + 0.5\mathcal{N}(\theta_2^\star, \sigma^2\mathcal{I}_{D\times D})$. We considered a univariate ($d = 1$) with $(\theta_1^\star, \theta_2^\star, \sigma) = (10, 60, 5)$ and a multivariate case ($d = 5$) with $\theta_1^\star = (10, 20, 33, 22, 25)$, $\theta_2^\star = (60, 60, 60, 60, 60)$ and $\sigma^2 = 5$. Furthermore, we explore the same Gaussian model $\mathbb{P}_\theta := \mathcal{N}(\theta, \sigma^2\mathcal{I}_{D\times D})$ with a contaminated Gaussian training DGP: $\mathbb{P}^\star_{\text{train}} := (1-\eta)\mathcal{N}(\theta^\star, \sigma^2) + \eta\mathcal{N}(\theta', \sigma^2)$ where we set $(\theta^\star, \theta', \sigma) = (25, 75, 5)$ and for $\eta \in \{0.0, 0.1, 0.2\}$ and an Exponential model $\mathbb{P}_\theta := \text{Exp}(\theta)$ with a contaminated Exponential DGP for the training DGP: $\mathbb{P}^\star_{\text{train}} := (1-\eta)\text{Exp}(0.05) + \eta\mathcal{N}(100, 0.5)$ $\mathbb{P}^\star_{\text{train}} := (1-\eta)\text{Exp}(\theta^\star) + \eta\mathcal{N}(\mu, \sigma)$ for $\eta \in \{0.0, 0.1, 0.2\}$.

Table A1. Average (standard derivation) *solve time* in seconds of algorithms across different DGPs.

DGP	RoBAS	Empirical MMD	KL-BDRO	DRO-BAS$_{\text{PE}}$	DRO-BAS$_{\text{PP}}$
1D $\mathcal{N}$ bimodal	275.33 (25.98)	2.17 (0.25)	3.90 (0.55)	0.29 (0.03)	0.30 (0.05)
5D $\mathcal{N}$ bimodal	262.94 (31.72)	2.18 (0.36)	2.63 (0.26)	1.24 (0.19)	1.38 (0.29)
$\mathcal{N}, \eta = 0.0$	286.27 (16.14)	2.22 (0.28)	4.01 (0.49)	0.29 (0.04)	0.30 (0.04)
$\mathcal{N}, \eta = 0.1$	284.62 (17.42)	2.23 (0.24)	3.99 (0.42)	0.29 (0.03)	0.30 (0.04)
$\mathcal{N}, \eta = 0.2$	260.67 (35.90)	2.20 (0.24)	3.74 (0.52)	0.28 (0.03)	0.28 (0.04)
Exp, $\eta = 0.0$	311.74 (37.48)	5.59 (0.68)	0.43 (0.06)	0.40 (0.06)	0.49 (0.06)
Exp, $\eta = 0.1$	296.65 (46.10)	5.39 (0.54)	0.44 (0.06)	0.42 (0.05)	0.50 (0.06)
Exp, $\eta = 0.2$	282.55 (53.23)	5.22 (0.47)	0.44 (0.06)	0.42 (0.05)	0.50 (0.06)

Table A2. Average (standard derivation) *sample time* in seconds of algorithms across different DGPs.

DGP	RoBAS	Empirical MMD	KL-BDRO	DRO-BAS$_{\text{PE}}$	DRO-BAS$_{\text{PP}}$
1D $\mathcal{N}$ bimodal	15.65 (4.76)	0.0 (0.0)	0.001 (0.002)	0.0002 (0.0)	0.0002 (0.0)
5D $\mathcal{N}$ bimodal	15.86 (4.84)	0.0 (0.0)	0.0104 (0.0004)	0.0006 (0.0003)	0.001 (0.0004)
$\mathcal{N}, \eta = 0.0$	15.85 (4.67)	0.0 (0.0)	0.0009 (0.002)	0.0002 (0.0)	0.0002 (0.0)
$\mathcal{N}, \eta = 0.1$	15.78 (4.59)	0.0 (0.0)	0.0009 (0.0017)	0.0002 (0.0)	0.0002 (0.0)
$\mathcal{N}, \eta = 0.2$	15.68 (4.58)	0.0 (0.0)	0.0008 (0.0014)	0.0002 (0.0)	0.0002 (0.0)
Exp, $\eta = 0.0$	38.10 (5.62)	0.0 (0.0)	0.1 (0.01)	0.0 (0.0)	0.0 (0.0)
Exp, $\eta = 0.1$	37.77 (4.78)	0.0 (0.0)	0.1 (0.01)	0.0 (0.0)	0.0 (0.0)
Exp, $\eta = 0.2$	37.82 (4.85)	0.0 (0.0)	0.1 (0.01)	0.0 (0.0)	0.0 (0.0)

Appendix C. Alternative RoBAS formulations

We now discuss alternative RoBAS formulations as the ones provided in Dellaporta et al. [12] through the DRO-BAS framework. In particular, in the DRO-BAS case the authors explored two distinct formulations of ambiguity sets: the BAS$_{\text{PP}}$ (3), mirroring

the definition of $\mathcal{B}_\epsilon^k(\mathbb{P}_n^{\text{pred(NPL)}})$, and the BAS_{PE} (4) which is based on a posterior expectation by considering the expected KL distance to the model family.

In the case of Robust Bayesian Ambiguity Sets, an interesting connection can be made between $\mathcal{B}_\epsilon(\mathbb{P}_n^{\text{pred(NPL)}})$ and the ambiguity set based on the expected squared MMD, namely:

$$\mathcal{A}_\epsilon^k(\text{DP}(c', \mathbb{F}')) = \left\{\mathbb{P} \in \mathcal{P}_k(\Xi) : \mathbb{E}_{\mathbb{Q}\sim\text{DP}(c',\mathbb{F}')}\left[\mathbb{D}_k^2(\mathbb{P}, \mathbb{P}_{\theta_k(\mathbb{Q})})\right]\right\}. \tag{A23}$$

The squared MMD has previously been used in Generalised Bayesian Inference to define a loss function by Chérief-Abdellatif and Alquier [49]. By using the properties of the MMD, we can prove the following equivalence:

Lemma A5. *Let* $\mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}$ *denote the expectation under* $\mathbb{E}_{\mathbb{Q}\sim\text{DP}(c',\mathbb{F}')}$ *and suppose:*

$$\begin{aligned} v(\text{DP}) &:= \mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}\left[\left\langle \mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}, \mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}\right\rangle\right] - \left\langle \mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}[\mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}], \mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}[\mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}]\right\rangle \\ &= \mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}\left[\left\|\mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}\right\|_{\mathcal{H}_k}^2\right] - \left\|\mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}[\mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}]\right\|_{\mathcal{H}_k}^2 \\ &< \infty. \end{aligned} \tag{A24}$$

Then for any $\mathbb{P} \in \mathcal{P}_k(\Xi)$*:*

$$\mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}\left[\mathbb{D}_k^2(\mathbb{P}, \mathbb{P}_{\theta_k(\mathbb{Q})})\right] = \mathbb{D}_k^2(\mathbb{P}, \mathbb{P}_n^{pred(NPL)}) + v(\text{DP}) \tag{A25}$$

and for any $\epsilon \geq 0$*:*

$$\mathcal{B}_\epsilon^k(\mathbb{P}_n^{pred(NPL)}) \equiv \mathcal{A}_{\epsilon^2+v(\text{DP})}^k(\text{DP}(c', \mathbb{F}')). \tag{A26}$$

Proof of Lemma A5. The proof is based on the reproducing property of the RKHS. Recall that $\mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}$ denotes the expectation under $\mathbb{E}_{\mathbb{Q}\sim\text{DP}(c',\mathbb{F}')}$. For any $\mathbb{P} \in \mathcal{P}_k(\Xi)$ and $\epsilon \geq 0$ we have:

$$\begin{aligned} &\mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}[\mathbb{D}_k^2(\mathbb{P}, \mathbb{P}_{\theta_k(\mathbb{Q})})] \\ &= \langle \mu_{\mathbb{P}}, \mu_{\mathbb{P}}\rangle_{\mathcal{H}_k} - 2\mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}\left[\left\langle \mu_{\mathbb{P}}, \mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}\right\rangle_{\mathcal{H}_k}\right] + \mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}\left[\left\langle \mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}, \mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}\right\rangle_{\mathcal{H}_k}\right] \\ &= \langle \mu_{\mathbb{P}}, \mu_{\mathbb{P}}\rangle_{\mathcal{H}_k} - 2\mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}\left[\left\langle \mu_{\mathbb{P}}, \mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}\right\rangle_{\mathcal{H}_k}\right] + \left\langle \mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}[\mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}], \mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}[\mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}]\right\rangle_{\mathcal{H}_k} \\ &\quad - \left\langle \mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}[\mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}], \mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}[\mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}]\right\rangle_{\mathcal{H}_k} + \mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}\left[\left\langle \mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}, \mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}\right\rangle_{\mathcal{H}_k}\right] \\ &= \mathbb{D}_k^2(\mathbb{Q}, \mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}[\mathbb{P}_{\theta_k(\mathbb{Q})}]) + v(\text{DP}). \end{aligned} \tag{A27}$$

This proves the statement of Equation (A25). To prove the statement of Equation (A26) recall that

$$\mathcal{B}_\epsilon^k(\mathbb{P}_n^{\text{pred(NPL)}}) := \{\mathbb{P} \in \mathcal{P}_k(\Xi) : \mathbb{D}_k(\mathbb{P}, \mathbb{P}_n^{\text{pred(NPL)}}) \leq \epsilon\} \tag{A28}$$

$$= \{\mathbb{P} \in \mathcal{P}_k(\Xi) : \mathbb{D}_k(\mathbb{P}, \mathbb{E}_{\text{DP}_{\tilde{\zeta}_{1:n}}}[\mathbb{P}_{\theta_k(\mathbb{Q})}]) \leq \epsilon\} \tag{A29}$$

Using (A27) we have that for $\epsilon \geq 0$ and some $\mathbb{P} \in \mathcal{B}_\epsilon^k(\mathbb{P}_n^{\text{pred(NPL)}})$:

$$\mathbb{D}_k(\mathbb{P}, \mathbb{E}_{\text{DP}_{\zeta_{1:n}}}[\mathbb{P}_{\theta_k(\mathbb{Q})}]) \leq \epsilon \tag{A30}$$

$$\iff \mathbb{D}_k^2(\mathbb{P}, \mathbb{E}_{\text{DP}_{\zeta_{1:n}}}[\mathbb{P}_{\theta_k(\mathbb{Q})}]) \leq \epsilon^2 \tag{A31}$$

$$\iff \mathbb{E}_{\text{DP}_{\zeta_{1:n}}}[\mathbb{D}_k^2(\mathbb{P}, \mathbb{P}_{\theta_k(\mathbb{Q})})] \leq \epsilon^2 + v(n) \tag{A32}$$

$$\iff \mathbb{P} \in \mathcal{A}^k_{\epsilon^2 + v(\text{DP})}(\text{DP}(c', \mathbb{F}')). \tag{A33}$$

hence, $\mathcal{B}_\epsilon^k(\mathbb{P}_n^{\text{pred(NPL)}}) \equiv \mathcal{A}^k_{\epsilon^2+v(\text{DP})}(\text{DP}(c', \mathbb{F}'))$, which completes the proof. □

The constant $v(\text{DP})$ (A24) is an effective variance term of the kernel mean embedding $\mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}$ in the RKHS under the nonparametric posterior $\text{DP} \equiv \text{DP}(c', \mathbb{F}')$. It follows directly from Lemma A5 that the primal DRO-RoBAS problem in (11) is equivalent to the following optimisation problem:

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{P} \,:\, \mathbb{E}_{\text{DP}_{\zeta_{1:n}}}[\mathbb{D}_k^2(\mathbb{P}, \mathbb{P}_{\theta_k(\mathbb{Q})})] \leq \epsilon^2 + v(\text{DP})} \mathbb{E}_{x \sim \mathbb{P}}\left[f(z, x)\right] \tag{A34}$$

or equivalently,

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{A}^k_{\epsilon^2 + v(\text{DP})}(\mathbb{P}_n^{\text{pred(NPL)}})} \mathbb{E}_{\xi \sim \mathbb{P}}[f_x(\xi)]. \tag{A35}$$

Since these two problems are equivalent, we refer to both of them as DRO-RoBAS.

An alternative formulation could be defined, based on the expected MMD distance, i.e., by considering the set

$$\left\{\mathbb{P} \in \mathcal{P}_k(\Xi) : \mathbb{E}_{\text{DP}_{\zeta_{1:n}}}\left[\mathbb{D}_k(\mathbb{P}, \mathbb{P}_{\theta_k(\mathbb{Q})})\right] \leq \epsilon\right\} \tag{A36}$$

for tolerance level $\epsilon \geq 0$. This mirrors the BAS_{PE} in Dellaporta et al. [12] and could lead to a differently shaped ambiguity set as it corresponds to an expected MMD ball. However, deriving the dual formulation for the corresponding worst-case DRO problem based on this set is challenging as it involves deriving the convex conjugate of the expected MMD distance $(\mathbb{E}_{\text{DP}_{\zeta_{1:n}}}[\mathbb{D}_k(\cdot, \mathbb{P}_{\theta_k(\mathbb{Q})})])^*$ or equivalently the support function $\delta^\star_{\mathcal{C}}(g)$ of the corresponding set $\mathcal{C} := \{\mu \in \mathcal{H}_k : \mathbb{E}_{\mathbb{Q} \sim \text{DP}(c', \mathbb{F}')}\left[\|\mu - \mu_{\mathbb{P}_{\theta_k(\mathbb{Q})}}\|_k\right] \leq \epsilon\}$. This is not trivial and it has not been tackled in this work, however it poses itself as important future work as it can open the way to ambiguity sets whose shape itself is informed by the posterior rather than just their nominal distribution as in RoBAS.

References

1. Black, B. *Parametric Distributionally Robust Optimisation Models for Resource and Inventory Planning Problems*; Lancaster University: Lancaster, UK, 2022.
2. Li, J.Y.; Kwon, R.H. Portfolio selection under model uncertainty: A penalized moment-based optimization approach. *J. Glob. Optim.* **2013**, *56*, 131–164. [CrossRef]
3. Zhang, J.; Menon, A.K.; Veit, A.; Bhojanapalli, S.; Kumar, S.; Sra, S. Coping with label shift via distributionally robust optimisation. In Proceedings of the International Conference on Learning Representations, Virtual, 3–7 May 2021.
4. Kuhn, D.; Esfahani, P.M.; Nguyen, V.A.; Shafieezadeh-Abadeh, S. Wasserstein Distributionally Robust Optimization: Theory and Applications in Machine Learning. In *Operations Research & Management Science in the Age of Analytics*; Informs: Catonsville, MD, USA, 2019; Chapter 6, pp. 130–166. [CrossRef]
5. Staib, M.; Jegelka, S. Distributionally robust optimization and generalization in kernel methods. *Adv. Neural Inf. Process. Syst.* **2019**, *32*, 9134–9144.

6. Zhu, J.J.; Jitkrittum, W.; Diehl, M.; Schölkopf, B. Kernel distributionally robust optimization: Generalized duality theorem and stochastic approximation. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*; PMLR: Cambridge, MA, USA, 2021; pp. 280–288.
7. Hu, Z.; Hong, L.J. Kullback-Leibler divergence constrained distributionally robust optimization. *Available Optim. Online* **2013**, *1*, 9.
8. Iyengar, G.; Lam, H.; Wang, T. Hedging against complexity: Distributionally robust optimization with parametric approximation. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*; PMLR: Cambridge, MA, USA, 2023; pp. 9976–10011.
9. Michel, P.; Hashimoto, T.; Neubig, G. Modeling the second player in distributionally robust optimization. In Proceedings of the International Conference on Learning Representations, Virtual, 3–7 May 2021.
10. Michel, P.; Hashimoto, T.; Neubig, G. Distributionally robust models with parametric likelihood ratios. In Proceedings of the International Conference on Learning Representations, Virtual, 26–28 October 2022.
11. Shapiro, A.; Zhou, E.; Lin, Y. Bayesian distributionally robust optimization. *SIAM J. Optim.* **2023**, *33*, 1279–1304. [CrossRef]
12. Dellaporta, C.; O'Hara, P.; Damoulas, T. Decision Making under the Exponential Family: Distributionally Robust Optimisation with Bayesian Ambiguity Sets. In Proceedings of the International Conference on Machine Learning, Vancouver, BC, Canada, 13–19 July 2025.
13. Grünwald, P. The safe Bayesian: Learning the learning rate via the mixability gap. In *Proceedings of the International Conference on Algorithmic Learning Theory*; Springer: Berlin/Heidelberg, Germany, 2012; pp. 169–183.
14. Walker, S.G. Bayesian inference with misspecified models. *J. Stat. Plan. Inference* **2013**, *143*, 1621–1633. [CrossRef]
15. Lyddon, S.; Walker, S.; Holmes, C.C. Nonparametric learning from Bayesian models with randomized objective functions. *Adv. Neural Inf. Process. Syst.* **2018**, *31*, 2075–2085.
16. Fong, E.; Lyddon, S.; Holmes, C. Scalable nonparametric sampling from multimodal posteriors with the posterior bootstrap. In *Proceedings of the International Conference on Machine Learning*; PMLR: Cambridge, MA, USA, 2019; pp. 1952–1962.
17. Rahimian, H.; Mehrotra, S. Distributionally robust optimization: A review. *arXiv* **2019**, arXiv:1908.05659.
18. Husain, H.; Nguyen, V.; van den Hengel, A. Distributionally Robust Bayesian Optimization with φ-divergences. In *Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2023; Volume 36.
19. Wang, S.; Wang, H.; Li, X.; Honorio, J. Learning against distributional uncertainty: On the trade-off between robustness and specificity. *IEEE J. Sel. Top. Signal Process.* **2025**, *19*, 1420–1435. [CrossRef]
20. Dellaporta, C.; Knoblauch, J.; Damoulas, T.; Briol, F.X. Robust Bayesian inference for simulator-based models via the MMD posterior bootstrap. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*; PMLR: Cambridge, MA, USA, 2022; pp. 943–970.
21. Bissiri, P.G.; Holmes, C.C.; Walker, S.G. A general framework for updating belief distributions. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **2016**, *78*, 1103–1130. [CrossRef]
22. Ghosh, A.; Basu, A. Robust Bayes estimation using the density power divergence. *Ann. Inst. Stat. Math.* **2016**, *68*, 413–437. [CrossRef]
23. Knoblauch, J.; Jewson, J.; Damoulas, T. An optimization-centric view on Bayes' rule: Reviewing and generalizing variational inference. *J. Mach. Learn. Res.* **2022**, *23*, 1–109.
24. Matsubara, T.; Knoblauch, J.; Briol, F.X.; Oates, C.J. Robust generalised Bayesian inference for intractable likelihoods. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **2022**, *84*, 997–1022. [CrossRef]
25. Altamirano, M.; Briol, F.X.; Knoblauch, J. Robust and scalable Bayesian online changepoint detection. In *Proceedings of the International Conference on Machine Learning*; PMLR: Cambridge, MA, USA, 2023; pp. 642–663.
26. Altamirano, M.; Briol, F.X.; Knoblauch, J. Robust and Conjugate Gaussian Process Regression. In *Proceedings of the International Conference on Machine Learning*; PMLR: Cambridge, MA, USA, 2024; pp. 1155–1185.
27. Laplante, W.; Altamirano, M.; Duncan, A.; Knoblauch, J.; Briol, F.X. Robust and conjugate spatio-temporal gaussian processes. *arXiv* **2025**, arXiv:2502.02450.
28. Müller, A. Integral probability metrics and their generating classes of functions. *Adv. Appl. Probab.* **1997**, *29*, 429–443. [CrossRef]
29. Gretton, A.; Borgwardt, K.M.; Rasch, M.J.; Schölkopf, B.; Smola, A. A kernel two-sample test. *J. Mach. Learn. Res.* **2012**, *13*, 723–773.
30. Briol, F.X.; Barp, A.; Duncan, A.B.; Girolami, M. Statistical inference for generative models with maximum mean discrepancy. *arXiv* **2019**, arXiv:1906.05944.
31. Gao, R. Finite-sample guarantees for Wasserstein distributionally robust optimization: Breaking the curse of dimensionality. *Oper. Res.* **2023**, *71*, 2291–2306. [CrossRef]
32. Chen, Y.; Kim, J.; Anderson, J. Distributionally robust decision making leveraging conditional distributions. In Proceedings of the IEEE Conference on Decision and Control, Cancun, Mexico, 6–9 December 2022; pp. 5652–5659.

33. Romao, L.; Hota, A.R.; Abate, A. Distributionally robust optimal and safe control of stochastic systems via kernel conditional mean embedding. In *Proceedings of the IEEE Conference on Decision and Control*; IEEE: New York, NY, USA, 2023; pp. 2016–2021.
34. Gao, R.; Kleywegt, A. Distributionally robust stochastic optimization with Wasserstein distance. *Math. Oper. Res.* **2023**, *48*, 603–655. [CrossRef]
35. Vázquez, F.G.; Rückmann, J.J.; Stein, O.; Still, G. Generalized semi-infinite programming: A tutorial. *J. Comput. Appl. Math.* **2008**, *217*, 394–419. [CrossRef]
36. Chérief-Abdellatif, B.E.; Alquier, P. Finite sample properties of parametric MMD estimation: Robustness to misspecification and dependence. *Bernoulli* **2022**, *28*, 181–213. [CrossRef]
37. Alquier, P.; Gerber, M. Universal robust regression via maximum mean discrepancy. *Biometrika* **2024**, *111*, 71–92. [CrossRef]
38. Seeger, M. Gaussian processes for machine learning. *Int. J. Neural Syst.* **2004**, *14*, 69–106. [CrossRef]
39. Mohajerin Esfahani, P.; Kuhn, D. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Math. Program.* **2018**, *171*, 115–166. [CrossRef]
40. Bertsimas, D.; Gupta, V.; Kallus, N. Robust saa. *arXiv* **2014**, arXiv:1408.4445.
41. Huber, P.J. Robust estimation of a location parameter. In *Breakthroughs in Statistics: Methodology and Distribution*; Springer: Berlin/Heidelberg, Germany, 1992; pp. 492–518.
42. Liu, J.; Shen, Z.; He, Y.; Zhang, X.; Xu, R.; Yu, H.; Cui, P. Towards out-of-distribution generalization: A survey. *arXiv* **2021**, arXiv:2108.13624.
43. Zhai, R.; Dan, C.; Kolter, Z.; Ravikumar, P. Doro: Distributional and outlier robust optimization. In *Proceedings of the International Conference on Machine Learning*; PMLR: Cambridge, MA, USA, 2021; pp. 12345–12355.
44. Nietert, S.; Goldfeld, Z.; Shafiee, S. Outlier-robust wasserstein dro. *Adv. Neural Inf. Process. Syst.* **2023**, *36*, 62792–62820.
45. Porteus, E.L. Stochastic inventory theory. *Handb. Oper. Res. Manag. Sci.* **1990**, *2*, 605–652.
46. Rahimi, A.; Recht, B. Random features for large-scale kernel machines. *Adv. Neural Inf. Process. Syst.* **2007**, *20*, 1177–1184.
47. Bach, F. Sharp analysis of low-rank kernel matrix approximations. In *Proceedings of the Conference on Learning Theory*; PMLR: Cambridge, MA, USA, 2013; pp. 185–209.
48. Jewson, J.; Smith, J.Q.; Holmes, C. Principles of Bayesian inference using general divergence criteria. *Entropy* **2018**, *20*, 442. [CrossRef]
49. Chérief-Abdellatif, B.E.; Alquier, P. MMD-Bayes: Robust Bayesian estimation via maximum mean discrepancy. In *Proceedings of the Symposium on Advances in Approximate Bayesian Inference*; PMLR: Cambridge, MA, USA, 2020; pp. 1–21.
50. McDiarmid, C. On the method of bounded differences. *Surv. Comb.* **1989**, *141*, 148–188.
51. Kingma, D.P. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.