Statistical, Robustness, and Computational Guarantees for Sliced Wasserstein Distances

Sloan Nietert
Cornell University
nietert@cs.cornell.edu

Ritwik Sadhu
Cornell University
rs2526@cornell.edu

Ziv Goldfeld
Cornell University
goldfeld@cornell.edu

Kengo Kato
Cornell University
kk976@cornell.edu

Abstract

Sliced Wasserstein distances preserve properties of classic Wasserstein distances while being more scalable for computation and estimation in high dimensions. The goal of this work is to quantify this scalability from three key aspects: (i) empirical convergence rates; (ii) robustness to data contamination; and (iii) efficient computational methods. For empirical convergence, we derive fast rates with explicit dependence of constants on dimension, subject to log-concavity of the population distributions. For robustness, we characterize minimax optimal, dimension-free robust estimation risks, and show an equivalence between robust sliced 1-Wasserstein estimation and robust mean estimation. This enables lifting statistical and algorithmic guarantees available for the latter to the sliced 1-Wasserstein setting. Moving on to computational aspects, we analyze the Monte Carlo estimator for the average-sliced distance, demonstrating that larger dimension can result in faster convergence of the numerical integration error. For the max-sliced distance, we focus on a subgradient-based local optimization algorithm that is frequently used in practice, albeit without formal guarantees, and establish an $O(\epsilon^{-4})$ computational complexity bound for it. Our theory is validated by numerical experiments, which altogether provide a comprehensive quantitative account of the scalability question.

1 Introduction

Sliced Wasserstein distances consider the average or maximum of Wasserstein distances between one-dimensional projections of the two distributions. Formally, for $1 \le p < \infty$, they are defined as

$$\underline{\mathsf{W}}_p(\mu,\nu) := \left[\int_{\mathbb{S}^{d-1}} \mathsf{W}_p^p(\mathfrak{p}^\theta_\sharp \mu, \mathfrak{p}^\theta_\sharp \nu) d\sigma(\theta)\right]^{1/p} \quad \text{and} \quad \overline{\mathsf{W}}_p(\mu,\nu) := \max_{\theta \in \mathbb{S}^{d-1}} \mathsf{W}_p(\mathfrak{p}^\theta_\sharp \mu, \mathfrak{p}^\theta_\sharp \nu), \tag{1}$$

where $\mathfrak{p}^\theta_\sharp \mu$ is the pushforward of μ under the projection $\mathfrak{p}^\theta : x \mapsto \theta^\intercal x$ from $\mathbb{R}^d$ to $\mathbb{R}$ and σ is the uniform distribution on the unit sphere $\mathbb{S}^{d-1}$ in $\mathbb{R}^d$. Sliced Wasserstein distances were introduced in [49] as a means to mitigate the computational burden of evaluating classic W_p, which rapidly becomes excessive as d grows. Indeed, sliced distances are readily computable using the closed-form expression for W_p between distributions on $\mathbb{R}$ (as the L^p norm between quantile functions). Further, $\underline{\mathsf{W}}_p$ and $\overline{\mathsf{W}}_p$ are metrics on $\mathcal{P}_p(\mathbb{R}^d)$ and generate the same topology as classic W_p [10, 44, 6, 42]. As such, the sliced distances have been applied to various statistical inference and machine learning tasks, including barycenter computation [49, 9], generative modeling [18, 17, 44, 58], autoencoders [29], differential privacy [50], Bayesian computation [41] and topological data analysis [12].

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

1.1 Statistical, Robustness, and Computational Aspects of Sliced Distances

In practice, the sliced Wasserstein distances in (1) must be approximated from two aspects: (i) empirically estimate the population measures μ and ν, and (ii) employ numerical integration or optimization methods to compute the average- or max-sliced distances, respectively. While these approximations are implemented in all but every application of sliced distances, formal guarantees concerning their accuracy are partial or even missing. For the estimation error, the question boils down to quantifying the rate at which $\underline{\mathsf{W}}_p(\hat{\mu}_n, \mu)$ and $\overline{\mathsf{W}}_p(\hat{\mu}_n, \mu)$ decay to 0, where $\hat{\mu}_n$ is the empirical distribution of n independent observations from μ.[1] These rates are known to adapt to the low-dimensionality of the projected distribution, but previously derived rates do not seem to be sharp [35], rely on high-level assumptions that may be hard to verify in practice [47], or hide dimension-dependent constants whose characterization is crucial for understanding the scalability of sliced distances [42]. More recently, [37] showed that near-parametric rates (i.e., up to polylogarithmic factors) are achievable for the average-sliced p-Wasserstein distance in the two-sample case, under the alternative ($\mu \neq \nu$). Limit distributions for sliced Wasserstein distances were studied in [37, 21, 60, 59], but these results inherently neglect constants and dependence on dimension.

Concerning robust estimation, while these aspects were studied for classic Wasserstein distances [4, 45, 40, 30, 53, 46], they were not considered under sliced W_p. Improvement in robustness to outliers due to projection-averaging was demonstrated for the Cramér-von Mises statistic in the context of multivariate two-sample testing [27]. It therefore stands to reason that similar gains would emerge for Wasserstein distances, which is especially appealing since robust estimation of classic W_p in high dimensions is hard. Indeed, [46] showed that when an ϵ-fraction of data is contaminated, W_p admits worst-case estimation risk $\sqrt{d}\epsilon^{1/p-1/2}$ over distributions with bounded covariance. Consequently, obtaining accurate estimates of W_p from contaminated data is infeasible in high dimensions when $\epsilon = \Omega(1)$, which further motivates exploring robustness under slicing.

From the computational standpoint, the average-sliced distance $\underline{\mathsf{W}}_p$ is typically computed using Monte Carlo (MC) integration [28, 42]. The accuracy of this approach strongly depends on the variance of the function $\theta \mapsto \mathsf{W}_p^p(\mathfrak{p}_\sharp^\theta \mu, \mathfrak{p}_\sharp^\theta \nu)$ when θ is uniformly distributed on $\mathbb{S}^{d-1}$, which may scale badly with d. A bound on the MC integration error in terms of this variance was provided in [42] but without further analysis to control it by basic properties of the population distribution or characterize its dependence on d. Accordingly, the accuracy of the MC-based approach for computing $\underline{\mathsf{W}}_p$ stands unresolved. Recently, [43] used the conditional central limit theorem [51] to derive a Gaussian approximation of $\underline{\mathsf{W}}_2$ that can be computed in closed form. The accuracy of this approximation may improve as $d \to \infty$, contingent on certain weak dependence assumptions on the data distribution. A popular approach for computing the max-sliced distance is the heuristic alternating optimization procedure from [28, 17], which, however, lacks formal convergence guarantees. More recently, computational aspects of the so-called "projection-robust" Wasserstein distance, which considers projections to k-dimensional subspaces, were explored in [48, 34, 25].[2] As maximization of projected distance is a non-convex and non-smooth optimization problems, these works considered convex relaxations [48] or entropic regularization [34, 25] to prove approximate convergence to a stationary point.

1.2 Contributions

The goal of this paper is to close the aforementioned gaps by (i) deriving fast empirical convergence rates for sliced distances with explicit dimension dependence; (ii) characterizing minimax optimal robust estimation rates with improved dependence on dimension; and (iii) providing formal guarantees for frequently used methods for computing both the average- and max-sliced W_p. Focusing on log-concave distributions, we show that both average- and max-sliced empirical distances converge as $n^{-1/\max\{2,p\}}$, which is sharp as it matches lower bounds from [8]. Furthermore, we characterize the constant in terms of d and elementary properties of the population distribution (e.g., mean, moments, covariance matrix). Our derivation leverages the machinery of [8] for analyzing empirical convergence of Wasserstein distances between log-concave measures on $\mathbb{R}$. To that end, we show that log-concavity is preserved under projections and derive lower bounds on the Cheeger constant of the projected distribution. Our results elucidate scaling rates of d with n for which (high-dimensional) empirical convergence holds true, thereby addressing the scalability of empirical estimates question.

[1] The two-sample setting, which concerns the convergence $\underline{\mathsf{W}}_p(\hat{\mu}_n, \hat{\nu}_n)$ and $\overline{\mathsf{W}}_p(\hat{\mu}_n, \hat{\nu}_n)$ towards the corresponding distance between the population measures, is also of interest.

[2] Despite the name "projection-robust", these works do not explore robust estimation.

For robustness guarantees, we formalize minimax risk for robust estimation under sliced W_p with total variation (TV) contamination and prove that $\overline{\mathsf{W}}_p$ enjoys a dimension-free risk of $\epsilon^{1/p-1/q}$ when clean distributions have bounded qth moments for $q > p$ and the corruption level is at most ϵ. $\underline{\mathsf{W}}_p$ admits a strictly smaller risk which scales at the same rate when $q = O(1)$. In contrast, the comparable risk for classic W_p in this setting acquires an extra $\sqrt{d}$ factor. Using the framework of generalized resilience [61], we extend these guarantees to the finite-sample setting with adversarial corruptions, obtaining matching rates up to an added empirical approximation term. Furthermore, when $p = 1$, we prove equivalence between standard mean resilience [55] and resilience w.r.t. $\overline{\mathsf{W}}_1$, allowing one to lift statistical and algorithmic guarantees for robust mean estimation to the sliced W_1 setting.

Lastly, we provide formal guarantees for popular methods for computing $\underline{\mathsf{W}}_p$ and $\overline{\mathsf{W}}_p$, which were until now lacking. Our analysis relies on showing that $w_p : \theta \mapsto \mathsf{W}_p(\mathfrak{p}_\sharp^\theta \mu, \mathfrak{p}_\sharp^\theta \nu)$ and its pth power are Lipschitz continuous on $\mathbb{S}^{d-1}$ and deriving sharp bounds on their Lipschitz constants. Having that, we analyze the MC estimator for the average-sliced distance, and use concentration of Lipschitz functions on the unit sphere to bound the variance of w_p^p. The obtained bound reveals that higher dimension can in fact shrink the MC error when the covariance matrices have bounded operator norms. We numerically verify this surprising observation on synthetic examples.

For the max-sliced distance, we analyze the heuristic algorithm from [17, 28], which utilizes alternating subgradient-based optimization. We observe that in addition to being Lipschitz continuous, the optimization objective for $p = 2$ is weakly convex with easily computable gradients. This lets us cast the algorithm from [17, 28] under the proximal stochastic subgradient optimization framework of [14], from which we obtain local solutions for $w_2(\theta)$ with $O(\epsilon^{-4})$ computational complexity. An empirical comparison with the more advanced approaches of [34, 25] for computing the projection-robust Wasserstein distance (with $k = 1$ to match the sliced framework) based on Riemannian optimization reveals that our subgradient-based method is significantly faster in terms of iteration complexity and computation time. We also consider global optimization by showing that $\overline{\mathsf{W}}_p$ computation matches the framework of [36] for Lipschitz function optimization over convex domains. Adapting their LIPO algorithm to our problem, we obtain a provably consistent algorithm for computing $\overline{\mathsf{W}}_p$. However, the number of function evaluations that LIPO requires grows exponentially with dimension, which renders the locally optimal subgradient method preferable when dimension is large.

2 Background and Preliminaries

Notation. We use $\|\cdot\|$ for the Euclidean norm in $\mathbb{R}^d$. The operator norm for matrices is $\|\cdot\|_{\mathrm{op}}$. The unit sphere in $\mathbb{R}^d$ is denoted by $\mathbb{S}^{d-1}$, while $\mathbb{B}^d$ is the unit ball. Let $\mathcal{P}(\mathbb{R}^d)$ denote the space of Borel probability measures on $\mathbb{R}^d$ equipped with the TV metric $\|\mu - \nu\|_{\mathrm{TV}} = \frac{1}{2}|\mu - \nu|(\mathbb{R}^d)$, and set $\mathcal{P}_p(\mathbb{R}^d) := \{\mu \in \mathcal{P}(\mathbb{R}^d) : \int \|x\|^p d\mu(x) < \infty\}$ for $1 \le p < \infty$. The support of $\mu \in \mathcal{P}(\mathbb{R}^d)$ is denoted as $\mathrm{spt}(\mu)$, and we write $\mu \le \nu$ for setwise inequality. For for a measurable map f, the pushforward of μ under f is denoted as $f_\sharp \mu = \mu \circ f^{-1}$, i.e., if $X \sim \mu$ then $f(X) \sim f_\sharp \mu$. For two numbers a and b, we use the notation $a \wedge b = \min\{a, b\}$ and $a \vee b = \max\{a, b\}$. The distance between a set S and a point x in a metric $(\mathcal{X}, d)$ space is defined as $\mathrm{dist}(x, S) := \inf_{y \in S} d(x, y)$.

Some of our results assume log-concavity of the population distribution. A probability measure $\mu \in \mathcal{P}(\mathbb{R}^d)$ is *log-concave* if for every nonempty compact sets $A, B \subset \mathbb{R}^d$ and $\lambda \in [0, 1]$, we have $\mu\big(\lambda A + (1-\lambda)B\big) \ge \mu(A)^\lambda \mu(B)^{1-\lambda}$. A probability density function f on $\mathbb{R}^d$ is called *log-concave* if for every $x, y \in \mathbb{R}^d$ and $\lambda \in [0, 1]$, it satisfies $f\big(\lambda x + (1-\lambda)y\big) \ge f(x)^\lambda f(y)^{1-\lambda}$. Any non-degenerate distribution is log-concave if and only if it has a log-concave density [11, Theorem 1.1]. For $\beta \in (0, 2]$, let $\psi_\beta(t) = e^{t^\beta} - 1$ for $t \ge 0$, and recall that the corresponding Orlicz (quasi-)norm of a real-valued random variable X is defined as $\|X\|_{\psi_\beta} := \inf\{c > 0 : \mathbb{E}[\psi_\beta(|X|/c)] \le 1\}$. A Borel probability measure $\mu \in \mathcal{P}(\mathbb{R}^d)$ is called *sub-Gaussian* if $\|\|X\|\|_{\psi_2} < \infty$ for $X \sim \mu$.

Classic and sliced Wasserstein distances. For $1 \le p < \infty$, the p-Wasserstein distance between $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$ is $\mathsf{W}_p(\mu, \nu) := \inf_{\pi \in \Pi(\mu,\nu)} \left[\int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^p \, d\pi(x, y)\right]^{1/p}$, where $\Pi(\mu, \nu)$ is the set of couplings of μ and ν. W_p is a metric on $\mathcal{P}_p(\mathbb{R}^d)$ and metrizes weak convergence plus convergence of pth moments. While for $d > 1$ the definition of W_p generally amounts to an infinite-dimensional optimization problem, the expression simplifies when distributions are supported in $\mathbb{R}$. This motivates

the notion of the average- and max-sliced Wasserstein distances from (1). Both sliced distances are also metrics on $\mathcal{P}_p(\mathbb{R}^d)$ that induce the same topology as W_p [10, 44, 6, 42].

To present the simple one-dimensional formulae for W_p, for $\mu \in \mathcal{P}(\mathbb{R}^d)$ and $\theta \in \mathbb{S}^{d-1}$, let $F_\mu(t;\theta) := \mu\big(\{x \in \mathbb{R}^d : \theta^\intercal x \le t\}\big)$ be the distribution function of $\mathfrak{p}_\sharp^\theta \mu$, and $F_\mu^{-1}(\tau;\theta) = \inf\{t \in \mathbb{R} : F_\mu(t;\theta) \ge \tau\}$, for $\tau \in (0,1)$, be the quantile function. The W_p between measures on $\mathbb{R}$ amounts to the L^p distance between their quantile functions: $\mathsf{W}_p^p(\mathfrak{p}_\sharp^\theta \mu, \mathfrak{p}_\sharp^\theta \nu) = \int_0^1 \big|F_\mu^{-1}(\tau;\theta) - F_\nu^{-1}(\tau;\theta)\big|^p d\tau$. For $p = 1$, the expression further simplifies to $\mathsf{W}_1(\mathfrak{p}_\sharp^\theta \mu, \mathfrak{p}_\sharp^\theta \nu) = \int_{\mathbb{R}} \big|F_\mu(t;\theta) - F_\nu(t;\theta)\big|\, dt$.

Sliced Wasserstein distances between empirical distributions can be computed via order statistics. Let $\hat{\mu}_n := n^{-1}\sum_{i=1}^n \delta_{X_i}$ and $\hat{\nu}_n := n^{-1}\sum_{i=1}^n \delta_{Y_i}$ be the empirical distributions of samples $X_1, \dots, X_n$ and $Y_1, \dots, Y_n$. For each $\theta \in \mathbb{S}^{d-1}$, denote $X_i(\theta) = \theta^\intercal X_i$, and let $X_{(1)}(\theta) \le \cdots \le X_{(n)}(\theta)$ be the order statistics; define $Y_{(1)}(\theta) \le \cdots \le Y_{(n)}(\theta)$ analogously. By Lemma 4.2 in [8], we have $\mathsf{W}_p^p(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \hat{\nu}_n) = n^{-1}\sum_{i=1}^n \big|X_{(i)}(\theta) - Y_{(i)}(\theta)\big|^p$. The sliced distances $\underline{\mathsf{W}}_p$ and $\overline{\mathsf{W}}_p$ are computed by integrating or maximizing the above over $\theta \in \mathbb{S}^{d-1}$.

3 Empirical Convergence Rates

We study empirical convergence rates of sliced Wasserstein distances for log-concave distributions. The next result gives sharp one-sample rates with explicit dependence on the effective dimension.

Theorem 1 (Empirical convergence rates)**.** *Let $1 \le p < \infty$ and $n \ge 2$. Suppose that $\mu \in \mathcal{P}(\mathbb{R}^d)$ is log-concave with covariance matrix Σ and set $k = \mathrm{rank}(\Sigma)$. Then,*

$$\mathbb{E}\big[\underline{\mathsf{W}}_p(\hat{\mu}_n, \mu)\big] \lesssim_p \frac{\|\Sigma\|_{\mathrm{op}}^{1/2}\sqrt{(\log n)^{\mathbb{1}\{p=2\}}}}{n^{1/(2\vee p)}}, \tag{2a}$$

$$\mathbb{E}\big[\overline{\mathsf{W}}_p(\hat{\mu}_n, \mu)\big] \lesssim_p \frac{\|\Sigma\|_{\mathrm{op}}^{1/2} k \log n}{n^{1/p}} + \frac{\|\Sigma\|_{\mathrm{op}}^{1/2}\sqrt{k \log n}}{n^{1/(2\vee p)}} + \frac{\|\Sigma\|_{\mathrm{op}}^{1/2}\sqrt{(\log n)^{\mathbb{1}\{p=2\}}}}{n^{1/(2\vee p)}}. \tag{2b}$$

The proof of Theorem 1, in Appendix D.1, employs the machinery of [8] for analyzing empirical convergence of log-concave distributions on $\mathbb{R}$ based on their Cheeger constant (see Appendix A). For (2a), we show that log-concavity is preserved under projections and lower bound the Cheeger constant of the projected distribution by $c/\|\Sigma\|_{\mathrm{op}}$, uniformly in $\theta \in \mathbb{S}^{d-1}$. For (2b), concentration and covering arguments enables approximating the expected max-sliced distance by $\sup_{\theta \in \mathbb{S}^{d-1}} \mathbb{E}\big[\mathsf{W}_p(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \mu)\big]$, for which the aforementioned (uniform in θ) bounds are applicable.

Remark 1 (Lower bounds)**.** *The rate in* (2a) *is sharp up to log factors over the log-concave class. Corollary 6.14 in [8] implies that $\mathbb{E}[\underline{\mathsf{W}}_p(\hat{\mu}_n, \mu)]^p \ge c_p(\mathrm{tr}(\Sigma)/d)^{p/2}/\big(n(\log n)^{p/2}\big)$, for $\mu = \mathcal{N}(0, \Sigma) \in \mathcal{P}(\mathbb{R}^d)$ and any $p > 2$. For $p = 2$, a similar computation yields $\mathbb{E}[\underline{\mathsf{W}}_2(\hat{\mu}_n, \mu)]^2 \gtrsim (\mathrm{tr}(\Sigma)/d)\log\log n/n$, while for $p \in [1, 2)$, $\mathbb{E}[\underline{\mathsf{W}}_p(\hat{\mu}_n, \mu)] \ge \mathbb{E}[\underline{\mathsf{W}}_1(\hat{\mu}_n, \mu)] \gtrsim \sqrt{\tau(\Sigma)^2/n}$ where $\tau(\Sigma) = \frac{1}{d}\sum_{i=1}^d \sqrt{\lambda_i(\Sigma)}$ is the average of root eigenvalues of Σ. Since $\overline{\mathsf{W}}_p(\hat{\mu}_n, \mu) \gtrsim \underline{\mathsf{W}}_p(\hat{\mu}_n, \mu)$, this also yields a lower bound for $\overline{\mathsf{W}}_p$, while [47] gives a $\sqrt{d/n}$ lower bound under the T_p inequality.*

Remark 2 (Comparison with [47, 35])**.** *In [47], empirical rates for $\overline{\mathsf{W}}_p$ were derived under a high-level $T_{p'}(\sigma^2)$ assumption on μ. Our rate of decay from* (2b) *is faster, while replacing their entropy-transport inequality condition with log-concavity. The bounds for $\overline{\mathsf{W}}_p$ in [35, Theorem 3.6] assume the projection Poincaré inequality and $M_q := (\mu\|x\|^q)^{1/q} < \infty$ for $q > p$, and matches* (2a) *as $q \to \infty$ in terms of the dependence on n. However, their dependence on d is implicit through M_q which typically grows prohibitively with q and d. Our log-concavity assumption is strictly stronger than the Poincaré inequality, but yields a bound in terms of $\|\Sigma\|_{\mathrm{op}}$ which, for example, is constant in d when $\Sigma = I_d$.*[3] *Finally, we note that the bound for $\overline{\mathsf{W}}_p$ in* (2b) *adapts to the effective dimensionality k of the data, contrasting previously available bounds that depend on the ambient dimension d.*

Remark 3 (Concentration bounds)**.** *Combining the expectation bounds from Theorem 1 with [35, Theorem 3.8] yields concentration bounds for empirical sliced distances. These are presented in Appendix B and are later used to derive formal guarantees for computing $\overline{\mathsf{W}}_p$.*

[3] A recent preprint [5], that was posted on arXiv after this paper was submitted, shows that an improved estimate holds with high probability (compared to the convergence in expectation studied herein) for isotropic log-concave random vectors; cf. Equation (1.13) therein.

When $p > 2$, the rates in Theorem 1 are slower than parametric. Nevertheless, in the two-sample case with $\mu \neq \nu$, parametric rates are attainable uniformly in p for compactly supported distributions.

Proposition 1 (Parametric rates under the alternative)**.** *Let $1 \leq p < \infty$, and suppose that μ, ν have compact supports with $\operatorname{diam}\big(\operatorname{spt}(\mu)\big) \vee \operatorname{diam}\big(\operatorname{spt}(\nu)\big) \leq R$. Then,*

$$\mathbb{E}\big[\big|\underline{\mathsf{W}}_p^p(\hat{\mu}_n, \hat{\nu}_n) - \underline{\mathsf{W}}_p^p(\mu, \nu)\big|\big] \lesssim_{p,R} n^{-1/2} \quad \textit{and} \quad \mathbb{E}\big[\big|\overline{\mathsf{W}}_p^p(\hat{\mu}_n, \hat{\nu}_n) - \overline{\mathsf{W}}_p^p(\mu, \nu)\big|\big] \lesssim_{p,R} dn^{-1/2}.$$

If further $\mu \neq \nu$, then the same (parametric) rate also holds for empirical $\underline{\mathsf{W}}_p$ and $\overline{\mathsf{W}}_p$.

Proposition 1 is proven in Appendix D.2 using a comparison inequality between W_p and W_1 and elementary bounds for W_1 using its integral representation and KR duality.

Remark 4 (Comparison to [37])**.** *Theorem 2 of [37] establishes a bound of $(\log n/n)^{1/2}$ on the the two-sample average-sliced Wasserstein distance, but under bounded moment assumptions instead of compact support.*

4 Robust Estimation

We examine robustness of sliced Wasserstein distances to outliers, showing that slicing enables dimension-free risk bounds that avoid $\operatorname{poly}(d)$ factors present for classic W_p (cf. [46]). We consider TV corruptions, where an unknown "clean" distribution μ is contaminated to obtain $\tilde{\mu}$ with $\|\mu - \tilde{\mu}\|_{\mathrm{TV}} \leq \epsilon$. Upon observing $\tilde{\mu}$, the goal is to return a distribution $T(\tilde{\mu})$ such that the error $\mathsf{D}\big(T(\tilde{\mu}), \mu\big)$ is small, where $\mathsf{D} \in \{\underline{\mathsf{W}}_p, \overline{\mathsf{W}}_p\}$. Without further assumptions, this error can be unbounded, so we require that μ belongs to a family $\mathcal{G} \subset \mathcal{P}(\mathbb{R}^d)$ encoding standard moment bounds. We consider the minimax risk for robust estimation under D with TV contamination, defined by

$$R(\mathsf{D}, \mathcal{G}, \epsilon) = \inf_{T:\mathcal{P}(\mathbb{R}^d)\to\mathcal{P}(\mathbb{R}^d)} \sup_{(\mu,\tilde{\mu})\in\mathcal{G}\times\mathcal{P}(\mathbb{R}^d);\ \|\tilde{\mu}-\mu\|_{\mathrm{TV}}\leq\epsilon} \mathsf{D}\big(T(\tilde{\mu}), \mu\big).$$

Fix $q > p$ and let $\mathcal{G}_q(\sigma) := \big\{\mu \in \mathcal{P}_q(\mathbb{R}^d) : \sup_{\theta\in\mathbb{S}^{d-1}} \mu|\theta^\intercal(x - \mu x)|^q \leq \sigma^q\big\}$ contain all distributions whose projections have bounded central qth moments. In particular, $\mathcal{G}_2(\sigma) = \{\mu \in \mathcal{P}_2(\mathbb{R}^d) : \|\Sigma_\mu\|_{\mathrm{op}} \leq \sigma\}$. The next theorem characterizes minimax robust estimation risk over this class.

Theorem 2 (Population-limit robust estimation)**.** *Fix $1 \leq p < q$, $\sigma \geq 0$, and $0 \leq \epsilon \leq 0.49$.*[4] *We have $R\big(\underline{\mathsf{W}}_p, \mathcal{G}_q(\sigma), \epsilon\big) \asymp \sigma\sqrt{(1 \vee d/q)(1 \wedge p/d)}\,\epsilon^{1/p-1/q}$ and $R\big(\overline{\mathsf{W}}_p, \mathcal{G}_q(\sigma), \epsilon\big) \asymp \sigma\epsilon^{1/p-1/q}$.*

Note that the $\sqrt{(1 \vee d/q)(1 \wedge p/d)}$ prefactor in the first bound is always less than 1. The proof in Appendix D.4 controls the risk via $\sup_{\mu,\nu\in\mathcal{G}_q(\sigma),\|\mu-\nu\|_{\mathrm{TV}}\leq\epsilon} \mathsf{D}(\mu, \nu)$, a modulus of continuity that captures the sensitivity of D to small perturbations that preserve membership to the clean family. We employ techniques based on generalized resilience [61, 55] to relate this modulus to similar quantities arising in the robust estimation of pth moment tensors, giving the above rates. The procedure that achieves these rates projects the observed contaminated distribution onto the corresponding family of clean distributions in TV norm.

Remark 5 (Comparison to [46])**.** *A related framework [46] considers robust estimation of W_p under input measure contamination. They obtain a rate of $\sigma\sqrt{d}\epsilon^{1/p-1/2}$ using similar methods under the weaker Huber ϵ-contamination model when $q = 2$ (see Corollary 1 therein). Evidently, slicing eliminates a $\sqrt{d}$ factor from the minimax estimation risk. In Appendix D.4, we interpolate between these regimes, proving that k-dimensional sliced distances admit risks bounded by $\sigma\sqrt{1 \vee k/q}\,\epsilon^{1/p-1/q}$.*

Theorem 2 characterizes population-limit robust estimation, i.e., when data is abundant. The next result, proven in Appendix D.5, extends to the finite-sample regime. For a radius $R > 0$, we write μ_R to denote the distribution of $X \sim \mu$ conditioned on $\|X - \mu x\| \leq R$.

Proposition 2 (Finite-sample robust estimation)**.** *Fix $1 \leq p < q, \sigma \geq 0$, and $0 < \epsilon \leq 0.49$, and let $\mathsf{D} \in \{\underline{\mathsf{W}}_p, \overline{\mathsf{W}}_p\}$. Then there exists a radius $R \asymp \sqrt{d/\epsilon}$ and a procedure which, given $n \geq (R^p + \epsilon^{-2})\, d\log(d/\epsilon)$ samples with at least $(1 - \epsilon)n$ drawn i.i.d. from any $\mu \in \mathcal{G}_q(\sigma)$, returns $\nu \in \mathcal{P}(\mathbb{R}^d)$ such that $\mathsf{D}(\nu, \mu) \lesssim R\big(\mathsf{D}, \mathcal{G}_q(\sigma), \epsilon\big) + \mathbb{E}\big[\mathsf{D}\big((\widehat{\mu}_R)_n, \mu_R\big)\big]$ with probability at least 0.99*[5].

[4] The upper bound on ϵ of 0.49 can be substituted with an any constant bounded away from 1/2.

[5] See Appendix D.5 for precise high-probability bounds and extension to the strong contamination model.

Evidently, the finite-sample error bound comprises the population-limit robust estimation risk (which is necessary) plus the empirical estimation error associated with the truncated distribution μ_R. The lower bounds on n ensures that the empirical distribution $(\widehat{\mu}_R)_n$ satisfies the same generalized resilience property appearing in the population-limit analysis. The truncated empirical convergence term can typically be bounded by the corresponding untruncated version. For example, when $\mu \in \mathcal{G}_2(\sigma)$ is log-concave and $\mathsf{D} = \overline{\mathsf{W}}_1$, we can bound this term by $O\big(\sigma d \log n/n + \sigma\sqrt{d \log n/n}\big)$, which follows from Theorem 1 and the fact that μ_R is also log-concave with $\|\Sigma_{\mu_R}\|_{\mathrm{op}} \le \|\Sigma_\mu\|_{\mathrm{op}} \le \sigma$ for any $R > 0$.

When $p = 1$, we prove in Appendix D.6 a precise connection to resilience, a sufficient condition for robust mean estimation, which may be of independent interest.

Proposition 3 (Connection to mean resilience)**.** *For $0 \le \epsilon < 1$, $\mu \in \mathcal{P}_1(\mathbb{R}^d)$ is (ρ, ϵ)-resilient, i.e. $\|\mu x - \nu x\| \le \rho$ for all $\nu \le \frac{1}{1-\epsilon}\mu$, if and only if $\overline{\mathsf{W}}_1(\mu, \nu) \le \Theta(\rho)$ for all $\nu \le \frac{1}{1-\epsilon}\mu$.*

This suggests borrowing from the existing family of robust mean estimation algorithms, primarily developed for the bounded covariance setting ($q = 2$). In Appendix D.7, we inspect an efficient spectral reweighting procedure and apply it for both $\overline{\mathsf{W}}_p$ and $\underline{\mathsf{W}}_p$ when $1 \le p < 2$.

Proposition 4 (Efficient computation via spectral reweighting)**.** *If $1 \le p < q = 2$ and $0 \le \epsilon \le 1/12$, the guarantee of Proposition 2 is achieved by an $\widetilde{O}(nd^2)$-time spectral reweighting algorithm.*

5 Formal Computational Guarantees

The computational tractability of empirical sliced Wasserstein distances relies on the simplified expressions for W_p between distribution on $\mathbb{R}$. However, even then, evaluating $\underline{\mathsf{W}}_p$ and $\overline{\mathsf{W}}_p$ requires computing the average or the maximum of one-dimensional distances over projection directions $\theta \in \mathbb{S}^{d-1}$. This section provides formal guarantees for two such popular computational methods: MC integration for $\underline{\mathsf{W}}_p$ and alternating subgradient-based optimization for $\overline{\mathsf{W}}_p$. Our analysis relies on the observation that $w_p(\theta) := \mathsf{W}_p(\mathfrak{p}_\sharp^\theta \mu, \mathfrak{p}_\sharp^\theta \nu)$ and its pth power are Lipschitz functions on $\mathbb{S}^{d-1}$.

Lemma 1 (Lipschitz continuity)**.** *The functions w_p and w_p^p are Lipschitz with constants bounded by $L_{\mu,\nu}^p = \sup_{\theta \in \mathbb{S}^{d-1}} \big[(\mu|\theta^\intercal x|^p)^{1/p} + (\nu|\theta^\intercal x|^p)^{1/p}\big]$ and $M_{\mu,\nu}^p = 3p2^p \sup_{\theta \in \mathbb{S}^{d-1}} (\mu|\theta^\intercal x|^p + \nu|\theta^\intercal x|^p)$, respectively.*

Lemma 1 (proven in Appendix D.8) sharpens the Lipschitz constants derived in [47, Lemma 2], which correspond to bounding $|\theta^\intercal x|$ by $\|x\|$ in the above expressions. The projected moments $(\mu|\theta^\intercal x|^p)^{1/p}$ typically has a milder dependence on d than $(\mu\|x\|^p)^{1/p}$, which is crucial for the subsequent analysis.

5.1 Average-Slicing: Monte Carlo Integration

The typical approach for computing the integral over the unit sphere in $\underline{\mathsf{W}}_p$ is MC averaging. Fix $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$ and let $\hat{\mu}_n$ and $\hat{\nu}_n$ be the associated empirical measures. Take $\Theta \sim \mathrm{Unif}(\mathbb{S}^{d-1})$ and consider i.i.d. copies thereof $\Theta_1, \ldots, \Theta_m$. The MC based estimate of $\underline{\mathsf{W}}_p^p$ is given by

$$\widehat{\underline{\mathsf{W}}}_{\mathsf{MC}}^p := \frac{1}{m}\sum_{j=1}^m \mathsf{W}_p^p(\mathfrak{p}_\sharp^{\Theta_j}\hat{\mu}_n, \mathfrak{p}_\sharp^{\Theta_j}\hat{\nu}_n) = \frac{1}{mn}\sum_{j=1}^m\sum_{i=1}^n \big|X_{(i)}(\Theta_j) - Y_{(i)}(\Theta_j)\big|^p,$$

where $X_{(1)}(\theta) \le \cdots \le X_{(n)}(\theta)$ is the order statistics, which is readily evaluated using sorting algorithms with $O(n \log n)$ average/worst-case complexity (e.g., `quick_sort` or `merge_sort`).

The next result bounds the effective error of $\widehat{\underline{\mathsf{W}}}_{\mathsf{MC}}^p$ in approximating the population distance $\underline{\mathsf{W}}_p^p(\mu, \nu)$.

Proposition 5 (Monte Carlo error bound)**.** *Let $1 \le p < \infty$, and assume $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$ are log-concave with covariance matrices Σ_μ and Σ_ν, respectively. The MC estimate above satisfies*

$$\mathbb{E}\Big[\Big|\widehat{\underline{\mathsf{W}}}_{\mathsf{MC}}^p - \underline{\mathsf{W}}_p^p(\mu, \nu)\Big|\Big] \lesssim_p \frac{\|\mu x - \nu x\|^p + \|\Sigma_\mu\|_{\mathrm{op}}^{p/2} + \|\Sigma_\nu\|_{\mathrm{op}}^{p/2}}{\sqrt{md}} + \frac{\big(\|\Sigma_\nu\|_{\mathrm{op}}^{p/2} + \|\Sigma_\mu\|_{\mathrm{op}}^{p/2}\big)(\log n)^{\mathbb{1}\{p=2\}}}{n^{(p\wedge 2)/2}}$$

where the hidden constant depends only on p.

Algorithm 1 Projected subgradient method for $\tilde{w}_2^2$

Input: $\theta_0 \in \mathbb{B}^d$, a sequence $\{\alpha_t\}_{t\geq 0} \subset \mathbb{R}_+$, and iteration count T
for $t = 0, \ldots, T$ **do**
  Calculate $\xi_t \in \partial \tilde{w}_2^2(\theta_t)$
  Set $x_{t+1} = \text{Proj}_{\mathbb{B}^d}(x_t - \alpha_t \xi_t)$
Sample $t^* \in \{0, \ldots, T\}$ according to the probability distribution $\mathbb{P}(t^* = t) = \frac{\alpha_t}{\sum_{t=0}^T \alpha_t}$.
Return x_{t^*}

Proposition 5 is proven in Appendix D.9 by separately bounding the MC and the empirical approximation errors. For the former, we use the Lipschitzness of w_p^p on $\mathbb{S}^{d-1}$ to show that it concentrates about its median. This enables controlling the variance of $\frac{1}{m}\sum_{j=1}^m w_p^p(\Theta_j)$, which, in turn, bounds the MC error. For the empirical approximation error, we reduce the analysis to one-sample empirical convergence under W_p^p for measures on $\mathbb{R}$ and obtain explicit rates by drawing upon the results of [8].

Remark 6 (Comparison to [42])**.** *Error bounds for the MC estimate $\underline{\widehat{\mathsf{W}}}_{\mathsf{MC}}$ were also provided in [42], but their results differ from ours in two key ways: they use implicit empirical approximation bounds and leave their MC error in terms of $\mathrm{Var}\big(w_p^p(\Theta)\big)$ without further analysis. Proposition 5 provides an explicit convergence rates and bounds the said variance in terms of basic characteristics of the population distributions, providing precise rates in n, m, d, and p.*

Remark 7 (Blessing of dimensionality)**.** *A cruder approximation of the Lipschitz constant that stems from [47, Lemma 2] would yield $\mu\|x\|^p + \nu\|x\|^p$ as the numerator of the first term. However, such a bound can have a significantly worse dimension dependence. Indeed, if, for instance, μ and ν are both mean zero log-concave with identity covariance matrices, then $\mu\|x\|^p + \nu\|x\|^p$ is $O_d(d^{p/2})$ while the numerator in our bound is $O_d(1)$. For such μ and ν, the bound decays to 0 as $d \to \infty$.*

5.2 Max-Slicing: Subgradient Methods and the LIPO Algorithm

Maximization of projected Wasserstein distance is a non-convex and non-smooth optimization problem. Therefore, past works that studied k-dimensional subspace projections relied on convex relaxations [48] or entropic regularization [34, 25] to prove approximate convergence to a stationary point. We show that regularization is not needed in the one-dimensional case of $\overline{\mathsf{W}}_p$ by proving an $O(\epsilon^{-4})$ computational complexity bound for convergence to stationarity of the simple subgradient-based optimization routine from [28, 17]. We also note that global solutions are attainable via generic algorithms for optimizing Lipschitz functions, but with rates that deteriorate exponentially with d.

Local guarantees for subgradient methods. First note that we may relax the $\overline{\mathsf{W}}_p$ optimization domain from $\mathbb{S}^{d-1}$ to the unit ball $\mathbb{B}^d$ without changing the value (indeed, for any $\theta \in \mathbb{B}^d$, $w_p(\theta) = \|\theta\| w_p(\theta/\|\theta\|)$. Together with [8, Lemma 4.2], we express the empirical max-sliced distance as:

$$\overline{\mathsf{W}}_p^p(\hat{\mu}_n, \hat{\nu}_n) = \max_{\theta\in\mathbb{B}^d} \min_{\pi\in\Pi(\hat{\mu}_n,\hat{\nu}_n)} \mathbb{E}_\pi\left[|\theta^\intercal(X-Y)|^p\right] = -\min_{\theta\in\mathbb{B}^d}\max_{\sigma\in S_n}\left(-\frac{1}{n}\sum_{i=1}^n |\theta^\intercal(X_i - Y_{\sigma(i)})|^p\right),$$

where S_n is the symmetric group. Here we used the fact that the optimal coupling is given by the order statistics, and hence it suffices to optimize over permutations. Denote $\rho(\sigma,\theta) := -\frac{1}{n}\sum_{i=1}^n |\theta^\intercal(X_i - Y_{\sigma(i)})|^p$ and $\tilde{w}_p^p(\theta) := \max_{\sigma\in S_n}\rho(\sigma,\theta)$. The subgradient of $\tilde{w}_p^p$ has the closed form $\partial\tilde{w}_p^p(\theta) = \mathrm{Conv}\left(\left\{\partial_\theta\rho(\sigma^*,\theta) : \sigma^* \in \mathrm{argmax}_{\sigma\in S_n}\hat{\rho}(\sigma,\theta)\right\}\right)$. We can compute an optimal $\sigma^* \in S_n$ via order statistics and evaluate the corresponding subgradient vector in $\partial_\theta\rho(\sigma^*,\theta)$. This gives direct access to subgradients of $\tilde{w}_p^p$ without approximation arguments or regularization.

A heuristic description of Algorithm 1 was given in [28, 17], but without formal guarantees. Proposition 6 below can be viewed as closing that gap by providing said guarantees. In particular, for $p = 2$ the objective function $\tilde{w}_2^2$ is weakly convex [34, Lemma 2.2] and Lipschitz (Lemma 1). Together with the computable subgradients, this enables applying the proximal stochastic subgradient method from [14]. Algorithm 1 describes the adaptation of this method to our problem, after replacing the stochastic subgradient sampling step therein with the direct subgradient calculation described above. The following proposition provides convergence guarantees for Algorithm 1.

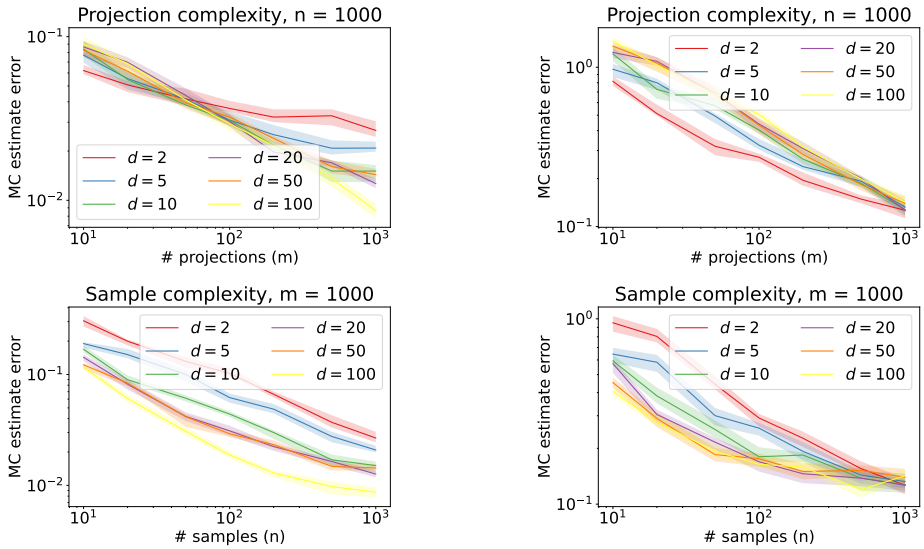

(a) $\left|\widehat{\underline{\mathsf{W}}}_{\mathsf{MC}}^2 - \underline{\mathsf{W}}_2^2(\mu,\nu)\right|$ under Model (1).

(b) $\left|\widehat{\underline{\mathsf{W}}}_{\mathsf{MC}}^2 - \underline{\mathsf{W}}_2^2(\mu,\nu)\right|$ under Model (2).

Figure 1: Projection and sample complexity for $\overline{\mathsf{W}}_2$.

Proposition 6 (Computational complexity of subgradient method)**.** *Fix any $\epsilon > 0$ and $n, d \in \mathbb{N}$ such that $d \geq (\log n)^2$. Let $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ be log-concave with covariance matrices Σ_μ and Σ_ν, respectively, and consider $M_{\mu,\nu}^2$ as defined in Lemma 1. Then, there exist universal constants c_1, c_2, c_3, c_4 such that the following holds: Algorithm 1 for the objective $\varphi(\theta) = \tilde{w}_2^2 + \delta_{\mathbb{B}^d}$, where $\delta_{\mathbb{B}^d} = -\infty \mathbb{1}_{(\mathbb{B}^d)^c}$, with step size $\alpha_t \propto \frac{1}{\sqrt{t+1}}$, outputs a point θ_{t^*} that is close to a near-stationary point θ^*, in the sense that $\mathbb{E}_{t^*}[\|\theta^* - \theta_{t^*}\|] \leq \frac{\epsilon}{2\rho_{\mu,\nu}}$, for $\rho_{\mu,\nu} = \|\mu x - \nu x\|^2 + c_1 d\left(\|\Sigma_\mu\|_{\mathrm{op}} + \|\Sigma_\nu\|_{\mathrm{op}}\right)$, and $\operatorname{dist}\big(0, \partial\tilde{w}_2^2(\theta^*)\big) \leq \epsilon$, within a number of computations $N \leq C_{\mu,\nu}\epsilon^{-4} n \log n$, where $C_{\mu,\nu} := c_2 \rho_{\mu,\nu}^2 \big(M_{\mu,\nu}^2 + c_3(\|\Sigma_\mu\|_{\mathrm{op}} + \|\Sigma_\nu\|_{\mathrm{op}})\big)^2$, with probability at least $1 - \frac{c_4}{n}$.*

Proposition 6 is proven in Appendix D.10 via the complexity bound from [14, Corollary 2]. As the algorithm is tuned for the empirical objective $\tilde{w}_p^p(\theta)$, the bound depends on the random Lipschitz and weak convexity constants $M_n = 4\sup_{\theta \in \mathbb{S}^{d-1}}(\hat{\mu}_n|\theta^\intercal x|^2 + \hat{\nu}_n|\theta^\intercal x|^2)$ and $\rho_n = 2\max_{i,j=1,\ldots,n}\|X_i - Y_j\|^2$. We use concentration bounds for M_n and ρ_n to obtain the deterministic bound above.

Remark 8 (Comparison to past works)**.** *Computation of projection-robust Wasserstein distances (i.e., when projections are k-dimensional) was studied in [48] and [34, 25] using a convex relaxations and entropic regularization, respectively. A similar $O(\epsilon^{-4})$ convergence rate is proven in [34] for their regularized method. Proposition 6 shows that regularization in not necessary to achieve this rate when projections are one-dimensional. The result of [34] was improved to $O(\epsilon^{-3})$ in [25] using Riemannian block coordinate descent (still with entropic regularization). While this rate is faster than in Proposition 6, our goal was to couple the simpler and abundantly used subgradient ascent approach with formal guarantees. In addition, the next section shows that empirically, our algorithm is much faster than those of [34, 25] for the $\overline{\mathsf{W}}_2$ in terms of complexity and computation time.*

Remark 9 (The non-quadratic case)**.** *For $p \neq 2$, we still have Lipschitzness of the objective function in θ (Lemma 1). Recent work on finding stationary points for non-smooth, non-convex, Lipschitz functions, such as [15], provide convergence guarantees for these cases. These guarantees appear to be of the same $\epsilon^{-1/4}$ order (cf. [15, Theorem 3.2]), but we leave a full exploration for future work.*

Remark 10 (Global guarantees via LIPO)**.** *We can attain global optimality, i.e., compute $\overline{\mathsf{W}}_p(\hat{\mu}_n, \hat{\nu}_n)$ itself, via the LIPO algorithm [36]. LIPO performs global optimization of Lipschitz functions over convex domains based on function evaluations, which are readily accessible in our problem via sorting. In Appendix C, we adapt LIPO to the max-sliced distance, prove consistency, and derive its complexity. While this approach attains global optimality, the number of evaluation grows exponentially with dimension. Hence, the subgradient method described above is preferable when dimension is large.*

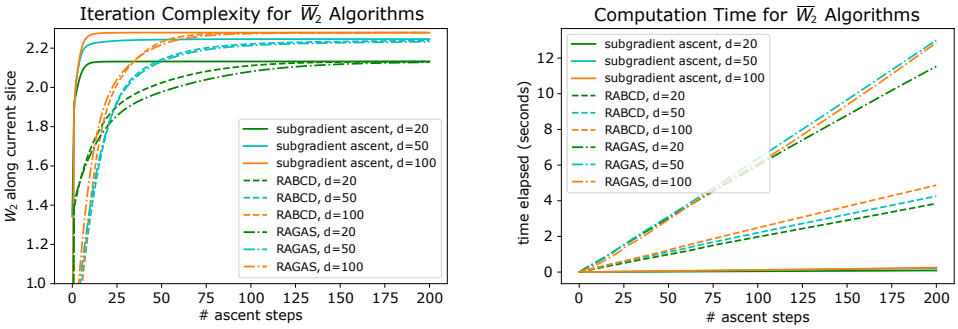

Figure 2: Errors and runtime versus step count for $\overline{\mathsf{W}}_2$ computation algorithms.

6 Empirical Results

Projection and sample complexity for $\underline{\mathsf{W}}_2$. We validate the convergence rates of the MC-based estimate of $\underline{\mathsf{W}}_2$ predicted by Propositions 5 in the following two models: (1) $\mu = \mathcal{N}(0, I_d)$, $\nu = 0.5\mathcal{N}(0, I_d) + 0.5\mathcal{N}(0, I_d + 0.5\mathbf{1}_d\mathbf{1}_d^\intercal/d)$, and (2) $\mu = \mathcal{N}(0, I_d)$, $\nu = \mathcal{N}(2\mathbf{1}_d, I_d)$, where $\mathbf{1}_d$ is a vector with all coordinates equal to 1. For Model (1), Proposition 5 predicts a decreasing error with dimension, and inverse square root decay in number of projections and number of samples. For Model (2), on the other hand, the errors should increase with d for sufficiently large n. Plots of the projection and sample complexities for each model (averaged over 100 runs) are given in Figure 1 at the top of the previous page, and are in line with the above discussion and our theoretical results. Additionally, confidence bands are plotted representing top and bottom 10% quantiles among 20 bootstrappped means from the same 100 runs. An additional experimental setup, comparing 10 component normal mixtures with different means and variances, can be found in Appendix E.

Comparison of $\overline{\mathsf{W}}_2$ algorithms. We compare the performance of the subgradient-based Algorithm 1 and the Riemannian optimization methods of [34, 25]. Consider the setup from [25, Section 6.1], where $\mu = \mathrm{Unif}([-1, 1]^d)$ and $\nu = T_\sharp\mu$, with $T(x) = x + \sum_{i=1}^{10} \mathrm{sign}(x_i)e_i$, is the fragmented hypercube distribution with $k^* = 10$. Figure 2 shows the errors and runtime by step count of Algorithm 1 (with a constant step size) and the Riemannian algorithms from [34, 25] (abbreviated RAGAS and RABCD, respectively) for different ambient dimensions. For these algorithms, we used the code from `https://github.com/mhhuang95/PRW_RBCD` with their default choice of parameters; we also tried optimizing these parameters but the observed trends remained the same. Sample size is fixed at $n = 500$ and computation times are averaged over 10 trials. Evidently, the subgradient ascent algorithm converges significantly faster and within fewer iterations than the other two methods, for all considered values of d. Despite our $O(\epsilon^{-4})$ iteration complexity bound, which is slower than the best known rates [25], this favorable empirical performance may be attributed to the fact that Algorithm 1 relies on the cheap sorting operation, as opposed to the burdensome computation of regularization operations in [34, 25]. It may also be the case that our bound can be improved, which we plan to explore in future work.

Robust estimation. To support Proposition 3, we perform robust estimation via a standard iterative filtering algorithm developed for mean estimation [19]. For $d \in \{10, 20, \ldots, 200\}$, we take $n = 10d\epsilon^{-2}$ samples, with $(1-\epsilon)n$ drawn i.i.d. from $\mathcal{N}(0, I_d)$ and ϵn from a product noise distribution used in [19], with $\epsilon = 0.1$. For each d, iterative filtering returns a candidate subset of clean samples, and Figure 3 (left) compares these subsets to the true clean samples both in $\overline{\mathsf{W}}_1$ (estimated via projected subgradient ascent) and in ℓ_2 distance between means. Note that the error in the latter never exceeds that in the former by more than a factor of 2 (Proposition 3 implies that mean and $\underline{\mathsf{W}}_1$ risk are equivalent up to constant factors). In Figure 3 (right), we set $\mu = (1-\epsilon)\delta_0 + \epsilon\,\mathrm{Unif}(\sqrt{d/\epsilon}\,\mathbb{S}^{d-1})$ with null contamination, and take n, d ranging as before. In this case, since many samples are 0, we can efficiently compute a lower bound on classic W_1 between the filtered and clean samples in high dimensions. As predicted by Theorem 2, we observe the $\sqrt{d}$ separation in estimation error between W_1 and $\overline{\mathsf{W}}_1$. In this case, only errors for the filtered samples are plotted, since the unfiltered samples have no contamination. See Appendix E for additional experiments on generative modeling with contaminated datasets, along with full details and code for all experiments.

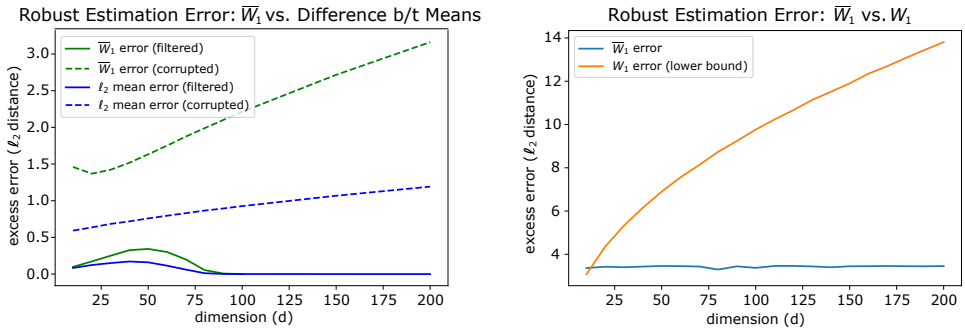

Figure 3: Robust estimation errors for the iterative filtering estimate in two scenarios: (left) comparing $\overline{\mathsf{W}}_1$ with difference between means and (right) comparing $\overline{\mathsf{W}}_1$ to W_1. Mean and $\overline{\mathsf{W}}_1$ errors are bounded as $d \to \infty$, while W_1 error scales like $\sqrt{d}$.

7 Summary and Concluding Remarks

This paper provided a quantitative study of the scalability of sliced Wasserstein distances to high dimensions. Three key aspects were covered:

- **Empirical convergence rates:** We established sharp, dimension-free rates for $\underline{\mathsf{W}}_p$ and $\overline{\mathsf{W}}_p$ with explicitly characterized dimension-dependent constants. Our bounds reveal the interplay between the number of samples n, dimension d, and the order of the distance p.
- **Robust estimation:** The minimax optimal robust estimation rate of $\underline{\mathsf{W}}_p$ and $\overline{\mathsf{W}}_p$, under contamination level ϵ, was characterized as $O(\sigma\epsilon^{1/p-1/q})$. This rate is dimension-free and improves upon corresponding results for W_p by a $\sqrt{d}$ factor. We showed that robust estimation of $\overline{\mathsf{W}}_1$ is equivalent to robust mean estimation, which enables lifting statistical/algorithmic results from means to $\overline{\mathsf{W}}_1$.
- **Computational guarantees:** The error of a MC-based estimator for $\underline{\mathsf{W}}_p$ was derived, showing that it can improve as $d \to \infty$, depending on the growth-rate of the mean and the operator norm of the covariance matrix. For $\overline{\mathsf{W}}_p$, we analyzed the subgradient-based local optimization algorithm from [28, 17], and proved $O(\epsilon^{-4})$ complexity using Lipschitzness and weak convexity of the objective.

In all three aspects, the benefit of slicing in terms of dependence on dimension was clearly evident, thus providing rigorous justification the perceived scalability of sliced distance. Going forward, we plan to pursue improved complexity bounds for the subgradient ascent algorithm for computing $\overline{\mathsf{W}}_2$, as our empirical results suggest it converges faster than Proposition 6 predicts. We are also interested in understanding conditions on μ, ν under which faster global guarantees for $\overline{\mathsf{W}}_p$ can be provided, e.g., by precluding the existence of nontrivial local optima for $\mathsf{W}_p(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \hat{\nu}_n)$ on $\mathbb{S}^{d-1}$, or matching the conditions of [36, Theorem 15], which results in polynomial and even exponential rates for LIPO. Extensions of our results to projection-robust Wasserstein distance, which considers projections to k-dimensional subspaces, are of interest, aiming to understand the effect of k on the results.

Acknowledgments and Disclosure of Funding

The authors thank Jason Gaitonde for helpful discussion on high-dimensional probability. S. Nietert was supported by a NSF Graduate Research Fellowship under Grant DGE-1650441. Z. Goldfeld is partially supported by NSF grants CCF-1947801, CCF-2046018, and DMS-2210368, and the 2020 IBM Academic Award. K. Kato is supported by NSF grants DMS-1952306, DMS-2014636, and DMS-2210368.

References

[1] R. Adamczak, R. Latała, A. E. Litvak, A. Pajor, and N. Tomczak-Jaegermann. Tail estimates for norms of sums of log-concave random vectors. *Proceedings of the London Mathematical Society*, 108(3):600–637, 2014.

[2] R. Adamczak, A. Litvak, A. Pajor, and N. Tomczak-Jaegermann. Quantitative estimates of the convergence of the empirical covariance matrix in log-concave ensembles. *Journal of the American Mathematical Society*, 23(2):535–561, 2010.

[3] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning (ICML)*, 2017.

[4] Y. Balaji, R. Chellappa, and S. Feizi. Robust optimal transport with applications in generative modeling and domain adaptation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

[5] D. Bartl and S. Mendelson. Structure preservation via the Wasserstein distance. *arXiv preprint arXiv:2209.07058*, 2022.

[6] E. Bayraktar and G. Guo. Strong equivalence between metrics of Wasserstein type. *Electronic Communications in Probability*, 26:1–13, 2021.

[7] S. G. Bobkov and C. Houdré. Isoperimetric constants for product probability measures. *The Annals of Probability*, 25(1):184 – 205, 1997.

[8] S. G. Bobkov and M. Ledoux. *One-dimensional empirical measures, order statistics, and Kantorovich transport distances*. American Mathematical Society, 2019.

[9] N. Bonneel, J. Rabin, G. Peyré, and H. Pfister. Sliced and Radon Wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 1(51):22–45, 2015.

[10] N. Bonnotte. *Unidimensional and evolution methods for optimal transportation*. PhD thesis, Paris 11, 2013.

[11] C. Borell. Convex measures on locally convex spaces. *Arkiv för Matematik*, 12(1-2):239–252, 1974.

[12] M. Carriere, M. Cuturi, and S. Oudot. Sliced Wasserstein kernel for persistence diagrams. In *International Conference on Machine Learning (ICML)*, 2017.

[13] Y. Chen. An almost constant lower bound of the isoperimetric coefficient in the KLS conjecture. *Geometric and Functional Analysis*, 31(1):34–61, 2021.

[14] D. Davis and D. Drusvyatskiy. Stochastic subgradient method converges at the rate $O(k^{-1/4})$ on weakly convex functions. *arXiv preprint arXiv:1802.02988*, 2018.

[15] D. Davis and D. Drusvyatskiy. A gradient sampling method with complexity guarantees for general Lipschitz functions. *arXiv preprint arXiv:2112.06969*, 2021.

[16] L. Deng. The MNIST database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.

[17] I. Deshpande, Y.-T. Hug, R. Sun, A. Pyrros, N. Siddiqui, S. Koyejo, Z. Zhao, D. Forsyth, and A. G. Schwing. Max-sliced Wasserstein distance and its use for GANs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[18] I. Deshpande, Z. Zhang, and A. G. Schwing. Generative modeling using the sliced Wasserstein distance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[19] I. Diakonikolas, G. Kamath, D. M. Kane, J. Li, A. Moitra, and A. Stewart. Being robust (in high dimensions) can be practical. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2017.

[20] D. L. Donoho and R. C. Liu. The "Automatic" Robustness of Minimum Distance Functionals. *The Annals of Statistics*, 16(2):552 – 586, 1988.

[21] Z. Goldfeld, K. Kato, G. Rioux, and R. Sadhu. Statistical inference with regularized optimal transport. *arXiv preprint arXiv:2205.04283*, 2022.

[22] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of Wasserstein GANs. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

[23] A. K. Gupta and S. Nadarajah, editors. *Handbook of beta distribution and its applications*, volume 174 of *Statistics: Textbooks and Monographs*. Marcel Dekker, Inc., New York, 2004.

[24] S. B. Hopkins, J. Li, and F. Zhang. Robust and heavy-tailed mean estimation made simple, via regret minimization. In *Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS)*, 2020.

[25] M. Huang, S. Ma, and L. Lai. A Riemannian block coordinate descent method for computing the projection robust Wasserstein distance. In *International Conference on Machine Learning*, pages 4446–4455. PMLR, 2021.

[26] R. Kannan, L. Lovász, and M. Simonovits. Isoperimetric problems for convex bodies and a localization lemma. *Discrete & Computational Geometry*, 13(3):541–559, 1995.

[27] I. Kim, S. Balakrishnan, and L. Wasserman. Robust multivariate nonparametric tests via projection averaging. *The Annals of Statistics*, 48(6):3417–3441, 2020.

[28] S. Kolouri, K. Nadjahi, U. Simsekli, R. Badeau, and G. K. Rohde. Generalized sliced Wasserstein distances. *arXiv preprint arXiv:1902.00434*, 2019.

[29] S. Kolouri, P. E. Pope, C. E. Martin, and G. K. Rohde. Sliced Wasserstein auto-encoders. In *International Conference on Learning Representations*, 2018.

[30] K. Le, H. Nguyen, Q. Nguyen, N. Ho, T. Pham, and H. Bui. On robust optimal transport: Computational complexity, low-rank approximation, and barycenter computation. *arXiv preprint arXiv:2102.06857*, 2021.

[31] M. Ledoux. Concentration of measure and logarithmic Sobolev inequalities. In *Seminaire de probabilites XXXIII*, pages 120–216. Springer, 1999.

[32] M. Ledoux and M. Talagrand. *Probability in Banach Spaces: isoperimetry and processes*, volume 23. Springer Science & Business Media, 1991.

[33] Y. T. Lee and S. S. Vempala. The Kannan-Lovász-Simonovits conjecture. *arXiv preprint arXiv:1807.03465*, 2018.

[34] T. Lin, C. Fan, N. Ho, M. Cuturi, and M. Jordan. Projection robust Wasserstein distance and Riemannian optimization. *Advances in neural information processing systems*, 33:9383–9397, 2020.

[35] T. Lin, Z. Zheng, E. Chen, M. Cuturi, and M. Jordan. On projection robust optimal transport: Sample complexity and model misspecification. In *International Conference on Artificial Intelligence and Statistics*, pages 262–270. PMLR, 2021.

[36] C. Malherbe and N. Vayatis. Global optimization of Lipschitz functions. In *International Conference on Machine Learning*, pages 2314–2323. PMLR, 2017.

[37] T. Manole, S. Balakrishnan, and L. Wasserman. Minimax confidence intervals for the sliced Wasserstein distance. *Electronic Journal of Statistics*, 16(1):2252–2345, 2022.

[38] O. Marchal and J. Arbel. On the sub-Gaussianity of the beta and Dirichlet distributions. *Electron. Commun. Probab.*, 22:Paper No. 54, 14, 2017.

[39] E. Milman. On the role of convexity in isoperimetry, spectral gap and concentration. *Inventiones mathematicae*, 177(1):1–43, 2009.

[40] D. Mukherjee, A. Guha, J. Solomon, Y. Sun, and M. Yurochkin. Outlier-robust optimal transport. In *International Conference on Machine Learning (ICML)*, 2021.

[41] K. Nadjahi, V. De Bortoli, A. Durmus, R. Badeau, and U. Şimşekli. Approximate Bayesian computation with the sliced-Wasserstein distance. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5470–5474. IEEE, 2020.

[42] K. Nadjahi, A. Durmus, L. Chizat, S. Kolouri, S. Shahrampour, and U. Simsekli. Statistical and topological properties of sliced probability divergences. *Advances in Neural Information Processing Systems*, 33:20802–20812, 2020.

[43] K. Nadjahi, A. Durmus, P. E. Jacob, R. Badeau, and U. Simsekli. Fast approximation of the sliced-Wasserstein distance using concentration of random projections. *Advances in Neural Information Processing Systems*, 34, 2021.

[44] K. Nadjahi, A. Durmus, U. Şimşekli, and R. Badeau. Asymptotic guarantees for learning generative models with the sliced-Wasserstein distance. *arXiv preprint arXiv:1906.04516*, 2019.

[45] J. S. Nath. Unbalanced optimal transport using integral probability metric regularization. *arXiv preprint arXiv:2011.05001*, 2020.

[46] S. Nietert, R. Cummings, and Z. Goldfeld. Outlier-robust optimal transport: duality, structure, and statistical analysis. In *International Conference on Artificial Intelligence and Statistics*, 2022.

[47] J. Niles-Weed and P. Rigollet. Estimation of Wasserstein distances in the spiked transport model. *arXiv preprint arXiv:1909.07513*, 2019.

[48] F.-P. Paty and M. Cuturi. Subspace robust Wasserstein distances. In *International Conference on Machine Learning*, pages 5072–5081. PMLR, 2019.

[49] J. Rabin, G. Peyré, J. Delon, and M. Bernot. Wasserstein barycenter and its application to texture mixing. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 435–446. Springer, 2011.

[50] A. Rakotomamonjy and R. Liva. Differentially private sliced Wasserstein distance. In *International Conference on Machine Learning*, pages 8810–8820. PMLR, 2021.

[51] G. Reeves. Conditional central limit theorems for gaussian projections. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pages 3045–3049. IEEE, 2017.

[52] M. C. Spruill. Asymptotic distribution of coordinates on high dimensional spheres. *Electron. Comm. Probab.*, 12:234–247, 2007.

[53] G. Staerman, P. Laforgue, P. Mozharovskyi, and F. d'Alché-Buc. When OT meets MoM: robust estimation of Wasserstein distance. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2021.

[54] J. Steinhardt. *Robust learning: Information theory and algorithms*. Stanford University, 2018.

[55] J. Steinhardt, M. Charikar, and G. Valiant. Resilience: A criterion for learning in the presence of arbitrary outliers. In A. R. Karlin, editor, *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, volume 94, pages 45:1–45:21, 2018.

[56] A. W. van der Vaart and J. A. Wellner. Weak convergence. In *Weak convergence and empirical processes*, pages 16–28. Springer, 1996.

[57] M. J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.

[58] J. Wu, Z. Huang, D. Acharya, W. Li, J. Thoma, D. P. Paudel, and L. V. Gool. Sliced Wasserstein generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3713–3722, 2019.

[59] J. Xi and J. Niles-Weed. Distributional convergence of the sliced Wasserstein process. *arXiv preprint arXiv:2206.00156*, 2022.

[60] X. Xu and Z. Huang. Central limit theorem for the sliced 1-Wasserstein distance and the max-sliced 1-Wasserstein distance. *arXiv preprint arXiv:2205.14624*, 2022.

[61] B. Zhu, J. Jiao, and J. Steinhardt. Generalized resilience and robust statistics. *arXiv preprint arXiv:1909.08755*, 2019.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] We state all assumptions needed for our results and compare them to existing work in several remarks.
   (c) Did you discuss any potential negative societal impacts of your work? [No] We provide a statistical analysis of an existing framework; applications of sliced optimal transport have societal implications but these are not addressed by this work.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] The result of each statement includes all needed assumptions.
   (b) Did you include complete proofs of all theoretical results? [Yes] Yes, see Appendixes A-D for all proofs.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See Appendix E for experimental details and code.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 6 and Appendix E for details.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] Yes, for experiments with significant statistical variation.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Appendix E for details.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] See Section 6 and Appendix E.
   (b) Did you mention the license of the assets? [No] The code which we used was not accompanied by licenses.
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We link to full code for experiments in Appendix E.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A] Our experiments use only synthetic data.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] Our experiments use only synthetic data.
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Cheeger Constant

Our empirical convergence rate analysis for the proof of Theorem 1 relies on controlling the Cheeger (isoperimetric) constant of the projected distributions. This section collects basic definitions and facts about Cheeger constants.

For $\mu \in \mathcal{P}(\mathbb{R}^d)$, define the boundary measure of a Borel subset $A \subset \mathbb{R}^d$ as

$$\mu^+(\partial A) := \liminf_{\epsilon \downarrow 0} \frac{\mu(A^\epsilon) - \mu(A)}{\epsilon},$$

where $A^\epsilon = \{x \in \mathbb{R}^d : d(x, A) \leq \epsilon\}$ is the ϵ-blowup of A, with $d(x, A) := \inf\{\|x - y\| : y \in A\}$. The *Cheeger constant* $h(\mu)$ of μ is defined as

$$h(\mu) := \inf_{A \subset \mathbb{R}^d} \frac{\mu^+(\partial A)}{\min\{\mu(A), \mu(A^c)\}},$$

which serves as a measure of bottleneckedness for μ. Indeed, a small $h(\mu)$ indicates the existence of a measurable $A \subset \mathbb{R}^d$ whose boundary measure is much smaller than the measure of A and A^c themselves. If μ has density f, then we also write $h(f) = h(\mu)$.

If $d = 1$ and μ has density f with distribution function F, then the Cheeger constant admits the simplified expression [7, Theorem 1.3]

$$h(\mu) = \text{essinf}_{x \in \mathbb{R}} \frac{f(x)}{\min\{F(x), 1 - F(x)\}}.$$

Furthermore, if F is strictly increasing around x, then for $t = F(x)$, we have

$$\frac{f(x)}{\min\{F(x), 1 - F(x)\}} = \frac{f(F^{-1}(t))}{\min\{t, 1 - t\}}$$

The numerator on the right-hand side (RHS) is denoted by $I(t) := f(F^{-1}(t))$; lower bounding this function plays a key role in our empirical convergence rate analysis. The main observation in that regard is that if f is log-concave in $d = 1$, then $\{x : 0 < F(x) < 1\}$ is an interval and f is positive on the interval, which implies $I(t) \geq h(f) \min\{t, 1 - t\}$ for $t \in (0, 1)$.

Consequently, lower bounding $I(t)$ reduces to controlling $h(f)$ from below. In general, it is known from [26] that if f is a log-concave density on $\mathbb{R}^d$ with covariance matrix Σ, then there exists a constant $c_d > 0$ that depends only on d such that

$$h(f) \geq \frac{c_d}{\|\Sigma\|_{\text{op}}^{1/2}}. \tag{3}$$

The KLS conjecture [26, 33] states that c_d can be chosen to be independent of d. The best available result up to date is due to [13], which shows that $c_d = 1/d^{o_d(1)}$ as $d \to \infty$.

The proof of the concentration inequalities in Proposition 7 below requires another property of log-concave distributions, namely the fact that they satisfy Poincaré inequalities. A probability measure $\mu \in \mathcal{P}(\mathbb{R}^d)$ is said to satisfy a Poincaré inequality with constant $M_\mu > 0$ if

$$\text{Var}_\mu(f) \leq M_\mu \mathbb{E}[\|\nabla f\|^2] \tag{4}$$

for any function $f : \mathbb{R}^d \to \mathbb{R}$ such that both sides of the above display are finite. The Maz'ya-Cheeger theorem (Theorem 1.1 in [39]) yields that $1/M_\mu \geq h(\mu)/2 > 0$, so that any (nondegenerate) log-concave distribution automatically satisfies a Poincaré inequality.

B Concentration Inequalities

We present concentration bounds for the empirical sliced distances as a corollary of Theorem 1. This result is utilized to provide global guarantees for computing $\overline{\mathsf{W}}_p$ via the LIPO algorithm [36] (cf. Proposition 8 in Appendix C).

Proposition 7 (Concentration inequalities)**.** *Let $1 \leq p < \infty$ and $n \geq 2$, and assume that $\mu \in \mathcal{P}(\mathbb{R}^d)$ is log-concave with non-singular covariance matrix Σ. Then, for any $t > 0$,*

$$\mathbb{P}\left(\underline{\mathsf{W}}_p(\hat{\mu}_n, \mu) \geq \frac{\left(\|\Sigma\|_{\mathrm{op}}(\log n)^{\mathbb{1}\{p=2\}}\right)^{1/2}}{n^{1/(2\vee p)}} + t\right) \leq 2\exp\left(-K\min\left\{n^{1/p}t, n^{2/(2\vee p)}t^2\right\}\right), \tag{5a}$$

$$\mathbb{P}\left(\overline{\mathsf{W}}_p(\hat{\mu}_n, \mu) \geq \alpha_{n,\mu} + t\right) \leq 2\exp\left(-K\min\left\{n^{1/p}t, n^{2/(2\vee p)}t^2\right\}\right), \tag{5b}$$

where $K \lesssim d^{o_d(1)}\max\{\|\Sigma\|_{\mathrm{op}}^{1/2}, \|\Sigma\|_{\mathrm{op}}\}$ and $\alpha_{n,\mu}$ is defined by the RHS of (2b) *with $k = d$.*

The proof of Proposition 7 combines the expectation bounds from Theorem 1 with the concentration inequality for empirical sliced Wasserstein distances from [35, Theorem 3.8]. The latter result holds under a Poincaré inequality assumption on the population distribution, which is always satisfied for log-concave measures (cf. [39, Theorem 1.1]), and is hence applicable for our setting.

The proof proceeds by lower bounding the Poincaré constant of μ, and then using a concentration result with expectation centering in [35] that relies on the Poincaré constant, combined with our expectation bounds (Theorem 1). By assumption, μ is log-concave with covariance matrix Σ_μ. This, in particular, implies that (3) holds, with $c_d = d^{o_d(1)}$ (cf. Theorem 1 in [13]). Combined with the Maz'ya-Cheeger inequality (Theorem 1.1 in [39]), this gives the following bound for the Poincaré constant M_μ of μ:

$$\frac{1}{M_\mu} \geq \frac{h(\mu)}{2} \geq \frac{1}{2d^{o_d(1)}\|\Sigma\|_{\mathrm{op}}}.$$

Now, by Theorem 3.8 in [35], we have

$$\mathbb{P}\big(|\rho(\hat{\mu}_n, \mu) - \mathbb{E}[\rho(\hat{\mu}_n, \mu)]| \geq t\big) \leq 2\exp\left(-K\min\left\{n^{1/p}t,\, n^{2/(2\vee p)}t^2\right\}\right),\ t > 0,$$

where $\rho = \underline{\mathsf{W}}_p$ or $\overline{\mathsf{W}}_p$, and K depends only on M_μ. A careful review of the proof of Theorem 3.8 and intermediate results in [31] yields that $1/\min\{2\sqrt{M_\mu}, 6e^5 M_\mu\}$ is a valid choice of K in the above display, so that $K \lesssim d^{o_d(1)}\max\{\|\Sigma\|_{\mathrm{op}}^{1/2}, \|\Sigma\|_{\mathrm{op}}\}$. Plugging (2a) and (2b) into the above display completes the proof.

C Global Guarantees for Max-Sliced W_p Computation via LIPO

We can compute $\overline{\mathsf{W}}_p(\hat{\mu}_n, \hat{\nu}_n) = \max_{\theta\in\mathbb{B}^d} \mathsf{W}_p(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \hat{\nu}_n)$ itself via the LIPO algorithm [36], which performs global optimization of Lipschitz functions over convex domains based on function evaluations. LIPO sequentially chooses the next evaluation point only if it can increase the function value, based on the Lipschitz condition. Setting $\hat{w}_p(\theta) := \mathsf{W}_p\big(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \hat{\nu}_n\big)$ and tuning LIPO to the (empirical) Lipschitz constant $\hat{L}_n^p := \sup_{\theta\in\mathbb{S}^{d-1}}\left[(\hat{\mu}_n|\theta^\intercal x|^p)^{1/p} + (\hat{\nu}_n|\theta^\intercal x|^p)^{1/p}\right]$ (see Lemma 1), if $\Theta_1, \ldots, \Theta_t$ are the t previous evaluation points, the next evaluation will be at Θ_{t+1} provided that

$$\min_{1\leq i\leq t}\left\{\hat{w}_p(\Theta_i) + \hat{L}_n^p\|\Theta_{t+1} - \Theta_i\|\right\} \geq \max_{1\leq i\leq t}\hat{w}_p(\Theta_i).$$

The output after k steps is $\max_{1\leq i\leq k}\hat{w}_p(\Theta_i)$. See [36, Figure 1] for the full pseudo-algorithm. We have the following global guarantee for the performance of LIPO.

Proposition 8 (LIPO error bound)**.** *Let $1 \leq p < \infty$ and assume that $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$ are log-concave with non-singular covariance matrices Σ_μ and Σ_ν, respectively. Let $\Theta_1, \ldots \Theta_k$ be a sequence of*

points generated by the LIPO for computing $\max_{\theta\in\mathbb{S}^{d-1}} \hat{w}_p(\theta)$*. Then for any* $t > 0$ *and* $n \geq C_p d^{p/2}$*, we have*

$$\mathbb{P}\left(\left|\overline{\mathsf{W}}_p(\mu,\nu) - \max_{1\leq i\leq k} \hat{w}_p(\Theta_i)\right| \leq 2L_{\mu,\nu}\left(\frac{\log(1/\delta)}{k}\right)^{1/d} + \alpha_n + 2t\right) \geq 1 - \delta - \beta - \gamma_n(t)$$

where $\alpha_n = \alpha_{n,\mu} + \alpha_{n,\nu}$ *with* $\alpha_{n,\mu}$ *given by the RHS of* (2b) *with* $k = d$*,* $\alpha_{n,\nu}$ *defined analogously,*

$$\begin{aligned} L_{\mu,\nu} &= (\|\Sigma_\mu\|_{\mathrm{op}}^{1/2} + \|\Sigma_\nu\|_{\mathrm{op}}^{1/2})\left((2p)^{1/p} \vee 2 + \frac{1}{2}\right) + \|\mu x\| + \|\nu x\|, \\ \beta &= \exp(-c_p\sqrt{d}), \quad \textit{and} \\ \gamma_n(t) &= 2\exp\left(-K_\mu \min\left(n^{1/p}t, n^{2/(2\vee p)}t^2\right)\right) + 2\exp\left(-K_\nu \min\left(n^{1/p}t, n^{2/(2\vee p)}t^2\right)\right), \end{aligned}$$

with $K_\mu \lesssim d^{o_d(1)} \max\{\|\Sigma_\mu\|_{\mathrm{op}}^{1/2},\, \|\Sigma_\mu\|_{\mathrm{op}}\}$ *and* $K_\nu \lesssim d^{o_d(1)} \max\{\|\Sigma_\nu\|_{\mathrm{op}}^{1/2},\, \|\Sigma_\nu\|_{\mathrm{op}}\}$.

The proof of Proposition 8 is given in Appendix D.11. The analysis separately bounds the empirical approximation error of the max-sliced objective and the error due to LIPO. The empirical error is treated using the concentration inequality from Proposition 7. For the LIPO analysis, we first argue that the (random) Lipschitz constant $\hat{L}_n^p$ concentrated about its mean and bound the latter by the population Lipschitz constant $L_{\mu,\nu}^p$. With this deterministic bound, the result follows from [36, Corollary 13]. Evidently, while Proposition 8 provides a global optimality guarantee, the resulting rate depends exponentially on dimension, which is too conservative in high-dimensional settings.

D Proofs of Results in the Main Text

Additional notation: We use $N(\epsilon, \mathcal{F}, d)$ to denote the ϵ-covering number of a function class or set $\mathcal{F}$ with respect to (w.r.t.) a metric d on $\mathcal{F}$, and $N_{[]}(\epsilon, \mathcal{F}, d)$ denotes the corresponding bracketing number.

D.1 Proof of Theorem 1

The proof relies on [8, Theorem 6.6], restated below, that bounds empirical convergence rates for W_p between distributions on $\mathbb{R}$.

Lemma 2 (Theorem 6.6 in [8])**.** *Fix* $1 \leq p < \infty$ *and* $n \geq 2$*. Let* $\mu \in \mathcal{P}(\mathbb{R})$ *have log-concave density* f *with distribution function* F*. Set* $I(t) = f\big(F^{-1}(t)\big)$ *for* $t \in (0,1)$*, where* F^{-1} *is the quantile function of* F*. Then,*

$$\mathbb{E}\big[\mathsf{W}_p^p(\hat{\mu}_n, \mu)\big] \leq \left(\frac{Cp^2}{n}\right)^{p/2} \int_{1/(n+1)}^{n/(n+1)} \frac{\big(t(1-t)\big)^{p/2}}{I^p(t)}\, dt, \tag{6}$$

where C *is a universal constant.*

We will apply Lemma 2 to $\mathsf{W}_p(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \mu)$ and bound the corresponding I-function from below *uniformly* over the projection parameter $\theta \in \mathbb{S}^{d-1}$. Recall that the distribution function of $\mathfrak{p}_\sharp^\theta \mu$ is denoted by $F_\mu(\cdot;\theta)$, which we abbreviate as F_θ throughout this proof and denote the corresponding density by f_θ. We first observe that since $\mu \in \mathcal{P}(\mathbb{R}^d)$ is log-concave, then so is $\mathfrak{p}_\sharp^\theta \mu$ for any $\theta \in \mathbb{S}^{d-1}$.

Let $h_\theta := h(\mathfrak{p}_\sharp^\theta \mu)$ denote the Cheeger constant of the projected distribution. From the discussion in Appendix A, we know that $1/\left(f_\theta\big(F_\theta^{-1}(t)\big)\right) \geq h_\theta \min\{t, 1-t\}$ for $t \in (0,1)$. Given that, the proof for the $\underline{\mathsf{W}}_p$ case is relatively straightforward from Lemma 2. Bounding $\mathbb{E}[\overline{\mathsf{W}}_p(\hat{\mu}_n, \mu)]$, however, requires extra work to treat the supremum over θ that appears inside the expectation.

$\underline{\mathsf{W}}_p$ case. Suppose that $\theta \in \mathbb{S}^{d-1}$ is such that $\theta^\intercal \Sigma \theta = 0$. Then, $\mathfrak{p}_\sharp^\theta \mu$ degenerates to a point mass, so that $\mathsf{W}_p^p\big(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \mu\big) = 0$.

Suppose $\theta^\intercal \Sigma \theta > 0$. Then, $\mathfrak{p}_\sharp^\theta \mu$ is nondegerate log-concave, so it has a log-concave density. Observe that $h_\theta \gtrsim 1/(\theta^\intercal \Sigma \theta)^{1/2}$. If $1 \le p < 2$, then

$$\int_0^1 \frac{\big(t(1-t)\big)^{p/2}}{t^p \wedge (1-t)^p}\, dt < \infty,$$

so that by Lemma 2, we have

$$\mathbb{E}\big[\mathsf{W}_p^p\big(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \mu\big)\big] \lesssim \left(\frac{\theta^\intercal \Sigma \theta}{n}\right)^{p/2}.$$

If $2 \le p < \infty$, then by Lemma 2, dividing $\int_{1/(n+1)}^{n/(n+1)}$ into $\int_{1/(n+1)}^{1/2} + \int_{1/2}^{n/(n+1)}$ and using the symmetry, we have

$$\mathbb{E}\big[\mathsf{W}_p^p\big(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \mu\big)\big] \lesssim \left(\frac{\theta^\intercal \Sigma \theta}{n}\right)^{p/2} \int_{1/(n+1)}^{1/2} t^{-p/2}\, dt.$$

Here

$$\int_{1/(n+1)}^{1/2} t^{-p/2}\, dt = \begin{cases} \big[\log(n+1) - \log 2\big] & p = 2 \\ \frac{1}{p/2-1}\big[(n+1)^{p/2-1} - 2^{p/2-1}\big] & p > 2 \end{cases},$$

so that

$$\mathbb{E}\big[\mathsf{W}_p^p\big(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \mu\big)\big] \lesssim \frac{(\theta^\intercal \Sigma \theta)^{p/2} (\log n)^{\mathbb{1}\{p=2\}}}{n}.$$

The result follows by noting that $\theta^\intercal \Sigma \theta \le \|\Sigma\|_{\mathrm{op}}$, integrating the display over $\theta \in \mathbb{S}^{d-1}$ and applying Fubini's theorem.

Remark 11 (Better bound for $\underline{\mathsf{W}}_2$)**.** *The above calculation actually yields the slightly better bound*

$$\mathbb{E}[\underline{\mathsf{W}}_2^2(\mu, \nu)] \lesssim \frac{k\|\Sigma\|_{\mathrm{op}} \log n}{nd}$$

for $p = 2$ by using the spectral decomposition $\Sigma = \sum_{i=1}^k \lambda_i a_i a_i^\intercal$, as follows:

$$\int_{\mathbb{S}^{d-1}} \theta^\intercal \Sigma \theta\, d\sigma(\theta) = \sum_{i=1}^k \int_{\mathbb{S}^{d-1}} (a_i^\intercal \theta)^2\, d\sigma(\theta) = \frac{1}{d} \sum_{i=1}^k \lambda_i \le \frac{k\|\Sigma\|_{\mathrm{op}}}{d}.$$

$\overline{\mathsf{W}}_p$ case. We divide the proof into two steps. In Step 1, we will prove the claim of the theorem when $k = d$, i.e., Σ is of full rank. In Step 2, we reduce the general case to the $d = k$ case.

Step 1. Assume $k = \mathrm{rank}(\Sigma) = d$. The main idea is to approximate $\mathbb{E}\big[\overline{\mathsf{W}}_p(\hat{\mu}_n, \mu)\big] = \mathbb{E}\big[\sup_{\theta \in \mathbb{S}^{d-1}} \mathsf{W}_p\big(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \mu\big)\big]$ by the maximum expected projected distance (roughly speaking, switch the expectation and the supremum). To that end we will employ a covering argument of the unit sphere along with Lipschitz continuity of $\mathsf{W}_p\big(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \mu\big)$ w.r.t. the samples and θ. These technical results are collected in the following lemma.

Lemma 3. *The following hold:*

(i) For any $\epsilon \in (0,1)$, we have $N(\epsilon, \mathbb{S}^{d-1}, \|\cdot\|) \le (5/\epsilon)^d$.

(ii) For any $\gamma \in \mathcal{P}_p(\mathbb{R})$, the map $u \mapsto \mathsf{W}_p(n^{-1}\sum_{i=1}^n \delta_{u_i}, \gamma)$ with $u = (u_1, \ldots, u_n)$ is $n^{-1/(2\vee p)}$-Lipschitz. Further, it is partially differentiable a.e. w.r.t. each u_i, and its partial derivative w.r.t. u_i is bounded by $n^{-1/p}$

(iii) The map $\theta \mapsto \mathsf{W}_p(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \mu)$ is L-Lipschitz with $L = (\sup_\theta \hat{\mu}_n |\theta^\intercal x|^p)^{1/p} + (\sup_\theta \mu |\theta^\intercal x|^p)^{1/p}$. If $\mu \in \mathcal{P}(\mathbb{R}^d)$ is centered and log-concave with non-singular covariance matrix Σ, then

$$\mathbb{E}[L] \le 2p\|\Sigma\|_{\mathrm{op}}^{1/2}\sqrt{d}$$

Proof of Lemma 3. (i) Follows from an elementary volumetric argument, which is omitted for brevity.

(ii) By the triangle inequality and definition of W_p, we have

$$\left|\mathsf{W}_p\left(n^{-1}\sum_{i=1}^n \delta_{u_i},\gamma\right)-\mathsf{W}_p\left(n^{-1}\sum_{i=1}^n \delta_{u_i'},\gamma\right)\right| \le \left(n^{-1}\sum_{i=1}^n |u_i-u_i'|^p\right)^{1/p}.$$

To bound the RHS by $n^{-1/(2\vee p)}\|u-u'\|$ we apply Jensen's inequality when $p\le 2$, and using the fact that $\sum_{i=1}^n a_i^{p/2} \le (\sum_{i=1}^n a_i)^{p/2}$ when $p\ge 2$.

The second statement follows from the fact that when coordinates other than u_i are kept fixed, the RHS of the above display is bounded by $\|u_i-u_i'\|$.

(iii) A simpler version is proven in [47, Lemma 2], but we include the argument for completeness. Applying Lemma 1 with $\nu=\hat\mu_n$, we obtain Lipschitz continuity with constant $L := \left(\sup_\theta \mu|\theta^\intercal x|^p\right)^{1/p} + \left(\sup_\theta \hat\mu_n|\theta^\intercal x|^p\right)^{1/p} \le (\mu\|x\|^p)^{1/p} + (\hat\mu_n\|x\|^p)^{1/p}$ so that $\mathbb{E}[L]\le 2(\mu\|x\|^p)^{1/p}$

Since μ is centered and log-concave with covariance matrix Σ, in particular

$$(\mu\|x\|^p)^{1/p} \le p(\mathbb{E}[\|X_1\|^2])^{1/2} = p\sqrt{\mathrm{tr}(\mathrm{Cov}(X_1))} \le p\sqrt{d}\|\Sigma\|_{\mathrm{op}}^{1/2}.$$

See, for example, Remark 1 after Theorem 3.1 in [1]. □

We are ready to prove the empirical convergence rate of the max-sliced distance. For the remainder of the proof we will assume, without loss of generality, that μ has mean 0, since $\overline{\mathsf{W}}_p(t^v_\sharp\mu, t^v_\sharp\nu) = \overline{\mathsf{W}}_p(\mu,\nu)$ for any location shift $t^v : x\mapsto x+v$ and any probability measures μ and ν, and any location shifted log-concave distribution is also log-concave with the same covariance matrix. Let $\tilde w_p(\theta) = \mathsf{W}_p(\mathfrak{p}^\theta_\sharp\hat\mu_n, \mathfrak{p}^\theta_\sharp\mu)$ and observe that

$$\mathbb{E}\big[\overline{\mathsf{W}}_p(\hat\mu_n,\mu)\big] \le \sup_{\theta\in\mathbb{S}^{d-1}} \mathbb{E}[\tilde w_p(\theta)] + \mathbb{E}\left[\sup_{\theta\in\mathbb{S}^{d-1}}\big(\tilde w_p(\theta)-\mathbb{E}\big[\tilde w_p(\theta)\big]\big)\right].$$

From the proof for the average-sliced case, we have

$$\sup_{\theta\in\mathbb{S}^{d-1}} \mathbb{E}\big[\tilde w_p(\theta)\big] \lesssim_p \frac{\left(\|\Sigma\|_{\mathrm{op}}(\log n)^{\mathbb{1}\{p=2\}}\right)^{1/2}}{n^{1/(2\vee p)}}. \tag{7}$$

Let $\theta_1,\dots,\theta_{N_\epsilon}$ be a minimal ϵ-net of $\mathbb{S}^{d-1}$, where $N_\epsilon = N(\epsilon,\mathbb{S}^{d-1},\|\cdot\|)$. Using Lemma 3, we have

$$\mathbb{E}\left[\sup_{\theta\in\mathbb{S}^{d-1}}\big(\tilde w_p(\theta)-\mathbb{E}\big[\tilde w_p(\theta)\big]\big)\right] \le \inf_{\epsilon>0}\mathbb{E}\left[\max_{1\le j\le N_\epsilon}\big(\tilde w_p(\theta_j)-\mathbb{E}\big[\tilde w_p(\theta_j)\big]\big)+2\epsilon L\right], \tag{8}$$

where L is a random variable with $\mathbb{E}[L]\le c_p\|\Sigma\|_{\mathrm{op}}^{1/2}d^{1/2}$.

To control the maximum inside the expectation on the RHS of (8), we use an approach based on maximal inequalities for sub-exponential random variables, similar to Theorem 3.5 in [35]. Briefly, we will first show that for each θ, $\tilde w_p(\theta)$ is a Lipschitz function of the projected observations $(\theta^\intercal X_1,\dots,\theta^\intercal X_n)$, with bounded gradient in each coordinate, which will imply sub-exponential concentration for each $\tilde w_p(\theta)$. The term $\max_{1\le j\le N_\epsilon}\big(\tilde w_p(\theta_j)-\mathbb{E}\big[\tilde w_p(\theta_j)\big]\big)$ will then be bounded via a maximal inequality as a direct consequence of this concentration (cf. Exercise 2.8 in [57]).

We will use the following refined concentration inequality for Lipschitz functions of random variables satisfying a Poincaré inequality, stated in [31]. Since explicit constants are not derived there, a proof is provided in Appendix D.3,

Lemma 4 (Concentration from Poincaré inequality; Corollary 4.6 in [31])**.** *Let $\mu\in\mathcal{P}(\mathbb{R}^d)$ satisfy the Poincaré inequality* (4) *with constant M_μ and $f:\mathbb{R}^{nd}\to\mathbb{R}$ be α-Lipschitz. For $x_1,\dots,x_n\in\mathbb{R}^d$, define the functions*

$$f_i(\cdot|x_1,\dots,x_{i-1},x_{i+1},\dots,x_n) := f(x_1,\dots,x_{i-1},\cdot,x_{i+1},\dots,x_n),\quad i=1,\dots,n,$$

and assume that $\max_{1\le i\le n}\sup_{x\in\mathbb{R}^d}\left\|\nabla f_i(x|X_1,\dots,X_{i-1},X_{i+1},\dots,X_n)\right\|\le\beta$ *a.s. Then,*

$$\mu^{\otimes n}(f\ge\mu^{\otimes n}f+t)\le\exp\left(-\min\left\{\frac{t}{2\beta M_\mu^{1/2}},\frac{t^2}{6e^5\alpha^2M_\mu}\right\}\right),\ t>0.$$

The random vector $(\theta^\intercal X_1, \dots, \theta^\intercal X_n)$ in $\mathbb{R}^n$ has i.i.d. coordinates with law $(\mathfrak{p}_\sharp^\theta \mu)^{\otimes n}$. The distribution $\mathfrak{p}_\sharp^\theta \mu$ is log-concave with variance $\theta^\intercal \Sigma \theta > 0$, which is bounded above by $\|\Sigma\|_{\mathrm{op}}$. This yields that $h_\theta := h(\mathfrak{p}_\sharp^\theta \mu) \gtrsim \|\Sigma\|_{\mathrm{op}}^{-1}$ for each $\theta \in \mathbb{S}^{d-1}$. By item (ii) in Lemma 3, the partial derivatives of $\tilde{w}_p(\theta)$ w.r.t. $\theta^\intercal X_i$, denoted $\nabla_i \tilde{w}_p(\theta)$, satisfy $\max_i \|\nabla_i \tilde{w}_p(\theta)\| \le n^{-1/p}$ a.s., and $\tilde{w}_p(\theta)$ is $n^{-1/(2\vee p)}$-Lipschitz in $(\theta^\intercal X_1, \dots, \theta^\intercal X_n)$. By the Maz'ya-Cheeger Theorem (cf. Theorem 1.1 in [39]), $M_{\mathfrak{p}_\sharp^\theta \mu}^{-1} \ge h_\theta/2 \ge \|\Sigma_\mu\|_{\mathrm{op}}^{-1}/2$. Combining these facts and applying Lemma 4, we have

$$\begin{aligned}\mathbb{P}\big(\tilde{w}_p(\theta) - \mathbb{E}[\tilde{w}_p(\theta)] > t\big) &\le \exp\left(- \min\left(\frac{t}{\sqrt{2} n^{-1/p} \|\Sigma\|_{\mathrm{op}}^{1/2}}, \frac{t^2}{3e^5 n^{-2/(2\vee p)} \|\Sigma\|_{\mathrm{op}}} \right) \right) \\ &\le \exp\left(- \frac{t^2}{\sqrt{2} n^{-1/p} \|\Sigma\|_{\mathrm{op}}^{1/2} t + 3e^5 n^{-2/(2\vee p)} \|\Sigma\|_{\mathrm{op}}} \right), \ t > 0.\end{aligned}$$

A simple union bound then gives

$$\mathbb{P}\Big(\max_{1 \le j \le N_\epsilon} \big(\tilde{w}_p(\theta_j) - \mathbb{E}\big[\tilde{w}_p(\theta_j)\big]\big) > t \Big) \le N_\epsilon \exp\left(- \frac{t^2}{\sqrt{2} n^{-1/p} \|\Sigma\|_{\mathrm{op}}^{1/2} t + 3e^5 n^{-2/(2\vee p)} \|\Sigma\|_{\mathrm{op}}} \right),$$

which, by an expectation bound for sub-exponential random variables (cf. Exercise 2.8 in [57]), yields

$$\begin{aligned}&\mathbb{E}\left[\max_{1 \le j \le N_\epsilon} \big(\tilde{w}_p(\theta_j) - \mathbb{E}\big[\tilde{w}_p(\theta_j)\big]\big) \right] \\ &\le \sqrt{6e^5 n^{-2/(2\vee p)} \|\Sigma\|_{\mathrm{op}}} (\sqrt{\pi} + \sqrt{\log N_\epsilon}) + 2\sqrt{2} n^{-1/p} \|\Sigma\|_{\mathrm{op}}^{1/2} (1 + \log N_\epsilon) \\ &\lesssim \|\Sigma\|_{\mathrm{op}}^{1/2} n^{-1/(2\vee p)} (1 + \sqrt{\log N_\epsilon}) + \|\Sigma\|_{\mathrm{op}}^{1/2} n^{-1/p} (1 + \log N_\epsilon). \qquad (9)\end{aligned}$$

By Lemma 3 (i), $\log N_\epsilon \le d \log(5/\epsilon)$. Thus, setting $\epsilon = n^{-1/2}$ and plugging (9) into (8), we have

$$\mathbb{E}\big[\overline{\mathsf{W}}_p(\hat{\mu}_n, \mu)\big] \lesssim_p \|\Sigma\|_{\mathrm{op}}^{1/2} \left(\frac{(\log n)^{\mathbb{1}_{\{p=2\}}}}{n^{1/(2\vee p)}} + \frac{(1 + \sqrt{d \log n})}{n^{1/(2\vee p)}} + \frac{(1 + d \log n)}{n^{1/p}} + \frac{d^{1/2}}{n^{1/2}} \right).$$

The last term on the RHS of the above display is of smaller order in n and d than the other two terms. Further, for $n \ge 2$, we have $(1 + \sqrt{d \log n}) \lesssim \sqrt{d \log n}$ and $(1 + d \log n) \lesssim d \log n$. This leads to the bound stated in Theorem 1 when $k = d$.

Step 2. Suppose now that $1 \le k < d$. Again assume without loss of generality that the mean of μ is zero. Observe that $\overline{\mathsf{W}}_p$ is invariant under common orthogonal transformations, i.e., for any $d \times d$ orthogonal matrix Q, $\overline{\mathsf{W}}_p(Q_\sharp \mu, Q_\sharp \nu) = \overline{\mathsf{W}}_p(\mu, \nu)$. With this in mind, we see that we may assume without loss of generality that Σ is diagonal whose first k diagonal entries are nonzero. Then, for $X = (X_1, \dots, X_d)^\intercal \sim \mu$ and $\theta = (\theta^1, \dots, \theta^d)^\intercal \in \mathbb{S}^{d-1}$, $\theta^\intercal X = \sum_{j=1}^k \theta^j X_j$ a.s. Thus, we have

$$\sup_{\theta \in \mathbb{S}^{d-1}} \mathsf{W}_p(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \mu) = \sup_{\substack{\theta = (\theta^1, \dots, \theta^d)^\intercal \in \mathbb{S}^{d-1} \\ \theta^{k+1} = \dots = \theta^d = 0}} \mathsf{W}_p(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \mu) \quad \text{a.s.}$$

The bound stated in Theorem 1 follows by the argument in Step 1 with d replaced by k. □

D.2 Proof of Proposition 1

Upper bound for $\underline{\mathsf{W}}_p^p$. Let $F_\mu(t, \theta) = \mathbb{P}(\theta^\intercal X \le t)$ for $X \sim \mu$, and analogously define $F_\nu(t, \theta)$. Then,

$$\begin{aligned}\mathbb{E}\big[\big|\underline{\mathsf{W}}_p^p(\hat{\mu}_n, \hat{\nu}_n) - \underline{\mathsf{W}}_p^p(\mu, \nu)\big|\big] &\le \mathbb{E}\left[\int_{\mathbb{S}^{d-1}} \Big| \mathsf{W}_p^p(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \hat{\nu}_n) - \mathsf{W}_p^p(\mathfrak{p}_\sharp^\theta \mu, \mathfrak{p}_\sharp^\theta \nu) \Big| \, d\sigma(\theta) \right] \\ &\le C_{p,R} \, \mathbb{E}\left[\int_{\mathbb{S}^{d-1}} \Big(\mathsf{W}_1(\mathfrak{p}_\sharp^\theta \hat{\mu}_n, \mathfrak{p}_\sharp^\theta \mu) + \mathsf{W}_1(\mathfrak{p}_\sharp^\theta \hat{\nu}_n, \mathfrak{p}_\sharp^\theta \nu) \Big) \, d\sigma(\theta) \right]\end{aligned}$$

$$
\begin{aligned}
&\leq \frac{C_{p,R}}{\sqrt{n}}\Bigg(\int_{\mathbb{S}^{d-1}}\int_{-\infty}^{\infty}\sqrt{F_\mu(t,\theta)(1-F_\mu(t,\theta))}\,dt\,d\sigma(\theta)\\
&\qquad+\int_{\mathbb{S}^{d-1}}\int_{-\infty}^{\infty}\sqrt{F_\nu(t,\theta)(1-F_\nu(t,\theta))}\,dt\,d\sigma(\theta)\Bigg)\\
&\leq \frac{RC_{p,R}}{\sqrt{n}},
\end{aligned}
$$

where the second inequality follows from a comparison between W_p and W_1 for compactly supported distributions (Lemma 4 in [21]), the third from the integral representation of W_1, and the final inequality from truncating the inner integrals to $[-R,R]$ and observing that $p(1-p)\leq 1/4$ for $p\in[0,1]$.

Upper bound for $\overline{\mathsf{W}}_p^p$. As earlier, observe that

$$
\begin{aligned}
\mathbb{E}\big[\big|\overline{\mathsf{W}}_p^p(\hat\mu_n,\hat\nu_n)-\overline{\mathsf{W}}_p^p(\mu,\nu)\big|\big] &\leq \mathbb{E}\left[\sup_{\theta\in\mathbb{S}^{d-1}}\left|\mathsf{W}_p^p(\mathfrak{p}_\sharp^\theta\hat\mu_n,\mathfrak{p}_\sharp^\theta\hat\nu_n)-\mathsf{W}_p^p(\mathfrak{p}_\sharp^\theta\mu,\mathfrak{p}_\sharp^\theta\nu)\right|\right]\\
&\leq C_{p,R}\,\mathbb{E}\left[\sup_{\theta\in\mathbb{S}^{d-1}}\Big(\mathsf{W}_1(\mathfrak{p}_\sharp^\theta\hat\mu_n,\mathfrak{p}_\sharp^\theta\mu)+\mathsf{W}_1(\mathfrak{p}_\sharp^\theta\hat\nu_n,\mathfrak{p}_\sharp^\theta\nu)\Big)\right]\\
&\leq C_{p,R}\left(\mathbb{E}\left[\sup_{\theta\in\mathbb{S}^{d-1}}\mathsf{W}_1(\mathfrak{p}_\sharp^\theta\hat\mu_n,\mathfrak{p}_\sharp^\theta\mu)\right]+\mathbb{E}\left[\sup_{\theta\in\mathbb{S}^{d-1}}\mathsf{W}_1(\mathfrak{p}_\sharp^\theta\hat\nu_n,\mathfrak{p}_\sharp^\theta\nu)\right]\right).
\end{aligned}
$$

Now, $\mathbb{E}\left[\sup_{\theta\in\mathbb{S}^{d-1}}\mathsf{W}_1(\mathfrak{p}_\sharp^\theta\hat\mu_n,\mathfrak{p}_\sharp^\theta\mu)\right]$ admits the following dual representation via KR duality:

$$
\mathbb{E}\left[\sup_{\theta\in\mathbb{S}^{d-1}}\mathsf{W}_1(\mathfrak{p}_\sharp^\theta\hat\mu_n,\mathfrak{p}_\sharp^\theta\mu)\right]=\mathbb{E}\left[\sup_{f\in\mathsf{Lip}_{1,0}(\mathbb{R}),\ \theta\in\mathbb{S}^{d-1}}(\hat\mu_n-\mu)(f\circ\mathfrak{p}^\theta)\right], \tag{10}
$$

where $\mathsf{Lip}_{1,0}(\mathbb{R})=\{f:\mathbb{R}\to\mathbb{R}\,:\,|f(x)-f(y)|\leq|x-y|\ \forall x,y\in\mathbb{R},f(0)=0\}$. By Lemma 8 in [21], the function class $\mathcal{G}=\{f\circ\mathfrak{p}^\theta:\theta\in\mathbb{S}^{d-1},f\in\mathsf{Lip}_{1,0}(\mathbb{R})\}$ is μ-Donsker, and we have

$$
\log N_{[\,]}(\epsilon,\mathcal{G},L^2(\mu))\lesssim R^{5/3}\epsilon^{-3/2}+5d\log(R/\epsilon),
$$

which, by the global maximal inequality (Theorem 2.14.2 in [56]), gives,

$$
\mathbb{E}\left[\sup_{f\in\mathsf{Lip}_{1,0}(\mathbb{R}),\ \theta\in\mathbb{S}^{d-1}}(\hat\mu_n-\mu)(f\circ\mathfrak{p}^\theta)\right]\lesssim\frac{R^{5/4}+d\log R}{\sqrt{n}}. \tag{11}
$$

Combining this with (10), and repeating the same argument for ν, we have the second statement.

The final statement of the theorem on $\underline{\mathsf{W}}_p$ and $\overline{\mathsf{W}}_p$ follows from the first two upper bounds combined with the elementary inequality $|a-b|\leq b^{1-p}|a^p-b^p|$ for $a\geq 0,b>0,p\geq 1$. □

D.3 Proof of Lemma 4

The proof of this lemma essentially recovers constants in Corollary 4.6 in [31], but a full argument is included for completeness. With some abuse of notation, let $\|\nabla f\|_\infty=\sup_x\|\nabla f(x)\|$. By Theorem 4.5 in [31], for any λ-Lipschitz function f with $\lambda\leq 2/\sqrt{M_\mu}$, we have

$$
\mathsf{Ent}_\mu(e^f)\leq B(\lambda)\mathbb{E}\big[\|\nabla f\|_\infty^2e^f\big], \tag{12}
$$

where $\mathsf{Ent}_\mu(f):=\mathbb{E}[f\log f]$ is the entropy functional of f, and

$$
B(\lambda)\leq\frac{M_\mu}{2}\left(\frac{2+\lambda\sqrt{M_\mu}}{2-\lambda\sqrt{M_\mu}}\right)e^{\sqrt{5M_\mu}\lambda}.
$$

Each function f_i in the statement of the proposition is β-Lipschitz. From (12) together with the tensorization property of the entropy functional (cf. Proposition 2.2 in [31]), we obtain

$$
\mathsf{Ent}_{\mu^{\otimes n}}\left(\frac{\lambda f}{\beta}\right)\leq\frac{\lambda^2}{\beta^2}\sum_{i=1}^n\mathbb{E}\left[\mathsf{Ent}_\mu\left(\frac{\lambda f_i}{\beta}\right)\right]\leq\frac{\lambda^2B(\lambda)}{\beta^2}\sum_{i=1}^n\mathbb{E}\big[\|\nabla f_i\|_\infty^2e^f\big],\quad\forall\lambda\in\big(0,2/\sqrt{M_\mu}\big].
$$

Further, we have $B(\lambda) \leq \frac{3e^5 M_\mu}{2}$ for $\lambda \leq 1/\sqrt{M_\mu}$ and $\sum_{i=1}^n \|\nabla f_i\|_\infty^2 \leq \beta$ $\mu^{\otimes n}$-a.e. by assumption. Therefore

$$\mathsf{Ent}_{\mu^{\otimes n}}\left(\frac{\lambda f}{\beta}\right) \leq \frac{3e^5 M_\mu \alpha^2}{2\beta^2}\lambda^2 \mathbb{E}\left[e^{\frac{\lambda f}{\beta}}\right], \quad \forall \lambda \in \left(0, 1/\sqrt{M_\mu}\right].$$

By Corollary 2.11 in [31], this yields

$$\mu^{\otimes n}\left(\frac{f}{\beta} > \mu\left(\frac{f}{\beta}\right) + r\right) \leq \exp\left(-\min\left\{\frac{r}{2\sqrt{M_\mu}}, \frac{r^2\beta^2}{6M_\mu e^5 \alpha^2}\right\}\right), \; r > 0,$$

from which the result follows by replacing r with r/β. □

D.4 Proof of Theorem 2

Denoting by $\wedge$ the setwise minimum of two measures, we first recall some useful facts. Throughout we write $\Theta \sim \mathrm{Unif}(\mathbb{S}^{d-1})$ for a random direction on the sphere sampled independently of any other randomness.

Fact 1. *For $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$ and $p \geq 1$, we have $\mathsf{W}_p(\mu, \nu) \leq \mathsf{W}_p(\mu - \mu \wedge \nu, \nu - \mu \wedge \nu)$.*

This follows by infimizing over transport plans which leave the shared mass $\mu \wedge \nu$ unmoved.

Fact 2. *For $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$, $p \geq 1$, and $c \geq 0$, we have $\mathsf{W}_p(c\mu, c\nu) = c^{1/p}\mathsf{W}_p(\mu, \nu)$.*

It is easy to check that these properties extend to $\underline{\mathsf{W}}_p$ and $\overline{\mathsf{W}}_p$. We also employ the following.

Lemma 5. *Fixing $p \geq 1$ and $\mu = \mathrm{Unif}(\mathbb{S}^{d-1})$, we have*

$$(\mu|x_1|^p)^{1/p} \asymp \sqrt{1 \wedge p/d}.$$

Lemma 6. *Fixing $p \geq 1$ and $\mu \in \mathcal{P}_p(\mathbb{R}^d)$, we have*

$$\mu(\|x\|^p)^{1/p} = \mathsf{W}_p(\mu, \delta_0) \asymp \sqrt{1 \vee d/p}\; \underline{\mathsf{W}}_p(\mu, \delta_0).$$

We defer proofs of the previous lemmas to Appendix D.4.4. In what follows, we will refer to any $\mathsf{D} : \mathcal{P}(\mathbb{R}^d)^2 \to [0, \infty]$ as a *statistical distance*, specifying additional properties as needed. Our risk bounds use the following standard lemma (see e.g. [20]), with a proof provided for completeness.

Lemma 7. *For any statistical distance D, corruption fraction $\epsilon \in [0, 1]$, and clean family $\mathcal{G} \subseteq \mathcal{P}(\mathbb{R}^d)$, define the modulus of continuity*

$$\mathfrak{m}(\mathsf{D}, \mathcal{G}, \epsilon) = \sup_{\substack{\mu, \nu \in \mathcal{G} \\ \|\mu - \nu\|_{\mathrm{TV}} \leq \epsilon}} \mathsf{D}(\mu, \nu). \tag{13}$$

We then have

$$\frac{1}{2}\mathfrak{m}(\mathsf{D}, \mathcal{G}, \epsilon) \leq R(\mathsf{D}, \mathcal{G}, \epsilon) \leq \mathfrak{m}(\mathsf{D}, \mathcal{G}, 2\epsilon).$$

Proof. For the lower bound, take any μ, ν feasible for (13). Then, if the statistician observes ϵ-contaminated measure $\tilde{\kappa} = \mu$, the clean measure could potentially be either $\kappa = \mu$ or $\kappa = \nu$. Hence any estimate $T(\tilde{\kappa})$ for the clean measure κ must incur error at least $D(\mu, \nu)/2$ in the worst case. For the upper bound, consider T which projects $\tilde{\kappa}$ onto $\mathcal{G}$ in TV. Then, $\|T(\tilde{\kappa}) - \kappa\|_{\mathrm{TV}} \leq \|T(\tilde{\kappa}) - \tilde{\kappa}\|_{\mathrm{TV}} + \|\tilde{\kappa} - \kappa\|_{\mathrm{TV}} \leq 2\epsilon$, and so $D(T(\tilde{\kappa}), \kappa) \leq \mathfrak{m}(\mathsf{D}, \mathcal{G}, 2\epsilon)$ by definition. □

By rescaling $\mathbb{R}^d$ appropriately, it is easy to check that $\mathfrak{m}(\mathsf{D}, \mathcal{G}_q(\sigma), \epsilon) = \sigma\mathfrak{m}(\mathsf{D}, \mathcal{G}_q(1), \epsilon)$ for our choices of D, so we will assume $\sigma = 1$ and write $\mathcal{G}_q = \mathcal{G}_q(1)$ from now on.

D.4.1 Lower bounds

Immediately, we can apply Lemma 7 to obtain the lower bounds of Theorem 2.

Proposition 9. *Fix $1 \leq p < q$ and corruption fraction $\epsilon \in [0, 1/2]$. Then we have*

$$R(\underline{\mathsf{W}}_p, \mathcal{G}_q, \epsilon) \gtrsim \sqrt{(1 \vee d/q)(1 \wedge p/d)}\; \epsilon^{1/p - 1/q}$$
$$R(\overline{\mathsf{W}}_p, \mathcal{G}_q, \epsilon) \gtrsim \epsilon^{1/p - 1/q}.$$

Proof. For $\overline{\mathsf{W}}_p$, we consider $\mu = \delta_0$ and $\nu = (1-\epsilon)\delta_0 + \epsilon\delta_y$ where $\|y\| = (2\epsilon)^{-1/q}$. Trivially, $\mu \in \mathcal{G}_q$, and

$$\begin{aligned}\sup_{\theta\in\mathbb{S}^{d-1}} \nu|\theta^\intercal(x-\nu y)|^q &= \sup_{\theta\in\mathbb{S}^{d-1}} (1-\epsilon)\epsilon^q|\theta^\intercal y|^q + \epsilon(1-\epsilon)^q|\theta^\intercal y|^q \\ &= \|y\|^q\left[(1-\epsilon)\epsilon^q + \epsilon(1-\epsilon)^q\right] \\ &\le \frac{1}{2}\epsilon^{-1}2\epsilon(1-\epsilon)^q \le 1,\end{aligned}$$

so $\nu \in \mathcal{G}_q$ as well. Moreover, we have

$$\overline{\mathsf{W}}_p(\mu,\nu) = \epsilon^{1/p}\overline{\mathsf{W}}_p(\delta_0,\delta_y) = 2^{-1/q}\epsilon^{1/p-1/q} \ge \frac{1}{2}\epsilon^{1/p-1/q},$$

and so Lemma 7 gives the desired risk bound for $\overline{\mathsf{W}}_p$.

For $\underline{\mathsf{W}}_p$, we fix $\mu = \delta_0$ and set $\nu = (1-\epsilon)\delta_0 + \epsilon\,\mathrm{Unif}(r\mathbb{S}^{d-1})$ with $r = \epsilon^{-1/q}\overline{\mathsf{W}}_q(\mathrm{Unif}(\mathbb{S}^{d-1}),\delta_0)^{-1}$. As before $\mu,\nu \in \mathcal{G}_q$, since

$$\begin{aligned}\overline{\mathsf{W}}_q(\nu,\delta_{\nu x}) &= \epsilon^{1/q}\,\overline{\mathsf{W}}_q(\mathrm{Unif}(r\mathbb{S}^{d-1}),\delta_0) \\ &= r\epsilon^{1/q}\,\overline{\mathsf{W}}_q(\mathrm{Unif}(\mathbb{S}^{d-1}),\delta_0) = 1.\end{aligned}$$

Furthermore, we have

$$\begin{aligned}\underline{\mathsf{W}}_p(\mu,\nu) &= \epsilon^{1/p}\,\overline{\mathsf{W}}_p(\mathrm{Unif}(r\mathbb{S}^{d-1}),\delta_0) \\ &= \epsilon^{1/p-1/q}\overline{\mathsf{W}}_q(\mathrm{Unif}(\mathbb{S}^{d-1}),\delta_0)^{-1}\,\overline{\mathsf{W}}_p(\delta_0,\mathrm{Unif}(\mathbb{S}^{d-1})) \\ &\asymp \sqrt{(1\vee d/q)(1\wedge p/d)}\,\epsilon^{1/p-1/q},\end{aligned}$$

where the last relation uses Lemma 5. Again, we obtain the desired risk bound via Lemma 7. □

D.4.2 Upper bounds

Next, we introduce an important notion of *(generalized) resilience* [55, 61]. We say that a distribution $\mu \in \mathcal{P}(\mathbb{R}^d)$ is (ρ,ϵ)-resilient w.r.t. a statistical distance D if $\mathsf{D}(\mu,\nu) \le \rho$ for all distributions $\nu \le \frac{1}{1-\epsilon}\mu$ (i.e. for all ϵ-deletions of μ). Standard (mean) resilience refers to resilience w.r.t. $\mathsf{D}_{\mathrm{mean}}(\mu,\nu) = \|\mu x - \nu x\|$. Writing $\mathcal{G}^{\mathsf{D}}_{\rho,\epsilon} \subset \mathcal{P}(\mathbb{R}^d)$ for the family of $\mu \in \mathcal{P}(\mathbb{R}^d)$ which are (ρ,ϵ)-resilient w.r.t. D, we have the following standard result.

Proposition 10. *Fix $0 \le \epsilon < 1/2$, $\rho \ge 0$, and D satisfying the triangle inequality. Then, we have $R(\mathsf{D},\mathcal{G}^{\mathsf{D}}_{\rho,2\epsilon},\epsilon) \le 2\rho$.*

Proof. Fix $\mu,\nu \in \mathcal{G}^{\mathsf{D}}_{\rho,2\epsilon}$ with $\|\mu-\nu\|_{\mathrm{TV}} \le 2\epsilon$. We consider the midpoint $\gamma = \frac{1}{1-\|\mu-\nu\|_{\mathrm{TV}}}\mu\wedge\nu \in \mathcal{P}(\mathbb{R}^d)$ and compute

$$\mathsf{D}(\mu,\nu) \le \mathsf{D}(\mu,\gamma) + \mathsf{D}(\nu,\gamma) \le 2\rho,$$

implying the desired risk bound via Lemma 7. □

We will also use the following standard result for one-dimensional (mean) resilience (see e.g. [55, Proposition 23]), which is a consequence of Markov's inequality.

Lemma 8. *Fix $0 \le \epsilon \le 1/2$ and $\mu \in \mathcal{P}(\mathbb{R})$ with $\mu|x-x_0|^p \le \sigma^p$ for some $x_0 \in \mathbb{R}$. Then, for all distributions $\nu \le \frac{1}{1-\epsilon}\mu$, we have $|\mu x - \nu x| \lesssim \sigma\epsilon^{1-1/p}$.*

Next, for $\mathsf{D} \in \{\underline{\mathsf{W}}_p, \overline{\mathsf{W}}_p\}$, we show that it suffices to prove resilience with respect to the simpler distances defined by

$$\begin{aligned}&\underline{\mathsf{D}}_p(\mu,\nu) = \underline{\mathsf{D}}_p(\mu-\nu) := \left|\mathbb{E}[(\mu-\nu)(|\Theta^\intercal x|^p)]\right| = |\underline{\mathsf{W}}^p_p(\mu,\delta_0) - \underline{\mathsf{W}}^p_p(\nu,\delta_0)| \\ \text{and}\quad &\overline{\mathsf{D}}_p(\mu,\nu) = \overline{\mathsf{D}}_p(\mu-\nu) := \sup_{\theta\in\mathbb{S}^{d-1}}\left|(\mu-\nu)(|\theta^\intercal x|^p)\right|,\end{aligned}$$

respectively. These distances encode a certain similarity of moment tensors, with $\overline{\mathsf{D}}_2(\mu,\nu) = \|\Sigma_\mu + (\mu x)(\mu x)^\intercal - \Sigma_\nu - (\nu x)(\nu x)^\intercal\|_{\mathrm{op}}$.

Recall that $\mathsf{D} = \mathsf{D}_{\mathcal{F}}$ is an integral probability metric (IPM) w.r.t. a class $\mathcal{F}$ of measurable functions on $\mathbb{R}^d$ if $\mathsf{D}(\mu,\nu) = \sup_{f\in\mathcal{F}}(\mu-\nu)(f)$. By design, we have the following.

Lemma 9. *The statistical distances $\underline{\mathsf{D}}_p$ and $\overline{\mathsf{D}}_p$ are IPMs w.r.t. the function classes $\underline{\mathcal{F}}_p = \{x \mapsto c_p s\|x\|^p : s \in \{\pm 1\}\}$ and $\overline{\mathcal{F}}_p = \{x \mapsto s|\theta^\intercal x|^p : s \in \{\pm 1\}, \theta \in \mathbb{S}^{d-1}\}$, respectively, where $c_p = \mathbb{E}[|\Theta_1|^p] \asymp \sqrt{1 \wedge p/d}$. Moreover, $\underline{\mathsf{D}}_p(\mu, \delta_0) = \underline{\mathsf{W}}_p^p(\mu, \delta_0)$ and $\overline{\mathsf{D}}_p(\mu, \delta_0) = \overline{\mathsf{W}}_p^p(\mu, \delta_0)$.*

Proof. For $\underline{\mathsf{D}}_p$, we compute

$$\begin{aligned}
\underline{\mathsf{D}}_p(\mu, \nu) &= \left|\mathbb{E}[(\mu - \nu)(|\Theta^\intercal x|^p)]\right| \\
&= \left|(\mu - \nu)(\mathbb{E}|\Theta^\intercal x|^p)\right| \\
&= \left|(\mu - \nu)(c_p\|x\|^p)\right| \\
&= \sup_{s \in \{\pm 1\}} (\mu - \nu)(c_p s\|x\|^p).
\end{aligned}$$

Likewise, for $\overline{\mathsf{D}}_p$, we check

$$\begin{aligned}
\overline{\mathsf{D}}_p(\mu, \nu) &= \sup_{\theta \in \mathbb{S}^{d-1}} \left|(\mu - \nu)(|\theta^\intercal x|^p)\right| \\
&= \sup_{s \in \{\pm 1\}, \theta \in \mathbb{S}^{d-1}} (\mu - \nu)(s|\theta^\intercal x|^p).
\end{aligned}$$

Computations when $\nu = \delta_0$ are trivial, since there is a single coupling between μ and ν. □

The third property is particularly relevant to resilience.

Lemma 10. *Let $\mathsf{D} = \mathsf{D}_{\mathcal{F}}$ be an IPM. Then μ is (ρ, ϵ)-resilient w.r.t. D if and only if μ is $(\epsilon(1 - \epsilon)^{-1}\rho, 1 - \epsilon)$-resilient w.r.t. D.*

Proof. Writing $\mu = (1 - \epsilon)\nu + \epsilon\alpha$ for some $\nu, \alpha \in \mathcal{P}(\mathbb{R}^d)$, we have

$$\begin{aligned}
\mathsf{D}(\nu, \mu) &= \mathsf{D}'(\epsilon^{-1}[\mu - (1 - \epsilon)\alpha] - \mu) \\
&= \frac{1 - \epsilon}{\epsilon}\mathsf{D}(\mu, \alpha) && \text{(homogeneity).}
\end{aligned}$$

□

We now formally translate resilience w.r.t. $\underline{\mathsf{D}}_p$ and $\overline{\mathsf{D}}_p$ to that which we desire.

Proposition 11. *Fix $0 \le \epsilon < 1$, $\rho \ge 0$, and $(\mathsf{D}, \mathsf{D}') \in \{(\overline{\mathsf{W}}_p, \overline{\mathsf{D}}_p), (\underline{\mathsf{W}}_p, \underline{\mathsf{D}}_p)\}$. If $\mu \in \mathcal{P}(\mathbb{R}^d)$ with $\mu x = 0$ is (ρ, ϵ)-resilient w.r.t. D', then μ is $(2\rho^{1/p} + 2\epsilon^{1/p}\mathsf{D}(\mu, \delta_0), \epsilon)$-resilient w.r.t. D.*

Proof. Fixing such μ and taking $\nu \le \frac{1}{1-\epsilon}\mu$, write $\mu = (1 - \epsilon)\nu + \epsilon\alpha$ for some $\alpha \in \mathcal{P}(\mathbb{R}^d)$ and write $\tau = \epsilon \wedge (1 - \epsilon)$, so that $\nu, \alpha \le \tau^{-1}\mu$. Then, we bound

$$\begin{aligned}
\mathsf{D}(\mu, \nu)^p &= \mathsf{D}((1 - \epsilon)\nu + \epsilon\alpha, \nu)^p \\
&\le \epsilon\,\mathsf{D}(\alpha, \nu)^p && \text{(Facts 1 and 2)} \\
&\le 2^p\epsilon \sup_{\kappa \le \tau^{-1}\mu} \mathsf{D}(\kappa, \delta_0)^p && \text{(triangle inequality for D)} \\
&= 2^p\epsilon \sup_{\kappa \le \tau^{-1}\mu} \mathsf{D}'(\kappa, \delta_0) && \text{(Lemma 9)} \\
&\le 2^p\epsilon \sup_{\kappa \le \tau^{-1}\mu} \mathsf{D}'(\kappa, \mu) + 2^p\epsilon\,\mathsf{D}(\mu, \delta_0)^p && \text{(triangle inequality for D}'\text{)} \\
&\le 2^p\rho + 2^p\epsilon\,\mathsf{D}(\mu, \delta_0)^p, && \text{(Lemma 10)}
\end{aligned}$$

giving the desired bound after taking pth roots. □

Equipped with this result, we are prepared to prove the upper bounds of Theorem 2. Given $\mu \in \mathcal{G}_q$, we must provide bounds on $\mathsf{D}(\mu, \delta_0)$ as well as the resilience of μ w.r.t. D'.

Lemma 11. *Fixing $1 \le p < q$ and $\mu \in \mathcal{G}_q$ with $\mu x = 0$, we have*

$$\begin{aligned}
\underline{\mathsf{W}}_p(\mu, \delta_0) &\lesssim \sqrt{(1 \vee d/q)(1 \wedge p/d)} \\
\overline{\mathsf{W}}_p(\mu, \delta_0) &\lesssim 1.
\end{aligned}$$

Lemma 12. *Fix $1 \le p < q$, corruption fraction $0 \le \epsilon \le 1/2$, and $\mu \in \mathcal{G}_q$ with $\mu x = 0$. Then, μ is $(C\sqrt{(1 \vee d/q)(1 \wedge p/d)}\epsilon^{1/p-1/q}, \epsilon)$-resilient w.r.t. $\underline{\mathsf{D}}_p^{1/p}$ and $(C\epsilon^{1/p-1/q}, \epsilon)$-resilient w.r.t. $\overline{\mathsf{D}}_p^{1/p}$, for some absolute constant $C > 0$.*

Together, these give the desired risk bounds.

Proposition 12. *Fix $1 \le p < q$ and corruption fraction $0 \le \epsilon \le 0.49$. Then we have*

$$R(\underline{\mathsf{W}}_p, \mathcal{G}_q, \epsilon) \lesssim \sqrt{(1 \vee d/q)(1 \wedge p/d)}\, \epsilon^{1/p-1/q}$$
$$R(\overline{\mathsf{W}}_p, \mathcal{G}_q, \epsilon) \lesssim \epsilon^{1/p-1/q}.$$

Proof. Fixing $\mu \in \mathcal{G}_q$, it suffices by Proposition 10 to prove that μ is $(C\epsilon^{1/p-1/q}, \epsilon)$-resilient w.r.t. $\overline{\mathsf{W}}_p$ and $(C\sqrt{(1 \vee d/q)(1 \wedge p/d)}\, \epsilon^{1/p-1/q}, \epsilon)$-resilient w.r.t. $\underline{\mathsf{W}}_p$ for all $0 \le \epsilon \le 0.98$, where $C > 0$ is some absolute constant. Since these distances are translation invariant, we can assume without loss of generality that $\mu x = 0$. By Lemmas 11 and 12, we know that μ is $(C\epsilon^{1/p-1/q}, \epsilon)$-resilient w.r.t. $\overline{\mathsf{D}}_p^{1/p}$ and $(C\sqrt{(1 \vee d/q)(1 \wedge p/d)}\, \epsilon^{1/p-1/q}, \epsilon)$-resilient w.r.t. $\underline{\mathsf{D}}_p^{1/p}$ for all $0 \le \epsilon \le 1/2$ and some absolute constant $C > 0$. For $1/2 \le \epsilon \le 0.98$, the same resiliency bounds are implied by Lemma 10, since $1 - \epsilon \ge 0.02 \ge \epsilon/49$. Finally, we apply Proposition 11 to obtain the desired risk bounds. □

We now prove the preceding lemmas.

Proof of Lemma 11. Fixing $\mu \in \mathcal{G}_q$ with $\mu x = 0$, we bound

$$\begin{aligned}
\underline{\mathsf{W}}_p(\mu, \delta_{\mu x}) &\asymp \sqrt{1 \wedge p/d}\, \mathsf{W}_p(\mu, \delta_{\mu x}) && \text{(Lemma 6)} \\
&\le \sqrt{1 \wedge p/d}\, \mathsf{W}_q(\mu, \delta_{\mu x}) && (q > p) \\
&= \sqrt{1 \wedge p/d}\, \frac{\mathsf{W}_q(\mu, \delta_{\mu x})}{\underline{\mathsf{W}}_q(\mu, \delta_{\mu x})} \underline{\mathsf{W}}_q(\mu, \delta_{\mu x}) \\
&\asymp \sqrt{(1 \wedge p/d)(1 \vee d/q)}\, \underline{\mathsf{W}}_q(\mu, \delta_{\mu x}) && \text{(Lemma 6)} \\
&\le \sqrt{(1 \wedge p/d)(1 \vee d/q)} && \mu \in \mathcal{G}_q. \qquad (14)
\end{aligned}$$

Similarly, we obtain

$$\overline{\mathsf{W}}_p(\mu, \delta_{\mu x}) \le \overline{\mathsf{W}}_q(\mu, \delta_{\mu x}) \le 1.$$

□

Proof of Lemma 12. By Lemma 9, $\underline{\mathsf{D}}_p$ and $\overline{\mathsf{D}}_p$ are IPMs with respect to the stated function classes $\underline{\mathcal{F}}_p$ and $\overline{\mathcal{F}}_p$, respectively. Note that if $\mathsf{D} = \mathsf{D}_{\mathcal{F}}$ is an IPM for any symmetric $\mathcal{F} = -\mathcal{F}$, then μ is (ρ, ϵ)-resilient w.r.t. D if and only if $f_\sharp \mu$ is (ρ, ϵ)-resilient (in mean) for all $f \in \mathcal{F}$.

Now, fix $\mu \in \mathcal{G}_q$ with $\mu x = 0$. For $\underline{\mathsf{D}}_p$, we observe that

$$\begin{aligned}
\mu((\|x\|^p)^{q/p}) &= \mu(\|x\|^q) \\
&\lesssim C^q (1 \vee d/q)^{q/2} \sup_{\theta \in \mathbb{S}^{d-1}} \mu(|\theta^\intercal x|^q) && \text{(Lemma 6)} \\
&\le C^q (1 \vee d/q)^{q/2} && (\mu \in \mathcal{G}_q) \\
&= \left[C^p (1 \vee d/q)^{p/2}\right]^{q/p},
\end{aligned}$$

for some absolute constant $C > 0$. For $f \in \underline{\mathcal{F}}_p$, we then have that $f_\sharp \mu$ has q/p-th moments bounded by $O(C^p (1 \wedge p/d)^{p/2} (1 \vee d/q)^{p/2})$, and is thus $(O(C^p (1 \wedge p/d)^{p/2} (1 \vee d/q)^{p/2} \epsilon^{1-p/q}), \epsilon)$-resilient, by Lemma 8. Taking pth roots gives the claim.

For $\overline{\mathsf{D}}_p$, note that for $\theta \in \mathbb{S}^{d-1}$, we have

$$\mu((|\theta^\intercal x|^p)^{q/p}) = \mu(|\theta^\intercal x|^q) \le 1.$$

For $f \in \overline{\mathcal{F}}_p$, we then have that $f_\sharp \mu$ has q/p-th moments bounded by 1 and is thus $O(\epsilon^{1-p/q}, \epsilon)$-resilient. Taking pth roots gives the claim. □

D.4.3 Higher-dimensional slicing

We now extend Proposition 12 to the k-dimensional sliced distances defined by

$$\underline{\mathsf{W}}_{p,k}(\mu,\nu) := \left[\int_{\mathrm{Gr}_k(\mathbb{R}^d)} \mathsf{W}_p^p(\mathfrak{p}^E_\sharp \mu, \mathfrak{p}^E_\sharp \nu) d\sigma_k(E)\right]^{1/p} \quad \text{and} \quad \overline{\mathsf{W}}_{p,k}(\mu,\nu) := \sup_{E\in \mathrm{Gr}_k(\mathbb{R}^d)} \mathsf{W}_p(\mathfrak{p}^E_\sharp \mu, \mathfrak{p}^E_\sharp \nu),$$

where $\mathrm{Gr}_k(\mathbb{R}^d)$ is the Grassmannian of k-dimensional linear subspaces of $\mathbb{R}^d$, σ_k is its standard Haar measure, and $\mathfrak{p}^E$ is the orthogonal projection onto $E \in \mathrm{Gr}_k(\mathbb{R}^d)$. These coincide with $\underline{\mathsf{W}}_p$ and $\overline{\mathsf{W}}_p$ when $k = 1$, and both equal W_p when $k = d$. We focus here on $\overline{\mathsf{W}}_p$, with $\underline{\mathsf{W}}_p$ inheriting the same risk bound, although stronger guarantees can be obtained in a similar manner to the proof of Proposition 12. First, we extend $\overline{\mathsf{D}}_p$ to this regime as

$$\overline{\mathsf{D}}_{p,k} := \sup_{\substack{U\in\mathbb{R}^{d\times k}\\ U^\intercal U = I_k}} \left|(\mu-\nu)(\|U^\intercal x\|^p)\right|,$$

and observe that all of the properties from Lemma 9 still hold. Moreover, for $\mu \in \mathcal{G}_q$ with $\mu x = 0$, we obtain the needed analog of Lemma 11, bounding

$$\begin{aligned}
\overline{\mathsf{W}}_{p,k}(\mu,\delta_0)^{1/p} &= \sup_{\substack{U\in\mathbb{R}^{d\times k}\\ U^\intercal U = I_k}} \mu(\|U^\intercal x\|^p)^{1/p} \\
&\le \sup_{\substack{U\in\mathbb{R}^{d\times k}\\ U^\intercal U = I_k}} \mu(\|U^\intercal x\|^q)^{1/q} && (q > p) \\
&\le \sqrt{1\vee k/q} \sup_{\substack{U\in\mathbb{R}^{d\times k}\\ U^\intercal U = I_k}} \sup_{\theta\in\mathbb{S}^{k-1}} \mu(|\theta^\intercal U^\intercal x|^q)^{1/q} && \text{(Lemma 6)} \\
&\le \sqrt{1\vee k/q} \sup_{\theta\in\mathbb{S}^{d-1}} \mu(|\theta^\intercal x|^q)^{1/q} && (\mathbb{S}^{k-1}\subset\mathbb{S}^{d-1}) \\
&\le \sqrt{1\vee k/q} && (\mu\in\mathcal{G}_q).
\end{aligned}$$

In the same way, we can extract this factor of $\sqrt{1\vee k/q}$ for the resiliency of μ w.r.t. $\overline{\mathsf{D}}_{p,k}^{1/p}$ to prove the needed analog of Lemma 12. Combining these results gives that $R(\underline{\mathsf{W}}_{p,k},\mathcal{G}_q,\epsilon) \le R(\overline{\mathsf{W}}_{p,k},\mathcal{G}_q,\epsilon) \lesssim \sqrt{1\vee k/q}\, R(\overline{\mathsf{W}}_p,\mathcal{G}_q,\epsilon) \asymp \sqrt{1\vee k/q}\,\epsilon^{1/p-1/q}$ for $0\le\epsilon\le 0.49$, as desired.

D.4.4 Proofs of auxiliary lemmas

Proof of Lemma 5. Let $\Theta \sim \mathrm{Unif}(\mathbb{S}^{d-1})$. When $d = 1$, we have $\mathbb{E}[|\Theta_1|^p] = 1$. Otherwise, we use that the probability density function of Θ_1 at $s \in [-1,1]$ is proportional to $(1-s^2)^{\frac{d-3}{2}}$ [52]. Equivalently, $(\Theta_1+1)/2 \sim \mathrm{Beta}\left(\frac{d-1}{2},\frac{d-1}{2}\right)$. We will first prove the desired statement for even integer $p = 2m$, where

$$\mathbb{E}\left[|\Theta_1|^{2m}\right] = \frac{(2m)!}{2^{2m}m!}\frac{\Gamma(d-1)\Gamma(\frac{d-1}{2}+m)}{\Gamma(\frac{d-1}{2})\Gamma(d-1+2m)}$$

(see, e.g., [38]). Simplifying, we obtain

$$\begin{aligned}
\mathbb{E}\left[|\Theta_1|^{2m}\right] &= \frac{\Gamma(2m+1)}{2^{2m}\Gamma(m+1)}\frac{\Gamma(d-1)\Gamma(\frac{d-1}{2}+m)}{\Gamma(\frac{d-1}{2})\Gamma(d-1+2m)} \\
&= \frac{\Gamma(d/2)\Gamma(m+1/2)}{\sqrt{\pi}2^{2m}\Gamma(m+d/2)}.
\end{aligned}$$

Employing Stirling's formula, we compute

$$\begin{aligned}
2^{2m}\,\mathbb{E}\left[|\Theta_1|^{2m}\right] &\asymp \frac{\Gamma(d/2)\Gamma(m+1/2)}{\Gamma(m+d/2)} \\
&\asymp \frac{(d/2)^{d/2-1/2}e^{-d/2}(m+1/2)^m e^{-m-1/2}}{(m+d/2)^{m+d/2-1/2}e^{-m-d/2}}
\end{aligned}$$

$$\asymp \frac{(d/2)^{d/2-1/2}(m+1/2)^m}{(m+d/2)^{m+d/2-1/2}}$$
$$= \left(\frac{d/2}{m+d/2}\right)^{\frac{d-1}{2}} \left(\frac{m+1/2}{m+d/2}\right)^m.$$

Consequently, we have

$$\mathbb{E}\left[|\Theta_1|^{2m}\right]^{1/2m} \asymp \left(1+\frac{m}{d/2}\right)^{-\frac{d-1}{4m}} \sqrt{\frac{m+1/2}{m+d/2}}$$
$$\asymp \sqrt{\frac{m+1/2}{m+d/2}}$$
$$\asymp 1 \wedge \sqrt{m/d},$$

as desired. When $p \geq 2$ is not an even integer, we use that $\mathbb{E}\left[|\Theta_1|^p\right]$ is monotonically increasing in p to obtain matching bounds by rounding p up and down to the nearest even integers. To obtain the needed lower bound when $p \in [1, 2)$, we derive

$$\mathbb{E}\left[|\Theta_1|\right] = \frac{4\Gamma(d-1)}{\Gamma\left(\frac{d-1}{2}\right)^2} \frac{\left(\frac{d-1}{2}\right)^{d-1}}{(d-1)^d},$$

using the formula for the mean absolute deviation of the beta distribution [23]. Applying Stirling's formula once more, we obtain

$$\mathbb{E}\left[|\Theta_1|\right] \asymp \frac{(d-1)^{d-3/2}\left(\frac{d-1}{2}\right)^{d-1}}{(\frac{d-1}{2})^{d-2}(d-1)^d} = \frac{(d-1)/2}{(d-1)^{3/2}} \asymp d^{-1/2},$$

as desired. □

Proof of Lemma 6. Taking $X \sim \mu$ and $\Theta \sim \mathrm{Unif}(\mathbb{S}^{d-1})$, we use rotational symmetry of the sphere to compute

$$\mathbb{E}\left[|\Theta^\intercal X|^p\right] = \mathbb{E}\left[|\Theta_1|^p\right] \mathbb{E}\left[\|X\|^p\right],$$

giving the first equality via Lemma 5. The inequality follows by comparing an average to a supremum, and the inequality is tight when these coincide, i.e. when μ is rotationally symmetric about 0. □

D.5 Proof of Proposition 2

The high-level structure of our proof follows a standard template for finite-sample robust mean and covariance estimation (see, e.g., [55, 61]). We first prove Proposition 2 under bounded support and then extend our result to the general setting. Throughout, we write $\mathbb{B}_r := \{x \in \mathbb{R}^d : \|x\| \leq r\}$ for the Euclidean ball of radius $r \geq 0$.

Bounded Support: For ease of presentation, we slightly extend our notion of resilience in a standard way. We say that $\mu \in \mathcal{P}(\mathbb{R}^d)$ is (ρ, ϵ)-resilient w.r.t. D about $\kappa \in \mathcal{P}(\mathbb{R}^d)$ if $\mathsf{D}(\nu, \kappa) \leq \rho$ for all $\nu \leq \frac{1}{1-\epsilon}\mu$. Namely, we will consider the resilience of an empirical measure $\kappa = \hat{\mu}_n$ about its population measure μ. If $\mathsf{D} = \mathsf{D}_{\mathcal{F}}$ is an IPM for symmetric $\mathcal{F} = -\mathcal{F}$, note that μ is (ρ, ϵ)-resilient w.r.t. D about κ if and only if $f_\sharp\mu$ is (ρ, ϵ)-resilient (in mean) about $f_\sharp\kappa$ for all $f \in \mathcal{F}$.

We first recall and derive some basic results for finite-sample resilience. The following lemma is a simplification of [55, Proposition 4], specified to the 1-dimensional case.

Lemma 13 (1-dimensional finite-sample resilience)**.** *Suppose that $\mu \in \mathcal{P}([-M, M])$ is (ρ, ϵ)-resilient in mean for $\epsilon \leq 0.999$. Then, with probability at least $1-\delta$, the empirical distribution $\hat{\mu}_n$ is (ρ', ϵ)-resilient in mean about μ with $\rho' = O\left(\rho\left(1+\sqrt{\frac{\log(1/\delta)}{\epsilon^2 n}}\right) + \frac{M\log(1/\delta)}{n}\right)$.*

The result is stated in [55] for $\epsilon < 1/2$, but the proof only uses that ϵ is bounded away from 1. We then extend this result to IPMs over uniformly bounded function classes.

Proposition 13 (Finite-sample resilience w.r.t. IPMs)**.** *Let $\mathsf{D}_{\mathcal{F}}$ be the IPM induced by a function class $\mathcal{F} = -\mathcal{F}$ on $\mathbb{R}^d$ with $\|f\|_\infty \le M$ for $f \in \mathcal{F}$, and fix any finite subset $\mathcal{H} \subseteq \mathcal{F}$ such that $\mathsf{D}_{\mathcal{F}}(\mu,\nu) \le \mathsf{D}_{\mathcal{H}}(\mu,\nu) + \rho$. Then if $\mu \in \mathcal{P}(\mathbb{R}^d)$ is (ρ,ϵ)-resilient w.r.t. $\mathsf{D}_{\mathcal{F}}$ for $\epsilon \le 0.999$, we have that $\hat{\mu}_n$ is (ρ',ϵ)-resilient w.r.t. $\mathsf{D}_{\mathcal{F}}$ about μ with probability $1-\delta$, where $\rho' = O\left(\rho + \frac{\rho}{\epsilon}\sqrt{\frac{\log(|\mathcal{H}|/\delta)}{n}} + \frac{M\log(|\mathcal{H}|/\delta)}{n}\right)$.*

Proof. For any $\nu_n \le (1-\epsilon)\hat{\mu}_n$, we have

$$\mathsf{D}_{\mathcal{F}}(\nu_n,\mu) \le \max_{f\in\mathcal{H}}(\nu_n - \mu)(f) + \rho.$$

Now, fixing any $f \in \mathcal{H}$, resilience of μ w.r.t. $\mathsf{D}_{\mathcal{F}}$ requires that $f_\sharp\mu$ is (ρ,ϵ)-resilient. Noting that $f_\sharp\nu_n \le \frac{1}{1-\epsilon} f_\sharp\hat{\mu}_n$, Lemma 13 gives that

$$|(\nu_n - \mu)(f)| = |(f_\sharp\nu_n)x - (f_\sharp\mu)x| \le O\left(\rho + \frac{\rho}{\epsilon}\sqrt{\frac{\log(|\mathcal{H}|/\delta)}{n}} + \frac{M\log(|\mathcal{H}|/\delta)}{n}\right)$$

with probability at least $1 - \delta/|\mathcal{H}|$. A union bound over $f \in \mathcal{H}$ gives the desired result. □

To apply this result, we approximate $\overline{\mathsf{D}}_p$ with an IPM over a finite function class.

Lemma 14 (Approximating $\overline{\mathsf{D}}_p$)**.** *For each $\gamma > 0$, there exists a net $\mathcal{N} \subset \mathbb{S}^{d-1}$ of size $(10R^p p/\gamma)^d$ such that for all $\mu,\nu \in \mathcal{P}(\mathbb{B}_R)$, we have*

$$\overline{\mathsf{D}}_p(\mu,\nu) = \sup_{\theta\in\mathbb{S}^{d-1}} |(\mu-\nu)(|\theta^\intercal x|^p)| \le \max_{\theta\in\mathcal{N}} |(\mu-\nu)(|\theta^\intercal x|^p)| + \gamma.$$

Proof. Let $\mathcal{N}$ be a $\gamma(2R^p p)^{-1}$-covering for $\mathbb{S}^{d-1}$ in ℓ_2 with $|\mathcal{N}| \le (10R^p p\gamma^{-1})^d$, the existence of which is guaranteed by Lemma 3 (i). Then, taking θ to be a direction achieving the LHS supremum, and $\tilde{\theta} \in \mathcal{N}$ to be its nearest neighbor in $\mathcal{N}$, we have

$$\begin{aligned}
|(\mu-\nu)(|\theta^\intercal x|^p)| &\le |(\mu-\nu)(|\tilde{\theta}^\intercal x|^p)| + 2\sup_{\kappa\in\mathcal{P}(\mathbb{B}_R)} \kappa||\theta^\intercal x|^p - |\tilde{\theta}^\intercal x|^p| \\
&\le |(\mu-\nu)(|\tilde{\theta}^\intercal x|^p)| + 2R^p p\|\theta - \tilde{\theta}\| \\
&\le |(\mu-\nu)(|\tilde{\theta}^\intercal x|^p)| + \gamma,
\end{aligned}$$

where the second inequality follows by Lipschitzness. Supremizing over θ gives the lemma. □

Combining, we obtain finite-sample resilience w.r.t. our distances of interest. Slightly abusing notation for brevity, we write $\underline{\rho}(\tau) = \underline{\rho}(\tau,p,d,q) = \sqrt{(1\vee d/q)(1\wedge p/d)}\tau^{1/p-1/q}$ and $\overline{\rho}(\tau) = \overline{\rho}(\tau,p,q) = \tau^{1/p-1/q}$ for our resilience bounds for the class $\mathcal{G}_q$ w.r.t. $\underline{\mathsf{W}}_p$ and $\overline{\mathsf{W}}_p$.

Lemma 15 (Finite-sample resilience under bounded support)**.** *Let $0 \le \epsilon \le 0.999$ and $q > p$. If $\mu \in \mathcal{G}_q$ with $\mathrm{diam}(\mathrm{spt}(\mu)) \le R/2$ and $n = \Omega\big((R^p + \epsilon^{-2})(d\log(R/\epsilon) + \log(1/\delta))\big)$, then $\hat{\mu}_n$ is $(O(\underline{\rho}(\epsilon)),\epsilon)$-resilient w.r.t. $\underline{\mathsf{W}}_p$ and $(O(\overline{\rho}(\epsilon)),\epsilon)$-resilient w.r.t. $\overline{\mathsf{W}}_p$ with probability $1-\delta$.*

Proof. Assume without loss of generality that $\mu x = 0$ and $\mu \in \mathcal{G}_q \cap \mathcal{P}(\mathbb{B}_R)$. By Lemma 12 (combined with Lemma 10 if $\epsilon \ge 1/2$), we have that μ is $(\underline{\rho}(\epsilon),\epsilon)$-resilient w.r.t. $\underline{\mathsf{D}}_p^{1/p}$ and $(\overline{\rho}(\epsilon),\epsilon)$-resilient w.r.t. $\overline{\mathsf{D}}_p^{1/p}$. For $\overline{\mathsf{D}}_p$, observe that for $\|x\| \le R$ and $\theta \in \mathbb{S}^{d-1}$, we have $|\theta^\intercal x|^p \le R^p$. Thus, applying Proposition 13 with $\mathcal{F} = \overline{\mathcal{F}}_p$, $M = R^p$, and $\mathcal{H}$ induced by the net from Lemma 14 with $\gamma = \overline{\rho}(\epsilon)^p$ gives that $\hat{\mu}_n$ is $(O(2^p\overline{\rho}(\epsilon)^p),\epsilon)$-resilient about μ w.r.t. $\overline{\mathsf{D}}_p$ with probability at least $1-\delta/2$ whenever

$$\begin{aligned}
n &\ge (\overline{\rho}(\epsilon)^p R^p + \epsilon^{-2})\log(2|\mathcal{H}|/\delta)/2^p \\
&= (\overline{\rho}(\epsilon)^p R^p + \epsilon^{-2})\log((20R^p p/\overline{\rho}(\epsilon)^p)^d/\delta)/2^p.
\end{aligned}$$

Plugging in our value for $\overline{\rho}(\epsilon)$ and applying some crude bounds shows the stated sample complexity of $n = \Omega\big((R^p + \epsilon^{-2})(d\log(R/\epsilon) + \log(1/\delta))\big)$ suffices. Since the resilience bound is centered about

μ and $\mu \in \mathcal{G}_q$, we deduce that $\overline{\mathsf{D}}_p(\hat{\mu}_n, \delta_0) \le \overline{\mathsf{D}}_p(\hat{\mu}_n, \mu) + \overline{\mathsf{D}}_p(\mu, \delta_0) \le O(2^p \overline{\rho}(\epsilon)^p) + 1 = O(2^p)$. Thus, by Proposition 11, we have that $\hat{\mu}_n$ is $(O(\overline{\rho}(\epsilon)), \epsilon)$-resilient w.r.t. $\overline{\mathsf{W}}_p$. An analogous argument shows that the same sample complexity suffices for $\underline{\mathsf{W}}_p$ (of course, far fewer samples are actually needed, but we shall not focus on this distinction). Applying a union bound gives that both resilience guarantees hold with probability $1 - \delta$. □

Finally, we use finite-sample resiliency to bound finite-sample robust estimation risk.

Proposition 14. *Let $0 \le \epsilon \le 0.499$ and $q > p$ and $\mathsf{D} \in \{\underline{\mathsf{W}}_p, \overline{\mathsf{W}}_p\}$. Then there exists an estimation procedure $\mathcal{A}$ with the following guarantee: for any $\mu \in \mathcal{G}_q$ with $\mathrm{diam}(\mathrm{spt}(\mu)) \le R/2$ and $n = \Omega\big((R^p + \epsilon^{-2})(d \log(R/\epsilon) + \log(1/\delta))\big)$, after observing any random distribution $\tilde{\mu}_n$ such that $\|\tilde{\mu}_n - \hat{\mu}_n\|_{\mathrm{TV}} \le \epsilon$ almost surely, $\mathcal{A}$ produces ν such that $\mathsf{D}(\nu, \mu) \lesssim R(\mathsf{D}, \mathcal{G}_q, \epsilon) + \mathsf{D}(\hat{\mu}_n, \mu)$ with probability $1 - \delta$.*

Proof. For $\mathsf{D} = \overline{\mathsf{W}}_p$, define

$$\Pi_{\epsilon,\rho}(\tilde{\mu}_n) = \big\{\kappa \in \mathcal{P}(\mathbb{R}^d) : \|\kappa - \tilde{\mu}_n\|_{\mathrm{TV}} \le \epsilon \text{ and } \kappa \text{ is } (\rho, 2\epsilon)\text{-resilient w.r.t. } \overline{\mathsf{W}}_p\big\}$$

Write $\rho_\star = \inf\{\rho \ge 0 : \Pi_{\epsilon,\rho}(\tilde{\mu}_n) \ne \emptyset\}$ for the smallest resilience parameter such that this set is non-empty. Since $2\epsilon \le 0.999$, we know by Lemma 15 that $\hat{\mu}_n$ is $(O(\overline{\rho}(2\epsilon), 2\epsilon)$-resilient w.r.t. $\overline{\mathsf{W}}_p$ (for an appropriate choice of constant in the sample complexity) with probability $1 - \delta/2$. Noting that $\|\tilde{\mu}_n - \hat{\mu}_n\|_{\mathrm{TV}} \le \epsilon$, we have $\rho_\star \lesssim \overline{\rho}(2\epsilon) \lesssim \overline{\rho}(\epsilon)$ with probability $1 - \delta/2$.

Now consider any algorithm which returns $\nu \in \Pi_{\epsilon, 2\rho_\star}(\tilde{\mu}_n)$. Then we have $\|\nu - \hat{\mu}_n\|_{\mathrm{TV}} \le \|\nu - \tilde{\mu}_n\|_{\mathrm{TV}} + \|\tilde{\mu}_n - \hat{\mu}_n\|_{\mathrm{TV}} \le 2\epsilon$. By considering their midpoint $\kappa = \frac{1}{1 - \|\nu - \hat{\mu}_n\|_{\mathrm{TV}}} \nu \wedge \hat{\mu}_n$ and applying $(O(\overline{\rho}(2\epsilon)), 2\epsilon)$-resilience of ν and $\hat{\mu}_n$ w.r.t. $\overline{\mathsf{W}}_p$, we deduce that $\overline{\mathsf{W}}_p(\nu, \hat{\mu}_n) \le \overline{\mathsf{W}}_p(\nu, \kappa) + \overline{\mathsf{W}}_p(\kappa, \hat{\mu}_n) \lesssim \overline{\rho}(\epsilon) \lesssim R(\overline{\mathsf{W}}_p, \mathcal{G}_q, \epsilon)$ with probability $1 - \delta$. By the triangle inequality for $\overline{\mathsf{W}}_p$, we thus have $\overline{\mathsf{W}}_p(\nu, \mu) \lesssim R(\overline{\mathsf{W}}_p, \mathcal{G}_q, \epsilon) + \overline{\mathsf{W}}_p(\hat{\mu}_n, \mu)$ with probability $1 - \delta$. An analogous argument gives the corresponding result for $\underline{\mathsf{W}}_p$. □

Reduction to bounded support: To prove Proposition 2, we provide a reduction from the general case to that of bounded support via Markov's inequality and a coupling argument.

Lemma 16 (High probability norm bound)**.** *If $X \sim \mu \in \mathcal{G}_q$, there exists $R \lesssim \delta^{-1/q}\sqrt{1 \vee d/q} \le \delta^{-1/q}\sqrt{d}$ such that $\|X - \mu x\| \le R$ with probability at least $1 - \delta$.*

Proof. Assume without loss of generality that $\mu x = 0$. We compute

$$\mu(\|x\|^q)^{1/q} \lesssim (1 \vee d/q)^{q/2} \sup_{\theta \in \mathbb{S}^{d-1}} \mu(|\theta^\intercal x|^q)^{1/q} \le (1 \vee d/q)^{q/2},$$

where the first inequality uses Lemma 6 and the second uses $\mu \in \mathcal{G}_q$. Markov's inequality then gives the claim. □

Lemma 17 (Switch of base measure)**.** *Fix $\mu \in \mathcal{P}(\mathbb{R}^d)$ and $A \subseteq \mathbb{R}^d$ with $\mu(A) \ge 1 - \epsilon$. Write μ_A for the distribution of $X \sim \mu$ conditioned on $X \in A$. Consider any random measure $\tilde{\mu}_n$ such that $\|\tilde{\mu}_n - \hat{\mu}_n\| \le \epsilon'$ almost surely, where $n \ge 3\log(1/\delta)/\epsilon$. Then there exists a coupling of $(\tilde{\mu}_n, \hat{\mu}_n)$ and $(\widehat{\mu_A})_n$ such that $\|\tilde{\mu}_n - (\widehat{\mu_A})_n\|_{\mathrm{TV}} \le 2\epsilon + \epsilon'$ with probability at least $1 - \delta$.*

Proof. Given n i.i.d. samples $X_1, \ldots, X_n$ from μ, Lemma 17 and a Chernoff bound give that at least $(1 - 2\epsilon)n$ of them satisfy $X_i \in A$, with probability at least $1 - \delta$. Define the coupled set of samples $Y_1, \ldots, Y_n$ by $Y_i = X_i$ if $X_i \in A$ and $Y_i \sim \mu_A$ i.i.d. otherwise, and choose $(\widehat{\mu_A})_n$ as their empirical measure (by design, the marginal distribution of $Y_1, \ldots, Y_n$ coincides with n samples from μ_A). Under this coupling, we then have

$$\|\tilde{\mu}_n - (\widehat{\mu_A})_n\|_{\mathrm{TV}} \le \|\tilde{\mu}_n - \hat{\mu}_n\|_{\mathrm{TV}} + \|\hat{\mu}_n - (\widehat{\mu_A})_n\|_{\mathrm{TV}} \le 2\epsilon + \epsilon'$$

with probability at least $1 - \delta$. □

Proof of Proposition 2. By Lemma 16, we have that for $X \sim \mu$, $\|X - \mu x\| \le R \asymp \sqrt{d/\epsilon}$ with probability at least $1 - \epsilon/400$. Letting A denote the ball of radius R around μx and applying Lemma 17 with failure probability 0.001, we can view $\tilde{\mu}_n$ as being a $\frac{201}{200}\epsilon$-corrupted version of n i.i.d. samples from the conditional distribution μ_R, with probability at least 0.999. Thus, applying the procedure from Proposition 14 with $R \asymp \sqrt{d/\epsilon}$, confidence probability 0.999 and corruption fraction $\frac{201}{200}\epsilon < 0.499$, we obtain ν with $\overline{\mathsf{W}}_p(\nu, \mu_R) \lesssim R(\mathsf{D}, \mathcal{G}_q, \epsilon) + \mathsf{D}((\widehat{\mu_R})_n, \mu_R)$ with unconditional probability 0.998. By resilience of μ and the fact that $\|\mu - \mu_R\|_{\mathrm{TV}} \le \epsilon/400$, the same recovery guarantees hold with base measure μ. Finally, we bound $\mathsf{D}((\widehat{\mu_R})_n, \mu_R)$ by its expectation via Markov's inequality to obtain $\overline{\mathsf{W}}_p(\nu, \mu) \lesssim R(\mathsf{D}, \mathcal{G}_q, \epsilon) + \mathbb{E}[\mathsf{D}((\widehat{\mu_R})_n, \mu_R)]$ with probability 0.99. □

D.6 Proof of Proposition 3

For any $\mu, \nu \in \mathcal{P}_1(\mathbb{R}^d)$, we have $\overline{\mathsf{W}}_p(\mu, \nu) \ge \|\mu x - \nu x\|$ (seen by taking θ in the direction of $\mu x - \nu x$). Hence, if $\overline{\mathsf{W}}_p(\mu, \nu) \le \rho$ for all $\nu \le \frac{1}{1-\epsilon}\mu$, then μ is (ρ, ϵ)-resilient in mean. (This direction holds for all $p \ge 1$). For the other direction, we mirror the proof of Theorem 2, first establishing a simple lemma.

Lemma 18. *Fix $X \sim \mu \in \mathcal{P}_1(\mathbb{R})$ and define the quantiles $\tau_\epsilon = \sup\{t \in \mathbb{R} : \Pr(X \ge t) \ge \epsilon\}$ and $\tilde{\tau}_\epsilon = \sup\{t \in \mathbb{R} : \Pr(|X| \ge t) \ge \epsilon\}$. Then, we have*

$$\mathbb{E}[|X| \mid |X| \ge \tilde{\tau}_\epsilon] \le 4\,\mathbb{E}[X \mid X \ge \tau_\epsilon] \vee \mathbb{E}[-X \mid X \le \tau_{1-\epsilon}]$$

Simply put, if $|X|$ has large tails, then one of X or $-X$ must have a large tail.

Proof. Writing $X_+ = X \vee 0$ and $X_- = -X \vee 0$, we bound

$$\begin{aligned}
&\mathbb{E}\big[|X| \,\big|\, |X| \ge \tilde{\tau}_\epsilon\big] = \mathbb{E}\big[X_+ \,\big|\, |X| \ge \tilde{\tau}_\epsilon\big] + \mathbb{E}\big[X_- \,\big|\, |X| \ge \tilde{\tau}_\epsilon\big] \\
&\le \mathbb{E}[X_+ | X \ge \tau_\epsilon] + \mathbb{E}[X_- | X \le \tau_{1-\epsilon}] \\
&\le \mathbb{E}[X - (\tau_\epsilon \wedge 0) \mid X \ge \tau_\epsilon] + \mathbb{E}[-X + (\tau_{1-\epsilon} \vee 0) \mid X \le \tau_{1-\epsilon}] \\
&= \mathbb{E}[X \mid X \ge \tau_\epsilon] + \mathbb{E}[-X \mid X \le \tau_{1-\epsilon}] - (\tau_\epsilon \wedge 0) + (\tau_{1-\epsilon} \vee 0) \\
&\le \mathbb{E}[X \mid X \ge \tau_\epsilon] + \mathbb{E}[-X \mid X \le \tau_{1-\epsilon}] + (-\tau_{1-\epsilon} \vee 0) + (\tau_\epsilon \vee 0) \\
&\le \mathbb{E}[X \mid X \ge \tau_\epsilon] + \mathbb{E}[-X \mid X \le \tau_{1-\epsilon}] + (\mathbb{E}[-X \mid X \le \tau_{1-\epsilon}] \vee 0) + (\mathbb{E}[X \mid X \ge \tau_\epsilon] \vee 0).
\end{aligned}$$

Now, it is easy to check that each summand is bounded by $\mathbb{E}[X \mid X \ge \tau_\epsilon] \vee \mathbb{E}[-X \mid X \le \tau_{1-\epsilon}]$ (since this maximum is non-negative), giving the lemma. □

Continuing, we take $\mu \in \mathcal{P}_1(\mathbb{R}^d)$ which is (ρ, ϵ)-mean-resilient and assume without loss of generality that $\mu x = 0$. For all $\nu \le \frac{1}{1-\epsilon}\mu$, we write $\mu = (1-\epsilon)\nu + \epsilon\alpha$ for $\alpha \in \mathcal{P}(\mathbb{R}^d)$ and bound

$$\begin{aligned}
\overline{\mathsf{W}}_p(\mu, \nu) &= \overline{\mathsf{W}}_p((1-\epsilon)\nu + \epsilon\alpha, \nu) \\
&\le \epsilon^{1/p}\overline{\mathsf{W}}_p(\alpha, \nu) && \text{(Fact 2)} \\
&\le \epsilon^{1/p}(\overline{\mathsf{W}}_p(\alpha, \delta_0) + \overline{\mathsf{W}}_p(\nu, \delta_0)) && \text{(triangle inequality)} \\
&\le 2\epsilon^{1/p} \sup_{\kappa \le \frac{1}{(1-\epsilon)\wedge\epsilon}\mu} \overline{\mathsf{W}}_p(\kappa, \delta_0) \\
&= 2\epsilon \sup_{\kappa \le \frac{1}{(1-\epsilon)\wedge\epsilon}\mu} \sup_{\theta \in \mathbb{S}^{d-1}} \mathbb{E}_\kappa\left[|\theta^\mathsf{T} X|^p\right]^{1/p} \\
&= 2\epsilon \sup_{\theta \in \mathbb{S}^{d-1}} \mathbb{E}_\kappa\left[|\theta^\mathsf{T} X|^p \,\middle|\, |\theta^\mathsf{T} X| \ge \tilde{\tau}_{\epsilon\wedge(1-\epsilon)}(\theta)\right]^{1/p}.
\end{aligned}$$

where $\tilde{\tau}_\epsilon(\theta) = \sup\{t \in \mathbb{R} : \mathbb{P}(|\theta^\mathsf{T} X| \ge t) \ge \epsilon\}$ for $X \sim \mu$. (Technically, the final inequality may fail if μ has a point mass at $\tilde{\tau}_{\epsilon\wedge(1-\epsilon)}(\theta)$; in this case, assume that ties are broken with independent randomness so that the conditioned event has probability $\epsilon \wedge (1-\epsilon)$). From now on, we will use that $p = 1$. Writing $\tau_\epsilon(\theta) = \sup\{t \in \mathbb{R} : \mathbb{P}(\theta^\mathsf{T} X \ge t) \ge \epsilon\}$ and breaking ties in the same way, we apply Lemma 18 to bound

$$\overline{\mathsf{W}}_1(\mu, \nu) \le 8\epsilon \sup_{\theta \in \mathbb{S}^{d-1}} \mathbb{E}_\mu[\theta^\mathsf{T} X \mid \theta^\mathsf{T} X \ge \tau_{\epsilon\wedge(1-\epsilon)}(\theta)]$$

$$
\begin{aligned}
&= 8\epsilon \sup_{\theta \in \mathbb{S}^{d-1}} \theta^\intercal \mathbb{E}_\mu[X \mid \theta^\intercal X \geq \tau_{\epsilon \wedge (1-\epsilon)}(\theta)] \\
&= 8\epsilon \left\| \mathbb{E}_\mu[X \mid \theta^\intercal X \geq \tau_{\epsilon \wedge (1-\epsilon)}(\theta)] \right\| \\
&= 8\epsilon \left\| \mathbb{E}_\mu X - \mathbb{E}_\mu[X \mid \theta^\intercal X \geq \tau_{\epsilon \wedge (1-\epsilon)}(\theta)] \right\|.
\end{aligned}
$$

Now, if $\epsilon \geq 1/2$, we can use resilience of μ to bound

$$
\overline{\mathsf{W}}_1(\mu, \nu) \leq 8\epsilon\rho \leq 8\rho.
$$

Otherwise, writing E for the event that $\theta^\intercal X \geq \tau_\epsilon(\theta)$, we have

$$
\begin{aligned}
\overline{\mathsf{W}}_1(\mu, \nu) &\leq 8\epsilon \left\| \epsilon \mathbb{E}_\mu[X|E] + (1-\epsilon)\mathbb{E}_\mu[X|E^c] - \mathbb{E}_\mu[X|E] \right\| \\
&= 8\epsilon(1-\epsilon) \left\| \mathbb{E}_\mu[X|E^c] - \mathbb{E}_\mu[X|E] \right\| \\
&= 8\epsilon(1-\epsilon) \left\| \mathbb{E}_\mu[X|E^c] - \epsilon^{-1}(\mathbb{E}_\mu[X] - (1-\epsilon)\mathbb{E}_\mu[X|E^c]) \right\| \\
&= 8(1-\epsilon) \left\| \mathbb{E}_\mu[X] - \mathbb{E}_\mu[X|E^c] \right\| \\
&\leq 8(1-\epsilon)\rho \leq 8\rho.
\end{aligned}
$$

Hence, μ is $(8\rho, \epsilon)$-resilient w.r.t. $\overline{\mathsf{W}}_1$.

Immediately, this allows $\overline{\mathsf{W}}_1$ to inherit a multitude of (population-limit and finite-sample) risk bounds from the robust mean estimation literature. See [54] for a detailed survey of robust statistics results based on resiliency. For example, $\mu \in \mathcal{G}_q$ is known to be $(O(\epsilon^{1-1/q}), \epsilon)$-mean-resilient, immediately implying Theorem 2 for $\overline{\mathsf{W}}_1$.

D.7 Proof of Proposition 4

When $q = 2$, we mirror the approach of Proposition 2 but perform projection onto the space of distributions with bounded covariance, instead of onto the space of resilient distributions. We require the following standard result (see, e.g., Lemma A.18 of [19]), establishing finite-sample covariance bounds under bounded support.

Lemma 19. *Let $\mu \in \mathcal{P}(\mathbb{R}^d)$ with $\|\Sigma_\mu\|_{\mathrm{op}} \leq \sigma^2$ and $\mathrm{diam}(\mathrm{spt}(\mu)) \leq R$. Then the empirical distribution $\hat{\mu}_n$ satisfies $\|\Sigma_{\hat{\mu}_n}\|_{\mathrm{op}} \lesssim \sigma^2$ with probability at least 0.999 for $n \gtrsim R^2 \log(d)$.*

Importantly, there are efficient filtering algorithms for projecting onto the set of distributions with bounded covariance (see, e.g., Theorem 3.1 [24]).

Lemma 20 (Spectral reweighting)**.** *Let $x_1, \ldots, x_n \in \mathbb{R}^d$ and $0 < \epsilon \leq 1/10$. Suppose the discrete measure $\mu_n = \frac{1}{n}\sum_{i=1}^n \delta_{x_i}$ admits an ϵ-deletion $\nu_n \leq \frac{1}{1-\epsilon}\mu_n$ such that $\|\Sigma_{\mu_n}\|_{\mathrm{op}} \leq \sigma^2$. Then, given $\{x_i\}_{i=1}^n$ and ϵ, there is an algorithm which finds $\nu \leq \frac{1}{1-3\epsilon}\mu_n$ such that $\|\Sigma_\nu\|_{\mathrm{op}} \lesssim \sigma^2$ with probability 0.999, in time $\widetilde{O}(nd^2)$.*

Combining, we prove the proposition. We remark that sample complexity is dominated by empirical convergence under $\mathsf{D} \in \{\underline{\mathsf{W}}_p, \overline{\mathsf{W}}_p\}$ of the truncated version of a distribution with bounded second moments. This can be improved significantly in many cases of interest, for example under log-concavity of the clean distribution.

Proof of Proposition 4. First, we consider the case of bounded support, when $\mathrm{diam}(\mathrm{spt}(\mu)) \leq R$, and with contamination fraction $\epsilon \in [0, 1/10]$. We mirror the argument of Proposition 14, but project onto the set of distributions with bounded covariance using spectral reweighting. Write $\tilde{\mu}_n$ for the empirical distribution of the ϵ-contaminated samples and $\hat{\mu}_n$ for that of the clean samples, with $\hat{\mu}_n \leq \frac{1}{1-\epsilon}\tilde{\mu}_n$. Combining Lemmas 19 and 20, we find that $\|\Sigma_{\hat{\mu}_n}\|_{\mathrm{op}} \lesssim \sigma^2$ and that the spectral reweighting algorithm returns $\nu \leq \frac{1}{1-3\epsilon}\hat{\mu}_n$ with $\|\Sigma_\nu\|_{\mathrm{op}} \lesssim \sigma^2$ in time $\widetilde{O}(nd^2)$, all with probability 0.998. By resilience of the class $\mathcal{G}_2(\sigma)$ w.r.t. D and Markov's inequality, we have

$$
\mathsf{D}(\nu, \mu) \leq \mathsf{D}(\nu, \hat{\mu}_n) + \mathsf{D}(\hat{\mu}_n, \mu) \lesssim R(\mathsf{D}, \mathcal{G}_q(\sigma), \epsilon) + \mathbb{E}[\mathsf{D}(\hat{\mu}_n, \mu)]
$$

with probability 0.995. For the unbounded case, we apply Lemma 16 and Lemma 17 as in the proof of Proposition 2 to reduce to $R \asymp \sqrt{d/\epsilon}$ and obtain the desired error bound with probability at least 0.99, so long as $0 < \epsilon \leq 1/12$ (any constant separated from 1/10 will do). $\square$

D.8 Proof of Lemma 1

We start by showing that the Lipschitz constant of w_p is upper bounded by $L^p_{\mu,\nu}$. Fix $\theta_1, \theta_2 \in \mathbb{S}^{d-1}$ and observe that

$$\begin{aligned}\big|w_p(\theta_1) - w_p(\theta_2)\big| &= \big|\mathsf{W}_p\big(\mathfrak{p}^{\theta_1}_\sharp \mu, \mathfrak{p}^{\theta_1}_\sharp \nu\big) - \mathsf{W}_p\big(\mathfrak{p}^{\theta_2}_\sharp \mu, \mathfrak{p}^{\theta_2}_\sharp \nu\big)\big| \\ &\le \mathsf{W}_p\big(\mathfrak{p}^{\theta_1}_\sharp \mu, \mathfrak{p}^{\theta_2}_\sharp \mu\big) + \mathsf{W}_p\big(\mathfrak{p}^{\theta_1}_\sharp \nu, \mathfrak{p}^{\theta_2}_\sharp \nu\big) \\ &\le \|\theta_1 - \theta_2\| \sup_{\theta \in \mathbb{S}^{d-1}} \Big(\big(\mu|\theta^\intercal x|^p\big)^{1/p} + \big(\nu|\theta^\intercal x|^p\big)^{1/p}\Big),\end{aligned}$$

where the last step uses the optimal transportation cost formulation of W_p. The RHS above is $L^p_{\mu,\nu}$ from the lemma, which concludes the proof of the first statement.

Next, we bound the Lipschitz constant of w^p_p. For $\theta_1, \theta_2 \in \mathbb{S}^{d-1}$ and $i = 1, 2$, let (X_i, Y_i) be a coupling of μ and ν so that $(\theta_i^\intercal X_i, \theta_i^\intercal Y_i)$ is optimal for $\mathsf{W}_p\big(\mathfrak{p}^{\theta_i}_\sharp \mu, \mathfrak{p}^{\theta_i}_\sharp \nu\big)$. These couplings are constructed as follows. For $i = 1, 2$, let (U_i, V_i) be an optimal couplings for $\mathsf{W}_p\big(\mathfrak{p}^{\theta_i}_\sharp \mu, \mathfrak{p}^{\theta_i}_\sharp \nu\big)$. Take $\mathrm{P}_i \in \mathbb{R}^{d\times d}$ as a unitary matrix whose first row is θ_i, and let $\mathrm{P}_{i,-1} \in \mathbb{R}^{(d-1)\times d}$ denote the matrix obtained by deleting the first row of P_i. Given $u_1, u_2 \in \mathbb{R}$ generate the random variables $W_i(u_i) \sim \mathcal{L}\big(P_{i,-1}X \big| \theta_i^\intercal X = u_i\big)$, for $i = 1, 2$, where $X \sim \mu$ and $\mathcal{L}(\cdot)$ designates the probability law of a random variable. Setting $\bar{U}_i := \big(U_i, W_i(U_i)\big)$ for $i = 1, 2$, observe that $\bar{U}_i \sim \mathcal{L}\big(\mathrm{P}_i X\big)$ and further that $X_i := \mathrm{P}_i^\intercal \bar{U}_i \sim \mu$. Constructing $\bar{V}_i$, for $i = 1, 2$, in an analogous fashion but with ν in place of μ, and defining Y_i similarly to X_i above, we obtain the desired (X_i, Y_i) couplings.

Then by optimality of the couplings, we have

$$\begin{aligned}w^p_p(\theta_1) - w^p_p(\theta_2) &\le \mathbb{E}\big[\big|\theta_1^\intercal(X_2 - Y_2)\big|^p - \big|\theta_2^\intercal(X_2 - Y_2)\big|^p\big], \\ w^p_p(\theta_2) - w^p_p(\theta_1) &\le \mathbb{E}\big[\big|\theta_2^\intercal(X_1 - Y_1)\big|^p - \big|\theta_1^\intercal(X_1 - Y_1)\big|^p\big].\end{aligned}$$

Combining these bounds, we obtain

$$\begin{aligned}|w^p_p(\theta_1) - w^p_p(\theta_2)| &\le p\|\theta_1 - \theta_2\| \, \mathbb{E}\left[\max_{i=1,2}\left|\frac{(\theta_1 - \theta_2)^\intercal(X_i - Y_i)}{\|\theta_1 - \theta_2\|}\right| \cdot \max_{i,j=1,2}\big|\theta_i^\intercal(X_j - Y_j)\big|^{p-1}\right] \\ &\le p\|\theta_1 - \theta_2\| \, \mathbb{E}\left[\max_{\substack{i=1,2\\ j=1,2,3}}\big|\theta_j'(X_i - Y_i)\big|^p\right] \\ &\le 3p2^p\|\theta_1 - \theta_2\| \sup_{\theta\in\mathbb{S}^{d-1}} \mathbb{E}\big[\big|\theta^\intercal X_1|^p + |\theta^\intercal Y_1|^p\big],\end{aligned}$$

where for the second inequality we have defined $\theta_3 := \frac{\theta_1 - \theta_2}{\|\theta_1 - \theta_2\|}$. This concludes the proof.

Remark 12 (Alternative Lipschitz constants)**.** *The Lipschitz constant for w^p_p can be alternatively derived as*

$$\begin{aligned}|w^p_p(\theta_1) - w^p_p(\theta_2)| &\le p\|\theta_1 - \theta_2\| \, \mathbb{E}\left[\max_{i=1,2}\left|\frac{(\theta_1 - \theta_2)^\intercal(X_i - Y_i)}{\|\theta_1 - \theta_2\|}\right| \cdot \max_{i,j=1,2}\big|\theta_i^\intercal(X_j - Y_j)\big|^{p-1}\right] \\ &\le p\|\theta_1 - \theta_2\| \, \mathbb{E}\left[\max_{\substack{i=1,2\\ j=1,2,3}}\big|\theta_j'(X_i - Y_i)\big|^p\right] \\ &\lesssim_p \|\theta_1 - \theta_2\| \left(\|\mu x - \nu x\| + \sup_{\theta\in\mathbb{S}^{d-1}} \mathbb{E}\big[\big|\theta^\intercal(X_1 - \mu x)\big|^p + \big|\theta^\intercal(Y_1 - \nu x)\big|^p\big]\right),\end{aligned}$$

where the terms corresponding to mean difference and covariance are separated.

D.9 Proof of Proposition 5

We decompose the error by introducing the Monte Carlo average for the population projected distances:

$$\mathbb{E}\Big[\Big|\underline{\widehat{\mathsf{W}}}^p_{\mathsf{MC}} - \underline{\mathsf{W}}^p_p(\mu, \nu)\Big|\Big]$$

$$\leq \mathbb{E}\left[\left|\widehat{\underline{\mathsf{W}}}_{\mathsf{MC}}^p - \frac{1}{m}\sum_{i=1}^m \mathsf{W}_p^p\big(\mathfrak{p}_\sharp^{\Theta_i}\mu, \mathfrak{p}_\sharp^{\Theta_i}\nu\big)\right|\right] + \mathbb{E}\left[\left|\frac{1}{m}\sum_{i=1}^m \mathsf{W}_p^p\big(\mathfrak{p}_\sharp^{\Theta_i}\mu, \mathfrak{p}_\sharp^{\Theta_i}\nu\big) - \underline{\mathsf{W}}_p^p(\mu,\nu)\right|\right]\Bigg\}. \tag{15}$$

For the first term, using the fact that $\Theta_1, \ldots, \Theta_n$ are i.i.d., we have

$$\mathbb{E}\left[\left|\widehat{\underline{\mathsf{W}}}_{\mathsf{MC}}^p - \frac{1}{m}\sum_{i=1}^m \mathsf{W}_p^p\big(\mathfrak{p}_\sharp^{\Theta_i}\mu, \mathfrak{p}_\sharp^{\Theta_i}\nu\big)\right|\right] \leq \mathbb{E}\Bigg\{\mathbb{E}\Bigg[\Big|\mathsf{W}_p^p\big(\mathfrak{p}_\sharp^{\Theta}\hat{\mu}_n, \mathfrak{p}_\sharp^{\Theta}\hat{\nu}_n\big) - \mathsf{W}_p^p\big(\mathfrak{p}_\sharp^{\Theta}\mu, \mathfrak{p}_\sharp^{\Theta}\nu\big)\Big| \,\Big|\, \Theta\Bigg]\Bigg\}. \tag{16}$$

Denote $(f \oplus g)(x,y) = f(x) + g(y)$, and let $c(x,y) = \|x-y\|^2$. Further, define the c-conjugate of a function f as $f^c(y) = \inf_x c(x,y) - f(x)$. For each $\theta \in \mathbb{S}^{d-1}$, observe that

$$\begin{aligned}
&\mathsf{W}_p^p\big(\mathfrak{p}_\sharp^\theta\hat{\mu}_n, \mathfrak{p}_\sharp^\theta\hat{\nu}_n\big) - \mathsf{W}_p^p\big(\mathfrak{p}_\sharp^\theta\mu, \mathfrak{p}_\sharp^\theta\nu\big)\\
&\quad\leq \sup_{\substack{(\varphi,\psi)\in L^1(\mu)\times L^1(\nu):\\ \varphi\oplus\psi\leq c}} \big\{(\mathfrak{p}_\sharp^\theta\hat{\mu}_n)\varphi + (\mathfrak{p}_\sharp^\theta\hat{\nu}_n)\psi\big\} - \sup_{\substack{(f,g)\in L^1(\mu)\times L^1(\nu):\\ f\oplus g\leq c}} \big\{(\mathfrak{p}_\sharp^\theta\mu)f + (\mathfrak{p}_\sharp^\theta\nu)g\big\}\\
&\quad\leq \sup_{(\varphi,\psi)\in L^1(\mu)\times L^1(\nu)} \big\{(\mathfrak{p}_\sharp^\theta\hat{\mu}_n)\varphi + (\mathfrak{p}_\sharp^\theta\mu)\varphi^c + (\mathfrak{p}_\sharp^\theta\hat{\nu}_n)\psi + (\mathfrak{p}_\sharp^\theta\nu)\psi^c\big\}\\
&\quad= \mathsf{W}_p^p\big(\mathfrak{p}_\sharp^\theta\hat{\mu}_n, \mathfrak{p}_\sharp^\theta\mu\big) + \mathsf{W}_p^p\big(\mathfrak{p}_\sharp^\theta\hat{\nu}_n, \mathfrak{p}_\sharp^\theta\nu\big).
\end{aligned}$$

Repeating this argument for $\mathsf{W}_p^p\big(\mathfrak{p}_\sharp^\theta\mu, \mathfrak{p}_\sharp^\theta\nu\big) - \mathsf{W}_p^p\big(\mathfrak{p}_\sharp^\Theta\hat{\mu}_n, \mathfrak{p}_\sharp^\Theta\hat{\nu}_n\big)$ we obtain

$$\Big|\mathsf{W}_p^p\big(\mathfrak{p}_\sharp^\theta\mu, \mathfrak{p}_\sharp^\theta\nu\big) - \mathsf{W}_p^p\big(\mathfrak{p}_\sharp^\theta\hat{\mu}_n, \mathfrak{p}_\sharp^\theta\hat{\nu}_n\big)\Big| \leq \mathsf{W}_p^p(\mathfrak{p}_\sharp^\theta\hat{\mu}_n, \mathfrak{p}_\sharp^\theta\mu) + \mathsf{W}_p^p(\mathfrak{p}_\sharp^\theta\hat{\nu}_n, \mathfrak{p}_\sharp^\theta\nu).$$

The proof of Theorem 1 implies that, for any $\theta \in \mathbb{S}^{d-1}$,

$$\begin{aligned}
\mathbb{E}\big[\mathsf{W}_p^p\big(\mathfrak{p}_\sharp^\theta\hat{\mu}_n, \mathfrak{p}_\sharp^\theta\mu\big)\big] &\leq C_p\frac{(\log n)^{\mathbb{1}\{p=2\}}\|\Sigma_\mu\|_{\mathrm{op}}^{p/2}}{n^{(p\wedge 2)/2}},\\
\mathbb{E}\big[\mathsf{W}_p^p\big(\mathfrak{p}_\sharp^\theta\hat{\nu}_n, \mathfrak{p}_\sharp^\theta\nu\big)\big] &\leq C_p\frac{(\log n)^{\mathbb{1}\{p=2\}}\|\Sigma_\nu\|_{\mathrm{op}}^{p/2}}{n^{(p\wedge 2)/2}}.
\end{aligned}$$

Inserting this back into (16), we have

$$\mathbb{E}\left[\left|\widehat{\underline{\mathsf{W}}}_{\mathsf{MC}}^p - \frac{1}{m}\sum_{i=1}^m \mathsf{W}_p^p\big(\mathfrak{p}_\sharp^{\Theta_i}\mu, \mathfrak{p}_\sharp^{\Theta_i}\nu\big)\right|\right] \leq \frac{C_p\big(\|\Sigma_\nu\|_{\mathrm{op}}^{p/2} + \|\Sigma_\mu\|_{\mathrm{op}}^{p/2}\big)(\log n)^{\mathbb{1}\{p=2\}}}{n^{(p\wedge 2)/2}} \tag{17}$$

For the second term in (15), recall that $w_p^p(\theta) := \mathsf{W}_p^p\big(\mathfrak{p}_\sharp^\theta\mu, \mathfrak{p}_\sharp^\theta\nu\big)$ for $\theta \in \mathbb{S}^{d-1}$, and bound

$$\mathbb{E}\left[\left|\frac{1}{m}\sum_{i=1}^m \mathsf{W}_p^p\big(\mathfrak{p}_\sharp^{\Theta_i}\mu, \mathfrak{p}_\sharp^{\Theta_i}\nu\big) - \underline{\mathsf{W}}_p^p(\mu,\nu)\right|\right] \leq \sqrt{\frac{1}{m}\mathrm{Var}\big(w_p^p(\Theta)\big)}$$

To control the variance we use concentration of Lipschitz functions on the unit sphere. By Remark 12 following the proof of Lemma 1, w_p^p is $\tilde{M}_{\mu,\nu}^p$-Lipschitz, with $\tilde{M}_{\mu,\nu}^p \lesssim_p \|\mu x - \nu x\|^p + \sup_{\theta\in\mathbb{S}^{d-1}}(\mu|\theta^\intercal(x-\mu x)|^p + \nu|\theta^\intercal(x-\nu x)|^p)$. Denoting the median by $\mathrm{med}(\cdot)$, we have for $d \geq 3$ (cf. e.g., [32, Chapter 1])

$$\mathbb{P}\Big(\big|w_p^p(\Theta) - \mathrm{med}\big(w_p^p(\theta)\big)\big| \geq t\Big) \leq 8\exp\left(-\frac{(d-2)t^2}{2(\tilde{M}_{\mu,\nu}^p)^2}\right).$$

Consequently,

$$\begin{aligned}
\mathrm{Var}\big(w_p^p(\Theta)\big) &\leq \mathbb{E}\Big[\big(w_p^p(\Theta) - \mathrm{med}\big(w_p^p(\Theta)\big)\big)^2\Big]\\
&= \int_0^\infty \mathbb{P}\Big(\big|w_p^p(\Theta) - \mathrm{med}\big(w_p^p(\Theta)\big)\big| \geq \sqrt{t}\Big)\,dt
\end{aligned}$$

$$\leq \frac{16(\tilde{M}^p_{\mu,\nu})^2}{d-2}.$$

Alternatively, for $d \leq 2$, letting Θ, Θ' be independent samples drawn uniformly from $\mathbb{S}^{d-1}$, we have

$$\mathrm{Var}\big(w_p^p(\Theta)\big) = \frac{1}{2}\mathbb{E}\big[\big|w_p^p(\Theta) - w_p^p(\Theta')\big|^2\big] \leq \frac{(\tilde{M}^p_{\mu,\nu})^2}{2}\mathbb{E}\big[|\Theta - \Theta'|^2\big] \leq (\tilde{M}^p_{\mu,\nu})^2.$$

Combining the two variance bounds, for any $d \geq 1$, we obtain

$$\mathbb{E}\left[\left|\frac{1}{m}\sum_{i=1}^m \mathsf{W}_p^p\big(\mathfrak{p}_\sharp^{\Theta_i}\mu, \mathfrak{p}_\sharp^{\Theta_i}\nu\big) - \underline{\mathsf{W}}_p^p(\mu,\nu)\right|\right] \lesssim \frac{4\tilde{M}^p_{\mu,\nu}}{\sqrt{md}}, \tag{18}$$

where the hidden constant is universal.

We now focus on bounding $\tilde{M}^p_{\mu,\nu}$, leveraging log-concavity of μ and ν. We present the derivation for $\sup_\theta \mu|\theta^\intercal(x-\mu x)|$; the one corresponding to ν is analogous. To control this term we use exponential concentration for 1-Lipschitz functions of log-concave random variables. Recalling that μ and ν being log-concave implies that so are $\mathfrak{p}_\sharp^\theta\mu$ and $\mathfrak{p}_\sharp^\theta\nu$, Theorem 1.2 in [39] yields

$$\mathbb{P}\big(\big|\theta^\intercal X - \mu(\theta^\intercal x)\big| > t\big) \leq e\exp(-D_\mu t),$$

where $X \sim \mu$ and $D_\mu \geq c/\sqrt{\|\Sigma_\mu\|_{\mathrm{op}}}$, with a universal constant c. Then,

$$\begin{aligned}
\sup_{\theta\in\mathbb{S}^{d-1}} \mu|\theta^\intercal(x-\mu x)|^p &\leq \int_0^\infty \mathbb{P}\big(\big|\theta^\intercal x - \mu(\theta^\intercal x)\big|^p > t\big)dt \\
&\leq \int_0^\infty e\exp(-D_\mu t^{1/p})\,dt \\
&= \Gamma(p+1)D_\mu^{-p} \\
&\leq \Gamma(p+1)\left(\frac{\sqrt{\|\Sigma_\mu\|_{\mathrm{op}}}}{c}\right)^p \\
&\leq C_p\|\Sigma_\nu\|_{\mathrm{op}}^{p/2},
\end{aligned}$$

for a constant C_p depending only on p. Similarly, we obtain

$$\sup_{\theta\in\mathbb{S}^{d-1}} \nu|\theta^\intercal(x-\nu x)|^p \leq C_p\|\Sigma_\nu\|_{\mathrm{op}}^{p/2},$$

which together implies

$$\tilde{M}^p_{\mu,\nu} \leq C_p'\left(\|\mu x - \nu x\|^p + \|\Sigma_\mu\|_{\mathrm{op}}^{p/2} + \|\Sigma_\nu\|_{\mathrm{op}}^{p/2}\right).$$

Inserting the above bound into (18) and combining with (17) yields the result.

D.10 Proof of Proposition 6

Observe that $\tilde{w}_2^2(\theta)$ is M_n-Lipschitz by Lemma 1 and ρ_n-weakly convex by Lemma 2.2 in [34], where $M_n = 4\sup_\theta(\hat{\mu}_n|\theta^\intercal x|^2 + \hat{\nu}_n|\theta^\intercal x|^2)$ and $\rho_n = 2\max_{i,j}\|X_i - Y_j\|^2$. By equation (2.10) in [14], there exists a choice of step sizes $\alpha_t = \frac{c_{\rho_n,M_n}}{\sqrt{t+1}}$, such that Algorithm 1 for the objective $\varphi(\theta) = \tilde{w}_2^2 + \delta_{\mathbb{B}^d}$, where $\delta_{\mathbb{B}^d} = \infty\mathbb{1}_{(\mathbb{B}^d)^c}$, outputs a point θ_{t^*} that is close to a near-stationary point θ^*, in the sense that $\mathbb{E}_{t^*}[\|\theta^* - \theta_{t^*}\|] \leq \frac{\epsilon}{2\rho_n}$ and $\mathrm{dist}\big(0, \partial\tilde{w}_2^2(\theta^*)\big) \leq \epsilon$, in number of steps

$$T \leq \left\lceil \frac{64\rho_n^2 M_n^2\big(1 \wedge \frac{M_n}{2\rho_n}\big)}{\epsilon^4}\right\rceil.$$

We derive high probability upper bounds on M_n and ρ_n to obtain a non-stochastic bound on the computational complexity of our algorithm.

We will first reduce our problem to the case where μ and ν are isotropic log-concave, where our assumptions will lead to concentration inequalities on the above quantities. Assume first that Σ_μ

and Σ_ν have rank d. Let $T^\mu(x) = \Sigma_\mu^{-1/2}(x - \mu x)$, and define T^ν analogously. Then, $\tilde{\mu} = T^\mu_\sharp \mu$ and $\tilde{\nu} = T^\nu_\sharp \nu$ are isotropic log-concave. Let $\tilde{\mu}_n$ and $\tilde{\nu}_n$ be empirical measures corresponding to $\tilde{\mu}$ and $\tilde{\nu}$, obtained by applying T_μ and T_ν to samples from μ and ν, respectively. For the first, we have

$$M_n \le M_{\mu,\nu}^2 + 24 \sup_{\theta \in \mathbb{S}^{d-1}} |\hat{\mu}_n \theta^\intercal (x - \mu x)|^2 + 24 \sup_{\theta \in \mathbb{S}^{d-1}} \hat{\nu}_n |\theta^\intercal (x - \nu x)|^2, \quad \text{and}$$

$$\rho_n \le 6\Big(\|\mu x - \nu x\|^2 + \|\Sigma_\mu\|_{\mathrm{op}} \max_i \|\Sigma_\mu^{-1/2}(X_i - \mu x)\|^2 + \|\Sigma_\nu\|_{\mathrm{op}} \max_j \|\Sigma_\nu^{-1/2}(Y_j - \nu x)\|^2 \Big).$$

Further, assuming that μ and ν are centered, we have

$$\begin{aligned}
&\sup_{\theta \in \mathbb{S}^{d-1}} \left| \hat{\mu}_n |\theta^\intercal x|^2 - \mu |\theta^\intercal x|^2 \right| \\
&\quad \le \sup_{\theta \in \mathbb{S}^{d-1}} \left| \hat{\mu}_n |\theta_{\Sigma_\mu}^\intercal x|^2 - \mu |\theta^\intercal x|^2 \right| \qquad \left[\theta_{\Sigma_\mu} = \frac{\Sigma_\mu^{-1/2}\theta}{\|\Sigma_\mu^{-1/2}\theta\|} \right] \\
&\quad \le \|\Sigma_\mu\|_{\mathrm{op}} \sup_{\theta \in \mathbb{S}^{d-1}} \left| \tilde{\mu}_n |\theta^\intercal x|^2 - \tilde{\mu} |\theta^\intercal x|^2 \right|,
\end{aligned}$$

and similarly,

$$\sup_{\theta \in \mathbb{S}^{d-1}} \left| \hat{\nu}_n |\theta^\intercal x|^2 - \nu |\theta^\intercal x|^2 \right| \le \|\Sigma_\nu\|_{\mathrm{op}} \sup_{\theta \in \mathbb{S}^{d-1}} \left| \tilde{\nu}_n |\theta^\intercal x|^2 - \tilde{\nu} |\theta^\intercal x|^2 \right|.$$

For isotropic $\tilde{\mu}$ and $\tilde{\nu}$, we have (cf. Theorem 4.2 in [2])

$$\mathbb{P}\left(\sup_{\theta \in \mathbb{S}^{d-1}} \left| \tilde{\mu}_n |\theta^\intercal x|^2 - \tilde{\mu} |\theta^\intercal x|^2 \right| \le \epsilon \right) \ge 1 - \exp\left(-c n^{1/4} \epsilon \sqrt{d} \right).$$

Choosing $\epsilon = 1/c$ above and noting that $\tilde{\mu} |\theta^\intercal x|^2 = \tilde{\nu} |\theta^\intercal x|^2 = 1$, we have

$$M_n \le M_{\mu,\nu}^2 + 4(1 + 1/c)\Big(\|\Sigma_\mu\|_{\mathrm{op}} + \|\Sigma_\nu\|_{\mathrm{op}} \Big) \tag{19}$$

with probability at least $1 - \frac{2}{n}$. Additionally, by Lemma 3.1 in [2], if $d \ge \big(\log n\big)^2$, there exists a universal constant $C > 0$ such that

$$\max\left\{ \max_i \|\Sigma_\mu^{-1/2}(X_i - \mu x)\|^2, \max_i \|\Sigma_\mu^{-1/2}(X_i - \mu x)\|^2 \right\} \le Cd,$$

implying

$$\rho_n \le 6\Big(\|\mu x - \nu x\|^2 + Cd \left(\|\Sigma_\mu\|_{\mathrm{op}} + \|\Sigma_\nu\|_{\mathrm{op}} \right) \Big) \tag{20}$$

with probability at least $1 - \frac{2}{n}$.

If Σ_μ and Σ_ν are not full rank, then the above results hold for $\mu * \mathrm{Unif}(B_d(0, \sigma))$ and $\nu * \mathrm{Unif}(B_d(0, \sigma))$ instead, which are log-concave measures with covariance matrices $\Sigma_\mu + \sigma^2 I_d/(d+1)$ and $\Sigma_\nu + \sigma^2 I_d/(d+1)$, respectively. Letting M_n^σ, ρ_n^σ and $M_{\mu,\nu}^{2,\sigma}$ denote M_n, ρ_n and $M_{\mu,\nu}^2$ for these perturbed measures, we observe that $|\rho_n - \rho| \le \sigma^2$, $|M_n^\sigma - M_n| \le 96\sigma^2$, and $|M_{\mu,\nu}^{2,\sigma} - M_{\mu,\nu}^2| \le 48\sigma^2$. Choosing $\sigma^2 = \|\Sigma_\mu\|_{\mathrm{op}} + \|\Sigma_\nu\|_{\mathrm{op}}$, we see that (19) and (20) hold for non-full dimensional μ and ν as well with adjustments to c and C.

Combining (19) and (20), and noting that sorting to obtain the optimal permutation σ^* and computing the subdifferential $\partial\rho(\sigma^*, \theta) = \nabla_\theta \rho(\sigma^*, \theta)$ takes $O(n \log n)$ operations, we have the result. □

D.11 Proof of Proposition 8

By Lemma 1, we have that $\hat{w}_p(\theta)$ is $\hat{L}_n$-Lipschitz with $\hat{L}_n := \sup_{\theta \in \mathbb{S}^{d-1}} \big\{ (\hat{\mu}_n |\theta^\intercal x|^p)^{1/p} + (\hat{\nu}_n |\theta^\intercal x|^p)^{1/p} \big\}$. This yields, via the LIPO convergence guarantee (Corollary 13 in [36]), that

$$\max_{\theta \in \mathbb{S}^{d-1}} \hat{w}_p(\theta) - \max_{1 \le i \le k} \hat{w}_p(\Theta_i) \le 2\hat{L}_n \left(\frac{\log(1/\delta)}{k} \right)^{1/d} \tag{21}$$

with probability at least $1 - \delta$.

As in the previous section, we will first reduce our problem to the case where μ and ν are isotropic log-concave. For $1 \le p \le 2$, $\hat{L}_n \le \sup_{\theta \in \mathbb{S}^{d-1}} \left\{ (\hat{\mu}_n |\theta^\intercal x|^2)^{1/2} + (\hat{\nu}_n |\theta^\intercal x|^2)^{1/2} \right\}$, so that it suffices to bound $\hat{L}_n$ for $p \ge 2$. We have

$$
\begin{aligned}
\hat{L}_n &= \sup_{\theta \in \mathbb{S}^{d-1}} \left\{ (\hat{\mu}_n |\theta^\intercal x|^p)^{1/p} + (\hat{\nu}_n |\theta^\intercal x|^p)^{1/p} \right\} \\
&\le \sup_{\theta \in \mathbb{S}^{d-1}} \left\{ (\hat{\mu}_n |\theta^\intercal (x - \mu x)|^p)^{1/p} + (\hat{\nu}_n |\theta^\intercal (x - \nu x)|^p)^{1/p} \right\} \\
&\qquad + \sup_\theta |\mu(\theta^\intercal x)| + \sup_\theta |\nu(\theta^\intercal x)| \\
&\le \|\Sigma_\mu\|_{\mathrm{op}}^{1/2} \left[\sup_{\theta \in \mathbb{S}^{d-1}} (\tilde{\mu} |\theta^\intercal x|^p)^{1/p} + \sup_{\theta \in \mathbb{S}^{d-1}} \left| (\tilde{\mu}_n |\theta^\intercal x|^p)^{1/p} - (|\tilde{\mu}(\theta^\intercal x)|^p)^{1/p} \right| \right] \\
&\qquad + \|\Sigma_\nu\|_{\mathrm{op}}^{1/2} \left[\sup_{\theta \in \mathbb{S}^{d-1}} (\tilde{\mu} |\theta^\intercal x|^p)^{1/p} + \sup_{\theta \in \mathbb{S}^{d-1}} \left| (\tilde{\nu}_n |\theta^\intercal x|^p)^{1/p} - (|\tilde{\nu}(\theta^\intercal x)|^p)^{1/p} \right| \right] \\
&\qquad + \sup_\theta |\mu(\theta^\intercal x)| + \sup_\theta |\nu(\theta^\intercal x)| \\
&\le \|\Sigma_\mu\|_{\mathrm{op}}^{1/2} \left[(2p)^{1/p} + \sup_{\theta \in \mathbb{S}^{d-1}} \left| (\tilde{\mu}_n |\theta^\intercal x|^p)^{1/p} - (|\tilde{\mu}(\theta^\intercal x)|^p)^{1/p} \right| \right] \\
&\qquad + \|\Sigma_\nu\|_{\mathrm{op}}^{1/2} \left[(2p)^{1/p} + \sup_{\theta \in \mathbb{S}^{d-1}} \left| (\tilde{\nu}_n |\theta^\intercal x|^p)^{1/p} - (|\tilde{\nu}(\theta^\intercal x)|^p)^{1/p} \right| \right] \\
&\qquad + \sup_\theta |\mu(\theta^\intercal x)| + \sup_\theta |\nu(\theta^\intercal x)|
\end{aligned}
$$

Now, applying [2, Theorem 4.2] with $\epsilon = 1/2$ and $t = 1$, we get

$$
\begin{aligned}
\mathbb{P}\left(\sup_{\theta \in \mathbb{S}^{d-1}} \left| (\tilde{\mu}_n |\theta^\intercal x|^p)^{1/p} - (|\tilde{\mu}(\theta^\intercal x)|^p)^{1/p} \right| > \frac{1}{2} \right) &\le \mathbb{P}\left(\sup_{\theta \in \mathbb{S}^{d-1}} |\tilde{\mu}_n |\theta^\intercal x|^p - |\tilde{\mu}(\theta^\intercal x)|^p| > \frac{1}{2^p} \right) \\
&\le 1 - e^{-c_p \sqrt{d}}
\end{aligned} \tag{22}
$$

under assumed constraints on n in the statement. An analogous bound holds for ν, which yields that

$$
\mathbb{P}\left(\hat{L}_n \ge (\|\Sigma_\mu\|_{\mathrm{op}}^{1/2} + \|\Sigma_\nu\|_{\mathrm{op}}^{1/2}) \left((2p)^{1/p} + \frac{1}{2} \right) + \sup_\theta |\mu(\theta^\intercal x)| + \sup_\theta |\nu(\theta^\intercal x)| \right) \le e^{-c_p \sqrt{d}}
$$

Recall that $\beta = \exp(-c_p \sqrt{d})$. Plugging (22) back into (21), we get

$$
\max_{\theta \in \mathbb{S}^{d-1}} \hat{w}_p(\theta) - \max_{1 \le i \le k} \hat{w}_p(\Theta_i) \le L_{\mu,\nu} \left(\frac{\log(1/\delta)}{k} \right)^{1/d} \tag{23}
$$

with probability $1 - \delta - \beta$.

Finally, we have $\max_{\theta \in \mathbb{S}^{d-1}} \hat{w}_p(\theta) = \overline{\mathsf{W}}_p(\hat{\mu}_n, \hat{\nu}_n)$, and

$$
|\overline{\mathsf{W}}_p(\hat{\mu}_n, \hat{\nu}_n) - \overline{\mathsf{W}}_p(\mu, \nu)| \le \overline{\mathsf{W}}_p(\hat{\mu}_n, \mu) + \overline{\mathsf{W}}_p(\hat{\nu}_n, \nu).
$$

By (5b) in Proposition 7, for any $t > 0$,

$$
\mathbb{P}\left(\overline{\mathsf{W}}_p(\hat{\mu}_n, \mu) \ge \alpha_{n,\mu} + t \right) \le 2 \exp\left(-K_\mu \min\left(n^{1/p} t, n^{2/(2 \vee p)} t^2 \right) \right),
$$

$$
\mathbb{P}\left(\overline{\mathsf{W}}_p(\hat{\nu}_n, \nu) \ge \alpha_{n,\nu} + t \right) \le 2 \exp\left(-K_\nu \min\left(n^{1/p} t, n^{2/(2 \vee p)} t^2 \right) \right),
$$

where $K_\mu \lesssim d^{o_d(1)} \max\{\|\Sigma_\mu\|_{\mathrm{op}}^{1/2}, \|\Sigma_\mu\|_{\mathrm{op}}\}$ and $K_\nu \lesssim d^{o_d(1)} \max\{\|\Sigma_\nu\|_{\mathrm{op}}^{1/2}, \|\Sigma_\nu\|_{\mathrm{op}}\}$. Setting

$$
\gamma_n(t) = 2 \exp\left(-K_\mu^{-1} \min\left(n^{1/p} t, n^{2/(2 \vee p)} t^2 \right) \right) + 2 \exp\left(-K_\nu^{-1} \min\left(n^{1/p} t, n^{2/(2 \vee p)} t^2 \right) \right),
$$

we then have

$$
\mathbb{P}\left(|\overline{\mathsf{W}}_p(\hat{\mu}_n, \hat{\nu}_n) - \overline{\mathsf{W}}_p(\mu, \nu)| > \alpha_{n,\mu} + \alpha_{n,\nu} + 2t \right) \le \gamma_n(t).
$$

Combining the above display with (23), we get the desired result.

E Additional Experiments and Details

Code for reproducing this paper's experiments can be found at `https://github.com/sbnietert/sliced-Wp`. Distance computations and plots for Figure 1 were performed on a cluster machine with 8 CPU cores and 64GB RAM in approximately 6 hours. Distance computations and plots for Figures 2 and 3 were performed on a cluster machine with 4 CPU cores and 20GB RAM in approximately 30 minutes. For Figure 3 (right), the lower bound on W_1 is computed by only considering couplings which leave the shared mass at 0 unmoved.

As an additional experimental setup along the lines of Figure 1, we consider Model (3): Gaussian mixtures $\mu = \frac{1}{10}\sum_{i=1}^{10}\mathcal{N}(\mu_{1,i}, \Sigma_{1,i})$ and $\nu = \frac{1}{10}\mathcal{N}(\mu_{2,i}, \Sigma_{2,i})$, where means $\mu_{1,i}$ and $\mu_{2,i}$ are respectively generated from $\mathcal{N}(\mathbf{1}_d, I_d)$ and $\mathcal{N}(3\,\mathbf{1}_d, I_d)$, and the covariance matrices of the mixtures are simulated as $\frac{1}{k}X^\intercal X$, where X is $k \times d$ data matrix generated from $\mathcal{N}(0, I_d)$ and k is a uniformly sampled integer from 1 to d. Conditioned on fixed random choices of μ and ν, we provide the corresponding projection and sample complexity plots in Figure 4, with general trends matching those of Figure 1. For both this experiment and Figure 1 in the main text, the population versions of the distances where no closed forms exist were calculated by setting the number of samples and Monte Carlo directions to 5000 and 2000 respectively. Computations and plots were performed on a cluster machine with 8 CPU cores and 64GB RAM in approximately 12 hours.

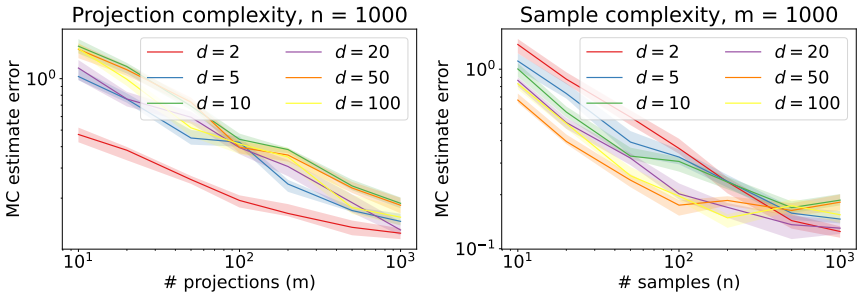

Figure 4: $\left|\widehat{\underline{\mathsf{W}}}^2_{\mathsf{MC}} - \underline{\mathsf{W}}^2_2(\mu,\nu)\right|$ under Model (3).

Finally, we consider how the robustness properties of sliced W_p may impact its application to generative modeling. Minimum distance estimation with respect to classic W_1 serves as a theoretical foundation for Wasserstein GANs [3, 22], a successful approach for training generative models. Later work extended this approach to average and max-sliced W_p [18, 17], albeit at a slightly less direct level (in these papers, sliced distances are computed in a feature embedding space rather than raw image space). In Figure 5, we display samples generated from open source implementations of the standard Wasserstein GAN with Gradient Penalty (WGAN-GP) [22] and an average-sliced WGAN [18] trained for 20 epochs over the MNIST dataset [16] of digit images with 10% random noise contamination, using default parameter settings. Computations were performed on a cluster machine with 4 CPU cores, a NVIDIA Tesla T4 GPU, and 20GB RAM in roughly 12 hours. While there are differences between the produced samples, the two GAN architectures seem too distinct to draw any strong conclusions. Moreover, the robustness guarantees from Section 4 hold after preprocessing that appears too expensive to perform for data of this scale, so it is not surprising that the sliced WGAN reproduces random noise. Translating methods and guarantees for standard

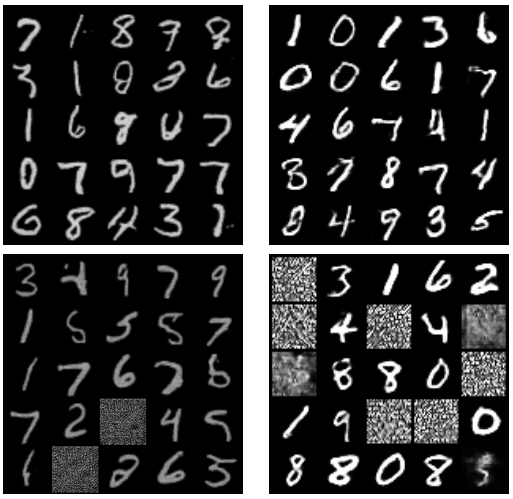

(a) Sliced WGAN (b) WGAN-GP

Figure 5: Preliminary GAN experiments with uncontaminated (top) vs. contaminated (bottom) MNIST data.

WGAN robustification (e.g., [46]) to the sliced setting and thorough empirical comparisons are an interesting avenue for future research beyond the scope of this paper.