

---

# Regret-Optimal Wasserstein-Robust Linear Regression

---

Sloan Nietert

EPFL

sloan.niertert@epfl.ch

Daniel Kuhn

EPFL

daniel.kuhn@epfl.ch

## Abstract

Distributionally robust optimization (DRO) minimizes worst-case risk under distributional uncertainty, but its decisions can be suboptimal in terms of their excess risk, or regret. We study this gap for multivariate least-squares regression, assuming that the observed source law and unknown target law differ by an adversarial 2-Wasserstein perturbation of size  $\varepsilon$ . We first identify the asymptotic complexity of robust decision-making in the  $\varepsilon \rightarrow 0$  limit, characterized by a problem-specific condition number  $\kappa$ . Here, a wide spectrum of simple algorithms, including DRO, achieve the instance-optimal rate of  $O(\kappa\varepsilon^2)$ . For fixed  $\varepsilon > 0$ , we prove that DRO still achieves the minimax rate if the problem has bounded residual effective rank. On the other hand, we identify significant failure modes where DRO is provably suboptimal by dimension-dependent factors. Fortunately, we prove that two alternative procedures, minimax regret optimization (MRO) and spectral-norm regularization (SNR), attain the minimax rate. We derive an exact convex reformulation of MRO and an accelerated proximal algorithm for SNR, along with high-probability finite-sample guarantees. We support these theoretical results with numerical experiments separating the performance of MRO and SNR from DRO.

## 1 Introduction

Distributionally robust optimization (DRO) is a flexible tool for hedging against model misspecification in ML, statistics, and operations research (Rahimian and Mehrotra, 2022; Lin et al., 2022; Kuhn et al., 2025). Wasserstein DRO is a well-studied variant that minimizes worst-case risk over all Wasserstein perturbations of the observed data distribution. For many tasks, this objective admits dual formulations, tractable implementations, and regularization interpretations (Kuhn et al., 2019; Gao and Kleywegt, 2022; Gao et al., 2024). Model misspecification comes in many flavors, e.g., sampling error, distribution shift, and data poisoning. In this work, we take an adversarial view: the learner observes a source law and seeks a predictor that performs well under a nearby but otherwise unrestricted target law. As we will see, robust risk at the source (the objective minimized by DRO) is fundamentally distinct from decision quality at the target. Our motivating question is the following:

*When does Wasserstein DRO make good decisions in the face of adversarial data perturbations, and when can we efficiently outperform it?*

We examine this for the setting of multivariate linear regression. Fix dimensions  $p, m \geq 1$  and let  $\mathcal{A} := \mathbb{R}^{m \times p}$  be the class of linear predictors. Set  $d := p + m$  and write  $\mathcal{P}_2(\mathbb{R}^d)$  for the laws of  $Z = (X, Y) \in \mathbb{R}^p \times \mathbb{R}^m$  with finite second moment. For  $P \in \mathcal{P}_2(\mathbb{R}^d)$ , the least-squares risk of  $\mathbf{A} \in \mathcal{A}$  is  $L(\mathbf{A}, P) := \mathbb{E}_P \|\mathbf{A}X - Y\|^2$ , and its *excess risk*, or (ex-ante) *regret*, is

$$R(\mathbf{A}, P) := L(\mathbf{A}, P) - \inf_{\mathbf{A}' \in \mathcal{A}} L(\mathbf{A}', P).$$

This regret is ex-ante, rather than ex-post, because we benchmark against the best decision given knowledge of  $P$  and not a realized pair  $(X, Y) \sim P$ . Both risk and regret depend on  $P$  only through

its second-moment matrix  $\Sigma_P := \mathbb{E}_P[ZZ^\top]$ , with blocks  $\Sigma_X^P$ ,  $\Sigma_Y^P$ , and  $\Sigma_{XY}^P = (\Sigma_{YX}^P)^\top$ . Write  $\mathcal{P}^* := \{P \in \mathcal{P}_2(\mathbb{R}^d) : \Sigma_X^P \succ \mathbf{0}\}$  for the class of *non-degenerate* laws. For  $P \in \mathcal{P}^*$ , its unique risk minimizer is  $\mathbf{A}_P := \Sigma_{YX}^P (\Sigma_X^P)^{-1}$  and regret has the closed form

$$R(\mathbf{A}, P) = \|(\mathbf{A} - \mathbf{A}_P)(\Sigma_X^P)^{1/2}\|_F^2.$$

We write  $\|\cdot\|_F$  and  $\|\cdot\|_{\text{op}}$  for the Frobenius norm and Euclidean operator norm, respectively.

In our population model, an adversary chooses a source–target pair  $P_s, P_t \in \mathcal{P}_2(\mathbb{R}^d)$  satisfying  $W_2(P_s, P_t) \leq \varepsilon$  and reveals only the source  $P_s$ . Given the radius  $\varepsilon > 0$ , the learner applies an estimator  $\hat{\mathbf{A}} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathcal{A}$  and incurs target regret  $R(\hat{\mathbf{A}}[P_s], P_t)$ . We measure the *robust regret* of  $\hat{\mathbf{A}}$  at a reference law  $P$  by allowing  $P$  to occupy either the source or target role:

$$R_\varepsilon(\hat{\mathbf{A}}, P) := \sup_{Q \in \mathcal{B}_\varepsilon(P)} \max\{R(\hat{\mathbf{A}}[P], Q), R(\hat{\mathbf{A}}[Q], P)\},$$

where  $\mathcal{B}_\varepsilon(P) := \{Q \in \mathcal{P}_2(\mathbb{R}^d) : W_2(P, Q) \leq \varepsilon\}$ . The first term varies the target with fixed source  $P$ , while the second varies the source with fixed target  $P$ .

**Robust estimators.** We primarily compare *empirical risk minimization* (ERM), *Wasserstein DRO*, *minimax regret optimization* (MRO), and *spectral-norm regularization* (SNR):<sup>1</sup>

$$\begin{aligned} \text{ERM}[P] &\in \arg \min_{\mathbf{A} \in \mathcal{A}} L(\mathbf{A}, P), \\ \text{DRO}_\varepsilon[P] &\in \arg \min_{\mathbf{A} \in \mathcal{A}} \sup_{Q \in \mathcal{B}_\varepsilon(P)} L(\mathbf{A}, Q), \\ \text{MRO}_\varepsilon[P] &\in \arg \min_{\mathbf{A} \in \mathcal{A}} \sup_{Q \in \mathcal{B}_\varepsilon(P)} R(\mathbf{A}, Q), \\ \text{SNR}_\varepsilon[P] &\in \arg \min_{\mathbf{A} \in \mathcal{A}} L(\mathbf{A}, P) + \varepsilon^2 \|\mathbf{A}\|_{\text{op}}^2. \end{aligned}$$

Here, ERM fits the source law directly, returning  $\mathbf{A}_P$  when  $P$  is non-degenerate. MRO replaces the risk objective in DRO with regret. Finally, SNR adds a squared operator-norm penalty to ERM.

**Local guarantees.** To quantify estimation complexity, we first study the local minimax benchmark

$$R_\varepsilon^*(P) := \inf_{\hat{\mathbf{A}}} R_\varepsilon(\hat{\mathbf{A}}, P) = \inf_{\mathbf{A} \in \mathcal{A}} \sup_{Q \in \mathcal{B}_\varepsilon(P)} R(\mathbf{A}, Q). \quad (1)$$

Equality follows by considering  $P$  in the source role for the “ $\geq$ ” direction and considering constant estimators for the “ $\leq$ ” direction. For non-degenerate  $P$ , let  $\Sigma_E^P$  denote the second-moment matrix of the residual  $Y - \mathbf{A}_P X$ . Our first result shows that the following limit exists and computes it exactly:

$$\kappa_P := \lim_{\varepsilon \downarrow 0} \varepsilon^{-2} R_\varepsilon^*(P) = 1 + \lambda_{\max} \left( \mathbf{A}_P \mathbf{A}_P^\top + \frac{\Sigma_E^P}{\lambda_{\min}(\Sigma_X^P)} \right).$$

Thus, the *condition number*  $\kappa_P$  governs the optimal local regret,  $\kappa_P \varepsilon^2 + o(\varepsilon^2)$ . It reflects the size of the regression coefficients and that of the residual second moment relative to the smallest feature eigenvalue. All four estimators attain this local rate up to absolute multiplicative constants as  $\varepsilon \downarrow 0$ .

**Global guarantees.** For fixed  $\varepsilon > 0$ , we tightly characterize estimator performance as a function of  $\varepsilon$  and the condition number. For a class  $\mathcal{G} \subseteq \mathcal{P}_2(\mathbb{R}^d)$ , define the *worst-case* and *minimax* regrets  $R_\varepsilon(\hat{\mathbf{A}}, \mathcal{G}) := \sup_{P \in \mathcal{G}} R_\varepsilon(\hat{\mathbf{A}}, P)$  and  $R_\varepsilon^*(\mathcal{G}) := \inf_{\hat{\mathbf{A}}} R_\varepsilon(\hat{\mathbf{A}}, \mathcal{G})$ , so that  $R_\varepsilon^*(P) = R_\varepsilon^*(\{P\})$ . For the well-conditioned family  $\mathcal{G}(\kappa) := \{P \in \mathcal{P}^* : \kappa_P \leq \kappa\}$ , our main findings are summarized in Table 1. Namely, the minimax rate is  $O(\kappa \varepsilon^2)$  for all  $\varepsilon > 0$ , and this is attained by MRO and SNR. In contrast, ERM can suffer arbitrarily large regret even for well-conditioned instances. DRO lies between these extremes; when  $m = O(1)$  or, more generally, when the residual second moment  $\Sigma_E^P$  has low effective rank, we prove that it achieves the minimax rate. For  $\kappa \geq 2$ , however, its worst-case robust regret is suboptimal by a multiplicative factor of  $\sqrt{\min\{p, m\}} + \sqrt{m/\kappa}$ .

For upper bounds, our analysis employs the Bures– $W_2$  geometry of second-moment matrices to prove a key regret stability lemma. In particular, we show that  $R(\mathbf{A}_P, Q) \lesssim \kappa_P W_2(P, Q)^2$  for every

<sup>1</sup> All four optimization problems admit minimizers; see Section D.1. All of our regret bounds (upper and lower) hold for any minimizer selection rule.

| Estimator $\hat{\mathbf{A}}$ | ERM         | DRO $_{\varepsilon}$                                        | MRO $_{\varepsilon}$       | SNR $_{\varepsilon}$                         |
|------------------------------|-------------|-------------------------------------------------------------|----------------------------|----------------------------------------------|
| Robust regret                | $\infty$    | $(\sqrt{\min\{p, m\}}\kappa + \sqrt{m\kappa})\varepsilon^2$ | $\kappa\varepsilon^2$      | $\kappa\varepsilon^2$                        |
| Minimax optimal?             | ×           | ×                                                           | ✓                          | ✓                                            |
| Computation                  | Closed form | Convex reformulation                                        | (New) convex reformulation | Accelerated proximal gradient method (FISTA) |

Table 1: Worst-case robust regret  $R_{\varepsilon}(\hat{\mathbf{A}}, \mathcal{G}(\kappa))$  for our running estimators, up to absolute multiplicative constants, for  $\varepsilon > 0$  and  $\kappa \geq 2$ . Computation uses second-moment matrices.

non-degenerate  $P$  and arbitrary  $Q \in \mathcal{P}_2(\mathbb{R}^d)$ . More generally, our stability lemma controls changes in square-root regret through the predictor’s operator norm, explaining the role of the SNR penalty. For lower bounds, we provide explicit constructions where these upper bounds are tight.

**Computation and finite-sample statistics.** We derive a novel, exact convex reformulation of MRO, building upon previous results for DRO (Shafieezadeh-Abadeh et al., 2018), along with an efficient proximal algorithm for SNR, based on FISTA (Beck and Teboulle, 2009). Given a fixed source second-moment matrix and radius  $\varepsilon > 0$ , our SNR algorithm reaches objective gap  $\delta > 0$  in  $O(\delta^{-1/2})$  iterations, each using a singular value decomposition. We also extend our results to two finite-sample models: adversarial corruption of independent and identically distributed (i.i.d.) target observations, and i.i.d. source observations with a Wasserstein-perturbed population target. Here, we prove high-probability regret bounds under fourth-moment assumptions, with improved dependence on the confidence level under sub-Gaussian tails. Finally, we provide synthetic numerical experiments illustrating the predicted separation of MRO and SNR from DRO.

## 1.1 Related Work

Wasserstein DRO (WDRO) for regression and other quadratic models has been studied extensively through the lens of duality, regularization, and robust risk minimization; see, e.g., Nguyen et al. (2022, 2021a); Taşkesen et al. (2023) in addition to the works cited above. Shafieezadeh-Abadeh et al. (2018) explicitly develop WDRO algorithms for our multivariate linear regression setting, applied to a sequential Kalman filtering problem. More broadly, WDRO has gained interest in ML from the perspective of adversarial training (Sinha et al., 2018), and Wasserstein-robust risk has been well explored by several communities (Lee and Raginsky, 2018; Blanchet and Murthy, 2019).

Our focus is different: we evaluate estimators by *excess* risk at the unknown target law, rather than by risk at the observed source or worst-case law. This distinction is essential, since regret-optimal procedures need not coincide with DRO. MRO, also called distributionally robust regret optimization, has garnered significant study. Agarwal and Zhang (2022) compare MRO and DRO with general ambiguity sets and show that heterogeneity in irreducible risk can cause DRO to incur larger regret. In operations research and control, related formulations have been studied for newsvendor models (Perakis and Roels, 2008; Chen and Xie, 2021), more general Wasserstein ambiguity sets (Cho and Yang, 2024; Fiechtner and Blanchet, 2025), and multistage stochastic control (Fiechtner and Blanchet, 2026). However, these works do not establish worst-case robust regret rates for estimator comparison.

Fiechtner and Blanchet (2025) introduce a first-order sensitivity parameter which nullifies for our setting. Theorem 3 in Version 5 of their paper (posted after the conference submission of this work) establishes a second-order sensitivity expansion whose coefficient in fact specializes to our condition number. We provide a self-contained Bures-geometric derivation, an explicit formula for multivariate least squares, and connect this local quantity to global stability bounds. Their all-radius optimality result for ERM concerns losses jointly quadratic in the decision and uncertainty, excluding our setting.

Our model is closely related to adversarially robust statistics under Wasserstein perturbations. Scalar-response linear regression has been studied under 1-Wasserstein (Zhu et al., 2022; Liu and Loh, 2023) and 2-Wasserstein contamination (Chao and Dobriban, 2023). These works establish robust guarantees under different distributional and design assumptions; our global comparisons are instead indexed by the local minimax condition number. Moreover, moving to multivariate responses is crucial for us to separate the worst-case regret of DRO from the minimax rates of MRO and SNR.

## 1.2 Preliminaries

**Notation.** Unsubscripted vector norms are Euclidean. Write  $\mathbb{S}_+^d$  for the cone of positive semidefinite  $d \times d$  matrices and  $\mathbb{O}^d := \{\mathbf{U} \in \mathbb{R}^{d \times d} : \mathbf{U}\mathbf{U}^\top = \mathbf{U}^\top\mathbf{U} = \mathbf{I}_d\}$  for the family of orthogonal matrices. For symmetric matrices  $\mathbf{M}, \mathbf{N}$  of the same size, write  $\mathbf{M} \preceq \mathbf{N}$  ( $\mathbf{M} \prec \mathbf{N}$ ) when  $\mathbf{N} - \mathbf{M}$  is positive semidefinite (positive definite), and define  $\succeq, \succ$  analogously. Let  $\lambda_{\min}(\mathbf{M})$  and  $\lambda_{\max}(\mathbf{M})$  denote the smallest and largest eigenvalues of a symmetric matrix  $\mathbf{M}$ . For a real matrix  $\mathbf{M}$ , let  $\mathbf{M}^\dagger$  denote its Moore–Penrose pseudoinverse. If square, write  $\text{tr}(\mathbf{M})$  for its trace. Let  $\mathbf{0}$  and  $\mathbf{I}$  denote the zero and identity matrices, respectively, indexed by dimension when needed. The notation  $\text{diag}$  constructs a diagonal matrix from scalar arguments and a block-diagonal matrix from square-matrix arguments. For  $\mathbf{M} \succeq \mathbf{0}$ ,  $\mathbf{M}^{1/2}$  denotes its unique positive semidefinite square root. Write  $a \lesssim b$  if  $a \leq Cb$  for an absolute constant  $C > 0$ . Write  $a \gtrsim b$  if  $b \lesssim a$  and  $a \asymp b$  if  $a \lesssim b$  and  $b \lesssim a$ .

**Bures–Wasserstein geometry.** We employ the 2-Wasserstein distance defined for  $P, Q \in \mathcal{P}_2(\mathbb{R}^d)$  by

$$W_2(P, Q) := \inf_{\pi \in \Pi(P, Q)} \mathbb{E}_{(Z, Z') \sim \pi} [\|Z - Z'\|^2]^{\frac{1}{2}},$$

where  $\Pi(P, Q)$  denotes the set of joint probability laws on  $\mathbb{R}^d \times \mathbb{R}^d$  with marginals  $P$  and  $Q$ . Since our least-squares risk only depends on (uncentered) second moments, it is convenient to switch from the full  $W_2$  geometry on probability measures to its induced geometry on second-moment matrices. In particular, we turn to the *Bures distance*, defined between positive semidefinite  $\mathbf{A}, \mathbf{B} \succeq \mathbf{0}$  by

$$\mathbf{B}(\mathbf{A}, \mathbf{B})^2 := \text{tr}(\mathbf{A}) + \text{tr}(\mathbf{B}) - 2 \text{tr}((\mathbf{A}^{\frac{1}{2}} \mathbf{B} \mathbf{A}^{\frac{1}{2}})^{\frac{1}{2}}).$$

We refer the reader to Bhatia et al. (2019) for a recent survey of this metric and recall several of its important properties below, deferring proofs to Appendix B. We often use an equivalent formulation of the Bures distance as the minimum Frobenius norm between square-root factors of  $\mathbf{A}$  and  $\mathbf{B}$ .

**Lemma 1.1** (Procrustes formulation of B). *For all  $\mathbf{A}, \mathbf{B} \in \mathbb{S}_+^d$ , we have*

$$\mathbf{B}(\mathbf{A}, \mathbf{B}) = \min_{\mathbf{U} \in \mathbb{O}^d} \|\mathbf{A}^{\frac{1}{2}} - \mathbf{B}^{\frac{1}{2}} \mathbf{U}\|_F = \min_{\substack{\mathbf{S}, \mathbf{T} \in \mathbb{R}^{d \times d} \\ \mathbf{A} = \mathbf{S}\mathbf{S}^\top, \mathbf{B} = \mathbf{T}\mathbf{T}^\top}} \|\mathbf{S} - \mathbf{T}\|_F.$$

Thus,  $\mathbf{B}(\mathbf{A}, \mathbf{B}) \leq \varepsilon$  if and only if there exists  $\Delta \in \mathbb{R}^{d \times d}$  with  $\|\Delta\|_F \leq \varepsilon$  such that  $\mathbf{A} = (\mathbf{B}^{1/2} + \Delta)(\mathbf{B}^{1/2} + \Delta)^\top$ . The minimum is attained by  $\mathbf{U} = \mathbf{I}_d$  if  $\mathbf{A}$  and  $\mathbf{B}$  are simultaneously diagonalizable.

More importantly, this metric is precisely that induced by  $W_2$  on the space of second-moment matrices. Namely, setting  $\mathcal{B}_\varepsilon(\Sigma) := \{\Sigma' \in \mathbb{S}_+^d : \mathbf{B}(\Sigma', \Sigma) \leq \varepsilon\}$ , we have the following.

**Lemma 1.2** ( $W_2$ –Bures equivalence). *For all  $P \in \mathcal{P}_2(\mathbb{R}^d)$  and  $\varepsilon \geq 0$ , we have*

$$\{\Sigma_Q : Q \in \mathcal{B}_\varepsilon(P)\} = \mathcal{B}_\varepsilon(\Sigma_P) = \left\{(\Sigma_P^{1/2} + \Delta)(\Sigma_P^{1/2} + \Delta)^\top : \|\Delta\|_F \leq \varepsilon\right\}.$$

Throughout, we refer to  $\mathbb{S}_+^d$  equipped with  $\mathbf{B}$  as *Bures space*.

## 2 Local Asymptotic Complexity and the Condition Number

We begin by exploring the fundamental limits of robust estimation in the asymptotic  $\varepsilon \rightarrow 0$  regime. Recall the local minimax robust regret  $R_\varepsilon^*(P)$ , defined in (1). In general, this quantity depends on the complex interaction between  $P$ ,  $R$ , and the  $W_2$  geometry. However, as  $\varepsilon \rightarrow 0$ , we show that the complexity of robust prediction is dictated by the following limiting quantity, which we compute in closed form under a mild non-degeneracy assumption.

**Definition 2.1.** Define the *condition number* of  $P \in \mathcal{P}^*$  by  $\kappa_P := \limsup_{\varepsilon \downarrow 0} \varepsilon^{-2} R_\varepsilon^*(P)$ . For convenience, we set  $\kappa_P := \infty$  for  $P \notin \mathcal{P}^*$ .<sup>2</sup>

**Theorem 2.2** (Condition number formula). *If  $P \in \mathcal{P}^*$ , then  $R_\varepsilon^*(P) = \kappa_P \varepsilon^2 + o(\varepsilon^2)$  as  $\varepsilon \downarrow 0$  and*

$$\kappa_P = 1 + \lambda_{\max}(\mathbf{A}_P \mathbf{A}_P^\top + \lambda_{\min}(\Sigma_X^P)^{-1} \Sigma_E^P).$$

<sup>2</sup>This definition and our resulting non-asymptotic analysis can be adapted to the setting where  $\Sigma_X^P$  is singular but the condition number limit is finite (which occurs when the optimal risk is 0). We omit this edge case for ease of presentation.

This implies the useful bounds

$$1 + \max \left\{ \|\mathbf{A}_P\|_{\text{op}}^2, \frac{\lambda_{\max}(\Sigma_E^P)}{\lambda_{\min}(\Sigma_X^P)} \right\} \leq \kappa_P \leq 1 + \|\mathbf{A}_P\|_{\text{op}}^2 + \frac{\lambda_{\max}(\Sigma_E^P)}{\lambda_{\min}(\Sigma_X^P)}.$$

We derive this formula in Appendix E.1 by separately considering perturbations to the feature coordinates  $X$  and the response coordinates  $Y$ . The additive 1 term arises due to response perturbations which leave  $X$  fixed, while the remaining eigenvalue term results from feature perturbations which leave  $Y$  fixed. This latter term is determined by both the size of the optimal coefficient matrix and the scale of the residual variation compared to that of the features.

The proof equates  $\kappa_P$  to a second-order sensitivity of  $R(\mathbf{A}_Q, P)$  as  $Q$  varies in the neighborhood of  $P$ . In Bures space, this amounts to a squared operator norm of a suitable Euclidean derivative, which we compute with some calculus and a scalarization trick that simplifies tensor calculations.

### 3 Non-Asymptotic Regret Bounds

Although  $\kappa$  is defined by an  $\varepsilon \rightarrow 0$  limit, we seek regret bounds which hold for all  $\varepsilon > 0$ . To derive these, we prove a key lemma allowing one to transfer regret bounds between nearby distributions.

**Lemma 3.1** (Regret stability). *For all  $\mathbf{A} \in \mathcal{A}$ ,  $P \in \mathcal{P}^*$ , and  $Q \in \mathcal{P}_2(\mathbb{R}^d)$ , we have*

$$\left| \sqrt{R(\mathbf{A}, P)} - \sqrt{R(\mathbf{A}, Q)} \right| \leq (2\sqrt{\kappa_P} + \|\mathbf{A}\|_{\text{op}})W_2(P, Q).$$

Moreover, for  $\mathbf{A} = \mathbf{A}_P$ , we have  $R(\mathbf{A}_P, Q) \leq 4\kappa_P W_2(P, Q)^2$ .

The proof in Appendix F.1 performs some finite-difference error analysis using a closed form for the regret in terms of Euclidean projections, valid even for degenerate  $Q$ . Here, (a variant of) the Davis–Kahan–Wedin sine theorem (Wedin, 1972) serves as a key technical tool, controlling the Frobenius discrepancy between projection matrices onto the row spaces of nearby matrices.

Regret stability implies several important regret bounds for our main estimators.

**Lemma 3.2** (Stability implications). *Let  $P \in \mathcal{P}^*$  and  $\varepsilon > 0$ . We then have*

$$R_\varepsilon(\text{MRO}_\varepsilon, P) \lesssim \kappa_P \varepsilon^2, \quad R_\varepsilon(\text{SNR}_\varepsilon, P) \lesssim \kappa_P \varepsilon^2.$$

In contrast, our ERM and DRO guarantees adapt to the largest condition number in a ball around  $P$ :

$$R_\varepsilon(\text{ERM}, P) \lesssim \sup_{Q \in \mathcal{B}_\varepsilon(P)} \kappa_Q \varepsilon^2 \quad R_\varepsilon(\text{DRO}_\varepsilon, P) \lesssim \sup_{Q \in \mathcal{B}_{2\varepsilon}(P)} \kappa_Q \varepsilon^2$$

If  $\varepsilon^2 \leq \lambda_{\min}(\Sigma_X^P)/16$ , then these two upper bounds are at most  $O(\kappa_P \varepsilon^2)$ .

We give the proof for MRO and SNR below, deferring the ERM and DRO arguments to Appendix F.2.

*Proof of Lemma 3.2 (MRO and SNR).* Let  $Q \in \mathcal{B}_\varepsilon(P)$  and fix source–target pair  $(P_s, P_t)$  equal to either  $(P, Q)$  or  $(Q, P)$ . For MRO, we set  $\hat{\mathbf{A}} := \text{MRO}_\varepsilon[P_s]$  and bound

$$R(\hat{\mathbf{A}}, P_t) \leq \sup_{Q' \in \mathcal{B}_\varepsilon(P_s)} R(\hat{\mathbf{A}}, Q') \leq \sup_{Q' \in \mathcal{B}_\varepsilon(P_s)} R(\mathbf{A}_P, Q') \leq 16\kappa_P \varepsilon^2.$$

The first inequality uses  $P_t \in \mathcal{B}_\varepsilon(P_s)$ , the second uses MRO optimality, and the last uses  $\mathcal{B}_\varepsilon(P_s) \subseteq \mathcal{B}_{2\varepsilon}(P)$  and Lemma 3.1. Supremizing over  $Q$  and source–target pairs gives  $R_\varepsilon(\text{MRO}_\varepsilon, P) \lesssim \kappa_P \varepsilon^2$ .

For SNR, take instead  $\hat{\mathbf{A}} := \text{SNR}_\varepsilon[P_s]$ . Since risk and regret differ by a constant independent of  $\mathbf{A}$ ,

$$R(\hat{\mathbf{A}}, P_s) + \varepsilon^2 \|\hat{\mathbf{A}}\|_{\text{op}}^2 \leq R(\mathbf{A}_P, P_s) + \varepsilon^2 \|\mathbf{A}_P\|_{\text{op}}^2 \leq 5\kappa_P \varepsilon^2,$$

where the first inequality uses SNR optimality and the second follows from  $P_s \in \mathcal{B}_\varepsilon(P)$ , Lemma 3.1, and  $\|\mathbf{A}_P\|_{\text{op}}^2 \leq \kappa_P$  (by Theorem 2.2). Applying Lemma 3.1 again yields

$$\begin{aligned} \sqrt{R(\hat{\mathbf{A}}, P_t)} &\leq \sqrt{R(\hat{\mathbf{A}}, P_s)} + (\|\hat{\mathbf{A}}\|_{\text{op}} + 2\sqrt{\kappa_P})\varepsilon \\ &\leq \sqrt{2} \sqrt{R(\hat{\mathbf{A}}, P_s) + \varepsilon^2 \|\hat{\mathbf{A}}\|_{\text{op}}^2} + 2\sqrt{\kappa_P} \varepsilon \leq 6\sqrt{\kappa_P} \varepsilon. \end{aligned}$$

The second inequality is Cauchy–Schwarz, and the last uses the objective comparison above. Squaring this SNR bound and supremizing over  $Q$  and source–target pairs gives  $R_\varepsilon(\text{SNR}_\varepsilon, P) \lesssim \kappa_P \varepsilon^2$ .  $\square$

By combining the DRO guarantee above with existing DRO regularizer bounds (Gao et al., 2024; Lee and Raginsky, 2018), we obtain an all-radius guarantee adapted to the residual effective rank. Writing  $r_{\text{eff}}(\mathbf{M}) := \text{tr}(\mathbf{M})/\|\mathbf{M}\|_{\text{op}}$  for  $\mathbf{M} \succeq \mathbf{0}$ , with  $r_{\text{eff}}(\mathbf{0}) := 0$ , we prove the following in Appendix F.3.

**Lemma 3.3** (Rank-adaptive DRO bound). *For  $P \in \mathcal{P}^*$  and  $\varepsilon > 0$ , set  $r := r_{\text{eff}}(\Sigma_E^P)$ . Then*

$$R_\varepsilon(\text{DRO}_\varepsilon, P) \lesssim \left( \sqrt{\min\{p, r\}}(\kappa_P - 1) + \sqrt{r(\kappa_P - 1)} + \kappa_P \right) \varepsilon^2.$$

In particular, DRO achieves regret  $O(\kappa_P \varepsilon^2)$  whenever the residual second moment has effective rank  $O(1)$ . In the worst case, however, the effective rank can be as large as  $m$ .

Combining the above, we derive worst-case rates for the family  $\mathcal{G}(\kappa) = \{P \in \mathcal{P}^* : \kappa_P \leq \kappa\}$ .

**Theorem 3.4** (Population-limit summary rates). *Fix  $\varepsilon > 0$  and  $\kappa \geq 1$ . Then the minimax rate for the class  $\mathcal{G} = \mathcal{G}(\kappa)$  is  $R_\varepsilon^*(\mathcal{G}) \asymp \kappa \varepsilon^2$ , and this is achieved by both  $\text{MRO}_\varepsilon$  and  $\text{SNR}_\varepsilon$ . In contrast, we have  $R_\varepsilon(\text{ERM}, \mathcal{G}) = \infty$  and*

$$R_\varepsilon(\text{DRO}_\varepsilon, \mathcal{G}) \asymp \left( \sqrt{\min\{p, m\}}(\kappa - 1) + \sqrt{m(\kappa - 1)} + \kappa \right) \varepsilon^2.$$

*In particular, if  $\kappa \geq 2$  and  $p \geq m$ , then  $R_\varepsilon(\text{DRO}_\varepsilon, \mathcal{G}) \asymp \sqrt{m} \kappa \varepsilon^2$ .*

In summary, MRO and SNR are always minimax optimal, ERM can suffer unbounded regret, and DRO is provably suboptimal by a dimension-dependent factor for  $\kappa \geq 2$ . The upper bounds in Appendix F.4 apply Lemmas 3.2 and 3.3. The minimax lower bound follows by rescaling the local result from Theorem 2.2; this works because the condition number is invariant to uniform rescaling.

For ERM, we consider a target law  $P$  with identically zero responses. A nearby source law with a nearly collapsed feature and a response correlated with that feature forces the ERM solution to diverge, yielding unbounded target regret even though  $\kappa_P = 1$ . The DRO lower bound requires more care. We employ two constructions with  $\lambda_{\min}(\Sigma_X^P) = \varepsilon^2$  and independent residual noise spread across many response coordinates. In each, the adversary contracts one feature coordinate while leaving the responses unchanged. We then exploit rotational symmetry and a strong duality formula to characterize the DRO minimizers explicitly. The noisy response coordinates determine the minimum value, and we tune the contraction so that attaining this value forces an inflated coefficient on the contracted feature and hence large target regret. This demonstrates how independent residual noise, which contributes an additive constant to target risk and thus cancels from regret, can nevertheless lead to suboptimal DRO decisions.

## 4 Efficient Computation

We next treat computation, given an observed second-moment matrix  $\tilde{\Sigma} \in \mathbb{S}_+^d$  and radius  $\varepsilon > 0$ . Since risk and regret only depend on laws via their second moments, we define  $L(\mathbf{A}, \Sigma) := L(\mathbf{A}, P)$  for any law  $P$  such that  $\Sigma_P = \Sigma$ , and likewise for  $R(\mathbf{A}, \Sigma)$ . By the Bures– $W_2$  equivalence (Lemma 1.2), the optimization problems defining our estimators admit equivalent formulations over second-moment matrices, used below. We defer sampling considerations to Section 5.

ERM admits the canonical minimum-Frobenius-norm solution  $\tilde{\Sigma}_{YX} \tilde{\Sigma}_X^\dagger$ , which can be computed via eigendecomposition of  $\tilde{\Sigma}_X$ . Alternatively, one may apply any standard convex solver to the objective  $\mathbf{A} \mapsto L(\mathbf{A}, \tilde{\Sigma}) = \text{tr}(\mathbf{A}^\top \mathbf{A} \tilde{\Sigma}_X) - 2 \text{tr}(\mathbf{A} \tilde{\Sigma}_{YX}) + \text{tr}(\tilde{\Sigma}_Y)$ .

Shafieezadeh-Abadeh et al. (2018) developed a convex reformulation for the DRO problem, accompanied by a Frank–Wolfe method with convergence guarantees (see Theorem 3.3 therein). The convex reformulation naïvely requires solving a large semidefinite program, so their Frank–Wolfe method is a noted improvement, albeit with some dependencies on  $\lambda_{\min}(\tilde{\Sigma})^{-1}$  in their convergence rate. We do not seek any further improvements, given the stronger performance of MRO and SNR.

The MRO objective  $\mathbf{A} \mapsto \sup_{\Sigma \in \mathcal{B}_\varepsilon(\tilde{\Sigma})} R(\mathbf{A}, \Sigma)$  is convex as the pointwise supremum of convex functions in  $\mathbf{A}$ . However, regret is convex in  $\Sigma$ , as the pointwise supremum over  $\mathbf{A}'$  of  $L(\mathbf{A}, \Sigma) - L(\mathbf{A}', \Sigma)$  which is linear in  $\Sigma$ . This appears problematic, since the inner problem requires maximizing over  $\Sigma$ . Nevertheless, we show that it admits a convex reformulation. Here  $\text{tr}_p(\mathbf{H})$  denotes the sum of the  $p$  largest eigenvalues of a symmetric matrix  $\mathbf{H}$ , counted with multiplicity.

**Theorem 4.1** (Convex reformulation for MRO). *Let  $\varepsilon > 0$ ,  $\tilde{\Sigma} \in \mathbb{S}_+^d$ , and  $\tilde{\mathbf{S}} \in \mathbb{R}^{d \times d}$  such that  $\tilde{\Sigma} = \tilde{\mathbf{S}}\tilde{\mathbf{S}}^\top$ . Decompose*

$$\tilde{\mathbf{S}} = \begin{pmatrix} \tilde{\mathbf{S}}_X \\ \tilde{\mathbf{S}}_Y \end{pmatrix}, \quad \tilde{\mathbf{S}}_X \in \mathbb{R}^{p \times d}, \quad \tilde{\mathbf{S}}_Y \in \mathbb{R}^{m \times d}.$$

*For  $\mathbf{A} \in \mathcal{A}$ , consider the MRO objective  $F(\mathbf{A}) := \sup_{\Sigma \in \mathcal{B}_\varepsilon(\tilde{\Sigma})} R(\mathbf{A}, \Sigma)$  and the residual factor  $\tilde{\mathbf{E}}_{\mathbf{A}} := \tilde{\mathbf{S}}_Y - \mathbf{A}\tilde{\mathbf{S}}_X \in \mathbb{R}^{m \times d}$ . Define the extended-real-valued function  $D : \mathcal{A} \times \mathbb{R} \rightarrow \mathbb{R} \cup \{+\infty\}$  by*

$$D(\mathbf{A}, \eta) := \eta \left( \varepsilon^2 - \|\tilde{\mathbf{S}}_X\|_F^2 + \text{tr}_p \left( \tilde{\mathbf{S}}_X^\top \tilde{\mathbf{S}}_X + \tilde{\mathbf{E}}_{\mathbf{A}}^\top ((\eta - 1)\mathbf{I}_m - \mathbf{A}\mathbf{A}^\top)^{-1} \tilde{\mathbf{E}}_{\mathbf{A}} \right) \right)$$

*if  $\eta > 1 + \|\mathbf{A}\|_{\text{op}}^2$  and  $+\infty$  otherwise. Then  $D$  is convex as an extended-value function on  $\mathcal{A} \times \mathbb{R}$  and  $F(\mathbf{A}) = \inf_{\eta \in \mathbb{R}} D(\mathbf{A}, \eta)$ . Consequently, MRO admits the lifted convex reformulation*

$$\inf_{\mathbf{A} \in \mathcal{A}} F(\mathbf{A}) = \inf_{\mathbf{A} \in \mathcal{A}} \sup_{\Sigma \in \mathcal{B}_\varepsilon(\tilde{\Sigma})} R(\mathbf{A}, \Sigma) = \inf_{\mathbf{A} \in \mathcal{A}, \eta \in \mathbb{R}} D(\mathbf{A}, \eta).$$

This proof in Section G.1 extends a known dual formulation for the inner DRO problem

$$\sup_{\Sigma \in \mathcal{B}_\varepsilon(\tilde{\Sigma})} L(\mathbf{A}, \Sigma) = \inf_{\eta > 1 + \|\mathbf{A}\|_{\text{op}}^2} \eta \left( \varepsilon^2 + \text{tr} \left( \tilde{\mathbf{E}}_{\mathbf{A}}^\top ((\eta - 1)\mathbf{I}_m - \mathbf{A}\mathbf{A}^\top)^{-1} \tilde{\mathbf{E}}_{\mathbf{A}} \right) \right), \quad (2)$$

which underlies the Frank–Wolfe implementation in Shafieezadeh-Abadeh et al. (2018). In both cases, the variable  $\eta$  is introduced as a Lagrange multiplier for the Bures ball constraint. However, our route to the MRO formula involves an additional maximization problem over rank- $p$  projection matrices. Initially outside the infimum, we employ a non-convex minimax theorem to bring the supremum inside, with the  $\text{tr}_p$  term arising due to the Ky Fan variational principle.

Theorem 4.1 suggests solving the MRO problem by applying a subgradient method to minimize  $D$  over  $\mathbf{A}, \eta$ . Indeed, subgradients can be computed via eigendecomposition, due to the Ky Fan variational formula for  $\text{tr}_p$  and Danskin’s theorem. Moreover, if  $\hat{\mathbf{A}}, \hat{\eta}$  are a feasible,  $\delta$ -approximate minimizing pair, then  $\hat{\mathbf{A}}$  is a  $\delta$ -approximate MRO solution. We defer a complete convergence analysis to future work, since the minimax rate of MRO is shared by the much simpler SNR, explored next.

In contrast, the SNR objective  $G(\mathbf{A}) := L(\mathbf{A}, \tilde{\Sigma}) + \varepsilon^2 \|\mathbf{A}\|_{\text{op}}^2$  is convex and enjoys simpler subgradient computations. Indeed, for unit singular vectors  $u \in \mathbb{R}^m, v \in \mathbb{R}^p$  satisfying  $\mathbf{A}v = \|\mathbf{A}\|_{\text{op}}u$ , we have<sup>3</sup>

$$2\mathbf{A}\tilde{\Sigma}_X - 2\tilde{\Sigma}_Y X + 2\varepsilon^2 \|\mathbf{A}\|_{\text{op}} uv^\top \in \partial G(\mathbf{A}).$$

Some simple bounds on the optimizer norm yield the following for a projected subgradient method.

**Proposition 4.2** (SNR subgradient convergence). *Let  $\tilde{\Sigma} \in \mathbb{S}_+^d$  and  $\varepsilon > 0$ . If  $\tilde{\Sigma}_Y = \mathbf{0}$ , then  $\mathbf{0}$  minimizes  $G$ . Otherwise, define radius  $B := (d \text{tr} \tilde{\Sigma}_Y)^{1/2} / \varepsilon$ , Lipschitz bound  $M := 2(B\|\tilde{\Sigma}_X\|_{\text{op}} + \|\tilde{\Sigma}_{YX}\|_F + \varepsilon^2 B)$ , and iteration count  $T := \lceil M^2 B^2 \delta^{-2} \rceil$ , for accuracy  $\delta > 0$ . Consider the update  $\mathbf{A}_{t+1} \leftarrow \text{proj}_B(\mathbf{A}_t - \lambda \mathbf{G}_t)$ , where  $\mathbf{A}_1 = \mathbf{0}$ ,  $\mathbf{G}_t \in \partial G(\mathbf{A}_t)$ ,  $\lambda := B/(M\sqrt{T})$ , and  $\text{proj}_B$  is the Euclidean projection onto the Frobenius ball of radius  $B$  centered at  $\mathbf{0}$ . Then, the averaged iterate  $\bar{\mathbf{A}}_T = T^{-1} \sum_{t=1}^T \mathbf{A}_t$  achieves  $G(\bar{\mathbf{A}}_T) \leq \min_{\mathbf{A}} G(\mathbf{A}) + \delta$ .*

The proof in Section G.2 is routine. Hence, we can solve SNR to accuracy  $\delta$  with  $O(\delta^{-2})$  iterations, each computing a top singular vector. To accelerate this rate, we use that our nonsmooth penalty admits an efficiently computable proximal operator, enabling the application of FISTA (Beck and Teboulle, 2009). In Section G.3, we prove an improved rate for the resulting Algorithm 1.<sup>4</sup>

**Proposition 4.3** (SNR FISTA convergence). *Fix  $\tilde{\Sigma} \in \mathbb{S}_+^d$ ,  $\varepsilon > 0$ , and  $\delta > 0$ . Then Algorithm 1 achieves  $G(\mathbf{A}_T) \leq \min_{\mathbf{A}} G(\mathbf{A}) + \delta$  after  $T = \lceil (4d\lambda_{\max}(\tilde{\Sigma}_X) \text{tr}(\tilde{\Sigma}_Y))^{1/2} \varepsilon^{-1} \cdot \delta^{-1/2} \rceil$  iterations.*

In particular, FISTA achieves accuracy  $\delta$  using  $O(\delta^{-1/2})$  full singular value decompositions.

<sup>3</sup>The first two terms for this subgradient representative are obtained directly by differentiating  $L(\cdot, \tilde{\Sigma})$ . The last term follows, for example, from the characterization of the spectral-norm subdifferential set in Example 2.8 of Lewis (1995).

<sup>4</sup>Faster linear rates are available if  $L(\cdot, \tilde{\Sigma})$  is strongly convex, but this requires  $\lambda_{\min}(\tilde{\Sigma}_X)$  to be bounded away from 0, which is not guaranteed.

---

**Algorithm 1** FISTA for SNR

---

**Require:** Observed second moments  $\tilde{\Sigma}_X \in \mathbb{R}^{p \times p}$ ,  $\tilde{\Sigma}_{YX} \in \mathbb{R}^{m \times p}$ , penalty  $\varepsilon > 0$ , iteration count  $T$

- 1: Compute  $M = 2\lambda_{\max}(\tilde{\Sigma}_X)$
- 2: **if**  $M = 0$  **then return**  $\mathbf{A}_T = \mathbf{0}$
- 3: Set  $\alpha = \varepsilon^2/M$ , initialize  $\mathbf{A}_0 = \mathbf{0}_{m \times p}$ ,  $\mathbf{Y}_0 = \mathbf{A}_0$ ,  $t_0 = 1$
- 4: **for**  $k = 0, 1, \dots, T-1$  **do**
- 5:    $\mathbf{B}_k = \mathbf{Y}_k - \frac{2}{M}(\mathbf{Y}_k \tilde{\Sigma}_X - \tilde{\Sigma}_{YX})$  ▷ gradient step for the smooth loss
- 6:    $\mathbf{A}_{k+1} = \text{ProxSpecSq}(\mathbf{B}_k, \alpha)$  ▷ proximal step
- 7:    $t_{k+1} = (1 + \sqrt{1 + 4t_k^2})/2$
- 8:    $\mathbf{Y}_{k+1} = \mathbf{A}_{k+1} + \frac{t_k - 1}{t_{k+1}}(\mathbf{A}_{k+1} - \mathbf{A}_k)$  ▷ Nesterov update
- 9: **return**  $\mathbf{A}_T$

---



---

**Algorithm 2** ProxSpecSq: proximal map of  $\alpha \|\cdot\|_{\text{op}}^2$ 


---

**Require:** Matrix  $\mathbf{B} \in \mathbb{R}^{m \times p}$ , parameter  $\alpha > 0$

- 1: Compute thin SVD  $\mathbf{B} = \mathbf{P} \text{diag}(\sigma_1, \dots, \sigma_q) \mathbf{Q}^\top$ ,  $\sigma_1 \geq \dots \geq \sigma_q \geq 0$ ,  $q = \min\{m, p\}$
- 2: **if**  $\sigma_1 = 0$  **then return**  $\mathbf{B}$
- 3: Define  $\sigma_{q+1} = 0$ . Find the unique  $k_* \in \{1, \dots, q\}$  such that
$$\sigma_{k_*} > \tau_{k_*} \geq \sigma_{k_*+1}, \quad \tau_k = \sum_{i=1}^k \sigma_i / (k + 2\alpha).$$
- 4: Set  $\tau = \tau_{k_*}$  and define clipped singular values  $s_i = \min\{\sigma_i, \tau\}$  for  $i \in \{1, \dots, q\}$
- 5: **return**  $\mathbf{P} \text{diag}(s_1, \dots, s_q) \mathbf{Q}^\top$

---

## 5 Finite-Sample Statistics

We now turn to the finite-sample setting. Let  $Z_1, \dots, Z_n \stackrel{\text{i.i.d.}}{\sim} P$  and  $P_n := n^{-1} \sum_{i=1}^n \delta_{Z_i}$ . We treat two models, each specifying an observed empirical source measure  $P_{s,n}$  and an unknown population target  $P_t$ . In the *sample-corruption model*,  $P_t = P$  and  $P_{s,n} = n^{-1} \sum_{i=1}^n \delta_{\tilde{Z}_i}$  for an adversarially perturbed sample  $\{\tilde{Z}_i\}_{i=1}^n$  only guaranteed to satisfy  $n^{-1} \sum_{i=1}^n \|\tilde{Z}_i - Z_i\|^2 \leq \varepsilon^2$ . Naturally, this implies  $W_2(P_{s,n}, P_n) \leq \varepsilon$ . In the *distribution-shift model*, the learner instead observes  $P_{s,n} = P_n$ , and  $P_t = Q$  for a perturbed population law  $Q \in \mathcal{B}_\varepsilon(P)$ . Let  $\mathbb{M}_{\varepsilon,n}(P)$  denote the family of joint laws of  $(P_{s,n}, P_t)$  arising from either model, over all admissible, possibly randomized, sample-corruption rules and all fixed  $Q \in \mathcal{B}_\varepsilon(P)$ . Seeking robust, high-probability regret bounds, we thus define

$$R_{\varepsilon,n,\delta}(\hat{\mathbf{A}}_n, P) := \sup_{\mathbb{P} \in \mathbb{M}_{\varepsilon,n}(P)} \inf \left\{ \tau \geq 0 : \mathbb{P} \left( R(\hat{\mathbf{A}}_n[P_{s,n}], P_t) \leq \tau \right) \geq 1 - \delta \right\},$$

where  $\delta \in (0, 1)$  and  $\hat{\mathbf{A}}_n$  is any (measurable) estimator mapping  $P_{s,n}$  to  $\hat{\mathbf{A}}_n[P_{s,n}] \in \mathcal{A}$ .

To control empirical convergence of second-moment matrices, we introduce the *Bures variance proxy*, a rescaled fourth-moment bound given by

$$\mathbb{V}_B(P) := \mathbb{E}_P \left[ \left\| (ZZ^\top - \Sigma_P) \Sigma_P^{\dagger/2} \right\|_F^2 \right] = \mathbb{E}_P [\|Z\|^2 Z^\top \Sigma_P^\dagger Z] - \text{tr}(\Sigma_P). \quad (3)$$

Writing  $\mathcal{P}_4(\mathbb{R}^d)$  for the set of laws with finite fourth moments, we prove the following in Section H.1.

**Proposition 5.1** ( $\mathbb{V}_B$  properties). *Let  $P \in \mathcal{P}_4(\mathbb{R}^d)$ . Then  $\mathbb{V}_B(P) < \infty$ , and  $\mathbb{V}_B(P) = 0$  if and only if  $ZZ^\top = \Sigma_P$  for almost all  $Z \sim P$ . If  $\Sigma_P \succ 0$ , then  $\mathbb{V}_B(P) \leq \mathbb{E}_P[\|Z\|^4] / \lambda_{\min}(\Sigma_P)$ . For Gaussian  $P = \mathcal{N}(0, \Sigma_P)$ , we have  $\mathbb{V}_B(P) = (\text{rank}(\Sigma_P) + 1) \text{tr}(\Sigma_P)$ , and if further  $P \in \mathcal{P}^*$  is non-degenerate, then  $\mathbb{V}_B(P) \geq p \text{tr}(\Sigma_P^P)$ . Crucially, we have  $\mathbb{E}[\mathbb{B}(\Sigma_P, \Sigma_{P_n})^2] \leq \mathbb{V}_B(P)/n$ .*

Hence, we can extend population-limit results to the finite-sample setting by treating sampling as an adversarial Bures perturbation and slightly inflating  $\varepsilon$ . However, this approach uses Markov's inequality, leading to sample complexity with inverse polynomial dependence on the failure probability. Alternatively, sub-Gaussian concentration controls the empirical condition number directly,

avoiding radius inflation. For  $K > 0$ , recall that a random vector  $W$  on  $\mathbb{R}^d$  is  $K$ -sub-Gaussian if  $\mathbb{E}[\exp((u^\top W)^2/K^2)] \leq 2$  for each unit vector  $u \in \mathbb{R}^d$ . We prove the following in Section H.2.

**Theorem 5.2** (Finite-sample regret bounds). *Let  $P \in \mathcal{P}^*$ ,  $\varepsilon > 0$ , and  $\delta \in (0, 1)$ . Suppose first that  $P \in \mathcal{P}_4(\mathbb{R}^d)$ , and fix  $V \geq V_B(P)$ . Then, for sample size  $n \geq V\delta^{-1}\varepsilon^{-2}$  and expanded radius  $\varepsilon_n := \varepsilon + \sqrt{V/(n\delta)}$ , we have*

$$R_{\varepsilon,n,\delta}(\text{MRO}_{\varepsilon_n}, P), R_{\varepsilon,n,\delta}(\text{SNR}_{\varepsilon_n}, P) \lesssim \kappa_P \varepsilon^2.$$

*Alternatively, let  $(X, Y) \sim P$ , write  $E = Y - \mathbf{A}_P X$ , and suppose that the whitened random vector  $((\Sigma_X^P)^{-1/2} X, (\Sigma_E^P)^{\dagger/2} E)$  is  $K$ -sub-Gaussian. Then, for  $n \gtrsim K^4(d + \log \frac{1}{\delta})(1 + \text{tr}(\Sigma_E^P)\varepsilon^{-2})$ ,*

$$R_{\varepsilon,n,\delta}(\text{MRO}_\varepsilon, P), R_{\varepsilon,n,\delta}(\text{SNR}_\varepsilon, P) \lesssim \kappa_P \varepsilon^2.$$

Notably, the sub-Gaussian guarantee improves the polynomial dependence on  $\delta^{-1}$  of the first sample-size bound to logarithmic, and it requires no radius inflation. Still, these sample complexities are not comparable in general, and we defer the question of optimal sample complexity to future work.

## 6 Experiments

Finally, we implement our estimators and compare their performance on synthetic finite-sample instances. All methods use the empirical source second-moment matrix  $\Sigma_{P_{s,n}}$ . For ERM, we compute the closed-form solution using NumPy. For MRO, we implement the lifted objective  $D(\mathbf{A}, \eta)$  from Theorem 4.1 in PyTorch, evaluating its  $\text{tr}_p$  term by eigendecomposition and computing subgradients by automatic differentiation. We use the same template for DRO, replacing  $D(\mathbf{A}, \eta)$  with the simpler dual objective in (2). For both methods, we use the normalized subgradient method with  $k$ th iteration step size  $k^{-1/2}$ , initialize  $\mathbf{A}$  at the ERM solution and  $\eta = 2 + \|\mathbf{A}\|_{\text{op}}^2$ , and repair feasibility after each step by taking  $\eta \leftarrow \max\{\eta, 1 + \|\mathbf{A}\|_{\text{op}}^2 + \zeta\}$  for a numerical buffer  $\zeta = 0.001$ . We return the final iterate after  $T = 800$  iterations. For SNR, we implement the FISTA method in Algorithm 1 using NumPy, with the same ERM initialization and  $T = 60$  iterations. The population benchmarks discussed below served as numerical diagnostics for selecting iteration counts and tolerances.

In Fig. 1, we consider two experimental settings based on the Lemma F.10 construction. Both use corruption budget  $\varepsilon = 1$  and dimensions  $p = m = d/2 \geq 2$ . We take target measure  $P = P_t$  as the law of  $(X, Y)$ , where  $X \sim \mathcal{N}(0, \mathbf{I}_m)$  and  $Y = X + U$ , and  $U \sim \mathcal{N}(0, \text{diag}(0, \mathbf{I}_{m-1}))$  is independent of  $X$ . Then  $\mathbf{A}_P = \Sigma_X^P = \mathbf{I}_m$ ,  $\Sigma_E^P = \text{diag}(0, \mathbf{I}_{m-1})$ , and hence  $\kappa_P = 3$ . Target population regret is therefore  $R(\mathbf{A}, P) = \|\mathbf{A} - \mathbf{I}_m\|_F^2$ . Drawing  $\{(X^{(i)}, Y^{(i)})\}_{i=1}^n$  i.i.d. from  $P$ , we replace each  $X_1^{(i)}$  by  $m^{-1/4} X_1^{(i)}$ , stopping if needed before the first replacement that would exceed a total squared perturbation budget of  $n\varepsilon^2$ . All other feature coordinates, responses, and unselected observations remain unchanged. Writing  $\tilde{X}^{(i)}$  for the resulting features, the observed empirical source  $P_{s,n} = n^{-1} \sum_{i=1}^n \delta_{(\tilde{X}^{(i)}, Y^{(i)})}$  satisfies  $W_2(P_{s,n}, P_n)^2 \leq \frac{1}{n} \sum_{i=1}^n \|\tilde{X}^{(i)} - X^{(i)}\|^2 \leq \varepsilon^2$ , obeying our Section 5 sample-corruption model. We use sample size  $n = 10m^2$  to match, for fixed confidence, the  $d(1 + \text{tr}(\Sigma_E^P)) \asymp m^2$  dimensional dependence in the sub-Gaussian bound of Theorem 5.2.

We also compute population benchmarks using the source law  $P_s$  of  $(\text{diag}(m^{-1/4}, \mathbf{I}_{m-1})X, Y)$  for  $(X, Y) \sim P$ , corresponding to contracting every observation. This choice is justified because the average cost of contracting every observation converges almost surely to  $(1 - m^{-1/4})^2 < \varepsilon^2$  as  $n \rightarrow \infty$ . The estimators then receive  $\Sigma_{P_s}$  instead of  $\Sigma_{P_{s,n}}$ . Appendix I uses rotational symmetry to reduce the population problems to optimization over the two-parameter family  $\mathbf{A}(\alpha, \beta) = \text{diag}(\alpha, \beta \mathbf{I}_{m-1})$ . Under this reduction, ERM and SNR admit closed-form solutions, while DRO and MRO admit constant-size semidefinite programs which we solve using CVXPY (Diamond and Boyd, 2016; Agrawal et al., 2018). These computations benchmark the estimators at exact population second moments and provide diagnostics for the finite-sample optimization routines. For the population MRO benchmark, we break ties by minimizing the source risk  $L(\cdot, P_s)$ .

For **Setting A**, we supply radius  $\rho = \varepsilon = 1$ , vary  $d$  between 20 and 400, and plot the target population regret  $R(\hat{\mathbf{A}}[P_{s,n}], P)$  for  $\hat{\mathbf{A}} \in \{\text{ERM}, \text{DRO}_\rho, \text{MRO}_\rho, \text{SNR}_\rho\}$ . This setting illustrates the dimension-free behavior of MRO and SNR and  $\sqrt{d}$  growth of DRO regret. For **Setting B**, we fix  $d = 400$  and vary the supplied radius  $\rho \in [0.1\varepsilon, 20\varepsilon]$ . For small  $\rho$ , DRO, MRO, and SNR perform similarly to




Figure 1: Target population regret  $R(\hat{\mathbf{A}}[P_{s,n}], P)$  for Settings A (left) and B (right). Solid curves show finite-sample averages, while dashed curves show population benchmarks. For Setting B, the dashed vertical line marks the corruption budget  $\varepsilon = 1$ .

ERM, and their relative ranking at  $\rho = \varepsilon$  matches Setting A. As the radius  $\rho$  varies, the best method depends on its scale; SNR is strongest for radii below the budget  $\varepsilon$ , MRO is strongest for most larger radii, and for very large radii excessive regularization or robustification causes ERM to dominate. This precise behavior is beyond the scope of our current theoretical results and warrants further study.

For both settings, the plots report target population regret averaged over 10 independent trials, with error bars showing the minimum and maximum over trials, and dashed curves showing the population benchmarks. Error bars are nearly imperceptible due to low variance at the chosen sample size.

See Appendix J for additional details and code, an alternative to Setting A where  $d$  is fixed but the effective rank of  $\Sigma_E^P$  is varied, and a final experiment with real-world network traffic data (Koumar et al., 2025). We emphasize that the observed separations between MRO/SNR and ERM/DRO need not hold universally, with ample room remaining for finer-grained instance-dependent analysis.

## 7 Concluding Remarks

In this paper, we introduced a new framework for regret-optimal Wasserstein-robust decision-making, focused on the setting of least-squares regression. We fully characterized the complexity of robust estimation in the  $\varepsilon \rightarrow 0$  limit, via the condition number, and extended these asymptotic guarantees to all  $\varepsilon > 0$  using regret stability. While DRO is a natural robust algorithm, we proved and verified empirically that it can be outperformed by MRO and our new procedure, SNR, for high dimensional tasks. Moreover, we developed a new convex reformulation for MRO and an efficient proximal gradient method for SNR. Moving forward, we hope that this initial work spurs more research into characterizing and efficiently achieving regret-optimal performance for robust stochastic optimization.

## Acknowledgements and AI Usage Statement

This work was supported as a part of NCCR Automation, a National Centre of Competence in Research, funded by the Swiss National Science Foundation (grant number 51NF40 225155). We also thank the Optimal Transport and Optimization Workshop (OT+OPT '26) attendees for their helpful feedback on an early version of this work.

We used ChatGPT (versions 5-6 Pro) for editing, checking and refining mathematical arguments, and tuning lower bound constructions. Version 5.5 suggested an improvement to Lemma 3.1 that allowed us to extend our analysis from MRO to SNR. It also suggested the argument used in Step 1 of the proof of Theorem 4.1, enabling our convex reformulation of MRO. We have independently verified all mathematical arguments and take full responsibility for the content and correctness of this work.

## References

- Agarwal, A. and Zhang, T. (2022). Minimax regret optimization for robust machine learning under distribution shift. In *Conference on Learning Theory (COLT)*.
- Agrawal, A., Verschueren, R., Diamond, S., and Boyd, S. (2018). A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1):42–60.
- Beck, A. and Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202.
- Bhatia, R. (2009). *Positive Definite Matrices*. Princeton University Press.
- Bhatia, R., Jain, T., and Lim, Y. (2019). On the Bures–Wasserstein distance between positive definite matrices. *Expositiones Mathematicae*, 37(2):165–191.
- Blanchet, J. and Murthy, K. (2019). Quantifying distributional model risk via optimal transport. *Mathematics of Operations Research*, 44(2):565–600.
- Bubeck, S. (2015). Convex optimization: Algorithms and complexity. *Foundations and Trends in Machine Learning*, 8(3–4):231–357.
- Chao, P. and Dobriban, E. (2023). Minimax statistical estimation under Wasserstein contamination. *arXiv preprint arXiv:2308.01853*.
- Chen, Z. and Xie, W. (2021). Regret in the newsvendor model with demand and yield randomness. *Production and Operations Management*, 30(11):4176–4197.
- Cho, Y. and Yang, I. (2024). Wasserstein distributionally robust regret minimization. *IEEE Control Systems Letters*, 8:820–825.
- Diamond, S. and Boyd, S. (2016). CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5.
- Fiechtner, L.-B. and Blanchet, J. (2025). Wasserstein distributionally robust regret optimization. *arXiv preprint arXiv:2504.10796*.
- Fiechtner, L.-B. and Blanchet, J. (2026). Distributionally robust regret optimal LQR with common stage-law ambiguity. In *Proceedings of the 65th IEEE Conference on Decision and Control (CDC)*. To appear; arXiv:2604.06158, version 3.
- Gao, R., Chen, X., and Kleywegt, A. J. (2024). Wasserstein distributionally robust optimization and variation regularization. *Operations Research*, 72(3):1177–1191.
- Gao, R. and Kleywegt, A. (2022). Distributionally robust stochastic optimization with Wasserstein distance. *Mathematics of Operations Research*, 48(2):603–655.
- Horn, R. A. and Johnson, C. R. (1991). *Topics in Matrix Analysis*. Cambridge University Press.
- Jacobsson, S., Durisi, G., Coldrey, M., Goldstein, T., and Studer, C. (2017). Quantized precoding for massive MU-MIMO. *IEEE Transactions on Communications*, 65(11):4670–4684.
- Koumar, J., Hynek, K., Čejka, T., and Šiška, P. (2025). CESNET-TimeSeries24: Time series dataset for network traffic anomaly detection and forecasting. *Scientific Data*, 12(1):338.
- Kuhn, D., Mohajerin Esfahani, P., Nguyen, V. A., and Shafieezadeh-Abadeh, S. (2019). Wasserstein distributionally robust optimization: Theory and applications in machine learning. In *Operations Research & Management Science in the Age of Analytics*, INFORMS TutORials in Operations Research, pages 130–166. INFORMS.
- Kuhn, D., Shafiee, S., and Wiesemann, W. (2025). Distributionally robust optimization. *Acta Numerica*, 34:579–804.
- Lee, J. and Raginsky, M. (2018). Minimax statistical learning with Wasserstein distances. In *Advances in Neural Information Processing Systems (NeurIPS)*.

- Lewis, A. S. (1995). The convex analysis of unitarily invariant matrix functions. *Journal of Convex Analysis*, 2(1–2):173–183.
- Lin, F., Fang, X., and Gao, Z. (2022). Distributionally robust optimization: A review on theory and applications. *Numerical Algebra, Control and Optimization*, 12(1):159–212.
- Liu, Z. and Loh, P.-L. (2023). Robust W-GAN-based estimation under Wasserstein contamination. *Information and Inference*, 12(1):312–362.
- Nguyen, V. A., Kuhn, D., and Mohajerin Esfahani, P. (2022). Distributionally robust inverse covariance estimation: The Wasserstein shrinkage estimator. *Operations Research*, 70(1):490–515.
- Nguyen, V. A., Shafiee, S., Filipović, D., and Kuhn, D. (2021a). Mean-covariance robust risk measurement. *arXiv preprint arXiv:2112.09959*.
- Nguyen, V. A., Shafieezadeh-Abadeh, S., Kuhn, D., and Mohajerin Esfahani, P. (2021b). Bridging Bayesian and minimax mean square error estimation via Wasserstein distributionally robust optimization. *Mathematics of Operations Research*, 48(1):1–37.
- Parikh, N. and Boyd, S. (2014). Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):127–239.
- Perakis, G. and Roels, G. (2008). Regret in the newsvendor model with partial information. *Operations Research*, 56(1):188–203.
- Petersen, K. B. and Pedersen, M. S. (2012). The Matrix Cookbook. Technical report, Technical University of Denmark.
- Pólik, I. and Terlaky, T. (2007). A survey of the S-lemma. *SIAM Review*, 49(3):371–418.
- Rahimian, H. and Mehrotra, S. (2022). Frameworks and results in distributionally robust optimization. *Open Journal of Mathematical Optimization*, 3:1–85. Article 4.
- Ricceri, B. (2019). Minimax theorems in a fully non-convex setting. *Journal of Nonlinear and Variational Analysis*, 3(1):45–52.
- Shafieezadeh-Abadeh, S., Nguyen, V. A., Kuhn, D., and Mohajerin Esfahani, P. (2018). Wasserstein distributionally robust Kalman filtering. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Sinha, A., Namkoong, H., and Duchi, J. C. (2018). Certifying some distributional robustness with principled adversarial training. In *International Conference on Learning Representations (ICLR)*.
- Sion, M. (1958). On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176.
- Taşkesen, B., Iancu, D. A., Koçyiğit, Ç., and Kuhn, D. (2023). Distributionally robust linear quadratic control. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Vershynin, R. (2026). *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge University Press, 2 edition.
- Wedin, P.-Å. (1972). Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12(1):99–111.
- Zhu, B., Jiao, J., and Steinhardt, J. (2022). Generalized resilience and robust statistics. *The Annals of Statistics*, 50(4):2256–2283.

## Appendix Contents

|          |                                                                                |           |
|----------|--------------------------------------------------------------------------------|-----------|
| <b>A</b> | <b>Linear Algebra Preliminaries</b>                                            | <b>14</b> |
| A.1      | Basic Matrix Inequalities and Identities . . . . .                             | 14        |
| A.2      | Scalar, Tensor, and Spectral Variational Facts . . . . .                       | 16        |
| <b>B</b> | <b>Bures–Wasserstein Geometry and Regression</b>                               | <b>17</b> |
| B.1      | Common Block Decompositions . . . . .                                          | 17        |
| B.2      | Regression in Bures Space . . . . .                                            | 17        |
| B.3      | Bures–Wasserstein Geometry . . . . .                                           | 19        |
| <b>C</b> | <b>DRO Duality and Symmetry Reductions</b>                                     | <b>21</b> |
| C.1      | Strong Duality for the Inner DRO Problem . . . . .                             | 21        |
| C.2      | Averaging over Symmetries . . . . .                                            | 23        |
| <b>D</b> | <b>Remaining Proofs for Section 1</b>                                          | <b>25</b> |
| D.1      | Existence of Minimizers (Footnote 1) . . . . .                                 | 25        |
| <b>E</b> | <b>Remaining Proofs for Section 2</b>                                          | <b>25</b> |
| E.1      | Condition Number Formula (Proof of Theorem 2.2) . . . . .                      | 25        |
| <b>F</b> | <b>Remaining Proofs for Section 3</b>                                          | <b>29</b> |
| F.1      | Regret Transfer (Proof of Lemma 3.1) . . . . .                                 | 29        |
| F.2      | Stability Implications for ERM and DRO (Remaining Proof for Lemma 3.2) . . . . | 31        |
| F.3      | Rank-Adaptive DRO Bound (Proof of Lemma 3.3) . . . . .                         | 34        |
| F.4      | Main Population-Limit Regret Bounds (Proof of Theorem 3.4) . . . . .           | 35        |
| <b>G</b> | <b>Remaining Proofs for Section 4</b>                                          | <b>41</b> |
| G.1      | Convex Reformulation for MRO (Proof of Theorem 4.1) . . . . .                  | 41        |
| G.2      | Subgradient Method Convergence for SNR (Proof of Proposition 4.2) . . . . .    | 45        |
| G.3      | FISTA Convergence for SNR (Proof of Proposition 4.3) . . . . .                 | 46        |
| <b>H</b> | <b>Remaining Proofs for Section 5</b>                                          | <b>46</b> |
| H.1      | Properties of the Bures Variance Proxy (Proof of Proposition 5.1) . . . . .    | 46        |
| H.2      | Finite-Sample Regret Bounds (Proof of Theorem 5.2) . . . . .                   | 48        |
| <b>I</b> | <b>Exact Reductions for Population Benchmarks in Section 6</b>                 | <b>52</b> |
| I.1      | Reduction to Scalar Optimization Using Rotational Symmetry . . . . .           | 52        |
| I.2      | Constant-Size SDP Reformulations for DRO and MRO . . . . .                     | 52        |
| I.3      | Closed form for SNR . . . . .                                                  | 55        |
| <b>J</b> | <b>Full Experimental Details and Additional Experiments</b>                    | <b>56</b> |
| J.1      | Additional Synthetic Experiment with Varied Residual Rank . . . . .            | 56        |
| J.2      | Network-Traffic Experiment with CESNET Dataset . . . . .                       | 57        |

## A Linear Algebra Preliminaries

This appendix collects linear algebraic tools used repeatedly in the proofs. We include short derivations when they are more convenient than a citation, and otherwise point to standard references.

**Additional notation.** Write  $\mathbb{S}^n$  for the space of symmetric  $n \times n$  matrices. For a matrix  $\mathbf{M} \in \mathbb{R}^{a \times b}$ , let  $\text{range}(\mathbf{M})$  denote its column space, so that  $\text{range}(\mathbf{M}^\top)$  is its row space, and write  $\text{rank}(\mathbf{M}) := \dim \text{range}(\mathbf{M})$ . Write  $\text{vec}(\mathbf{M}) \in \mathbb{R}^{ab}$  for its column-stacked vectorization and  $\otimes$  for the Kronecker product. For matrices of the same size, we write  $\langle \mathbf{M}, \mathbf{N} \rangle_{\text{F}} := \text{tr}(\mathbf{M}^\top \mathbf{N})$  for their Frobenius inner product. We order the singular values of  $\mathbf{M} \in \mathbb{R}^{a \times b}$  as  $\sigma_1(\mathbf{M}) \geq \dots \geq \sigma_{\min\{a,b\}}(\mathbf{M}) \geq 0$ , including zeros, and write  $\sigma_{\min}(\mathbf{M}) := \sigma_{\min\{a,b\}}(\mathbf{M})$ . For  $\mathbf{H} \in \mathbb{S}^d$  and  $k \in \{1, \dots, d\}$ , let  $\text{tr}_k(\mathbf{H})$  denote the sum of its  $k$  largest eigenvalues and write  $\mathbf{H}_+$  for its positive semidefinite part, obtained by replacing negative eigenvalues by zero. Write  $(t)_+ := \max\{t, 0\}$  for scalars. We write  $0_n$  for the zero vector in  $\mathbb{R}^n$  and  $e_i$  for the  $i$ th standard basis vector, with its ambient dimension specified when needed.

### A.1 Basic Matrix Inequalities and Identities

**Lemma A.1** (Basic norm and projection inequalities). *For dimension-compatible matrices,*

$$\begin{aligned} \|\mathbf{ABC}\|_{\text{F}} &\leq \|\mathbf{A}\|_{\text{op}} \|\mathbf{B}\|_{\text{F}} \|\mathbf{C}\|_{\text{op}}, \\ \|\mathbf{AB}\|_{\text{op}} &\leq \|\mathbf{A}\|_{\text{op}} \|\mathbf{B}\|_{\text{op}}, \\ \|\mathbf{AB}\|_{\text{F}} &\leq \|\mathbf{A}\|_{\text{op}} \|\mathbf{B}\|_{\text{F}}, \\ \|\mathbf{A}\|_{\text{op}} &\leq \|\mathbf{A}\|_{\text{F}}, \quad \|\mathbf{A}\|_{\text{F}} \leq \sqrt{\text{rank}(\mathbf{A})} \|\mathbf{A}\|_{\text{op}}. \end{aligned}$$

If  $\mathbf{\Pi}$  is an orthogonal projection, then

$$\|\mathbf{M}\mathbf{\Pi}\|_{\text{F}} \leq \|\mathbf{M}\|_{\text{F}}, \quad \|\mathbf{M}\|_{\text{F}}^2 = \|\mathbf{M}\mathbf{\Pi}\|_{\text{F}}^2 + \|\mathbf{M}(\mathbf{I} - \mathbf{\Pi})\|_{\text{F}}^2.$$

Finally,

$$|\langle \mathbf{A}, \mathbf{B} \rangle_{\text{F}}| \leq \|\mathbf{A}\|_{\text{F}} \|\mathbf{B}\|_{\text{F}}.$$

*Proof.* The first inequality holds for all unitarily invariant matrix norms, including the Frobenius norm; see, e.g., Horn and Johnson, 1991, Section 3.5, Exercise 4. The second is a standard submultiplicativity property; see, e.g., Horn and Johnson, 1991, Section 3.5, Exercise 3. The third inequality follows from the first by taking  $\mathbf{C} = \mathbf{I}$ . The comparisons between Frobenius and operator norms follow from the singular-value identities  $\|\mathbf{A}\|_{\text{op}} = \sigma_1(\mathbf{A})$  and  $\|\mathbf{A}\|_{\text{F}}^2 = \sum_i \sigma_i(\mathbf{A})^2$ . The projection identities follow from  $\mathbf{\Pi}(\mathbf{I} - \mathbf{\Pi}) = \mathbf{0}$  and the Pythagorean theorem for the Frobenius inner product. The last inequality is Cauchy–Schwarz for the Frobenius inner product.  $\square$

**Lemma A.2** (Pseudoinverse and projection identities). *Let  $\mathbf{M} \in \mathbb{R}^{r \times d}$ . Then  $\mathbf{M}^\dagger \mathbf{M}$  is the orthogonal projection onto the row space of  $\mathbf{M}$ , and  $\mathbf{M}\mathbf{M}^\dagger$  is the orthogonal projection onto the column space of  $\mathbf{M}$ . We always have  $\mathbf{M}^\dagger = \mathbf{M}^\top (\mathbf{M}\mathbf{M}^\top)^\dagger$ . If  $\mathbf{M}$  has full row rank, then*

$$\mathbf{M}^\dagger = \mathbf{M}^\top (\mathbf{M}\mathbf{M}^\top)^{-1}, \quad \mathbf{M}\mathbf{M}^\dagger = \mathbf{I}_r, \quad \|\mathbf{M}^\dagger\|_{\text{op}} = \sigma_{\min}(\mathbf{M})^{-1}.$$

If  $\mathbf{B} \succeq \mathbf{0}$ , then  $\mathbf{B}^\dagger \mathbf{B}^2 = \mathbf{B}$ .

*Proof.* Take a compact singular value decomposition  $\mathbf{M} = \mathbf{U}\mathbf{D}\mathbf{V}^\top$  over its nonzero singular values. Then  $\mathbf{M}^\dagger = \mathbf{V}\mathbf{D}^{-1}\mathbf{U}^\top$ , so  $\mathbf{M}^\dagger \mathbf{M} = \mathbf{V}\mathbf{V}^\top$ ,  $\mathbf{M}\mathbf{M}^\dagger = \mathbf{U}\mathbf{U}^\top$ . These are exactly the claimed row-space and column-space projections. Further,  $\mathbf{M}\mathbf{M}^\top = \mathbf{U}\mathbf{D}^2\mathbf{U}^\top$ , and so

$$\mathbf{M}^\top (\mathbf{M}\mathbf{M}^\top)^\dagger = \mathbf{V}\mathbf{D}\mathbf{U}^\top (\mathbf{U}\mathbf{D}^{-2}\mathbf{U}^\top) = \mathbf{V}\mathbf{D}^{-1}\mathbf{U}^\top = \mathbf{M}^\dagger$$

If  $\mathbf{M}$  has full row rank, then  $\mathbf{M}\mathbf{M}^\top$  is invertible and the displayed formulas follow directly. The claim for  $\mathbf{B} \succeq \mathbf{0}$  follows from the eigenvalue decomposition of  $\mathbf{B}$ .  $\square$

**Lemma A.3** (Schur complement). *Let*

$$\mathbf{Q} = \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^\top & \mathbf{C} \end{pmatrix}$$

be symmetric. If  $\mathbf{A} \succ \mathbf{0}$ , then

$$\mathbf{Q} \succeq \mathbf{0} \iff \mathbf{C} - \mathbf{B}^\top \mathbf{A}^{-1} \mathbf{B} \succeq \mathbf{0}.$$

The analogous statement holds with  $\mathbf{C} \succ \mathbf{0}$  and Schur complement  $\mathbf{A} - \mathbf{B} \mathbf{C}^{-1} \mathbf{B}^\top$ .

*Proof.* Use the congruence

$$\begin{pmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{B}^\top \mathbf{A}^{-1} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^\top & \mathbf{C} \end{pmatrix} \begin{pmatrix} \mathbf{I} & -\mathbf{A}^{-1} \mathbf{B} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} = \begin{pmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \mathbf{C} - \mathbf{B}^\top \mathbf{A}^{-1} \mathbf{B} \end{pmatrix}.$$

Congruence by an invertible matrix preserves positive (semi)definiteness.  $\square$

**Lemma A.4** (Rectangular resolvent identities). *Let  $\mathbf{M} \in \mathbb{R}^{r \times d}$  and  $\eta > \|\mathbf{M}\|_{\text{op}}^2$ . Then*

$$\begin{aligned} (\eta \mathbf{I}_d - \mathbf{M}^\top \mathbf{M})^{-1} \mathbf{M}^\top &= \mathbf{M}^\top (\eta \mathbf{I}_r - \mathbf{M} \mathbf{M}^\top)^{-1}, \\ \mathbf{M} (\eta \mathbf{I}_d - \mathbf{M}^\top \mathbf{M})^{-1} &= (\eta \mathbf{I}_r - \mathbf{M} \mathbf{M}^\top)^{-1} \mathbf{M}, \\ \mathbf{M}^\top \mathbf{M} (\eta \mathbf{I}_d - \mathbf{M}^\top \mathbf{M})^{-1} &= \mathbf{M}^\top (\eta \mathbf{I}_r - \mathbf{M} \mathbf{M}^\top)^{-1} \mathbf{M}, \\ \eta (\eta \mathbf{I}_d - \mathbf{M}^\top \mathbf{M})^{-1} &= \mathbf{I}_d + \mathbf{M}^\top (\eta \mathbf{I}_r - \mathbf{M} \mathbf{M}^\top)^{-1} \mathbf{M}. \end{aligned}$$

*Proof.* We start by computing

$$\mathbf{M}^\top (\eta \mathbf{I}_r - \mathbf{M} \mathbf{M}^\top) = (\eta \mathbf{I}_d - \mathbf{M}^\top \mathbf{M}) \mathbf{M}^\top.$$

The terms in parentheses are both invertible due to our lower bound on  $\eta$ , so suitable left/right multiplication gives the first identity. The second identity follows by taking transposes, and the third follows by left multiplication of the second identity by  $\mathbf{M}^\top$ . The fourth identity follows by combining the third with the identity  $\eta(\eta \mathbf{I} - \mathbf{C})^{-1} = \mathbf{I} + \mathbf{C}(\eta \mathbf{I} - \mathbf{C})^{-1}$ , which can be confirmed by multiplying both expressions on the right by  $\eta \mathbf{I} - \mathbf{C}$ .  $\square$

**Lemma A.5** (Lipschitz continuity of the smallest singular value). *Let  $\mathbf{S}, \bar{\mathbf{S}} \in \mathbb{R}^{r \times d}$ . Then*

$$|\sigma_{\min}(\bar{\mathbf{S}}) - \sigma_{\min}(\mathbf{S})| \leq \|\bar{\mathbf{S}} - \mathbf{S}\|_{\text{op}}.$$

*Proof.* This is the standard singular-value perturbation inequality; see, e.g., Theorem 3.3.16(c) of Horn and Johnson (1991).  $\square$

**Lemma A.6** (Row-space projection perturbation). *Let  $\mathbf{M}_0 \in \mathbb{R}^{r \times d}$  have full row rank, and let  $\mathbf{M}_1 = \mathbf{M}_0 + \mathbf{\Delta}$  for some  $\mathbf{\Delta} \in \mathbb{R}^{r \times d}$ . Let  $\mathbf{\Pi}_0, \mathbf{\Pi}_1 \in \mathbb{R}^{d \times d}$  denote orthogonal projections onto the row spaces of  $\mathbf{M}_0, \mathbf{M}_1$ , respectively. Then,*

$$\|(\mathbf{I}_d - \mathbf{\Pi}_0) \mathbf{\Pi}_1\|_{\text{F}} \leq \|(\mathbf{I}_d - \mathbf{\Pi}_1) \mathbf{\Pi}_0\|_{\text{F}} \leq \frac{\|\mathbf{\Delta}\|_{\text{F}}}{\sigma_{\min}(\mathbf{M}_0)}.$$

*Proof.* Although we have not seen this precise statement in the literature, the left-hand side connects to the principal angles between row spaces and has been well examined, with the closely related Davis–Kahan–Wedin sine theorem providing a similar bound (see, e.g., Section 3 of Wedin, 1972).

For the proof, define the complementary projections  $\mathbf{\Pi}_i^\perp := \mathbf{I}_d - \mathbf{\Pi}_i$  for  $i \in \{0, 1\}$ . To prove the first inequality, we compute

$$\begin{aligned} \|\mathbf{\Pi}_0^\perp \mathbf{\Pi}_1\|_{\text{F}}^2 &= \text{tr}((\mathbf{I}_d - \mathbf{\Pi}_0) \mathbf{\Pi}_1) && (\mathbf{\Pi}_0, \mathbf{\Pi}_1 \text{ projections}) \\ &= \text{tr}(\mathbf{\Pi}_1) - \text{tr}(\mathbf{\Pi}_0 \mathbf{\Pi}_1) \\ &\leq \text{tr}(\mathbf{\Pi}_0) - \text{tr}(\mathbf{\Pi}_0 \mathbf{\Pi}_1) \\ &= \|\mathbf{\Pi}_1^\perp \mathbf{\Pi}_0\|_{\text{F}}^2. \end{aligned}$$

The inequality uses that  $\text{tr}(\mathbf{\Pi}_1) \leq r = \text{tr}(\mathbf{\Pi}_0)$ , since  $\mathbf{M}_0$  has full row rank. Thus, it suffices to bound  $\|\mathbf{\Pi}_1^\perp \mathbf{\Pi}_0\|_{\text{F}}^2$ .

Next, since  $\mathbf{M}_0$  is full row rank, we can take a thin singular value decomposition

$$\mathbf{M}_0 = \mathbf{U} \mathbf{D} \mathbf{V}^\top,$$

where  $\mathbf{U} \in \mathbb{O}^r$  is orthogonal,  $\mathbf{D} \in \mathbb{R}^{r \times r}$  is diagonal with strictly positive entries, and  $\mathbf{V} \in \mathbb{R}^{d \times r}$  has orthonormal columns. By this choice,  $\mathbf{V} = \mathbf{M}_0^\top \mathbf{U} \mathbf{D}^{-1}$  and  $\mathbf{\Pi}_0 = \mathbf{V} \mathbf{V}^\top$ , implying that

$$\|\mathbf{\Pi}_1^\perp \mathbf{\Pi}_0\|_F^2 = \text{tr}(\mathbf{V}^\top \mathbf{\Pi}_1^\perp \mathbf{V}) = \|\mathbf{\Pi}_1^\perp \mathbf{V}\|_F^2.$$

Since  $\mathbf{\Pi}_1^\perp$  is the orthogonal projection onto the orthogonal complement of the row space  $R_1 \subseteq \mathbb{R}^d$  of  $\mathbf{M}_1$ , we have

$$\|\mathbf{\Pi}_1^\perp \mathbf{V}\|_F = \min\{\|\mathbf{V} - \mathbf{W}\|_F : \text{columns of } \mathbf{W} \text{ lie in } R_1\}.$$

Consider the specific competitor  $\mathbf{W} := \mathbf{M}_1^\top \mathbf{U} \mathbf{D}^{-1}$ , whose columns lie in  $R_1$ . We then have

$$\mathbf{V} - \mathbf{W} = (\mathbf{M}_0^\top - \mathbf{M}_1^\top) \mathbf{U} \mathbf{D}^{-1} = -\mathbf{\Delta}^\top \mathbf{U} \mathbf{D}^{-1}.$$

Therefore,

$$\|\mathbf{\Pi}_1^\perp \mathbf{V}\|_F \leq \|\mathbf{\Delta}^\top \mathbf{U} \mathbf{D}^{-1}\|_F \leq \|\mathbf{\Delta}\|_F \|\mathbf{D}^{-1}\|_{\text{op}} = \frac{\|\mathbf{\Delta}\|_F}{\sigma_{\min}(\mathbf{M}_0)}.$$

Combining the steps above, we obtain

$$\|(\mathbf{I}_d - \mathbf{\Pi}_0) \mathbf{\Pi}_1\|_F = \|\mathbf{\Pi}_0^\perp \mathbf{\Pi}_1\|_F \leq \|\mathbf{\Pi}_1^\perp \mathbf{\Pi}_0\|_F = \|\mathbf{\Pi}_1^\perp \mathbf{V}\|_F \leq \frac{\|\mathbf{\Delta}\|_F}{\sigma_{\min}(\mathbf{M}_0)},$$

as desired.  $\square$

## A.2 Scalar, Tensor, and Spectral Variational Facts

**Lemma A.7** (Elementary scalar identities). *For  $a_1, \dots, a_k \geq 0$ ,*

$$\left( \sum_{i=1}^k a_i \right)^2 \leq k \sum_{i=1}^k a_i^2.$$

*For  $A, B \geq 0$ ,*

$$\inf_{\delta > 0} (A + \delta) \left( 1 + \frac{B}{\delta} \right) = (\sqrt{A} + \sqrt{B})^2,$$

*with optimizer  $\delta = \sqrt{AB}$  when  $AB > 0$ . Finally, for any norm  $\|\cdot\|$ , we have*

$$\inf_x \max\{\|x - x_0\|^2, \|x - x_1\|^2\} = \frac{1}{4} \|x_0 - x_1\|^2.$$

*Proof.* The first inequality is Cauchy–Schwarz. For the second identity, expand

$$(A + \delta) \left( 1 + \frac{B}{\delta} \right) = A + B + \delta + \frac{AB}{\delta}.$$

Basic calculus confirms that this is minimized at  $\delta = \sqrt{AB}$  if  $AB > 0$  or otherwise in the limit as  $\delta \rightarrow 0$ . In both cases, the infimum value is  $(\sqrt{A} + \sqrt{B})^2$ , as claimed.

The last identity follows because any  $x$  with both distances at most  $r$  must satisfy  $\|x_0 - x_1\| \leq 2r$ , while the midpoint achieves radius  $\|x_0 - x_1\|/2$ .  $\square$

**Lemma A.8** (Kronecker trace identity). *For matrices  $\mathbf{A}, \mathbf{B}, \mathbf{C}$  with dimensions such that the following products are well-defined,*

$$\text{tr}(\mathbf{A} \mathbf{C} \mathbf{B} \mathbf{C}^\top) = \text{vec}(\mathbf{C})^\top (\mathbf{B}^\top \otimes \mathbf{A}) \text{vec}(\mathbf{C}).$$

*Proof.* This follows from the matrix vectorization identities  $\text{vec}(\mathbf{A} \mathbf{C} \mathbf{B}) = (\mathbf{B}^\top \otimes \mathbf{A}) \text{vec}(\mathbf{C})$  and  $\text{tr}(\mathbf{U} \mathbf{V}^\top) = \text{vec}(\mathbf{V})^\top \text{vec}(\mathbf{U})$ ; see, e.g., Petersen and Pedersen, 2012, Section 10.2.  $\square$

**Lemma A.9** (Variational matrix formulas). *For any matrix  $\mathbf{A}$ ,*

$$\|\mathbf{A}\|_{\mathbb{F}} = \sup_{\|\mathbf{B}\|_{\mathbb{F}} \leq 1} \langle \mathbf{A}, \mathbf{B} \rangle_{\mathbb{F}}.$$

*For any  $\mathbf{H} \in \mathbb{S}^d$  and  $k \in \{1, \dots, d\}$ , we have the Ky Fan variational formula*

$$\text{tr}_k(\mathbf{H}) = \max_{\substack{\mathbf{\Pi}^2 = \mathbf{\Pi} = \mathbf{\Pi}^\top \\ \text{tr}(\mathbf{\Pi}) = k}} \text{tr}(\mathbf{H}\mathbf{\Pi}),$$

*and equivalently*

$$\text{tr}_k(\mathbf{H}) = \min_{\tau \in \mathbb{R}, \mathbf{Y} \succeq \mathbf{0}, \mathbf{Y} \succeq \mathbf{H} - \tau \mathbf{I}_d} k\tau + \text{tr}(\mathbf{Y}).$$

*Proof.* The Frobenius formula simply is Hilbert space duality. The projection formula is Ky Fan's maximum principle; see, e.g., Horn and Johnson, 1991, Section 3.4. For the epigraph formula, diagonalize  $\mathbf{H} = \sum_{i=1}^d \lambda_i u_i u_i^\top$ , for an orthonormal eigenbasis  $\{u_i\}_{i=1}^d$ . For fixed  $\tau$ , any feasible  $\mathbf{Y}$  satisfies  $u_i^\top \mathbf{Y} u_i \geq \max\{\lambda_i - \tau, 0\}$  for all  $i$ . Hence,

$$\text{tr}(\mathbf{Y}) \geq \sum_{i=1}^d (\lambda_i - \tau)_+.$$

Equality is attained by  $\mathbf{Y} = (\mathbf{H} - \tau \mathbf{I}_d)_+$ . Thus the inner minimum equals

$$k\tau + \sum_{i=1}^d (\lambda_i - \tau)_+,$$

whose minimum over  $\tau$  is the sum of the  $k$  largest eigenvalues of  $\mathbf{H}$ .  $\square$

## B Bures–Wasserstein Geometry and Regression

This appendix develops the second-moment formulation of least-squares regression in Bures space and collects some basic facts of Bures–Wasserstein geometry used throughout the proofs.

### B.1 Common Block Decompositions

Throughout the proofs, we frequently apply two types of matrix block decompositions. Recall that  $d = p + m$ . If  $\mathbf{\Sigma} \in \mathbb{S}^d$  is symmetric, we write

$$\mathbf{\Sigma} = \begin{pmatrix} \mathbf{\Sigma}_X & \mathbf{\Sigma}_{XY} \\ \mathbf{\Sigma}_{YX} & \mathbf{\Sigma}_Y \end{pmatrix}, \quad \mathbf{\Sigma}_X \in \mathbb{S}^p, \quad \mathbf{\Sigma}_{XY} = \mathbf{\Sigma}_{YX}^\top \in \mathbb{R}^{p \times m}, \quad \mathbf{\Sigma}_Y \in \mathbb{S}^m.$$

This matches our existing notation for second moments when  $\mathbf{\Sigma} = \mathbf{\Sigma}_P$ . We employ this decomposition when  $\mathbf{\Sigma}$  is a second-moment matrix, a nearby symmetric matrix, or a tangent direction to the positive semidefinite cone.

If  $\mathbf{S} \in \mathbb{R}^{d \times k}$  is introduced as a square-root factor for a matrix  $\mathbf{\Sigma} \in \mathbb{S}_+^d$ , i.e.,  $\mathbf{\Sigma} = \mathbf{S}\mathbf{S}^\top$ , or as the difference between two such factors, then we often decompose

$$\mathbf{S} = \begin{pmatrix} \mathbf{S}_X \\ \mathbf{S}_Y \end{pmatrix}, \quad \mathbf{S}_X \in \mathbb{R}^{p \times k}, \quad \mathbf{S}_Y \in \mathbb{R}^{m \times k}.$$

This decomposition arises naturally for Bures distance computations. For intuition when working with such blocks, note that if  $Z = (X, Y) = (\mathbf{S}_X W, \mathbf{S}_Y W)$  for  $W \sim \mathcal{N}(0, \mathbf{I}_k)$ , then  $\mathbb{E}[ZZ^\top] = \mathbf{S}\mathbf{S}^\top$ .

### B.2 Regression in Bures Space

Since all risks considered in the paper depend only on second moments, we extend the basic definitions from the main body to Bures space. For  $\mathbf{\Sigma} \in \mathbb{S}_+^d$  and  $\mathbf{A} \in \mathcal{A} = \mathbb{R}^{m \times p}$ , define

$$\begin{aligned} \mathbf{A}_\Sigma &:= \mathbf{\Sigma}_{YX} \mathbf{\Sigma}_X^\dagger, \\ L(\mathbf{A}, \mathbf{\Sigma}) &:= \text{tr}(\mathbf{A}^\top \mathbf{A} \mathbf{\Sigma}_X) - 2 \text{tr}(\mathbf{A} \mathbf{\Sigma}_{XY}) + \text{tr}(\mathbf{\Sigma}_Y), \\ R(\mathbf{A}, \mathbf{\Sigma}) &:= L(\mathbf{A}, \mathbf{\Sigma}) - L(\mathbf{A}_\Sigma, \mathbf{\Sigma}), \\ \Sigma_E &:= (\mathbf{A}_\Sigma \quad -\mathbf{I}_m) \mathbf{\Sigma} (\mathbf{A}_\Sigma \quad -\mathbf{I}_m)^\top. \end{aligned}$$

For  $P \in \mathcal{P}^*$ , these agree with the distributional definitions in the main body when  $\Sigma = \Sigma_P$ . We note that the residual second moment can be simplified as follows:

$$\begin{aligned}\Sigma_E &= (\mathbf{A}_\Sigma \quad -\mathbf{I}_m) \begin{pmatrix} \Sigma_X & \Sigma_{XY} \\ \Sigma_{YX} & \Sigma_Y \end{pmatrix} (\mathbf{A}_\Sigma \quad -\mathbf{I}_m)^\top \\ &= \mathbf{A}_\Sigma \Sigma_X \mathbf{A}_\Sigma^\top - \mathbf{A}_\Sigma \Sigma_{XY} - \Sigma_{YX} \mathbf{A}_\Sigma^\top + \Sigma_Y \\ &= \Sigma_{YX} \Sigma_X^\dagger \Sigma_X \Sigma_X^\dagger \Sigma_{XY} - 2 \Sigma_{YX} \Sigma_X^\dagger \Sigma_{XY} + \Sigma_Y \\ &= \Sigma_Y - \Sigma_{YX} \Sigma_X^\dagger \Sigma_{XY},\end{aligned}$$

where the last equality uses that  $\Sigma_X^\dagger \Sigma_X \Sigma_X^\dagger = \Sigma_X^\dagger$ .

The next lemma records several useful identities for these Bures-space regression quantities in terms of a square-root factor for  $\Sigma$ .

**Lemma B.1** (Regression factor identities and regret formula). *Let  $\Sigma \in \mathbb{S}_+^d$ , and fix any  $\mathbf{S} \in \mathbb{R}^{d \times r}$  such that  $\Sigma = \mathbf{S}\mathbf{S}^\top$ . Let  $\Pi_X = \mathbf{S}_X^\dagger \mathbf{S}_X$  be the orthogonal projection onto the row space of  $\mathbf{S}_X$ . Then*

$$\mathbf{A}_\Sigma = \Sigma_{YX} \Sigma_X^\dagger = \mathbf{S}_Y \mathbf{S}_X^\dagger.$$

For the chosen factor  $\mathbf{S}$ , the corresponding residual factor

$$\mathbf{E} := \mathbf{S}_Y - \mathbf{A}_\Sigma \mathbf{S}_X = \mathbf{S}_Y (\mathbf{I} - \Pi_X)$$

satisfies

$$\mathbf{E} \mathbf{S}_X^\top = \mathbf{0}, \quad \mathbf{E} \Pi_X = \mathbf{0}, \quad \mathbf{E} \mathbf{E}^\top = \Sigma_E.$$

Finally, for every  $\mathbf{A} \in \mathcal{A}$ , we have  $L(\mathbf{A}, \Sigma) = \|\mathbf{A} \mathbf{S}_X - \mathbf{S}_Y\|_{\mathbb{F}}^2$  and

$$R(\mathbf{A}, \Sigma) = \|(\mathbf{S}_Y - \mathbf{A} \mathbf{S}_X) \Pi_X\|_{\mathbb{F}}^2 = \|(\mathbf{A} - \mathbf{A}_\Sigma) \mathbf{S}_X\|_{\mathbb{F}}^2.$$

*Proof.* By Lemma A.2,  $\mathbf{S}_X^\dagger = \mathbf{S}_X^\top (\mathbf{S}_X \mathbf{S}_X^\top)^\dagger$ , and so

$$\mathbf{S}_Y \mathbf{S}_X^\dagger = \mathbf{S}_Y \mathbf{S}_X^\top (\mathbf{S}_X \mathbf{S}_X^\top)^\dagger = \Sigma_{YX} \Sigma_X^\dagger = \mathbf{A}_\Sigma.$$

Thus,  $\mathbf{A}_\Sigma \mathbf{S}_X = \mathbf{S}_Y \mathbf{S}_X^\dagger \mathbf{S}_X = \mathbf{S}_Y \Pi_X$ , and hence  $\mathbf{E} = \mathbf{S}_Y (\mathbf{I} - \Pi_X)$ . Since  $\Pi_X \mathbf{S}_X^\top = \mathbf{S}_X^\top$  and  $(\mathbf{I} - \Pi_X) \Pi_X = \mathbf{0}$ , we obtain

$$\mathbf{E} \mathbf{S}_X^\top = \mathbf{0}, \quad \mathbf{E} \Pi_X = \mathbf{0}.$$

Moreover,  $(\mathbf{A}_\Sigma \quad -\mathbf{I}_m) \mathbf{S} = -\mathbf{E}$ , and so

$$\mathbf{E} \mathbf{E}^\top = (\mathbf{A}_\Sigma \quad -\mathbf{I}_m) \Sigma (\mathbf{A}_\Sigma \quad -\mathbf{I}_m)^\top = \Sigma_E.$$

Finally,

$$L(\mathbf{A}, \Sigma) = \text{tr}((\mathbf{A} \mathbf{S}_X - \mathbf{S}_Y)(\mathbf{A} \mathbf{S}_X - \mathbf{S}_Y)^\top) = \|\mathbf{A} \mathbf{S}_X - \mathbf{S}_Y\|_{\mathbb{F}}^2.$$

We decompose this norm using  $\Pi_X$  and  $\mathbf{I} - \Pi_X$ :

$$\|\mathbf{A} \mathbf{S}_X - \mathbf{S}_Y\|_{\mathbb{F}}^2 = \|(\mathbf{A} \mathbf{S}_X - \mathbf{S}_Y) \Pi_X\|_{\mathbb{F}}^2 + \|\mathbf{S}_Y (\mathbf{I} - \Pi_X)\|_{\mathbb{F}}^2.$$

The second term is independent of  $\mathbf{A}$  and the first is zero for  $\mathbf{A} = \mathbf{A}_\Sigma$ . Thus,

$$R(\mathbf{A}, \Sigma) = \|(\mathbf{A} \mathbf{S}_X - \mathbf{S}_Y) \Pi_X\|_{\mathbb{F}}^2.$$

Finally, since  $\mathbf{S}_X \Pi_X = \mathbf{S}_X$  and  $\mathbf{A}_\Sigma \mathbf{S}_X = \mathbf{S}_Y \Pi_X$ ,

$$(\mathbf{A} \mathbf{S}_X - \mathbf{S}_Y) \Pi_X = (\mathbf{A} - \mathbf{A}_\Sigma) \mathbf{S}_X,$$

which gives the final equality.  $\square$

Next, set

$$\mathcal{C}^* := \{\Sigma_P : P \in \mathcal{P}^*\} = \{\Sigma \in \mathbb{S}_+^d : \Sigma_X \succ \mathbf{0}\}.$$

Using the second-moment formulation of the risk together with Lemma 1.2, the condition number and our main-body estimators admit equivalent Bures-space formulations. For  $\Sigma \in \mathcal{C}^*$ , define

$$R_\varepsilon^*(\Sigma) := \inf_{\mathbf{A} \in \mathcal{A}} \sup_{\Sigma' \in \mathcal{B}_\varepsilon(\Sigma)} R(\mathbf{A}, \Sigma'), \quad \kappa(\Sigma) := \limsup_{\varepsilon \downarrow 0} \varepsilon^{-2} R_\varepsilon^*(\Sigma). \quad (4)$$

Then, by (1) and Lemma 1.2,  $\kappa_P = \kappa(\Sigma_P)$  for  $P \in \mathcal{P}^*$ . Likewise, the estimators can be written as

$$\begin{aligned}\text{ERM}[\Sigma] &\in \arg \min_{\mathbf{A} \in \mathcal{A}} L(\mathbf{A}, \Sigma), \\ \text{DRO}_\varepsilon[\Sigma] &\in \arg \min_{\mathbf{A} \in \mathcal{A}} \sup_{\Sigma' \in \mathcal{B}_\varepsilon(\Sigma)} L(\mathbf{A}, \Sigma'), \\ \text{MRO}_\varepsilon[\Sigma] &\in \arg \min_{\mathbf{A} \in \mathcal{A}} \sup_{\Sigma' \in \mathcal{B}_\varepsilon(\Sigma)} R(\mathbf{A}, \Sigma'), \\ \text{SNR}_\varepsilon[\Sigma] &\in \arg \min_{\mathbf{A} \in \mathcal{A}} L(\mathbf{A}, \Sigma) + \varepsilon^2 \|\mathbf{A}\|_{\text{op}}^2.\end{aligned}$$

### B.3 Bures–Wasserstein Geometry

We now turn to proofs of the Section 1.2 lemmas on the Bures metric.

*Proof of Lemma 1.1.* Given  $\mathbf{A}, \mathbf{B} \in \mathbb{S}_+^d$ , the identity

$$\text{B}(\mathbf{A}, \mathbf{B}) = \min_{\mathbf{U} \in \mathbb{O}^d} \|\mathbf{A}^{\frac{1}{2}} - \mathbf{B}^{\frac{1}{2}} \mathbf{U}\|_{\text{F}} \quad (5)$$

$$= \min_{\substack{\mathbf{S}, \mathbf{T} \in \mathbb{R}^{d \times d} \\ \mathbf{A} = \mathbf{S}\mathbf{S}^\top, \mathbf{B} = \mathbf{T}\mathbf{T}^\top}} \|\mathbf{S} - \mathbf{T}\|_{\text{F}}. \quad (6)$$

is well-known; see, e.g., Bhatia et al. (2019, Theorem 1).

Suppose  $\text{B}(\mathbf{A}, \mathbf{B}) \leq \varepsilon$ . Fix  $\mathbf{U} \in \mathbb{O}^d$  minimizing (5), so that  $\|\mathbf{A}^{1/2} - \mathbf{B}^{1/2} \mathbf{U}\|_{\text{F}} \leq \varepsilon$ . Setting  $\Delta := \mathbf{A}^{1/2} \mathbf{U}^\top - \mathbf{B}^{1/2}$ , we then have  $\|\Delta\|_{\text{F}} = \|\Delta \mathbf{U}\|_{\text{F}} \leq \varepsilon$  and

$$(\mathbf{B}^{\frac{1}{2}} + \Delta)(\mathbf{B}^{\frac{1}{2}} + \Delta)^\top = \mathbf{A}^{\frac{1}{2}} \mathbf{U}^\top \mathbf{U} \mathbf{A}^{\frac{1}{2}} = \mathbf{A},$$

as desired. On the other hand, if such  $\Delta$  exists, then  $\mathbf{S} := \mathbf{B}^{\frac{1}{2}} + \Delta$  and  $\mathbf{T} := \mathbf{B}^{\frac{1}{2}}$  are feasible square-root factors for (6), and so

$$\text{B}(\mathbf{A}, \mathbf{B}) \leq \|\mathbf{S} - \mathbf{T}\|_{\text{F}} = \|\Delta\|_{\text{F}} \leq \varepsilon.$$

If  $\mathbf{A}$  and  $\mathbf{B}$  are simultaneously diagonalizable, then  $\mathbf{A}^{1/2}$  and  $\mathbf{B}^{1/2}$  commute. In this case, we have

$$\begin{aligned}\text{B}(\mathbf{A}, \mathbf{B})^2 &= \text{tr}(\mathbf{A}) + \text{tr}(\mathbf{B}) - 2 \text{tr}((\mathbf{A}^{\frac{1}{2}} \mathbf{B} \mathbf{A}^{\frac{1}{2}})^{\frac{1}{2}}) \\ &= \text{tr}(\mathbf{A}) + \text{tr}(\mathbf{B}) - 2 \text{tr}(\mathbf{A}^{\frac{1}{2}} \mathbf{B}^{\frac{1}{2}}) \\ &= \|\mathbf{A}^{\frac{1}{2}} - \mathbf{B}^{\frac{1}{2}}\|_{\text{F}}^2,\end{aligned}$$

implying that the minimum in (5) is attained by  $\mathbf{U} = \mathbf{I}_d$ , as claimed.  $\square$

*Proof of Lemma 1.2.* The second equality follows directly from Lemma 1.1. We now prove the first.

Proposition 1 of Nguyen et al. (2021a) establishes a related claim for the augmented Gelbrich distance over mean–covariance pairs, assuming non-singular covariance matrices. Here we give a direct argument for the Bures case with uncentered and potentially singular second-moment matrices.

**⊆ direction.** Fix  $Q \in \mathcal{B}_\varepsilon(P)$  and any coupling of random variables  $(Z, Z')$  with  $Z \sim P$  and  $Z' \sim Q$ . Setting  $\mathbf{C} := \mathbb{E}[Z(Z')^\top]$ , the block second-moment matrix

$$\begin{pmatrix} \Sigma_P & \mathbf{C} \\ \mathbf{C}^\top & \Sigma_Q \end{pmatrix}$$

is PSD. Hence, there exists a contraction  $\mathbf{K} \in \mathbb{R}^{d \times d}$  with  $\|\mathbf{K}\|_{\text{op}} \leq 1$  such that  $\mathbf{C} = \Sigma_P^{1/2} \mathbf{K} \Sigma_Q^{1/2}$  (see, e.g., Bhatia, 2009, Proposition 1.3.2). Therefore, by duality between the operator norm  $\|\cdot\|_{\text{op}}$  and nuclear norm  $\|\cdot\|_*$  (the sum of singular values), we have

$$\text{tr}(\mathbf{C}) = \text{tr}(\mathbf{K} \Sigma_Q^{1/2} \Sigma_P^{1/2}) \leq \|\Sigma_Q^{1/2} \Sigma_P^{1/2}\|_* = \text{tr}\left((\Sigma_P^{1/2} \Sigma_Q \Sigma_P^{1/2})^{\frac{1}{2}}\right).$$

It then follows from definition of  $B$  that

$$\mathbb{E} \|Z - Z'\|^2 = \text{tr}(\Sigma_P) + \text{tr}(\Sigma_Q) - 2 \text{tr}(C) \geq B(\Sigma_P, \Sigma_Q)^2.$$

Infimizing over couplings gives  $W_2(P, Q) \geq B(\Sigma_P, \Sigma_Q)$ . Thus  $\Sigma_Q \in \mathcal{B}_\varepsilon(\Sigma_P)$ .

$\supseteq$  **direction.** Fix  $\Sigma \in \mathcal{B}_\varepsilon(\Sigma_P)$ . It suffices to identify  $Q \in \mathcal{B}_\varepsilon(P)$  such that  $\Sigma = \Sigma_Q$ . Abbreviate  $B_P := \Sigma_P^{1/2}$ . By the Bures bound and Lemma 1.1, there exists a factor  $B \in \mathbb{R}^{d \times d}$  such that  $\|B_P - B\|_F \leq \varepsilon$  and  $BB^\top = \Sigma$ . Let  $\Pi \in \mathbb{R}^{d \times d}$  be the orthogonal projection onto  $\text{range}(\Sigma_P) = \text{range}(B_P)$ . By Lemma A.2, we have  $\Pi = B_P B_P^\dagger = B_P^\dagger \Sigma_P B_P^\dagger$ , where  $\dagger$  denotes the Moore–Penrose pseudoinverse. Let  $Z \sim P$  and  $W \sim \mathcal{N}(0, I_d - \Pi)$  be independent. Now consider the random variable  $U := B_P^\dagger Z + W$ , constructed so that

$$\begin{aligned} \mathbb{E}[UU^\top] &= B_P^\dagger \Sigma_P B_P^\dagger + I_d - \Pi = \Pi + I_d - \Pi = I_d \\ B_P U &= \Pi Z + B_P W = Z. \end{aligned}$$

Indeed,  $Z$  is supported on the range of  $\Sigma_P$  and  $W$  is supported on the null space of  $B_P$ . We take  $Q$  as the law of  $Z' := BU$ . Indeed, since  $U$  has identity second moment, we have  $\Sigma_Q = BB^\top = \Sigma$ . Moreover, the coupling  $(Z, Z')$  certifies that

$$W_2(P, Q)^2 \leq \mathbb{E}[\|Z - Z'\|^2] = \mathbb{E}[\|(B_P - B)U\|^2] = \|B_P - B\|_F^2 \leq \varepsilon^2.$$

Thus,  $Q \in \mathcal{B}_\varepsilon(P)$ , as desired.  $\square$

Lastly, we apply Lemma 1.1 to prove a simple upper bound on the Bures distance.

**Lemma B.2** (Bures distance upper bound). *For  $A, B \in \mathbb{S}_+^d$  with  $\text{range}(B) \subseteq \text{range}(A)$ , we have*

$$B(A, B)^2 \leq \text{tr}((B - A)A^\dagger(B - A)).$$

*Proof.* We assume without loss of generality that  $A \succ 0$ . Indeed, if  $A = 0$ , then the range constraint forces  $B = 0$  and the claim holds trivially. Otherwise, since  $\text{range}(B) \subseteq \text{range}(A)$ , both matrices vanish on  $\text{range}(A)^\perp$ . We may therefore restrict to  $\text{range}(A)$ , where  $A$  is positive definite and  $A^\dagger$  agrees with its inverse.

Set  $C := A^{-1/2}BA^{-1/2} \succeq 0$ . By construction, we have

$$B = A^{1/2}CA^{1/2} = (A^{1/2}C^{1/2})(A^{1/2}C^{1/2})^\top.$$

Using the square-root characterization of the Bures distance (Lemma 1.1) with the particular square-root factors  $A^{1/2}$  and  $A^{1/2}C^{1/2}$ , we obtain

$$B(A, B)^2 \leq \|A^{1/2}(C^{1/2} - I_d)\|_F^2.$$

Since  $C^{1/2} - I_d = (C - I_d)(C^{1/2} + I_d)^{-1}$  and  $\|(C^{1/2} + I_d)^{-1}\|_{\text{op}} \leq 1$ , we further bound

$$\begin{aligned} \|A^{1/2}(C^{1/2} - I_d)\|_F &\leq \|A^{1/2}(C - I_d)\|_F \|(C^{1/2} + I_d)^{-1}\|_{\text{op}} \\ &\leq \|A^{1/2}(C - I_d)\|_F. \end{aligned} \quad (\text{Lemma A.1})$$

By our choice of  $C$ , we have  $A^{1/2}(C - I_d) = (B - A)A^{-1/2}$ , and thus

$$\|A^{1/2}(C - I_d)\|_F^2 = \text{tr}((B - A)A^{-1}(B - A)).$$

Combining the above gives the claimed bound on  $B(A, B)$ .  $\square$

Finally, we establish Lipschitz continuity of the square root of the smallest feature eigenvalue.

**Lemma B.3** (Bures continuity of the smallest feature eigenvalue). *For  $\Sigma, \bar{\Sigma} \in \mathbb{S}_+^d$ , we have*

$$\left| \sqrt{\lambda_{\min}(\bar{\Sigma}_X)} - \sqrt{\lambda_{\min}(\Sigma_X)} \right| \leq B(\bar{\Sigma}, \Sigma).$$

*In particular,  $\lambda_{\min}(\Sigma_X)$  is continuous with respect to the Bures metric.*

*Proof.* By Lemma 1.1, there exist square-root factors  $\mathbf{S}, \bar{\mathbf{S}} \in \mathbb{R}^{d \times d}$  of  $\Sigma$  and  $\bar{\Sigma}$  such that

$$\|\bar{\mathbf{S}} - \mathbf{S}\|_F = B(\bar{\Sigma}, \Sigma).$$

Using the Section B.1 block decompositions, we have  $\Sigma_X = \mathbf{S}_X \mathbf{S}_X^\top$  and  $\bar{\Sigma}_X = \bar{\mathbf{S}}_X \bar{\mathbf{S}}_X^\top$ . Hence, by Lemma A.5,

$$\begin{aligned} \left| \sqrt{\lambda_{\min}(\bar{\Sigma}_X)} - \sqrt{\lambda_{\min}(\Sigma_X)} \right| &= |\sigma_{\min}(\bar{\mathbf{S}}_X) - \sigma_{\min}(\mathbf{S}_X)| \\ &\leq \|\bar{\mathbf{S}}_X - \mathbf{S}_X\|_{\text{op}} \\ &\leq \|\bar{\mathbf{S}} - \mathbf{S}\|_F \\ &= B(\bar{\Sigma}, \Sigma). \end{aligned}$$

The continuity claim follows immediately.  $\square$

## C DRO Duality and Symmetry Reductions

This appendix establishes a dual formulation for the inner DRO maximization problem and an averaging principle for DRO, MRO, and SNR under orthogonal symmetries. We use  $\sup \emptyset = -\infty$ .

### C.1 Strong Duality for the Inner DRO Problem

We first establish a helper lemma, which will be applied both for the inner DRO problem below and for the fixed-projection problem in Section G.1. Very similar arguments appear in Shafieezadeh-Abadeh et al. (2018); Nguyen et al. (2021b), but we provide a self-contained proof to account for the case where  $r < d$  (untreated in those works).

**Lemma C.1** (Quadratic-ball duality). *Let  $\mathbf{M} \in \mathbb{R}^{k \times d}$  and  $\tilde{\mathbf{W}} \in \mathbb{R}^{d \times r}$ , where  $k, r \geq 1$ . Set  $b := \|\mathbf{M}\|_{\text{op}}^2$ . For every  $\eta > b$ , we have*

$$\sup_{\mathbf{W} \in \mathbb{R}^{d \times r}} \left\{ \|\mathbf{M}\mathbf{W}\|_F^2 - \eta \|\mathbf{W} - \tilde{\mathbf{W}}\|_F^2 \right\} = \eta \operatorname{tr} \left( (\eta \mathbf{I}_k - \mathbf{M}\mathbf{M}^\top)^{-1} \mathbf{M} \tilde{\mathbf{W}} \tilde{\mathbf{W}}^\top \mathbf{M}^\top \right). \quad (7)$$

Moreover, for every  $t \in \mathbb{R}$ , with  $\sup \emptyset = -\infty$ , we have

$$\begin{aligned} \sup_{\substack{\mathbf{W} \in \mathbb{R}^{d \times r} \\ \|\mathbf{W} - \tilde{\mathbf{W}}\|_F^2 \leq t}} \|\mathbf{M}\mathbf{W}\|_F^2 &= \sup_{\mathbf{W} \in \mathbb{R}^{d \times r}} \inf_{\eta \geq 0} \left\{ \|\mathbf{M}\mathbf{W}\|_F^2 + \eta(t - \|\mathbf{W} - \tilde{\mathbf{W}}\|_F^2) \right\} \\ &= \inf_{\eta \geq 0} \sup_{\mathbf{W} \in \mathbb{R}^{d \times r}} \left\{ \|\mathbf{M}\mathbf{W}\|_F^2 + \eta(t - \|\mathbf{W} - \tilde{\mathbf{W}}\|_F^2) \right\} \\ &= \inf_{\eta > b} \eta \left( t + \operatorname{tr} \left( (\eta \mathbf{I}_k - \mathbf{M}\mathbf{M}^\top)^{-1} \mathbf{M} \tilde{\mathbf{W}} \tilde{\mathbf{W}}^\top \mathbf{M}^\top \right) \right). \end{aligned} \quad (8)$$

*Proof.* Fix  $t \in \mathbb{R}$ . For  $\eta \geq 0$ , define

$$\Phi_t(\eta) := \sup_{\mathbf{W} \in \mathbb{R}^{d \times r}} \left\{ \|\mathbf{M}\mathbf{W}\|_F^2 + \eta(t - \|\mathbf{W} - \tilde{\mathbf{W}}\|_F^2) \right\}.$$

We first compute this supremum for  $\eta > b$ . Expanding the quadratic expression gives

$$\|\mathbf{M}\mathbf{W}\|_F^2 - \eta \|\mathbf{W} - \tilde{\mathbf{W}}\|_F^2 = -\operatorname{tr}(\mathbf{W}^\top (\eta \mathbf{I}_d - \mathbf{M}^\top \mathbf{M}) \mathbf{W}) + 2\eta \langle \mathbf{W}, \tilde{\mathbf{W}} \rangle_F - \eta \|\tilde{\mathbf{W}}\|_F^2.$$

Since  $\eta \mathbf{I}_d - \mathbf{M}^\top \mathbf{M} \succ \mathbf{0}$ , this expression is strictly concave in  $\mathbf{W}$ . Its unique maximizer satisfies

$$(\eta \mathbf{I}_d - \mathbf{M}^\top \mathbf{M}) \mathbf{W} = \eta \tilde{\mathbf{W}},$$

and therefore equals  $\eta(\eta \mathbf{I}_d - \mathbf{M}^\top \mathbf{M})^{-1} \tilde{\mathbf{W}}$ . Substitution yields

$$\begin{aligned} \Phi_t(\eta) &= \eta t + \eta^2 \operatorname{tr} \left( \tilde{\mathbf{W}}^\top (\eta \mathbf{I}_d - \mathbf{M}^\top \mathbf{M})^{-1} \tilde{\mathbf{W}} \right) - \eta \|\tilde{\mathbf{W}}\|_F^2 \\ &= \eta t + \eta \operatorname{tr} \left( \tilde{\mathbf{W}}^\top [\eta(\eta \mathbf{I}_d - \mathbf{M}^\top \mathbf{M})^{-1} - \mathbf{I}_d] \tilde{\mathbf{W}} \right) \\ &= \eta t + \eta \operatorname{tr} \left( (\eta \mathbf{I}_k - \mathbf{M}\mathbf{M}^\top)^{-1} \mathbf{M} \tilde{\mathbf{W}} \tilde{\mathbf{W}}^\top \mathbf{M}^\top \right), \end{aligned} \quad (9)$$

where the last equality uses Lemma A.4 and cyclicity of trace. Subtracting  $\eta t$  proves (7).

We next prove the final equality in (8). Given (9), it suffices to show

$$\inf_{\eta \geq 0} \Phi_t(\eta) = \inf_{\eta > b} \Phi_t(\eta).$$

For  $0 \leq \eta < b$ , varying one column of  $\mathbf{W}$  along a top eigenvector of  $\mathbf{M}^\top \mathbf{M}$  makes the expression defining  $\Phi_t(\eta)$  unbounded above, since its leading quadratic coefficient is  $b - \eta > 0$ . Thus  $\Phi_t(\eta) = +\infty$ . At the endpoint  $\eta = b$ , observe that as  $\eta \downarrow b$ , the functions  $\mathbf{W} \mapsto \|\mathbf{M}\mathbf{W}\|_{\mathbb{F}}^2 - \eta \|\mathbf{W} - \tilde{\mathbf{W}}\|_{\mathbb{F}}^2$  increase pointwise to their values at  $b$ , so their suprema converge to the supremum at  $b$ . Further,  $\eta \mapsto \eta t$  is continuous, so we obtain

$$\lim_{\eta \downarrow b} \Phi_t(\eta) = \Phi_t(b),$$

allowing  $+\infty$ . Hence excluding all  $\eta \leq b$  does not change the infimum.

We now return to the initial constrained supremum in (8). For each fixed  $\mathbf{W}$ , we have

$$\inf_{\eta \geq 0} \left\{ \|\mathbf{M}\mathbf{W}\|_{\mathbb{F}}^2 + \eta(t - \|\mathbf{W} - \tilde{\mathbf{W}}\|_{\mathbb{F}}^2) \right\} = \begin{cases} \|\mathbf{M}\mathbf{W}\|_{\mathbb{F}}^2, & \|\mathbf{W} - \tilde{\mathbf{W}}\|_{\mathbb{F}}^2 \leq t, \\ -\infty, & \text{otherwise.} \end{cases}$$

This proves the first equality in (8).

For the remaining equality, suppose first that  $t > 0$  and set

$$\gamma := \max_{\substack{\mathbf{W} \in \mathbb{R}^{d \times r} \\ \|\mathbf{W} - \tilde{\mathbf{W}}\|_{\mathbb{F}}^2 \leq t}} \|\mathbf{M}\mathbf{W}\|_{\mathbb{F}}^2.$$

This maximum exists and is finite because the feasible set is compact and the objective is continuous. By definition of  $\gamma$ , the quadratic implication

$$t - \|\mathbf{W} - \tilde{\mathbf{W}}\|_{\mathbb{F}}^2 \geq 0 \implies \gamma - \|\mathbf{M}\mathbf{W}\|_{\mathbb{F}}^2 \geq 0$$

holds for every  $\mathbf{W}$ . The first inequality is strictly satisfied at  $\mathbf{W} = \tilde{\mathbf{W}}$ . After vectorizing  $\mathbf{W}$ , the S-lemma (Pólik and Terlaky, 2007, Theorem 2.2 and Section 8.1) thus gives some  $\eta \geq 0$  such that

$$\gamma - \|\mathbf{M}\mathbf{W}\|_{\mathbb{F}}^2 - \eta(t - \|\mathbf{W} - \tilde{\mathbf{W}}\|_{\mathbb{F}}^2) \geq 0 \quad \forall \mathbf{W} \in \mathbb{R}^{d \times r}.$$

Equivalently,  $\Phi_t(\eta) \leq \gamma$ . Conversely,  $\Phi_t(\eta) \geq \gamma$  for every  $\eta \geq 0$ , since the additional term in the definition of  $\Phi_t$  is nonnegative on the feasible set. Consequently,

$$\gamma = \inf_{\eta \geq 0} \Phi_t(\eta).$$

This yields the second equality in (8) when  $t > 0$ .

If  $t = 0$ , the only feasible point is  $\mathbf{W} = \tilde{\mathbf{W}}$ , so the constrained supremum is  $\|\mathbf{M}\tilde{\mathbf{W}}\|_{\mathbb{F}}^2$ . Moreover, the formula (9) established for  $\Phi_t(\eta)$  implies that

$$\inf_{\eta \geq 0} \Phi_0(\eta) \leq \lim_{\eta \rightarrow \infty} \Phi_0(\eta) = \|\mathbf{M}\tilde{\mathbf{W}}\|_{\mathbb{F}}^2,$$

using that  $\eta(\eta \mathbf{I}_k - \mathbf{M}\mathbf{M}^\top)^{-1} \rightarrow \mathbf{I}_k$ . Since  $\Phi_0(\eta)$  is at least the constrained supremum for every  $\eta \geq 0$ , equality follows.

Finally, if  $t < 0$ , the constrained feasible set is empty and its supremum equals  $-\infty$ . Meanwhile, the formula (9) for  $\Phi_t$  gives

$$\Phi_t(\eta) = \eta t + O(1) \rightarrow -\infty \quad \text{as } \eta \rightarrow \infty,$$

so  $\inf_{\eta \geq 0} \Phi_t(\eta) = -\infty$  as well.  $\square$

We now apply this to the inner maximization problem in the DRO formulation. For  $\mathbf{A} \in \mathcal{A}$ , write

$$\mathbf{M}_{\mathbf{A}} := (\mathbf{A} \quad -\mathbf{I}_m) \in \mathbb{R}^{m \times d},$$

so that  $L(\mathbf{A}, P) = \mathbb{E}_P \|\mathbf{M}_{\mathbf{A}} Z\|^2$ . Direct multiplication gives

$$\mathbf{M}_{\mathbf{A}} \mathbf{M}_{\mathbf{A}}^\top = \mathbf{A} \mathbf{A}^\top + \mathbf{I}_m, \quad \|\mathbf{M}_{\mathbf{A}}\|_{\text{op}}^2 = 1 + \|\mathbf{A}\|_{\text{op}}^2.$$

We state the dual result for a generic quadratic loss before specializing to regression.

**Lemma C.2** (Inner DRO strong duality). *Let  $P \in \mathcal{P}_2(\mathbb{R}^d)$ ,  $\varepsilon > 0$ , and  $\mathbf{M} \in \mathbb{R}^{k \times d}$  for  $1 \leq k \leq d$ . Then*

$$\sup_{Q \in \mathcal{B}_\varepsilon(P)} \mathbb{E}_Q \|\mathbf{M}Z\|^2 = \inf_{\eta > \|\mathbf{M}\|_{\text{op}}^2} \eta \left( \varepsilon^2 + \text{tr}((\eta \mathbf{I}_k - \mathbf{M}\mathbf{M}^\top)^{-1} \mathbf{M} \Sigma_P \mathbf{M}^\top) \right).$$

*For regression, taking  $k = m$  and  $\mathbf{M} = \mathbf{M}_\mathbf{A}$  gives*

$$\sup_{Q \in \mathcal{B}_\varepsilon(P)} L(\mathbf{A}, Q) = \inf_{\eta > 1 + \|\mathbf{A}\|_{\text{op}}^2} \eta \left( \varepsilon^2 + \text{tr}(((\eta - 1)\mathbf{I}_m - \mathbf{A}\mathbf{A}^\top)^{-1} \mathbf{M}_\mathbf{A} \Sigma_P \mathbf{M}_\mathbf{A}^\top) \right).$$

*Proof.* By Lemma 1.2, the second moments of laws in  $\mathcal{B}_\varepsilon(P)$  are precisely the matrices  $\mathbf{W}\mathbf{W}^\top$ , where  $\mathbf{W} \in \mathbb{R}^{d \times d}$  satisfies  $\|\mathbf{W} - \Sigma_P^{1/2}\|_{\text{F}}^2 \leq \varepsilon^2$ . Since  $\mathbb{E}_Q \|\mathbf{M}Z\|^2 = \text{tr}(\mathbf{M} \Sigma_Q \mathbf{M}^\top)$ , we obtain

$$\begin{aligned} \sup_{Q \in \mathcal{B}_\varepsilon(P)} \mathbb{E}_Q \|\mathbf{M}Z\|^2 &= \sup_{\substack{\mathbf{W} \in \mathbb{R}^{d \times d} \\ \|\mathbf{W} - \Sigma_P^{1/2}\|_{\text{F}}^2 \leq \varepsilon^2}} \|\mathbf{M}\mathbf{W}\|_{\text{F}}^2 \\ &= \inf_{\eta > \|\mathbf{M}\|_{\text{op}}^2} \eta \left( \varepsilon^2 + \text{tr}((\eta \mathbf{I}_k - \mathbf{M}\mathbf{M}^\top)^{-1} \mathbf{M} \Sigma_P \mathbf{M}^\top) \right), \end{aligned}$$

where the second equality applies Lemma C.1 with  $r = d$ ,  $\tilde{\mathbf{W}} = \Sigma_P^{1/2}$ , and  $t = \varepsilon^2$ .

Finally, the regression specialization follows by taking  $\mathbf{M} = \mathbf{M}_\mathbf{A}$  and using  $\mathbf{M}_\mathbf{A} \mathbf{M}_\mathbf{A}^\top = \mathbf{A}\mathbf{A}^\top + \mathbf{I}_m$  together with  $L(\mathbf{A}, Q) = \mathbb{E}_Q \|\mathbf{M}_\mathbf{A}Z\|^2$ .  $\square$

We note that the infimum over the open dual domain need not be attained; its value may be approached at the left endpoint. We next specialize to the case where  $\mathbf{A}\mathbf{A}^\top$  is diagonal.

**Corollary C.3** (Diagonal case). *Let  $P \in \mathcal{P}_2(\mathbb{R}^d)$ ,  $\varepsilon > 0$ , and  $\mathbf{A} \in \mathcal{A}$  satisfy*

$$\mathbf{A}\mathbf{A}^\top = \text{diag}(a_1^2, \dots, a_m^2).$$

*For  $i = 1, \dots, m$ , define the  $i$ th coordinate residual second moment*

$$r_i^2 := \mathbb{E}_P[(e_i^\top (\mathbf{A}X - Y))^2] = e_i^\top \mathbf{M}_\mathbf{A} \Sigma_P \mathbf{M}_\mathbf{A}^\top e_i.$$

*We then have*

$$\sup_{Q \in \mathcal{B}_\varepsilon(P)} L(\mathbf{A}, Q) = \inf_{\eta > 1 + \max_i a_i^2} \eta \left( \varepsilon^2 + \sum_{i=1}^m \frac{r_i^2}{\eta - 1 - a_i^2} \right).$$

*Proof.* In the regression formula of Lemma C.2,

$$((\eta - 1)\mathbf{I}_m - \mathbf{A}\mathbf{A}^\top)^{-1} = \text{diag}\left(\frac{1}{\eta - 1 - a_1^2}, \dots, \frac{1}{\eta - 1 - a_m^2}\right).$$

Taking the trace of its product with  $\mathbf{M}_\mathbf{A} \Sigma_P \mathbf{M}_\mathbf{A}^\top$  gives the stated sum.  $\square$

## C.2 Averaging over Symmetries

All symmetry reductions used below preserve the feature and response subspaces. We therefore formulate the averaging principle directly for block-diagonal orthogonal transformations. Recall that for a measurable map  $T$ ,  $T_\# P$  denotes the pushforward law of  $T(Z)$  when  $Z \sim P$ .

**Lemma C.4** (Averaging over symmetries). *Let  $U$  be a compact subgroup of*

$$\{\text{diag}(\mathbf{U}_X, \mathbf{U}_Y) : \mathbf{U}_X \in \mathbb{O}^p, \mathbf{U}_Y \in \mathbb{O}^m\}.$$

*Suppose  $P, \tilde{P} \in \mathcal{P}_2(\mathbb{R}^d)$  satisfy  $\mathbf{U}_\# P = P$  and  $\mathbf{U}_\# \tilde{P} = \tilde{P}$  for every  $\mathbf{U} \in U$ . Let  $\varepsilon > 0$ , and fix any minimizer  $\hat{\mathbf{A}} \in \mathcal{A}$  for the  $\text{DRO}_\varepsilon[\tilde{P}]$ ,  $\text{MRO}_\varepsilon[\tilde{P}]$ , or  $\text{SNR}_\varepsilon[\tilde{P}]$  optimization problem.*

*Let  $\mu$  be normalized Haar measure on  $U$  and define*

$$\bar{\mathbf{A}} := \int_U \mathbf{U}_Y^\top \hat{\mathbf{A}} \mathbf{U}_X \, d\mu(\mathbf{U}).$$

Then  $\bar{\mathbf{A}}$  is a minimizer of the same optimization problem, is invariant in the sense that

$$\mathbf{U}_Y^\top \bar{\mathbf{A}} \mathbf{U}_X = \bar{\mathbf{A}} \quad \forall \mathbf{U} \in U,$$

and satisfies  $R(\bar{\mathbf{A}}, P) \leq R(\hat{\mathbf{A}}, P)$ . Consequently, restricting any of these optimization problems to invariant coefficient matrices preserves its optimal value. Moreover, to lower bound the target regret of every optimizer, it suffices to lower bound the target regret of every invariant optimizer.

*Proof.* For  $\mathbf{U} = \text{diag}(\mathbf{U}_X, \mathbf{U}_Y) \in U$ , define the linear coefficient transformation

$$T_{\mathbf{U}}(\mathbf{A}) := \mathbf{U}_Y^\top \mathbf{A} \mathbf{U}_X.$$

Orthogonality gives

$$\begin{aligned} L(\mathbf{A}, \mathbf{U}_\# Q) &= \mathbb{E}_Q \|\mathbf{A} \mathbf{U}_X X - \mathbf{U}_Y Y\|^2 \\ &= \mathbb{E}_Q \|\mathbf{U}_Y^\top \mathbf{A} \mathbf{U}_X X - Y\|^2 \\ &= L(T_{\mathbf{U}}(\mathbf{A}), Q). \end{aligned}$$

The map  $T_{\mathbf{U}}$  is a bijection from  $\mathcal{A}$  to itself, so

$$\inf_{\mathbf{B} \in \mathcal{A}} L(\mathbf{B}, \mathbf{U}_\# Q) = \inf_{\mathbf{B} \in \mathcal{A}} L(\mathbf{B}, Q).$$

Subtracting this identity from the preceding risk identity yields

$$R(\mathbf{A}, \mathbf{U}_\# Q) = R(T_{\mathbf{U}}(\mathbf{A}), Q).$$

Also, left and right multiplication by orthogonal matrices gives  $\|T_{\mathbf{U}}(\mathbf{A})\|_{\text{op}} = \|\mathbf{A}\|_{\text{op}}$ .

The source ambiguity set is invariant under each  $\mathbf{U} \in U$ . Indeed,

$$W_2(\mathbf{U}_\# Q, \tilde{P}) = W_2(\mathbf{U}_\# Q, \mathbf{U}_\# \tilde{P}) = W_2(Q, \tilde{P}).$$

Applying this also to  $\mathbf{U}^{-1}$  proves  $\mathbf{U}_\# \mathcal{B}_\varepsilon(\tilde{P}) = \mathcal{B}_\varepsilon(\tilde{P})$ . It follows that each of the three objectives

$$\sup_{Q \in \mathcal{B}_\varepsilon(\tilde{P})} L(\mathbf{A}, Q), \quad \sup_{Q \in \mathcal{B}_\varepsilon(\tilde{P})} R(\mathbf{A}, Q), \quad L(\mathbf{A}, \tilde{P}) + \varepsilon^2 \|\mathbf{A}\|_{\text{op}}^2$$

is invariant under  $T_{\mathbf{U}}$ . Each objective is also convex in  $\mathbf{A}$ .

The integral defining  $\bar{\mathbf{A}}$  is well-defined because its integrand is continuous on the compact group  $U$ . For any fixed  $\mathbf{V} = \text{diag}(\mathbf{V}_X, \mathbf{V}_Y) \in U$ ,

$$\begin{aligned} \mathbf{V}_Y^\top \bar{\mathbf{A}} \mathbf{V}_X &= \int_U (\mathbf{U}_Y \mathbf{V}_Y)^\top \hat{\mathbf{A}} (\mathbf{U}_X \mathbf{V}_X) d\mu(\mathbf{U}) \\ &= \bar{\mathbf{A}}, \end{aligned}$$

where the last equality uses right invariance of Haar measure under  $\mathbf{U} \mapsto \mathbf{U}\mathbf{V}$ . Thus  $\bar{\mathbf{A}}$  is invariant.

Let  $J$  denote whichever of the three objectives is minimized by  $\hat{\mathbf{A}}$ . By convexity, invariance, and Jensen's inequality,

$$J(\bar{\mathbf{A}}) \leq \int_U J(T_{\mathbf{U}}(\hat{\mathbf{A}})) d\mu(\mathbf{U}) = J(\hat{\mathbf{A}}).$$

Since  $\hat{\mathbf{A}}$  is a minimizer,  $\bar{\mathbf{A}}$  is also a minimizer.

Finally, invariance of  $P$  gives  $R(T_{\mathbf{U}}(\hat{\mathbf{A}}), P) = R(\hat{\mathbf{A}}, P)$ . The regret is convex in its coefficient argument, so another application of Jensen's inequality yields

$$R(\bar{\mathbf{A}}, P) \leq \int_U R(T_{\mathbf{U}}(\hat{\mathbf{A}}), P) d\mu(\mathbf{U}) = R(\hat{\mathbf{A}}, P).$$

This proves all claims.  $\square$

## D Remaining Proofs for Section 1

### D.1 Existence of Minimizers (Footnote 1)

We treat the estimators individually.

**ERM.** Solving the normal equation  $\nabla_{\mathbf{A}} L(\mathbf{A}, P) = \mathbf{0}$  reveals that  $\Sigma_{YX}^P (\Sigma_X^P)^\dagger$  is always a minimizer of  $L(\cdot, P)$ .

**DRO.** Since  $L(\cdot, Q)$  is continuous and convex, the DRO objective  $F(\mathbf{A}) = \sup_{Q \in \mathcal{B}_\varepsilon(P)} L(\mathbf{A}, Q)$  is convex and lower semicontinuous as the supremum of continuous convex functions. It thus suffices to show that  $F$  is coercive. To prove this, consider the feasible measure  $Q = P * \mathcal{N}(0, \varepsilon^2 \text{diag}(\mathbf{I}_p, \mathbf{0}_{m \times m})/p)$ , which certifies that

$$F(\mathbf{A}) \geq L(\mathbf{A}, Q) = L(\mathbf{A}, P) + \frac{\varepsilon^2}{p} \mathbb{E}_{W \sim \mathcal{N}(0, \mathbf{I}_p)} \|\mathbf{A}W\|^2 \geq \frac{\varepsilon^2}{p} \|\mathbf{A}\|_{\text{F}}^2.$$

Thus  $\|\mathbf{A}\|_{\text{F}} \rightarrow \infty$  implies  $F(\mathbf{A}) \rightarrow \infty$ , and so  $F$  is coercive, as desired.

**MRO.** Similar to the above, continuity and convexity of  $R(\cdot, Q)$  imply that the MRO objective  $G(\mathbf{A}) = \sup_{Q \in \mathcal{B}_\varepsilon(P)} R(\mathbf{A}, Q)$  is convex and lower semicontinuous. It thus remains to show that  $G$  is coercive. To prove this, consider the same feasible measure  $Q$  as for DRO. Since  $Q$  is non-degenerate, we have

$$G(\mathbf{A}) \geq R(\mathbf{A}, Q) = \|(\mathbf{A} - \mathbf{A}_Q)(\Sigma_X^Q)^{1/2}\|_{\text{F}}^2 \geq \frac{\varepsilon^2}{p} \|\mathbf{A} - \mathbf{A}_Q\|_{\text{F}}^2.$$

Thus  $\|\mathbf{A}\|_{\text{F}} \rightarrow \infty$  implies  $G(\mathbf{A}) \rightarrow \infty$ , and so  $G$  is coercive, as desired.

**SNR.** The regularized objective  $H(\mathbf{A}) = L(\mathbf{A}, P) + \varepsilon^2 \|\mathbf{A}\|_{\text{op}}^2$  is clearly convex and lower semicontinuous. Moreover, since the risk is non-negative, it is also coercive.

## E Remaining Proofs for Section 2

### E.1 Condition Number Formula (Proof of Theorem 2.2)

We formulate the proof in Bures space, adopting the conventions of Appendix B. Initially, we find it simpler to analyze the alternative condition number  $\tilde{\kappa} : \mathcal{C}^* \rightarrow [0, \infty]$ , defined for  $\Sigma \in \mathcal{C}^*$  by

$$\tilde{\kappa}(\Sigma) := \limsup_{\varepsilon \downarrow 0} \sup_{\Sigma' \in \mathcal{B}_\varepsilon(\Sigma)} R(\mathbf{A}_{\Sigma'}, \Sigma), \quad (10)$$

which we will later show coincides with  $\kappa$ .

Our first step is to differentiate the parameter map  $g(\Sigma) := \Sigma_{YX} \Sigma_X^{-1}$ . To avoid boundary issues, we extend this map to the open domain  $\mathcal{C} := \{\Sigma \in \mathbb{S}^d : \Sigma_X \succ 0\}$ .

**Lemma E.1** (Parameter map derivative). *The map  $g : \mathcal{C} \rightarrow \mathcal{A}$  is continuously differentiable, with*

$$Dg(\Sigma)[\mathbf{H}] = (\mathbf{H}_{YX} - g(\Sigma)\mathbf{H}_X)\Sigma_X^{-1} \quad \forall \mathbf{H} \in \mathbb{S}^d.$$

*Proof.* The maps  $\Sigma \mapsto \Sigma_{YX}$  and  $\Sigma \mapsto \Sigma_X$  are linear, while matrix inversion is smooth on the positive definite cone. Hence  $g$  is continuously differentiable on  $\mathcal{C}$ . By the product rule and the standard derivative of the inverse map,

$$\begin{aligned} Dg(\Sigma)[\mathbf{H}] &= \mathbf{H}_{YX} \Sigma_X^{-1} - \Sigma_{YX} \Sigma_X^{-1} \mathbf{H}_X \Sigma_X^{-1} \\ &= (\mathbf{H}_{YX} - g(\Sigma)\mathbf{H}_X) \Sigma_X^{-1}, \end{aligned}$$

as claimed.  $\square$

We now characterize  $\tilde{\kappa}$  in terms of this derivative. We prove a slightly strengthened, uniform claim.

**Lemma E.2** (Derivative expression for  $\tilde{\kappa}$ ). *Fixing  $\Sigma \in \mathcal{C}^*$  with square root  $\mathbf{S} := \Sigma^{1/2}$ , we have*

$$\tilde{\kappa}(\Sigma) = \sup_{\Delta \in \mathbb{R}^{d \times d}: \|\Delta\|_F \leq 1} \|Dg(\Sigma)[\mathbf{S}\Delta^\top + \Delta\mathbf{S}]\Sigma_X^{1/2}\|_F^2. \quad (11)$$

Moreover, the expansion

$$\sup_{\Sigma' \in \mathcal{B}_\varepsilon(\Sigma)} R(g(\Sigma'), \Sigma) = \tilde{\kappa}(\Sigma)\varepsilon^2 + o(\varepsilon^2) \quad \varepsilon \rightarrow 0 \quad (12)$$

holds uniformly as  $\Sigma$  varies in any compact subset of  $\mathcal{C}^*$ .

*Proof.* Write  $\Sigma(\Delta) := (\mathbf{S} + \Delta)(\mathbf{S} + \Delta)^\top = \Sigma + \mathbf{S}\Delta^\top + \Delta\mathbf{S} + \Delta\Delta^\top$ , and let  $K$  be any compact subset of  $\mathcal{C}^*$ . By Lemma E.1,  $g$  is continuously differentiable on  $\mathcal{C} \supset K$ , and so

$$g(\Sigma + \mathbf{H}) - g(\Sigma) = Dg(\Sigma)[\mathbf{H}] + o(\|\mathbf{H}\|_F)$$

as  $\mathbf{H} \rightarrow \mathbf{0}$ , uniformly over  $\Sigma \in K$ . Further,

$$\|Dg(\Sigma)\|_{\text{op}} := \sup_{\|\mathbf{H}\|_F \leq 1} \|Dg(\Sigma)[\mathbf{H}]\|_F$$

is bounded for  $\Sigma \in K$ . Further

$$\|\Sigma(\Delta) - \Sigma\|_F \leq 2\|\mathbf{S}\|_F\|\Delta\|_F + \|\Delta\|_F^2 = 2\sqrt{\text{tr}(\Sigma)}\|\Delta\|_F + \|\Delta\|_F^2 = O(\|\Delta\|_F)$$

as  $\Delta \rightarrow \mathbf{0}$ , uniformly for  $\Sigma \in K$ . Combining gives

$$\begin{aligned} g(\Sigma(\Delta)) - g(\Sigma) &= Dg(\Sigma)[\mathbf{S}\Delta^\top + \Delta\mathbf{S} + \Delta\Delta^\top] + o(\|\Delta\|_F) \\ &= Dg(\Sigma)[\mathbf{S}\Delta^\top + \Delta\mathbf{S}] + (1 + \|Dg(\Sigma)\|_{\text{op}})o(\|\Delta\|_F) \\ &= Dg(\Sigma)[\mathbf{S}\Delta^\top + \Delta\mathbf{S}] + o(\|\Delta\|_F) \end{aligned}$$

as  $\Delta \rightarrow \mathbf{0}$ , uniformly over  $\Sigma \in K$ . Hence,

$$\begin{aligned} \|(g(\Sigma(\Delta)) - g(\Sigma))\Sigma_X^{1/2}\|_F^2 &= \|(Dg(\Sigma)[\mathbf{S}\Delta^\top + \Delta\mathbf{S}] + o(\|\Delta\|_F))\Sigma_X^{1/2}\|_F^2 \\ &= \|Dg(\Sigma)[\mathbf{S}\Delta^\top + \Delta\mathbf{S}]\Sigma_X^{1/2}\|_F^2 + E(\Sigma, \Delta) \end{aligned}$$

where, by Cauchy–Schwarz and submultiplicativity of the Frobenius norm

$$\begin{aligned} |E(\Sigma, \Delta)| &\leq 2\|Dg(\Sigma)\|_{\text{op}}\|\mathbf{S}\Delta^\top + \Delta\mathbf{S}\|_F o(\|\Delta\|_F)\|\Sigma_X\|_F + o(\|\Delta\|_F^2)\|\Sigma_X\|_F \\ &\leq 4\|Dg(\Sigma)\|_{\text{op}}\sqrt{\text{tr}(\Sigma)}\|\Delta\|_F o(\|\Delta\|_F)\|\Sigma_X\|_F + o(\|\Delta\|_F^2)\|\Sigma_X\|_F \\ &= o(\|\Delta\|_F^2) \end{aligned}$$

as  $\Delta \rightarrow \mathbf{0}$ , uniformly over  $\Sigma \in K$ . Combining this with Lemma 1.1 gives

$$\begin{aligned} \sup_{\Sigma' \in \mathcal{B}_\varepsilon(\Sigma)} R(g(\Sigma'), \Sigma) &= \sup_{\Delta \in \mathbb{R}^{d \times d}: \|\Delta\|_F \leq \varepsilon} R(g(\Sigma(\Delta)), \Sigma) \\ &= \sup_{\Delta \in \mathbb{R}^{d \times d}: \|\Delta\|_F \leq \varepsilon} \|(g(\Sigma(\Delta)) - g(\Sigma))\Sigma_X^{1/2}\|_F^2 \\ &= \sup_{\Delta \in \mathbb{R}^{d \times d}: \|\Delta\|_F \leq \varepsilon} \|(Dg(\Sigma)[\mathbf{S}\Delta^\top + \Delta\mathbf{S}] + o(\|\Delta\|_F))\Sigma_X^{1/2}\|_F^2 \\ &= \sup_{\Delta \in \mathbb{R}^{d \times d}: \|\Delta\|_F \leq 1} \|Dg(\Sigma)[\mathbf{S}\Delta^\top + \Delta\mathbf{S}]\Sigma_X^{1/2}\|_F^2 \varepsilon^2 + o(\varepsilon^2), \end{aligned}$$

as  $\varepsilon \rightarrow 0$ , uniformly over  $\Sigma \in K$ . The lemma then follows by the definition of  $\tilde{\kappa}$ .  $\square$

Next, we substitute the derivative formula from Lemma E.1 into (11) and simplify. We use the variational form of the Frobenius norm  $\|\mathbf{X}\|_F = \sup_{\mathbf{Y}: \|\mathbf{Y}\|_F \leq 1} \langle \mathbf{X}, \mathbf{Y} \rangle_F$  to streamline calculations.

**Lemma E.3** (Closed form for  $\tilde{\kappa}$ ). *For  $\Sigma \in \mathcal{C}^*$ , we have*

$$\tilde{\kappa}(\Sigma) = 1 + \lambda_{\max}(\mathbf{A}_\Sigma \mathbf{A}_\Sigma^\top + \lambda_{\min}(\Sigma_X)^{-1} \Sigma_E).$$

*In particular, we have*

$$1 + \max \left\{ \|\mathbf{A}_\Sigma\|_{\text{op}}^2, \frac{\lambda_{\max}(\Sigma_E)}{\lambda_{\min}(\Sigma_X)} \right\} \leq \tilde{\kappa}(\Sigma) \leq 1 + \|\mathbf{A}_\Sigma\|_{\text{op}}^2 + \frac{\lambda_{\max}(\Sigma_E)}{\lambda_{\min}(\Sigma_X)}.$$

*Proof.* Abbreviate  $\mathbf{A} := g(\Sigma) = \Sigma_{YX} \Sigma_X^{-1}$  and  $\mathbf{S} := \Sigma^{1/2}$ . Combining Lemmas E.1 and E.2 gives

$$\begin{aligned} \tilde{\kappa}(\Sigma) &= \sup_{\Delta \in \mathbb{R}^{d \times d}: \|\Delta\|_F \leq 1} \|Dg(\Sigma)[\mathbf{S} \Delta^\top + \Delta \mathbf{S}] \Sigma_X^{1/2}\|_F^2 \\ &= \sup_{\Delta \in \mathbb{R}^{d \times d}: \|\Delta\|_F \leq 1} \left\| (\mathbf{S}_Y \Delta_X^\top + \Delta_Y \mathbf{S}_X^\top - \mathbf{A}(\mathbf{S}_X \Delta_X^\top + \Delta_X \mathbf{S}_X^\top)) \Sigma_X^{-1/2} \right\|_F^2 \\ &= \sup_{\Delta \in \mathbb{R}^{d \times d}: \|\Delta\|_F \leq 1} \left\| ((\Delta_Y - \mathbf{A} \Delta_X) \mathbf{S}_X^\top + (\mathbf{S}_Y - \mathbf{A} \mathbf{S}_X) \Delta_X^\top) \Sigma_X^{-1/2} \right\|_F^2. \end{aligned}$$

To simplify the objective, write  $\mathbf{Q} := \mathbf{S}_X^\top \Sigma_X^{-1/2} \in \mathbb{R}^{d \times p}$ , noting that  $\mathbf{Q}^\top \mathbf{Q} = \mathbf{I}_p$  and that  $\mathbf{Q} \mathbf{Q}^\top \in \mathbb{R}^{d \times d}$  is the orthogonal projection onto the row space of  $\mathbf{S}_X$ . Further define the residual factor  $\mathbf{E} := \mathbf{S}_Y - \mathbf{A} \mathbf{S}_X \in \mathbb{R}^{m \times d}$ , which by Lemma B.1 satisfies  $\mathbf{E} \mathbf{E}^\top = \Sigma_E$  and  $\mathbf{E} \mathbf{S}_X^\top = \mathbf{0}$  (implying  $\mathbf{E} \mathbf{Q} = \mathbf{0}$ ). Substituting these definitions into the equation above and rearranging, we obtain

$$\begin{aligned} \tilde{\kappa}(\Sigma) &= \sup_{\|\Delta\|_F \leq 1} \left\| (\Delta_Y - \mathbf{A} \Delta_X) \mathbf{Q} + \mathbf{E} \Delta_X^\top \Sigma_X^{-1/2} \right\|_F^2 \\ &= \sup_{\|\Delta\|_F \leq 1} \left\| \underbrace{\mathbf{E} \Delta_X^\top \Sigma_X^{-1/2} - \mathbf{A} \Delta_X \mathbf{Q}}_{\text{feature perturbations}} + \underbrace{\Delta_Y \mathbf{Q}}_{\text{response perturbations}} \right\|_F^2. \end{aligned}$$

The first component depends on the perturbation  $\Delta_X$  to the feature factor  $\mathbf{S}_X$ , while the second component depends on the perturbation  $\Delta_Y$  to the response factor  $\mathbf{S}_Y$ .

To simplify further, we use the variational form of the Frobenius norm. For  $\mathbf{B} \in \mathbb{R}^{m \times p}$ , define

$$\phi(\mathbf{B}; \Delta) := \langle \mathbf{B}, \mathbf{E} \Delta_X^\top \Sigma_X^{-1/2} - \mathbf{A} \Delta_X \mathbf{Q} + \Delta_Y \mathbf{Q} \rangle_F.$$

Since the Frobenius unit ball is symmetric, we have

$$\tilde{\kappa}(\Sigma) = \sup_{\|\Delta\|_F \leq 1} \left( \sup_{\|\mathbf{B}\|_F \leq 1} \phi(\mathbf{B}; \Delta) \right)^2 = \sup_{\|\Delta\|_F \leq 1} \sup_{\|\mathbf{B}\|_F \leq 1} \phi(\mathbf{B}; \Delta)^2.$$

We now swap these suprema and maximize over  $\Delta$  for fixed  $\mathbf{B}$ . Grouping terms according to the perturbed row block and using cyclicity of the trace gives

$$\phi(\mathbf{B}; \Delta) = \langle \Sigma_X^{-1/2} \mathbf{B}^\top \mathbf{E} - \mathbf{A}^\top \mathbf{B} \mathbf{Q}^\top, \Delta_X \rangle_F + \langle \mathbf{B} \mathbf{Q}^\top, \Delta_Y \rangle_F.$$

Since the Frobenius norm is self-dual,

$$\begin{aligned} \sup_{\|\Delta\|_F \leq 1} \phi(\mathbf{B}; \Delta)^2 &= \left\| \begin{pmatrix} \Sigma_X^{-1/2} \mathbf{B}^\top \mathbf{E} - \mathbf{A}^\top \mathbf{B} \mathbf{Q}^\top \\ \mathbf{B} \mathbf{Q}^\top \end{pmatrix} \right\|_F^2 \\ &= \|\Sigma_X^{-1/2} \mathbf{B}^\top \mathbf{E} - \mathbf{A}^\top \mathbf{B} \mathbf{Q}^\top\|_F^2 + \|\mathbf{B} \mathbf{Q}^\top\|_F^2 \\ &= \underbrace{\|\Sigma_X^{-1/2} \mathbf{B}^\top \mathbf{E}\|_F^2 + \|\mathbf{A}^\top \mathbf{B}\|_F^2}_{\text{feature contribution}} + \underbrace{\|\mathbf{B}\|_F^2}_{\text{response contribution}}. \end{aligned}$$

The last equality uses  $\mathbf{E} \mathbf{Q} = \mathbf{0}$  to eliminate the cross term within the feature contribution and applies  $\mathbf{Q}^\top \mathbf{Q} = \mathbf{I}_p$  to simplify the other norms. For fixed  $\mathbf{B}$ , restricting to  $\Delta_Y = \mathbf{0}$  leaves exactly the first bracket, while restricting to  $\Delta_X = \mathbf{0}$  leaves the final term.

By homogeneity, the outer supremum may be restricted to  $\|\mathbf{B}\|_F = 1$ . The response contribution is then 1 for every  $\mathbf{B}$ . Using  $\mathbf{E}\mathbf{E}^\top = \boldsymbol{\Sigma}_E$  and Lemma A.8, we obtain

$$\begin{aligned}\tilde{\kappa}(\boldsymbol{\Sigma}) &= 1 + \sup_{\|\mathbf{B}\|_F=1} \left\{ \text{tr}(\boldsymbol{\Sigma}_X^{-1} \mathbf{B}^\top \boldsymbol{\Sigma}_E \mathbf{B}) + \text{tr}(\mathbf{A}\mathbf{A}^\top \mathbf{B}\mathbf{B}^\top) \right\} \\ &= 1 + \sup_{\|\mathbf{B}\|_F=1} \text{vec}(\mathbf{B})^\top (\boldsymbol{\Sigma}_X^{-1} \otimes \boldsymbol{\Sigma}_E + \mathbf{I}_p \otimes \mathbf{A}\mathbf{A}^\top) \text{vec}(\mathbf{B}) \\ &= 1 + \left\| \boldsymbol{\Sigma}_X^{-1} \otimes \boldsymbol{\Sigma}_E + \mathbf{I}_p \otimes \mathbf{A}\mathbf{A}^\top \right\|_{\text{op}}.\end{aligned}$$

The remaining operator norm measures worst-case sensitivity to feature perturbations alone.

To compute it, eigendecompose  $\boldsymbol{\Sigma}_X = \mathbf{U} \text{diag}(s_1, \dots, s_p) \mathbf{U}^\top$ , where  $\mathbf{U} \in \mathbb{O}^p$  and  $0 < s_1 \leq \dots \leq s_p$ . Conjugating by  $\mathbf{U} \otimes \mathbf{I}_m$  gives a block-diagonal matrix with blocks  $s_i^{-1} \boldsymbol{\Sigma}_E + \mathbf{A}\mathbf{A}^\top$ . Since  $s_i^{-1} \leq s_1^{-1}$  and  $\boldsymbol{\Sigma}_E \succeq \mathbf{0}$ , the first block dominates all others in Loewner order. Consequently,

$$\tilde{\kappa}(\boldsymbol{\Sigma}) = \underbrace{1}_{\text{response contribution}} + \underbrace{\lambda_{\max}(\mathbf{A}\mathbf{A}^\top + s_1^{-1} \boldsymbol{\Sigma}_E)}_{\text{feature contribution}}.$$

This proves the formula since  $s_1 = \lambda_{\min}(\boldsymbol{\Sigma}_X)$ .

Finally, positive semidefiniteness of the two matrices inside  $\lambda_{\max}$  and subadditivity of  $\lambda_{\max}$  give

$$\max\{\|\mathbf{A}\|_{\text{op}}^2, s_1^{-1} \lambda_{\max}(\boldsymbol{\Sigma}_E)\} \leq \tilde{\kappa}(\boldsymbol{\Sigma}) - 1 \leq \|\mathbf{A}\|_{\text{op}}^2 + s_1^{-1} \lambda_{\max}(\boldsymbol{\Sigma}_E),$$

establishing the remaining bounds.  $\square$

We next show that  $\kappa$  and  $\tilde{\kappa}$  coincide. The upper bound is relatively direct in light of (12).

**Lemma E.4** (Condition number comparison, upper bound). *For all  $\boldsymbol{\Sigma} \in \mathcal{C}^*$ ,  $\kappa(\boldsymbol{\Sigma}) \leq \tilde{\kappa}(\boldsymbol{\Sigma})$ .*

*Proof.* We upper bound  $\kappa(\boldsymbol{\Sigma})$  in (4) by considering the feasible parameter  $\mathbf{A} = \mathbf{A}_{\boldsymbol{\Sigma}} = g(\boldsymbol{\Sigma})$ . This gives

$$\begin{aligned}\kappa(\boldsymbol{\Sigma}) &\leq \limsup_{\varepsilon \downarrow 0} \varepsilon^{-2} \sup_{\boldsymbol{\Sigma}' \in \mathcal{B}_\varepsilon(\boldsymbol{\Sigma})} R(g(\boldsymbol{\Sigma}), \boldsymbol{\Sigma}') \\ &\leq \limsup_{\varepsilon \downarrow 0} \varepsilon^{-2} \sup_{\boldsymbol{\Sigma}' \in \mathcal{B}_\varepsilon(\boldsymbol{\Sigma})} \sup_{\boldsymbol{\Sigma}'' \in \mathcal{B}_\varepsilon(\boldsymbol{\Sigma}')} R(g(\boldsymbol{\Sigma}''), \boldsymbol{\Sigma}').\end{aligned}$$

Noting that the convergence in (12) is uniform over  $\boldsymbol{\Sigma}' \in \mathcal{B}_\varepsilon(\boldsymbol{\Sigma})$ , that expansion implies

$$\begin{aligned}\kappa(\boldsymbol{\Sigma}) &\leq \limsup_{\varepsilon \downarrow 0} \varepsilon^{-2} \left( \sup_{\boldsymbol{\Sigma}' \in \mathcal{B}_\varepsilon(\boldsymbol{\Sigma})} \tilde{\kappa}(\boldsymbol{\Sigma}') \varepsilon^2 + o(\varepsilon^2) \right) \\ &= \limsup_{\varepsilon \downarrow 0} \sup_{\boldsymbol{\Sigma}' \in \mathcal{B}_\varepsilon(\boldsymbol{\Sigma})} \tilde{\kappa}(\boldsymbol{\Sigma}') \\ &= \tilde{\kappa}(\boldsymbol{\Sigma}).\end{aligned}$$

The final equality uses continuity of  $\tilde{\kappa}$ , which is evident from our closed form in Lemma E.3.  $\square$

We prove a slightly stronger result in the other direction (using  $\liminf$  instead of  $\limsup$  below).

**Lemma E.5** (Condition number comparison, lower bound). *For all  $\boldsymbol{\Sigma} \in \mathcal{C}^*$ , we have*

$$\kappa(\boldsymbol{\Sigma}) \geq J_{\boldsymbol{\Sigma}} := \liminf_{\varepsilon \downarrow 0} \varepsilon^{-2} \inf_{\mathbf{A} \in \mathcal{A}} \sup_{\boldsymbol{\Sigma}' \in \mathcal{B}_\varepsilon(\boldsymbol{\Sigma})} R(\mathbf{A}, \boldsymbol{\Sigma}') \geq \tilde{\kappa}(\boldsymbol{\Sigma}).$$

*Proof.* The first inequality is immediate from (4), since  $\limsup \geq \liminf$ .

For the second inequality, fix  $\boldsymbol{\Delta}_* \in \mathbb{R}^{d \times d}$  with  $\|\boldsymbol{\Delta}_*\|_F = 1$  achieving the supremum in formula (11) for  $\tilde{\kappa}(\boldsymbol{\Sigma})$ , noting that a maximizer exists by compactness of the Frobenius ball and continuity

of the objective. Now, for each  $t \in \mathbb{R}$ , define  $\Sigma^t := (\Sigma^{1/2} + t\Delta_\star)(\Sigma^{1/2} + t\Delta_\star)^\top$ , which satisfies  $B(\Sigma^t, \Sigma) \leq |t|$  by Lemma 1.1. In particular,  $\Sigma^{\pm\epsilon} \in \mathcal{B}_\epsilon(\Sigma)$ . We thus have

$$\begin{aligned} J_\Sigma &:= \liminf_{\epsilon \downarrow 0} \epsilon^{-2} \inf_{\mathbf{A} \in \mathcal{A}} \sup_{\Sigma' \in \mathcal{B}_\epsilon(\Sigma)} R(\mathbf{A}, \Sigma') \\ &\geq \liminf_{\epsilon \downarrow 0} \epsilon^{-2} \inf_{\mathbf{A} \in \mathcal{A}} \max\{R(\mathbf{A}, \Sigma^{+\epsilon}), R(\mathbf{A}, \Sigma^{-\epsilon})\} \\ &= \liminf_{\epsilon \downarrow 0} \epsilon^{-2} \inf_{\mathbf{A} \in \mathcal{A}} \max\left\{\|(\mathbf{A} - \mathbf{A}_{+\epsilon})(\Sigma_X^{+\epsilon})^{1/2}\|_F^2, \|(\mathbf{A} - \mathbf{A}_{-\epsilon})(\Sigma_X^{-\epsilon})^{1/2}\|_F^2\right\}, \end{aligned} \quad (13)$$

where  $\mathbf{A}_t := g(\Sigma^t)$ . Since  $\Sigma_X^t \rightarrow \Sigma_X$  as  $t \rightarrow 0$ , we have for fixed  $\lambda \in (0, 1)$  that  $\Sigma_X^t \succ \lambda \Sigma_X$  for sufficiently small  $t$ . Consequently, we have

$$\begin{aligned} J_\Sigma &\geq \sup_{\lambda \in (0, 1)} \lambda \liminf_{\epsilon \downarrow 0} \epsilon^{-2} \inf_{\mathbf{A} \in \mathcal{A}} \max\left\{\|(\mathbf{A} - \mathbf{A}_{+\epsilon})\Sigma_X^{1/2}\|_F^2, \|(\mathbf{A} - \mathbf{A}_{-\epsilon})\Sigma_X^{1/2}\|_F^2\right\} \\ &= \liminf_{\epsilon \downarrow 0} \epsilon^{-2} \inf_{\mathbf{A} \in \mathcal{A}} \max\left\{\|(\mathbf{A} - \mathbf{A}_{+\epsilon})\Sigma_X^{1/2}\|_F^2, \|(\mathbf{A} - \mathbf{A}_{-\epsilon})\Sigma_X^{1/2}\|_F^2\right\}. \end{aligned}$$

By Lemma A.7, the infimum is attained by the midpoint  $\bar{\mathbf{A}}_\epsilon = (\mathbf{A}_{+\epsilon} + \mathbf{A}_{-\epsilon})/2$ . This gives,

$$J_\Sigma \geq \frac{1}{4} \liminf_{\epsilon \downarrow 0} \epsilon^{-2} \|(\mathbf{A}_{+\epsilon} - \mathbf{A}_{-\epsilon})\Sigma_X^{1/2}\|_F^2$$

By construction, we have

$$\Sigma^{+\epsilon} - \Sigma^{-\epsilon} = 2\epsilon(\Sigma^{1/2}\Delta_\star^\top + \Delta_\star\Sigma^{1/2}),$$

which, since  $g$  is continuously differentiable, gives

$$\begin{aligned} \mathbf{A}_{+\epsilon} - \mathbf{A}_{-\epsilon} &= g(\Sigma^{+\epsilon}) - g(\Sigma^{-\epsilon}) \\ &= Dg(\Sigma)[\Sigma^{+\epsilon} - \Sigma] - Dg(\Sigma)[\Sigma^{-\epsilon} - \Sigma] + o(\epsilon) \\ &= Dg(\Sigma)[\Sigma^{+\epsilon} - \Sigma^{-\epsilon}] + o(\epsilon) \\ &= 2\epsilon Dg(\Sigma)[\Sigma^{1/2}\Delta_\star^\top + \Delta_\star\Sigma^{1/2}] + o(\epsilon) \end{aligned}$$

as  $\epsilon \rightarrow 0$ . Substituting this expansion into the preceding lower bound for  $J_\Sigma$  and using optimality of  $\Delta_\star$ , we obtain

$$\begin{aligned} J_\Sigma &\geq \|Dg(\Sigma)[\Sigma^{1/2}\Delta_\star^\top + \Delta_\star\Sigma^{1/2}]\Sigma_X^{1/2}\|_F^2 \\ &= \sup_{\Delta \in \mathbb{R}^{d \times d}: \|\Delta\|_F \leq 1} \|Dg(\Sigma)[\Sigma^{1/2}\Delta^\top + \Delta\Sigma^{1/2}]\Sigma_X^{1/2}\|_F^2 \\ &= \tilde{\kappa}(\Sigma), \end{aligned}$$

as desired.  $\square$

Combining Lemma E.4 and Lemma E.5 with (4) gives

$$\kappa(\Sigma) = \tilde{\kappa}(\Sigma) = \lim_{\epsilon \downarrow 0} \epsilon^{-2} \inf_{\mathbf{A} \in \mathcal{A}} \sup_{\Sigma' \in \mathcal{B}_\epsilon(\Sigma)} R(\mathbf{A}, \Sigma') = \lim_{\epsilon \downarrow 0} \epsilon^{-2} R_\epsilon^*(\Sigma)$$

for  $\Sigma \in \mathcal{C}^*$ . The theorem then follows by the Lemma E.3 formula and inequalities.  $\square$

## F Remaining Proofs for Section 3

### F.1 Regret Transfer (Proof of Lemma 3.1)

We first establish a refined bound by comparing projected residual factors. The argument separates the direct effect of the perturbation from the change in the feature row space, and it allows the perturbed feature second-moment matrix to be singular.

**Lemma F.1** (Regret stability, refined). *For all  $\epsilon \geq 0$ ,  $\mathbf{A} \in \mathcal{A}$ ,  $P \in \mathcal{P}^*$ , and  $Q \in \mathcal{B}_\epsilon(P)$ , we have*

$$\left| \sqrt{R(\mathbf{A}, P)} - \sqrt{R(\mathbf{A}, Q)} \right| \leq \left( \max\left\{\|\mathbf{A} - \mathbf{A}_P\|_{\text{op}}, \sqrt{1 + \|\mathbf{A}\|_{\text{op}}^2}\right\} + \sqrt{\frac{\lambda_{\max}(\Sigma_E^P)}{\lambda_{\min}(\Sigma_X^P)}} \right) \epsilon.$$

*Proof.* We adopt the Bures-space conventions and block decompositions of Appendix B. Set  $\Sigma := \Sigma_P$  and  $\bar{\Sigma} := \Sigma_Q$ . By Lemmas 1.1 and 1.2, there are square-root factors  $\mathbf{S}, \bar{\mathbf{S}} \in \mathbb{R}^{d \times d}$  such that

$$\Sigma = \mathbf{S}\mathbf{S}^\top, \quad \bar{\Sigma} = \bar{\mathbf{S}}\bar{\mathbf{S}}^\top, \quad \Delta := \bar{\mathbf{S}} - \mathbf{S}, \quad \|\Delta\|_F \leq \varepsilon.$$

Let  $\Pi_X$  and  $\bar{\Pi}_X$  denote the orthogonal projections onto the row spaces of  $\mathbf{S}_X$  and  $\bar{\mathbf{S}}_X$ , respectively, and define the residual factors

$$\mathbf{E}_A := \mathbf{S}_Y - \mathbf{A}\mathbf{S}_X, \quad \bar{\mathbf{E}}_A := \bar{\mathbf{S}}_Y - \mathbf{A}\bar{\mathbf{S}}_X, \quad \mathbf{E} := \mathbf{E}_{\mathbf{A}\Sigma}.$$

By Lemma B.1, we have

$$\sqrt{R(\mathbf{A}, \Sigma)} = \|\mathbf{E}_A \Pi_X\|_F, \quad \sqrt{R(\mathbf{A}, \bar{\Sigma})} = \|\bar{\mathbf{E}}_A \bar{\Pi}_X\|_F,$$

and the optimal residual factor  $\mathbf{E}$  satisfies  $\mathbf{E}\Pi_X = \mathbf{0}$  and  $\mathbf{E}\mathbf{E}^\top = \Sigma_E$ .

We compare these projected residuals directly. Since  $\mathbf{S}_X \Pi_X = \mathbf{S}_X$  and  $\bar{\mathbf{S}}_X \bar{\Pi}_X = \bar{\mathbf{S}}_X$ , we have

$$\begin{aligned} \mathbf{E}_A \Pi_X - \bar{\mathbf{E}}_A \bar{\Pi}_X &= (\mathbf{S}_Y - \mathbf{A}\mathbf{S}_X) \Pi_X - (\bar{\mathbf{S}}_Y - \mathbf{A}\bar{\mathbf{S}}_X) \bar{\Pi}_X \\ &= \mathbf{S}_Y (\Pi_X - \bar{\Pi}_X) - \Delta_Y \bar{\Pi}_X + \mathbf{A} \Delta_X. \end{aligned}$$

Since our desired bound does not depend on  $\mathbf{S}_Y$  or  $\Sigma_Y$  directly, we substitute  $\mathbf{S}_Y = \mathbf{A}_\Sigma \mathbf{S}_X + \mathbf{E}$  for the first term to obtain

$$\begin{aligned} \mathbf{S}_Y (\Pi_X - \bar{\Pi}_X) &= (\mathbf{A}_\Sigma \mathbf{S}_X + \mathbf{E}) (\Pi_X - \bar{\Pi}_X) \\ &= \mathbf{A}_\Sigma \mathbf{S}_X - \mathbf{A}_\Sigma \mathbf{S}_X \bar{\Pi}_X - \mathbf{E} \bar{\Pi}_X \quad (\mathbf{S}_X \Pi_X = \mathbf{S}_X, \mathbf{E} \Pi_X = \mathbf{0}) \\ &= \mathbf{A}_\Sigma \mathbf{S}_X (\mathbf{I}_d - \bar{\Pi}_X) - \mathbf{E} \bar{\Pi}_X \\ &= \mathbf{A}_\Sigma \bar{\mathbf{S}}_X (\mathbf{I}_d - \bar{\Pi}_X) - \mathbf{A}_\Sigma \Delta_X (\mathbf{I}_d - \bar{\Pi}_X) - \mathbf{E} \bar{\Pi}_X \\ &= -\mathbf{A}_\Sigma \Delta_X (\mathbf{I}_d - \bar{\Pi}_X) - \mathbf{E} \bar{\Pi}_X, \end{aligned}$$

where the last equality uses  $\bar{\mathbf{S}}_X (\mathbf{I}_d - \bar{\Pi}_X) = \mathbf{0}$ . Substituting this into the previous equation gives

$$\begin{aligned} \mathbf{E}_A \Pi_X - \bar{\mathbf{E}}_A \bar{\Pi}_X &= -\mathbf{A}_\Sigma \Delta_X (\mathbf{I}_d - \bar{\Pi}_X) - \mathbf{E} \bar{\Pi}_X - \Delta_Y \bar{\Pi}_X + \mathbf{A} \Delta_X \\ &= (\mathbf{A} - \mathbf{A}_\Sigma) \Delta_X (\mathbf{I}_d - \bar{\Pi}_X) + (\mathbf{A} \Delta_X - \Delta_Y) \bar{\Pi}_X - \mathbf{E} \bar{\Pi}_X. \end{aligned}$$

Let  $\mathbf{T}$  denote the sum of the first two terms in the last line, and write

$$c := \max \left\{ \|\mathbf{A} - \mathbf{A}_\Sigma\|_{\text{op}}, \sqrt{1 + \|\mathbf{A}\|_{\text{op}}^2} \right\}.$$

Note that  $1 + \|\mathbf{A}\|_{\text{op}}^2 = \lambda_{\max}(\mathbf{I}_m + \mathbf{A}\mathbf{A}^\top) = \|(\mathbf{A} \quad -\mathbf{I}_m)\|_{\text{op}}^2$ . Moreover, the two terms of  $\mathbf{T}$  have orthogonal right projections, so we obtain

$$\begin{aligned} \|\mathbf{T}\|_F^2 &= \|(\mathbf{A} - \mathbf{A}_\Sigma) \Delta_X (\mathbf{I}_d - \bar{\Pi}_X)\|_F^2 + \|(\mathbf{A} \Delta_X - \Delta_Y) \bar{\Pi}_X\|_F^2 \\ &\leq \|\mathbf{A} - \mathbf{A}_\Sigma\|_{\text{op}}^2 \|\Delta_X (\mathbf{I}_d - \bar{\Pi}_X)\|_F^2 + \|(\mathbf{A} \quad -\mathbf{I}_m)\|_{\text{op}}^2 \|\Delta \bar{\Pi}_X\|_F^2 \quad (\text{Lemma A.1}) \\ &\leq c^2 (\|\Delta_X (\mathbf{I}_d - \bar{\Pi}_X)\|_F^2 + \|\Delta_X \bar{\Pi}_X\|_F^2 + \|\Delta_Y \bar{\Pi}_X\|_F^2) \\ &\leq c^2 (\|\Delta_X\|_F^2 + \|\Delta_Y\|_F^2) = c^2 \|\Delta\|_F^2. \end{aligned}$$

The last inequality uses the orthogonal decomposition of  $\Delta_X$  and the contraction property of  $\bar{\Pi}_X$ .

It remains to bound  $\mathbf{E} \bar{\Pi}_X$ , the component of the original optimal residual in the perturbed feature row space. Since  $\mathbf{E} \Pi_X = \mathbf{0}$ , the row-space perturbation bound (Lemma A.6) gives

$$\begin{aligned} \|\mathbf{E} \bar{\Pi}_X\|_F &= \|\mathbf{E} (\mathbf{I}_d - \Pi_X) \bar{\Pi}_X\|_F \\ &\leq \|\mathbf{E}\|_{\text{op}} \|(\mathbf{I}_d - \Pi_X) \bar{\Pi}_X\|_F \\ &\leq \frac{\|\mathbf{E}\|_{\text{op}}}{\sigma_{\min}(\mathbf{S}_X)} \|\Delta_X\|_F \\ &\leq \sqrt{\frac{\lambda_{\max}(\Sigma_E)}{\lambda_{\min}(\Sigma_X)}} \varepsilon. \end{aligned}$$

Here  $\mathbf{S}_X$  has full row rank because  $\Sigma_X \succ \mathbf{0}$ ; no rank assumption on  $\bar{\mathbf{S}}_X$  is needed.

Finally, the reverse triangle inequality yields

$$\begin{aligned}
\left| \sqrt{R(\mathbf{A}, \boldsymbol{\Sigma})} - \sqrt{R(\mathbf{A}, \bar{\boldsymbol{\Sigma}})} \right| &\leq \|\mathbf{E}_{\mathbf{A}} \boldsymbol{\Pi}_X - \bar{\mathbf{E}}_{\mathbf{A}} \bar{\boldsymbol{\Pi}}_X\|_F \\
&\leq \|\mathbf{T}\|_F + \|\mathbf{E} \bar{\boldsymbol{\Pi}}_X\|_F \\
&\leq \left( c + \sqrt{\frac{\lambda_{\max}(\boldsymbol{\Sigma}_E)}{\lambda_{\min}(\boldsymbol{\Sigma}_X)}} \right) \varepsilon.
\end{aligned}$$

Translating back to  $P$  and  $Q$  proves the claim.  $\square$

*Proof of Lemma 3.1.* By Theorem 2.2,

$$\sqrt{\kappa_P} \geq \max \left\{ \sqrt{1 + \|\mathbf{A}_P\|_{\text{op}}^2}, \sqrt{\frac{\lambda_{\max}(\boldsymbol{\Sigma}_E^P)}{\lambda_{\min}(\boldsymbol{\Sigma}_X^P)}} \right\}.$$

Moreover, the triangle inequality for  $\|\cdot\|_{\text{op}}$  and subadditivity of the square root imply

$$\begin{aligned}
\max \left\{ \|\mathbf{A} - \mathbf{A}_P\|_{\text{op}}, \sqrt{1 + \|\mathbf{A}\|_{\text{op}}^2} \right\} &\leq \|\mathbf{A}\|_{\text{op}} + \max \{ \|\mathbf{A}_P\|_{\text{op}}, 1 \} \\
&\leq \|\mathbf{A}\|_{\text{op}} + \sqrt{\kappa_P}.
\end{aligned}$$

Applying Lemma F.1 with  $\varepsilon := W_2(P, Q)$  thus gives

$$\left| \sqrt{R(\mathbf{A}, P)} - \sqrt{R(\mathbf{A}, Q)} \right| \leq (\|\mathbf{A}\|_{\text{op}} + 2\sqrt{\kappa_P}) W_2(P, Q).$$

For  $\mathbf{A} = \mathbf{A}_P$ , we have  $R(\mathbf{A}_P, P) = 0$ , and so Lemma F.1 gives

$$\begin{aligned}
R(\mathbf{A}_P, Q) &\leq \left( \sqrt{1 + \|\mathbf{A}_P\|_{\text{op}}^2} + \sqrt{\frac{\lambda_{\max}(\boldsymbol{\Sigma}_E^P)}{\lambda_{\min}(\boldsymbol{\Sigma}_X^P)}} \right)^2 W_2(P, Q)^2 \\
&\leq 4\kappa_P W_2(P, Q)^2,
\end{aligned}$$

proving the second claim.  $\square$

## F.2 Stability Implications for ERM and DRO (Remaining Proof for Lemma 3.2)

We employ a few helper lemmas, adopting the notation and block decompositions from Appendix B.

**Lemma F.2** (Perturbed condition number). *Fix  $\boldsymbol{\Sigma} \in \mathcal{C}^*$  and  $\bar{\boldsymbol{\Sigma}} \in \mathbb{S}_+^d$  with  $B(\bar{\boldsymbol{\Sigma}}, \boldsymbol{\Sigma}) \leq \sqrt{\lambda_{\min}(\boldsymbol{\Sigma}_X)}/2$ . Then  $\bar{\boldsymbol{\Sigma}} \in \mathcal{C}^*$  and  $\kappa(\bar{\boldsymbol{\Sigma}}) \leq 19\kappa(\boldsymbol{\Sigma})$ .*

*Proof.* By Theorem 2.2, the condition number is determined up to constants by the operator norm of the optimal coefficient matrix, the maximum eigenvalue of the residual second-moment matrix, and the minimum eigenvalue of the feature second-moment matrix. We now control each of these quantities for  $\bar{\boldsymbol{\Sigma}}$  in terms of those for  $\boldsymbol{\Sigma}$ , beginning with the minimum feature eigenvalue.

Set  $\varepsilon := B(\bar{\boldsymbol{\Sigma}}, \boldsymbol{\Sigma})$ ,  $\lambda := \lambda_{\min}(\boldsymbol{\Sigma}_X) > 0$  and  $\bar{\lambda} := \lambda_{\min}(\bar{\boldsymbol{\Sigma}}_X)$ . By Lemma B.3,

$$\sqrt{\bar{\lambda}} \geq \sqrt{\lambda} - \varepsilon \geq \sqrt{\lambda}/2,$$

since  $\varepsilon \leq \sqrt{\lambda}/2$ . Thus  $\bar{\lambda} \geq \lambda/4$  and  $\bar{\boldsymbol{\Sigma}} \in \mathcal{C}^*$ .

Next, let  $\mathbf{A} := \mathbf{A}_{\boldsymbol{\Sigma}}$ ,  $\bar{\mathbf{A}} := \mathbf{A}_{\bar{\boldsymbol{\Sigma}}}$ , and  $k := \kappa(\boldsymbol{\Sigma})$ . We then bound

$$\begin{aligned}
\sqrt{\bar{\lambda}} \|\bar{\mathbf{A}} - \mathbf{A}\|_{\text{op}} &\leq \sqrt{\bar{\lambda}} \|\bar{\mathbf{A}} - \mathbf{A}\|_F && \text{(Lemma A.1)} \\
&\leq \|(\bar{\mathbf{A}} - \mathbf{A}) \bar{\boldsymbol{\Sigma}}_X^{1/2}\|_F \\
&= \sqrt{R(\bar{\mathbf{A}}, \bar{\boldsymbol{\Sigma}})} && \text{(regret formula)} \\
&\leq 2\sqrt{k} \varepsilon. && \text{(Lemma 3.1)}
\end{aligned}$$

Since  $\varepsilon/\sqrt{\lambda} \leq 1$  and  $\|\mathbf{A}\|_{\text{op}} \leq \sqrt{k}$  by Theorem 2.2, it follows that

$$\|\bar{\mathbf{A}}\|_{\text{op}} \leq \|\mathbf{A}\|_{\text{op}} + \frac{2\sqrt{k}\varepsilon}{\sqrt{\lambda}} \leq 3\sqrt{k}.$$

We now bound the residual relative to the smallest feature eigenvalue. By Lemma 1.1, there are square-root factors  $\mathbf{S}, \bar{\mathbf{S}} \in \mathbb{R}^{d \times d}$  of  $\Sigma, \bar{\Sigma}$  with  $\Delta := \bar{\mathbf{S}} - \mathbf{S}$  satisfying  $\|\Delta\|_F \leq \varepsilon$ . Define the corresponding residual factors and projection

$$\mathbf{E} := \mathbf{S}_Y - \mathbf{A}\mathbf{S}_X, \quad \bar{\mathbf{E}} := \bar{\mathbf{S}}_Y - \bar{\mathbf{A}}\bar{\mathbf{S}}_X, \quad \bar{\Pi}_X := \bar{\mathbf{S}}_X^\dagger \bar{\mathbf{S}}_X.$$

By Lemma B.1,  $\mathbf{E}\mathbf{E}^\top = \Sigma_E$ ,  $\bar{\mathbf{E}}\bar{\mathbf{E}}^\top = \bar{\Sigma}_E$ ,  $\bar{\mathbf{E}} = \bar{\mathbf{S}}_Y(\mathbf{I}_d - \bar{\Pi}_X)$ , and  $\bar{\mathbf{S}}_X(\mathbf{I}_d - \bar{\Pi}_X) = \mathbf{0}$ . Further,

$$\bar{\mathbf{S}}_Y = \mathbf{S}_Y + \Delta_Y = \mathbf{E} + \mathbf{A}\mathbf{S}_X + \Delta_Y = \mathbf{E} + \mathbf{A}\bar{\mathbf{S}}_X - \mathbf{A}\Delta_X + \Delta_Y.$$

We thus compute

$$\begin{aligned} \bar{\mathbf{E}} &= \bar{\mathbf{S}}_Y(\mathbf{I}_d - \bar{\Pi}_X) \\ &= (\mathbf{E} + \mathbf{A}\bar{\mathbf{S}}_X - \mathbf{A}\Delta_X + \Delta_Y)(\mathbf{I}_d - \bar{\Pi}_X) \\ &= (\mathbf{E} + \Delta_Y - \mathbf{A}\Delta_X)(\mathbf{I}_d - \bar{\Pi}_X). \end{aligned}$$

Since the right-hand projection is a contraction, we have

$$\begin{aligned} \|\bar{\mathbf{E}}\|_{\text{op}} &\leq \|\mathbf{E}\|_{\text{op}} + \|\Delta_Y - \mathbf{A}\Delta_X\|_{\text{op}} \\ &\leq \sqrt{\|\Sigma_E\|_{\text{op}}} + \sqrt{1 + \|\mathbf{A}\|_{\text{op}}^2} \varepsilon \\ &\leq \sqrt{k}(\sqrt{\lambda} + \varepsilon). \end{aligned}$$

Here the second inequality uses  $\|(-\mathbf{A} \quad \mathbf{I}_m)\|_{\text{op}} = \sqrt{1 + \|\mathbf{A}\|_{\text{op}}^2}$ , and the last uses Theorem 2.2. Therefore, we bound

$$\frac{\|\bar{\Sigma}_E\|_{\text{op}}}{\bar{\lambda}} = \frac{\|\bar{\mathbf{E}}\|_{\text{op}}^2}{\bar{\lambda}} \leq k \frac{(\sqrt{\lambda} + \varepsilon)^2}{(\sqrt{\lambda} - \varepsilon)^2} \leq 9k,$$

where the last inequality uses  $\varepsilon \leq \sqrt{\lambda}/2$ .

Finally, Theorem 2.2 gives

$$\kappa(\bar{\Sigma}) \leq 1 + \|\bar{\mathbf{A}}\|_{\text{op}}^2 + \frac{\|\bar{\Sigma}_E\|_{\text{op}}}{\bar{\lambda}} \leq 1 + 18k \leq 19k$$

using  $k \geq 1$ . □

**Lemma F.3** (DRO strong duality). *Let  $\varepsilon \geq 0$ . For all  $P \in \mathcal{P}^*$  and  $\hat{P} \in \mathcal{B}_\varepsilon(P)$ , we have*

$$\inf_{\mathbf{A} \in \mathcal{A}} \sup_{Q \in \mathcal{B}_\varepsilon(\hat{P})} L(\mathbf{A}, Q) = \max_{Q \in \mathcal{B}_\varepsilon(\hat{P})} \inf_{\mathbf{A} \in \mathcal{A}} L(\mathbf{A}, Q).$$

*In particular, a maximizer exists for the right-hand side.*

*Proof.* By Lemma 1.2, it suffices to prove

$$\inf_{\mathbf{A} \in \mathcal{A}} \max_{\Sigma \in K} L(\mathbf{A}, \Sigma) = \max_{\Sigma \in K} \inf_{\mathbf{A} \in \mathcal{A}} L(\mathbf{A}, \Sigma), \quad (14)$$

where  $K := \mathcal{B}_\varepsilon(\Sigma_{\hat{P}})$ . The factor representation in Lemma 1.2 shows that  $K$  is the continuous image of a compact Frobenius ball and hence compact. It is convex because Wasserstein balls are convex under mixtures, as follows by mixing transport couplings, and the second-moment map is affine. Since  $L$  is continuous, convex in  $\mathbf{A}$ , and affine in  $\Sigma$ , Sion's minimax theorem (Sion, 1958) gives equality of the infimum-supremum values. Finally,  $\Sigma \mapsto \inf_{\mathbf{A} \in \mathcal{A}} L(\mathbf{A}, \Sigma)$  is upper semicontinuous as an infimum of continuous functions, so its supremum on  $K$  is attained. □

We now prove the remaining excess risk bounds for ERM and DRO.

**Lemma F.4** (Stability implications, ERM). *Let  $P \in \mathcal{P}_2(\mathbb{R}^d)$  and  $\varepsilon > 0$ . We then have*

$$R_\varepsilon(\text{ERM}, P) \leq 4\varepsilon^2 \sup_{Q \in \mathcal{B}_\varepsilon(P)} \kappa_Q.$$

*If additionally  $\varepsilon^2 \leq \lambda_{\min}(\Sigma_X^P)/4$ , then  $R_\varepsilon(\text{ERM}, P) \leq 76\kappa_P\varepsilon^2$ .*

*Proof.* The first bound is trivial if  $\mathcal{B}_\varepsilon(P)$  contains a law outside of  $\mathcal{P}^*$ , since that law has infinite condition number. We may therefore assume  $\mathcal{B}_\varepsilon(P) \subseteq \mathcal{P}^*$ .

For any  $Q \in \mathcal{B}_\varepsilon(P)$ , the ERM solutions at  $P$  and  $Q$  are unique. Applying Lemma 3.1 then gives

$$\begin{aligned} R(\text{ERM}[P], Q) &= R(\mathbf{A}_P, Q) \leq 4\kappa_P\varepsilon^2, \\ R(\text{ERM}[Q], P) &= R(\mathbf{A}_Q, P) \leq 4\kappa_Q\varepsilon^2. \end{aligned}$$

Since  $P \in \mathcal{B}_\varepsilon(P)$ , both bounds are at most  $4\varepsilon^2 \sup_{Q' \in \mathcal{B}_\varepsilon(P)} \kappa_{Q'}$ . Taking the maximum over the two orientations and then the supremum over  $Q \in \mathcal{B}_\varepsilon(P)$  proves the first bound.

If  $\varepsilon^2 \leq \lambda_{\min}(\Sigma_X^P)/4$ , then Lemmas 1.2 and F.2 give  $\kappa_Q \leq 19\kappa_P$  for every  $Q \in \mathcal{B}_\varepsilon(P)$ . Substitution yields the second bound.  $\square$

**Lemma F.5** (Stability implications, DRO). *Let  $P \in \mathcal{P}_2(\mathbb{R}^d)$  and  $\varepsilon > 0$ . We then have*

$$R_\varepsilon(\text{DRO}_\varepsilon, P) \leq 16\varepsilon^2 \sup_{Q \in \mathcal{B}_{2\varepsilon}(P)} \kappa_Q.$$

*If  $\varepsilon^2 \leq \lambda_{\min}(\Sigma_X^P)/16$ , this implies that  $R_\varepsilon(\text{DRO}_\varepsilon, P) \leq 304\kappa_P\varepsilon^2$ .*

*Proof.* The first bound is trivial if  $\mathcal{B}_{2\varepsilon}(P)$  contains a law outside of  $\mathcal{P}^*$ , since that law has infinite condition number. We may therefore assume  $\mathcal{B}_{2\varepsilon}(P) \subseteq \mathcal{P}^*$ .

Fix  $Q \in \mathcal{B}_\varepsilon(P)$  and let  $(P_s, P_t)$  be either  $(P, Q)$  or  $(Q, P)$ . Set  $\hat{\mathbf{A}} := \text{DRO}_\varepsilon[P_s]$ . We first identify  $\hat{\mathbf{A}}$  as the ERM solution at a least-favorable law, then transfer its zero regret there to the target.

By Lemma F.3, there exists a maximizer  $\hat{Q}$  of

$$\max_{Q' \in \mathcal{B}_\varepsilon(P_s)} \inf_{\mathbf{A} \in \mathcal{A}} L(\mathbf{A}, Q').$$

DRO optimality and the minimax identity (Lemma F.3) give

$$\begin{aligned} L(\hat{\mathbf{A}}, \hat{Q}) &\leq \sup_{Q' \in \mathcal{B}_\varepsilon(P_s)} L(\hat{\mathbf{A}}, Q') \\ &= \inf_{\mathbf{A} \in \mathcal{A}} \sup_{Q' \in \mathcal{B}_\varepsilon(P_s)} L(\mathbf{A}, Q') \\ &= \max_{Q' \in \mathcal{B}_\varepsilon(P_s)} \inf_{\mathbf{A} \in \mathcal{A}} L(\mathbf{A}, Q') \\ &= \inf_{\mathbf{A} \in \mathcal{A}} L(\mathbf{A}, \hat{Q}). \end{aligned}$$

The reverse inequality is automatic, so  $\hat{\mathbf{A}}$  minimizes  $L(\cdot, \hat{Q})$ . Moreover,  $\hat{Q} \in \mathcal{B}_\varepsilon(P_s) \subseteq \mathcal{B}_{2\varepsilon}(P) \subseteq \mathcal{P}^*$ . Its risk minimizer is therefore unique, implying  $\hat{\mathbf{A}} = \mathbf{A}_{\hat{Q}}$ .

Since both  $\hat{Q}$  and  $P_t$  belong to  $\mathcal{B}_\varepsilon(P_s)$ , we have  $W_2(\hat{Q}, P_t) \leq 2\varepsilon$ . Applying Lemma 3.1 with reference law  $\hat{Q}$  therefore gives

$$R(\hat{\mathbf{A}}, P_t) = R(\mathbf{A}_{\hat{Q}}, P_t) \leq 4\kappa_{\hat{Q}} W_2(\hat{Q}, P_t)^2 \leq 16\varepsilon^2 \sup_{Q' \in \mathcal{B}_{2\varepsilon}(P)} \kappa_{Q'}.$$

Supremizing over source–target pairs  $Q \in \mathcal{B}_\varepsilon(P)$  yields the claimed bound for  $R_\varepsilon(\text{DRO}_\varepsilon, P)$ .

Finally, if  $\varepsilon^2 \leq \lambda_{\min}(\Sigma_X^P)/16$ , then  $2\varepsilon \leq \sqrt{\lambda_{\min}(\Sigma_X^P)}/2$ . By Lemmas 1.2 and F.2,  $\kappa_{Q'} \leq 19\kappa_P$  for all  $Q' \in \mathcal{B}_{2\varepsilon}(P)$ . Substitution then gives  $R_\varepsilon(\text{DRO}_\varepsilon, P) \leq 304\kappa_P\varepsilon^2$ .  $\square$

### F.3 Rank-Adaptive DRO Bound (Proof of Lemma 3.3)

To start, we establish an additional robust regret bound for DRO by more directly exploiting its min-max problem formulation. Similar bounds exist in the literature, typically stated for more generic loss functions (Gao et al., 2024; Lee and Raginsky, 2018).

**Lemma F.6** (DRO regularizer bound). *For  $P \in \mathcal{P}^*$ , define*

$$\Gamma_P := \sqrt{\text{tr}((\mathbf{I}_m + \mathbf{A}_P \mathbf{A}_P^\top) \Sigma_E^P)}.$$

*Then, for every  $\varepsilon \geq 0$ ,*

$$R_\varepsilon(\text{DRO}_\varepsilon, P) \leq 4\Gamma_P \varepsilon + 5\kappa_P \varepsilon^2. \quad (15)$$

*Proof.* Fix  $P \in \mathcal{P}^*$ , and abbreviate

$$\mathbf{A} := \mathbf{A}_P, \quad \mathbf{S} := \Sigma_P^{1/2}, \quad \mathbf{M} := (\mathbf{A} \quad -\mathbf{I}_m), \quad \mathbf{E} := \mathbf{S}_Y - \mathbf{A} \mathbf{S}_X.$$

By Lemma B.1,  $\mathbf{E} \mathbf{E}^\top = \Sigma_E^P$ . Hence

$$\|\mathbf{M}^\top \mathbf{E}\|_F^2 = \text{tr}((\mathbf{I}_m + \mathbf{A} \mathbf{A}^\top) \Sigma_E^P) = \Gamma_P^2.$$

Also,  $\|\mathbf{M}\|_{\text{op}}^2 = 1 + \|\mathbf{A}\|_{\text{op}}^2 \leq \kappa_P$  by Theorem 2.2.

We first bound changes in the risk of the fixed predictor  $\mathbf{A}$  in the neighborhood of  $P$ . For any  $\delta \geq 0$  and  $Q \in \mathcal{B}_\delta(P)$ , Lemma 1.2 provides  $\Delta \in \mathbb{R}^{d \times d}$  such that

$$\Sigma_Q = (\mathbf{S} + \Delta)(\mathbf{S} + \Delta)^\top, \quad \|\Delta\|_F \leq \delta.$$

Since the residual factor at  $Q$  for  $\mathbf{A}$  is  $\mathbf{E} - \mathbf{M} \Delta$ , we have by Lemma B.1 that

$$\begin{aligned} L(\mathbf{A}, Q) - L(\mathbf{A}, P) &= \|\mathbf{E} - \mathbf{M} \Delta\|_F^2 - \|\mathbf{E}\|_F^2 \\ &= -2\langle \mathbf{M}^\top \mathbf{E}, \Delta \rangle_F + \|\mathbf{M} \Delta\|_F^2. \end{aligned}$$

Cauchy-Schwarz and  $\|\mathbf{M} \Delta\|_F^2 \leq \|\mathbf{M}\|_{\text{op}}^2 \|\Delta\|_F^2 \leq \kappa_P \delta^2$  then give the two-sided bound

$$-2\Gamma_P \delta \leq L(\mathbf{A}, Q) - L(\mathbf{A}, P) \leq 2\Gamma_P \delta + \kappa_P \delta^2. \quad (16)$$

For the lower bound, we have simply discarded the nonnegative quadratic term.

**Target role.** We now bound regret when  $P$  is the target law. Fix  $Q \in \mathcal{B}_\varepsilon(P)$  and set  $\hat{\mathbf{A}} := \text{DRO}_\varepsilon[Q]$ . Since  $P \in \mathcal{B}_\varepsilon(Q)$ , DRO optimality and a comparison with  $\mathbf{A}$  imply that

$$\begin{aligned} L(\hat{\mathbf{A}}, P) &\leq \sup_{Q' \in \mathcal{B}_\varepsilon(Q)} L(\hat{\mathbf{A}}, Q') \\ &\leq \sup_{Q' \in \mathcal{B}_\varepsilon(Q)} L(\mathbf{A}, Q') \\ &\leq \sup_{Q' \in \mathcal{B}_{2\varepsilon}(P)} L(\mathbf{A}, Q') \\ &\leq L(\mathbf{A}, P) + 4\Gamma_P \varepsilon + 4\kappa_P \varepsilon^2. \end{aligned}$$

The third inequality uses  $\mathcal{B}_\varepsilon(Q) \subseteq \mathcal{B}_{2\varepsilon}(P)$ , and the last applies (16) with  $\delta = 2\varepsilon$ . Subtracting  $L(\mathbf{A}, P) = \inf_{\mathbf{B} \in \mathcal{A}} L(\mathbf{B}, P)$  gives

$$R(\text{DRO}_\varepsilon[Q], P) \leq 4\Gamma_P \varepsilon + 4\kappa_P \varepsilon^2.$$

We note that, using the terminology of Gao et al. (2024), the RHS is an upper bound on the Wasserstein regularizer of the loss  $(x, y) \mapsto \|\mathbf{A}_P x - y\|^2$  at  $P$ , evaluated at radius  $2\varepsilon$ .

**Source role.** Now set  $\hat{\mathbf{A}} := \text{DRO}_\varepsilon[P]$ . Comparing with  $\mathbf{A}$ , we bound regret at  $Q \in \mathcal{B}_\varepsilon(P)$  by

$$\begin{aligned} R(\hat{\mathbf{A}}, Q) &= L(\hat{\mathbf{A}}, Q) - L(\mathbf{A}, Q) + R(\mathbf{A}, Q) \\ &\leq \sup_{Q' \in \mathcal{B}_\varepsilon(P)} L(\mathbf{A}, Q') - L(\mathbf{A}, Q) + R(\mathbf{A}, Q) \\ &\leq 4\Gamma_P \varepsilon + \kappa_P \varepsilon^2 + R(\mathbf{A}, Q) \\ &\leq 4\Gamma_P \varepsilon + 5\kappa_P \varepsilon^2. \end{aligned}$$

The first inequality uses feasibility of  $Q$  and DRO optimality. The second applies the upper bound in (16) to  $Q'$  and its lower bound to  $Q$ , both with  $\delta = \varepsilon$ . The last uses  $R(\mathbf{A}_P, Q) \leq 4\kappa_P W_2(P, Q)^2$  from Lemma 3.1.

Taking the maximum of these two role-specific bounds and then the supremum over  $Q \in \mathcal{B}_\varepsilon(P)$  gives (15).  $\square$

Next, we combine Lemma F.6 with Lemma F.5 to obtain the claimed result.

*Proof of Lemma 3.3.* Fixing  $P \in \mathcal{P}^*$ , we abbreviate

$$\lambda := \lambda_{\min}(\Sigma_X^P), \quad \mu := \lambda_{\max}(\Sigma_E^P), \quad r := r_{\text{eff}}(\Sigma_E^P), \quad s := \kappa_P - 1.$$

The condition number formula (Theorem 2.2) gives

$$\|\mathbf{A}_P\|_{\text{op}}^2 \leq s, \quad \mu \leq s\lambda.$$

Also  $\text{tr}(\Sigma_E^P) = r\mu$ , including when  $\Sigma_E^P = \mathbf{0}$  under the convention  $r = 0$ .

If  $\varepsilon^2 \leq \lambda/16$ , the small-radius bound in Lemma F.5 gives

$$R_\varepsilon(\text{DRO}_\varepsilon, P) \lesssim \kappa_P \varepsilon^2,$$

which suffices. We therefore assume  $\varepsilon^2 > \lambda/16$  below.

To bound  $\Gamma_P$ , we estimate its coefficient-dependent trace term in two ways:

$$\begin{aligned} \text{tr}(\mathbf{A}_P \mathbf{A}_P^\top \Sigma_E^P) &\leq \|\mathbf{A}_P\|_{\text{op}}^2 \text{tr}(\Sigma_E^P) = r\mu \|\mathbf{A}_P\|_{\text{op}}^2, \\ \text{tr}(\mathbf{A}_P \mathbf{A}_P^\top \Sigma_E^P) &\leq \mu \|\mathbf{A}_P\|_{\text{F}}^2 \leq p\mu \|\mathbf{A}_P\|_{\text{op}}^2, \end{aligned}$$

with the inequalities using Lemma A.1. Taking the smaller bound and using  $\|\mathbf{A}_P\|_{\text{op}}^2 \leq s$  yields

$$\begin{aligned} \Gamma_P^2 &= \text{tr}(\Sigma_E^P) + \text{tr}(\mathbf{A}_P \mathbf{A}_P^\top \Sigma_E^P) \\ &\leq \mu(r + s \min\{p, r\}) \\ &\leq \lambda(rs + s^2 \min\{p, r\}). \end{aligned}$$

Using  $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$  for  $a, b \geq 0$ , we further bound

$$\Gamma_P \leq \sqrt{\lambda} \left( \sqrt{rs} + s \sqrt{\min\{p, r\}} \right).$$

Since  $\sqrt{\lambda} < 4\varepsilon$ , Lemma F.6 now gives

$$\begin{aligned} R_\varepsilon(\text{DRO}_\varepsilon, P) &\leq 4\Gamma_P \varepsilon + 5\kappa_P \varepsilon^2 \\ &\leq \left( 16s \sqrt{\min\{p, r\}} + 16\sqrt{rs} + 5\kappa_P \right) \varepsilon^2 \\ &\lesssim \left( \sqrt{\min\{p, r\}} (\kappa_P - 1) + \sqrt{r(\kappa_P - 1)} + \kappa_P \right) \varepsilon^2, \end{aligned}$$

as desired.  $\square$

#### F.4 Main Population-Limit Regret Bounds (Proof of Theorem 3.4)

We start with upper bounds, which follow from our existing results. Indeed, Lemma 3.2 implies that  $R_\varepsilon(\text{SNR}_\varepsilon, \mathcal{G}), R_\varepsilon(\text{MRO}_\varepsilon, \mathcal{G}) \lesssim \kappa \varepsilon^2$ . Moreover, Lemma 3.3 and the basic effective rank bound  $r_{\text{eff}}(\Sigma_E^P) \leq m$  imply that

$$R_\varepsilon(\text{DRO}_\varepsilon, \mathcal{G}) \lesssim \left( \sqrt{\min\{p, m\}} (\kappa - 1) + \sqrt{m(\kappa - 1)} + \kappa \right) \varepsilon^2,$$

as desired. Of course, no upper bound was claimed for ERM.

For the lower bounds, we scale up the local minimax result to obtain the minimax lower bound for  $\mathcal{G}$ . Then, we give an explicit construction establishing unbounded regret for ERM. Two further constructions, whose analysis is more subtle, yield the dimension-dependent terms in the DRO rate.

We start with a helper lemma that demonstrates a certain invariance to rescaling.

**Lemma F.7** (Rescaling invariance). *Let  $P, Q \in \mathcal{P}_2(\mathbb{R}^d)$  and  $\mathbf{A} \in \mathcal{A}$ . For  $\lambda > 0$ , let  $P_\lambda$  and  $Q_\lambda$  denote the pushforwards of  $P$  and  $Q$ , respectively, under  $z \mapsto \lambda z$ . We then have*

$$\begin{aligned} W_2(P_\lambda, Q_\lambda) &= \lambda W_2(P, Q), \\ L(\mathbf{A}, P_\lambda) &= \lambda^2 L(\mathbf{A}, P), \\ R(\mathbf{A}, P_\lambda) &= \lambda^2 R(\mathbf{A}, P), \\ \kappa_{P_\lambda} &= \kappa_P. \end{aligned}$$

*Moreover, for  $\varepsilon \geq 0$ ,  $R_{\lambda\varepsilon}^*(P_\lambda) = \lambda^2 R_\varepsilon^*(P)$ , and the optimization problems defining  $\text{DRO}_{\lambda\varepsilon}[P_\lambda]$  and  $\text{DRO}_\varepsilon[P]$  have the same minimizer set.*

*Proof.* Scaling both coordinates of a transport coupling multiplies its expected squared cost by  $\lambda^2$ . Applying the same argument to the inverse scaling proves the Wasserstein identity. The risk and regret identities follow from  $\Sigma_{P_\lambda} = \lambda^2 \Sigma_P$ . The condition number formula is invariant under this scaling, as is membership in  $\mathcal{P}^*$ .

The Wasserstein identity implies that  $\mathcal{B}_{\lambda\varepsilon}(P_\lambda) = \{Q_\lambda : Q \in \mathcal{B}_\varepsilon(P)\}$ . Combining this with regret homogeneity and (1) proves the local minimax identity. Likewise,

$$\sup_{Q' \in \mathcal{B}_{\lambda\varepsilon}(P_\lambda)} L(\mathbf{A}, Q') = \lambda^2 \sup_{Q \in \mathcal{B}_\varepsilon(P)} L(\mathbf{A}, Q),$$

so the DRO minimizer sets coincide.  $\square$

We now turn to the minimax lower bound for  $\mathcal{G}(\kappa)$ .

**Lemma F.8** (Minimax lower bound). *For every  $\varepsilon > 0$  and  $\kappa \geq 1$ , we have  $R_\varepsilon^*(\mathcal{G}(\kappa)) \geq \kappa\varepsilon^2$ .*

*Proof.* Let  $P \in \mathcal{P}_2(\mathbb{R}^d)$  be the law of  $(X, Y)$  where

$$X \sim \mathcal{N}(0, \mathbf{I}_p), \quad Y = \sqrt{\kappa - 1} X_1 e_1, \quad e_1 \in \mathbb{R}^m.$$

By construction,  $\mathbf{A}_P$  has only one potentially nonzero entry,  $\sqrt{\kappa - 1}$  in position  $(1, 1)$ , and the optimal residual is zero. Thus, Theorem 2.2 gives  $\kappa_P = \kappa$ .

For  $t > 0$ , let  $P_t$  be the pushforward of  $P$  under  $z \mapsto tz$ . By Lemma F.7,  $P_t \in \mathcal{G}(\kappa)$  and

$$R_\varepsilon^*(\mathcal{G}(\kappa)) \geq R_\varepsilon^*(P_t) = t^2 R_{\varepsilon/t}^*(P).$$

Further, the local expansion in Theorem 2.2 gives

$$\lim_{t \rightarrow \infty} t^2 R_{\varepsilon/t}^*(P) = \kappa\varepsilon^2.$$

Since the preceding lower bound holds for every  $t > 0$ , the claim follows.  $\square$

Next, we prove the lower bound for ERM.

**Lemma F.9** (ERM lower bound). *For every  $\varepsilon > 0$  and  $\kappa \geq 1$ ,  $R_\varepsilon(\text{ERM}, \mathcal{G}(\kappa)) = \infty$ , independently of the choice of ERM minimizer.*

*Proof.* Consider target  $P \in \mathcal{P}_2(\mathbb{R}^d)$  defined as the law of  $(\varepsilon W/2, 0_m)$  where  $W \sim \mathcal{N}(0, \mathbf{I}_p)$ . Then  $\Sigma_X^P = (\varepsilon^2/4)\mathbf{I}_p$ ,  $\mathbf{A}_P = \mathbf{0}$ , and  $\Sigma_E^P = \mathbf{0}$ . Hence  $\kappa_P = 1$  and  $P \in \mathcal{G}(\kappa)$ .

For  $\delta \in (0, 1)$ , consider source  $Q_\delta \in \mathcal{P}_2(\mathbb{R}^d)$  taken as the law of

$$\tilde{X} = \frac{\varepsilon}{2}(\delta W_1, W_2, \dots, W_p), \quad \tilde{Y} = \frac{\varepsilon}{2}W_1 e_1, \quad e_1 \in \mathbb{R}^m.$$

The coupling using the same  $W$  gives

$$W_2(P, Q_\delta)^2 \leq \frac{\varepsilon^2}{4}((1 - \delta)^2 + 1) \leq \varepsilon^2.$$

The source feature second-moment matrix is positive definite, and its unique ERM solution has entry  $1/\delta$  in position  $(1, 1)$  and zeros elsewhere. Therefore,

$$R_\varepsilon(\text{ERM}, \mathcal{G}(\kappa)) \geq R(\text{ERM}[Q_\delta], P) = \frac{\varepsilon^2}{4\delta^2}.$$

Letting  $\delta \downarrow 0$  proves the claim.  $\square$

For the DRO result, we combine two lower bounds in addition to the minimax lower bound. Although DRO solutions are generally difficult to characterize, we are able to streamline analysis in both cases by exploiting rotational symmetry. We fix  $\varepsilon = 1$  for now to simplify presentation.

**Lemma F.10** (DRO lower bound, A). *For every  $\kappa \geq 1$ ,*

$$R_1(\text{DRO}_1, \mathcal{G}(\kappa)) \gtrsim \sqrt{\min\{p, m\}} (\kappa - 1),$$

*independently of the choice of DRO minimizer.*

*Proof.* The claim is immediate for  $\kappa = 1$ . Write  $q := \min\{p, m\}$ . If  $q < 65$ , then Lemma F.8 gives

$$R_1(\text{DRO}_1, \mathcal{G}(\kappa)) \geq R_1^*(\mathcal{G}(\kappa)) \gtrsim \kappa \gtrsim \sqrt{q} (\kappa - 1).$$

We may therefore assume  $\kappa > 1$  and  $q \geq 65$ .

**Construction.** Set  $b := \sqrt{(\kappa - 1)/2} > 0$ . Let  $X \sim \mathcal{N}(0, \mathbf{I}_p)$  and

$$E \sim \mathcal{N}(0, \text{diag}(0, b^2 \mathbf{I}_{q-1}, \mathbf{0}_{m-q}))$$

be independent. Take target  $P$  as the law of  $(X, Y)$  where

$$Y = \begin{pmatrix} bX_{1:q} \\ 0_{m-q} \end{pmatrix} + E$$

Here and below, empty blocks of size zero should be omitted. By this construction,

$$\mathbf{A}_P = \begin{pmatrix} b\mathbf{I}_q & \mathbf{0}_{q \times (p-q)} \\ \mathbf{0}_{(m-q) \times q} & \mathbf{0}_{(m-q) \times (p-q)} \end{pmatrix}, \quad \Sigma_X^P = \mathbf{I}_p,$$

and  $\Sigma_E^P = \text{diag}(0, b^2 \mathbf{I}_{q-1}, \mathbf{0}_{m-q})$ . Thus,

$$\mathbf{A}_P \mathbf{A}_P^\top + \lambda_{\min}(\Sigma_X^P)^{-1} \Sigma_E^P = \text{diag}(b^2, 2b^2 \mathbf{I}_{q-1}, \mathbf{0}_{m-q}),$$

and Theorem 2.2 gives  $\kappa_P = 1 + 2b^2 = \kappa$ .

For the perturbed source, take  $\tilde{P}$  as the law of  $(\tilde{X}, Y)$ , where  $\tilde{X}_1 = cX_1$  and  $\tilde{X}_{2:p} = X_{2:p}$ , for  $c \in (0, 1]$  to be tuned. This construction ensures  $W_2(\tilde{P}, P)^2 \leq \mathbb{E}[\|X - \tilde{X}\|^2] \leq 1$ .

**Average over symmetries.** Consider the compact subgroup

$$U := \{\text{diag}(1, \mathbf{H}, s\mathbf{I}_{p-q}, 1, \mathbf{H}, t\mathbf{I}_{m-q}) : \mathbf{H} \in \mathbb{O}^{q-1}, s, t \in \{-1, 1\}\} \subseteq \mathbb{O}^d.$$

Both  $P$  and  $\tilde{P}$  are invariant under pushforwards by each  $\mathbf{U} \in U$ . Moreover, the invariant coefficient matrices (in the sense of Lemma C.4) are exactly

$$\mathbf{A}(\alpha, \beta) := \begin{pmatrix} \text{diag}(\alpha, \beta \mathbf{I}_{q-1}) & \mathbf{0}_{q \times (p-q)} \\ \mathbf{0}_{(m-q) \times q} & \mathbf{0}_{(m-q) \times (p-q)} \end{pmatrix}, \quad \alpha, \beta \in \mathbb{R}.$$

Indeed, varying  $s$  and  $t$  eliminates the extra feature columns and response rows. Within the leading  $q \times q$  block, taking  $\mathbf{H} = -\mathbf{I}_{q-1}$  eliminates the off-diagonal blocks, and commuting with every  $\mathbf{H} \in \mathbb{O}^{q-1}$  forces the remaining  $(q-1) \times (q-1)$  block to be a scalar multiple of the identity. By Lemma C.4, it therefore suffices to lower bound the target regret of invariant DRO optimizers.

**Construct candidate DRO optimizer.** Define the invariant objective by

$$J(\alpha, \beta) := \sup_{Q \in \mathcal{B}_1(\tilde{P})} L(\mathbf{A}(\alpha, \beta), Q).$$

We have

$$\mathbf{A}(\alpha, \beta) \mathbf{A}(\alpha, \beta)^\top = \text{diag}(\alpha^2, \beta^2 \mathbf{I}_{q-1}, \mathbf{0}_{m-q}),$$

and the coordinate residual second moments at the source are

$$\mathbb{E}[(\mathbf{A}(\alpha, \beta) \tilde{X} - Y)_i^2] = \begin{cases} (\alpha c - b)^2, & i = 1, \\ (\beta - b)^2 + b^2, & 2 \leq i \leq q, \\ 0, & q < i \leq m. \end{cases}$$

Thus Corollary C.3 with  $\varepsilon = 1$  gives

$$J(\alpha, \beta) = \inf_{\eta > 1 + \max\{\alpha^2, \beta^2\}} \eta \left( 1 + \frac{(\alpha c - b)^2}{\eta - 1 - \alpha^2} + \frac{(q-1)[(\beta - b)^2 + b^2]}{\eta - 1 - \beta^2} \right). \quad (17)$$

Now consider the related problem associated with coordinates  $2, \dots, q$ :

$$V(\beta) := \inf_{\eta > 1 + \beta^2} \eta \left( 1 + \frac{(q-1)[(\beta - b)^2 + b^2]}{\eta - 1 - \beta^2} \right).$$

By Lemma A.7, with  $\delta = \eta - (1 + \beta^2)$ , this infimum is attained at

$$\eta(\beta) := 1 + \beta^2 + \sqrt{(1 + \beta^2)(q-1)[(\beta - b)^2 + b^2]},$$

with value

$$V(\beta) = \left( \sqrt{1 + \beta^2} + \sqrt{(q-1)[(\beta - b)^2 + b^2]} \right)^2.$$

This function is strictly convex and coercive. Let  $\beta_\star$  denote its unique minimizer, set  $\eta_\star := \eta(\beta_\star)$ , and let  $\alpha_\star := \sqrt{\eta_\star - 1} > 0$ . We set the free scaling parameter  $c$  to  $b/\alpha_\star > 0$ . Below we shall verify that  $c \leq 1$ , ensuring that  $\tilde{P} \in \mathcal{B}_1(P)$ .

**Certify optimality.** We next show that  $J$  is uniquely minimized at  $(\alpha_\star, \beta_\star)$ . For every feasible  $\eta$  for (17),  $\eta/(\eta - 1 - \alpha^2) \geq 1$ . The remaining terms are bounded from below by  $V(\beta)$ , so

$$J(\alpha, \beta) \geq V(\beta) + (\alpha c - b)^2 \geq V(\beta_\star).$$

At  $(\alpha_\star, \beta_\star)$ , the first-coordinate residual vanishes because  $\alpha_\star c = b$ . Also  $\alpha_\star^2 > \beta_\star^2$ , so

$$J(\alpha_\star, \beta_\star) = \inf_{\eta > \eta_\star} \eta \left( 1 + \frac{(q-1)[(\beta_\star - b)^2 + b^2]}{\eta - 1 - \beta_\star^2} \right).$$

When minimizing instead over the larger range  $\eta > 1 + \beta_\star^2$ , our construction ensures that the minimum is attained at  $\eta = \eta_\star$ . Although  $\eta_\star$  is excluded from the current feasible range, the objective is continuous there, so the infimum is approached as  $\eta \downarrow \eta_\star$  and equals  $V(\beta_\star)$ . The preceding lower bound then forces every minimizing pair  $(\alpha, \beta)$  to satisfy  $\beta = \beta_\star$  and  $\alpha c = b$ . Since  $c = b/\alpha_\star > 0$ , this implies that  $(\alpha_\star, \beta_\star)$  is the unique minimizing pair.

**Bound  $\alpha_\star$  and  $\beta_\star$  from below, verify  $c$  is admissible.** Define

$$h(\beta) := \sqrt{V(\beta)} = \sqrt{1 + \beta^2} + \sqrt{q-1} \sqrt{(\beta - b)^2 + b^2},$$

so that  $\beta_\star \in \arg \min_{\beta} h(\beta)$ . This function is strictly convex and coercive. For  $\beta \leq b/2$ ,

$$h'(\beta) = \frac{\beta}{\sqrt{1 + \beta^2}} + \sqrt{q-1} \frac{\beta - b}{\sqrt{(\beta - b)^2 + b^2}} \leq 1 - \sqrt{\frac{q-1}{5}} < 0.$$

The inequality uses  $\beta - b \leq -b/2$ . Hence  $\beta_\star > b/2$ . It follows that

$$\begin{aligned} \alpha_\star^2 &= \eta_\star - 1 \\ &\geq \sqrt{(1 + \beta_\star^2)(q-1)[(\beta_\star - b)^2 + b^2]} \\ &\geq \frac{b^2}{2} \sqrt{q-1}. \end{aligned} \quad (18)$$

Since  $q \geq 65$ , this implies  $\alpha_\star \geq 2b$  and hence  $c = b/\alpha_\star \leq 1/2$ .

**Lower bound regret.** Since  $c \leq 1$ ,  $\tilde{P} \in \mathcal{B}_1(P)$ , and its unique invariant DRO optimizer satisfies

$$\begin{aligned} R(\mathbf{A}(\alpha_\star, \beta_\star), P) &= (\alpha_\star - b)^2 + (q-1)(\beta_\star - b)^2 \\ &\geq (\alpha_\star - b)^2 \\ &\geq \frac{1}{4} \alpha_\star^2 \\ &\geq \frac{b^2}{8} \sqrt{q-1} \gtrsim \sqrt{q} (\kappa - 1). \end{aligned}$$

The first equality uses  $\Sigma_X^P = \mathbf{I}_p$  and the regret formula, the second inequality uses that  $\alpha_* \geq 2b$ , the third inequality uses (18), and the last inequality uses  $b^2 = (\kappa - 1)/2$ .

By Lemma C.4, every DRO optimizer at  $\tilde{P}$  has target regret at least that of  $\mathbf{A}(\alpha_*, \beta_*)$ . Since  $P \in \mathcal{G}(\kappa)$  and  $\tilde{P} \in \mathcal{B}_1(P)$ ,

$$R_1(\text{DRO}_1, \mathcal{G}(\kappa)) \geq R(\text{DRO}_1[\tilde{P}], P) \geq R(\mathbf{A}(\alpha_*, \beta_*), P) \gtrsim \sqrt{q}(\kappa - 1).$$

This proves the claim for every choice of DRO minimizer.  $\square$

**Lemma F.11** (DRO lower bound, B). *For every  $\kappa \geq 1$ ,*

$$R_1(\text{DRO}_1, \mathcal{G}(\kappa)) \gtrsim \sqrt{m(\kappa - 1)},$$

*independently of the choice of DRO minimizer.*

*Proof.* The claim is immediate for  $\kappa = 1$ . If  $m = 1$ , then Lemma F.8 gives

$$R_1(\text{DRO}_1, \mathcal{G}(\kappa)) \geq R_1^*(\mathcal{G}(\kappa)) \gtrsim \kappa \geq \sqrt{\kappa - 1}.$$

We may therefore assume  $\kappa > 1$  and  $m \geq 2$ .

**Construction.** Set

$$s := \kappa - 1 > 0, \quad N := (m - 1)s > 0, \quad a := \frac{1}{2} \min\{\sqrt{s}, N^{1/4}\} > 0.$$

Let  $X \sim \mathcal{N}(0, \mathbf{I}_p)$  and  $E \sim \mathcal{N}(0, \text{diag}(0, s\mathbf{I}_{m-1}))$  be independent. Fix target  $P \in \mathcal{P}_2(\mathbb{R}^d)$  as the law of  $(X, Y)$ , where

$$Y = aX_1e_1 + E, \quad e_1 \in \mathbb{R}^m.$$

By this construction,

$$\mathbf{A}_P = (ae_1 \quad \mathbf{0}_{m \times (p-1)}), \quad \Sigma_X^P = \mathbf{I}_p, \quad \Sigma_E^P = \text{diag}(0, s\mathbf{I}_{m-1}).$$

Thus

$$\mathbf{A}_P \mathbf{A}_P^\top + \lambda_{\min}(\Sigma_X^P)^{-1} \Sigma_E^P = \text{diag}(a^2, s\mathbf{I}_{m-1}).$$

Since  $a^2 \leq s$  and  $m \geq 2$ , Theorem 2.2 gives  $\kappa_P = 1 + s = \kappa$ .

For the perturbed source, take  $\tilde{P}$  as the law of  $(\tilde{X}, Y)$ , where  $\tilde{X}_1 = cX_1$  and  $\tilde{X}_{2:p} = X_{2:p}$ , for  $c \in (0, 1]$  to be specified. Coupling the same  $X, E$  gives

$$W_2(\tilde{P}, P)^2 \leq \mathbb{E} \|X - \tilde{X}\|^2 = (1 - c)^2 \leq 1.$$

**Average over symmetries.** Consider the compact subgroup

$$U := \{\text{diag}(1, t\mathbf{I}_{p-1}, 1, \mathbf{H}) : t \in \{-1, 1\}, \mathbf{H} \in \mathbb{O}^{m-1}\} \subseteq \mathbb{O}^d.$$

Both  $P$  and  $\tilde{P}$  are invariant under pushforwards by each  $\mathbf{U} \in U$ . The invariant coefficient matrices (in the sense of Lemma C.4) are exactly

$$\mathbf{A}(\alpha) := (\alpha e_1 \quad \mathbf{0}_{m \times (p-1)}), \quad \alpha \in \mathbb{R}.$$

Indeed, varying  $t$  eliminates columns  $2, \dots, p$ , and taking  $\mathbf{H} = -\mathbf{I}_{m-1}$  eliminates rows  $2, \dots, m$ . By Lemma C.4, it therefore suffices to lower bound the target regret of invariant DRO optimizers.

**Construct candidate DRO optimizer.** Define the invariant objective by

$$J(\alpha) := \sup_{Q \in \mathcal{B}_1(\tilde{P})} L(\mathbf{A}(\alpha), Q).$$

We have

$$\mathbf{A}(\alpha) \mathbf{A}(\alpha)^\top = \text{diag}(\alpha^2, \mathbf{0}_{m-1}),$$

and the coordinate residual second moments at the source are

$$\mathbb{E}[(\mathbf{A}(\alpha)\tilde{X} - Y)_i^2] = \begin{cases} (\alpha c - a)^2, & i = 1, \\ s, & 2 \leq i \leq m. \end{cases}$$

Thus Corollary C.3 with  $\varepsilon = 1$  gives

$$J(\alpha) = \inf_{\eta > 1 + \alpha^2} \eta \left( 1 + \frac{(\alpha c - a)^2}{\eta - 1 - \alpha^2} + \frac{N}{\eta - 1} \right). \quad (19)$$

Now consider the related problem associated with coordinates  $2, \dots, m$ :

$$V := \inf_{\eta > 1} \eta \left( 1 + \frac{N}{\eta - 1} \right).$$

By Lemma A.7, with  $\delta = \eta - 1$ , this infimum is attained at  $\eta_* := 1 + \sqrt{N}$  and  $V = (1 + \sqrt{N})^2$ . We now set  $\alpha_* := \sqrt{\eta_* - 1} = N^{1/4}$  and  $c := a/\alpha_*$ . Our choice of  $a$  gives  $0 < a \leq \alpha_*/2$  and hence  $0 < c \leq 1/2$ . Thus the source construction is admissible.

**Certify optimality.** We next show that  $J$  is uniquely minimized at  $\alpha_*$ . For every feasible  $\eta$  in (19),  $\eta/(\eta - 1 - \alpha^2) \geq 1$ . The remaining terms are bounded from below by  $V$ , so

$$J(\alpha) \geq V + (\alpha c - a)^2 \geq V.$$

At  $\alpha = \alpha_*$ , the first-coordinate residual vanishes because  $\alpha_* c = a$ . Since  $1 + \alpha_*^2 = \eta_*$ , we obtain

$$J(\alpha_*) = \inf_{\eta > \eta_*} \eta \left( 1 + \frac{N}{\eta - 1} \right).$$

When minimizing instead over the larger range  $\eta > 1$ , the bulk minimum is attained at  $\eta = \eta_*$ . Although  $\eta_*$  is excluded from the restricted range, the objective is continuous there, so the infimum is approached as  $\eta \downarrow \eta_*$  and equals  $V$ . Thus  $J(\alpha_*) = V$ .

If  $\alpha$  is any minimizer of  $J$ , the preceding lower bound gives  $0 = J(\alpha) - V \geq (\alpha c - a)^2$ . Hence,  $\alpha c = a$  and, since  $c > 0$ ,  $\alpha = a/c = \alpha_*$ . Thus,  $\mathbf{A}(\alpha_*)$  is the unique invariant DRO optimizer.

**Lower bound regret.** Since  $\Sigma_X^P = \mathbf{I}_p$ , the regret formula and our choice of  $c = a/\alpha_*$  give

$$\begin{aligned} R(\mathbf{A}(\alpha_*), P) &= \|\mathbf{A}(\alpha_*) - \mathbf{A}_P\|_F^2 \\ &= (\alpha_* - a)^2 \\ &= \alpha_*^2(1 - c)^2 \\ &\geq \frac{1}{4}\alpha_*^2 = \frac{1}{4}\sqrt{(m-1)(\kappa-1)} \gtrsim \sqrt{m(\kappa-1)}. \end{aligned}$$

The inequality uses  $c \leq 1/2$ , and the final comparison uses  $m-1 \geq m/2$  for  $m \geq 2$ .

By Lemma C.4, every DRO optimizer at  $\tilde{P}$  has target regret at least that of  $\mathbf{A}(\alpha_*)$ . Since  $P \in \mathcal{G}(\kappa)$  and  $\tilde{P} \in \mathcal{B}_1(P)$ ,

$$R_1(\text{DRO}_1, \mathcal{G}(\kappa)) \geq R(\text{DRO}_1[\tilde{P}], P) \geq R(\mathbf{A}(\alpha_*), P) \gtrsim \sqrt{m(\kappa-1)}.$$

This proves the claim for every choice of DRO minimizer.  $\square$

Combining the results above, we obtain the following unified lower bound.

**Lemma F.12** (DRO lower bound). *For all  $\varepsilon > 0$  and  $\kappa \geq 1$ , we have*

$$R_\varepsilon(\text{DRO}_\varepsilon, \mathcal{G}(\kappa)) \gtrsim (\sqrt{\min\{p, m\}(\kappa-1)} + \sqrt{m(\kappa-1)} + \kappa)\varepsilon^2,$$

*no matter how ties are broken to select DRO solutions.*

*Proof.* Scale the hard source–target pairs from Lemmas F.10 and F.11 by  $\varepsilon$ . By Lemma F.7, their condition numbers are unchanged, their Wasserstein distances are multiplied by  $\varepsilon$ , and the DRO minimizer sets are unchanged when the supplied radius changes from 1 to  $\varepsilon$ . Their target regrets are multiplied by  $\varepsilon^2$ . Combining these bounds with Lemma F.8 and using  $\max\{a, b, c\} \geq (a + b + c)/3$  gives the claim.  $\square$

## G Remaining Proofs for Section 4

### G.1 Convex Reformulation for MRO (Proof of Theorem 4.1)

Throughout, we use the notation from Theorem 4.1 and the block decompositions from Section B.1. The observed matrix  $\tilde{\mathbf{S}}$ , its square-root factor  $\tilde{\mathbf{S}}$ , and the radius  $\varepsilon > 0$  are fixed, and their dependence is suppressed. We use  $\sup \emptyset = -\infty$ . Recalling  $\tilde{\mathbf{E}}_{\mathbf{A}} = \tilde{\mathbf{S}}_Y - \mathbf{A}\tilde{\mathbf{S}}_X$ , our goal is to prove

$$F(\mathbf{A}) = \inf_{\eta > b_{\mathbf{A}}} \eta \left( \varepsilon^2 - \|\tilde{\mathbf{S}}_X\|_{\mathbb{F}}^2 + \text{tr}_p(\mathbf{H}_{\mathbf{A},\eta}) \right) \quad \forall \mathbf{A} \in \mathcal{A},$$

where  $b_{\mathbf{A}} := 1 + \|\mathbf{A}\|_{\text{op}}^2$  and, for  $\eta > b_{\mathbf{A}}$ ,

$$\mathbf{H}_{\mathbf{A},\eta} := \tilde{\mathbf{S}}_X^\top \tilde{\mathbf{S}}_X + \tilde{\mathbf{E}}_{\mathbf{A}}^\top ((\eta - 1)\mathbf{I}_m - \mathbf{A}\mathbf{A}^\top)^{-1} \tilde{\mathbf{E}}_{\mathbf{A}}.$$

We must also show that the displayed objective, namely  $D(\mathbf{A}, \eta)$ , is jointly convex in  $(\mathbf{A}, \eta)$  on its convex domain  $\eta > b_{\mathbf{A}}$ . Infimizing over  $\mathbf{A}$  then gives the claimed convex reformulation of MRO.

By Lemmas 1.2 and B.1, the worst-case regret in the  $\varepsilon$ -neighborhood of  $\tilde{\mathbf{S}}$  is

$$F(\mathbf{A}) = \sup_{\mathbf{\Sigma} \in \mathcal{B}_\varepsilon(\tilde{\mathbf{S}})} R(\mathbf{A}, \mathbf{\Sigma}) = \sup_{\substack{\mathbf{S} \in \mathbb{R}^{d \times d} \\ \|\mathbf{S} - \tilde{\mathbf{S}}\|_{\mathbb{F}} \leq \varepsilon}} \|(\mathbf{S}_Y - \mathbf{A}\mathbf{S}_X)\mathbf{\Pi}_{\mathbf{S}_X}\|_{\mathbb{F}}^2,$$

where  $\mathbf{\Pi}_{\mathbf{S}_X} \in \mathbb{R}^{d \times d}$  denotes the orthogonal projection onto the row space of  $\mathbf{S}_X \in \mathbb{R}^{p \times d}$ . This projection has rank at most  $p$  and depends on  $\mathbf{S}_X$ . We will relax this dependence by introducing an independent projection from the family

$$\Pi_p := \{\mathbf{\Pi} \in \mathbb{S}_+^d : \mathbf{\Pi}^2 = \mathbf{\Pi}, \text{tr}(\mathbf{\Pi}) = p\}, \quad (20)$$

while retaining the constraint  $\mathbf{S}_X = \mathbf{S}_X \mathbf{\Pi}$ . This means that the row space of  $\mathbf{S}_X$  is contained in  $\text{range}(\mathbf{\Pi})$ , a condition satisfied by  $\mathbf{\Pi}_{\mathbf{S}_X}$  by definition.

For  $\mathbf{A} \in \mathcal{A}$ , a squared perturbation budget  $t > 0$ , and  $\mathbf{\Pi} \in \Pi_p$ , define

$$f_{\mathbf{\Pi}}(\mathbf{A}, t) := \sup_{\substack{\mathbf{S} \in \mathbb{R}^{d \times d} \\ \mathbf{S}_X = \mathbf{S}_X \mathbf{\Pi} \\ \|\mathbf{S} - \tilde{\mathbf{S}}\|_{\mathbb{F}}^2 \leq t}} \|(\mathbf{S}_Y - \mathbf{A}\mathbf{S}_X)\mathbf{\Pi}\|_{\mathbb{F}}^2, \quad (21)$$

$$V_{\mathbf{A}}(t) := \sup_{\mathbf{\Pi} \in \Pi_p} f_{\mathbf{\Pi}}(\mathbf{A}, t). \quad (22)$$

Every  $\mathbf{\Pi}_{\mathbf{S}_X}$  is dominated in Loewner order by some  $\mathbf{\Pi} \in \Pi_p$ . For such an extension,  $\mathbf{S}_X = \mathbf{S}_X \mathbf{\Pi}$ , and the projected squared residual can only increase. Consequently,

$$F(\mathbf{A}) \leq V_{\mathbf{A}}(\varepsilon^2).$$

The remainder of the proof proceeds in three steps.

1. **Fixed-projection duality.** Consider the Lagrangian reformulation

$$f_{\mathbf{\Pi}}(\mathbf{A}, t) = \sup_{\substack{\mathbf{S} \in \mathbb{R}^{d \times d} \\ \mathbf{S}_X = \mathbf{S}_X \mathbf{\Pi}}} \inf_{\eta \geq 0} \left\{ \|(\mathbf{S}_Y - \mathbf{A}\mathbf{S}_X)\mathbf{\Pi}\|_{\mathbb{F}}^2 + \eta(t - \|\mathbf{S} - \tilde{\mathbf{S}}\|_{\mathbb{F}}^2) \right\}. \quad (23)$$

We first prove that the supremum and infimum can be exchanged. Specifically, for  $\eta \geq 0$ , define the fixed-projection Lagrangian envelope

$$\Lambda_{\mathbf{\Pi}}(\mathbf{A}, t, \eta) := \sup_{\substack{\mathbf{S} \in \mathbb{R}^{d \times d} \\ \mathbf{S}_X = \mathbf{S}_X \mathbf{\Pi}}} \left\{ \|(\mathbf{S}_Y - \mathbf{A}\mathbf{S}_X)\mathbf{\Pi}\|_{\mathbb{F}}^2 + \eta(t - \|\mathbf{S} - \tilde{\mathbf{S}}\|_{\mathbb{F}}^2) \right\}. \quad (24)$$

We show that

$$f_{\mathbf{\Pi}}(\mathbf{A}, t) = \inf_{\eta > b_{\mathbf{A}}} \Lambda_{\mathbf{\Pi}}(\mathbf{A}, t, \eta)$$

via a change of coordinates and Lemma C.1. The same lemma also gives the closed form

$$\Lambda_{\mathbf{\Pi}}(\mathbf{A}, t, \eta) = \eta \left( t - \|\tilde{\mathbf{S}}_X\|_{\mathbb{F}}^2 + \text{tr}(\mathbf{H}_{\mathbf{A},\eta} \mathbf{\Pi}) \right), \quad \eta > b_{\mathbf{A}}. \quad (25)$$

2. **Projection–multiplier minimax exchange.** Applying the non-convex minimax theorem Ricceri (2019, Theorem 1.1), we obtain

$$\begin{aligned} V_{\mathbf{A}}(t) &= \sup_{\mathbf{\Pi} \in \Pi_p} \inf_{\eta > b_{\mathbf{A}}} \Lambda_{\mathbf{\Pi}}(\mathbf{A}, t, \eta) \\ &= \inf_{\eta > b_{\mathbf{A}}} \sup_{\mathbf{\Pi} \in \Pi_p} \Lambda_{\mathbf{\Pi}}(\mathbf{A}, t, \eta) \\ &= \inf_{\eta > b_{\mathbf{A}}} \eta \left( t - \|\tilde{\mathbf{S}}_X\|_F^2 + \text{tr}_p(\mathbf{H}_{\mathbf{A}, \eta}) \right). \end{aligned}$$

The last equality uses the Ky Fan variational formula from Lemma A.9. At  $t = \varepsilon^2$ , this matches our claimed expression for  $F(\mathbf{A})$ . It remains to show that  $F(\mathbf{A}) \geq V_{\mathbf{A}}(\varepsilon^2)$ .

3. **Exactness and joint convexity.** For fixed  $\mathbf{A}$ , the final expression above reveals that  $V_{\mathbf{A}}$  is concave in the squared budget  $t$ , and its definition implies that it is finite for every  $t > 0$ . Hence,  $V_{\mathbf{A}}$  is continuous on  $(0, \infty)$ . For  $0 < s < \varepsilon^2$ , fix  $\mathbf{\Pi} \in \Pi_p$  and any  $\mathbf{S}$  feasible for the supremum in (21) defining  $f_{\mathbf{\Pi}}(\mathbf{A}, s)$ . Using the spare perturbation budget, we approximate  $\mathbf{S}$  by factors within distance  $\varepsilon$  of  $\tilde{\mathbf{S}}$  whose feature row-space projection is exactly  $\mathbf{\Pi}$ . Their actual regrets approach  $\|(\mathbf{S}_Y - \mathbf{A}\mathbf{S}_X)\mathbf{\Pi}\|_F^2$ . Supremizing over  $\mathbf{S}$  and then  $\mathbf{\Pi}$  gives

$$V_{\mathbf{A}}(s) \leq F(\mathbf{A}) \leq V_{\mathbf{A}}(\varepsilon^2).$$

Letting  $s \uparrow \varepsilon^2$  and using continuity proves  $F(\mathbf{A}) = V_{\mathbf{A}}(\varepsilon^2)$ . Joint convexity follows by expressing  $D$  as the supremum over  $\mathbf{\Pi} \in \Pi_p$  of the envelopes  $\Lambda_{\mathbf{\Pi}}(\mathbf{A}, \varepsilon^2, \eta)$ , each jointly convex in  $(\mathbf{A}, \eta)$ .

We now establish these three steps in order.

**Lemma G.1** (Fixed-projection duality). *For every  $\mathbf{A} \in \mathcal{A}$ ,  $t > 0$ , and  $\mathbf{\Pi} \in \Pi_p$ , the Lagrangian envelope  $\Lambda_{\mathbf{\Pi}}$  satisfies (25) whenever  $\eta > b_{\mathbf{A}}$ . Moreover,*

$$f_{\mathbf{\Pi}}(\mathbf{A}, t) = \inf_{\eta \geq 0} \Lambda_{\mathbf{\Pi}}(\mathbf{A}, t, \eta) = \inf_{\eta > b_{\mathbf{A}}} \Lambda_{\mathbf{\Pi}}(\mathbf{A}, t, \eta). \quad (26)$$

*Proof.* Fix  $\mathbf{A} \in \mathcal{A}$ ,  $t > 0$ , and  $\mathbf{\Pi} \in \Pi_p$ , and write  $\mathbf{M} := (\mathbf{A} \quad -\mathbf{I}_m)$  so that  $\|\mathbf{M}\|_{\text{op}}^2 = b_{\mathbf{A}}$ .

We first reparameterize the supremum in (21) defining  $f_{\mathbf{\Pi}}$  using reduced coordinates which eliminate the row-space constraint. We then apply Lemma C.1 to establish strong duality and derive (25).

**Step 1: Reduced coordinates.** For any  $\mathbf{S} \in \mathbb{R}^{d \times d}$  satisfying  $\mathbf{S}_X = \mathbf{S}_X \mathbf{\Pi}$ , feasible for the supremum in (21), consider substituting  $\mathbf{S}_Y$  with  $\mathbf{S}_Y \mathbf{\Pi} + \tilde{\mathbf{S}}_Y(\mathbf{I}_d - \mathbf{\Pi})$ . This preserves the objective and cannot increase  $\|\mathbf{S} - \tilde{\mathbf{S}}\|_F^2$ , so it preserves feasibility. Thus, we may freely restrict the feasible set to those  $\mathbf{S}$  which further satisfy  $\mathbf{S}_Y(\mathbf{I}_d - \mathbf{\Pi}) = \tilde{\mathbf{S}}_Y(\mathbf{I}_d - \mathbf{\Pi})$ . Now fix any  $\mathbf{U} \in \mathbb{R}^{d \times p}$  whose columns form an orthonormal basis of  $\text{range}(\mathbf{\Pi})$ , so that  $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_p$  and  $\mathbf{U}\mathbf{U}^\top = \mathbf{\Pi}$ . Define the reduced coordinates of feasible  $\mathbf{S}$  by

$$\mathbf{W}_X := \mathbf{S}_X \mathbf{U} \in \mathbb{R}^{p \times p}, \quad \mathbf{W}_Y := \mathbf{S}_Y \mathbf{U} \in \mathbb{R}^{m \times p}, \quad \mathbf{W} := \begin{pmatrix} \mathbf{W}_X \\ \mathbf{W}_Y \end{pmatrix} = \mathbf{S}\mathbf{U} \in \mathbb{R}^{d \times p}.$$

On the restricted feasible set, the correspondence between  $\mathbf{S}$  and  $\mathbf{W}$  is bijective, since  $\mathbf{S}_X = \mathbf{S}_X \mathbf{\Pi} = \mathbf{W}_X \mathbf{U}^\top$ ,  $\mathbf{S}_Y \mathbf{\Pi} = \mathbf{W}_Y \mathbf{U}^\top$ , and  $\mathbf{S}_Y(\mathbf{I}_d - \mathbf{\Pi})$  is fixed. Further define

$$\tilde{\mathbf{W}}_X := \tilde{\mathbf{S}}_X \mathbf{U}, \quad \tilde{\mathbf{W}}_Y := \tilde{\mathbf{S}}_Y \mathbf{U}, \quad \tilde{\mathbf{W}} := \begin{pmatrix} \tilde{\mathbf{W}}_X \\ \tilde{\mathbf{W}}_Y \end{pmatrix} = \tilde{\mathbf{S}}\mathbf{U}.$$

We then have

$$\begin{aligned} \|\mathbf{S} - \tilde{\mathbf{S}}\|_F^2 &= \|(\mathbf{S}_X - \tilde{\mathbf{S}}_X)\mathbf{\Pi}\|_F^2 + \|(\mathbf{S}_X - \tilde{\mathbf{S}}_X)(\mathbf{I}_d - \mathbf{\Pi})\|_F^2 + \|(\mathbf{S}_Y - \tilde{\mathbf{S}}_Y)\mathbf{\Pi}\|_F^2 \\ &= \|(\mathbf{W}_X - \tilde{\mathbf{W}}_X)\mathbf{U}^\top\|_F^2 + \|\tilde{\mathbf{S}}_X(\mathbf{I}_d - \mathbf{\Pi})\|_F^2 + \|(\mathbf{W}_Y - \tilde{\mathbf{W}}_Y)\mathbf{U}^\top\|_F^2 \\ &= \|\mathbf{W} - \tilde{\mathbf{W}}\|_F^2 + \underbrace{\|\tilde{\mathbf{S}}_X(\mathbf{I}_d - \mathbf{\Pi})\|_F^2}_{:= c_{\mathbf{\Pi}}}, \end{aligned}$$

using  $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_p$ . Returning to (21), this yields

$$\begin{aligned}
f_{\Pi}(\mathbf{A}, t) &= \sup_{\substack{\mathbf{S} \in \mathbb{R}^{d \times d} \\ \mathbf{S}_X = \mathbf{S}_X \Pi \\ \mathbf{S}_Y (\mathbf{I}_d - \Pi) = \tilde{\mathbf{S}}_Y (\mathbf{I}_d - \Pi) \\ \|\mathbf{S} - \tilde{\mathbf{S}}\|_F^2 \leq t}} \|(\mathbf{S}_Y - \mathbf{A} \mathbf{S}_X) \Pi\|_F^2 \\
&= \sup_{\substack{\mathbf{W} \in \mathbb{R}^{d \times p} \\ \|\mathbf{W} - \tilde{\mathbf{W}}\|_F^2 \leq t - c_{\Pi}}} \|(\mathbf{W}_Y - \mathbf{A} \mathbf{W}_X) \mathbf{U}^\top\|_F^2 \\
&= \sup_{\substack{\mathbf{W} \in \mathbb{R}^{d \times p} \\ \|\mathbf{W} - \tilde{\mathbf{W}}\|_F^2 \leq t - c_{\Pi}}} \|\mathbf{W}_Y - \mathbf{A} \mathbf{W}_X\|_F^2. \tag{27}
\end{aligned}$$

Likewise, for all  $\eta \geq 0$ , an analogous argument gives

$$\Lambda_{\Pi}(\mathbf{A}, t, \eta) = \sup_{\mathbf{W} \in \mathbb{R}^{d \times p}} \left\{ \|\mathbf{M} \mathbf{W}\|_F^2 + \eta(t - c_{\Pi} - \|\mathbf{W} - \tilde{\mathbf{W}}\|_F^2) \right\}. \tag{28}$$

**Step 2: Apply quadratic-ball duality.** Since  $\mathbf{W}_Y - \mathbf{A} \mathbf{W}_X = -\mathbf{M} \mathbf{W}$ , the reduced formulations in (27) and (28) match the constrained and dual problems in Lemma C.1, with  $k = m$ ,  $r = p$ , center  $\tilde{\mathbf{W}}$ , and budget parameter  $t - c_{\Pi}$ . Since  $\|\mathbf{M}\|_{\text{op}}^2 = b_{\mathbf{A}}$ , that lemma gives

$$f_{\Pi}(\mathbf{A}, t) = \inf_{\eta \geq 0} \Lambda_{\Pi}(\mathbf{A}, t, \eta) = \inf_{\eta > b_{\mathbf{A}}} \Lambda_{\Pi}(\mathbf{A}, t, \eta).$$

For  $\eta > b_{\mathbf{A}}$ , (7) also gives

$$\Lambda_{\Pi}(\mathbf{A}, t, \eta) = \eta(t - c_{\Pi}) + \eta \operatorname{tr} \left( (\eta \mathbf{I}_m - \mathbf{M} \mathbf{M}^\top)^{-1} \mathbf{M} \tilde{\mathbf{W}} \tilde{\mathbf{W}}^\top \mathbf{M}^\top \right).$$

Recall that

$$\mathbf{M} \mathbf{M}^\top = \mathbf{A} \mathbf{A}^\top + \mathbf{I}_m, \quad \mathbf{M} \tilde{\mathbf{W}} = (\mathbf{A} \tilde{\mathbf{S}}_X - \tilde{\mathbf{S}}_Y) \mathbf{U} = -\tilde{\mathbf{E}}_{\mathbf{A}} \mathbf{U}.$$

Since  $\mathbf{U} \mathbf{U}^\top = \Pi$ , substitution yields

$$\Lambda_{\Pi}(\mathbf{A}, t, \eta) = \eta(t - c_{\Pi}) + \eta \operatorname{tr} \left( ((\eta - 1) \mathbf{I}_m - \mathbf{A} \mathbf{A}^\top)^{-1} \tilde{\mathbf{E}}_{\mathbf{A}} \Pi \tilde{\mathbf{E}}_{\mathbf{A}}^\top \right), \quad \eta > b_{\mathbf{A}}. \tag{29}$$

Combining (29) with the identity  $c_{\Pi} = \|\tilde{\mathbf{S}}_X\|_F^2 - \operatorname{tr}(\tilde{\mathbf{S}}_X^\top \tilde{\mathbf{S}}_X \Pi)$ , cyclicity of the trace, and the definition of  $\mathbf{H}_{\mathbf{A}, \eta}$ , we obtain (25).  $\square$

**Lemma G.2** (Projection–multiplier minimax exchange). *For every  $\mathbf{A} \in \mathcal{A}$  and  $t > 0$ ,*

$$\begin{aligned}
V_{\mathbf{A}}(t) &= \inf_{\eta > b_{\mathbf{A}}} \sup_{\Pi \in \Pi_p} \Lambda_{\Pi}(\mathbf{A}, t, \eta) \\
&= \inf_{\eta > b_{\mathbf{A}}} \eta \left( t - \|\tilde{\mathbf{S}}_X\|_F^2 + \operatorname{tr}_p(\mathbf{H}_{\mathbf{A}, \eta}) \right). \tag{30}
\end{aligned}$$

*Proof.* Fix  $\mathbf{A} \in \mathcal{A}$  and  $t > 0$ , and write

$$I := (b_{\mathbf{A}}, \infty), \quad I_j := [b_{\mathbf{A}} + 1/j, b_{\mathbf{A}} + j], \quad j \geq 2.$$

These nested compact intervals exhaust  $I$ . By Lemma G.1 and the definition of  $V_{\mathbf{A}}$ ,

$$V_{\mathbf{A}}(t) = \sup_{\Pi \in \Pi_p} \inf_{\eta \in I} \Lambda_{\Pi}(\mathbf{A}, t, \eta).$$

We now apply the minimax theorem of Ricceri (2019, Theorem 1.1) to the joint objective  $g(\Pi, \eta) := -\Lambda_{\Pi}(\mathbf{A}, t, \eta)$ , with exhaustion sequence  $\{I_j\}_{j \geq 2}$ . This requires checking the following assumptions.

**Compactness, continuity, and convexity.** The set  $\Pi_p$  is closed, since its defining constraints in (20) are continuous. Moreover,  $\|\Pi\|_F^2 = \operatorname{tr}(\Pi) = p$  for each  $\Pi \in \Pi_p$ , so it is in fact compact. This implies that each  $\Pi \mapsto \inf_{\eta \in I_j} g(\Pi, \eta)$  is relatively inf-compact (see the definition in Ricceri, 2019) because the domain  $\Pi_p$  is compact. Equation (25), established in Lemma G.1, shows that  $(\Pi, \eta) \mapsto \Lambda_{\Pi}(\mathbf{A}, t, \eta)$  is jointly continuous on  $\Pi_p \times I$ . Finally, for fixed  $\Pi$ ,  $\eta \mapsto \Lambda_{\Pi}(\mathbf{A}, t, \eta)$  is convex as a supremum of affine functions in (24).

**Connectedness of best-response sets.** Fixing  $\eta \in I$ , we next prove that the set of projections  $\Pi \in \Pi_p$  maximizing  $\Lambda_\Pi(\mathbf{A}, t, \eta)$  is connected. Since  $\eta > 0$ , (25) gives

$$\arg \max_{\Pi \in \Pi_p} \Lambda_\Pi(\mathbf{A}, t, \eta) = \arg \max_{\Pi \in \Pi_p} \text{tr}(\mathbf{H}_{\mathbf{A}, \eta} \Pi).$$

List the eigenvalues of  $\mathbf{H}_{\mathbf{A}, \eta} \in \mathbb{S}_+^d$  in decreasing order, counted with multiplicity, and let  $\alpha$  be the  $p$ th eigenvalue. Let  $\mathbf{P}_{>}$  project onto the direct sum of eigenspaces with eigenvalues strictly greater than  $\alpha$ , set  $r_{>} := \text{rank}(\mathbf{P}_{>}) < p$ , and write  $\mathcal{V} := \text{null}(\mathbf{H}_{\mathbf{A}, \eta} - \alpha \mathbf{I}_d)$ . The equality conditions in the Ky Fan principle (Lemma A.9) show that every maximizing projection is  $\mathbf{P}_{>} + \mathbf{Q}$ , where  $\mathbf{Q}$  projects onto an  $r$ -dimensional subspace of  $\mathcal{V}$ , with  $r := p - r_{>}$ .

Set  $n := \dim \mathcal{V}$ . If  $r = n$ , the maximizing set is a singleton. Otherwise, choose  $\mathbf{V} \in \mathbb{R}^{d \times n}$  whose columns form an orthonormal basis of  $\mathcal{V}$ , and set  $\mathbf{Q}_0 := \text{diag}(\mathbf{I}_r, \mathbf{0}_{n-r})$ . The maximizing set is then

$$\{\mathbf{P}_{>} + \mathbf{V} \mathbf{O} \mathbf{Q}_0 \mathbf{O}^\top \mathbf{V}^\top : \mathbf{O} \in \mathbb{O}^n, \det(\mathbf{O}) = 1\}.$$

Indeed, allowing all  $\mathbf{O} \in \mathbb{O}^n$  generates every projection onto an  $r$ -dimensional subspace of  $\mathcal{V}$ . The determinant restriction loses none of them, since negating one column of  $\mathbf{O}$  changes its determinant sign but leaves  $\mathbf{O} \mathbf{Q}_0 \mathbf{O}^\top$  unchanged because  $\mathbf{Q}_0$  is diagonal. The displayed set is therefore a continuous image of the special orthogonal group in dimension  $n$ , which is path-connected. Hence the maximizing set is connected.

Lastly, fixing  $\Pi \in \Pi_p$  and  $j \geq 2$ , we must show that the set of  $\eta \in I_j$  minimizing  $\Lambda_\Pi(\mathbf{A}, t, \eta)$  is nonempty and connected. This follows directly by continuity and convexity of the objective in  $\eta$ .

To conclude, Ricceri's theorem then gives

$$V_{\mathbf{A}}(t) = \inf_{\eta > b_{\mathbf{A}}} \sup_{\Pi \in \Pi_p} \Lambda_\Pi(\mathbf{A}, t, \eta).$$

Substituting (25) and applying Lemma A.9 gives

$$\sup_{\Pi \in \Pi_p} \Lambda_\Pi(\mathbf{A}, t, \eta) = \eta \left( t - \|\tilde{\mathbf{S}}_X\|_{\text{F}}^2 + \text{tr}_p(\mathbf{H}_{\mathbf{A}, \eta}) \right),$$

which proves the second equality in (30).  $\square$

**Lemma G.3** (Exactness and joint convexity). *For each  $\mathbf{A} \in \mathcal{A}$ , the function  $t \mapsto V_{\mathbf{A}}(t)$  is finite, concave, and continuous on  $(0, \infty)$ , and*

$$F(\mathbf{A}) = V_{\mathbf{A}}(\varepsilon^2).$$

Moreover,  $D$ , extended by  $+\infty$  outside  $\eta > b_{\mathbf{A}}$ , is jointly convex on  $\mathcal{A} \times \mathbb{R}$ .

*Proof.* We prove the claims separately.

**Finiteness, concavity, and continuity.** Fixing  $\mathbf{A} \in \mathcal{A}$ , the supremum in (22) defining  $V_{\mathbf{A}}(t)$  is feasible for each  $t > 0$ . Indeed, we may take  $\mathbf{S} = \tilde{\mathbf{S}}$  and let  $\Pi \in \Pi_p$  project onto a  $p$ -dimensional subspace containing the row space of  $\tilde{\mathbf{S}}_X$ . Moreover, since each  $\Pi \in \Pi_p$  is dominated in Loewner order by  $\mathbf{I}_d$ , we have

$$\begin{aligned} 0 &\leq V_{\mathbf{A}}(t) \\ &\leq \sup_{\substack{\mathbf{S} \in \mathbb{R}^{d \times d} \\ \|\mathbf{S} - \tilde{\mathbf{S}}\|_{\text{F}}^2 \leq t}} \|\mathbf{S}_Y - \mathbf{A} \mathbf{S}_X\|_{\text{F}}^2 \\ &\leq b_{\mathbf{A}} \sup_{\substack{\mathbf{S} \in \mathbb{R}^{d \times d} \\ \|\mathbf{S} - \tilde{\mathbf{S}}\|_{\text{F}}^2 \leq t}} \|\mathbf{S}\|_{\text{F}}^2 && \text{(Lemma A.1)} \\ &\leq b_{\mathbf{A}} (\|\tilde{\mathbf{S}}\|_{\text{F}} + \sqrt{t})^2 < \infty. \end{aligned}$$

By Lemma G.2,  $t \mapsto V_{\mathbf{A}}(t)$  is the infimum of affine functions of  $t$ , hence concave. Being finite and concave on the open interval  $(0, \infty)$ , this function is continuous there.

**Exactness of the projection relaxation.** The inequality  $F(\mathbf{A}) \leq V_{\mathbf{A}}(\varepsilon^2)$  was established above. For the reverse inequality, fix  $0 < s < \varepsilon^2$  and a pair  $(\mathbf{\Pi}, \mathbf{S}) \in \Pi_p \times \mathbb{R}^{d \times d}$  satisfying

$$\mathbf{S}_X = \mathbf{S}_X \mathbf{\Pi}, \quad \|\mathbf{S} - \tilde{\mathbf{S}}\|_{\text{F}}^2 \leq s.$$

Now choose  $\mathbf{U} \in \mathbb{R}^{d \times p}$  with  $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_p$  and  $\mathbf{\Pi} = \mathbf{U} \mathbf{U}^\top$ . Since  $\mathbf{S}_X = \mathbf{S}_X \mathbf{\Pi}$ , we may write  $\mathbf{S}_X = \mathbf{C} \mathbf{U}^\top$  with  $\mathbf{C} := \mathbf{S}_X \mathbf{U} \in \mathbb{R}^{p \times p}$ . Then, for  $\tau > 0$ , define

$$\mathbf{S}_X^{(\tau)} := (\mathbf{C} + \tau \mathbf{I}_p) \mathbf{U}^\top, \quad \mathbf{S}^{(\tau)} := \begin{pmatrix} \mathbf{S}_X^{(\tau)} \\ \tilde{\mathbf{S}}_Y \end{pmatrix}.$$

The polynomial  $\det(\mathbf{C} + \tau \mathbf{I}_p)$  has leading coefficient 1, so it is nonzero for sufficiently small  $\tau > 0$ . For such  $\tau$ ,  $\mathbf{S}_X^{(\tau)}$  has full row rank and  $\mathbf{\Pi}_{\mathbf{S}_X^{(\tau)}} = \mathbf{\Pi}$ . Moreover,

$$\|\mathbf{S}^{(\tau)} - \tilde{\mathbf{S}}\|_{\text{F}} \leq \|\mathbf{S} - \tilde{\mathbf{S}}\|_{\text{F}} + \|\mathbf{S}^{(\tau)} - \mathbf{S}\|_{\text{F}} \leq \sqrt{s} + \tau \sqrt{p} < \varepsilon$$

for sufficiently small  $\tau$ . By Lemma 1.2,  $\mathbf{\Sigma}^{(\tau)} := \mathbf{S}^{(\tau)} (\mathbf{S}^{(\tau)})^\top$  thus belongs to  $\mathcal{B}_\varepsilon(\tilde{\mathbf{\Sigma}})$  for sufficiently small  $\tau$ . Hence, Lemma B.1 gives

$$F(\mathbf{A}) \geq \lim_{\tau \downarrow 0} R(\mathbf{A}, \mathbf{\Sigma}^{(\tau)}) = \lim_{\tau \downarrow 0} \|(\mathbf{S}_Y - \mathbf{A} \mathbf{S}_X^{(\tau)}) \mathbf{\Pi}\|_{\text{F}}^2 = \|(\mathbf{S}_Y - \mathbf{A} \mathbf{S}_X) \mathbf{\Pi}\|_{\text{F}}^2.$$

Taking the supremum over  $(\mathbf{\Pi}, \mathbf{S})$  gives

$$V_{\mathbf{A}}(s) \leq F(\mathbf{A}) \leq V_{\mathbf{A}}(\varepsilon^2).$$

Letting  $s \uparrow \varepsilon^2$  and using continuity proves  $F(\mathbf{A}) = V_{\mathbf{A}}(\varepsilon^2)$ .

**Joint convexity.** For fixed  $\mathbf{S}, \mathbf{\Pi}$ , the expression

$$\|(\mathbf{S}_Y - \mathbf{A} \mathbf{S}_X) \mathbf{\Pi}\|_{\text{F}}^2 + \eta(\varepsilon^2 - \|\mathbf{S} - \tilde{\mathbf{S}}\|_{\text{F}}^2)$$

is convex in  $\mathbf{A}$  and affine in  $\eta$ , with no mixed term. It is therefore jointly convex in  $(\mathbf{A}, \eta)$ . The constraint  $\mathbf{S}_X = \mathbf{S}_X \mathbf{\Pi}$  is independent of these variables, so taking the supremum in (24) preserves joint convexity. Thus  $\Lambda_{\mathbf{\Pi}}(\cdot, \varepsilon^2, \cdot)$  is jointly convex on  $\mathcal{A} \times [0, \infty)$  for every fixed  $\mathbf{\Pi}$ . By (25) and the Ky Fan formula,

$$D(\mathbf{A}, \eta) = \sup_{\mathbf{\Pi} \in \Pi_p} \Lambda_{\mathbf{\Pi}}(\mathbf{A}, \varepsilon^2, \eta) \quad \text{when } \eta > b_{\mathbf{A}}.$$

Hence  $D$  is jointly convex on this domain. The domain itself is convex, being the strict epigraph of  $\mathbf{A} \mapsto 1 + \|\mathbf{A}\|_{\text{op}}^2$ . Extending  $D$  by  $+\infty$  outside this domain preserves convexity.  $\square$

*Proof of Theorem 4.1.* By Lemmas G.2 and G.3,

$$F(\mathbf{A}) = V_{\mathbf{A}}(\varepsilon^2) = \inf_{\eta > b_{\mathbf{A}}} D(\mathbf{A}, \eta) \quad \forall \mathbf{A} \in \mathcal{A},$$

and  $D$  is jointly convex as an extended-value function. Taking the infimum over  $\mathbf{A}$  gives

$$\inf_{\mathbf{A} \in \mathcal{A}} F(\mathbf{A}) = \inf_{\mathbf{A} \in \mathcal{A}, \eta \in \mathbb{R}} D(\mathbf{A}, \eta),$$

as claimed.  $\square$

## G.2 Subgradient Method Convergence for SNR (Proof of Proposition 4.2)

If  $\tilde{\mathbf{\Sigma}}_Y = \mathbf{0}$ , then  $G(\mathbf{0}) = 0$ . Since  $G$  is always non-negative,  $\mathbf{0}$  is a minimizer, as claimed.

Otherwise, recall that

$$\begin{aligned} B &:= (d \operatorname{tr} \tilde{\mathbf{\Sigma}}_Y)^{1/2} / \varepsilon \\ M &:= 2(B \|\tilde{\mathbf{\Sigma}}_X\|_{\text{op}} + \|\tilde{\mathbf{\Sigma}}_{YX}\|_{\text{F}} + \varepsilon^2 B) \\ T &:= \lceil M^2 B^2 \delta^{-2} \rceil. \end{aligned}$$

First, we note a simple bound on the norm of a minimizer of  $G(\mathbf{A}) = L(\mathbf{A}, \tilde{\mathbf{\Sigma}}) + \varepsilon^2 \|\mathbf{A}\|_{\text{op}}^2$ .

**Lemma G.4** (SNR solution bound). *Each minimizer  $\mathbf{A}_\star$  of  $G$  satisfies  $\|\mathbf{A}_\star\|_F \leq B$ .*

*Proof.* By optimality, we have

$$\varepsilon^2 \|\mathbf{A}_\star\|_{\text{op}}^2 \leq G(\mathbf{A}_\star) \leq G(\mathbf{0}) = \text{tr}(\tilde{\Sigma}_Y),$$

and so  $\|\mathbf{A}_\star\|_F \leq \sqrt{d} \|\mathbf{A}_\star\|_{\text{op}} \leq (d \text{tr}(\tilde{\Sigma}_Y))^{1/2} / \varepsilon = B$ .  $\square$

Thus, writing  $\mathcal{K}$  for the Frobenius ball of radius  $B$  centered at  $\mathbf{A}_1 = \mathbf{0}$ , we have  $G(\mathbf{A}_\star) = \min_{\mathbf{A} \in \mathcal{K}} G(\mathbf{A})$ . Moreover, the projection operator appearing in the proposition is exactly the Frobenius projection onto  $\mathcal{K}$ .

Next, we claim that  $G$  is  $M$ -Lipschitz under the Frobenius norm over the domain  $\mathcal{K}$ . Indeed, we can write the subdifferential set  $\partial G(\mathbf{A})$  as the convex hull

$$\text{conv} \left\{ 2\mathbf{A}\tilde{\Sigma}_X - 2\tilde{\Sigma}_{YX} + 2\varepsilon^2 \|\mathbf{A}\|_{\text{op}} uv^\top : u \in \mathbb{R}^m, v \in \mathbb{R}^p, \|u\| = \|v\| = 1, \mathbf{A}v = \|\mathbf{A}\|_{\text{op}} u \right\},$$

Here, the operator-norm subgradient is handled via the characterization of the spectral-norm subdifferential set in Example 2.8 of Lewis (1995). Thus the Lipschitz constant of  $G$  over  $\mathcal{K}$  is at most

$$\sup_{\substack{\mathbf{A} \in \mathcal{K} \\ u \in \mathbb{R}^m, \|u\| \leq 1 \\ v \in \mathbb{R}^p, \|v\| \leq 1}} \left\| 2\mathbf{A}\tilde{\Sigma}_X - 2\tilde{\Sigma}_{YX} + 2\varepsilon^2 \|\mathbf{A}\|_{\text{op}} uv^\top \right\|_F \leq M,$$

using the triangle inequality and Lemma A.1.

Noting that the Frobenius geometry is Euclidean, Proposition 4.2 then follows by standard guarantees for the projected subgradient method (e.g., Theorem 3.2 of Bubeck, 2015).  $\square$

### G.3 FISTA Convergence for SNR (Proof of Proposition 4.3)

First, we observe that Algorithm 2, given input  $\mathbf{B}$  and parameter  $\alpha$ , evaluates the proximal operator of  $\alpha \|\cdot\|_{\text{op}}^2$  at  $\mathbf{B}$ . In other words, it solves the minimization problem

$$\min_{\mathbf{C} \in \mathbb{R}^{m \times p}} \alpha \|\mathbf{C}\|_{\text{op}}^2 + \frac{1}{2} \|\mathbf{C} - \mathbf{B}\|_F^2.$$

Indeed, by the discussion of orthogonally invariant matrix functions in Parikh and Boyd (2014, Sections 6.7.2–6.7.3), the solution to the proximal problem above can be obtained by first computing an SVD of  $\mathbf{B}$  and then applying the  $\ell_\infty^2$  proximal operator to the vector of its singular values. The clipping procedure of Algorithm 2 is the squared- $\ell_\infty$  proximal operator derived in Jacobsson et al. (2017, Theorem 4).

Given that Algorithm 2 correctly implements the proximal operator for  $\alpha \|\cdot\|_{\text{op}}^2$ , Algorithm 1 is simply a standard application of the FISTA algorithm (Beck and Teboulle, 2009), applied to the objective  $G$ , decomposed as the sum of the smooth term  $L(\cdot, \tilde{\Sigma})$  and the proximable term  $\varepsilon^2 \|\cdot\|_{\text{op}}^2$ . In particular, letting  $\mathbf{A}_\star$  denote any minimizer of  $G$ , Theorem 4.4 of Beck and Teboulle (2009) gives

$$G(\mathbf{A}_T) - G(\mathbf{A}_\star) \leq \frac{2M \|\mathbf{A}_\star\|_F^2}{(T+1)^2},$$

using that  $L(\cdot, \tilde{\Sigma})$  is  $M = 2\lambda_{\max}(\tilde{\Sigma}_X)$  smooth. By Lemma G.4,  $\|\mathbf{A}_\star\|_F^2 \leq d \text{tr}(\tilde{\Sigma}_Y) / \varepsilon^2$ , and so  $\mathbf{A}_T$  has additive approximation error  $\delta$  once  $T \geq (2Md \text{tr}(\tilde{\Sigma}_Y))^{1/2} \varepsilon^{-1} \cdot \delta^{-1/2}$ , as claimed.

## H Remaining Proofs for Section 5

### H.1 Properties of the Bures Variance Proxy (Proof of Proposition 5.1)

Let  $Z \sim P$  and abbreviate  $\Sigma := \Sigma_P$ . Since  $\mathbb{V}_B(P)$  is the expectation of a quartic in the entries of  $Z$ , it is finite for  $P \in \mathcal{P}_4(\mathbb{R}^d)$ . By non-negativity of the Frobenius norm,  $\mathbb{V}_B(P)$  vanishes if and only if

$$(ZZ^\top - \Sigma)\Sigma_P^{\dagger/2} = 0 \iff ZZ^\top = \Sigma$$

almost surely. The equivalence holds because  $Z$  lies in the range of  $\Sigma$  almost surely.

Expanding the squared norm in the definition of  $V_B$  and using  $\Sigma\Sigma^\dagger Z = Z$  almost surely gives

$$\left\| (ZZ^\top - \Sigma)\Sigma^{\dagger/2} \right\|_F^2 = \|Z\|^2 Z^\top \Sigma^\dagger Z - 2\|Z\|^2 + \text{tr}(\Sigma).$$

Taking expectations and using  $\mathbb{E}_P \|Z\|^2 = \text{tr}(\Sigma)$  proves the equality in (3).

For  $\Sigma \neq 0$ , let  $\lambda_{\min}^+(\Sigma)$  denote its smallest positive eigenvalue. The fourth-moment bound follows from the inequality

$$Z^\top \Sigma^\dagger Z \leq \frac{\|Z\|^2}{\lambda_{\min}^+(\Sigma)}$$

on  $\text{range}(\Sigma)$ , implying that

$$V_B(P) \leq \mathbb{E}_P [\|Z\|^2 Z^\top \Sigma^\dagger Z] \leq \frac{\mathbb{E} \|Z\|^4}{\lambda_{\min}^+(\Sigma)}.$$

This is finite whenever  $P \in \mathcal{P}_4(\mathbb{R}^d)$  and  $\Sigma \neq 0$ . The  $\Sigma = 0$  case requires  $Z = 0$  almost surely and hence  $V_B(P) = 0$ . When  $\Sigma \succ 0$ ,  $\lambda_{\min}^+(\Sigma) = \lambda_{\min}(\Sigma)$ , giving the stated bound.

Now suppose that  $P = \mathcal{N}(0, \Sigma)$  is Gaussian, and assume without loss of generality that  $\Sigma$  is diagonalizable in the standard basis, with eigenvalues  $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d$ . Restricting to  $\text{range}(\Sigma)$ , let  $r := \text{rank}(\Sigma)$  and write

$$Z = \sum_{j=1}^r \sqrt{\lambda_j} G_j e_j, \quad G_j \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1).$$

This gives

$$Z^\top \Sigma^\dagger Z = \sum_{j=1}^r G_j^2, \quad \|Z\|^2 = \sum_{j=1}^r \lambda_j G_j^2,$$

and so we compute

$$\begin{aligned} \mathbb{E}_P [\|Z\|^2 Z^\top \Sigma^\dagger Z] &= \sum_{i=1}^r \lambda_i \mathbb{E} \left[ G_i^2 \sum_{j=1}^r G_j^2 \right] \\ &= \sum_{i=1}^r \lambda_i (3 + (r-1)) \\ &= (r+2) \text{tr}(\Sigma). \end{aligned}$$

Therefore,  $V_B(P) = (r+2) \text{tr}(\Sigma) - \text{tr}(\Sigma) = (r+1) \text{tr}(\Sigma)$ .

If further  $P \in \mathcal{P}^*$  is non-degenerate, we then have  $r = \text{rank}(\Sigma) \geq \text{rank}(\Sigma_X) = p$ . Moreover,

$$\text{tr}(\Sigma) \geq \text{tr}(\Sigma_Y) = \text{tr}(\mathbf{A}_P \Sigma_X \mathbf{A}_P^\top) + \text{tr}(\Sigma_E) \geq \text{tr}(\Sigma_E).$$

Combining with the above gives  $V_B(P) \geq (p+1) \text{tr}(\Sigma_E) \geq p \text{tr}(\Sigma_E)$ , as claimed.

It remains to prove the empirical Bures bound. Since each  $Z_i \in \text{range}(\Sigma)$  almost surely, we have  $\text{range}(\Sigma_{P_n}) \subseteq \text{range}(\Sigma)$  almost surely. Therefore Lemma B.2 gives

$$B(\Sigma, \Sigma_{P_n})^2 \leq \text{tr}((\Sigma_{P_n} - \Sigma)\Sigma^\dagger(\Sigma_{P_n} - \Sigma)).$$

Next, set  $\mathbf{H}_i := Z_i Z_i^\top - \Sigma$ , so that  $\Sigma_{P_n} - \Sigma = \sum_{i=1}^n \mathbf{H}_i / n$ . The  $\mathbf{H}_i$  matrices are independent and mean zero, so the cross terms vanish in expectation:

$$\mathbb{E} \text{tr}((\Sigma_{P_n} - \Sigma)\Sigma^\dagger(\Sigma_{P_n} - \Sigma)) = \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \text{tr}(\mathbf{H}_i \Sigma^\dagger \mathbf{H}_i) = \frac{1}{n} V_B(P),$$

where the last equality uses (3). This proves the claim.  $\square$

## H.2 Finite-Sample Regret Bounds (Proof of Theorem 5.2)

We first establish some empirical estimates needed for the sub-Gaussian guarantee.

**Lemma H.1** (Empirical conditioning and ERM error). *Fix  $P \in \mathcal{P}^*$  and  $\delta \in (0, 1)$ . Let  $(X, Y) \sim P$  and  $E = Y - \mathbf{A}_P X$ . Suppose that the random vector  $((\Sigma_X^P)^{-1/2} X, (\Sigma_E^P)^{\dagger/2} E)$  is  $K$ -sub-Gaussian. Then for  $n \gtrsim K^4(d + \log(1/\delta))$ , the following hold simultaneously with probability at least  $1 - \delta$ :*

$$\begin{aligned} \frac{3}{4} \Sigma_X^P &\preceq \Sigma_X^{P_n} \preceq \frac{5}{4} \Sigma_X^P, \\ \kappa_{P_n} &\leq 4\kappa_P, \\ R(\text{ERM}[P_n], P) &\lesssim \frac{K^4 \text{tr}(\Sigma_E^P)(d + \log(1/\delta))}{n}. \end{aligned}$$

In particular,  $P_n \in \mathcal{P}^*$  on this event.

*Proof.* Define the random vectors

$$\bar{X} := (\Sigma_X^P)^{-1/2} X, \quad \bar{E} := (\Sigma_E^P)^{\dagger/2} E, \quad W := (\bar{X}, \bar{E}).$$

We will consider both  $P$  and  $P_n$  as their base probability law. Let  $\Pi_E$  project onto  $\text{range}(\Sigma_E^P)$ . Then  $E \in \text{range}(\Sigma_E^P)$  almost surely, and the population normal equation gives  $\mathbb{E}_P[EX^\top] = \mathbf{0}$ . Therefore,

$$\Sigma_W^P := \mathbb{E}_P[WW^\top] = \text{diag}(\mathbf{I}_p, \Pi_E).$$

The vector  $W$  is supported on  $\mathbb{R}^p \times \text{range}(\Sigma_E^P)$  and has identity second moment on this subspace, whose dimension is at most  $d$ . The feature directions have unit second moment, so  $K \geq 1$ . Applying Vershynin (2026, Theorem 4.7.1 and Remark 4.7.3) on this subspace gives an event  $\mathcal{E}$  of probability at least  $1 - \delta$  on which

$$\|\mathbb{E}_{P_n}[WW^\top] - \Sigma_W^P\|_{\text{op}} \leq \alpha_n := CK^2 \sqrt{\frac{d + \log(1/\delta)}{n}} \leq \frac{1}{4}.$$

Here  $C$  is an absolute constant, and the sample-size assumption absorbs the linear term in the cited concentration bound. The second-moment formulation does not require  $W$  to be centered.

Define the empirical blocks

$$\begin{aligned} \mathbf{H}_n &:= \mathbb{E}_{P_n}[\bar{X}\bar{X}^\top] \in \mathbb{R}^{p \times p}, \\ \mathbf{C}_n &:= \mathbb{E}_{P_n}[\bar{E}\bar{X}^\top] \in \mathbb{R}^{m \times p}, \\ \mathbf{J}_n &:= \mathbb{E}_{P_n}[\bar{E}\bar{E}^\top] \in \mathbb{R}^{m \times m}. \end{aligned}$$

On  $\mathcal{E}$  we have

$$\|\mathbf{H}_n - \mathbf{I}_p\|_{\text{op}}, \|\mathbf{C}_n\|_{\text{op}}, \|\mathbf{J}_n - \Pi_E\|_{\text{op}} \leq \alpha_n \leq \frac{1}{4}. \quad (31)$$

Condition on this event for the remainder of the proof. Since

$$\Sigma_X^{P_n} = (\Sigma_X^P)^{1/2} \mathbf{H}_n (\Sigma_X^P)^{1/2},$$

we obtain the claimed bounds for the empirical feature second-moment matrix, with the lower bound implying that  $P_n \in \mathcal{P}^*$ . Also,  $\|\mathbf{H}_n^{-1}\|_{\text{op}} \leq (1 - \alpha_n)^{-1} \leq 4/3$ .

The identities  $Y = \mathbf{A}_P X + E$  and  $E = (\Sigma_E^P)^{1/2} \bar{E}$  imply that

$$\Sigma_{YX}^{P_n} = \mathbf{A}_P \Sigma_X^{P_n} + \mathbb{E}_{P_n}[EX^\top] = \mathbf{A}_P \Sigma_X^{P_n} + (\Sigma_E^P)^{1/2} \mathbf{C}_n (\Sigma_X^P)^{1/2}.$$

Combining the above with the formula for  $\mathbf{A}_{P_n}$  gives

$$\begin{aligned} \mathbf{A}_{P_n} - \mathbf{A}_P &= \Sigma_{YX}^{P_n} (\Sigma_X^{P_n})^{-1} - \mathbf{A}_P \\ &= (\Sigma_E^P)^{1/2} \mathbf{C}_n (\Sigma_X^P)^{1/2} (\Sigma_X^{P_n})^{-1} \\ &= (\Sigma_E^P)^{1/2} \mathbf{C}_n \mathbf{H}_n^{-1} (\Sigma_X^P)^{-1/2}. \end{aligned} \quad (32)$$

Consequently,

$$\|\mathbf{A}_{P_n} - \mathbf{A}_P\|_{\text{op}} \leq \frac{\alpha_n}{1 - \alpha_n} \sqrt{\frac{\lambda_{\max}(\Sigma_E^P)}{\lambda_{\min}(\Sigma_X^P)}} \leq \frac{1}{3} \sqrt{\frac{\lambda_{\max}(\Sigma_E^P)}{\lambda_{\min}(\Sigma_X^P)}}.$$

Using (32), the empirical fitted residual can be written as

$$Y - \mathbf{A}_{P_n} X = E + (\mathbf{A}_P - \mathbf{A}_{P_n})X = (\boldsymbol{\Sigma}_E^P)^{1/2}(\bar{E} - \mathbf{C}_n \mathbf{H}_n^{-1} \bar{X}).$$

Using the definitions of  $\mathbf{C}_n$ ,  $\mathbf{H}_n$  and  $\mathbf{J}_n$ , along with symmetry of  $\mathbf{H}_n^{-1}$ , we compute

$$\begin{aligned} & \mathbb{E}_{P_n} [(\bar{E} - \mathbf{C}_n \mathbf{H}_n^{-1} \bar{X})(\bar{E} - \mathbf{C}_n \mathbf{H}_n^{-1} \bar{X})^\top] \\ &= \mathbf{J}_n - \mathbb{E}_{P_n} [\bar{E} \bar{X}^\top] \mathbf{H}_n^{-1} \mathbf{C}_n^\top - \mathbf{C}_n \mathbf{H}_n^{-1} \mathbb{E}_{P_n} [\bar{X} \bar{E}^\top] + \mathbf{C}_n \mathbf{H}_n^{-1} \mathbb{E}_{P_n} [\bar{X} \bar{X}^\top] \mathbf{H}_n^{-1} \mathbf{C}_n^\top \\ &= \mathbf{J}_n - 2\mathbf{C}_n \mathbf{H}_n^{-1} \mathbf{C}_n^\top + \mathbf{C}_n \mathbf{H}_n^{-1} \mathbf{H}_n \mathbf{H}_n^{-1} \mathbf{C}_n^\top \\ &= \mathbf{J}_n - \mathbf{C}_n \mathbf{H}_n^{-1} \mathbf{C}_n^\top. \end{aligned}$$

Consequently, we have

$$\begin{aligned} \boldsymbol{\Sigma}_E^{P_n} &= \mathbb{E}_{P_n} [(Y - \mathbf{A}_{P_n} X)(Y - \mathbf{A}_{P_n} X)^\top] \\ &= (\boldsymbol{\Sigma}_E^P)^{1/2} (\mathbf{J}_n - \mathbf{C}_n \mathbf{H}_n^{-1} \mathbf{C}_n^\top) (\boldsymbol{\Sigma}_E^P)^{1/2} \\ &\preceq (\boldsymbol{\Sigma}_E^P)^{1/2} \mathbf{J}_n (\boldsymbol{\Sigma}_E^P)^{1/2} \\ &\preceq \frac{5}{4} \boldsymbol{\Sigma}_E^P. \end{aligned}$$

The last inequality follows from  $\mathbf{J}_n \preceq \boldsymbol{\Pi}_E + \alpha_n \mathbf{I}_m$  and  $(\boldsymbol{\Sigma}_E^P)^{1/2} \boldsymbol{\Pi}_E = (\boldsymbol{\Sigma}_E^P)^{1/2}$ . Combining the above with Theorem 2.2 gives

$$\begin{aligned} \kappa_{P_n} &\leq 1 + \|\mathbf{A}_{P_n}\|_{\text{op}}^2 + \frac{\lambda_{\max}(\boldsymbol{\Sigma}_E^{P_n})}{\lambda_{\min}(\boldsymbol{\Sigma}_X^P)} \\ &\leq 1 + 2\|\mathbf{A}_P\|_{\text{op}}^2 + 2\|\mathbf{A}_{P_n} - \mathbf{A}_P\|_{\text{op}}^2 + \frac{\frac{5}{4}\lambda_{\max}(\boldsymbol{\Sigma}_E^P)}{\frac{3}{4}\lambda_{\min}(\boldsymbol{\Sigma}_X^P)} \\ &\leq 1 + 2\|\mathbf{A}_P\|_{\text{op}}^2 + \left(\frac{2}{9} + \frac{5}{3}\right) \frac{\lambda_{\max}(\boldsymbol{\Sigma}_E^P)}{\lambda_{\min}(\boldsymbol{\Sigma}_X^P)} \\ &\leq 4\kappa_P. \end{aligned}$$

Finally, (32) gives

$$\begin{aligned} R(\text{ERM}[P_n], P) &= \|(\mathbf{A}_{P_n} - \mathbf{A}_P)(\boldsymbol{\Sigma}_X^P)^{1/2}\|_{\text{F}}^2 \\ &= \|(\boldsymbol{\Sigma}_E^P)^{1/2} \mathbf{C}_n \mathbf{H}_n^{-1}\|_{\text{F}}^2 \\ &\leq \frac{\alpha_n^2}{(1 - \alpha_n)^2} \text{tr}(\boldsymbol{\Sigma}_E^P) \\ &\lesssim \frac{K^4 \text{tr}(\boldsymbol{\Sigma}_E^P)(d + \log(1/\delta))}{n}, \end{aligned}$$

as claimed.  $\square$

We now prove the two guarantees separately. For each, we establish a high-probability event depending only on the i.i.d. sample from  $P$ , under which the claimed bounds hold uniformly over the admissible choices in both perturbation models.

**Fourth-moment guarantee.** Abbreviate  $\rho := \varepsilon_n$ . By Proposition 5.1,  $\mathbb{E}[\text{B}(\boldsymbol{\Sigma}_{P_n}, \boldsymbol{\Sigma}_P)^2] \leq V/n$ , so Markov's inequality gives  $\text{B}(\boldsymbol{\Sigma}_{P_n}, \boldsymbol{\Sigma}_P) \leq \sqrt{V/(\delta n)}$  with probability at least  $1 - \delta$ . On this event,

$$\text{B}(\boldsymbol{\Sigma}_{P_{s,n}}, \boldsymbol{\Sigma}_P) + \text{B}(\boldsymbol{\Sigma}_{P_t}, \boldsymbol{\Sigma}_P) \leq \varepsilon + \sqrt{\frac{V}{\delta n}} = \rho \quad (33)$$

under either adversary model. Indeed, under the sample-corruption model, we have  $P_t = P$  and  $W_2(P_{s,n}, P_n) \leq \varepsilon$ , so the claim follows by the  $W_2$ -Bures equivalence (Lemma 1.2). Likewise, under the distribution-shift model, we have  $P_{s,n} = P_n$  and  $\text{B}(\boldsymbol{\Sigma}_{P_t}, \boldsymbol{\Sigma}_P) \leq W_2(P_t, P) \leq \varepsilon$ .

Equation (33) implies that  $\text{B}(\boldsymbol{\Sigma}_{P_{s,n}}, \boldsymbol{\Sigma}_P), \text{B}(\boldsymbol{\Sigma}_{P_t}, \boldsymbol{\Sigma}_P), \text{B}(\boldsymbol{\Sigma}_{P_{s,n}}, \boldsymbol{\Sigma}_{P_t}) \leq \rho$ , with the last case holding by the triangle inequality. Note that  $\rho \leq 2\varepsilon$  under our assumed sample complexity. We

now proceed as in the population-limit proof. In particular, we use the equivalent second-moment formulations of the estimators and regret stability, justified by Lemma 1.2.

For the MRO guarantee, set  $\hat{\mathbf{A}} := \text{MRO}_\rho[P_{s,n}]$  and bound

$$\begin{aligned}
R(\hat{\mathbf{A}}, P_t) &\leq \sup_{\Sigma \in \mathcal{B}_\rho(\Sigma_{P_{s,n}})} R(\hat{\mathbf{A}}, \Sigma) && (\mathbf{B}(\Sigma_{P_{s,n}}, \Sigma_{P_t}) \leq \rho) \\
&\leq \sup_{\Sigma \in \mathcal{B}_\rho(\Sigma_{P_{s,n}})} R(\mathbf{A}_P, \Sigma) && (\text{MRO optimality}) \\
&\leq \sup_{\Sigma \in \mathcal{B}_{2\rho}(\Sigma_P)} R(\mathbf{A}_P, \Sigma) && (\mathbf{B}(\Sigma_{P_{s,n}}, \Sigma_P) \leq \rho) \\
&\leq 4\kappa_P(2\rho)^2 && (\text{Lemma 3.1}) \\
&\leq 64\kappa_P\varepsilon^2, && (\rho \leq 2\varepsilon)
\end{aligned}$$

as claimed.

For SNR, set  $\hat{\mathbf{A}} := \text{SNR}_\rho[P_{s,n}]$ . Since risk and regret differ by a constant independent of the coefficient matrix, SNR optimality implies

$$R(\hat{\mathbf{A}}, P_{s,n}) + \rho^2 \|\hat{\mathbf{A}}\|_{\text{op}}^2 \leq R(\mathbf{A}_P, P_{s,n}) + \rho^2 \|\mathbf{A}_P\|_{\text{op}}^2 \leq 5\kappa_P \rho^2.$$

The last inequality follows from  $\mathbf{B}(\Sigma_{P_{s,n}}, \Sigma_P) \leq \rho$ , Lemma 3.1, and  $\|\mathbf{A}_P\|_{\text{op}}^2 \leq \kappa_P$  from Theorem 2.2. To convert this to a regret bound under  $P_t$ , we apply regret stability twice, using  $P$  as the reference law in both applications:

$$\begin{aligned}
\sqrt{R(\hat{\mathbf{A}}, P_t)} &\leq \sqrt{R(\hat{\mathbf{A}}, P)} + (2\sqrt{\kappa_P} + \|\hat{\mathbf{A}}\|_{\text{op}}) \mathbf{B}(\Sigma_{P_t}, \Sigma_P) \\
&\leq \sqrt{R(\hat{\mathbf{A}}, P_{s,n})} + (2\sqrt{\kappa_P} + \|\hat{\mathbf{A}}\|_{\text{op}}) \rho,
\end{aligned}$$

where the last inequality uses (33). By Cauchy–Schwarz and the SNR objective comparison above,

$$\sqrt{R(\hat{\mathbf{A}}, P_t)} \leq \sqrt{2} \sqrt{R(\hat{\mathbf{A}}, P_{s,n}) + \rho^2 \|\hat{\mathbf{A}}\|_{\text{op}}^2} + 2\sqrt{\kappa_P} \rho \leq 6\sqrt{\kappa_P} \rho \leq 12\sqrt{\kappa_P} \varepsilon.$$

Squaring gives the desired SNR bound.

Thus, both estimators incur target regret  $O(\kappa_P \varepsilon^2)$  under the specified event. The bounds hold uniformly over all admissible sample corruptions and population targets in the respective models, including randomized, sample-dependent choices. The event has probability at least  $1 - \delta$  under every admissible rule, since the reference sample has law  $P^{\otimes n}$ . Taking the supremum over  $\mathbb{M}_{\varepsilon,n}(P)$  therefore proves the first assertion.

**Sub-Gaussian guarantee.** Our assumption on the sample size allows us to apply Lemma H.1. Fix a realization on its event of probability at least  $1 - \delta$ . Then  $P_n \in \mathcal{P}^*$ ,  $\kappa_{P_n} \leq 4\kappa_P$ , and  $\Sigma_X^P \preceq \frac{4}{3} \Sigma_X^{P_n}$ . Moreover, we have

$$e_n := R(\mathbf{A}_{P_n}, P) \lesssim \frac{K^4 \text{tr}(\Sigma_E^P)(d + \log(1/\delta))}{n} \lesssim \varepsilon^2,$$

where the last inequality uses the theorem’s sample-size assumption.

We first establish a comparison that will be used in both perturbation models. For each  $\mathbf{A} \in \mathcal{A}$ , the regret formula and the squared triangle inequality imply that

$$\begin{aligned}
R(\mathbf{A}, P) &= \|(\mathbf{A} - \mathbf{A}_P)(\Sigma_X^P)^{1/2}\|_{\text{F}}^2 \\
&\leq 2\|(\mathbf{A} - \mathbf{A}_{P_n})(\Sigma_X^P)^{1/2}\|_{\text{F}}^2 + 2e_n \\
&\leq \frac{8}{3}\|(\mathbf{A} - \mathbf{A}_{P_n})(\Sigma_X^{P_n})^{1/2}\|_{\text{F}}^2 + 2e_n \\
&= \frac{8}{3}R(\mathbf{A}, P_n) + 2e_n,
\end{aligned} \tag{34}$$

where the last inequality uses  $\Sigma_X^P \preceq \frac{4}{3} \Sigma_X^{P_n}$ .

In the sample-corruption model, fix  $\hat{\mathbf{A}}_n \in \{\text{MRO}_\varepsilon, \text{SNR}_\varepsilon\}$  and any admissible observed source  $P_{s,n}$ . Set  $\hat{\mathbf{A}} := \hat{\mathbf{A}}_n[P_{s,n}]$ . Since  $W_2(P_{s,n}, P_n) \leq \varepsilon$  and  $P_n \in \mathcal{P}^*$ , Lemma 3.2 gives

$$R(\hat{\mathbf{A}}, P_n) \leq R_\varepsilon(\hat{\mathbf{A}}_n, P_n) \lesssim \kappa_{P_n} \varepsilon^2 \lesssim \kappa_P \varepsilon^2.$$

Applying (34), we obtain

$$R(\hat{\mathbf{A}}, P_t) = R(\hat{\mathbf{A}}, P) \leq \frac{8}{3} R(\hat{\mathbf{A}}, P_n) + 2e_n \lesssim \kappa_P \varepsilon^2,$$

using  $e_n \lesssim \varepsilon^2$  and  $\kappa_P \geq 1$ . This proves the desired target regret bound uniformly over admissible corruptions.

In the distribution-shift model,  $P_{s,n} = P_n$  and  $P_t = Q \in \mathcal{B}_\varepsilon(P)$ . Since  $Q$  need not belong to  $\mathcal{B}_\varepsilon(P_n)$ , we will first transfer regret from  $P_n$  to  $P$ , then from  $P$  to  $Q$ . For the second transfer, we also need to control the coefficient norm.

For  $\hat{\mathbf{A}} = \text{SNR}_\varepsilon[P_n]$ , comparison with the empirical risk minimizer  $\mathbf{A}_{P_n}$  gives

$$R(\hat{\mathbf{A}}, P_n) + \varepsilon^2 \|\hat{\mathbf{A}}\|_{\text{op}}^2 \leq R(\mathbf{A}_{P_n}, P_n) + \varepsilon^2 \|\mathbf{A}_{P_n}\|_{\text{op}}^2 = \varepsilon^2 \|\mathbf{A}_{P_n}\|_{\text{op}}^2 \leq \kappa_{P_n} \varepsilon^2.$$

Here the first inequality uses SNR optimality, and the last uses Theorem 2.2. Since both terms on the left are nonnegative, this bounds the empirical regret and coefficient norm separately.

For MRO, define the objective  $F_n(\mathbf{A}) := \sup_{Q' \in \mathcal{B}_\varepsilon(P_n)} R(\mathbf{A}, Q')$ . We first show that

$$F_n(\mathbf{A}) \geq \varepsilon^2(1 + \|\mathbf{A}\|_{\text{op}}^2) \quad \forall \mathbf{A} \in \mathcal{A}. \quad (35)$$

By the  $W_2$ -Bures equivalence (Lemma 1.2), we may apply Theorem 4.1 with observed second moment  $\Sigma_{P_n}$ . Define the square-root factor  $\mathbf{S}_n := \Sigma_{P_n}^{1/2}$  and let  $\mathbf{S}_{n,X} \in \mathbb{R}^{p \times d}$  be its feature row block. For every  $\eta > 1 + \|\mathbf{A}\|_{\text{op}}^2$ , the matrix inside the  $\text{tr}_p$  term in that theorem is  $\mathbf{S}_{n,X}^\top \mathbf{S}_{n,X}$  plus a positive semidefinite matrix. The  $\text{tr}_p$  term is therefore at least

$$\text{tr}_p(\mathbf{S}_{n,X}^\top \mathbf{S}_{n,X}) = \text{tr}(\mathbf{S}_{n,X}^\top \mathbf{S}_{n,X}) = \|\mathbf{S}_{n,X}\|_F^2,$$

where the first equality uses  $\text{rank}(\mathbf{S}_{n,X}^\top \mathbf{S}_{n,X}) \leq p$ . Consequently, the dual objective is at least  $\eta \varepsilon^2$ , and the reformulation gives

$$F_n(\mathbf{A}) \geq \inf_{\eta > 1 + \|\mathbf{A}\|_{\text{op}}^2} \eta \varepsilon^2 = \varepsilon^2(1 + \|\mathbf{A}\|_{\text{op}}^2).$$

This proves (35).

Now take  $\hat{\mathbf{A}} = \text{MRO}_\varepsilon[P_n]$ . Since  $P_n \in \mathcal{B}_\varepsilon(P_n)$ , MRO optimality gives

$$R(\hat{\mathbf{A}}, P_n) \leq F_n(\hat{\mathbf{A}}) \leq F_n(\mathbf{A}_{P_n}) \leq 4\kappa_{P_n} \varepsilon^2.$$

The last inequality follows from Lemma 3.1 with reference law  $P_n$  and coefficient  $\mathbf{A}_{P_n}$ . Combining this upper bound with (35) also gives

$$\varepsilon^2(1 + \|\hat{\mathbf{A}}\|_{\text{op}}^2) \leq F_n(\hat{\mathbf{A}}) \leq 4\kappa_{P_n} \varepsilon^2,$$

and hence  $\|\hat{\mathbf{A}}\|_{\text{op}}^2 \leq 4\kappa_{P_n}$ .

Thus, for either choice of  $\hat{\mathbf{A}}_n \in \{\text{MRO}_\varepsilon, \text{SNR}_\varepsilon\}$  and  $\hat{\mathbf{A}} := \hat{\mathbf{A}}_n[P_n]$ , we have

$$R(\hat{\mathbf{A}}, P_n) \lesssim \kappa_P \varepsilon^2, \quad \|\hat{\mathbf{A}}\|_{\text{op}}^2 \lesssim \kappa_P,$$

using  $\kappa_{P_n} \leq 4\kappa_P$ . The empirical-to-population comparison (34) thus yields

$$R(\hat{\mathbf{A}}, P) \leq \frac{8}{3} R(\hat{\mathbf{A}}, P_n) + 2e_n \lesssim \kappa_P \varepsilon^2.$$

Finally, for every  $Q \in \mathcal{B}_\varepsilon(P)$ , Lemma 3.1 gives

$$\begin{aligned} \sqrt{R(\hat{\mathbf{A}}, Q)} &\leq \sqrt{R(\hat{\mathbf{A}}, P)} + (2\sqrt{\kappa_P} + \|\hat{\mathbf{A}}\|_{\text{op}}) W_2(P, Q) \\ &\leq \sqrt{R(\hat{\mathbf{A}}, P)} + (2\sqrt{\kappa_P} + \|\hat{\mathbf{A}}\|_{\text{op}}) \varepsilon \\ &\lesssim \sqrt{\kappa_P} \varepsilon. \end{aligned}$$

Squaring proves that  $R(\hat{\mathbf{A}}, P_t) \lesssim \kappa_P \varepsilon^2$  in the distribution-shift model.

On the chosen sample event, both estimators satisfy the claimed bounds for every admissible sample corruption and every population target  $Q \in \mathcal{B}_\varepsilon(P)$  in the respective models. The event has probability at least  $1 - \delta$  under every admissible rule, since the reference sample has law  $P^{\otimes n}$ . Taking the supremum over  $\mathbb{M}_{\varepsilon,n}(P)$  therefore proves the second assertion.  $\square$

## I Exact Reductions for Population Benchmarks in Section 6

In this section, we derive closed forms and exact constant-size semidefinite program (SDP) reformulations for the population benchmarks in Section 6. Write  $P := P_t$  and  $\tilde{P} := P_s$  for the target and source laws, respectively, and let  $\rho > 0$  denote the supplied radius, which may differ from the corruption budget  $\varepsilon = 1$ .

Set  $q := m - 1$ . Recall that  $m = p = d/2 \geq 2$  and that our  $P$  is the law of  $(X, Y)$ , where

$$X \sim \mathcal{N}(0, \mathbf{I}_m), \quad Y = X + U, \quad U \sim \mathcal{N}(0, \text{diag}(0, \mathbf{I}_q)),$$

and  $X$  and  $U$  are independent. For  $\alpha, \beta \in \mathbb{R}$ , consider the diagonal coefficient matrix given by

$$\mathbf{A}(\alpha, \beta) := \text{diag}(\alpha, \beta \mathbf{I}_q).$$

With this notation, our source measure  $\tilde{P}$  is the law of  $(\mathbf{A}(c, 1)X, Y)$ , where  $c := m^{-1/4}$ . Since  $\mathbf{A}_P = \mathbf{I}_m$  and  $\Sigma_X^P = \mathbf{I}_m$ , the target regret of  $\mathbf{A}(\alpha, \beta)$  is given by

$$R(\mathbf{A}(\alpha, \beta), P) = (\alpha - 1)^2 + q(\beta - 1)^2. \quad (36)$$

Moreover, the source risk is

$$L(\mathbf{A}(\alpha, \beta), \tilde{P}) = (c\alpha - 1)^2 + q((\beta - 1)^2 + 1). \quad (37)$$

Finally, since  $\Sigma_X^{\tilde{P}} = \text{diag}(c^2, \mathbf{I}_q) \succ \mathbf{0}$ , the population ERM solution is unique and satisfies

$$\text{ERM}[\tilde{P}] = \mathbf{A}(c^{-1}, 1), \quad R(\text{ERM}[\tilde{P}], P) = (c^{-1} - 1)^2.$$

### I.1 Reduction to Scalar Optimization Using Rotational Symmetry

As in Lemma F.10, both  $P$  and  $\tilde{P}$  are invariant under the compact group of orthogonal transformations

$$\{\text{diag}(1, \mathbf{H}, 1, \mathbf{H}) : \mathbf{H} \in \mathbb{O}^q\}.$$

Its invariant coefficient matrices (in the sense of Lemma C.4) are exactly those of the form  $\mathbf{A}(\alpha, \beta)$  for some  $\alpha, \beta \in \mathbb{R}$ . Indeed, taking  $\mathbf{H} = -\mathbf{I}_q$  eliminates the off-diagonal blocks, and commuting with every  $\mathbf{H} \in \mathbb{O}^q$  forces the remaining  $q \times q$  block to be a scalar multiple of  $\mathbf{I}_q$ . Hence, by Lemma C.4, each objective

$$\begin{aligned} J_{\text{DRO}}(\mathbf{A}) &:= \sup_{Q \in \mathcal{B}_\rho(\tilde{P})} L(\mathbf{A}, Q), \\ J_{\text{MRO}}(\mathbf{A}) &:= \sup_{Q \in \mathcal{B}_\rho(\tilde{P})} R(\mathbf{A}, Q), \\ J_{\text{SNR}}(\mathbf{A}) &:= L(\mathbf{A}, \tilde{P}) + \rho^2 \|\mathbf{A}\|_{\text{op}}^2 \end{aligned}$$

admits a minimizer of the form  $\mathbf{A}(\alpha, \beta)$ .

### I.2 Constant-Size SDP Reformulations for DRO and MRO

Using this symmetry reduction, we now derive SDP reformulations for the DRO and MRO problems. Fix  $\mathbf{A} = \mathbf{A}(\alpha, \beta)$  for now and write  $\gamma := 1 + \max\{\alpha^2, \beta^2\}$ .

Our first goal is to derive a clean dual formulation for  $J_{\text{DRO}}(\mathbf{A})$  and  $J_{\text{MRO}}(\mathbf{A})$ . We then introduce epigraph variables and simplify the resulting convex programs using Schur complements and a variational formula for  $\text{tr}_p$ .

**Step 1: Initial dual problems.** A convenient square-root factor of the source second moment is

$$\tilde{\mathbf{S}} := \begin{pmatrix} \text{diag}(c, \mathbf{I}_q) & \mathbf{0}_{m \times m} \\ \mathbf{I}_m & \text{diag}(0, \mathbf{I}_q) \end{pmatrix}, \quad \tilde{\mathbf{S}}\tilde{\mathbf{S}}^\top = \Sigma_{\tilde{P}}.$$

Set  $\tilde{\mathbf{E}}_{\mathbf{A}} := \tilde{\mathbf{S}}_Y - \mathbf{A}\tilde{\mathbf{S}}_X$  and  $\mathbf{B}_{\mathbf{A}, \eta} := (\eta - 1)\mathbf{I}_m - \mathbf{A}\mathbf{A}^\top$ . Then, the  $W_2$ -Bures equivalence (Lemma 1.2) and the DRO and MRO dual formulas in (2) and Theorem 4.1 give

$$\begin{aligned} J_{\text{DRO}}(\mathbf{A}) &= \inf_{\eta > \gamma} \left\{ \eta \rho^2 + \text{tr}(\eta \tilde{\mathbf{E}}_{\mathbf{A}}^\top \mathbf{B}_{\mathbf{A}, \eta}^{-1} \tilde{\mathbf{E}}_{\mathbf{A}}) \right\}, \\ J_{\text{MRO}}(\mathbf{A}) &= \inf_{\eta > \gamma} \left\{ \eta(\rho^2 - \|\tilde{\mathbf{S}}_X\|_{\text{F}}^2) + \text{tr}_m(\eta \tilde{\mathbf{S}}_X^\top \tilde{\mathbf{S}}_X + \eta \tilde{\mathbf{E}}_{\mathbf{A}}^\top \mathbf{B}_{\mathbf{A}, \eta}^{-1} \tilde{\mathbf{E}}_{\mathbf{A}}) \right\}. \end{aligned}$$

Here, we used  $p = m$  and positive homogeneity of  $\text{tr}_m$ . To unify these formulas, define

$$\mathbf{K}_\eta := \eta \tilde{\mathbf{S}}^\top \tilde{\mathbf{S}} + \eta \tilde{\mathbf{E}}_\mathbf{A}^\top \mathbf{B}_{\mathbf{A},\eta}^{-1} \tilde{\mathbf{E}}_\mathbf{A}.$$

Substituting into the above gives

$$J_{\text{DRO}}(\mathbf{A}) = \inf_{\eta > \gamma} \left\{ \eta(\rho^2 - \|\tilde{\mathbf{S}}\|_{\text{F}}^2) + \text{tr}(\mathbf{K}_\eta) \right\}, \quad (38)$$

$$J_{\text{MRO}}(\mathbf{A}) = \inf_{\eta > \gamma} \left\{ \eta(\rho^2 - \|\tilde{\mathbf{S}}_X\|_{\text{F}}^2) + \text{tr}_m(\mathbf{K}_\eta - \eta \tilde{\mathbf{S}}_Y^\top \tilde{\mathbf{S}}_Y) \right\}. \quad (39)$$

The matrices here are structured, due to our choice of source law, leading to further simplifications.

First, writing  $\mathbf{M}_\mathbf{A} := (\mathbf{A} \quad -\mathbf{I}_m)$ , Lemma A.4 implies that

$$\mathbf{K}_\eta = \eta^2 \tilde{\mathbf{S}}^\top (\eta \mathbf{I}_d - \mathbf{M}_\mathbf{A}^\top \mathbf{M}_\mathbf{A})^{-1} \tilde{\mathbf{S}}.$$

Next, note that the source random variables can be written as

$$(\tilde{X}_1, Y_1) = (cG_1, G_1), \quad (\tilde{X}_i, Y_i) = (G_i, G_i + H_i), \quad i = 2, \dots, m,$$

with independent standard Gaussian latent variables  $G_i$  and  $H_i$ . The corresponding factors are

$$\mathbf{S}_1 := \begin{pmatrix} c \\ 1 \end{pmatrix}, \quad \mathbf{S}_b := \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix},$$

with the subscript “b” denoting the “bulk” of the coordinates. For a scalar coefficient  $a \in \mathbb{R}$  and a fixed block factor  $\mathbf{S} \in \mathbb{R}^{2 \times k}$ , define

$$\mathbf{M}(a) := (a \quad -1), \quad \mathbf{K}(a, \mathbf{S}, \eta) := \eta^2 \mathbf{S}^\top (\eta \mathbf{I}_2 - \mathbf{M}(a)^\top \mathbf{M}(a))^{-1} \mathbf{S}, \quad \eta > 1 + a^2.$$

Write  $k_1 := \mathbf{K}(\alpha, \mathbf{S}_1, \eta)$ , identifying the  $1 \times 1$  matrix with its entry, and  $\mathbf{K}_b := \mathbf{K}(\beta, \mathbf{S}_b, \eta)$ . In latent-coordinate order  $G_1, (G_2, H_2), \dots, (G_m, H_m), H_1$ , the matrices  $\mathbf{K}_\eta$  and  $\tilde{\mathbf{S}}_Y^\top \tilde{\mathbf{S}}_Y$  become

$$\text{diag}(k_1, \underbrace{\mathbf{K}_b, \dots, \mathbf{K}_b}_{q \text{ times}}, 0), \quad \text{diag}(1, \underbrace{1_2 1_2^\top, \dots, 1_2 1_2^\top}_{q \text{ times}}, 0), \quad 1_2 = (1, 1)^\top.$$

The final zero comes from the unused latent coordinate  $H_1$ . We also compute

$$\|\tilde{\mathbf{S}}\|_{\text{F}}^2 = c^2 + 1 + 3q, \quad \|\tilde{\mathbf{S}}_X\|_{\text{F}}^2 = c^2 + q.$$

Substituting these decompositions into (38) and (39), we arrive at

$$J_{\text{DRO}}(\mathbf{A}) = \inf_{\eta > \gamma} \left\{ \eta(\rho^2 - c^2 - 1 - 3q) + k_1 + q \text{tr}(\mathbf{K}_b) \right\}, \quad (40)$$

$$J_{\text{MRO}}(\mathbf{A}) = \inf_{\eta > \gamma} \left\{ \eta(\rho^2 - c^2 - q) + \text{tr}_m \left( \text{diag}(k_1 - \eta, \underbrace{\mathbf{K}_b - \eta 1_2 1_2^\top, \dots, \mathbf{K}_b - \eta 1_2 1_2^\top}_{q \text{ copies}}) \right) \right\}. \quad (41)$$

Here, we removed the final zero from the matrix inside  $\text{tr}_m$ . This is permissible because this matrix coincides up to a permutation of rows and columns with  $\mathbf{K}_\eta - \eta \tilde{\mathbf{S}}_Y^\top \tilde{\mathbf{S}}_Y$ . By the definition of  $\mathbf{K}_\eta$ , that matrix is greater in Loewner order than  $\eta \tilde{\mathbf{S}}_X^\top \tilde{\mathbf{S}}_X \succeq \mathbf{0}$  and is thus positive semidefinite. The remaining blocks have dimension  $2m - 1 \geq m$ , so removing the zero does not change the eigenvalue sum.

**Step 2: Introduce epigraph variables and simplify.** We now introduce epigraph variables  $t_1 \geq k_1$  and  $\mathbf{T}_b \succeq \mathbf{K}_b$ . To make these lower bound constraints affine in  $\mathbf{A}$  and  $\eta$ , we use that for  $\eta > 1 + a^2$ , a Schur complement (Lemma A.3) gives

$$\mathbf{T} \succeq \mathbf{K}(a, \mathbf{S}, \eta) \iff \begin{pmatrix} \eta \mathbf{I}_2 - \mathbf{M}(a)^\top \mathbf{M}(a) & \eta \mathbf{S} \\ \eta \mathbf{S}^\top & \mathbf{T} \end{pmatrix} \succeq \mathbf{0}.$$

The upper-left block is quadratic in  $a$ . A second Schur complement replaces this quadratic term by an additional scalar block, yielding the affine constraint  $\mathcal{Q}(a, \mathbf{S}, \eta, \mathbf{T}) \succeq \mathbf{0}$ , where

$$\mathcal{Q}(a, \mathbf{S}, \eta, \mathbf{T}) := \begin{pmatrix} \eta \mathbf{I}_2 & \mathbf{M}(a)^\top & \eta \mathbf{S} \\ \mathbf{M}(a) & 1 & \mathbf{0}_{1 \times k} \\ \eta \mathbf{S}^\top & \mathbf{0}_{k \times 1} & \mathbf{T} \end{pmatrix} \in \mathbb{S}^{k+3}.$$

In particular, for  $\eta > 1 + a^2$ ,

$$\mathcal{Q}(a, \mathbf{S}, \eta, \mathbf{T}) \succeq \mathbf{0} \iff \mathbf{T} \succeq \mathbf{K}(a, \mathbf{S}, \eta).$$

For general  $\eta \in \mathbb{R}$ , its Schur complement at the scalar block 1 also shows that

$$\mathcal{Q}(a, \mathbf{S}, \eta, \mathbf{T}) \succeq \mathbf{0} \implies \eta \geq 1 + a^2. \quad (42)$$

Define the two instances of this constraint by

$$\begin{aligned} \mathcal{Q}_1(\alpha, \eta, t_1) &:= \mathcal{Q}(\alpha, \mathbf{S}_1, \eta, t_1) \in \mathbb{S}^4, \\ \mathcal{Q}_b(\beta, \eta, \mathbf{T}_b) &:= \mathcal{Q}(\beta, \mathbf{S}_b, \eta, \mathbf{T}_b) \in \mathbb{S}^5. \end{aligned}$$

Retaining the constraint that  $\eta > \gamma$  in (40) and (41), we replace  $k_1, \mathbf{K}_b$  by the epigraph variables  $t_1, \mathbf{T}_b$  and minimize over them subject to  $\mathcal{Q}_1(\alpha, \eta, t_1) \succeq \mathbf{0}$  and  $\mathcal{Q}_b(\beta, \eta, \mathbf{T}_b) \succeq \mathbf{0}$ . Both traces are nondecreasing in Loewner order, and equality is feasible at  $t_1 = k_1, \mathbf{T}_b = \mathbf{K}_b$ . Thus this substitution preserves the values of (40) and (41). For MRO, we additionally reparameterize using variables

$$z_1 := t_1 - \eta, \quad \mathbf{Z}_b := \mathbf{T}_b - \eta \mathbf{1}_2 \mathbf{1}_2^\top, \quad \mathbf{Z} := \text{diag}(z_1, \underbrace{\mathbf{Z}_b, \dots, \mathbf{Z}_b}_{q \text{ times}}).$$

By Lemma A.9, we have

$$\text{tr}_m(\mathbf{Z}) = \min_{\substack{\tau \in \mathbb{R}, y_1 \geq 0, y_1 \geq z_1 - \tau \\ \mathbf{Y}_b \in \mathbb{S}^2, \mathbf{Y}_b \succeq \mathbf{0}, \mathbf{Y}_b \succeq \mathbf{Z}_b - \tau \mathbf{I}_2}} m\tau + y_1 + q \text{tr}(\mathbf{Y}_b). \quad (43)$$

Substituting these representations into (40) and (41) and taking the infimum over  $(\alpha, \beta)$  results in the programs below with the additional constraint that  $\eta > \gamma$ . We show next that this constraint can be dropped and that the resulting minima are attained.

**Proposition I.1** (Constant-size SDP formulations). *A population DRO optimizer is obtained by extracting coefficient variables  $(\alpha, \beta)$  from any solution to*

$$\begin{aligned} \min_{\alpha, \beta, \eta, t_1, \mathbf{T}_b} \quad & t_1 + q \text{tr}(\mathbf{T}_b) + \eta(\rho^2 - c^2 - 1 - 3q) \\ \text{s.t.} \quad & \mathcal{Q}_1(\alpha, \eta, t_1) \succeq \mathbf{0}, \\ & \mathcal{Q}_b(\beta, \eta, \mathbf{T}_b) \succeq \mathbf{0}, \\ & \eta \geq 0, \quad \mathbf{T}_b \in \mathbb{S}^2. \end{aligned} \quad (44)$$

Likewise, a population MRO optimizer is obtained by taking  $(\alpha, \beta)$  from any solution to

$$\begin{aligned} \min_{\alpha, \beta, \eta, t_1, \mathbf{T}_b, \tau, y_1, \mathbf{Y}_b} \quad & \eta(\rho^2 - c^2 - q) + m\tau + y_1 + q \text{tr}(\mathbf{Y}_b) \\ \text{s.t.} \quad & \mathcal{Q}_1(\alpha, \eta, t_1) \succeq \mathbf{0}, \\ & \mathcal{Q}_b(\beta, \eta, \mathbf{T}_b) \succeq \mathbf{0}, \\ & \eta \geq 0, \\ & y_1 \geq 0, \quad y_1 \geq z_1 - \tau, \\ & \mathbf{Y}_b \succeq \mathbf{0}, \quad \mathbf{Y}_b \succeq \mathbf{Z}_b - \tau \mathbf{I}_2, \\ & \mathbf{T}_b, \mathbf{Y}_b \in \mathbb{S}^2. \end{aligned} \quad (45)$$

Both minima are attained, and the regret of the returned coefficient pairs at the target law  $P$  can be computed via (36).

*Proof.* By (42), both SDPs impose  $\eta \geq \gamma = \gamma(\alpha, \beta)$ . The preceding argument establishes the desired optimal value identities with the additional restriction  $\eta > \gamma$ . It remains to remove this strict constraint and prove that minimizers are attained.

For fixed  $(\alpha, \beta)$ , let  $\eta > \gamma$ . By taking the remaining variables sufficiently large, we can obtain a feasible point with this strictly feasible value of  $\eta$ . Moreover, the feasible set is convex, and the objective is continuous. Thus, the presence of the strict constraint does not affect the infimum.

For attainment, fix an optimal coefficient pair  $(\alpha, \beta)$ , which exists by Section I.1. Choose a feasible objective value  $C$  for either SDP in the statement, and restrict to feasible points with objective at most  $C$ . This does not change the infimum. We show that the restricted feasible set is compact.

First,  $\eta \geq \gamma \geq 1$ . Extracting principal submatrices from the  $\mathcal{Q}_1, \mathcal{Q}_b \succeq \mathbf{0}$  constraints gives

$$\begin{pmatrix} \eta \mathbf{I}_2 & \eta \mathbf{S}_1 \\ \eta \mathbf{S}_1^\top & t_1 \end{pmatrix} \succeq \mathbf{0}, \quad \begin{pmatrix} \eta \mathbf{I}_2 & \eta \mathbf{S}_b \\ \eta \mathbf{S}_b^\top & \mathbf{T}_b \end{pmatrix} \succeq \mathbf{0}.$$

Since  $\eta > 0$ , Lemma A.3 implies that

$$t_1 \geq \eta \mathbf{S}_1^\top \mathbf{S}_1 = \eta(c^2 + 1), \quad \mathbf{T}_b \succeq \eta \mathbf{S}_b^\top \mathbf{S}_b = \eta \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}.$$

For MRO, we therefore have

$$z_1 = t_1 - \eta \geq \eta c^2, \quad \mathbf{Z}_b = \mathbf{T}_b - \eta \mathbf{1}_2 \mathbf{1}_2^\top \succeq \eta \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix},$$

implying that the combined matrix  $\mathbf{Z} = \text{diag}(z_1, \mathbf{Z}_b, \dots, \mathbf{Z}_b)$ , with  $q$  copies of  $\mathbf{Z}_b$ , satisfies  $\mathbf{Z} \succeq \mathbf{0}$  and  $\text{tr}_m(\mathbf{Z}) \geq \eta(c^2 + q)$ . The SDP objective formulas and (43) now show that both objectives are at least  $\eta\rho^2$ . By our choice of  $C$ , the multiplier  $\eta$  must therefore lie in the bounded interval  $[\gamma, C/\rho^2]$ .

For DRO, boundedness of the objective and multiplier bounds  $t_1 + q \text{tr}(\mathbf{T}_b)$ . Since  $t_1 \geq 0$ ,  $\mathbf{T}_b \succeq \mathbf{0}$ , and  $q \geq 1$ , both  $t_1$  and  $\mathbf{T}_b$  must be bounded.

For MRO, boundedness of the objective and multiplier bounds  $m\tau + y_1 + q \text{tr}(\mathbf{Y}_b)$  from above. In the other direction, the constraints  $y_1 \geq 0$  and  $\mathbf{Y}_b \succeq \mathbf{0}$  imply that

$$m\tau + y_1 + q \text{tr}(\mathbf{Y}_b) \geq m\tau.$$

Further, the SDP constraints for  $y_1$  and  $\mathbf{Y}_b$  and the early observation that  $\mathbf{Z} \succeq \mathbf{0}$  imply that

$$\begin{aligned} m\tau + y_1 + q \text{tr}(\mathbf{Y}_b) &\geq m\tau + z_1 - \tau + q \text{tr}(\mathbf{Z}_b - \tau \mathbf{I}_2) \\ &\geq \text{tr}(\mathbf{Z}) - (1 + 2q - m)\tau \\ &\geq -q\tau. \end{aligned}$$

Since  $m, q > 0$ , these bounds control  $\tau$  in both directions. They then bound  $y_1$  and  $\text{tr}(\mathbf{Y}_b)$  separately, since  $y_1 \geq 0$  and  $\mathbf{Y}_b \succeq \mathbf{0}$ . Moreover,

$$0 \leq z_1 \leq y_1 + \tau, \quad \mathbf{0} \preceq \mathbf{Z}_b \preceq \mathbf{Y}_b + \tau \mathbf{I}_2,$$

so the shifted blocks  $z_1$  and  $\mathbf{Z}_b$  are bounded. The identities  $t_1 = z_1 + \eta$  and  $\mathbf{T}_b = \mathbf{Z}_b + \eta \mathbf{1}_2 \mathbf{1}_2^\top$  then bound the remaining variables  $t_1$  and  $\mathbf{T}_b$ .

Thus, for either SDP, the restricted feasible set is bounded. It is also nonempty and closed, because the constraints are closed and the objective is continuous. It is thus compact, and so the continuous objective attains its minimum on this set.  $\square$

### I.3 Closed form for SNR

For SNR, we obtain an explicit solution and closed form for the regret under  $P$ .

**Lemma I.2** (Closed-form SNR solution). *Let  $\lambda := \rho^2$ . Then the unique minimizer of  $\mathbf{A} \mapsto L(\mathbf{A}, \tilde{P}) + \lambda \|\mathbf{A}\|_{\text{op}}^2$  is  $\mathbf{A}(\alpha_\lambda, \beta_\lambda)$ , where*

$$(\alpha_\lambda, \beta_\lambda) = \begin{cases} \left( \frac{c}{c^2 + \lambda}, 1 \right), & 0 \leq \lambda \leq c(1 - c), \\ (t_\lambda, t_\lambda), \quad t_\lambda := \frac{q + c}{q + c^2 + \lambda}, & \lambda \geq c(1 - c). \end{cases} \quad (46)$$

Consequently, its exact clean regret is

$$R(\text{SNR}_\rho[\tilde{P}], P) = \begin{cases} \left( \frac{c}{c^2 + \lambda} - 1 \right)^2, & 0 \leq \lambda \leq c(1 - c), \\ m \left( \frac{q + c}{q + c^2 + \lambda} - 1 \right)^2, & \lambda \geq c(1 - c). \end{cases} \quad (47)$$

The two expressions agree at  $\lambda = c(1 - c)$ .

*Proof.* By our discussion in Section I.1, it suffices to minimize over  $\mathbf{A}(\alpha, \beta)$ . Dropping the additive constant  $q$  from (37), the reduced objective is

$$F_\lambda(\alpha, \beta) = (c\alpha - 1)^2 + q(\beta - 1)^2 + \lambda \max\{\alpha^2, \beta^2\}. \quad (48)$$

Recall that  $c = m^{-1/4} \in (0, 1]$ . Because  $\tilde{\Sigma}_X^P = \text{diag}(c^2, \mathbf{I}_q) \succ 0$ , the risk is strongly convex in the full matrix  $\mathbf{A}$ . Hence the regularized objective has a unique minimizer.

Introduce an epigraph variable  $t \geq 0$  with constraints  $|\alpha| \leq t$  and  $|\beta| \leq t$ . For fixed  $t$ , the quadratic part of (48) is minimized by projecting its unconstrained minimizers  $\alpha = c^{-1}$  and  $\beta = 1$  onto  $[-t, t]$ . Since both are positive, the constrained minimizers are

$$\alpha(t) := \min\{t, c^{-1}\}, \quad \beta(t) := \min\{t, 1\}.$$

Thus the epigraph problem reduces to minimizing

$$\phi_\lambda(t) := \lambda t^2 + (1 - ct)_+^2 + q(1 - t)_+^2, \quad t \geq 0.$$

Here  $\lambda = \rho^2 > 0$  and  $0 < c < 1$ , so  $\phi_\lambda$  is strictly convex. For  $t > c^{-1}$  it equals  $\lambda t^2$  and increases, so its minimizer lies in  $[0, c^{-1}]$ . Moreover,

$$\frac{1}{2}\phi'_\lambda(t) = \lambda t - c(1 - ct)_+ - q(1 - t)_+, \quad \frac{1}{2}\phi'_\lambda(1) = \lambda - c(1 - c).$$

If  $0 < \lambda \leq c(1 - c)$ , the minimizer lies in  $[1, c^{-1}]$ , where the derivative vanishes at

$$t = \frac{c}{c^2 + \lambda}.$$

The corresponding coefficients are  $\alpha = t$  and  $\beta = 1$ . If  $\lambda \geq c(1 - c)$ , the minimizer lies in  $[0, 1]$ , where the derivative vanishes at

$$t = \frac{q + c}{q + c^2 + \lambda}.$$

In this case  $\alpha = \beta = t$ . At  $\lambda = c(1 - c)$ , both expressions give  $t = 1$ . This proves (46). Substitution into (36) proves (47).  $\square$

## J Full Experimental Details and Additional Experiments

The synthetic experiments for Fig. 1 in Section 6 and Section J.1 below were run in 8 hours on an m8g.1 large Amazon EC2 instance with 2 vCPUs and 8 GB of RAM, while the experiment for Section J.2 was run in less than one hour on an M1 MacBook Air with an 8-core CPU and 16 GB RAM. Code for the experiments and plots is available at <https://github.com/sbnietert/regret-optimal-regression>.

### J.1 Additional Synthetic Experiment with Varied Residual Rank

For **Setting C**, we modify Setting A by fixing  $p = m = 200$  ( $d = 400$ ),  $n = 10m^2$ , and  $\rho = \varepsilon = 1$ , and varying the residual effective rank over 10 equally spaced values  $r \in [1, m - 1]$ . Specifically, for each  $r$ , the target law  $P_r$  is the law of  $(X, Y)$ , where

$$X \sim \mathcal{N}(0, \mathbf{I}_m), \quad Y = X + U_r, \quad U_r \sim \mathcal{N}\left(0, \text{diag}\left(0, 1, \frac{r-1}{m-2} \mathbf{I}_{m-2}\right)\right),$$

and  $X$  and  $U_r$  are independent. Then  $\mathbf{A}_{P_r} = \Sigma_X^{P_r} = \mathbf{I}_m$ , and the residual second moment satisfies

$$\text{tr}(\Sigma_E^{P_r}) = r, \quad \|\Sigma_E^{P_r}\|_{\text{op}} = 1, \quad r_{\text{eff}}(\Sigma_E^{P_r}) = r.$$

In particular, each  $\kappa_{P_r} = 3$ . The endpoint  $r = m - 1$  recovers the target law of Setting A at  $m = 200$ .

For each  $r$ , we draw  $n$  i.i.d. target observations and apply the same corruption as in Setting A. We fit each estimator  $\hat{\mathbf{A}} \in \{\text{ERM}, \text{DRO}_\rho, \text{MRO}_\rho, \text{SNR}_\rho\}$  to the resulting empirical source law  $P_{s,n}$ , using the implementations described in Section 6. In Fig. 2, we plot the target population regret  $R(\hat{\mathbf{A}}[P_{s,n}], P_r) = \|\hat{\mathbf{A}}[P_{s,n}] - \mathbf{I}_m\|_{\text{F}}^2$  against  $r$ . Curves report averages over 10 independent trials at each value of  $r$ , with error bars showing the minimum and maximum over trials (again imperceptible due to low variance).




Figure 2: Target population regret  $R(\hat{\mathbf{A}}[P_{s,n}], P_r)$  for Setting C, as a function of the target residual effective rank  $r = r_{\text{eff}}(\Sigma_E^{P_r})$ .

## J.2 Network-Traffic Experiment with CESNET Dataset

Finally, we test our estimators on a real-world network-traffic dataset, albeit with synthetically introduced Wasserstein perturbations. We employ the CESNET-TimeSeries24 dataset (Koumar et al., 2025), which contains network-traffic time series data from the CESNET3 network connecting public and research institutions in the Czech Republic. Specifically, it records hourly, institution-level counts of packets and bytes transmitted over the network between October 2023 and July 2024, and we attempt to predict byte totals from packet counts. We use 268 institutions with nonconstant counts and finite, nonnegative measurements for at least 99% of the 6,717 observed hourly timestamps, and intersect their valid timestamps, obtaining  $n = 6,228$  paired observations. For each response dimension  $m \in \{8, 16, 24, 32, 48, 64, 96, 128, 192, 256, 268\}$ , we take our target response vector to be the first  $m$  byte-total channels and its corresponding feature vector to be the first  $q = m/2$  matching packet-count channels plus an appended intercept term of 1, so  $p = q + 1$ . We center and normalize each physical channel according to its mean and standard deviation under the full empirical law of the raw observations, keeping these transformations and the observation set fixed throughout. Denote the resulting standardized target law by  $P_m \in \mathcal{P}_2(\mathbb{R}^{p+m})$ , treated as a population measure to isolate the complexity of robust regression from that of handling sampling error.

This setup was selected so that some of the feature and response coordinates are highly correlated, while others are nearly independent, yielding a range of effective ranks. Coordinate-wise normalization is a relatively common preprocessing procedure, and it helps to keep the condition number of the target law from growing too quickly with  $m$ . In particular, the effective rank of  $\Sigma_E^{P_m}$  grows roughly linearly with  $m$ , while the target condition number  $\kappa_{P_m}$  grows more slowly, as shown in Fig. 4.

To obtain the corresponding corrupted source law  $\tilde{P}_m$ , we apply the frozen transformation above to the same  $n$  observations, but with the first packet-count channel first multiplied by a gain of 0.5 (all other channels are unchanged). By coupling observations according to their timestamps, one can directly upper bound  $W_2(P_m, \tilde{P}_m) \leq \varepsilon := 1.08679$ , independent of  $m$ . We then fit each estimator  $\hat{\mathbf{A}} \in \{\text{ERM}, \text{DRO}_\varepsilon, \text{MRO}_\varepsilon, \text{SNR}_\varepsilon\}$  to  $\tilde{P}_m$ , using the implementations described in Section 6. In Fig. 3, we plot the target regret  $R(\hat{\mathbf{A}}[\tilde{P}_m], P_m)$  against  $m$ . We observe qualitative similarities to the separation observed in our synthetic experiments, with MRO and SNR outperforming DRO and ERM as the dimension grows (although the gap is less pronounced). We emphasize that ample room remains for more refined empirical study, e.g., to identify the most realistic modes of data corruption for specific applications and to compare the performance of these estimators beyond the worst case.




Figure 3: Target regret  $R(\hat{\mathbf{A}}[\tilde{P}_m], P_m)$  for our CESNET experiment, as a function of the response dimension  $m$ .




Figure 4: Effective rank  $\text{tr}(\Sigma_E^{P_m}) / \|\Sigma_E^{P_m}\|_{\text{op}}$  and condition number  $\kappa_{P_m}$  of the target law  $P_m$  in our CESNET experiment, as a function of the response dimension  $m$ .