

Wasserstein Distributionally Robust Optimization with Heterogeneous Data Sources

Yves Rychener

Risk Analytics and Optimization Laboratory, École polytechnique fédérale de Lausanne, Lausanne CH-1015, Switzerland
yves.rychener@epfl.ch

Adrián Esteban-Pérez

Department of Technology and Operations Management, Rotterdam School of Management, Erasmus University, Rotterdam
3062 PA, The Netherlands estebanperez@rsm.nl

Juan M. Morales

Department of Mathematical Analysis, Statistics and Operations Research and Applied Mathematics (Area of Statistics and Operations Research), University of Málaga, Málaga 29071, Spain juan.morales@uma.es

Daniel Kuhn

Risk Analytics and Optimization Laboratory, École polytechnique fédérale de Lausanne, Lausanne CH-1015, Switzerland
daniel.kuhn@epfl.ch

We study decision problems under uncertainty, where the decision-maker has access to K data sources that carry *biased* information about the underlying risk factors. The biases are measured by the mismatch between the risk factor distribution and the K data-generating distributions with respect to an optimal transport (OT) distance. In this situation the decision-maker can exploit the information contained in the biased samples by solving a distributionally robust optimization (DRO) problem, where the ambiguity set is defined as the intersection of K OT neighborhoods, each of which is centered at the empirical distribution on the samples generated by a biased data source. We show that if the decision-maker has a prior belief about the biases, then the out-of-sample performance of the DRO solution can improve with K —irrespective of the magnitude of the biases. We also show that, under standard convexity assumptions, the proposed DRO problem is computationally tractable if either K or the dimension of the risk factors is kept constant.

Key words: data-driven decision-making; distributionally robust optimization; optimal transport

1. Introduction

Stochastic optimization is the canonical framework for modeling decision problems under uncertainty (Shapiro et al. 2021). A basic single-stage stochastic program seeks a decision $\theta \in \Theta \subseteq \mathbb{R}^n$ that minimizes the expected value $\mathbb{E}_{\mathbb{P}}[\ell(\theta, \xi)]$ of an uncertainty-affected loss function $\ell(\theta, \xi)$ with respect

to the distribution \mathbb{P} of the random vector $\xi \in \Xi \subseteq \mathbb{R}^d$. However, in most real decision problems, the distribution \mathbb{P} is unobservable, implying that an essential input of the stochastic optimization model is unknown. When the decision-maker has access to training samples $\{\widehat{\xi}_i\}_{i=1}^N$ from \mathbb{P} , then the stochastic program can be addressed with the sample average approximation (SAA) (Shapiro et al. 2021, § 5), which minimizes the expected loss $\mathbb{E}_{\widehat{\mathbb{P}}}[\ell(\theta, \xi)]$ under the empirical distribution $\widehat{\mathbb{P}} = \frac{1}{N} \sum_{i=1}^N \delta_{\widehat{\xi}_i}$, or with methods from distributionally robust optimization (DRO) (Rahimian and Mehrotra 2022), which minimize the worst-case expected loss $\sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[\ell(\theta, \xi)]$ with respect to all distributions in an ambiguity set \mathcal{P} . Intuitively, \mathcal{P} should contain all distributions that are sufficiently likely to have generated the training samples. A popular choice is to define \mathcal{P} as a ball around $\widehat{\mathbb{P}}$ with respect to a Wasserstein distance (Mohajerin Esfahani and Kuhn 2018, Zhao and Guan 2018, Blanchet and Murthy 2019, Kuhn et al. 2019, Gao and Kleywegt 2023). If the radius of this ball scales with $N^{-\frac{1}{2}}$, then the optimal value of the DRO problem constitutes, in a precise statistical sense, a safe estimator for the optimal value of the underlying true stochastic program (Blanchet et al. 2021, Gao 2023). The optimal solution of the SAA problem typically performs well *in sample* (that is, under the empirical distribution $\widehat{\mathbb{P}}$) but poorly out of sample (that is, under the unknown true distribution \mathbb{P}). This phenomenon is often referred to as the optimizer’s curse (Smith and Winkler 2006). Empirical evidence suggests that DRO mitigates the optimizer’s curse.

Unfortunately, all existing approaches to data-driven optimization become ineffective at small sample sizes N . This is troubling because training data is scarce or even absent in many relevant decision problems under uncertainty. For example, a manufacturer introducing a new product has no historical sales data to predict the product’s demand distribution, which would help to plan production. Similarly, an investor trading in new securities has no historical market data to predict the securities’ return distribution, which would help to design a portfolio strategy. In addition, a retailer expanding into a new market has no historical data on customer behavior to inform assortment planning, a logistics manager has no historical data on resource availabilities after a supply chain disruption to inform inventory decisions, and government authorities have no historical data on

how people react to a new pandemic, which would help to determine effective health interventions. Broadly speaking, whenever an organization faces structural disruptions or implements strategic changes, it enters into a new regime characterized by a lack of relevant data.

In the remainder of the paper we refer to the distribution \mathbb{P} of the uncertain problem parameters ξ as the *target distribution*, and we assume that there is insufficient (or no) data from \mathbb{P} . In this situation, one could try to leverage data from one or several *source distributions* \mathbb{P}_k , $k \in [K]$, that are not too dissimilar from \mathbb{P} . For example, a retailer expanding into a new market with an unknown demand distribution may have access to demand data from other markets with a comparable customer structure. Similarly, an investor trading in new securities with an unknown return distribution may have access to return data from other securities issued by similar companies.

Suppose from now on that there are $N_k \geq 1$ training samples from \mathbb{P}_k , which can be used to construct an empirical distribution $\widehat{\mathbb{P}}_k$ that approximates \mathbb{P}_k , $k \in [K]$. In addition, assume temporarily that \mathbb{P}_1 coincides with the target distribution \mathbb{P} (implying that $\widehat{\mathbb{P}}_1$ only suffers from statistical errors but is *not* subject to a distribution shift), whereas \mathbb{P}_k for $k \geq 2$ strictly differs from \mathbb{P} (implying that $\widehat{\mathbb{P}}_k$ suffers both from statistical errors and a distribution shift). We now discuss different approaches for estimating the optimal value of the true stochastic program. A naïve approach would be to ignore all biased training samples and to solve the SAA problem under $\widehat{\mathbb{P}}_1$ only. As we assumed the target data to be scarce, however, this results in a fragile estimator with an unacceptably high variance. Alternatively, one could rely on data pooling, which amounts to solving the SAA problem under the mixture distribution $\sum_{k=1}^K N_k / (\sum_{k'=1}^K N_{k'}) \widehat{\mathbb{P}}_k$. If there is abundant source data, the resulting estimator displays a small variance. As the pooled dataset is dominated by biased samples, however, this estimator typically suffers from an unacceptably large bias. In particular, it is likely to overestimate the variance of \mathbb{P} (see Example 1 below). A better—geometry-aware—method for aggregating the empirical source distributions is to calculate their Wasserstein barycenter (Agueh and Carlier 2011). Recall that the barycenter $\bar{\xi} = \sum_{k=1}^K \lambda_k \xi_k$ of K points $\xi_k \in \mathbb{R}^d$ with weights $\lambda_k \geq 0$, $k \in [K]$, is the unique solution of $\min_{\xi \in \mathbb{R}^d} \sum_{k=1}^K \lambda_k \|\xi - \xi_k\|_2^2$. Analogously, the

p -Wasserstein barycenter of K distributions $\widehat{\mathbb{P}}_k$ with weights $\lambda_k \geq 0$, $k \in [K]$, is defined as a solution $\bar{\mathbb{P}}$ of $\min_{\mathbb{P} \in \mathcal{P}(\Xi)} \sum_{k=1}^K \lambda_k W_p(\mathbb{P}, \widehat{\mathbb{P}}_k)^p$, where W_p denotes the p -Wasserstein distance for some $p \geq 1$. Wasserstein barycenters are widely used in machine learning (Claici et al. 2018, Dognin et al. 2019, Montesuma and Mboula 2021, Srivastava et al. 2018, Zhuang et al. 2022) and computer vision (Barré et al. 2020, Solomon et al. 2015). The statistical properties of empirical Wasserstein barycenters are studied in (Le Gouic and Loubes 2017, Le Gouic et al. 2022, Panaretos and Zemel 2020), and efficient methods for their computation are proposed in (Altschuler and Boix-Adsera 2021, Cuturi and Doucet 2014). Any p -Wasserstein barycenter $\bar{\mathbb{P}}$ of the empirical source distributions is *discrete* and can thus be identified with an aggregated dataset. Solving the SAA problem under $\bar{\mathbb{P}}$ yields a more promising estimator than data pooling. The following example suggests indeed that, unlike naïve mixtures, Wasserstein barycenters preserve stylized features of the source distributions.

EXAMPLE 1 (MIXTURES VERSUS WASSERSTEIN BARYCENTERS). Assume that all source distributions are univariate Gaussians with different means and identical variances, that is, $\mathbb{P}_k \sim \mathcal{N}(\mu_k, \sigma^2)$ for all $k \in [K]$. If $\sum_{k=1}^K \lambda_k = 1$ and $\lambda_k \geq 0$ for all $k \in [K]$, then the mixture distribution $\sum_{k=1}^K \lambda_k \mathbb{P}_k$ fails to be Gaussian, and it has mean $\sum_{k=1}^K \lambda_k \mu_k$ and variance $\sigma^2 + \sum_{k=1}^K \lambda_k \mu_k^2 - (\sum_{k=1}^K \lambda_k \mu_k)^2 > \sigma^2$. Using (Gelbrich 1990, Theorem 2.1), one can further show that the 2-Wasserstein barycenter of the source distributions with weights λ_k , $k \in [K]$, is the Gaussian distribution with mean $\sum_{k=1}^K \lambda_k \mu_k$ and variance σ^2 . As the source distributions are supposed to capture stylized features of the target distribution \mathbb{P} , and as all source distributions are Gaussians and share the same variance σ^2 , it is reasonable to expect that \mathbb{P} is a Gaussian distribution with variance σ^2 . Thus, the 2-Wasserstein barycenter is the more plausible model for \mathbb{P} than the mixture of the source distributions. \blacktriangle

Despite the encouraging insights from Example 1, the 2-Wasserstein barycenter of two empirical source distributions is highly sensitive to data perturbation (Zhuang et al. 2022, Example 2). In addition, its variance provides only a *biased* estimator for the variance of the 2-Wasserstein barycenter of the true source distributions (see Proposition 2 below). Even worse, the 1-Wasserstein barycenter either coincides with one of the source distributions or is highly degenerate (see Corollary 1 below).

So far we have discussed several SAA models, which differ only in how the source data enters the empirical distribution. Recall that SAA methods are generally susceptible to the optimizer’s curse and may thus lead to disappointment in out-of-sample tests. Following Mohajerin Esfahani and Kuhn (2018), it is therefore natural to robustify each of these SAA models against all distributions in a Wasserstein ball centered at the corresponding empirical distribution. If this empirical distribution converges (in an appropriate sense) to the target distribution, then the resulting DRO models offer the usual statistical guarantees; see *e.g.*, (Kuhn et al. 2019, § 3). In particular, if the target distribution \mathbb{P} is known to coincide with the 2-Wasserstein barycenter of the source distributions \mathbb{P}_k with weights λ_k , $k \in [K]$, and if the empirical distribution is set to the 2-Wasserstein barycenter of the *empirical* source distributions $\widehat{\mathbb{P}}_k$ with the *same* weights λ_k , $k \in [K]$, then the optimal value of the DRO problem provides an upper confidence bound on the optimal value of the true stochastic program (Lau and Liu 2022, Theorem 4.7). However, this guarantee is only available if the weights λ_k , $k \in [K]$, are known. In addition, the ambiguity set inherits all shortcomings of the Wasserstein barycenter at its center—such as outlier sensitivity, degeneracy and biasedness.

In view of the challenges outlined above, we propose a new approach to data-driven decision-making with $K > 1$ data sources. Recall that any useful source distribution \mathbb{P}_k , $k \in [K]$, must be close to the target distribution \mathbb{P} . It is thus reasonable to assume that the decision-maker has a belief about the discrepancy between \mathbb{P} and \mathbb{P}_k with respect to some Wasserstein distance. As the empirical distribution $\widehat{\mathbb{P}}_k$ approximates \mathbb{P}_k , any such belief naturally translates into an upper bound on the Wasserstein distance between \mathbb{P} and $\widehat{\mathbb{P}}_k$. Hence, the target distribution should reside in the intersection of K Wasserstein balls, each of which is centered at one of the empirical source distributions. This prompts us to introduce the class of *multi-source DRO* models whose ambiguity sets are obtained by intersecting *multiple* Wasserstein balls associated with *different data sources*.

The main contributions of this paper can be summarized as follows.

- We introduce multi-source DRO and show that it is dually related to the problem of computing the Wasserstein barycenter of the K empirical source distributions.

- We show that if the loss function is piecewise concave in $\xi \in \mathbb{R}^d$, then the worst-case expected loss over the intersection of the Wasserstein balls associated with the K data sources matches the optimal value of a finite convex program. However, this convex program involves exponentially many variables and constraints and is NP-hard. Leveraging recent results on Wasserstein barycenters by Altschuler and Boix-Adsera (2021), we then prove that it becomes tractable if d or K is constant.

- We develop frequentist as well as Bayesian performance guarantees for multi-source DRO. Specifically, we provide guidance for calibrating the Wasserstein radii such that the multi-source DRO problem provides an upper confidence bound on the stochastic program under the target distribution. We find that a *frequentist* decision-maker without any information about the distances between the target and the source distributions does *not* benefit from source data. In contrast, we show that a Bayesian decision-maker with prior beliefs about the distances between the target and the source distributions can obtain stronger performance guarantees by using source data.

- Numerical experiments focusing on portfolio selection and assortment planning show that multi-source DRO can outperform various single-source DRO schemes on synthetic as well as real data.

Our tractability results imply that solving a multi-source DRO problem is not fundamentally harder than computing the Wasserstein barycenter of the source distributions. Thus, our multi-source DRO problems belong to the same complexity class as the single-source DRO problems by Lau and Liu (2022), which use the Wasserstein barycenter of the source distributions as the nominal distribution. However, multi-source DRO is less sensitive to perturbations of the data.

Intersections of $K = 2$ ambiguity sets constructed from source and target data were first studied in the context of distributionally robust linear regression (Taskesen et al. 2021). This work critically relies on the quadratic nature of the least-squares loss, which implies that only the first and second moments of the covariates and the responses are needed to evaluate the expected prediction loss. More recently and independently of our work, intersections of $K = 2$ Wasserstein ambiguity sets were studied in the context of logistic regression (Awasthi et al. 2022). The dual DRO problem derived in this work is a special case of our Theorem 2. As the logistic loss fails to be (piecewise)

concave, this dual problem does not admit an exact reformulation as a finite convex program. Selvi et al. (2024) study a generalization of the model by Awasthi et al. (2022) that not only accounts for $K = 2$ different data sources but also for adversarial attacks. They leverage tools from adjustable robust optimization to solve their model approximately. Finally, Tanoumand et al. (2023) explore the intersection of a Wasserstein ball with an ambiguity set derived from a goodness-of-fit test based on the linear-convex ordering of random vectors. Their rationale for intersecting ambiguity sets is to reduce the conservativeness of the DRO solutions. However, they construct both ambiguity sets from a single dataset, and therefore they do not cater for multi-source DRO.

Decision-making under uncertainty with multiple data sources is reminiscent of domain adaptation in machine learning (Zhao et al. 2020, Farahani et al. 2021, Montesuma and Mboula 2021) and also resembles federated learning with heterogeneous clients (Mohri et al. 2019, Ro et al. 2021). Domain adaptation leverages data from one or several source distributions to improve the performance of a machine learning model on the target distribution (Singhal et al. 2023). This can be achieved, for example, by using tools from optimal transport. Specifically, an optimal transport map between the source and the target distribution can be used to convert a predictor trained on the source domain to one on the target domain (Courty et al. 2014, 2016, Redko et al. 2017). However, these methods are tailored to prediction tasks and cannot be readily extended to decision-making tasks.

Heterogeneous data sources are also common in federated learning, where different devices share data to train a machine learning model in a distributed manner—often under privacy constraints (see, *e.g.*, (Zhang et al. 2021) for a survey). Distributionally robust federated learning models trained on a mixture of the client distributions with uncertain mixture weights are studied in (Deng et al. 2020, Sagawa et al. 2019, Xiong et al. 2023, Wang et al. 2023). A distributionally robust model with a Wasserstein ball around the empirical mixture distribution is proposed in (Nguyen et al. 2022).

Data-driven optimization with heterogeneous data sources has not yet received much attention in operations research and management science. It is known that if many unrelated data-driven decision problems must be solved simultaneously, each having only access to a small amount of data,

then shrinking the empirical distribution of each individual problem towards an anchor distribution constructed from the pooled data can improve upon naïve SAA approaches (Gupta and Kallus 2022). However, as all empirical distributions are shrunk towards the same anchor, this approach is difficult to justify in the presence of distribution shifts. Also, we focus on a single decision problem.

The rest of the paper is structured as follows. Section 2 introduces optimal transport barycenters as natural generalizations of Wasserstein barycenters. It further shows that 1-Wasserstein barycenters are ill-conceived for practical use and that empirical 2-Wasserstein barycenters are biased. Section 3 formally introduces the multi-source DRO problem and shows that it is equivalent to a finite convex program of exponential size if the loss function is piecewise concave. Section 4 shows that this convex program is tractable if either d or K is constant. Section 5 establishes statistical performance guarantees for multi-source DRO, and Section 6 reports on numerical experiments.

Notation. For any Borel set $\Xi \subseteq \mathbb{R}^d$, we use $\mathcal{M}_+(\Xi)$, $\mathcal{P}(\Xi)$ and $\mathcal{P}_p(\Xi)$ to denote the convex cone of all finite Borel measures on Ξ , the convex set of all probability distributions in $\mathcal{M}_+(\Xi)$ and the convex set of all probability distributions in $\mathcal{P}(\Xi)$ with finite p -th moment, respectively. For any $\mathbb{P} \in \mathcal{P}(\Xi)$ and Borel measurable transformation $\Phi : \Xi \rightarrow \Xi'$ between Borel sets $\Xi \subseteq \mathbb{R}^d$ and $\Xi' \subseteq \mathbb{R}^{d'}$, we denote by $\Phi_{\#}\mathbb{P}$ the pushforward distribution of \mathbb{P} under Φ . Thus, if $\xi \in \Xi$ follows \mathbb{P} , then $\Phi(\xi) \in \Xi'$ follows $\Phi_{\#}\mathbb{P}$. The conjugate f^* of a convex function $f : \mathbb{R}^d \rightarrow [-\infty, \infty]$ is defined through $f^*(u) = \sup_{\xi \in \mathbb{R}^d} \langle u, \xi \rangle - f(\xi)$. The indicator function δ_{Ξ} of a closed convex set $\Xi \subseteq \mathbb{R}^d$ is defined through $\delta_{\Xi}(\xi) = 0$ if $\xi \in \Xi$, $\delta_{\Xi}(\xi) = \infty$ otherwise. The conjugate of δ_{Ξ} is termed the *support function* of Ξ and is denoted by σ_{Ξ} . The *perspective* of a closed convex function $f(\xi)$ is defined as $tf(\xi/t)$ for $t \geq 0$, where $0f(\xi/0)$ is interpreted as $\sigma_{\text{dom}(f^*)}(\xi)$. For any $n \in \mathbb{N}$ we set $[n] = \{1, \dots, n\}$.

2. Multi-Margin OT and OT Barycenters

This section introduces OT barycenters as natural generalizations of Wasserstein barycenters and shows that they can be computed by solving multi-margin OT problems. Throughout the paper, we let $\Xi \subseteq \mathbb{R}^d$ be a convex closed support set. The family of multi-margin transportation plans associated with K probability distributions $\mathbb{P}_1, \dots, \mathbb{P}_K \in \mathcal{P}(\Xi)$ is defined as

$$\Pi(\mathbb{P}_1, \dots, \mathbb{P}_K) = \left\{ \pi \in \mathcal{P}(\Xi^K) : \pi \text{ has marginals } \mathbb{P}_1, \dots, \mathbb{P}_K, \text{ respectively} \right\}$$

(Pass 2015). If $K = 2$, then any $\pi \in \Pi(\mathbb{P}_1, \mathbb{P}_2)$ is simply referred to as a transportation plan or a coupling of \mathbb{P}_1 and \mathbb{P}_2 . Optimal transport (OT) seeks a transportation plan of minimum cost with respect to a cost function $c: \Xi \times \Xi \rightarrow \mathbb{R}_+$. We adopt the following standard assumption regarding c .

ASSUMPTION 1 (Transportation Cost Function). *The function $c(\xi_1, \xi_2)$ is lower semi-continuous, obeys the identity of indiscernibles (i.e., $c(\xi_1, \xi_2) = 0$ if and only if $\xi_1 = \xi_2$) and is lower bounded by $d^p(\xi_1, \xi_2)$ for some metric d on Ξ with compact sublevel sets and some exponent $p \in \mathbb{N}$.*

The OT cost associated with two distributions $\mathbb{P}_1, \mathbb{P}_2 \in \mathcal{P}(\Xi)$ is then defined as

$$C(\mathbb{P}_1, \mathbb{P}_2) = \min_{\pi \in \Pi(\mathbb{P}_1, \mathbb{P}_2)} \int_{\Xi^2} c(\xi_1, \xi_2) d\pi(\xi_1, \xi_2).$$

Intuitively, the OT problem on the right-hand side of the above expression seeks the cheapest transportation plan π for morphing \mathbb{P}_1 into \mathbb{P}_2 , where the cost of moving unit mass from ξ_1 to ξ_2 amounts to $c(\xi_1, \xi_2)$. The OT problem is guaranteed to be solvable under Assumption 1 (Villani 2009, Theorem 4.1). If $c(\xi_1, \xi_2) = d^p(\xi_1, \xi_2)$ for some metric d on Ξ and some exponent $p \in \mathbb{N}$, then $C(\mathbb{P}_1, \mathbb{P}_2)^{1/p}$ is termed the p -Wasserstein distance and is denoted by $W_p(\mathbb{P}_1, \mathbb{P}_2)$.

We are now ready to define the OT barycenter of multiple distributions. It generalizes the notion of the Wasserstein barycenter introduced by Agueh and Carlier (2011).

DEFINITION 1 (OT Barycenter). *An OT Barycenter \mathbb{P}^* of K distributions $\mathbb{P}_1, \dots, \mathbb{P}_K \in \mathcal{P}(\Xi)$ with associated weights $\lambda_1, \dots, \lambda_K \geq 0$ is any solution of the minimization problem*

$$\min_{\mathbb{P} \in \mathcal{P}(\Xi)} \sum_{k=1}^K \lambda_k C(\mathbb{P}, \mathbb{P}_k). \quad (\text{OT-BC})$$

If $C(\mathbb{P}, \mathbb{P}_k) = W_p^p(\mathbb{P}, \mathbb{P}_k)$, then the OT barycenters are referred to as p -Wasserstein barycenters.

ASSUMPTION 2 (OT Cost). *There exists $\xi_0 \in \Xi$ with $\int_{\Xi} c(\xi_0, \xi_k) d\mathbb{P}_k(\xi_k) < \infty$ for all $k \in [K]$.*

The following lemma shows that OT barycenters must exist under Assumptions 1 and 2.

LEMMA 1 (OT Barycenters). *Assumptions 1 and 2 ensure that problem (OT-BC) is solvable.*

Proof of Lemma 1 If $\lambda_k = 0$ for every $k \in [K]$, then any $\mathbb{P} \in \mathcal{P}(\Xi)$ solves (OT-BC), and thus the claim trivially holds. From now on we may thus assume without loss of generality that $\lambda_1 > 0$. In this case, the objective function of (OT-BC) is bounded below by $\lambda_1 C(\mathbb{P}, \mathbb{P}_1)$. Assumption 2 further implies that $\mathbb{P} = \delta_{\xi_0}$ is feasible in (OT-BC), and thus $\bar{C} = \sum_{k=1}^K \lambda_k C(\delta_{\xi_0}, \mathbb{P}_k) < \infty$ provides an upper bound on the optimal value of (OT-BC). Taken together, these insights reveal that we can restrict the feasible set of problem (OT-BC) to $B = \{\mathbb{P} \in \mathcal{P}(\Xi) : C(\mathbb{P}, \mathbb{P}_1) \leq \bar{C}/\lambda_1\}$ without affecting its optimal value. The set B is weakly compact by (Shafieezadeh Abadeh et al. 2023, Proposition 2.5). Furthermore, $C(\mathbb{P}, \mathbb{P}_k)$ is weakly lower semi-continuous in \mathbb{P} for every $k \in [K]$ (Clément and Desch 2008, Lemma 5.2), and thus the objective function of problem (OT-BC) is weakly lower semi-continuous. The minimum of (OT-BC) is thus attained thanks to Weierstrass' theorem. ■

The next theorem shows that OT barycenters can be found by solving multi-margin OT problems. It extends an existing result for 2-Wasserstein barycenters (Agueh and Carlier 2011, Anderes et al. 2016) to general OT barycenters. We include a short proof to keep the paper self-contained.

THEOREM 1 (OT Barycenters and Multi-Margin Transportation Plans). *Suppose that Assumptions 1 and 2 hold, and define $\phi : \Xi^K \rightarrow [0, \infty]$ and $\Phi : \Xi^K \rightarrow \Xi$ as the optimal value function and a measurable selector of the set-valued solution map of the minimization problem*

$$\min_{\xi \in \Xi} \sum_{k=1}^K \lambda_k c(\xi, \xi_k) \quad (1)$$

parametrized by $\xi_1, \dots, \xi_K \in \Xi$. Then there is an OT barycenter of $\mathbb{P}_1, \dots, \mathbb{P}_K \in \mathcal{P}(\Xi)$ representable as $\mathbb{P}^ = \Phi_{\#} \pi^*$, where π^* is an optimal solution of the multi-margin OT problem*

$$\min_{\pi \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_K)} \int_{\Xi^K} \phi(\xi_1, \dots, \xi_K) d\pi(\xi_1, \dots, \xi_K). \quad (2)$$

If the measurable selector Φ is unique, then every OT barycenter is representable in this form.

Proof of Theorem 1 As in Lemma 1, we may assume without loss of generality that $\lambda_1 > 0$. By Assumption 1, the transportation cost function c is lower semi-continuous and has bounded sublevel sets. The minimum of (1) is thus attained by Weierstrass' Theorem. In addition, the argmin

multifunction of problem (1) admits a measurable selector thanks to (Rockafellar and Wets 2009, Corollary 14.6 and Theorem 14.37). Thus, the functions ϕ and Φ are well-defined.

Assumptions 1 and 2 imply via Lemma 1 that (OT-BC) is solvable. In the following we will construct a sequence of optimization problems that are all equivalent to (OT-BC). As usual, we say that two optimization problems are equivalent if any feasible solution of the first problem can be used to construct a feasible solution to the second problem with the same (or a smaller) objective function value and vice versa. This correspondence holds in particular for all optimal solutions (if they exist). Thus, if one of the two equivalent problems is solvable, then so is the other one.

In a first step we note that problem (OT-BC) is equivalent to

$$\begin{aligned} \min \quad & \sum_{k=1}^K \lambda_k \int_{\Xi^2} c(\xi, \xi_k) d\pi_k(\xi, \xi_k) \\ \text{s.t.} \quad & \mathbb{P} \in \mathcal{P}(\Xi), \pi_k \in \Pi(\mathbb{P}, \mathbb{P}_k) \quad \forall k \in [K]. \end{aligned} \tag{BC-1}$$

To see this, assume first that \mathbb{P} is feasible in (OT-BC). By (Villani 2009, Theorem 4.1), there exists an optimal transportation plan $\pi_k \in \Pi(\mathbb{P}, \mathbb{P}_k)$ with $C(\mathbb{P}, \mathbb{P}_k) = \int_{\Xi^2} c(\xi, \xi_k) d\pi_k(\xi, \xi_k)$ for every $k = 1, \dots, K$. By the definition of the OT cost, it is then clear that $(\mathbb{P}, \pi_1, \dots, \pi_K)$ is feasible in (BC-1) and attains the same objective function value as \mathbb{P} in (OT-BC). Conversely, if $(\mathbb{P}, \pi_1, \dots, \pi_K)$ is feasible in (BC-1), then \mathbb{P} is feasible in (OT-BC) with the same objective function value.

In a second step we prove that problem (BC-1) is equivalent to

$$\begin{aligned} \min \quad & \int_{\Xi^{K+1}} \sum_{k=1}^K \lambda_k c(\xi, \xi_k) d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) \\ \text{s.t.} \quad & \mathbb{P} \in \mathcal{P}(\Xi), \bar{\pi} \in \Pi(\mathbb{P}, \mathbb{P}_1, \dots, \mathbb{P}_K). \end{aligned} \tag{BC-2}$$

To see this, assume first that $(\mathbb{P}, \pi_1, \dots, \pi_K)$ is feasible in (BC-1). An iterative application of the gluing lemma (Villani 2009, Chapter 1) implies that the K couplings $\pi_k \in \Pi(\mathbb{P}, \mathbb{P}_k), k \in [K]$, can be merged to a multi-margin transportation plan $\bar{\pi} \in \Pi(\mathbb{P}, \mathbb{P}_1, \dots, \mathbb{P}_K)$ satisfying $P_{0,k\#} \bar{\pi} = \pi_k$, where the truncation operator $P_{0,k} : \Xi \times \Xi^k \rightarrow \Xi^2$ is defined through $P_{0,k}(\xi, \xi_1, \dots, \xi_K) = (\xi, \xi_k)$ for every $k \in [K]$. By the measure-theoretic change-of-variables formula, we thus have

$$\sum_{k=1}^K \lambda_k \int_{\Xi^2} c(\xi, \xi_k) d\pi_k(\xi, \xi_k) = \int_{\Xi^{K+1}} \sum_{k=1}^K \lambda_k c(\xi, \xi_k) d\bar{\pi}(\xi, \xi_1, \dots, \xi_K).$$

Hence, $(\mathbb{P}, \bar{\pi})$ is feasible in (BC-2) and attains the same objective function value as $(\mathbb{P}, \pi_1, \dots, \pi_K)$. Conversely, if $(\mathbb{P}, \bar{\pi})$ is feasible in (BC-2), then $(\mathbb{P}, \pi_1, \dots, \pi_K)$ with $\pi_k = P_{0,k\#} \bar{\pi}$ for every $k \in [K]$ is feasible in (BC-1) and attains the same objective function value as $(\mathbb{P}, \bar{\pi})$.

Next, we show that problem (BC-2) is equivalent to

$$\begin{aligned} \min \quad & \int_{\Xi^K} \int_{\Xi} \sum_{k=1}^K \lambda_k c(\xi, \xi_k) d\pi_{\xi|\xi_1, \dots, \xi_K}(\xi|\xi_1, \dots, \xi_K) d\pi(\xi_1, \dots, \xi_K) \\ \text{s.t.} \quad & \pi \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_K), \pi_{\xi|\xi_1, \dots, \xi_K} \in \mathcal{P}_{\xi|\xi_1, \dots, \xi_K}, \end{aligned} \quad (\text{BC-3})$$

where $\mathcal{P}_{\xi|\xi_1, \dots, \xi_K}$ denotes the family of all transition kernels $\pi_{\xi|\xi_1, \dots, \xi_K}$ satisfying the following two conditions. First, $\pi_{\xi|\xi_1, \dots, \xi_K}(\cdot|\xi_1, \dots, \xi_K)$ is a probability distribution on Ξ for all $(\xi_1, \dots, \xi_K) \in \Xi^K$. In addition, $\pi_{\xi|\xi_1, \dots, \xi_K}(B|\cdot)$ is a Borel-measurable function on Ξ^K for all Borel sets $B \subseteq \Xi$.

To prove the equivalence of (BC-2) and (BC-3), assume first that $(\mathbb{P}, \bar{\pi})$ is feasible in (BC-2). Next, define $\pi \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_K)$ as the marginal distribution of (ξ_1, \dots, ξ_K) under $\bar{\pi}$, and define $\pi_{\xi|\xi_1, \dots, \xi_K} \in \mathcal{P}_{\xi|\xi_1, \dots, \xi_K}$ as the regular conditional distribution of ξ given (ξ_1, \dots, ξ_K) under $\bar{\pi}$, which exists thanks to (Dudley 2002, Theorem 10.2.2). Thus, $(\pi, \pi_{\xi|\xi_1, \dots, \xi_K})$ is feasible in (BC-3). By the law of total probability, $(\mathbb{P}, \bar{\pi})$ and $(\pi, \pi_{\xi|\xi_1, \dots, \xi_K})$ attain the same objective function values. Conversely, select any $(\pi, \pi_{\xi|\xi_1, \dots, \xi_K})$ feasible in (BC-3), and define $\bar{\pi} \in \mathcal{P}(\Xi^{K+1})$ through

$$\bar{\pi}(B) = \int_{\Xi^K} \int_{\Xi} \mathbb{1}_{(\xi, \xi_1, \dots, \xi_K) \in B} d\pi_{\xi|\xi_1, \dots, \xi_K}(\xi|\xi_1, \dots, \xi_K) d\pi(\xi_1, \dots, \xi_K)$$

for every Borel set $B \subseteq \Xi^{K+1}$. In addition, set $\mathbb{P} = P_{0\#} \bar{\pi}$, where $P_0 : \Xi \times \Xi^K \rightarrow \Xi$ is defined through $P_0(\xi, \xi_1, \dots, \xi_K) = \xi$. We thus have $\bar{\pi} \in \Pi(\mathbb{P}, \mathbb{P}_1, \dots, \mathbb{P}_K)$, implying that $(\mathbb{P}, \bar{\pi})$ is feasible in (BC-2). One also easily verifies that $(\pi, \pi_{\xi|\xi_1, \dots, \xi_K})$ and $(\mathbb{P}, \bar{\pi})$ attain the same objective function values.

We are now ready to prove that problem (BC-3) is equivalent to

$$\min_{\pi \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_K)} \int_{\Xi^K} \min_{\pi_{\xi} \in \mathcal{P}(\Xi)} \int_{\Xi} \sum_{k=1}^K \lambda_k c(\xi, \xi_k) d\pi_{\xi}(\xi) d\pi(\xi_1, \dots, \xi_K). \quad (\text{BC-4})$$

Note first that (BC-4) is well defined. Indeed the inner minimization problem over the distribution π_{ξ} in (BC-4) can be reformulated as

$$\min_{\pi_{\xi} \in \mathcal{P}(\Xi)} \int_{\Xi} \sum_{k=1}^K \lambda_k c(\xi, \xi_k) d\pi_{\xi}(\xi) = \min_{\xi_0 \in \Xi} \sum_{k=1}^K \lambda_k c(\xi_0, \xi_k) = \phi(\xi_1, \dots, \xi_K), \quad (3)$$

where the first equality holds because the integral with respect to π_ξ is minimized by the Dirac distribution $\delta_{\Phi(\xi_1, \dots, \xi_K)}$ that places unit mass on the minimizer $\Phi(\xi_1, \dots, \xi_K)$ of the integrand, while the second equality follows from the definition of ϕ . By Assumption 2, the optimal value of the inner minimization problem in (BC-4) is thus integrable under every $\pi \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_K)$.

The equivalence of (BC-3) and (BC-4) follows from a standard interchangeability principle; see, *e.g.*, (Shapiro 2017, Proposition 2.1). As the inner problem in (BC-4) is solved by a Dirac distribution, however, it is easier to prove this equivalence directly instead of leveraging an abstract result. To this end, assume first that $(\pi, \pi_{\xi|\xi_1, \dots, \xi_K})$ is feasible in (BC-3). Hence, π is feasible in (BC-4). Note also that, for every fixed (ξ_1, \dots, ξ_K) , the probability distribution $\pi_{\xi|\xi_1, \dots, \xi_K}(\cdot|\xi_1, \dots, \xi_K)$ is feasible (but generally suboptimal) in the inner problem in (BC-4). Thus, the objective function value of π in (BC-4) is smaller than or equal to that of $(\pi, \pi_{\xi|\xi_1, \dots, \xi_K})$ in (BC-3). Conversely, if π is feasible in (BC-4), then $(\pi, \pi_{\xi|\xi_1, \dots, \xi_K})$ is feasible in (BC-3), where $\pi_{\xi|\xi_1, \dots, \xi_K}$ is defined through $\pi_{\xi|\xi_1, \dots, \xi_K}(\cdot|\xi_1, \dots, \xi_K) = \delta_{\Phi(\xi_1, \dots, \xi_K)}(\cdot)$. Note that $\pi_{\xi|\xi_1, \dots, \xi_K}$ defined in this way satisfies all properties of a transition kernel because Φ is measurable. By construction, the objective function values of $(\pi, \pi_{\xi|\xi_1, \dots, \xi_K})$ in problem (BC-4) and of π in problem (BC-3) match.

Substituting (3) into (BC-4) shows that (BC-4) is equivalent to (2). Hence, the original OT barycenter problem (OT-BC) is equivalent to (2), as well, and (2) inherits solvability from (OT-BC). Define now $\mathbb{P}^* = \Phi_{\#} \pi^*$, where π^* is any minimizer of the multi-margin OT problem (2). It remains to be shown that \mathbb{P}^* solves (OT-BC). Indeed, for any Borel set $B \subseteq \Xi$ we have

$$\begin{aligned} \mathbb{P}^*(B) &= \Phi_{\#} \pi^*(B) = \pi^*(\Phi^{-1}(B)) \\ &= \int_{\Xi^K} \delta_{\Phi(\xi_1, \dots, \xi_K)}(B) d\pi^*(\xi_1, \dots, \xi_K) \\ &= \int_{\Xi^K} \int_B d\pi_{\xi|\xi_1, \dots, \xi_K}^*(\xi|\xi_1, \dots, \xi_K) d\pi^*(\xi_1, \dots, \xi_K) \\ &= \int_{B \times \Xi^K} d\bar{\pi}^*(\xi, \xi_1, \dots, \xi_K) = P_{0\#} \bar{\pi}^*(B), \end{aligned}$$

where the transition kernel $\pi_{\xi|\xi_1, \dots, \xi_K}^*$ is defined through $\pi_{\xi|\xi_1, \dots, \xi_K}^*(\cdot|\xi_1, \dots, \xi_K) = \delta_{\Phi(\xi_1, \dots, \xi_K)}(\cdot)$, and the joint distribution $\bar{\pi}^*$ is obtained by combining the marginal distribution π^* with the conditional

distribution $\pi_{\xi|\xi_1, \dots, \xi_K}^*$. As the Borel set $B \subseteq \Xi$ was chosen arbitrarily, we have thus shown that $\mathbb{P}^* = P_{0\#}\bar{\pi}^*$. The equivalences between the different optimization problems established above lead to the following conclusion. As π^* solves (2) and (BC-4) by assumption, $(\pi^*, \pi_{\xi|\xi_1, \dots, \xi_K}^*)$ solves (BC-3), $(\mathbb{P}^*, \bar{\pi}^*)$ solves (BC-2), and \mathbb{P}^* solves (OT-BC). In addition, it is now easy to see that if Φ is unique, then every solution of (OT-BC) must be of the form $\mathbb{P}^* = \Phi_{\#}\pi^*$ for some solution π^* of (2). \blacksquare

Agueh and Carlier (2011) prove that the 2-Wasserstein barycenter is unique if the distributions $\mathbb{P}_1, \dots, \mathbb{P}_K$ assign probability zero to all sets of Hausdorff dimension less than d . Conversely, Theorem 1 suggests that the OT barycenter is *not* unique if either the multi-margin transportation plan π^* or the minimizer map Φ fails to be unique. We illustrate each case with an example. These examples are based on discrete distributions, which are supported on sets of Hausdorff dimension 0.

EXAMPLE 2 (NON-UNIQUE MULTI-MARGIN TRANSPORTATION PLAN). The OT barycenter of the distributions $\mathbb{P}_1 = \frac{1}{2}(\delta_{(1,1)} + \delta_{(0,0)})$ and $\mathbb{P}_2 = \frac{1}{2}(\delta_{(0,1)} + \delta_{(1,0)})$ in $\mathcal{P}_2(\mathbb{R}^2)$ with $\lambda_1 = \lambda_2 = 1$, $d(\xi, \xi') = \|\xi - \xi'\|_2$ and $p = 2$ is not unique. Indeed, while the minimizer map $\Phi(\xi_1, \xi_2) = \frac{1}{2}(\xi_1 + \xi_2)$ of problem (1) is unique, both $\pi^{*1} = \frac{1}{2}(\delta_{((1,1),(1,0))} + \delta_{((0,0),(0,1))})$ as well as $\pi^{*2} = \frac{1}{2}(\delta_{((1,1),(0,1))} + \delta_{((0,0),(1,0))})$ solve the multi-margin OT problem (2). By Theorem 1, both $\mathbb{P}^{*1} = \frac{1}{2}(\delta_{(1,0.5)} + \delta_{(0,0.5)})$ as well as $\mathbb{P}^{*2} = \frac{1}{2}(\delta_{(0.5,1)} + \delta_{(0.5,0)})$ thus constitute minimizers of the OT barycenter problem (OT-BC). \blacktriangle

EXAMPLE 3 (NON-UNIQUE MINIMIZER MAP). The OT barycenter of two distributions $\mathbb{P}_1 \neq \mathbb{P}_2$ in $\mathcal{P}_2(\mathbb{R}^2)$ with $\lambda_1 = \lambda_2 = 1$ and $d(\xi, \xi') = \|\xi - \xi'\|_2$ is not unique. Indeed, while the multi-margin transportation plan π^* that solves (2) may or may not be unique in this situation, both $\Phi^{*1}(\xi_1, \xi_2) = \xi_1$ and $\Phi^{*2}(\xi_1, \xi_2) = \xi_2$ constitute minimizer maps for problem (1). As $\pi^* \in \Pi(\mathbb{P}_1, \mathbb{P}_2)$, Theorem 1 thus implies that both $\mathbb{P}_1 = \Phi_{\#}^{*1}\pi^*$ and $\mathbb{P}_2 = \Phi_{\#}^{*2}\pi^*$ solve the OT barycenter problem (OT-BC). \blacktriangle

Theorem 1 also implies that all 2-Wasserstein barycenters of two distributions that can be obtained by changing their weights reside on the 2-Wasserstein geodesic between these distributions.

PROPOSITION 1 (2-Wasserstein Barycenters Trace Out Geodesics). *If $\mathbb{P}_1, \mathbb{P}_2 \in \mathcal{P}_2(\mathbb{R}^d)$ are continuous distributions, then the 2-Wasserstein barycenter of \mathbb{P}_1 and \mathbb{P}_2 with respective weights λ and $1 - \lambda$ traces out the 2-Wasserstein geodesic between \mathbb{P}_1 and \mathbb{P}_2 as λ is swept from 0 to 1.*

Proof of Proposition 1 As \mathbb{P}_1 and \mathbb{P}_2 are continuous, the 2-Wasserstein geodesic between \mathbb{P}_1 and \mathbb{P}_2 coincides with the family of all probability distributions of the form $(T_\lambda)_\# \pi^*$ for $\lambda \in [0, 1]$, where $T_\lambda(\xi_1, \xi_2) = \lambda \xi_1 + (1 - \lambda) \xi_2$, and π^* solves the OT problem between \mathbb{P}_1 and \mathbb{P}_2 with squared Euclidean transportation cost (Panaretos and Zemel 2020, p. 46). Using the notation of Theorem 1 with $K = 2$, $\lambda_1 = \lambda$ and $\lambda_2 = 1 - \lambda$, one readily verifies that $\Phi(\xi_1, \xi_2) = T_\lambda(\xi_1, \xi_2)$ and

$$\phi(\xi_1, \xi_2) = \min_{\xi \in \mathbb{R}^d} \lambda \|\xi_1 - \xi\|_2^2 + (1 - \lambda) \|\xi_2 - \xi\|_2^2 = \lambda(1 - \lambda) \|\xi_1 - \xi_2\|_2^2$$

are the unique minimizer and the optimal value of problem (1), respectively. As ϕ is a nonnegative multiple of the squared Euclidean transportation cost, π^* solves the multi-margin OT problem (2). Theorem 1 thus implies that $(T_\lambda)_\# \pi^*$ constitutes a 2-Wasserstein barycenter of \mathbb{P}_1 and \mathbb{P}_2 with weights λ and $1 - \lambda$, respectively. As \mathbb{P}_1 and \mathbb{P}_2 are continuous, this barycenter is unique (Agueh and Carlier 2011). Hence, the barycenters of \mathbb{P}_1 and \mathbb{P}_2 trace out the geodesic from \mathbb{P}_1 to \mathbb{P}_2 . ■

To close this section, we show that if distances between distributions are measured by a Wasserstein distance of order 1 or 2, then the resulting OT barycenters may have undesirable properties. We first show that the 1-Wasserstein barycenter of $K = 2$ distributions coincides with the distribution that has the larger weight and is degenerate if both distributions have the same weight.

COROLLARY 1 (1-Wasserstein Barycenter of Two Distributions). *If $\mathbb{P}_1, \mathbb{P}_2 \in \mathcal{P}(\Xi)$ satisfy $\int_\Xi \|\xi\| d\mathbb{P}_k(\xi) < \infty$ for $k = 1, 2$ and if $d(\xi, \xi') = \|\xi - \xi'\|$, then a 1-Wasserstein barycenter of \mathbb{P}_1 and \mathbb{P}_2 with respective weights $\lambda_1, \lambda_2 \geq 0$ is given by \mathbb{P}_1 if $\lambda_1 \geq \lambda_2$ and by \mathbb{P}_2 if $\lambda_1 \leq \lambda_2$.*

Proof of Corollary 1 The claim follows directly from Theorem 1. Indeed, if $\lambda_1 \geq \lambda_2$, then $\Phi(\xi_1, \xi_2) = \xi_1$ is a measurable selector of the arg min multifunction of problem (1). Thus, if π^* solves problem (2), then $\xi_{1\#} \pi^* = \mathbb{P}_1$ solves the OT barycenter problem (OT-BC) thanks to Theorem 1. ■

Corollary 1 suggests that 1-Wasserstein barycenters have little conceptual appeal. In contrast, 2-Wasserstein barycenters are widely used in computer graphics (Solomon et al. 2015), machine learning (Dognin et al. 2019, Montesuma and Mboula 2021), Bayesian inference (Srivastava et al. 2018) and statistics (Claici et al. 2018, Barré et al. 2020) etc. In all of these applications, however, the

reference distributions $\mathbb{P}_1, \dots, \mathbb{P}_K$ are typically unknown. Given N_k independent samples $\{\widehat{\xi}_{k,j}\}_{j=1}^{N_k}$ from \mathbb{P}_k , $k = 1, \dots, K$, one can approximate the unobservable distribution \mathbb{P}_k with the corresponding empirical distribution $\widehat{\mathbb{P}}_k = \frac{1}{N_k} \sum_{j=1}^{N_k} \delta_{\widehat{\xi}_{k,j}}$. The unknown 2-Wasserstein barycenter \mathbb{P}^* of $\mathbb{P}_1, \dots, \mathbb{P}_K$ can thus be estimated by the empirical 2-Wasserstein barycenter $\widehat{\mathbb{P}}$ of $\widehat{\mathbb{P}}_1, \dots, \widehat{\mathbb{P}}_K$, which is itself a discrete distribution. The next proposition shows that this estimator is biased. In this proposition, we use $\mathbb{P}\text{-VAR}(\xi) = \mathbb{E}_{\mathbb{P}}[\|\xi - \mathbb{E}_{\mathbb{P}}[\xi]\|_2^2]$ to denote the trace of the covariance matrix of $\xi \in \mathbb{R}^d$ under the distribution $\mathbb{P} \in \mathcal{P}_2(\Xi)$. For brevity, we will sometimes refer to $\mathbb{P}\text{-VAR}(\xi)$ as the variance of ξ . It is easy to verify that $\mathbb{P}\text{-VAR}(\xi) = \min_{\xi_0 \in \Xi} \mathbb{E}_{\mathbb{P}}[\|\xi - \xi_0\|_2^2]$, where the minimization problem over ξ_0 is solved by $\mathbb{E}_{\mathbb{P}}[\xi]$. This formula shows that $\mathbb{P}\text{-VAR}(\xi)$ is concave in \mathbb{P} .

PROPOSITION 2 (Empirical OT Barycenters are Biased). *Set $c(\xi_1, \xi_2) = \|\xi_1 - \xi_2\|_2^2$ and $\Xi = \mathbb{R}^d$, and select $\lambda_1, \dots, \lambda_K > 0$. Define \mathbb{P}^* as the OT barycenter of K mutually different continuous distributions $\mathbb{P}_1, \dots, \mathbb{P}_K \in \mathcal{P}_2(\Xi)$. Similarly, define $\widehat{\mathbb{P}}$ as an OT barycenter of $\widehat{\mathbb{P}}_1, \dots, \widehat{\mathbb{P}}_K$, where $\widehat{\mathbb{P}}_k$ denotes the empirical distribution corresponding to N_k random samples from \mathbb{P}_k , $k \in [K]$. Then, we have $\mathbb{E}[\widehat{\mathbb{P}}\text{-VAR}(\xi)] < \mathbb{P}^*\text{-VAR}(\xi)$, where the expectation is with respect to all samples.*

As the transportation cost function $c(\xi_1, \xi_2)$ is quadratic and the distributions $\mathbb{P}_1, \dots, \mathbb{P}_K$ are square-integrable, Assumptions 1 and 2 both hold under the conditions of Proposition 2. Hence, \mathbb{P}^* exists thanks to Lemma 1, and it is unique thanks to (Agueh and Carlier 2011, Proposition 3.5). The empirical OT barycenter $\widehat{\mathbb{P}}$ trivially exists because $\widehat{\mathbb{P}}_1, \dots, \widehat{\mathbb{P}}_K$ are discrete.

Proof of Proposition 2 The proof consists of two parts. In the first part we use Theorem 1 to show that the multi-margin transportation plan π^* that induces \mathbb{P}^* has a variance-maximizing property. In the second part we exploit this variance-maximizing property to prove that $\widehat{\mathbb{P}}$ is biased.

Part I Assume without loss of generality that $\sum_{k=1}^K \lambda_k = 1$, and let $\pi \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_K)$ be any multi-martin transportation plan feasible in (2). Next, construct a joint distribution \mathbb{Q}_{π} of the random vector $(\xi, \xi_1, \dots, \xi_K) \in \Xi^{K+1}$ under which the marginal distribution of (ξ_1, \dots, ξ_K) is given by π , and the distribution of ξ conditional on (ξ_1, \dots, ξ_K) is given by $\sum_{k=1}^K \lambda_k \delta_{\xi_k}$. Note that the marginal

distribution $\sum_{k=1}^K \lambda_k \mathbb{P}_k$ and the variance \mathbb{Q}_π -VAR(ξ) of ξ under \mathbb{Q}_π are actually independent of π .

Using the same notation as in Theorem 1, one readily verifies that

$$\mathbb{Q}_\pi\text{-VAR}(\xi|\xi_1, \dots, \xi_K) = \min_{\xi_0 \in \Xi} \sum_{k=1}^K \lambda_k \|\xi_0 - \xi_k\|_2^2 = \phi(\xi_1, \dots, \xi_K)$$

and that the minimization problem over ξ_0 is uniquely solved by $\mathbb{E}_{\mathbb{Q}_\pi}[\xi|\xi_1, \dots, \xi_K] = \Phi(\xi_1, \dots, \xi_K)$.

By the law of total variance, we thus have

$$\begin{aligned} \mathbb{Q}_\pi\text{-VAR}(\xi) &= \mathbb{E}_{\mathbb{Q}_\pi}[\mathbb{Q}_\pi\text{-VAR}(\xi|\xi_1, \dots, \xi_K)] + \mathbb{Q}_\pi\text{-VAR}(\mathbb{E}_{\mathbb{Q}_\pi}[\xi|\xi_1, \dots, \xi_K]) \\ &= \mathbb{E}_{\mathbb{Q}_\pi}[\phi(\xi_1, \dots, \xi_K)] + \mathbb{Q}_\pi\text{-VAR}(\Phi(\xi_1, \dots, \xi_K)) \\ &= \mathbb{E}_\pi[\phi(\xi_1, \dots, \xi_K)] + \Phi_{\#}\pi\text{-VAR}(\xi), \end{aligned} \tag{4}$$

where the last equality holds because π is the marginal of (ξ_1, \dots, ξ_K) under \mathbb{Q}_π , which in turn implies that the pushforward distribution $\Phi_{\#}\pi$ is the marginal of $\Phi(\xi_1, \dots, \xi_K)$ under π . As the minimizer function Φ is unique, Theorem 1 implies that the OT barycenter \mathbb{P}^* must coincide with $\Phi_{\#}\pi^*$ for some solution π^* of the multi-margin OT problem (2), which minimizes $\mathbb{E}_\pi[\phi(\xi_1, \dots, \xi_K)]$ across all $\pi \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_K)$. Recall now that \mathbb{Q}_π -VAR(ξ) is independent of π . We may thus conclude from (4) that π^* must maximize the variance $\Phi_{\#}\pi$ -VAR(ξ) across all $\pi \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_K)$.

Part II We now use the insights of Part I to prove that $\widehat{\mathbb{P}}$ is biased. To this end, let $\widehat{\pi} \in \Pi(\widehat{\mathbb{P}}_1, \dots, \widehat{\mathbb{P}}_K)$ be a multi-margin transportation plan with $\widehat{\mathbb{P}} = \Phi_{\#}\widehat{\pi}$. Note that $\widehat{\pi}$ exists thanks to Theorem 1. Note also that the discrete distributions $\widehat{\mathbb{P}}$ and $\widehat{\pi}$ constitute random objects because they depend on the random samples underlying the empirical distributions $\widehat{\mathbb{P}}_1, \dots, \widehat{\mathbb{P}}_K$. By (Rockafellar and Wets 2009, Theorem 14.37), we may assume without loss of generality that both the locations as well as the probabilities of the atoms of $\widehat{\mathbb{P}}$ and $\widehat{\pi}$ depend measurably on these samples.¹ Next, define the distribution $\pi = \mathbb{E}[\widehat{\pi}]$ through $\pi(B) = \mathbb{E}[\widehat{\pi}(B)]$ for every Borel set $B \subseteq \Xi^K$, which is well-defined because $\widehat{\pi}(B)$ is a nonnegative measurable function of the samples. It is easy to verify that $\pi \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_K)$, that is, π is feasible in (2). In summary, we may conclude that

$$\mathbb{E}[\widehat{\mathbb{P}}\text{-VAR}(\xi)] = \mathbb{E}[\Phi_{\#}\widehat{\pi}\text{-VAR}(\xi)] = \mathbb{E}[\widehat{\pi}\text{-VAR}(\Phi(\xi_1, \dots, \xi_K))]$$

¹ Strictly speaking, this is an assumption about $\widehat{\mathbb{P}}$ and should be mentioned in the proposition statement. As it can be imposed without loss of generality, we hide it in the proof for better readability of the proposition statement.

$$\begin{aligned} &\leq \mathbb{E}[\widehat{\pi}] - \text{VAR}(\Phi(\xi_1, \dots, \xi_K)) = \pi - \text{VAR}(\Phi(\xi_1, \dots, \xi_K)) = \Phi_{\#}\pi - \text{VAR}(\xi) \\ &< \Phi_{\#}\pi^* - \text{VAR}(\xi) = \mathbb{P}^* - \text{VAR}(\xi). \end{aligned}$$

Here, the first inequality follows from Jensen's inequality, which applies because the variance of any random vector is concave in that random vector's distribution. The second inequality follows from the variance-maximizing property of π^* established in Part I. This inequality is strict if $\pi^* \neq \pi$. Indeed, the multi-margin OT plan π^* is the unique minimizer of $\mathbb{E}_{\pi}[\phi(\xi_1, \dots, \xi_K)]$ (by (Agueh and Carlier 2011, Theorem 4.1)) and therefore also the unique maximizer of $\Phi_{\#}\pi - \text{VAR}(\xi)$ (by Part I).

It remains to be shown that $\pi^* \neq \pi = \mathbb{E}[\widehat{\pi}]$. By (Agueh and Carlier 2011, Theorem 4.1) and (Gangbo and Świąch 1998, Theorem 2.1), there exists a bijective transformation $T : \Xi \rightarrow \Xi$ with $T_{\#}\mathbb{P}_1 = \mathbb{P}_2$ and $\pi^*(\xi_2 = T(\xi_1)) = 1$. Select now two Borel sets $A, B \subseteq \Xi$ with $A \cap B = \emptyset$ and $\mathbb{P}_1(A), \mathbb{P}_1(B) > 0$. Such sets exist because \mathbb{P}_1 is continuous. We thus have

$$\pi^*(\xi_1 \in A, \xi_2 \in T(B)) = \pi^*(\xi_1 \in A, \xi_1 \in B) = 0,$$

where the first equality holds because $\pi^*(\xi_2 = T(\xi_1)) = 1$ and T is bijective, while the second equality holds because $A \cap B = \emptyset$. Next, observe that the event in which $\widehat{\mathbb{P}}_1$ is supported on A and $\widehat{\mathbb{P}}_2$ is supported on $T(B)$ has a strictly positive probability equal to $\mathbb{P}_1(A)^{N_1} \mathbb{P}_2(T(B))^{N_2} = \mathbb{P}_1(A)^{N_1} \mathbb{P}_1(B)^{N_2} > 0$. In this event, we have $\widehat{\pi}(\xi_1 \in A, \xi_2 \in T(B)) = 1$ because $\widehat{\mathbb{P}}_1$ and $\widehat{\mathbb{P}}_2$ are the marginal distributions of ξ_1 and ξ_2 under $\widehat{\pi}$, respectively. This reasoning implies that

$$\mathbb{E}[\widehat{\pi}](\xi_1 \in A, \xi_2 \in T_2(B)) = \mathbb{E}[\widehat{\pi}(\xi_1 \in A, \xi_2 \in T_2(B))] \geq \mathbb{P}_1(A)^{N_1} \mathbb{P}_2(T(B))^{N_2} > 0.$$

We thus have $\mathbb{E}[\widehat{\pi}](\xi_1 \in A, \xi_2 \in T_2(B)) > \pi^*(\xi_1 \in A, \xi_2 \in T_2(B))$, which implies that $\pi^* \neq \mathbb{E}[\widehat{\pi}]$. ■

Having covered the basic theory of OT barycenters, we now turn to data-driven decision-making.

3. Multi-Source DRO

From now on we consider a stochastic optimization problem of the form

$$\min_{\theta \in \Theta} \int_{\Xi} \ell(\theta, \xi) d\mathbb{P}(\xi), \tag{5}$$

which minimizes the expected value of an uncertainty-affected loss function $\ell : \Xi \times \Theta \rightarrow \mathbb{R}$ over a feasible set $\Theta \subseteq \mathbb{R}^n$. We henceforth assume that $\ell(\theta, \xi)$ is lower semi-continuous in θ for every fixed $\xi \in \Xi$ and upper semi-continuous in ξ for every fixed $\theta \in \Theta$. We also assume that Θ is closed. In most applications of stochastic programming, the target distribution \mathbb{P} , which is needed to evaluate the expected loss of any fixed decision θ , is unknown. Given samples from \mathbb{P} , however, one can approximate \mathbb{P} with an estimated distribution $\widehat{\mathbb{P}}$ constructed from the data. Solving (5) under $\widehat{\mathbb{P}}$ instead of \mathbb{P} may lead to overfitting effects if data is scarce. That is, the optimal solution computed under $\widehat{\mathbb{P}}$ may underperform under \mathbb{P} , a phenomenon known as the *optimizer's curse* (Smith and Winkler 2006). As a possible remedy, it has been proposed to solve a DRO problem of the form

$$\inf_{\theta \in \Theta} \sup_{\mathbb{P} \in \mathbb{B}_\varepsilon(\widehat{\mathbb{P}})} \int_{\Xi} \ell(\theta, \xi) d\mathbb{P}(\xi), \quad (6)$$

where $\mathbb{B}_\varepsilon(\widehat{\mathbb{P}}) = \{\mathbb{P} \in \mathcal{P}(\Xi) : C(\mathbb{P}, \widehat{\mathbb{P}}) \leq \varepsilon\}$ denotes the set of all distributions whose OT distance to $\widehat{\mathbb{P}}$ is at most $\varepsilon \geq 0$ (Pflug et al. 2012). The DRO problem (6) pretends that an evil adversary can morph the estimated distribution $\widehat{\mathbb{P}}$ into some undesirable shape at a finite transportation budget ε . It turns out that preparing for destructive actions of a fictitious adversary can mitigate the optimizer's curse. The statistical and computational properties of OT-based DRO are studied in (Mohajerin Esfahani and Kuhn 2018, Zhao and Guan 2018, Blanchet and Murthy 2019, Gao and Kleywegt 2023).

In this paper we address more challenging decision-making situations in which data from the target distribution \mathbb{P} is extremely scarce or even unavailable. Instead, we assume to have access to data from $K > 1$ source distributions $\mathbb{P}_k \in \mathcal{P}(\Xi)$, $k \in [K]$, which are believed to be similar to \mathbb{P} . Each source distribution \mathbb{P}_k can thus be approximated by an estimated distribution $\widehat{\mathbb{P}}_k \in \mathcal{P}(\Xi)$ constructed from the respective source data. In this situation, one might solve the DRO problem (6) for a reference distribution $\widehat{\mathbb{P}}$ that is defined as a mixture or as an OT barycenter of the reference distributions $\widehat{\mathbb{P}}_k$, $k \in [K]$ (Lau and Liu 2022). However, mixtures often fail to retain stylized features of the target distribution (see Example 1), and OT barycenters are plagued by the shortcomings highlighted in Section 2. This prompts us to introduce the multi-source DRO problem

$$\inf_{\theta \in \Theta} \sup_{\mathbb{P} \in \bigcap_{k=1}^K \mathbb{B}_{\varepsilon_k}(\widehat{\mathbb{P}}_k)} \int_{\Xi} \ell(\theta, \xi) d\mathbb{P}(\xi), \quad (\text{MS-DRO})$$

which seeks to minimize the worst-case expected loss across all distributions in the intersection of K OT ambiguity sets. This ambiguity set contains all distributions that can be obtained by reshaping the reference distribution $\widehat{\mathbb{P}}_k$ at a transportation cost of at most ε_k for every $k \in [K]$.

In the remainder of this section we identify conditions under which the multi-source DRO problem (MS-DRO) is equivalent to a finite convex program susceptible to numerical solution. To this end, we first dualize the inner maximization problem over \mathbb{P} to obtain an equivalent minimization problem that can be combined with the outer minimization over θ . For ease of exposition, we temporarily hide the dependence of the loss function on θ , in which case the inner maximization problem in (MS-DRO) simplifies to the following *primal uncertainty quantification problem*

$$\sup_{\mathbb{P} \in \bigcap_{k=1}^K \mathbb{B}_{\varepsilon_k}(\widehat{\mathbb{P}}_k)} \int_{\Xi} \ell(\xi) \, d\mathbb{P}(\xi). \quad (\text{P-UQ})$$

Our convex reformulation results rely on the assumption that all reference distributions are discrete.

ASSUMPTION 3 (Reference Distributions). *The k -th reference distribution is representable as $\widehat{\mathbb{P}}_k = \sum_{j=1}^{N_k} p_{k,j} \delta_{\widehat{\xi}_{k,j}}$ with atoms $\{\widehat{\xi}_{k,j}\}_{j=1}^{N_k}$ and strictly positive probabilities $\{p_{k,j}\}_{j=1}^{N_k}$ for all $k \in [K]$.*

In view of Assumption 3 it is useful to define $\mathcal{A} = \times_{k=1}^K [N_k]$ as a set of K -dimensional multi-indices. Each $\alpha \in \mathcal{A}$ uniquely identifies a combination of atoms from the K discrete reference distributions. In addition, we will also need the following growth condition on the loss function.

ASSUMPTION 4 (Growth Condition). *There exist $\xi' \in \Xi$, $g' > 0$ and $p' \in \mathbb{N}$ such that $\ell(\xi) \leq g'(1 + d^{p'}(\xi, \xi'))$ for all $\xi \in \Xi$, where the d is the same metric as in Assumption 1.*

We are now ready to derive the dual of the uncertainty quantification problem (P-UQ). In the following we use \mathcal{E} to denote the set of all $\varepsilon = (\varepsilon_1, \dots, \varepsilon_K) \in \mathbb{R}_+^K$ for which (P-UQ) is feasible.

THEOREM 2 (Strong Duality). *If Assumptions 1, 3 and 4 hold with $p' \leq p$, then the primal uncertainty quantification problem (P-UQ) admits the dual*

$$\begin{aligned} \inf \quad & \sum_{k=1}^K \varepsilon_k \lambda_k + \sum_{k=1}^K \sum_{j=1}^{N_k} p_{k,j} \gamma_{k,j} \\ \text{s.t.} \quad & \lambda_k \in \mathbb{R}_+, \gamma_k \in \mathbb{R}^{N_k} \quad \forall k \in [K] \\ & \sup_{\xi \in \Xi} \ell(\xi) - \sum_{k=1}^K \lambda_k c(\xi, \widehat{\xi}_{k, \alpha_k}) \leq \sum_{k=1}^K \gamma_{k, \alpha_k} \quad \forall \alpha \in \mathcal{A}. \end{aligned} \quad (\text{D-UQ})$$

Strong (i.e., gap-free) duality holds for all $\varepsilon \notin \partial\mathcal{E}$.

If each reference distribution has the same number of atoms, that is, if $N_k = N$ for all $k \in [K]$, then $|\mathcal{A}| = N^K$. Hence, the number of constraints in problem (D-UQ) grows exponentially with K .

The proof of Theorem 2 relies on a decomposability result for multi-margin transportation plans.

LEMMA 2 (Multi-Margin Transportation Plans). *If Assumption 3 holds, then for any $\bar{\pi} \in \Pi(\mathbb{P}, \widehat{\mathbb{P}}_1, \dots, \widehat{\mathbb{P}}_K)$ there exist finite Borel measures $\nu_\alpha \in \mathcal{M}_+(\Xi)$, $\alpha \in \mathcal{A}$, with $\bar{\pi} = \sum_{\alpha \in \mathcal{A}} \nu_\alpha \otimes \delta_\alpha$, where δ_α is a shorthand for the Dirac distribution on Ξ^K that concentrates unit mass at $(\widehat{\xi}_{1,\alpha_1}, \dots, \widehat{\xi}_{K,\alpha_K})$.*

Proof of Lemma 2 Fix an arbitrary $\bar{\pi} \in \Pi(\mathbb{P}, \widehat{\mathbb{P}}_1, \dots, \widehat{\mathbb{P}}_K)$. Next, for any $\alpha \in \mathcal{A}$, construct a non-negative measure $\nu_\alpha \in \mathcal{M}_+(\Xi)$ by setting $\nu_\alpha(B) = \bar{\pi}(\xi \in B, \xi_k = \widehat{\xi}_{k,\alpha_k} \forall k = 1, \dots, K)$ for all Borel sets $B \subseteq \Xi$. By construction, ν_α thus constitutes a rescaled probability measure. As any probability measure is a Radon measure, we may conclude that ν_α is indeed a Radon measure for every $\alpha \in \mathcal{A}$. It remains to be shown that the measures ν_α , $\alpha \in \mathcal{A}$, constructed in this way satisfy $\bar{\pi} = \sum_{\alpha \in \mathcal{A}} \nu_\alpha \otimes \delta_\alpha$.

To this end, note that for any Borel sets $B, B_1, \dots, B_K \subseteq \Xi$ we have

$$\begin{aligned} \bar{\pi}(\xi \in B, \xi_k \in B_k \forall k) &= \sum_{\substack{\alpha \in \mathcal{A}: \\ \widehat{\xi}_{k,\alpha_k} \in B_k \forall k}} \bar{\pi}(\xi \in B, \xi_k = \widehat{\xi}_{k,\alpha_k} \forall k) \\ &= \sum_{\substack{\alpha \in \mathcal{A}: \\ \widehat{\xi}_{k,\alpha_k} \in B_k \forall k}} \nu_\alpha(\xi \in B) = \sum_{\alpha \in \mathcal{A}} \nu_\alpha(\xi \in B) \delta_\alpha(\xi_k \in B_k \forall k), \end{aligned}$$

where the first equality follows from the law of total probability, while the second and the third equalities exploit the definitions of ν_α and δ_α , respectively. The identity $\bar{\pi} = \sum_{\alpha \in \mathcal{A}} \nu_\alpha \otimes \delta_\alpha$ now follows because the sets of the form $B \times B_1 \times \dots \times B_K$ generate the Borel σ -algebra on Ξ^{K+1} . ■

Proof of Theorem 2. Recall that $C(\mathbb{P}, \widehat{\mathbb{P}}_k)$ denotes the optimal value of a minimization problem over transportation plans and is solvable thanks to Assumption 1; see (Villani 2009, Theorem 4.1). Consequently, the constraint $C(\mathbb{P}, \widehat{\mathbb{P}}_k) \leq \varepsilon_k$ is equivalent to the requirement that there exists $\pi_k \in \Pi(\mathbb{P}, \widehat{\mathbb{P}}_k)$ with $\int_{\Xi^2} c(\xi, \xi_k) d\pi_k(\xi, \xi_k) \leq \varepsilon_k$. This in turn implies that problem (P-UQ) is equivalent to

$$\begin{aligned} &\sup \int_{\Xi} \ell(\xi) d\mathbb{P}(\xi) \\ &\text{s.t. } \mathbb{P} \in \mathcal{P}(\Xi), \pi_k \in \Pi(\mathbb{P}, \widehat{\mathbb{P}}_k) \quad \forall k \in [K] \\ &\int_{\Xi^2} c(\xi, \xi_k) d\pi_k(\xi, \xi_k) \leq \varepsilon_k \quad \forall k \in [K]. \end{aligned} \tag{7}$$

Next, fix any distribution $\mathbb{P} \in \mathcal{P}(\Xi)$ and transportation plans $\pi_k \in \Pi(\mathbb{P}, \widehat{\mathbb{P}}_k)$ for $k \in [K]$. An iterative application of the Gluing Lemma (Villani 2009, Chapter 1) then shows that there must exist a multi-margin transportation plan $\bar{\pi} \in \Pi(\mathbb{P}, \widehat{\mathbb{P}}_1, \dots, \widehat{\mathbb{P}}_K)$ with $P_{0,k\#}\bar{\pi} = \pi_k$ for all $k \in [K]$, where $P_{0,k} : \Xi^{K+1} \rightarrow \Xi^2$ is defined as usual through $P_{0,k}(\xi, \xi_1, \dots, \xi_K) = (\xi, \xi_k)$. Appending $\bar{\pi}$ as a decision variable to problem (7) and enforcing the constraints that link $\bar{\pi}$ to π_k , we obtain

$$\begin{aligned} & \sup \int_{\Xi} \ell(\xi) \, d\mathbb{P}(\xi) \\ \text{s.t.} \quad & \mathbb{P} \in \mathcal{P}(\Xi), \bar{\pi} \in \Pi(\mathbb{P}, \widehat{\mathbb{P}}_1, \dots, \widehat{\mathbb{P}}_K), \pi_k \in \Pi(\mathbb{P}, \widehat{\mathbb{P}}_k) \quad \forall k \in [K] \\ & P_{0,k\#}\bar{\pi} = \pi_k \quad \forall k \in [K] \\ & \int_{\Xi \times \Xi} c(\xi, \xi_k) \, d\pi_k(\xi, \xi_k) \leq \varepsilon_k \quad \forall k \in [K]. \end{aligned}$$

By the measure-theoretic change-of-variables formula, problem (P-UQ) is thus equivalent to

$$\begin{aligned} & \sup \int_{\Xi} \ell(\xi) \, d\mathbb{P}(\xi) \\ \text{s.t.} \quad & \mathbb{P} \in \mathcal{P}(\Xi), \bar{\pi} \in \Pi(\mathbb{P}, \widehat{\mathbb{P}}_1, \dots, \widehat{\mathbb{P}}_K) \\ & \int_{\Xi^{K+1}} c(\xi, \xi_k) \, d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) \leq \varepsilon_k \quad \forall k \in [K]. \end{aligned} \tag{8}$$

By Lemma 2, every multi-margin transportation plan $\bar{\pi} \in \Pi(\mathbb{P}, \widehat{\mathbb{P}}_1, \dots, \widehat{\mathbb{P}}_K)$ can be represented as $\bar{\pi} = \sum_{\alpha \in \mathcal{A}} \nu_{\alpha} \otimes \delta_{\alpha}$ for some finite Borel measures $\nu_{\alpha} \in \mathcal{M}_+(\Xi)$, $\alpha \in \mathcal{A}$, where δ_{α} is a shorthand for the Dirac point mass at $(\widehat{\xi}_{1,\alpha_1}, \dots, \widehat{\xi}_{K,\alpha_K})$. Thus, problem (8) is equivalent to

$$\begin{aligned} & \sup \sum_{\alpha \in \mathcal{A}} \int_{\Xi} \ell(\xi) \, d\nu_{\alpha}(\xi) \\ \text{s.t.} \quad & \nu_{\alpha} \in \mathcal{M}_+(\Xi) \quad \forall \alpha \in \mathcal{A} \\ & \sum_{\substack{\alpha \in \mathcal{A}: \\ \alpha_k = j}} \int_{\Xi} d\nu_{\alpha}(\xi) = p_{k,j} \quad \forall j \in [N_k], \forall k \in [K] \\ & \sum_{\alpha \in \mathcal{A}} \int_{\Xi} c(\xi, \widehat{\xi}_{k,\alpha_k}) \, d\nu_{\alpha}(\xi) \leq \varepsilon_k \quad \forall k \in [K], \end{aligned} \tag{9}$$

which can be viewed as a semi-infinite linear program over $|\mathcal{A}| = \prod_{k=1}^K N_k$ Borel measures.

Assigning Lagrange multipliers $\gamma_{k,j} \in \mathbb{R}$ and $\lambda_k \in \mathbb{R}_+$ to the probability matching and OT constraints, respectively, the Lagrangian dual of problem (9) can be represented as

$$\begin{aligned} \inf_{\lambda \geq 0, \gamma} \sup_{\{\nu_\alpha\}_{\alpha \in \mathcal{M}_+(\Xi)}} & \sum_{k=1}^K \lambda_k \varepsilon_k + \sum_{k=1}^K \sum_{j=1}^{N_k} \gamma_{k,j} p_{k,j} - \sum_{k=1}^K \lambda_k \left(\sum_{\alpha \in \mathcal{A}} \int_{\Xi} c(\xi, \widehat{\xi}_{k, \alpha_k}) d\nu_\alpha(\xi) \right) \\ & - \sum_{k=1}^K \sum_{j=1}^{N_k} \gamma_{k,j} \sum_{\substack{\alpha \in \mathcal{A}: \\ \alpha_k = j}} \int_{\Xi} d\nu_\alpha(\xi) + \sum_{\alpha \in \mathcal{A}} \int_{\Xi} \ell(\xi) d\nu_\alpha(\xi). \end{aligned} \quad (10)$$

Routine rearrangements can then be used to show that problem (10) is equivalent to

$$\inf_{\lambda \geq 0, \gamma} \sum_{k=1}^K \varepsilon_k \lambda_k + \sum_{k=1}^K \sum_{j=1}^{N_k} p_{k,j} \gamma_{k,j} + \sum_{\alpha \in \mathcal{A}} \sup_{\nu_\alpha \in \mathcal{M}_+(\Xi)} \int_{\Xi} \left(\ell(\xi) - \sum_{k=1}^K \lambda_k c(\xi, \widehat{\xi}_{k, \alpha_k}) - \sum_{k=1}^K \gamma_{k, \alpha_k} \right) d\nu_\alpha(\xi).$$

Note that the supremum over ν_α evaluates to infinity unless the following robust constraint holds.

$$\ell(\xi) - \sum_{k=1}^K \lambda_k c(\xi, \widehat{\xi}_{k, \alpha_k}) - \sum_{k=1}^K \gamma_{k, \alpha_k} \leq 0 \quad \forall \xi \in \Xi$$

If this robust constraint holds, on the other hand, then the supremum over ν_α vanishes. Hence, the dual problem (10) is indeed equivalent to (D-UQ).

It remains to be shown that strong duality holds. Note first that, as Assumptions 1 and 4 hold for some $p' \leq p$, problem (P-UQ) and its equivalent reformulation (9) cannot be unbounded, and problem (D-UQ) is guaranteed to be feasible. In addition, the optimal values of (9) and (D-UQ) are ostensibly concave in ε . Using now standard perturbation arguments, one can show that these functions coincide and that strong duality holds for all $\varepsilon \neq \partial \mathcal{E}$ (Rockafellar 1974, Shapiro 2001). ■

Problem (D-UQ) can be viewed as a robust optimization problem with exponentially many robust constraints indexed by $\alpha \in \mathcal{A}$. Each constraint involves an embedded maximization problem over all uncertainty realizations $\xi \in \Xi$, which is generically hard. Under the following convexity assumption, however, the robust constraints can be reformulated in terms of finitely many convex constraints.

ASSUMPTION 5 (Convexity). *The support set Ξ is convex, the transportation cost function $c(\xi, \xi')$ is convex in ξ for every fixed $\xi' \in \Xi$, and the loss function is representable as $\ell(\xi) = \max_{l \in [L]} \ell_l(\xi)$ for some concave component functions $\ell_l : \Xi \rightarrow \mathbb{R}$, $l \in [L]$.*

The next theorem follows from now standard reformulation techniques developed in (Ben-Tal et al. 2015, Zhen et al. 2023) and is therefore stated without proof.

THEOREM 3 (Reformulation as a Finite Convex Program). *If Assumption 5 holds, then the dual uncertainty quantification problem (D-UQ) is equivalent to the finite convex program*

$$\begin{aligned}
\inf \quad & \sum_{k=1}^K \varepsilon_k \lambda_k + \sum_{k=1}^K \sum_{j=1}^{N_k} p_{k,j} \gamma_{k,j} \\
\text{s.t.} \quad & \lambda_k \in \mathbb{R}_+, \gamma_k \in \mathbb{R}^{N_k}, u_{\alpha,l} \in \mathbb{R}^d, v_{\alpha,l} \in \mathbb{R}^d, w_{\alpha,l,k} \in \mathbb{R}^d \quad \forall \alpha \in \mathcal{A}, \forall k \in [K], \forall l \in [L] \\
& [-\ell_l]^*(u_{\alpha,l}) + \sigma_{\Xi}(v_{\alpha,l}) + \sum_{k=1}^K \lambda_k c^* \left(\frac{w_{\alpha,l,k}}{\lambda_k}, \widehat{\xi}_{k,\alpha_k} \right) \leq \sum_{k=1}^K \gamma_{k,\alpha_k} \quad \forall \alpha \in \mathcal{A}, \forall l \in [L] \\
& u_{\alpha,l} + v_{\alpha,l} + \sum_{k=1}^K w_{\alpha,l,k} = 0 \quad \forall \alpha \in \mathcal{A}, \forall l \in [L],
\end{aligned} \tag{11}$$

where $c^*(u, \widehat{\xi})$ is the convex conjugate of the function $c(\xi, \widehat{\xi})$ with respect to its first argument ξ .

We emphasize that problem (11) is indeed convex because the conjugate of the convex component function $-\ell_l$, the support function of the convex set Ξ and the perspective of the partial conjugate of the convex cost function c are all convex. However, the number of decision variables and constraints of problem (11) scales with $|\mathcal{A}|$ and is thus exponential in the number K of reference distributions.

The following corollary derives a simplification of problem (11) under polyhedrality conditions.

COROLLARY 2 (Piecewise Affine Loss Function). *Suppose that $\Xi = \{\xi : C\xi \leq g\}$ for some $C \in \mathbb{R}^{m \times d}$ and $g \in \mathbb{R}^m$, $c(\xi, \xi') = \|\xi - \xi'\|^p$ for some norm $\|\cdot\|$ on \mathbb{R}^d and exponent $p \in \mathbb{N} \cup \{\infty\}$ and that $\ell(\xi) = \max_{l \in [L]} \langle a_l, \xi \rangle + b_l$ for some $a_l \in \mathbb{R}^d$ and $b_l \in \mathbb{R}$, $l \in [L]$. Then, problem (11) simplifies to*

$$\begin{aligned}
\inf \quad & \sum_{k=1}^K \varepsilon_k \lambda_k + \sum_{k=1}^K \sum_{j=1}^{N_k} p_{k,j} \gamma_{k,j} \\
\text{s.t.} \quad & \lambda_k \in \mathbb{R}_+, \gamma_k \in \mathbb{R}^{N_k}, z_{\alpha,l} \in \mathbb{R}_+^m, w_{\alpha,l,k} \in \mathbb{R}^d \quad \forall \alpha \in \mathcal{A}, \forall k \in [K], \forall l \in [L] \\
& b_l + \sum_{k=1}^K \left(\langle w_{\alpha,l,k}, \widehat{\xi}_{k,\alpha_k} \rangle + \varphi(q) \lambda_k \left\| \frac{w_{\alpha,l,k}}{\lambda_k} \right\|_*^q \right) + \langle z_{\alpha,l}, g \rangle \leq \sum_{k=1}^K \gamma_{k,\alpha_k} \quad \forall \alpha \in \mathcal{A}, \forall l \in [L] \\
& a_l - C^T z_{\alpha,l} = \sum_{k=1}^K w_{\alpha,l,k} \quad \forall \alpha \in \mathcal{A}, \forall l \in [L]
\end{aligned}$$

where $\|\cdot\|_*$ is the norm dual to $\|\cdot\|$, $q \in \mathbb{N} \cup \{\infty\}$ is the conjugate index of p satisfying $1/p + 1/q = 1$, and $\varphi(q) = (q-1)^{q-1}/q^q$ for $q > 1$ and $\varphi(1) = 1$.

Note that Assumption 5 is automatically satisfied under the conditions of Corollary 2. Note also that Corollary 2 extends (Mohajerin Esfahani and Kuhn 2018, Corollary 5.1) and (Zhen et al. 2023, § 6) to $K > 1$. The proof of Corollary 2 relies on now standard techniques and is thus omitted.

In the remainder of this section we will show that problem (P-UQ) admits a worst-case distribution with (sparse) discrete support.

PROPOSITION 3 (Sparse Worst-Case Distribution). *If Assumptions 1, 3 and 4 hold, then problem (P-UQ) is solved by a discrete distribution supported on at most $1 + \sum_{k=1}^K N_k$ points.*

Proposition 3 can be seen as a generalization of (Gao and Kleywegt 2023, Corollary 2), which focuses on the special case when $K = 1$; see also (Yue et al. 2022, Theorem 4). It is perhaps surprising that problem (P-UQ) admits a worst-case distribution with only $\mathcal{O}(\sum_{k=1}^K N_k)$ atoms even though the convex reformulation (11) of (P-UQ) involves $\mathcal{O}(\prod_{k=1}^K N_k)$ variables and constraints.

Proof of Proposition 3 Assumptions 1, 3 and 4 imply via a straightforward generalization of (Yue et al. 2022, Theorem 3) that (P-UQ) is solvable. In addition, from the proof of Theorem 2 we know that (P-UQ) is equivalent to problem (8). By using the definition of $\Pi(\mathbb{P}, \widehat{\mathbb{P}}_1, \dots, \widehat{\mathbb{P}}_K)$ and eliminating the decision variable \mathbb{P} , one then easily verifies that problem (8) is equivalent to

$$\begin{aligned}
 & \sup \int_{\Xi^{K+1}} \ell(\xi) \, d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) \\
 & \text{s.t. } \bar{\pi} \in \mathcal{M}_+(\Xi^{K+1}) \\
 & \int_{\Xi^{K+1}} d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) = 1 \\
 & \int_{\Xi^{K+1}} \mathbf{1}_{\xi_k = \widehat{\xi}_{k,j}} \, d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) = p_{k,j} \quad \forall j \in [N_k], \forall k \in [K] \\
 & \int_{\Xi^{K+1}} c(\xi, \xi_k) \, d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) \leq \varepsilon_k \quad \forall k \in [K].
 \end{aligned} \tag{12}$$

Note that problem (12) optimizes over the convex cone of nonnegative Borel measures on Ξ^{K+1} . The normalization constraint ensures that each feasible Borel measure $\bar{\pi}$ is in fact a probability distribution. As the probabilities $\{p_{k,j}\}_{j=1}^{N_k}$ associated with the atoms of the k -th reference distribution $\widehat{\mathbb{P}}_k$ sum to 1, the probability constraints ensure together with the normalization constraint that any feasible probability distribution $\bar{\pi}$ is supported on $Z = \Xi \times \widehat{\Xi}_1 \times \dots \times \widehat{\Xi}_K$, where $\widehat{\Xi}_k = \{\widehat{\xi}_{k,j}\}_{j=1}^{N_k}$

for every $k \in [K]$. We may thus restrict the integration domains of all integrals in (12) to Z without restricting the problem's feasible set. As $\sum_{j=1}^{N_k} p_{k,j} = 1$ for all $k \in [K]$, it then becomes evident that the normalization constraint makes one of the probability constraints redundant for every $k \in [K]$.

In summary, these arguments imply that problem (12) can be reformulated as

$$\begin{aligned}
& \sup \int_Z \ell(\xi) d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) \\
& \text{s.t. } \bar{\pi} \in \mathcal{M}_+(Z) \\
& \int_Z d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) = 1 \\
& \int_Z \mathbb{1}_{\xi_k = \hat{\xi}_{k,j}} d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) = p_{k,j} \quad \forall j \in [N_k - 1], \forall k \in [K] \\
& \int_Z c(\xi, \xi_k) d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) \leq \varepsilon_k \quad \forall k \in [K].
\end{aligned} \tag{13}$$

Consider now any maximizer $\bar{\pi}^*$ of problem (13), which exists because (P-UQ) is solvable by assumption and because (13) is equivalent to (P-UQ). Next, define $\varepsilon_k^* = \int_Z c(\xi, \xi_k) d\bar{\pi}^*(\xi, \xi_1, \dots, \xi_K)$, and note that $\varepsilon_k^* \leq \varepsilon_k$ for every $k \in [K]$. By construction, problem (13) is thus equivalent to

$$\begin{aligned}
& \sup \int_Z \ell(\xi) d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) \\
& \text{s.t. } \bar{\pi} \in \mathcal{M}_+(Z) \\
& \int_Z d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) = 1 \\
& \int_Z \mathbb{1}_{\xi_k = \hat{\xi}_{k,j}} d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) = p_{k,j} \quad \forall j \in [N_k - 1], \forall k \in [K] \\
& \int_Z c(\xi, \xi_k) d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) = \varepsilon_k^* \quad \forall k \in [K],
\end{aligned} \tag{14}$$

which is also solved by $\bar{\pi}^*$. Problem (14) can be viewed as an infinite-dimensional linear program in standard form. Indeed, the decision variable $\bar{\pi}$ is subject to a continuum of nonnegativity constraints as well as to a finite number of linear equality constraints. Standard-form linear programs in finitely many dimensions are solved by basic feasible solutions. Intuitively, in the context of the infinite-dimensional linear program (14), a basic feasible solution is a discrete Borel measure whose number of atoms equals the number $N = 1 + \sum_{k=1}^K N_k$ of linear equality constraints. This intuition can be

formalized. Specifically, by (Pinelis 2016, Corollary 5 and Proposition 6(v)) and (Yue et al. 2022, Proposition 1), problem (14) is equivalent to

$$\begin{aligned}
 & \sup \int_Z \ell(\xi) d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) \\
 & \text{s.t. } \bar{\pi} \in \mathcal{D}_N(Z) \\
 & \int_Z d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) = 1 \\
 & \int_Z \mathbb{1}_{\xi_k = \hat{\xi}_{k,j}} d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) = p_{k,j} \quad \forall j \in [N_k - 1], \forall k \in [K] \\
 & \int_Z c(\xi, \xi_k) d\bar{\pi}(\xi, \xi_1, \dots, \xi_K) = \varepsilon_k^* \quad \forall k \in [K],
 \end{aligned} \tag{15}$$

where $\mathcal{D}_N(Z)$ denotes the set of all discrete Borel measures in $\mathcal{M}_+(Z)$ that are supported on at most N points. It remains to be shown that the supremum of problem (15) is attained. By using similar techniques as in (Yue et al. 2022, § 5), however, it is easy to show that the objective function and the feasible set of problem (15) are weakly upper semi-continuous and weakly compact, respectively. The solvability of problem (15) thus follows from Weierstrass' theorem.

In summary, we have shown that problems (12)–(15) all admit a minimizer $\bar{\pi}^*$ supported on at most N points. Hence, the marginal distribution \mathbb{P}^* of ξ under $\bar{\pi}^*$ solves the uncertainty quantification problem (P-UQ) and is also a discrete distribution supported on at most N points. ■

To close this section, we highlight the striking similarity between our dual uncertainty quantification problem (D-UQ) and the following dual of the OT barycenter problem (OT-BC) with discrete marginals $\mathbb{P}_k = \hat{\mathbb{P}}_k$, $k \in [K]$, which is derived in (Altschuler and Boix-Adsera 2021, p. 5).

$$\begin{aligned}
 & \sup \sum_{k=1}^K \sum_{j=1}^{N_k} p_{k,j} \gamma_{k,j} \\
 & \text{s.t. } \gamma_k \in \mathbb{R}^{N_k} \quad \forall k \in [K] \\
 & \inf_{\xi \in \Xi} \sum_{k=1}^K \lambda_k c(\xi, \hat{\xi}_{k,\alpha_k}) \geq \sum_{k=1}^K \gamma_{k,\alpha_k} \quad \forall \alpha \in \mathcal{A}
 \end{aligned}$$

This problem is equivalent to (D-UQ) if the the Lagrange multipliers λ_k , $k \in [K]$, are identified with the (fixed) weights of the reference distributions in (OT-BC) and if the loss function $\ell(\xi)$ is set to 0.

4. Computational Complexity of Multi-Source DRO

We now demonstrate that the primal uncertainty quantification problem (P-UQ), which constitutes a critical component of multi-source DRO, is generically intractable. We then show that (P-UQ)

becomes tractable under widely used convexity assumptions if either the number K of OT ambiguity sets or the dimension d of the uncertain parameter ξ is kept constant. The main results of this section are inspired by Altschuler and Boix-Adsera (2021). We thus relegate their proofs to Appendix A.

We first prove that the following feasibility version of (P-UQ) is already NP-hard.

UQ FEASIBILITY

Instance. Discrete distributions $\widehat{\mathbb{P}}_k \in \mathcal{P}(\mathbb{R}^d)$ and radii $\varepsilon_k \geq 0$ for all $k \in [K]$.

Goal. For $\Xi = \mathbb{R}^d$, $\ell(\xi) = 0$ and $c(\xi, \xi') = \|\xi - \xi'\|_2^2$, decide if the primal uncertainty quantification problem (P-UQ) is feasible and, if not, find an infeasibility certificate given by a recession direction $(\lambda^\infty, \gamma^\infty)$ of the dual feasible set along which the objective function of (D-UQ) strictly decreases.

Note that if $(\lambda^\infty, \gamma^\infty)$ is an infeasibility certificate, then $(\lambda^0, \gamma^0) + t \cdot (\lambda^\infty, \gamma^\infty)$ is feasible in (D-UQ) for every (λ^0, γ^0) feasible in (D-UQ) and for every $t \geq 0$, and $\sum_{k=1}^K \varepsilon_k \lambda_k^\infty + \sum_{k=1}^K \sum_{j=1}^{N_k} p_{k,j} \gamma_{k,j}^\infty < 0$. The UQ FEASIBILITY problem is well-defined by virtue of Theorem 2, which asserts that (P-UQ) and (D-UQ) are strong duals for all radii $\varepsilon_k \geq 0$, $k \in [K]$. Indeed, this implies that an infeasibility certificate $(\lambda^\infty, \gamma^\infty)$ is guaranteed to exist whenever (P-UQ) is infeasible.

THEOREM 4 (NP-Hardness). *The UQ FEASIBILITY problem is NP-hard.*

We now show that (P-UQ) becomes tractable under standard regularity conditions if either K or d is fixed. This finding mirrors the known tractability results for OT barycenters by Altschuler and Boix-Adsera (2021). For simplicity of exposition, we assume throughout this discussion that $N_k = N$ for all $k \in [K]$. We also assume that $\log U$ represents an upper bound on the number of bits needed to encode any radius ε_k , any probability $p_{k,j}$ and any component of $\widehat{\xi}_{k,j}$ for all $j \in [N]$ and $k \in [K]$. Our tractability results rely on three minimal assumptions. All of them stipulate that certain elementary computations involving the problem data can be carried out in polynomial time. We tacitly assume that the polynomials bounding the worst-case runtimes of these computations are *known*. The first assumption concerns the generalized Moreau envelope problem shown below.

$$\ell^*(\lambda) = \sup_{\xi \in \Xi} \ell(\xi) - \sum_{k=1}^K \lambda_k c(\xi, \widehat{\xi}_k) \quad (16)$$

ASSUMPTION 6 (Generalized Moreau Envelope). *There exists an oracle for solving (16). If $\ell^*(\lambda) < \infty$, then the oracle outputs an optimal solution of (16). If $\ell^*(\lambda) = \infty$, then the oracle outputs a closed halfspace $\Lambda \subseteq \mathbb{R}^K$ such that $\lambda \notin \Lambda$ and $\lambda' \in \Lambda$ for every $\lambda' \in \mathbb{R}^K$ with $\ell^*(\lambda') < \infty$. The oracle runs in time polynomial in $d, K, N, \log U$ and the bit sizes of $\lambda_1, \dots, \lambda_K \geq 0$ and $\widehat{\xi}_1, \dots, \widehat{\xi}_K \in \mathbb{R}^d$.*

Assumption 6 will allow us to construct an efficient separation oracle for the feasible set of problem (D-UQ). It is satisfied, for example, if c is a convex function with efficiently computable subgradients, ℓ is the pointwise maximum of finitely many concave functions with efficiently computable supergradients and Ξ is a compact convex set that admits an efficient separation oracle. In this case, the generalized Moreau envelope problem (16) is susceptible to the ellipsoid algorithm.

We emphasize that Assumption 6 is restrictive in that it requires the supremum of problem (16) to be attained and a corresponding maximizer to be computable *exactly* and in polynomial time. As we will see in Remark 1 below, the tractability results of this section remain valid even if we have only access to an *approximate* Moreau envelope oracle. For ease of exposition, however, we present the key ideas of this section under the simplifying assumption that an exact oracle is available.

The next assumption requires the loss function to admit an efficient zeroth-order oracle.

ASSUMPTION 7 (Computational Properties of ℓ). *The loss $\ell(\xi)$ can be computed in time polynomial in the bit length of $\xi \in \mathbb{R}^d$.*

The last assumption concerns the transportation cost function c . Before we can state this assumption, we first need to introduce the notion of a c -diagram.

DEFINITION 2 (c -Diagram). *For any transportation cost function c satisfying Assumption 1, the c -diagram associated with points $\xi_1, \dots, \xi_N \in \Xi$ and weights $\gamma_1, \dots, \gamma_N \in \mathbb{R}$ is the family of cells*

$$\Xi_i = \{\xi \in \Xi : c(\xi, \xi_i) - \gamma_i < c(\xi, \xi_j) - \gamma_j \quad \forall j \in [N], j \neq i\} \quad \forall i \in [N].$$

Note that if $c(\xi, \xi') = \|\xi - \xi'\|_2^2$ is the quadratic transportation cost function and all weights $\gamma_i, i \in [N]$, are equal, then the c -diagram coincides with the Voronoi diagram induced by $\xi_i, i \in [N]$.

ASSUMPTION 8 (**Computational Properties of c**). *The following hold.*

(i) *The cost $c(\xi, \xi')$ can be computed in time polynomial in the bit length of $\xi, \xi' \in \mathbb{R}^d$.*

(ii) *If $\{\Xi_i\}_{i=1}^N$ is any c -diagram, then $\cup_{i=1}^N \text{cl}(\Xi_i) = \Xi$.*

(iii) *If $\{\Xi_{k,i}\}_{i=1}^N$, $k \in [K]$, are K different c -diagrams, then the set $\mathcal{A}' = \{\alpha \in \mathcal{A} : \mathcal{C}_\alpha \neq \emptyset\}$ defined in terms of the cells $\mathcal{C}_\alpha = \cap_{k=1}^K \Xi_{k,\alpha_k}$, $\alpha \in \mathcal{A}$, can be computed in time $\text{poly}(N, K, \log U)$.*

Assumption 8(i) is unrestrictive and usually trivial to check. Assumptions 8(ii) and (iii) will be instrumental for constructing a separation oracle for the feasible set of (D-UQ) that runs in time $\text{poly}(N, K, \log U)$. One readily verifies that Assumption 8(ii) holds whenever c is convex and satisfies Assumption 1. By (Altschuler and Boix-Adsera 2021, Lemma 18), Assumption 8(iii) is satisfied if $c(\xi, \xi') = \|\xi - \xi'\|_2^2$. In addition, it is also satisfied, for example, if $c(\xi, \xi')$ is given by $\|\xi - \xi'\|_1$, $\|\xi - \xi'\|_\infty$ or $\|\xi - \xi'\|_2$; see (Altschuler and Boix-Adsera 2021, Section 5.1).

We are now ready to state our main tractability result for problem (D-UQ).

THEOREM 5 (**Tractability of (D-UQ)**). *Suppose that Assumptions 6, 7 and 8(i) hold, and there is $\delta > 0$ such that all δ -optimal solutions of (D-UQ) belong to the ball of radius R around 0, where R is computable in time $\text{poly}(N, d, K, \log U)$. Then, the optimal value of (D-UQ) can be computed to any accuracy $\delta > 0$ in time $\text{poly}(N, d, \log U, \log 1/\delta)$ (with exponential dependence on K). If additionally Assumptions 8(ii) and 8(iii) hold, then the optimal value of (D-UQ) can be computed to any accuracy $\delta > 0$ in time $\text{poly}(N, K, \log U, \log 1/\delta)$ (with exponential dependence on d).*

Theorem 5 is proved by showing that—under appropriate conditions—the feasible set of problem (D-UQ) admits two different separation oracles with exponential dependence on d or K , respectively. Thus, given either of the two separation oracles, (D-UQ) can be solved with the ellipsoid algorithm (Ben-Tal and Nemirovski 2001, Theorem 5.2.1). The polynomial time solvability of problem (D-UQ) for fixed K is expected in view of Theorem 3, which reformulates (D-UQ) as a convex program with $\mathcal{O}(N^K \cdot d \cdot L)$ decision variables and constraints. However, the insight that (D-UQ) can sometimes even be solved in time polynomial in K for fixed d is perhaps more surprising and mirrors recent results by Altschuler and Boix-Adsera (2021) for Wasserstein barycenters. The assumption

that R can be computed in time $\text{poly}(d, K, N, \log U)$ seems difficult to check. One can show, however, that this assumption holds if there exists a discrete distribution $\mathbb{P}_S \in \mathcal{P}(\Xi)$ with $C(\mathbb{P}_S, \widehat{\mathbb{P}}_k) < \varepsilon_k$ for all $k \in [K]$, which can be encoded in $\text{poly}(d, K, N, \log U)$ bits, and if Assumption 4 holds with a known growth constant g' , which can be encoded in $\log U$ bits. Details are omitted for brevity.

REMARK 1 (WEAK SEPARATION ORACLES). Assumption 6 is too restrictive for many problems of practical interest. Indeed, it is easy to construct instances of (D-UQ) for which the optimal values of the embedded maximization problems are irrational or not even attained. In these cases, the Moreau envelope oracle cannot run in polynomial time or does not even exist, respectively. However, Theorem 5 remains valid if the Moreau envelope oracle of Assumption 6 outputs only a δ -optimal solution of the underlying optimization problem for any prescribed tolerance $\delta > 0$ (Grötschel et al. 2012, Corollary 4.2.7). This relaxation complicates the proof of Theorem 5 but is standard. \blacktriangle

5. Statistical Guarantees

We now investigate the relation between the multi-source DRO problem (MS-DRO) and the original stochastic optimization problem (5). To this end, we distinguish five scenarios characterized by increasingly realistic informational assumptions about the unknown target distribution \mathbb{P} .

Scenario 1. Assume that we know K source distributions \mathbb{P}_k , $k \in [K]$, which are all different from \mathbb{P} . In addition, even though \mathbb{P} is unknown, the p -Wasserstein distance between \mathbb{P} and \mathbb{P}_k has a *known* upper bound $r_k \geq 0$ for all $k \in [K]$. As an example, imagine that a retailer expands into a new country, where the demands of the offered products follow an unknown distribution \mathbb{P} . However, the retailer is already active in K other countries and knows the respective demand distributions \mathbb{P}_k , $k \in [K]$. While demand distributions invariably differ across countries, it is reasonable to assume that each \mathbb{P}_k provides *some* useful information about \mathbb{P} . Specifically, the retailer might know that $W_p(\mathbb{P}, \mathbb{P}_k) \leq r_k$, where W_p denotes the p -Wasserstein distance on $\mathcal{P}_p(\Xi)$ induced by a metric d on Ξ . Hence, \mathbb{P} is known to belong to $\bigcap_{k=1}^K \mathbb{B}_{r_k}^p(\mathbb{P}_k)$, where $\mathbb{B}_{r_k}^p(\mathbb{P}_k) = \{\mathbb{P} \in \mathcal{P}(\Xi) : W_p(\mathbb{P}, \mathbb{P}_k) \leq r_k\}$ is the p -Wasserstein ball of radius r_k around \mathbb{P}_k . The multi-source DRO problem with ambiguity set $\bigcap_{k=1}^K \mathbb{B}_{r_k}^p(\mathbb{P}_k)$ thus provides a deterministic upper bound on the corresponding stochastic program.

Scenario 2. As in scenario 1, we continue to assume that $W_p(\mathbb{P}, \mathbb{P}_k)$ has a *known* upper bound $r_k \geq 0$. In contrast to scenario 1, however, we no longer assume that \mathbb{P}_k is known. Instead, we only assume to have access to N_k independent samples $\widehat{\xi}_{k,j}$, $j \in [N_k]$, from \mathbb{P}_k . In the remainder of this section, we denote the corresponding empirical distribution by $\widehat{\mathbb{P}}_k = \frac{1}{N_k} \sum_{j=1}^{N_k} \delta_{\widehat{\xi}_{k,j}}$. Note that $\widehat{\mathbb{P}}_k$ constitutes a random object governed by the joint distribution $\mathbb{P}_k^{N_k}$ of the given independent samples from \mathbb{P}_k . It is well known that $\widehat{\mathbb{P}}_k$ approximates \mathbb{P}_k in p -Wasserstein distance provided that \mathbb{P}_k is light-tailed.

DEFINITION 3 (Light-Tailed Distributions). *A distribution $\mathbb{P} \in \mathcal{P}(\Xi)$ is (a, A) -light tailed for some positive constants $a, A > 0$ if $\mathbb{E}_{\mathbb{P}}[\exp(\|\xi\|^a)] \leq A$.*

Most of the subsequent results will depend on the following assumption.

ASSUMPTION 9 (Light-Tailed Reference Distributions). *There exist $a > p$ and $A > 0$ such that all reference distributions \mathbb{P}_k , $k \in [K]$, are (a, A) -light tailed.*

LEMMA 3 (Concentration Inequalities Fournier and Guillin (2015)). *Suppose that $p \neq d/2$ and that Assumption 9 holds. Then, there exist constants $c_1, c_2 > 0$ that depend only on a, A and d such that $\mathbb{P}_k^{N_k}[\mathbb{P}_k \in \mathbb{B}_{\varepsilon_k}^p(\widehat{\mathbb{P}}_k)] \geq 1 - \beta(\varepsilon_k, N_k)$ for all $\varepsilon_k \in \mathbb{R}_+$ and $N_k \in \mathbb{N}$, where*

$$\beta(\varepsilon_k, N_k) = \begin{cases} c_1 \exp(-c_2 N_k \varepsilon_k^{\max\{d/p, 2\}}) & \text{if } \varepsilon_k \leq 1, \\ c_1 \exp(-c_2 N_k \varepsilon_k^{a/p}) & \text{if } \varepsilon_k > 1. \end{cases}$$

Lemma 3 readily extends to $p = d/2$ at the expense of a more intricate formula for $\beta(\varepsilon_k, N_k)$. Given a fixed significance level $\beta_k \in (0, 1]$, one can solve the equation $\beta(\varepsilon_k, N_k) = \beta_k$ for ε_k to obtain

$$\varepsilon(\beta_k, N_k) = \begin{cases} \left(\frac{\log(c_1/\beta_k)}{c_2 N_k} \right)^{\min\{p/d, 1/2\}} & \text{if } N_k \geq \frac{\log(c_1/\beta_k)}{c_2}, \\ \left(\frac{\log(c_1/\beta_k)}{c_2 N_k} \right)^{p/a} & \text{if } N_k < \frac{\log(c_1/\beta_k)}{c_2}. \end{cases}$$

Lemma 3 then implies that the p -Wasserstein ball $\mathbb{B}_{\varepsilon_k}^p(\widehat{\mathbb{P}}_k)$ represents a $(1 - \beta_k)$ -confidence set for \mathbb{P}_k whenever its radius satisfies $\varepsilon_k \geq \varepsilon(\beta_k, N_k)$; see also (Kuhn et al. 2019, Theorem 18).

The following proposition leverages Lemma 3 to construct a confidence set for the unknown distribution \mathbb{P} . This is possible even though there is no data from \mathbb{P} . In the following, we use $\mathbb{P}_{[K]} = \otimes_{k=1}^K \mathbb{P}_k^{N_k}$ as a shorthand for the joint distribution of the samples $\widehat{\xi}_{k,j}$ for all $j \in [N_k]$ and $k \in [K]$.

PROPOSITION 4 (Known Distribution Shifts). *Suppose that $p \neq d/2$ and that Assumption 9 holds. If $W_p(\mathbb{P}_k, \mathbb{P}) \leq r_k$ and if $\varepsilon_k \geq r_k + \varepsilon(\beta_k, N_k)$ for some $\beta_k \in (0, 1]$ and for all $k \in [K]$, then*

$$\mathbb{P}_{[K]} \left(\mathbb{P} \in \bigcap_{k=1}^K \mathbb{B}_{\varepsilon_k}^p(\widehat{\mathbb{P}}_k) \right) \geq 1 - \sum_{k=1}^K \beta_k.$$

The ambiguity set $\bigcap_{k=1}^K \mathbb{B}_{\varepsilon_k}^p(\widehat{\mathbb{P}}_k)$ inherits the randomness of the empirical distributions. In particular, it may be empty with a small positive probability. Depending on the realizations of $\widehat{\mathbb{P}}_k$, $k \in [K]$, one should thus select sufficiently large radii $\varepsilon_k \geq r_k + \varepsilon(\beta_k, N_k)$ that lead to a nonempty ambiguity set.

Proof of Proposition 4 By the definition of the p -Wasserstein balls, we have

$$\begin{aligned} \mathbb{P}_{[K]} \left(\mathbb{P} \in \bigcap_{k=1}^K \mathbb{B}_{\varepsilon_k}^p(\widehat{\mathbb{P}}_k) \right) &= \mathbb{P}_{[K]} \left(W_p(\mathbb{P}, \widehat{\mathbb{P}}_k) \leq \varepsilon_k \quad \forall k \in [K] \right) \\ &\geq 1 - \sum_{k=1}^K \mathbb{P}_k^{N_k} \left(W_p(\mathbb{P}, \widehat{\mathbb{P}}_k) > \varepsilon_k \right) \\ &\geq 1 - \sum_{k=1}^K \mathbb{P}_k^{N_k} \left(W_p(\mathbb{P}, \mathbb{P}_k) + W_p(\mathbb{P}_k, \widehat{\mathbb{P}}_k) > \varepsilon_k \right) \geq 1 - \sum_{k=1}^K \beta_k, \end{aligned}$$

where the first inequality exploits De Morgan's laws and the union bound, whereas the second inequality follows from the triangle inequality. The third inequality holds thanks to Lemma 3, which applies because $\varepsilon_k - W_p(\mathbb{P}, \mathbb{P}_k) \geq \varepsilon_k - r_k \geq \varepsilon(\beta_k, N_k)$ for all $k \in [K]$. Thus, the claim follows. ■

Proposition 4 implies that the optimal value of the multi-source DRO problem (MS-DRO) with ambiguity set $\bigcap_{k=1}^K \mathbb{B}_{\varepsilon_k}^p(\widehat{\mathbb{P}}_k)$ provides a $(1 - \sum_{k=1}^K \beta_k)$ -upper confidence bound on the optimal value of the original stochastic optimization problem (5). Note that different choices of the parameters β_k , $k \in [K]$, yield different bounds with the same confidence level. Hence, these bounds can be tuned.

Scenario 3. As in scenario 2, we continue to assume that \mathbb{P} as well as all reference distributions \mathbb{P}_k , $k \in [K]$, are unknown. In contrast to scenario 2, however, we now assume that $\mathbb{P} = \mathbb{P}_1$ and that we have *no* information about the p -Wasserstein distance between \mathbb{P} and \mathbb{P}_k for $k \geq 2$. Instead, we only have access to independent samples from all source distributions and can thus construct empirical distributions $\widehat{\mathbb{P}}_k$, $k \in [K]$, as in scenario 2. It is now natural to set $\varepsilon_k = W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) + \varepsilon(\beta_1, N_1)$, $k \in [K]$, for some $\beta_1 \in (0, 1]$. In this case, one can show that the ambiguity set $\bigcap_{k=1}^K \mathbb{B}_{\varepsilon_k}^p(\widehat{\mathbb{P}}_k)$ contains \mathbb{P} with probability at least $1 - \beta_1$. To see this, note first that the inequality $W_p(\mathbb{P}, \widehat{\mathbb{P}}_1) \leq r_1$ implies

$$W_p(\mathbb{P}, \widehat{\mathbb{P}}_k) \leq W_p(\mathbb{P}, \widehat{\mathbb{P}}_1) + W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) \leq W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) + r_1 = r_k,$$

where the equality follows from the definitions of r_1 and r_k . Therefore, we have $\mathbb{B}_{\varepsilon_1}^p(\widehat{\mathbb{P}}_1) \subseteq \mathbb{B}_{\varepsilon_k}^p(\widehat{\mathbb{P}}_k)$ for all $k \in [K]$, and thus the ambiguity set $\bigcap_{k=1}^K \mathbb{B}_{\varepsilon_k}^p(\widehat{\mathbb{P}}_k)$ collapses to the crisp Wasserstein ball $\mathbb{B}_{\varepsilon_1}^p(\widehat{\mathbb{P}}_1)$. Lemma 3 then implies that $\mathbb{P} \in \bigcap_{k=1}^K \mathbb{B}_{\varepsilon_k}^p(\widehat{\mathbb{P}}_k)$ with probability at least $1 - \beta_1$. Thus, (MS-DRO) provides again an upper confidence bound on (5). However, the data from the reference distributions \mathbb{P}_k with $k \neq 1$ is wasted. It is unclear how to make use of this data without additional assumptions.

Scenario 4. As in scenario 3, we assume that $\mathbb{P} = \mathbb{P}_1$ and that the reference distributions \mathbb{P}_K , $k \in [K]$, are only accessible via data. In contrast to scenario 3, however, we now adopt a Bayesian perspective, that is, we assume that the decision-maker has a prior belief about the Wasserstein distances between the reference distributions. Arguably, the main motivation for using data from reference distributions is the belief that \mathbb{P} is in some sense “close” to \mathbb{P}_k (and thus also to $\widehat{\mathbb{P}}_k$).

In the following, we treat not only the samples $\{\widehat{\xi}_{k,j}\}_{j=1}^{N_k}$, $k \in [K]$, but also the reference distributions \mathbb{P}_k , $k \in [K]$, as random objects on an abstract probability space $(\Omega, \mathcal{F}, \mu)$. We assume that the distributions \mathbb{P}_k , $k \in [K]$, satisfy Assumption 9 μ -almost surely. The only property of μ to be used below is that, conditional on any fixed realizations of \mathbb{P}_k , $k \in [K]$, all samples are mutually independent under μ , and the conditional distribution of $\widehat{\xi}_{k,j}$ under μ coincides with \mathbb{P}_k for all $j \in [N_K]$ and $k \in [K]$. We also define F_k as the cumulative distribution function of $W_p(\mathbb{P}_1, \mathbb{P}_k)$ under μ . This function captures the decision-maker’s prior beliefs about the proximity between \mathbb{P}_1 and \mathbb{P}_k .

In Bayesian statistics, the prior beliefs are updated when new evidence emerges. In the example at hand, the decision-maker observes the p -Wasserstein distance between $\widehat{\mathbb{P}}_1$ and $\widehat{\mathbb{P}}_k$, which we denote as $\widehat{r}_k \in \mathbb{R}_+$. The prior beliefs can then be updated by conditioning on the event $W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) \leq \widehat{r}_k$.

REMARK 2. Note that conditioning on the event $W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) = \widehat{r}_k$ would provide a stronger update of the prior. We nevertheless prefer to condition on the less informative event $W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) \leq \widehat{r}_k$ for two reasons. First, it enhances the transparency of the subsequent derivations and obviates the need to mobilize advanced measure-theoretic tools. In addition, we aim to derive an upper confidence bound on $W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_k)$. As \mathbb{P}_1 is close to $\widehat{\mathbb{P}}_1$ by virtue of Lemma 3, such a confidence bound is within reach if $\widehat{\mathbb{P}}_1$ and $\widehat{\mathbb{P}}_k$ are known to be close. The extra information that $\widehat{\mathbb{P}}_1$ and $\widehat{\mathbb{P}}_k$ are “not *too* close,” which is available in the event $W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) = \widehat{r}_k$, cannot meaningfully strengthen this bound. \blacktriangle

PROPOSITION 5 (Bayesian Measure Concentration). *If μ and F_k , $k \in [K]$, are constructed as above, then we have $\mu(\mathbb{P}_1 \in \mathbb{B}_{\varepsilon_k}^p(\widehat{\mathbb{P}}_k) | W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) \leq \widehat{r}_k) \geq 1 - \beta_{F_k}(\varepsilon_k, \widehat{r}_k, N_1, N_k)$ for all $\varepsilon_k, \widehat{r}_k \in \mathbb{R}_+$ with $\varepsilon_k \geq \widehat{r}_k$ and for all $N_1, N_k \in \mathbb{N}$, $k \in [K]$, where*

$$\beta_{F_k}(\varepsilon_k, \widehat{r}_k, N_1, N_k) = \frac{\beta(\varepsilon_k - \widehat{r}_k, N_1) \cdot \left(\int_0^{\varepsilon_k} \beta(\varepsilon_k - r_k, N_k) dF_k(r_k) + 1 - F_k(\varepsilon_k) \right)}{\mu\left(W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) \leq \widehat{r}_k\right)}. \quad (17)$$

Proof of Proposition 5 Bayes' rule implies that

$$\begin{aligned} & \mu\left(W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_k) > \varepsilon_k \mid W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) \leq \widehat{r}_k\right) \\ &= \frac{\mu\left(W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) \leq \widehat{r}_k \mid W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_k) > \varepsilon_k\right) \cdot \mu\left(W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_k) > \varepsilon_k\right)}{\mu\left(W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) \leq \widehat{r}_k\right)}. \end{aligned} \quad (18)$$

The evidence in the denominator of the above expression is strictly positive for otherwise it would have been (μ -almost surely) impossible to observe \widehat{r}_k . The likelihood term in the numerator satisfies

$$\begin{aligned} & \mu\left(W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) \leq \widehat{r}_k \mid W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_k) > \varepsilon_k\right) \\ & \leq \mu\left(W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_k) - W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_1) \leq \widehat{r}_k \mid W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_k) > \varepsilon_k\right) \\ & \leq \mu\left(W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_1) > \varepsilon_k - \widehat{r}_k \mid W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_k) > \varepsilon_k\right) \\ & = \mathbb{E}_\mu \left[\mu\left(W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_1) > \varepsilon_k - \widehat{r}_k \mid \mathbb{P}_1, \dots, \mathbb{P}_K\right) \mid W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_k) > \varepsilon_k \right] \\ & \leq \beta(\varepsilon_k - \widehat{r}_k, N_1), \end{aligned}$$

where the equality follows from the law of total expectation and the observation that $\widehat{\mathbb{P}}_1$ and $\widehat{\mathbb{P}}_k$ are independent under μ conditional on any fixed realizations of the reference distributions $\mathbb{P}_{k'}$, $k' \in [K]$. The last inequality follows from Lemma 3, which applies because, conditional on any fixed realizations of the reference distributions, $\widehat{\mathbb{P}}_1$ is governed by the distribution $\mathbb{P}_1^{N_1}$ under μ and because the reference distributions satisfy Assumption 9 μ -almost surely. Similarly, the prior probability in the numerator of (18) can be re-expressed as

$$\begin{aligned} \mu\left(W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_k) > \varepsilon_k\right) & \leq \mu\left(W_p(\mathbb{P}_1, \mathbb{P}_k) + W_p(\mathbb{P}_k, \widehat{\mathbb{P}}_k) > \varepsilon_k\right) \\ & = \int_0^\infty \mu\left(W_p(\mathbb{P}_1, \mathbb{P}_k) + W_p(\mathbb{P}_k, \widehat{\mathbb{P}}_k) > \varepsilon_k \mid W_p(\mathbb{P}_1, \mathbb{P}_k) = r_k\right) dF_k(r_k) \\ & = \int_0^\infty \mu\left(W_p(\mathbb{P}_k, \widehat{\mathbb{P}}_k) > \varepsilon_k - r_k \mid W_p(\mathbb{P}_1, \mathbb{P}_k) = r_k\right) dF_k(r_k). \end{aligned}$$

If $\varepsilon_k \geq r_k$, then the integrand of the last expression can be further simplified to

$$\begin{aligned} & \mu \left(W_p(\mathbb{P}_k, \widehat{\mathbb{P}}_k) > \varepsilon_k - r_k \mid W_p(\mathbb{P}_1, \mathbb{P}_k) = r_k \right) \\ &= \mathbb{E}_\mu \left[\mu \left(W_p(\mathbb{P}_k, \widehat{\mathbb{P}}_k) > \varepsilon_k - r_k \mid \mathbb{P}_1, \dots, \mathbb{P}_K \right) \mid W_p(\mathbb{P}_1, \mathbb{P}_k) = r_k \right] \leq \beta(\varepsilon_k - r_k, N_k), \end{aligned}$$

where the first equality exploits the law of total expectation and the trivial observation that $W_p(\mathbb{P}_1, \mathbb{P}_k)$ reduces to a constant when we condition on $\mathbb{P}_{k'}$, $k' \in [K]$. The inequality follows from Lemma 3, which applies because, conditional on any fixed realizations of the reference distributions, $\widehat{\mathbb{P}}_k$ is governed by the distribution $\mathbb{P}_k^{N_k}$ under μ and because the reference distributions satisfy Assumption 9 μ -almost surely. Combining the above estimates shows that the prior satisfies

$$\mu \left(W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_k) > \varepsilon_k \right) \leq \int_0^{\varepsilon_k} \beta(\varepsilon_k - r_k, N_k) dF_k(r_k) + 1 - F_k(\varepsilon_k).$$

Substituting our estimates for the likelihood and the prior into (18) finally yields

$$\mu \left(W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_k) > \varepsilon_k \mid W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) \leq \widehat{r}_k \right) \leq \frac{\beta(\varepsilon_k - \widehat{r}_k, N_1) \cdot \left(\int_0^{\varepsilon_k} \beta(\varepsilon_k - r_k, N_k) dF_k(r_k) + 1 - F_k(\varepsilon_k) \right)}{\mu \left(W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) \leq \widehat{r}_k \right)}.$$

The last expression coincides with $\beta_F(\varepsilon_k, \widehat{r}_k, N_1, N_k)$. We may therefore conclude that

$$\begin{aligned} \mu \left(\mathbb{P}_1 \in \mathbb{B}_{\varepsilon_k}^p(\widehat{\mathbb{P}}_k) \mid W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) \leq \widehat{r}_k \right) &= 1 - \mu \left(W_p(\mathbb{P}_1, \widehat{\mathbb{P}}_k) > \varepsilon_k \mid W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) \leq \widehat{r}_k \forall k \in [K] \right) \\ &\geq 1 - \beta_{F_k}(\varepsilon_k, \widehat{r}_k, N_1, N_k), \end{aligned}$$

and thus the claim follows. ■

We highlight that $\beta_{F_k}(\varepsilon_k, \widehat{r}_k, N_1, N_k)$ is non-increasing in ε_k . Indeed, the denominator in (17) is positive and independent of ε_k . In addition, both terms in the numerator are nonnegative and non-increasing in ε_k whenever $\varepsilon_k \geq \widehat{r}_k$. Note that the derivative of the second term can be expressed as $\int_0^{\varepsilon_k} \partial_\varepsilon \beta(\varepsilon_k - r_k, N_k) dF_k(r_k)$ by virtue of the Reynolds theorem. Hence, it is strictly negative.

Proposition 5 and the union bound imply that, given any fixed $\beta_k \in (0, 1]$, $k \in [K]$, we have

$$\mu \left(\mathbb{P} \in \bigcap_{k=1}^K \mathbb{B}_{\varepsilon_k}^p(\widehat{\mathbb{P}}_k) \mid W_p(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k) \leq \widehat{r}_k \right) \geq 1 - \sum_{k=1}^K \beta_k$$




Figure 1 Normalized Bayesian significance level as a function of the normalized Wasserstein radius $c_2^{p/a}(\varepsilon_k - \hat{r}_k)$ when there is no prior (solid blue line), a weak prior (dashed orange line) and a strong prior (dotted green line).

whenever $\varepsilon_k \geq \varepsilon_{F_k}(\beta_k, \hat{r}_k, N_1, N_k)$ for all $k \in [K]$, where

$$\varepsilon_{F_k}(\beta_k, \hat{r}_k, N_1, N_k) = \inf \{ \varepsilon_k \geq 0 : \beta_{F_k}(\varepsilon_k, \hat{r}_k, N_1, N_k) \leq \beta_k \}.$$

Thus, (MS-DRO) provides a meaningful upper confidence bound on (5) that does not waste data.

Figure 1 shows the normalized Bayesian significance level $\beta_{F_k}(\varepsilon_k, \hat{r}_k, N_1, N_k) / \beta_{F_k}(\hat{r}_k, \hat{r}_k, N_1, N_k)$ as a function of the normalized Wasserstein radius $c_2^{p/a}(\varepsilon_k - \hat{r}_k)$ for the 2-Wasserstein distance on $\mathcal{P}_2(\mathbb{R}^5)$ when there are $N_1 = 5$ target samples and $N_k = 50$ source samples. Here, \mathbb{P}_1 and \mathbb{P}_k are light-tailed with critical exponent $a = 5$, and the 2-Wasserstein distance between \mathbb{P}_1 and \mathbb{P}_k is believed to follow a normal distribution with mean \hat{r}_k and variance $\sigma^2 = 0.2c_2^{-2p/a}$ (strong prior), $\sigma^2 = 0.5c_2^{-2p/a}$ (weak prior) or $+\infty$ (no prior). The chosen normalizations ensure that the curves in Figure 1 are independent of the unknown constants c_1, c_2 and $\mu(W_2(\hat{\mathbb{P}}_1, \hat{\mathbb{P}}_k) \leq \hat{r}_k)$. As expected, we observe that a stronger prior leads to stronger statistical guarantees.

Scenario 5. Consider now a special case of scenario 4, which arises when the decision-maker observes *no* data from the target distribution $\mathbb{P}_1 = \mathbb{P}$ at all (that is, when $N_1 = 0$). From the proof of Proposition 5 we already know that the prior probability obeys the estimate

$$\mu(W_p(\mathbb{P}_1, \hat{\mathbb{P}}_k) > \varepsilon_k) \leq \int_0^{\varepsilon_k} \beta(\varepsilon_k - r_k, N_k) dF_k(r_k) + 1 - F_k(\varepsilon_k). \tag{19}$$

As there is no data from \mathbb{P}_1 , this prior probability cannot be updated. However, given any significance level $\beta_k \in (0, 1]$, one can increase ε_k until the right hand side of (19) drops below β_k . Constructing ε_k

Scenario	Access to \mathbb{P}	Access to \mathbb{P}_k	Information on distribution shift
1	\times	\checkmark	$W_p(\mathbb{P}, \mathbb{P}_k) \leq r_k$
2	\times	samples from \mathbb{P}_k	$W_p(\mathbb{P}, \mathbb{P}_k) \leq r_k$
3	samples from \mathbb{P}	samples from \mathbb{P}_k	\times
4	samples from \mathbb{P}	samples from \mathbb{P}_k	$W_p(\mathbb{P}, \mathbb{P}_k) \sim F_k$
5	\times	samples from \mathbb{P}_k	$W_p(\mathbb{P}, \mathbb{P}_k) \sim F_k$

Table 1 Comparison of the different scenarios

in this manner for every $k \in [K]$ enables us again to show that (MS-DRO) provides a $(1 - \sum_{k=1}^K \beta_k)$ -upper confidence bound on the optimal value of the original stochastic program.

REMARK 3 (RELATIONSHIPS BETWEEN SCENARIOS). Note first that scenario 1 can be obtained as a special case of scenario 2 if the number of training samples from each source tends to infinity. Next, it is easy to see that scenario 2 emerges as a special case of scenario 5 under Dirac priors, that is, if $F_k(\varepsilon_k) = 0$ for $\varepsilon_k < r_k$ and $F_k(\varepsilon_k) = 1$ for $\varepsilon_k \geq r_k$, $k \in [K]$. Similarly, scenario 4 is obtained from scenario 5 by a Bayesian update after observing samples from \mathbb{P}_1 . Finally, scenario 3 emerges as a special case of scenario 4 as the priors become less informative. The information available to the decision-maker in the different scenarios is summarized in Table 1. \blacktriangle

6. Experiments

We now test the proposed DRO scheme with heterogeneous data sources in the context of a portfolio selection problem (see Section 6.1) as well as an assortment optimization problem (see Section 6.2).

6.1. Portfolio Selection

Consider the problem of investing a given amount of capital into d assets with uncertain rates of return $\xi \in \mathbb{R}^d$. The portfolio weight vector $\theta \in \mathbb{R}^d$ collects the percentage weights of the capital allocated to the available assets. In the absence of short-sales, it ranges over the d -dimensional unit simplex Θ . The goal is to minimize the expected value of the portfolio loss $-\langle \theta, \xi \rangle$ adjusted by the portfolio risk, measured by the conditional value-at-risk (CVaR) of $-\langle \theta, \xi \rangle$ at level $\eta \in (0, 1]$. Denoting the investor's risk aversion as $\rho \geq 0$, the portfolio selection problem can thus be expressed as

$$\inf_{\theta \in \Theta} \mathbb{E}_{\mathbb{P}}[-\langle \theta, \xi \rangle] + \rho \cdot \mathbb{P}\text{-CVaR}_{\eta}(-\langle \theta, \xi \rangle),$$

where \mathbb{P} denotes the target distribution of the asset returns. Replacing the CVaR by its definition due to Rockafellar and Uryasev (2000), this problem simplifies to

$$\inf_{\theta \in \Theta, \tau \in \mathbb{R}} \int_{\Xi} \max_{t \in \{1, 2\}} \{a_t \langle \theta, \xi \rangle + b_t \tau\} d\mathbb{P}(\xi), \quad (20)$$

where $a_1 = -1$, $a_2 = -1 - \frac{\rho}{\eta}$, $b_1 = \rho$, and $b_2 = \rho(1 - \frac{1}{\eta})$; see, *e.g.*, (Mohajerin Esfahani and Kuhn 2018, § 7.1). We henceforth assume that the target distribution \mathbb{P} is unknown but close to $K = 2$ source distributions \mathbb{P}_1 and \mathbb{P}_2 . While \mathbb{P}_1 and \mathbb{P}_2 are unknown, too, the investor has access to respective empirical distributions $\widehat{\mathbb{P}}_1$ and $\widehat{\mathbb{P}}_2$ and can thus construct the following instance of (MS-DRO).

$$\inf_{\theta \in \Theta, \tau \in \mathbb{R}} \sup_{\mathbb{P} \in \bigcap_{k=1}^2 \mathbb{B}_{\varepsilon_k}(\widehat{\mathbb{P}}_k)} \int_{\Xi} \max_{t \in \{1,2\}} \{a_t \langle \theta, \xi \rangle + b_t \tau\} d\mathbb{P}(\xi). \quad (21)$$

The ambiguity set in (21) is defined as the intersection of two 1-Wasserstein balls induced by the transportation cost function $c(\xi, \xi') = \|\xi - \xi'\|_1$. By Corollary 2, problem (21) can be reformulated as a tractable linear program and is thus susceptible to highly efficient optimization algorithms.

6.1.1. Sensitivity Analysis In the first experiment we use synthetic data to analyze the impact of the radii ε_1 and ε_2 on the optimal portfolio. Throughout this experiment we assume that there are $d = 10$ assets. We further assume that the asset returns are mutually independent under both source distributions and that $\xi_i \sim \mathcal{N}(i\%, 1\%)$ under \mathbb{P}_1 , whereas $\xi_i \sim \mathcal{N}((11-i)\%, 1\%)$ under \mathbb{P}_2 for all $i \in [d]$. We also measure risk by the CVaR at level $\eta = 20\%$ and set $\rho = 10$. In addition, we construct two empirical distributions $\widehat{\mathbb{P}}_1$ and $\widehat{\mathbb{P}}_2$ from $N_1 = N_2 = 30$ independent training samples from \mathbb{P}_1 and \mathbb{P}_2 , respectively. Finally, we set the radii of the two 1-Wasserstein balls to

$$\varepsilon_1 = \lambda \cdot (1 + m) \cdot W_1(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_2) \quad \text{and} \quad \varepsilon_2 = (1 - \lambda) \cdot (1 + m) \cdot W_1(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_2),$$

where $m \in \mathbb{R}_+$ captures the absolute and $\lambda \in [0, 1]$ the relative degree of ambiguity of the two source distributions. This construction ensures that $\bigcap_{k=1}^2 \mathbb{B}_{\varepsilon_k}(\widehat{\mathbb{P}}_k)$ is nonempty. Figure 2 visualizes the dependence of the optimal portfolio weights obtained from problem (21) on λ and m . Specifically, Figure 2a shows that, as λ increases (decreases), assets with low (high) indices are assigned more weight. Indeed, the ambiguity set of problem (21) degenerates to $\{\widehat{\mathbb{P}}_2\}$ as λ tends to 1 and to $\{\widehat{\mathbb{P}}_1\}$ as λ tends to 0, and assets with high (low) indices perform well under $\mathbb{P}_1 \approx \widehat{\mathbb{P}}_1$ ($\mathbb{P}_2 \approx \widehat{\mathbb{P}}_2$). Figure 2b shows that, as m increases, the portfolio weights all become approximately equal, that is, the equally-weighted portfolio appears to be optimal under high ambiguity. This observation is in agreement with theoretical findings for distributionally robust single-source portfolio selection problems due to Pflug et al. (2012); see also (Mohajerin Esfahani and Kuhn 2018, Proposition 7.2).




Figure 2 Impact of m and λ on the portfolio weights. The assets are ordered from bottom (asset 1) to top (asset 10).

6.1.2. Backtest on Synthetic Data In the second experiment we compare the optimal portfolios resulting from the multi-source DRO problem (21) against several other distributionally robust portfolios in terms of their out-of-sample means, standard deviations and Sharpe ratios. In contrast to Section 6.1.1, we now assume that the asset returns are representable as $\xi_i = \psi + \zeta_i$, where ψ is a systematic risk factor affecting all assets, whereas ζ_i is an idiosyncratic risk factor affecting only the i -th asset. More specifically, we assume that $\psi \sim \mathcal{N}(0, 2\%)$ and $\zeta_i \sim \mathcal{N}(r_i, 1\%)$, $i \in [10]$, follow independent normal distributions. We use two different models for the expected returns. In model 1, we set $r_i = +0.4\%$ if $i \leq 2$ and $r_i = -0.4\%$ if $i > 2$. In model 2, on the other hand, we set $r_i = +0.2\%$ if $i \leq 5$ and $r_i = -0.4\%$ if $i > 5$. Thus, assets 1 and 2 have high expected returns and assets 6–10 have low expected returns in *both* models. However, assets 3, 4 and 5 have low expected returns in model 1 and high expected returns in model 2. The distributions \mathbb{P}_1 and \mathbb{P}_2 are then defined as the distribution of ξ in model 1 and model 2, respectively. Throughout this experiment we assume that \mathbb{P}_1 coincides with the target distribution \mathbb{P} and that we have access to $N_1 = 5$ samples from \mathbb{P}_1 and to $N_2 = 30$ samples from \mathbb{P}_2 . All other problem parameters are chosen as in Section 6.1.1.

The purpose of this stylized example is to show that if data from the target distribution \mathbb{P}_1 is scarce, then one can leverage data from the source distribution \mathbb{P}_2 to improve the performance of the optimal portfolio. Intuitively, the (abundant) source data provides an initial guess for the unknown target distribution, and the (scarce) target data allows us to improve this initial guess.

We compare the optimal portfolios of the multi-source DRO problem (21) against several baselines. Every baseline portfolio is obtained by solving a single-source version of problem (21), which involves

only a single 1-Wasserstein ball. The center of this ball is set to the empirical distribution on the *target data*, the *source data*, the *pooled data* (that is, the union of the target and source data), or a 1-*Wasserstein barycenter* of the empirical distributions corresponding to the target and the source data. The latter approach is inspired by (Lau and Liu 2022). As in Section 6.1.1, we set the Wasserstein radii in problem (21) to $\varepsilon_1 = \lambda \cdot (1 + m) \cdot W_1(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_2)$ and $\varepsilon_2 = (1 - \lambda) \cdot (1 + m) \cdot W_1(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_2)$, where λ ranges from 0 to 1 in steps of 0.1, and $m \in \{0.002, 0.005, 0.01, 0.02\}$. The Wasserstein radius in the single-source DRO problems used to construct the baseline portfolios is chosen from

$$\{0.005, 0.01, 0.05, 0.1, 0.2, 0.5, 1, 2, 5, 10, 20, 50, 100, 200, 500\}.$$

In all cases, we select hyperparameters from the underlying search grids that maximize the average portfolio return across five (additional) independent validation samples from \mathbb{P}_1 .

For all portfolios under consideration, we compute the expected value, the standard deviation and the Sharpe ratio of the corresponding portfolio return under the target distribution \mathbb{P}_1 . Table 2 reports the means and standard errors of these performance metrics across 10 independent simulation runs. We observe that the multi-source DRO portfolios achieve the highest expected returns and Sharpe ratios. Optimizing only in view of the scarce target data results in noisy portfolios that are susceptible to the optimizer’s curse and thus perform poorly. Similarly, optimizing only in view of the source data results in biased portfolios because of the distribution shift from \mathbb{P}_2 to \mathbb{P}_1 . Indeed, under the source distribution \mathbb{P}_2 , it would be optimal to distribute the available capital evenly across the assets 1–5. However, this strategy invests 3/5 of the available capital in assets that destroy wealth under the target distribution \mathbb{P}_1 . Optimizing in view of the pooled data does not significantly improve performance because the target samples are outnumbered by the source samples and have thus only a minor impact on the optimal portfolio. These shortcomings are mitigated by optimizing in view of a 1-Wasserstein barycenter between the empirical distributions $\widehat{\mathbb{P}}_1$ and $\widehat{\mathbb{P}}_2$. By Corollary 1, however, if $\widehat{\mathbb{P}}_1$ and $\widehat{\mathbb{P}}_2$ are assigned the same weights, then either of them constitutes a 1-Wasserstein barycenter. In the experiment we choose the one that generates the smaller worst-case risk on the validation dataset. The portfolios obtained via multi-source DRO significantly outperform all

Table 2 Out-of-sample performance of optimal portfolios on synthetic data (mean (std. error) over 10 replications)

	Single-source DRO on target data	Single-source DRO on source data	Single-source DRO on pooled data	Single-source DRO on barycenter	Multi-source DRO
Sharpe ratio	0.002 (0.017)	-0.024 (0.009)	-0.021 (0.012)	0.007 (0.014)	0.034 (0.024)
Expected value	0.01 (0.067)	-0.097 (0.036)	-0.084 (0.047)	0.03 (0.059)	0.143 (0.1)
Standard deviation	4.035 (0.018)	4.018 (0.016)	4.02 (0.017)	4.034 (0.018)	4.084 (0.022)

baselines in terms of Sharpe ratio and expected return. On the other hand, their returns display the highest standard deviations. This is neither surprising nor troubling. The optimal portfolio under complete knowledge of the target distribution \mathbb{P}_1 allocates the available capital evenly to assets 1 and 2. While this portfolio evidently maximizes expected return, it suffers from poor diversification. It is easy to show that the lowest possible portfolio standard deviation is obtained by investing equal amounts of money in all assets (even those with negative expected returns).

6.1.3. Backtest on Real Data In the third experiment the asset universe consists of $d = 3$ regional sector indices (*Agriculture and Food Chain*, *Commodity Producers* and *Infrastructure*) provided by MSCI, which follow the Global Industry Classification Standard (GICS). Each index includes large and mid cap companies from a particular geographical region (Europe, USA or Pacific) denominated in US dollars. We believe that the European and US economies are sufficiently similar so that the European index return distributions provide a strong prior for those of the US and vice versa. However, we believe that the index return distributions of the Pacific area are not informative for those of Europe and the US. Note that the three sector indices were chosen because they cover basic industries and thus exist in all considered geographical regions. We then define \mathbb{P}_1 and \mathbb{P}_2 as the index return distributions corresponding to (any) two different regions. These distributions are unknown, but we can construct respective empirical distributions $\widehat{\mathbb{P}}_1$ and $\widehat{\mathbb{P}}_2$ from historical return data available from the MSCI online database (<https://www.msci.com/end-of-day-data-regional>). As in Section 6.1.2, we compare the optimal portfolios of the multi-source DRO problem (21) against those of different single-source DRO baselines. Also, we use the same risk aversion parameters and the same search grids for the hyperparameters (λ and m for multi-source DRO and ε for single-source DRO) as in Section 6.1.2.

Table 3 Out-of-sample Sharpe ratios of optimal portfolios on real data (mean (std. error) over 10 replications)

Source	Target	N_1	Single-s. DRO on target distribution	Single-s. DRO on source distribution	Single-s. DRO on pooled data	Single-s. DRO on barycenter	Multi-s. DRO
Europe	Pacific	5	0.059 (0.004)	0.054 (0.006)	0.052 (0.005)	0.052 (0.003)	0.039 (0.006)
		10	0.047 (0.005)	0.04 (0.004)	0.043 (0.0)	0.036 (0.004)	0.035 (0.006)
		20	0.052 (0.004)	0.049 (0.008)	0.049 (0.006)	0.042 (0.006)	0.035 (0.006)
		50	0.051 (0.004)	0.047 (0.005)	0.047 (0.003)	0.045 (0.004)	0.028 (0.004)
	USA	5	0.101 (0.01)	0.107 (0.006)	0.103 (0.007)	0.112 (0.008)	0.113 (0.003)
		10	0.095 (0.007)	0.103 (0.007)	0.104 (0.007)	0.108 (0.006)	0.112 (0.004)
		20	0.094 (0.008)	0.092 (0.009)	0.103 (0.007)	0.107 (0.008)	0.116 (0.004)
		50	0.105 (0.006)	0.095 (0.008)	0.105 (0.006)	0.113 (0.005)	0.115 (0.004)
Pacific	Europe	5	0.111 (0.009)	0.088 (0.0)	0.088 (0.0)	0.096 (0.008)	0.123 (0.014)
		10	0.112 (0.012)	0.088 (0.0)	0.088 (0.0)	0.111 (0.012)	0.126 (0.014)
		20	0.107 (0.01)	0.088 (0.0)	0.088 (0.0)	0.106 (0.01)	0.111 (0.012)
		50	0.113 (0.01)	0.089 (0.0)	0.088 (0.0)	0.101 (0.008)	0.145 (0.011)
	USA	5	0.109 (0.006)	0.115 (0.006)	0.115 (0.006)	0.113 (0.006)	0.116 (0.006)
		10	0.088 (0.006)	0.11 (0.007)	0.11 (0.007)	0.109 (0.008)	0.113 (0.006)
		20	0.099 (0.007)	0.107 (0.007)	0.106 (0.006)	0.114 (0.006)	0.114 (0.006)
		50	0.111 (0.005)	0.119 (0.004)	0.119 (0.004)	0.119 (0.004)	0.119 (0.004)
USA	Europe	5	0.117 (0.014)	0.091 (0.001)	0.09 (0.001)	0.116 (0.014)	0.144 (0.013)
		10	0.117 (0.014)	0.091 (0.001)	0.091 (0.001)	0.117 (0.014)	0.151 (0.014)
		20	0.119 (0.011)	0.09 (0.001)	0.095 (0.005)	0.108 (0.009)	0.15 (0.011)
		50	0.11 (0.01)	0.091 (0.001)	0.091 (0.002)	0.11 (0.01)	0.141 (0.013)
	Pacific	5	0.046 (0.007)	0.049 (0.002)	0.051 (0.004)	0.042 (0.004)	0.046 (0.006)
		10	0.07 (0.004)	0.06 (0.004)	0.058 (0.004)	0.058 (0.005)	0.046 (0.006)
		20	0.048 (0.004)	0.053 (0.003)	0.049 (0.003)	0.047 (0.003)	0.041 (0.004)
		50	0.047 (0.003)	0.052 (0.003)	0.048 (0.003)	0.047 (0.003)	0.046 (0.002)

The training, test and validation datasets cover monthly return data over the periods from January 2008 to December 2014, from January 2015 to December 2015, and from January 2016 to December 2021, respectively. In each simulation run, we randomly sample $N_1 \in \{5, 10, 20, 50\}$ return vectors from the training dataset of the target distribution \mathbb{P}_1 and $N_2 = 5$ return vectors from the training dataset of the source distribution \mathbb{P}_2 . These samples are used to construct $\widehat{\mathbb{P}}_1$ and $\widehat{\mathbb{P}}_2$, respectively. In addition, we sample 5 return vectors from the validation dataset of the target distribution, which are used to tune the hyperparameters. For each portfolio under consideration, we then compute the empirical Sharpe ratio on the test samples of the target distribution. Table 3 reports the means and standard errors of these empirical Sharpe ratios across 10 independent simulation runs for different choices of source and target regions and for different choices of the target sample size N_1 . We observe that multi-source DRO outperforms single-source DRO when the source data originates from a similar market (*e.g.*, European data can help to improve portfolios in the US but *not* in the Pacific area). When the source and target distributions differ too much, it is better to use single-source DRO based solely on the target data.

6.2. Assortment Optimization

A firm operates two stores in different geographical regions, each of which sells d different products. For simplicity we assume that the products are perfect substitutes. The firm aims to open a third store in a new region. Due to space constraints, only a subset of $B \leq d$ products can be offered in the new store. If $p_i \geq 0$ denotes the (given) price and $\xi_i \geq 0$ the (uncertain) demand of the i -th product for every $i \in [d]$, then the firm faces the following stylized assortment optimization problem.

$$\max_{\theta \in \{0,1\}^d} \left\{ \mathbb{E}_{\mathbb{P}} \left[\sum_{i=1}^d \theta_i p_i \xi_i \right] : \sum_{i=1}^d \theta_i \leq B \right\} \quad (22)$$

This problem seeks a subset of at most B products, encoded by the binary vector θ , that maximizes expected revenue. The expectation is evaluated with respect to the joint demand distribution \mathbb{P} in the region of the new store. We henceforth assume that \mathbb{P} is unknown but close to the $K = 2$ demand distributions \mathbb{P}_1 and \mathbb{P}_2 in the regions of the existing stores. While \mathbb{P}_1 and \mathbb{P}_2 are unknown, too, the firm has access to respective empirical distributions $\hat{\mathbb{P}}_1$ and $\hat{\mathbb{P}}_2$ and can thus solve an instance of (MS-DRO) corresponding to (22). The underlying ambiguity set is defined as the intersection of two 1-Wasserstein balls induced by the transportation cost function $c(\xi, \xi') = \|\xi - \xi'\|_1$. By Corollary 2, this instance of (MS-DRO) can be reformulated as a mixed-integer linear program.

In our experiments we set $d = 10$ and $B = 3$, and we define $p_i = 0.01 \times i$ for all $i \in [d]$. The training, validation and test datasets are constructed from historical product demand data from Kaggle (<https://www.kaggle.com/datasets/felixzhao/productdemandforecasting>). This dataset contains order quantities for multiple products in three different geographical regions designated by J, S and C, respectively. Among the products that are available in all three regions, we select the 10 most popular ones in terms of average demand. Kernel density estimates of the marginal demand distributions of all products are shown in Figure 3. We assume that the new store is located in region J and that the two existing stores are located in regions S and C. This choice is motivated by the observation that the demand distributions in region J are sometimes more similar to those in region S (see products 4, 6, 7, 9 and 10 in Figure 3) and sometimes more similar to those in region C




Figure 3 Kernel density estimates of the marginal demand distributions in region J (blue), S (orange) and C (green).

(see products 1, 2, 3, 5 and 8 in Figure 3). Hence, none of the two source distributions provides a consistently better approximation for \mathbb{P} , and thus it makes sense to use data from both sources.

As in Section 6.1.2, we compare the optimal solutions of the multi-source DRO problem against those of different single-source DRO baselines. We use the same parametrization for the radii of the Wasserstein balls and the same search grids for all hyperparameters (λ and m in multi-source DRO and ε in single-source DRO) as in Section 6.1.2. In all cases, we select hyperparameters from the underlying search grids that maximize the average revenue across ten independent validation samples. The validation samples are drawn from the same distribution as the training samples. For methods that use data from \mathbb{P}_1 as well as \mathbb{P}_2 , we select five validation samples from each distribution.

We compute the expected revenues of all optimal assortments on a dataset of 3,372 test samples. Table 4 reports the means and standard errors of these expected revenues across 20 independent simulation runs for different choices of the source sample sizes N_1 and N_2 . We observe that multi-source DRO consistently outperforms all baselines. The dominance of multi-source DRO is linked to the structure of the two source distributions, none of which provides a globally better approximation for the target distribution. Thus, using both data sources is better than using only one. In addition, multi-source DRO makes better use of the two data sources than traditional DRO with a Wasserstein ball centered at a mixture or at a Wasserstein barycenter of the empirical source distributions.

Acknowledgments This research was supported by the Swiss National Science Foundation under the NCCR Automation, grant agreement 51NF40_180545, and by the Spanish Ministry of Science and Innovation (AEI/10.13039/501100011033) through project PID2020-115460GB-I00. Finally, the authors thankfully acknowledge the computer resources, technical expertise, and assistance provided by the SCBI (Supercomputing and Bioinformatics) center of the University of Málaga.

Table 4 Out-of-sample expected revenue of optimal product portfolios (mean (std. error) over 20 replications)

N_1	N_2	Single-s. DRO on distribution $\widehat{\mathbb{P}}_1$	Single-s. DRO on distribution $\widehat{\mathbb{P}}_2$	Single-s. DRO on pooled data	Single-s. DRO on barycenter	Multi-s. DRO
25	25	504.52 (37.51)	431.53 (31.17)	481.08 (23.25)	480.16 (30.72)	502.20 (22.90)
	50	481.51 (41.26)	460.61 (29.33)	488.86 (22.87)	476.58 (37.78)	516.93 (15.66)
	75	478.79 (42.02)	475.65 (29.29)	470.58 (32.40)	501.66 (30.18)	515.83 (17.53)
	100	485.10 (37.71)	443.66 (34.65)	479.74 (29.72)	485.10 (37.71)	505.23 (22.31)
50	25	482.26 (37.96)	473.96 (31.94)	489.10 (15.0)	489.93 (30.76)	515.55 (21.24)
	50	481.97 (30.82)	442.18 (32.37)	485.77 (14.60)	473.80 (33.48)	519.13 (10.93)
	75	510.30 (21.27)	450.55 (32.32)	485.77 (14.60)	482.96 (30.69)	506.92 (17.70)
	100	471.92 (26.69)	430.48 (30.53)	482.18 (22.24)	443.68 (29.12)	519.13 (10.93)
75	25	490.24 (18.25)	465.64 (27.43)	482.43 (14.03)	486.91 (17.97)	514.70 (14.54)
	50	489.10 (15.00)	462.31 (26.60)	489.10 (15.00)	477.40 (18.45)	499.11 (15.23)
	75	477.50 (28.54)	450.55 (32.32)	479.10 (13.26)	465.20 (31.46)	492.44 (15.23)
	100	479.49 (34.36)	468.98 (28.15)	492.19 (23.03)	486.41 (30.31)	515.79 (12.25)
100	25	481.08 (23.25)	487.36 (27.41)	489.10 (15.00)	457.08 (27.38)	515.79 (12.25)
	50	461.92 (24.31)	460.61 (29.33)	482.43 (14.03)	462.11 (25.48)	505.79 (14.60)
	75	477.50 (28.54)	467.34 (24.41)	485.77 (14.60)	489.10 (15.00)	509.12 (14.03)
	100	470.83 (27.45)	465.64 (27.43)	485.77 (14.60)	481.08 (23.25)	509.12 (14.03)

Appendix A: Proofs of Section 4

Proof of Theorem 4 Throughout this proof we focus on OT barycenter and uncertainty quantification problems with quadratic transportation cost functions $c(\xi, \xi') = \|\xi - \xi'\|_2^2$ and full support sets $\Xi = \mathbb{R}^d$. Computing the optimal value of the OT barycenter problem (OT-BC) with uniform weights $\lambda_1 = \dots = \lambda_K = 1$ is known to be NP-hard even if all distributions $\mathbb{P}_1, \dots, \mathbb{P}_K$ are discrete (Altschuler and Boix-Adsera 2022, Theorem 1.1). We can thus prove NP-hardness of the UQ FEASIBILITY problem by reduction from the OT barycenter problem. To this end, fix any instance of the OT barycenter problem specified by K discrete distributions $\widehat{\mathbb{P}}_1, \dots, \widehat{\mathbb{P}}_K$ with at most N atoms and uniform weights. Next, define \mathcal{E} as the set of all radii $\varepsilon = (\varepsilon_1, \dots, \varepsilon_K)$ for which the uncertainty quantification problem (P-UQ) with reference distributions $\widehat{\mathbb{P}}_1, \dots, \widehat{\mathbb{P}}_K$ and vanishing loss function $\ell(\xi) = 0$ is feasible. As $\lambda_k = 1$ for every $k \in [K]$, (OT-BC) is equivalent to

$$\min_{\varepsilon \in \mathcal{E}} \sum_{k=1}^K \varepsilon_k. \quad (23)$$

Indeed, if \mathbb{P} is feasible in (OT-BC), then ε defined through $\varepsilon_k = C(\mathbb{P}, \widehat{\mathbb{P}}_k)$, $k \in [K]$, is feasible in (23) with the same objective value. Thus, the minimum of (23) is smaller or equal to that of (OT-BC). Conversely, if ε is feasible in (23), then the definition of \mathcal{E} implies that there exists $\mathbb{P} \in \mathcal{P}(\Xi)$ with $C(\mathbb{P}, \widehat{\mathbb{P}}_k) \leq \varepsilon_k$ for every $k \in [K]$, and \mathbb{P} has the same or a lower objective value than $\sum_{k=1}^K \varepsilon_k$. Thus, the minimum of (OT-BC) is smaller or equal to that of (23). Next, note that \mathcal{E} is convex and closed because the OT distance is convex and weakly lower semicontinuous in its arguments. Note also that the UQ FEASIBILITY problem with fixed discrete distributions $\widehat{\mathbb{P}}_1, \dots, \widehat{\mathbb{P}}_K$ serves as a separation oracle for (23). Indeed, for any $\bar{\varepsilon} \in \mathbb{R}^K$, the UQ FEASIBILITY problem either confirms that $\bar{\varepsilon} \in \mathcal{E}$ (which happens if and only if (P-UQ) is feasible) or enables

us to construct a halfspace \mathcal{H} that contains \mathcal{E} but not $\bar{\varepsilon}$. If $\bar{\varepsilon} \notin \mathcal{E}$ and $\bar{\varepsilon}_k < 0$ for some $k \in [K]$, then we can simply set $\mathcal{H} = \{\varepsilon \in \mathbb{R}^d : \varepsilon_k \geq 0\}$. If $\bar{\varepsilon} \notin \mathcal{E}$ and $\bar{\varepsilon}_k \geq 0$ for every $k \in [K]$, on the other hand, then we can set

$$\mathcal{H} = \left\{ \varepsilon \in \mathbb{R}^d : \sum_{k=1}^K \varepsilon_k \lambda_k^\infty + \sum_{k=1}^K \sum_{j=1}^{N_k} p_{k,j} \gamma_{k,j}^\infty \geq 0 \right\},$$

where $(\lambda^\infty, \gamma^\infty)$ is the recession direction of the feasible set of (D-UQ) returned by the UQ FEASIBILITY oracle. As the objective function of (D-UQ) with $\varepsilon = \bar{\varepsilon}$ decreases along $(\lambda^\infty, \gamma^\infty)$, it is clear that $\bar{\varepsilon} \notin \mathcal{H}$. To show that $\mathcal{E} \subseteq \mathcal{H}$, assume for the sake of contradiction that there exists $\tilde{\varepsilon} \in \mathcal{E}$ with $\tilde{\varepsilon} \notin \mathcal{H}$. Thus, $(\lambda^0, \gamma^0) + t \cdot (\lambda^\infty, \gamma^\infty)$ is feasible in (D-UQ) for any $t \geq 0$ and for any feasible solution (λ^0, γ^0) of (D-UQ), and we have

$$\lim_{t \rightarrow \infty} \sum_{k=1}^K \tilde{\varepsilon}_k (\lambda_k^0 + t \lambda_k^\infty) + \sum_{k=1}^K \sum_{j=1}^{N_k} p_{k,j} (\gamma_{k,j}^0 + t \gamma_{k,j}^\infty) = -\infty.$$

This implies that (D-UQ) is unbounded, which in turn implies via weak duality that (9) is infeasible. However, this contradicts our assumption that $\tilde{\varepsilon} \in \mathcal{E}$. We may thus conclude that $\mathcal{E} \subseteq \mathcal{H}$.

Next, we construct a ball that is guaranteed to contain all minimizers of (23). To this end, note that every distribution on \mathbb{R}^d is feasible in the OT barycenter problem (OT-BC). For example, $\widehat{\mathbb{P}}_1$ is feasible and has a finite objective function value $R = \sum_{k=1}^K C(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k)$, which provides an upper bound on the optimal value of (OT-BC). As (OT-BC) and (23) are equivalent, this implies that all minimizers of (23) must be contained in the simplex $\{\varepsilon \in \mathbb{R}_+^K : \sum_{k=1}^K \varepsilon_k \leq R\}$, which in turn resides within the ball of radius R around 0. Note that since $C(\widehat{\mathbb{P}}_1, \widehat{\mathbb{P}}_k)$ represents the optimal value of a linear program with $\text{poly}(d, K)$ variables and constraints, R can be encoded with $\text{poly}(d, K, \log U)$ bits, where $\log U$ denotes the bit complexity of each parameter needed to describe the reference distributions (*i.e.*, the probabilities as well as the locations of their atoms). Trivially, \mathcal{E} further contains balls of any positive radius because $\varepsilon \in \mathcal{E}$ implies that $\varepsilon' \in \mathcal{E}$ for every $\varepsilon' \geq \varepsilon$.

Given the separation oracle for \mathcal{E} constructed above as well as the ball of radius R that contains all minimizers of (23), we can compute the optimal value of (23) to any absolute accuracy $\delta > 0$ with the ellipsoid algorithm using $\text{poly}(d, K, \log U, \log \frac{1}{\delta})$ calls to the UQ FEASIBILITY oracle (Ben-Tal and Nemirovski 2001, Theorem 5.2.1). Recall that the OT barycenter problem (OT-BC) can be reformulated as a multi-margin OT problem, which is equivalent to a linear program with $\mathcal{O}(N^K)$ variables and constraints and with objective function coefficients that can be computed via $\text{poly}(d, K, \log U)$ arithmetic operations (see Theorem 1). Hence, the (equal) optimal values of (OT-BC) and (23) can be encoded with $\text{poly}(d, K, N, \log U)$ bits. Selecting the accuracy δ of the ellipsoid algorithm to satisfy $\log(1/\delta) = \text{poly}(d, k, N, \log U)$, we can thus

compute the exact optimal value of (23) by rounding the output of the ellipsoid algorithm to the nearest multiple of δ . In summary, we have thus conceived an algorithm for computing the optimal value of the NP-hard OT barycenter problem (OT-BC), which requires only $\text{poly}(d, K, N, \log U)$ calls to the UQ FEASIBILITY oracle. Consequently, the UQ FEASIBILITY oracle must be NP-hard, too. \blacksquare

Recall that, for ease of exposition, we have assumed throughout Section 4 that $N_k = N$ for all $k \in [K]$. We have also assumed that $\log U$ represents an upper bound on the number of bits needed to encode any radius ε_k , any probability $p_{k,j}$ and any component of $\widehat{\xi}_{k,j}$ for all $j \in [N]$ and $k \in [K]$.

Before proving Theorem 5, we first construct two different separation oracles for the feasible set \mathcal{D} of problem (D-UQ). A separation oracle for \mathcal{D} decides whether a given point $(\lambda, \gamma) \in \mathbb{R}^{K+KN}$ belongs to \mathcal{D} and, if not, constructs a closed halfspace $\mathcal{H} \subseteq \mathbb{R}^{K+KN}$ with $(\lambda, \gamma) \notin \mathcal{H}$ and $\mathcal{D} \subseteq \mathcal{H}$. The following lemma constructs the first of the two separation oracles. This oracle runs in polynomial time provided that K is fixed.

LEMMA 4 (Polynomial-Time Separation Oracle for Fixed K). *If Assumptions 6, 7 and 8(i) hold, then there is a separation oracle for \mathcal{D} that runs in time polynomial in $d, N, \log U$ and the bit size of (λ, γ) .*

Proof of Lemma 4 We construct a separation oracle for \mathcal{D} as follows. Given a test point (λ, γ) , we first check if $\lambda \geq 0$. Otherwise, if $\lambda_k < 0$ for some $k \in [K]$, the closed halfspace $\mathcal{H} = \{(\lambda', \gamma') \in \mathbb{R}^{K+KN} : \lambda'_k \geq 0\}$ contains \mathcal{D} but not (λ, γ) . For any fixed multi-index $\alpha \in \mathcal{A}$ we then use the generalized Moreau envelope oracle from Assumption 6 to compute the supremum on the left hand side of the α -th robust constraint in (D-UQ), which we denote here by $\ell_\alpha^*(\lambda)$. If $\ell_\alpha^*(\lambda) = \infty$, then the oracle outputs a closed halfspace $\Lambda \subseteq \mathbb{R}^K$ such that $\lambda \notin \Lambda$ and $\lambda' \in \Lambda$ for every $\lambda' \in \mathbb{R}^K$ with $\ell^*(\lambda') < \infty$. Thus, the closed halfspace $\mathcal{H} = \Lambda \times \mathbb{R}^{KN}$ contains \mathcal{D} but not (λ, γ) . If $\sum_{k=1}^K \gamma_{k,\alpha_k} < \ell_\alpha^*(\lambda) < \infty$, on the other hand, then the oracle outputs a maximizer $\xi^* \in \Xi$ for the optimization problem embedded in the α -th robust constraint in (D-UQ). The closed halfspace

$$\mathcal{H} = \left\{ (\lambda', \gamma') \in \mathbb{R}^{K+KN} : \ell(\xi^*) - \sum_{k=1}^K \lambda'_k c(\xi^*, \widehat{\xi}_{k\alpha_k}) \leq \sum_{k=1}^K \gamma'_{k,\alpha_k} \right\}$$

then contains every feasible solution $(\lambda', \gamma') \in \mathcal{D}$ because $\ell(\xi^*) - \sum_{k=1}^K \lambda'_k c(\xi^*, \widehat{\xi}_{k\alpha_k}) \leq \ell_\alpha^*(\lambda') \leq \sum_{k=1}^K \gamma'_{k,\alpha_k}$, where the first inequality holds because ξ^* is feasible but not necessarily optimal in the embedded optimization problem associated with λ' , and the second inequality follows from the definition of the feasible set \mathcal{D} . On the other hand, $(\lambda, \gamma) \notin \mathcal{H}$ because $\ell(\xi^*) - \sum_{k=1}^K \lambda_k c(\xi^*, \widehat{\xi}_{k\alpha_k}) = \ell_\alpha^*(\lambda) > \sum_{k=1}^K \gamma_{k,\alpha_k}$, where the equality holds because ξ^* is optimal in the embedded optimization problem associated with λ . If $\ell_\alpha^*(\lambda) \leq \sum_{k=1}^K \gamma_{k,\alpha_k}$, finally, then we have verified that the α -th robust constraint in (D-UQ) is satisfied. If $\lambda \geq 0$ and all $|\mathcal{A}| = N^K$

robust constraints are satisfied, then we can conclude that $(\lambda, \gamma) \in \mathcal{D}$. The procedure outlined above thus correctly decides whether $(\lambda, \gamma) \in \mathcal{D}$ and, if not, outputs a halfspace \mathcal{H} that contains \mathcal{D} but not (λ, γ) .

The runtime of the proposed procedure is dominated by the N^K calls to the Moreau envelope oracle. By Assumptions 6, 7 and 8(i), the runtime of each call grows polynomially with d , K , N , $\log U$ and the bit size of (λ, γ) . Thus, if K is fixed, the oracle runs in time polynomial in d , N , $\log U$ and the bit size of (λ, γ) . ■

Perhaps surprisingly, one can construct a second oracle that runs in time polynomial in K if d is fixed. This construction relies on the following lemma inspired by (Altschuler and Boix-Adsera 2021, Lemma 17).

LEMMA 5 (Constraint Elimination). *If Assumptions 1 and 8(ii) hold and $\lambda_k \geq 0$ for all $k \in [K]$, then*

$$\max_{\alpha \in \mathcal{A}} \sup_{\xi \in \Xi} \ell(\xi) - \sum_{k=1}^K \left(\lambda_k c(\xi, \widehat{\xi}_{k, \alpha_k}) + \gamma_{k, \alpha_k} \right) = \max_{\alpha \in \mathcal{A}'} \sup_{\xi \in \Xi} \ell(\xi) - \sum_{k=1}^K \left(\lambda_k c(\xi, \widehat{\xi}_{k, \alpha_k}) + \gamma_{k, \alpha_k} \right),$$

where $\mathcal{A}' = \{\alpha \in \mathcal{A} : \mathcal{C}_\alpha \neq \emptyset\}$ is defined in terms of the cells $\mathcal{C}_\alpha = \bigcap_{k=1}^K \Xi_{k, \alpha_k}$, $\alpha \in \mathcal{A}$, with

$$\Xi_{k, i} = \{\xi \in \Xi : \lambda_k c(\xi, \xi_{k, i}) + \gamma_{k, i} < \lambda_k c(\xi, \xi_{k, j}) + \gamma_{k, j} \quad \forall j \in [N], j \neq i\} \quad \forall k \in [K], \forall i \in [N].$$

Proof of Lemma 5 Since $\mathcal{A}' \subseteq \mathcal{A}$, the maximum over \mathcal{A}' is smaller than or equal to the maximum over \mathcal{A} .

To prove the reverse inequality, note that the maximum over \mathcal{A} can be expressed as

$$\begin{aligned} \sup_{\xi \in \Xi} \max_{\alpha \in \mathcal{A}} \ell(\xi, \theta) - \sum_{k=1}^K \left(\lambda_k c(\xi, \widehat{\xi}_{k, \alpha_k}) + \gamma_{k, \alpha_k} \right) &= \max_{\alpha' \in \mathcal{A}'} \sup_{\xi \in \mathcal{C}_{\alpha'}} \max_{\alpha \in \mathcal{A}} \ell(\xi, \theta) - \sum_{k=1}^K \left(\lambda_k c(\xi, \widehat{\xi}_{k, \alpha_k}) + \gamma_{k, \alpha_k} \right) \\ &= \max_{\alpha' \in \mathcal{A}'} \sup_{\xi \in \mathcal{C}_{\alpha'}} \ell(\xi, \theta) - \sum_{k=1}^K \left(\lambda_k c(\xi, \widehat{\xi}_{k, \alpha'_k}) - \gamma_{k, \alpha'_k} \right) \\ &\leq \max_{\alpha' \in \mathcal{A}'} \sup_{\xi \in \Xi} \ell(\xi, \theta) - \sum_{k=1}^K \left(\lambda_k c(\xi, \widehat{\xi}_{k, \alpha'_k}) + \gamma_{k, \alpha'_k} \right), \end{aligned}$$

where the first equality holds because $\bigcup_{\alpha' \in \mathcal{A}'} \text{cl}(\mathcal{C}_{\alpha'}) = \Xi$ thanks to Assumption 8(ii) and because ℓ and c are upper and lower semicontinuous, respectively (see Assumption 1). The second equality exploits the construction of the cells $\mathcal{C}_{\alpha'}$, $\alpha' \in \mathcal{A}'$, and the inequality holds because $\mathcal{C}_{\alpha'} \subseteq \Xi$ for every $\alpha' \in \mathcal{A}'$. ■

Lemma 5 implies that, among all N^K robust constraints in (D-UQ), it suffices to retain only those indexed by $\alpha \in \mathcal{A}'$. The others can be eliminated without affecting the solution of (D-UQ). By Assumption 8(iii), the reduced index set can be computed in time $\text{poly}(N, K, \log U)$ for any fixed dimension d . This implies in particular that $|\mathcal{A}'| = \text{poly}(N, K, \log U)$, which is much smaller than $|\mathcal{A}| = N^K$ for large K . The following lemma constructs the second separation oracle, which runs in polynomial time provided that d is fixed.

LEMMA 6 (Polynomial-Time Separation Oracle for Fixed d). *If Assumptions 6, 7 and 8 hold, then there is a separation oracle for \mathcal{D} that runs in time polynomial in K , N , $\log U$ and the bit size of (λ, γ) .*

Proof of Lemma 6 The desired separation oracle is constructed as in the proof of Lemma 4. However, the Moreau envelope oracle from Assumption 6 is only called for robust constraints indexed by some $\alpha \in \mathcal{A}'$. As $|\mathcal{A}'| = \text{poly}(N, K, \log U)$ and because of the assumed efficiency of the Moreau envelope oracle, the time needed for all $|\mathcal{A}'|$ oracle calls grows polynomially with $d, K, N, \log U$ and the bit size of (λ, γ) . Also, by Assumption 8(iii), the reduced index set \mathcal{A}' can be computed in time $\text{poly}(K, N, \log U)$. The overall runtime of the constructed separation oracle thus grows polynomially with $K, N, \log U$ and the bit size of (λ, γ) . ■

We emphasize that all known algorithms for computing the reduced index set \mathcal{A}' have a runtime that grows either exponentially with d or with K ; see the discussion in (Altschuler and Boix-Adsera 2021).

The proof of Theorem 5 follows now almost immediately from Lemmas 4 and 6.

Proof of Theorem 5 The assumptions of the theorem imply via Lemmas 4 and 6 that the feasible set of (D-UQ) admits a separation oracle that runs in time $\text{poly}(d, N, \log U)$ and $\text{poly}(K, N \log U)$, respectively. Thus, problem (D-UQ) can be solved efficiently with the ellipsoid algorithm (Ben-Tal and Nemirovski 2001, Theorem 5.2.1). The initial ellipsoid may be set to the ball of radius R around 0, which can be computed in time $\text{poly}(d, K, N \log U)$. As the transportation cost function is nonnegative, it is clear that if (λ, γ) is feasible in (D-UQ), then any $(\lambda', \gamma') \geq (\lambda, \gamma)$ is also feasible in (D-UQ). Hence, we may assume without loss of generality that the intersection of the initial ellipsoid and the feasible set of problem (D-UQ) contains some ball of radius 1. The ellipsoid algorithm with the first (second) separation oracle thus computes the optimal value of (D-UQ) to within any accuracy δ in time $\text{poly}(d, N, \log U, \log \frac{1}{\delta})$ ($\text{poly}(K, N \log U, \log \frac{1}{\delta})$). ■

References

- Agueh M, Carlier G (2011) Barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis* 43(2):904–924.
- Altschuler JM, Boix-Adsera E (2021) Wasserstein barycenters can be computed in polynomial time in fixed dimension. *Journal of Machine Learning Research* 22(44):1–19.
- Altschuler JM, Boix-Adsera E (2022) Wasserstein barycenters are NP-hard to compute. *SIAM Journal on Mathematics of Data Science* 4(1):179–203.
- Andres E, Borgwardt S, Miller J (2016) Discrete Wasserstein barycenters: Optimal transport for discrete data. *Mathematical Methods of Operations Research* 84(2):389–409.
- Awasthi P, Jung C, Morgenstern J (2022) Distributionally robust data join. *arXiv preprint arXiv:2202.05797*.

-
- Barré M, Giron C, Mazzolini M, d'Aspremont A (2020) Averaging atmospheric gas concentration data using Wasserstein barycenters. *arXiv preprint arXiv:2010.02762*.
- Ben-Tal A, Den Hertog D, Vial JP (2015) Deriving robust counterparts of nonlinear uncertain inequalities. *Mathematical Programming* 149(1):265–299.
- Ben-Tal A, Nemirovski A (2001) *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications* (SIAM).
- Blanchet J, Murthy K (2019) Quantifying distributional model risk via optimal transport. *Mathematics of Operations Research* 44(2):565–600.
- Blanchet J, Murthy K, Nguyen VA (2021) Statistical analysis of Wasserstein distributionally robust estimators. *Emerging Optimization Methods and Modeling Techniques with Applications*, 227–254 (INFORMS).
- Claici S, Chien E, Solomon J (2018) Stochastic Wasserstein barycenters. *International Conference on Machine Learning*, 999–1008.
- Clément P, Desch W (2008) Wasserstein metric and subordination. *Studia Mathematica* 1(189):35–52.
- Courty N, Flamary R, Tuia D (2014) Domain adaptation with regularized optimal transport. *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD*, 274–289.
- Courty N, Flamary R, Tuia D, Rakotomamonjy A (2016) Optimal transport for domain adaptation. *IEEE transactions on pattern analysis and machine intelligence* 39(9):1853–1865.
- Cuturi M, Doucet A (2014) Fast computation of Wasserstein barycenters. *International Conference on Machine Learning*, 685–693.
- Deng Y, Kamani MM, Mahdavi M (2020) Distributionally robust federated averaging. *Advances in Neural Information Processing Systems* 15111–15122.
- Dognin P, Melnyk I, Mroueh Y, Ross J, dos Santos C, Sercu T (2019) Wasserstein barycenter model ensembling. *International Conference on Learning Representations*.
- Dudley RM (2002) *Real Analysis and Probability* (Cambridge University Press).
- Farahani A, Voghoei S, Rasheed K, Arabnia HR (2021) A brief review of domain adaptation. *Advances in Data Science and Information Engineering* 877–894.

- Fournier N, Guillin A (2015) On the rate of convergence in Wasserstein distance of the empirical measure. *Probability Theory and Related Fields* 162(3):707–738.
- Gangbo W, Świąch A (1998) Optimal maps for the multidimensional Monge-Kantorovich problem. *Communications on Pure and Applied Mathematics* 51(1):23–45.
- Gao R (2023) Finite-sample guarantees for Wasserstein distributionally robust optimization: Breaking the curse of dimensionality. *Operations Research* 71(6):2291–2306.
- Gao R, Kleywegt A (2023) Distributionally robust stochastic optimization with Wasserstein distance. *Mathematics of Operations Research* 48(2):1526–5471.
- Gelbrich M (1990) On a formula for the L^2 Wasserstein metric between measures on Euclidean and Hilbert spaces. *Mathematische Nachrichten* 147(1):185–203.
- Grötschel M, Lovász L, Schrijver A (2012) *Geometric Algorithms and Combinatorial Optimization* (Springer).
- Gupta V, Kallus N (2022) Data pooling in stochastic optimization. *Management Science* 68(3):1595–1615.
- Kuhn D, Mohajerin Esfahani P, Nguyen VA, Shafieezadeh-Abadeh S (2019) Wasserstein distributionally robust optimization: Theory and applications in machine learning. *Operations Research & Management Science in the Age of Analytics*, 130–166 (INFORMS).
- Lau TTK, Liu H (2022) Wasserstein distributionally robust optimization via Wasserstein barycenters. *arXiv preprint arXiv:2203.12136*.
- Le Gouic T, Loubes JM (2017) Existence and consistency of Wasserstein barycenters. *Probability Theory and Related Fields* 168(3):901–917.
- Le Gouic T, Paris Q, Rigollet P, Stromme AJ (2022) Fast convergence of empirical barycenters in Alexandrov spaces and the Wasserstein space. *Journal of the European Mathematical Society* 25(6):2229–2250.
- Mohajerin Esfahani P, Kuhn D (2018) Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming* 171(1):115–166.
- Mohri M, Sivek G, Suresh AT (2019) Agnostic federated learning. *International Conference on Machine Learning*, 4615–4625.

-
- Montesuma EF, Mboula FMN (2021) Wasserstein barycenter for multi-source domain adaptation. *The IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16785–16793.
- Nguyen TA, Nguyen TD, Le LT, Dinh CT, Tran NH (2022) On the generalization of Wasserstein robust federated learning. *arXiv preprint arXiv:2206.01432*.
- Panaretos VM, Zemel Y (2020) *An Invitation to Statistics in Wasserstein Space* (Springer).
- Pass B (2015) Multi-marginal optimal transport: Theory and applications. *ESAIM: Mathematical Modelling and Numerical Analysis* 49(6):1771–1790.
- Pflug GC, Pichler A, Wozabal D (2012) The $1/N$ investment strategy is optimal under high model ambiguity. *Journal of Banking & Finance* 36(2):410–417.
- Pinelis I (2016) On the extreme points of moments sets. *Mathematical Methods of Operations Research* 83(3):325–349.
- Rahimian H, Mehrotra S (2022) Frameworks and results in distributionally robust optimization. *Open Journal of Mathematical Optimization* 3:1–85.
- Redko I, Habrard A, Sebban M (2017) Theoretical analysis of domain adaptation with optimal transport. *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 737–753.
- Ro J, Chen M, Mathews R, Mohri M, Suresh AT (2021) Communication-efficient agnostic federated averaging. *arXiv preprint arXiv:2104.02748*.
- Rockafellar RT (1974) *Convex Duality and Optimization* (SIAM).
- Rockafellar RT, Uryasev S (2000) Optimization of conditional value-at-risk. *Journal of Risk* 2:21–42.
- Rockafellar RT, Wets RJB (2009) *Variational Analysis* (Springer).
- Sagawa S, Koh PW, Hashimoto TB, Liang P (2019) Distributionally robust neural networks. *International Conference on Learning Representations*.
- Selvi A, Kreacic E, Ghassemi M, Potluru V, Balch T, Veloso M (2024) Distributionally and adversarially robust logistic regression via intersecting Wasserstein balls. *Private communication*.

- Shafieezadeh Abadeh S, Aolaritei L, Dörfler F, Kuhn D (2023) New perspectives on regularization and computation in optimal transport-based distributionally robust optimization. *arXiv preprint arXiv:2303.03900*.
- Shapiro A (2001) On duality theory of conic linear problems. Goberna MÁ, López MA, eds., *Semi-Infinite Programming*, 135–165 (Kluwer Academic Publishers).
- Shapiro A (2017) Interchangeability principle and dynamic equations in risk averse stochastic programming. *Operations Research Letters* 45(4):377–381.
- Shapiro A, Dentcheva D, Ruszczyński A (2021) *Lectures on Stochastic Programming: Modeling and Theory* (SIAM).
- Singhal P, Walambe R, Ramanna S, Kotecha K (2023) Domain adaptation: Challenges, methods, datasets, and applications. *IEEE Access* 11:6973–7020.
- Smith JE, Winkler RL (2006) The optimizer’s curse: Skepticism and postdecision surprise in decision analysis. *Management Science* 52(3):311–322.
- Solomon J, De Goes F, Peyré G, Cuturi M, Butscher A, Nguyen A, Du T, Guibas L (2015) Convolutional Wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Transactions on Graphics* 34(4):1–11.
- Srivastava S, Li C, Dunson DB (2018) Scalable Bayes via barycenter in Wasserstein space. *Journal of Machine Learning Research* 19(1):312–346.
- Tanoumand N, Bodur M, Naoum-Sawaya J (2023) Data-driven distributionally robust optimization: Intersecting ambiguity sets, performance analysis and tractability. *Optimization Online* .
- Taskesen B, Yue MC, Blanchet J, Kuhn D, Nguyen VA (2021) Sequential domain adaptation by synthesizing distributionally robust experts. *International Conference on Machine Learning*, 10162–10172.
- Villani C (2009) *Optimal Transport: Old and New* (Springer).
- Wang Z, Bühlmann P, Guo Z (2023) Distributionally robust machine learning with multi-source data. *arXiv preprint arXiv:2309.02211*.
- Xiong X, Guo Z, Cai T (2023) Distributionally robust transfer learning. *arXiv preprint arXiv:2309.06534*.

-
- Yue MC, Kuhn D, Wiesemann W (2022) On linear optimization over Wasserstein balls. *Mathematical Programming* 195(1):1107–1122.
- Zhang C, Xie Y, Bai H, Yu B, Li W, Gao Y (2021) A survey on federated learning. *Knowledge-Based Systems* 216:106775.
- Zhao C, Guan Y (2018) Data-driven risk-averse stochastic optimization with Wasserstein metric. *Operations Research Letters* 46(2):262–267.
- Zhao S, Li B, Xu P, Keutzer K (2020) Multi-source domain adaptation in the deep learning era: A systematic survey. *arXiv preprint arXiv:2002.12169*.
- Zhen J, Kuhn D, Wiesemann W (2023) A unified theory of robust and distributionally robust optimization via the primal-worst-equals-dual-best principle. *Operations Research* (forthcoming).
- Zhuang Y, Chen X, Yang Y (2022) Wasserstein k -means for clustering probability distributions. *Advances in Neural Information Processing Systems* 11382–11395.