

# Bridging Bayesian and Minimax Mean Square Error Estimation via Wasserstein Distributionally Robust Optimization

VIET ANH NGUYEN, SOROOSH SHAFIEEZADEH-ABADEH,  
DANIEL KUHN, AND PEYMAN MOHAJERIN ESFAHANI

ABSTRACT. We introduce a distributionally robust minimum mean square error estimation model with a Wasserstein ambiguity set to recover an unknown signal from a noisy observation. The proposed model can be viewed as a zero-sum game between a statistician choosing an estimator—that is, a measurable function of the observation—and a fictitious adversary choosing a prior—that is, a pair of signal and noise distributions ranging over independent Wasserstein balls—with the goal to minimize and maximize the expected squared estimation error, respectively. We show that if the Wasserstein balls are centered at normal distributions, then the zero-sum game admits a Nash equilibrium, where the players’ optimal strategies are given by an *affine* estimator and a *normal* prior, respectively. We further prove that this Nash equilibrium can be computed by solving a tractable convex program. Finally, we develop a Frank-Wolfe algorithm that can solve this convex program orders of magnitude faster than state-of-the-art general purpose solvers. We show that this algorithm enjoys a linear convergence rate and that its direction-finding subproblems can be solved in quasi-closed form.

## 1. INTRODUCTION

Consider the problem of estimating an unknown parameter  $x \in \mathbb{R}^n$  based on a linear measurement  $y \in \mathbb{R}^m$  corrupted by additive noise  $w \in \mathbb{R}^m$ . This setup is formalized through the linear measurement model

$$(1.1) \quad y = Hx + w,$$

where the observation matrix  $H \in \mathbb{R}^{m \times n}$  is assumed to be known. We further assume that the distribution  $\mathbb{P}_w$  of  $w$  has finite second moments and that  $w$  is independent of  $x$ . Thus, the conditional distribution  $\mathbb{P}_{y|x}$  of  $y$  given  $x$  is obtained by shifting  $\mathbb{P}_w$  by  $Hx$ . Note that the linear measurement model is fundamental for numerous applications in engineering (*e.g.*, linear systems theory [33, 55]), econometrics (*e.g.*, linear regression [67, 69], time series analysis [12, 35]), machine learning and signal processing (*e.g.*, Kalman filtering [43, 52, 57]) or information theory (*e.g.*, multiple-input multiple-output systems [14, 50]) etc.

An estimator of  $x$  given  $y$  is a measurable function  $\psi : \mathbb{R}^m \rightarrow \mathbb{R}^n$  that grows at most linearly. Thus, there exists  $C > 0$  such that  $|\psi(y)| \leq C(1 + \|y\|)$  for all  $y \in \mathbb{R}^m$ . The function value  $\psi(y)$  is the prediction of  $x$  based on the measurement  $y$  under the estimator  $\psi$ . In the following we denote the family of all estimators by  $\mathcal{F}$ . The quality of an estimator is measured by a risk function  $R : \mathcal{F} \times \mathbb{R}^n \rightarrow \mathbb{R}$ , which quantifies the mismatch between the parameter  $x$  and its prediction  $\psi(y)$ . A popular risk function is the mean square error (MSE)

$$R(\psi, x) = \mathbb{E}_{\mathbb{P}_{y|x}} [\|x - \psi(y)\|^2],$$

---

*Date:* November 12, 2019.

The authors are with the Department of Management Science and Engineering, Stanford University ([viet-anh.nguyen@stanford.edu](mailto:viet-anh.nguyen@stanford.edu)), the Risk Analytics and Optimization Chair, EPFL, Switzerland ([soroosh.shafiee](mailto:soroosh.shafiee), [daniel.kuhn@epfl.ch](mailto:daniel.kuhn@epfl.ch)) and the Delft Center for Systems and Control, Delft University of Technology, The Netherlands ([P.MohajerinEsfahani@tudelft.nl](mailto:P.MohajerinEsfahani@tudelft.nl)).

which defines the estimation error as the expected squared Euclidean distance between  $\psi(y)$  and  $x$ . If  $x$  was known, then  $R(\psi, x)$  could be minimized directly, and the constant estimator  $\psi^*(y) \equiv x$  would be optimal. In practice, however,  $x$  is unobservable. Otherwise there would be no need to solve an estimation problem in the first place. With  $x$  unknown, it is impossible to minimize the MSE directly. The statistics literature proposes two complementary workarounds for this problem: the Bayesian approach and the minimax approach.

The Bayesian statistician treats  $x$  as a random vector governed by a *prior* distribution  $\mathbb{P}_x$  that captures her beliefs about  $x$  before seeing  $y$  [43, § 11.4] and solves the minimum MSE (MMSE) estimation problem

$$(1.2) \quad \underset{\psi \in \mathcal{F}}{\text{minimize}} \quad \mathbb{E}_{\mathbb{P}_x} [R(\psi, x)].$$

If the distribution  $\mathbb{P}_x$  of  $x$  has finite second moments, then (1.2) is solvable. In this case, the optimal estimator, which is usually termed the Bayesian MMSE estimator, is of the form  $\psi_{\mathcal{B}}^*(y) = \mathbb{E}_{\mathbb{P}_{x|y}}[x]$ , where the conditional distribution  $\mathbb{P}_{x|y}$  of  $x$  given  $y$  is obtained from  $\mathbb{P}_x$  and  $\mathbb{P}_{y|x}$  via Bayes' theorem. However, the Bayesian MMSE estimator suffers from two conceptual shortcomings. First,  $\psi_{\mathcal{B}}^*$  is highly sensitive to the prior distribution  $\mathbb{P}_x$ , which is troubling if the statistician has little confidence in her beliefs. Second, computing  $\psi_{\mathcal{B}}^*$  requires precise knowledge of the noise distribution  $\mathbb{P}_w$ , which is typically unobservable and thus uncertain at least to some extent. Moreover,  $\psi_{\mathcal{B}}^*$  may generically have a complicated functional form, and evaluating  $\psi_{\mathcal{B}}^*(y)$  to high precision for a particular measurement  $y$  (*e.g.*, via Monte Carlo simulation) may be computationally challenging if the dimension of  $x$  is high.

These shortcomings are mitigated if we restrict the space  $\mathcal{F}$  of all measurable estimators in (1.2) to the space

$$(1.3) \quad \mathcal{A} = \{ \psi \in \mathcal{F} : \exists A \in \mathbb{R}^{n \times m}, b \in \mathbb{R}^n \text{ with } \psi(y) = Ay + b \ \forall y \in \mathbb{R}^m \}$$

of all *affine* estimators. In this case the distributions  $\mathbb{P}_x$  and  $\mathbb{P}_w$  need not be fully known. Instead, in order to evaluate the optimal affine estimator  $\psi_{\mathcal{A}}^*(y) = A^*y + b^*$ , it is sufficient to know the mean vectors  $\mu_x$  and  $\mu_w$  as well as the covariance matrices  $\Sigma_x$  and  $\Sigma_w$  of the distributions  $\mathbb{P}_x$  and  $\mathbb{P}_w$ , respectively. If  $H\Sigma_x H^\top + \Sigma_w \succ 0$ , which is the case if the noise covariance matrix has full rank, then the coefficients of the best affine estimator can be computed in closed form. Using (1.1) together with the independence of  $x$  and  $w$  one can show that

$$(1.4) \quad A^* = \Sigma_x H^\top (H\Sigma_x H^\top + \Sigma_w)^{-1} \quad \text{and} \quad b^* = \mu_x - A^*(H\mu_x + \mu_w).$$

If the random vector  $(x, y)$  follows a normal distribution, then the best affine estimator is also optimal among all measurable estimators. In general, however, we do not know how much optimality is sacrificed by restricting attention to affine estimators. Moreover, the uncertainty about  $\mathbb{P}_x$  and  $\mathbb{P}_w$  transpires through to their first- and second-order moments. As the coefficients (1.4) tend to be highly sensitive to these moments, their uncertainty remains worrying.

The minimax approach models the statistician's prior knowledge concerning  $x$  via a convex closed uncertainty set  $\mathbb{X} \subseteq \mathbb{R}^n$  as commonly used in robust optimization. The minimax MSE estimation problem is then formulated as a zero-sum game between the statistician, who selects the estimator  $\psi \in \mathcal{F}$  with the goal to minimize the MSE, and nature, who chooses the parameter value  $x \in \mathbb{X}$  with the goal to maximize the MSE.

$$(1.5a) \quad \underset{\psi \in \mathcal{F}}{\text{minimize}} \quad \max_{x \in \mathbb{X}} R(\psi, x)$$

By construction, any minimizer  $\psi_{\mathcal{M}}^*$  of (1.5a) incurs the smallest possible estimation error under the worst parameter realization within the uncertainty set  $\mathbb{X}$ . For this reason  $\psi_{\mathcal{M}}^*$  is called a minimax estimator. Note that the MSE  $R(\psi, x)$  generically displays a complicated non-concave dependence on  $x$  for any fixed  $\psi$ , which implies that nature's inner maximization problem in (1.5a) is usually non-convex. Thus, we should not expect the zero-sum game (1.5a) between the statistician and nature to admit a Nash equilibrium. However, the

inner maximization problem can be convexified by allowing nature to play mixed (randomized) strategies, that is, by reformulating (1.5a) as the (equivalent) convex-concave saddle point problem

$$(1.5b) \quad \underset{\psi \in \mathcal{F}}{\text{minimize}} \quad \max_{\mathbb{Q}_x \in \mathcal{M}(\mathbb{X})} \mathbb{E}_{\mathbb{Q}_x} [R(\psi, x)],$$

where  $\mathcal{M}(\mathbb{X})$  stands for the family of all distributions supported on  $\mathbb{X}$  with finite second-order moments. As  $\mathbb{E}_{\mathbb{Q}_x}[R(\psi, x)]$  is convex in  $\psi$  for any fixed  $\mathbb{Q}_x$  and concave (linear) in  $\mathbb{Q}_x$  for any fixed  $\psi$ , while  $\mathcal{F}$  and  $\mathcal{M}(\mathbb{X})$  are both convex sets, the zero-sum game (1.5b) admits a Nash equilibrium  $(\psi_{\mathcal{M}}^*, \mathbb{Q}_x^*)$  under mild technical conditions. Note that  $\psi_{\mathcal{M}}^*$  is again a minimax estimator. Moreover,  $\psi_{\mathcal{M}}^*$  is the statistician's best response to nature's choice  $\mathbb{Q}_x^*$  and vice versa. Using the terminology introduced above, this means that  $\psi_{\mathcal{M}}^*$  is the Bayesian MMSE estimator corresponding to the prior  $\mathbb{Q}_x^*$ . For this reason,  $\mathbb{Q}_x^*$  is usually referred to as the *least favorable prior*. Even though the minimax approach exonerates the statistician from narrowing down her beliefs to a single prior distribution  $\mathbb{Q}_x$ , it still requires precise information about  $\mathbb{P}_w$ , which may not be available in practice. On the other hand, as it robustifies the estimator against *any* distribution on  $\mathbb{X}$ , the minimax approach is often regarded as overly pessimistic. Moreover, as in the case of the Bayesian MMSE estimation problem,  $\psi_{\mathcal{M}}^*$  may generically have a complicated functional form, and evaluating  $\psi_{\mathcal{M}}^*(y)$  to high precision may be computationally challenging if the dimension of  $x$  is high. A simple remedy to mitigate these computational challenges would be to restrict  $\mathcal{F}$  to the family  $\mathcal{A}$  of affine estimators. The loss of optimality incurred by this approximation for different choices of  $\mathbb{X}$  is discussed in [41, § 4] and the references therein.

In this paper we bridge the Bayesian and the minimax approaches by leveraging tools from distributionally robust optimization. Specifically, we study distributionally robust estimation problems of the form

$$(1.6) \quad \underset{\psi \in \mathcal{F}}{\text{minimize}} \quad \max_{\mathbb{Q}_x \in \mathcal{Q}_x} \mathbb{E}_{\mathbb{Q}_x} [R(\psi, x)],$$

where  $\mathcal{Q}_x \subseteq \mathcal{M}(\mathbb{R}^n)$  is an *ambiguity set* of multiple (possibly infinitely many) plausible prior distributions of  $x$ . Note that if the ambiguity set collapses to the singleton  $\mathcal{Q}_x = \{\mathbb{P}_x\}$  for some  $\mathbb{P}_x \in \mathcal{M}(\mathbb{R}^n)$ , then the distributionally robust estimation problem (1.6) reduces to the Bayesian MMSE estimation problem (1.2). Similarly, under the ambiguity set  $\mathcal{Q}_x = \mathcal{M}(\mathbb{X})$  for some convex closed uncertainty set  $\mathbb{X} \subseteq \mathbb{R}^n$ , problem (1.6) reduces to the minimax mean square error estimation problem (1.5b). By providing considerable freedom in tailoring the ambiguity set  $\mathcal{Q}_x$ , the distributionally robust approach thus allows the statistician to reconcile the specificity of the Bayesian approach with the conservativeness of the minimax approach.

The estimation model (1.6) still relies on the premise that the noise distribution  $\mathbb{P}_w$  is precisely known, and this assumption is not tenable in practice. However, nothing prevents us from further robustifying (1.6) against uncertainty in  $\mathbb{P}_w$ . To this end, we define  $\mathcal{M}(\mathbb{R}^{n+m})$  as the family of all joint distributions of  $x$  and  $w$  with finite second-order moments. Moreover, we define the *average risk*  $\mathcal{R} : \mathcal{F} \times \mathcal{M}(\mathbb{R}^{n+m}) \rightarrow \mathbb{R}$  through

$$\mathcal{R}(\psi, \mathbb{P}) = \mathbb{E}_{\mathbb{P}} [\|x - \psi(Hx + w)\|^2].$$

If  $\mathbb{P} = \mathbb{P}_x \times \mathbb{P}_w$  for some marginal distributions  $\mathbb{P}_x \in \mathcal{M}(\mathbb{R}^n)$  and  $\mathbb{P}_w \in \mathcal{M}(\mathbb{R}^m)$ , which implies that  $x$  and  $w$  are independent under  $\mathbb{P}$ , and if  $\mathbb{P}_{y|x}$  is defined as  $\mathbb{P}_w$  shifted by  $Hx$ , then  $\mathcal{R}(\psi, \mathbb{P}) = \mathbb{E}_{\mathbb{P}_x} [R(\psi, x)]$ . Thus, the average risk  $\mathcal{R}(\psi, \mathbb{P})$  corresponds indeed to the risk  $R(\psi, x)$  averaged under the marginal distribution  $\mathbb{P}_x$ . In the remainder of this paper we will study generalized distributionally robust estimation problems of the form

$$(1.7) \quad \underset{\psi \in \mathcal{F}}{\text{minimize}} \quad \sup_{\mathbb{Q} \in \mathbb{B}(\widehat{\mathbb{P}})} \mathcal{R}(\psi, \mathbb{Q}),$$

where the ambiguity set  $\mathbb{B}(\widehat{\mathbb{P}}) \subseteq \mathcal{M}(\mathbb{R}^{n+m})$  captures distributional uncertainty in both  $\mathbb{P}_x$  and  $\mathbb{P}_w$ . Specifically, we will model  $\mathbb{B}(\widehat{\mathbb{P}})$  as a set of factorizable distributions  $\mathbb{Q} = \mathbb{Q}_x \times \mathbb{Q}_w$  close to a nominal distribution  $\widehat{\mathbb{P}} = \widehat{\mathbb{P}}_x \times \widehat{\mathbb{P}}_w$  in the sense that  $\mathbb{Q}_x$  and  $\mathbb{Q}_w$  are close to  $\widehat{\mathbb{P}}_x$  and  $\widehat{\mathbb{P}}_w$  in Wasserstein distance, respectively.

**Definition 1.1** (Wasserstein distance). *For any  $d \in \mathbb{N}$ , the type-2 Wasserstein distance between two distributions  $\mathbb{Q}_1, \mathbb{Q}_2 \in \mathcal{M}(\mathbb{R}^d)$  is defined as*

$$\mathbb{W}(\mathbb{Q}_1, \mathbb{Q}_2) = \inf_{\pi \in \Pi(\mathbb{Q}_1, \mathbb{Q}_2)} \left( \int_{\mathbb{R}^d \times \mathbb{R}^d} \|\xi_1 - \xi_2\|^2 \pi(d\xi_1, d\xi_2) \right)^{\frac{1}{2}},$$

where  $\Pi(\mathbb{Q}_1, \mathbb{Q}_2)$  denotes the set of all joint distributions or couplings  $\pi \in \mathcal{M}(\mathbb{R}^d \times \mathbb{R}^d)$  of the random vectors  $\xi_1 \in \mathbb{R}^d$  and  $\xi_2 \in \mathbb{R}^d$  with marginal distributions  $\mathbb{Q}_1$  and  $\mathbb{Q}_2$ , respectively.

The dependence of the Wasserstein distance on  $d$  is notationally suppressed to avoid clutter. Note that  $\mathbb{W}(\mathbb{Q}_1, \mathbb{Q}_2)^2$  is naturally interpreted as the optimal value of a transportation problem that determines the minimum cost of moving the distribution  $\mathbb{Q}_1$  to  $\mathbb{Q}_2$ , where the cost of moving a unit probability mass from  $\xi_1$  to  $\xi_2$  is given by the squared Euclidean distance  $\|\xi_1 - \xi_2\|^2$ . For this reason, the optimization variable  $\pi$  is sometimes referred to as a transportation plan and the Wasserstein distance as the earth mover's distance.

Formally, we define the *Wasserstein ambiguity set* as

$$(1.8) \quad \mathbb{B}(\widehat{\mathbb{P}}) = \left\{ \mathbb{Q}_x \times \mathbb{Q}_w : \begin{array}{ll} \mathbb{Q}_x \in \mathcal{M}(\mathbb{R}^n), & \mathbb{W}(\mathbb{Q}_x, \widehat{\mathbb{P}}_x) \leq \rho_x \\ \mathbb{Q}_w \in \mathcal{M}(\mathbb{R}^m), & \mathbb{W}(\mathbb{Q}_w, \widehat{\mathbb{P}}_w) \leq \rho_w \end{array} \right\},$$

where  $\widehat{\mathbb{P}}_x$  and  $\widehat{\mathbb{P}}_w$  represent prescribed nominal distributions that could be constructed via statistical analysis or expert judgement, while the Wasserstein radii  $\rho_x \geq 0$  and  $\rho_w \geq 0$  constitute hyperparameters that quantify the statistician's uncertainty about the nominal distributions of  $x$  and  $w$ . We emphasize that the distributionally robust estimation model (1.7) generalizes all preceding models. Indeed, if  $\rho_w = 0$ , then (1.7) reduces to the first distributionally robust model (1.6), which in turn encompasses both the MMSE estimation problem (1.2) (for  $\rho_x = 0$ ) and the minimax estimation problem (1.5b) (for  $\rho_x = \infty$ ) as special cases.

The distributionally robust estimation model (1.7) is conceptually attractive because the hyperparameters  $\rho_x$  and  $\rho_w$  allow the statistician to specify her level of trust in the nominal prior distribution  $\widehat{\mathbb{P}}_x$  and the nominal noise distribution  $\widehat{\mathbb{P}}_w$ . In the remainder of the paper we will show that (1.7) is also computationally attractive. This is maybe surprising because mixtures of factorizable distributions are generally not factorizable, which implies that the Wasserstein ambiguity set  $\mathbb{B}(\widehat{\mathbb{P}})$  is non-convex.

We remark that one could also work with an alternative ambiguity set of the form

$$(1.9) \quad \mathbb{B}'(\widehat{\mathbb{P}}) = \left\{ \mathbb{Q}_x \times \mathbb{Q}_w : \mathbb{Q}_x \in \mathcal{M}(\mathbb{R}^n), \mathbb{Q}_w \in \mathcal{M}(\mathbb{R}^m), \mathbb{W}(\mathbb{Q}_x \times \mathbb{Q}_w, \widehat{\mathbb{P}}_x \times \widehat{\mathbb{P}}_w) \leq \rho \right\},$$

which involves only a single hyperparameter  $\rho \geq 0$  and is therefore less expressive but maybe easier to calibrate than  $\mathbb{B}(\widehat{\mathbb{P}})$ . The following lemma is instrumental to understanding the relation between  $\mathbb{B}(\widehat{\mathbb{P}})$  and  $\mathbb{B}'(\widehat{\mathbb{P}})$ . The proof of this result is relegated to the appendix.

**Lemma 1.2** (Pythagoras' theorem for Wasserstein distances). *For any  $\mathbb{Q}_x^1, \mathbb{Q}_x^2 \in \mathcal{M}(\mathbb{R}^n)$  and  $\mathbb{Q}_w^1, \mathbb{Q}_w^2 \in \mathcal{M}(\mathbb{R}^m)$  we have  $\mathbb{W}(\mathbb{Q}_x^1 \times \mathbb{Q}_w^1, \mathbb{Q}_x^2 \times \mathbb{Q}_w^2)^2 = \mathbb{W}(\mathbb{Q}_x^1, \mathbb{Q}_x^2)^2 + \mathbb{W}(\mathbb{Q}_w^1, \mathbb{Q}_w^2)^2$ .*

If we denote the ambiguity sets (1.8) and (1.9) temporarily by  $\mathbb{B}_{\rho_x, \rho_w}(\widehat{\mathbb{P}})$  and  $\mathbb{B}'_{\rho}(\widehat{\mathbb{P}})$  in order to make their dependence on the hyperparameters explicit, then Lemma 1.2 implies that

$$\mathbb{B}'_{\rho}(\widehat{\mathbb{P}}) = \bigcup_{\rho_x^2 + \rho_w^2 \leq \rho^2} \mathbb{B}_{\rho_x, \rho_w}(\widehat{\mathbb{P}}).$$

This relation suggests that  $\mathbb{B}'_{\rho}(\widehat{\mathbb{P}})$  could be substantially larger than  $\mathbb{B}_{\rho_x, \rho_w}(\widehat{\mathbb{P}})$  for any fixed  $\rho, \rho_x, \rho_w \geq 0$  with  $\rho_x^2 + \rho_w^2 = \rho^2$  and thus lead to substantially more conservative estimators.

In the following we summarize the key contributions of this paper.

- (1) We construct a safe approximation for the distributionally robust MMSE estimation problem (1.7) by restricting attention to *affine* estimators and by maximizing the average risk over an *outer* approximation of the Wasserstein ambiguity set, which is described through first- and second-order moment conditions. We then prove that this safe approximation is equivalent to a tractable convex program.
- (2) We also study a *dual* estimation problem, which is obtained by interchanging the minimization and maximization operations in the *primal* problem (1.7) and thus lower bounds the optimal value of (1.7). We then construct a safe approximation for this dual problem by restricting the Wasserstein ambiguity set to contain only *normal* distributions. Assuming that the nominal distribution is normal, we prove that this safe approximation is again equivalent to a tractable convex program.
- (3) By construction, the primal and dual estimation problems are upper and lower bounded by their respective safe approximations. We prove, however, that the optimal values of the safe approximations collapse if the nominal distribution is normal. This result has three important implications.
  - (a) The primal and dual estimation problems and their safe approximations are all equivalent. This implies via contributions (1) and (2) that both original estimation problems are tractable.
  - (b) The primal estimation problem is solved by an *affine* estimator, and the dual estimation problem is solved by a *normal* distribution. In other words, we have discovered a new class of adaptive distributionally robust optimization problems for which affine decision rules are optimal.
  - (c) The affine estimator and the normal distribution that solve the primal and dual estimation problems, respectively, form a *Nash equilibrium* for the zero-sum game between the statistician and nature. Thus, the optimal normal distribution constitutes a *least favorable prior*, and the optimal affine estimator represents the corresponding *Bayesian MMSE estimator*.

We leverage these insights to prove that the optimal affine estimator can be constructed easily from the least favorable prior without the need to solve another optimization problem.

- (4) We argue that our main results remain valid if the nominal distribution is any elliptical distribution.
- (5) We develop a tailor-made Frank-Wolfe algorithm that can solve the dual estimation problem orders of magnitude faster than state-of-the-art general purpose solvers. We show that this algorithm enjoys a linear convergence rate. Moreover, we prove that the direction-finding subproblems can be solved in quasi-closed form, which means that the algorithm offers a favorable iteration complexity.

We highlight that the Wasserstein ambiguity set (1.8) is non-convex as it contains only distributions under which the signal and the noise are independent. To our best knowledge, we describe the first distributionally robust optimization model with independence conditions that admits a tractable reformulation.

In the following we briefly survey the landscape of existing MMSE estimation models that have a robustness flavor. Several authors have addressed the *minimax* MMSE estimation problem (1.5a) from the perspective of classical robust optimization [3, 4, 23, 24, 25, 42]. To guarantee computational tractability, in all of these papers the estimators are restricted to be affine functions of the measurements. In this case, the minimax MMSE estimation problem can be reformulated as a tractable semidefinite program (SDP) if the uncertainty set  $\mathbb{X}$  is an ellipsoid [24, 25] or an intersection of two ellipsoids [3]. Similar SDP reformulations are available if the observation matrix  $H$  is also subject to uncertainty and ranges over a spectral norm ball [25] or displays a block circulant structure, with each block ranging over a Frobenius norm ball [4]. If the uncertainty set is described by an intersection of several ellipsoids, then the minimax MMSE estimation problem admits an (inexact) SDP relaxation [23]. Even though the restriction to affine estimators may incur a loss of optimality, affine estimators are known to be near-optimal in all of the above minimax estimation models [42].

Another stream of literature investigates the *distributionally robust* estimation model (1.6) with an ambiguous signal distribution and a crisp noise distribution. When focusing on affine estimators only, this model can be reformulated as a tractable SDP if the uncertainty in the signal distribution is characterized through

spectral constraints on its covariance matrix [26]. This tractability result also extends to uncertain observation matrices. Similar SDP reformulations are available for the distributionally robust estimation model (1.7) when both the signal and the noise distribution are ambiguous and their covariance matrices are subject to spectral constraints [22]. Extensions to uncertain block circulant observation matrices are discussed in [2].

Some authors have studied less structured distributionally robust estimation problems where the signal  $x$  and the measurement  $y$  are governed by a distribution that may *not* obey the linear measurement model (1.1). In this case, the zero-sum game between the statistician and nature admits a Nash equilibrium if nature may choose any distribution that has a bounded Kullback-Leibler divergence with respect to a *normal* nominal distribution [47]. Intriguingly, the (affine) Bayesian MMSE estimator for the nominal distribution is optimal in this model and thus enjoys strong robustness properties. On the downside, there is no hope to improve this estimator's performance by tuning the size of the Kullback-Leibler ambiguity set. The underlying distributionally robust estimation model also serves as a fundamental building block for a robust Kalman filter [48]. Extensions to general  $\tau$ -divergence ambiguity sets that contain only normal distributions are reported in [72]. We emphasize that all papers surveyed so far merely derive SDP reformulations or SDP relaxations that can be addressed with general purpose solvers, but none of them develops a customized solution algorithm.

The present paper extends the distributionally robust MMSE estimation model introduced in [64], which accommodates a simple Wasserstein ambiguity set for the distribution of the signal-measurement pairs and makes no structural assumptions about the measurement noise. Note, however, that the linear measurement model (1.1) abounds in the literature on control theory, signal processing and information theory, implying that there are numerous applications where the measurement noise is *known* to be additive and independent of the signal. As we will see in Section 7, ignoring this structural information may result in weak estimators that sacrifice predictive performance. In Sections 2–4 we will further see that constructing an explicit Nash equilibrium is considerably more difficult in the presence of structural information. Finally, we describe here an accelerated Frank-Wolfe algorithm that improves the sublinear convergence rate established in [64] to a linear rate. We emphasize that, in contrast to the robust MMSE estimators derived in [47, 72] that are insensitive to the radii of the underlying divergence-based ambiguity sets, the estimators constructed here change with the Wasserstein radii  $\rho_x$  and  $\rho_w$ . Thus, using a Wasserstein ambiguity set to robustify the nominal MMSE estimation problem has a regularizing effect and leads to a parametric family of estimators that can be tuned to attain maximum prediction accuracy. Similar connections between robustification and regularization have previously been discovered in the context of statistical learning [63] and covariance estimation [54].

The paper is structured as follows. Sections 2 and 3 develop conservative approximations for the primal and dual distributionally robust MMSE estimation problems, respectively, both of which are equivalent to tractable convex programs. Section 4 shows that if the nominal distribution is normal, then both approximations are exact and can be used to find a Nash equilibrium for the zero-sum game between the statistician and nature. Extensions to elliptical nominal distributions are discussed in Section 5. Section 6 develops an efficient Frank-Wolfe algorithm for the dual MMSE estimation problem, and Section 7 reports on numerical results.

**Notation.** For any  $A \in \mathbb{R}^{d \times d}$  we use  $\text{Tr}[A]$  to denote the trace and  $\|A\| = \sqrt{\text{Tr}[A^\top A]}$  to denote the Frobenius norm of  $A$ . By slight abuse of notation, the Euclidean norm of  $v \in \mathbb{R}^d$  is also denoted by  $\|v\|$ . Moreover,  $I_d$  stands for the identity matrix in  $\mathbb{R}^{d \times d}$ . For any  $A, B \in \mathbb{R}^{d \times d}$ , we use  $\langle A, B \rangle = \text{Tr}[A^\top B]$  to denote the inner product and  $A \otimes B$  to denote the Kronecker product of  $A$  and  $B$ . The space of all symmetric matrices in  $\mathbb{R}^{d \times d}$  is denoted by  $\mathbb{S}^d$ . We use  $\mathbb{S}_+^d$  ( $\mathbb{S}_{++}^d$ ) to represent the cone of symmetric positive semidefinite (positive definite) matrices in  $\mathbb{S}^d$ . For any  $A, B \in \mathbb{S}^d$ , the relation  $A \succeq B$  ( $A \succ B$ ) means that  $A - B \in \mathbb{S}_+^d$  ( $A - B \in \mathbb{S}_{++}^d$ ). The unique positive semidefinite square root of a matrix  $A \in \mathbb{S}_+^d$  is denoted by  $A^{\frac{1}{2}}$ . For any  $A \in \mathbb{S}^d$ ,  $\lambda_{\min}(A)$  and  $\lambda_{\max}(A)$  denote the minimum and maximum eigenvalues of  $A$ , respectively.

## 2. THE GELBRICH MMSE ESTIMATION PROBLEM

The distributionally robust estimation problem (1.7) poses two fundamental challenges. First, checking feasibility of the inner maximization problem in (1.7) requires computing the Wasserstein distances  $W(\widehat{\mathbb{P}}_x, \mathbb{Q}_x)$  and  $W(\widehat{\mathbb{P}}_w, \mathbb{Q}_w)$ , which is #P-hard even if  $\widehat{\mathbb{P}}_x$  and  $\widehat{\mathbb{P}}_w$  are simple two-point distributions, while  $\mathbb{Q}_x$  and  $\mathbb{Q}_w$  are uniform distributions on hypercubes [68]. Efficient algorithms for computing Wasserstein distances are available only if both involved distributions are discrete [15, 59, 66], and analytical formulas are only known in exceptional cases (*e.g.*, if both distributions are Gaussian [32] or belong to the same family of elliptical distributions [31]). The second challenge is that the outer minimization problem in (1.7) constitutes an infinite-dimensional functional optimization problem. In order to bypass these computational challenges, we first seek a conservative approximation for (1.7) by relaxing the ambiguity set  $\mathbb{B}(\widehat{\mathbb{P}})$  and restricting the feasible set  $\mathcal{F}$ . We begin by constructing an outer approximation for the ambiguity set. To this end, we introduce a new distance measure on the space of mean vectors and covariance matrices.

**Definition 2.1** (Gelbrich distance). *For any  $d \in \mathbb{N}$ , the Gelbrich distance between two tuples of mean vectors and covariance matrices  $(\mu_1, \Sigma_1), (\mu_2, \Sigma_2) \in \mathbb{R}^d \times \mathbb{S}_+^d$  is defined as*

$$\mathbb{G}((\mu_1, \Sigma_1), (\mu_2, \Sigma_2)) = \sqrt{\|\mu_1 - \mu_2\|^2 + \text{Tr} [\Sigma_1 + \Sigma_2 - 2 \left( \Sigma_2^{\frac{1}{2}} \Sigma_1 \Sigma_2^{\frac{1}{2}} \right)^{\frac{1}{2}}]}.$$

The dependence of the Gelbrich distance on  $d$  is notationally suppressed in order to avoid clutter. One can show that  $\mathbb{G}$  constitutes a metric on  $\mathbb{R}^d \times \mathbb{S}_+^d$ , that is,  $\mathbb{G}$  is symmetric, non-negative, vanishes if and only if  $(\mu_1, \Sigma_1) = (\mu_2, \Sigma_2)$  and satisfies the triangle inequality [32, pp. 239].

**Proposition 2.2** (Commuting covariance matrices [32, p. 239]). *If  $\mu_1, \mu_2 \in \mathbb{R}^d$  are identical and  $\Sigma_1, \Sigma_2 \in \mathbb{S}_+^d$  commute ( $\Sigma_1 \Sigma_2 = \Sigma_2 \Sigma_1$ ), then the Gelbrich distance simplifies to  $\mathbb{G}((\mu_1, \Sigma_1), (\mu_2, \Sigma_2)) = \|\sqrt{\Sigma_1} - \sqrt{\Sigma_2}\|$ .*

While the Gelbrich distance is non-convex, the squared Gelbrich distance is convex in all of its arguments.

**Proposition 2.3** (Convexity and continuity of the squared Gelbrich distance). *The squared Gelbrich distance  $\mathbb{G}((\mu_1, \Sigma_1), (\mu_2, \Sigma_2))^2$  is jointly convex and continuous in  $\mu_1, \mu_2 \in \mathbb{R}^d$  and  $\Sigma_1, \Sigma_2 \in \mathbb{S}_+^d$ .*

*Proof of Proposition 2.3.* By [51, Proposition 2], the squared Gelbrich distance  $\mathbb{G}((\mu_1, \Sigma_1), (\mu_2, \Sigma_2))^2$  coincides with the optimal value of the semidefinite program

$$\begin{aligned} \min \quad & \|\mu_1 - \mu_2\|_2^2 + \text{Tr} [\Sigma_1 + \Sigma_2 - 2C] \\ \text{s. t.} \quad & C \in \mathbb{R}^{d \times d} \\ & \begin{bmatrix} \Sigma_1 & C \\ C^\top & \Sigma_2 \end{bmatrix} \succeq 0, \end{aligned}$$

see also [18, Section 3]. Less general results that hold when one of the matrices  $\Sigma_1$  or  $\Sigma_2$  is positive definite are reported in [32, 45, 56]. The convexity of the squared Gelbrich distance then follows from [7, Proposition 3.3.1], which guarantees that convexity is preserved under partial minimization. Moreover, the continuity of the squared Gelbrich distance follows from the continuity of the matrix square root established in Lemma A.2.  $\square$

Our interest in the Gelbrich distance stems mainly from the next proposition, which lower bounds the Wasserstein distance between two distributions in terms of their first- and second-order moments. We will later see that this bound becomes tight when  $\mathbb{Q}_1$  and  $\mathbb{Q}_2$  are normal or—more generally—elliptical distributions of the same type.

**Proposition 2.4** (Moment bound on the Wasserstein distance [31, Theorem 2.1]). *For any distributions  $\mathbb{Q}_1, \mathbb{Q}_2 \in \mathcal{M}(\mathbb{R}^d)$  with mean vectors  $\mu_1, \mu_2 \in \mathbb{R}^d$  and covariance matrices  $\Sigma_1, \Sigma_2 \in \mathbb{S}_+^d$ , respectively, we have*

$$\mathbb{W}(\mathbb{Q}_1, \mathbb{Q}_2) \geq \mathbb{G}((\mu_1, \Sigma_1), (\mu_2, \Sigma_2)).$$

Proposition 2.4 prompts us to construct an outer approximation for the Wasserstein ambiguity set  $\mathbb{B}(\widehat{\mathbb{P}})$  by using the Gelbrich distance. Specifically, we define the *Gelbrich ambiguity set* centered at  $\widehat{\mathbb{P}} = \widehat{\mathbb{P}}_x \times \widehat{\mathbb{P}}_w$  as

$$\mathbb{G}(\widehat{\mathbb{P}}) = \left\{ \begin{array}{l} \mathbb{Q}_x \in \mathcal{M}(\mathbb{R}^n), \quad \mu_x = \mathbb{E}_{\mathbb{Q}_x}[x], \quad \Sigma_x = \mathbb{E}_{\mathbb{Q}_x}[xx^\top] - \mu_x\mu_x^\top \\ \mathbb{Q}_x \times \mathbb{Q}_w : \quad \mathbb{Q}_w \in \mathcal{M}(\mathbb{R}^m), \quad \mu_w = \mathbb{E}_{\mathbb{Q}_w}[w], \quad \Sigma_w = \mathbb{E}_{\mathbb{Q}_w}[ww^\top] - \mu_w\mu_w^\top \\ \mathbb{G}((\mu_x, \Sigma_x), (\widehat{\mu}_x, \widehat{\Sigma}_x)) \leq \rho_x, \quad \mathbb{G}((\mu_w, \Sigma_w), (\widehat{\mu}_w, \widehat{\Sigma}_w)) \leq \rho_w \end{array} \right\},$$

where  $\widehat{\mu}_x$  and  $\widehat{\mu}_w$  denote the mean vectors and  $\widehat{\Sigma}_x$  and  $\widehat{\Sigma}_w$  the covariance matrices of  $\widehat{\mathbb{P}}_x$  and  $\widehat{\mathbb{P}}_w$ , respectively.

**Corollary 2.5** (Relation between Gelbrich and Wasserstein ambiguity sets). *For any  $\widehat{\mathbb{P}} = \widehat{\mathbb{P}}_x \times \widehat{\mathbb{P}}_w$  with  $\widehat{\mathbb{P}}_x \in \mathcal{M}(\mathbb{R}^n)$  and  $\widehat{\mathbb{P}}_w \in \mathcal{M}(\mathbb{R}^m)$  we have  $\mathbb{B}(\widehat{\mathbb{P}}) \subseteq \mathbb{G}(\widehat{\mathbb{P}})$ .*

*Proof of Corollary 2.5.* Select any  $\mathbb{Q} = \mathbb{Q}_x \times \mathbb{Q}_w \in \mathbb{B}(\widehat{\mathbb{P}})$  and define  $\mu_x$  and  $\mu_w$  as the mean vectors and  $\Sigma_x$  and  $\Sigma_w$  as the covariance matrices of  $\mathbb{Q}_x$  and  $\mathbb{Q}_w$ , respectively. By Proposition 2.4 we then have

$$\mathbb{G}((\mu_x, \Sigma_x), (\widehat{\mu}_x, \widehat{\Sigma}_x)) \leq \mathbb{W}(\mathbb{Q}_x, \widehat{\mathbb{P}}_x) \leq \rho_x \quad \text{and} \quad \mathbb{G}((\mu_w, \Sigma_w), (\widehat{\mu}_w, \widehat{\Sigma}_w)) \leq \mathbb{W}(\mathbb{Q}_w, \widehat{\mathbb{P}}_w) \leq \rho_w,$$

which in turn implies that  $\mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}})$ . We may thus conclude that  $\mathbb{B}(\widehat{\mathbb{P}}) \subseteq \mathbb{G}(\widehat{\mathbb{P}})$ .  $\square$

By restricting  $\mathcal{F}$  to the set  $\mathcal{A}$  of all affine estimators while relaxing  $\mathbb{B}(\widehat{\mathbb{P}})$  to the Gelbrich ambiguity set  $\mathbb{G}(\widehat{\mathbb{P}})$ , we obtain the following conservative approximation of the distributionally robust estimation problem (1.7).

$$(2.1) \quad \underset{\psi \in \mathcal{A}}{\text{minimize}} \quad \sup_{\mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}})} \mathcal{R}(\psi, \mathbb{Q})$$

From now on we will call (1.7) and (2.1) the Wasserstein and Gelbrich MMSE estimation problems, and we will refer to their minimizers as Wasserstein and Gelbrich MMSE estimators, respectively. As the average risk  $\mathcal{R}(\psi, \mathbb{Q})$  of a fixed affine estimator  $\psi \in \mathcal{A}$  is convex and quadratic in the mean vector  $\mu$  and affine in the covariance matrix  $\Sigma$  of the distribution  $\mathbb{Q}$ , the inner maximization problem in (2.1) is non-convex. Thus, one might suspect that the Gelbrich MMSE estimation problem is intractable. Below we will show, however, that (2.1) is equivalent to a finite convex program that can be solved in polynomial time. To this end, we first show that, under mild conditions, problem (2.1) is stable with respect to changes of its input parameters.

**Proposition 2.6** (Regularity of the Gelbrich MMSE estimation problem). *The Gelbrich MMSE estimation problem (2.1) enjoys the following regularity properties.*

- (i) **Conservativeness:** *Problem (2.1) upper bounds the Wasserstein MMSE estimation problem (1.7).*
- (ii) **Solvability:** *If  $\rho_x > 0$  and  $\rho_w > 0$ , then the minimum of (2.1) is attained.*
- (iii) **Stability:** *If  $\rho_x > 0$  and  $\rho_w > 0$ , then the minimum of (2.1) is continuous in  $(\widehat{\mu}_x, \widehat{\mu}_w, \widehat{\Sigma}_x, \widehat{\Sigma}_w)$ .*

The proof of Proposition 2.6 is lengthy and technical and is therefore relegated to the appendix. We are now ready to prove the main result of this section.



**Theorem 2.7** (Gelbrich MMSE estimation problem). *The Gelbrich MMSE estimation problem (2.1) is equivalent to the finite convex optimization problem*

$$(2.2) \quad \begin{aligned} \inf \quad & \gamma_x(\rho_x^2 - \text{Tr}[\widehat{\Sigma}_x]) + \gamma_x^2 \langle [\gamma_x I_n - (I_n - AH)^\top (I_n - AH)]^{-1}, \widehat{\Sigma}_x \rangle \\ & + \gamma_w(\rho_w^2 - \text{Tr}[\widehat{\Sigma}_w]) + \gamma_w^2 \langle (\gamma_w I_m - A^\top A)^{-1}, \widehat{\Sigma}_w \rangle \\ \text{s. t.} \quad & A \in \mathbb{R}^{n \times m}, \quad \gamma_x, \gamma_w \in \mathbb{R}_+ \\ & \gamma_x I_n - (I_n - AH)^\top (I_n - AH) \succ 0, \quad \gamma_w I_m - A^\top A \succ 0. \end{aligned}$$

Moreover, if  $\rho_x > 0$  and  $\rho_w > 0$ , then (2.2) admits an optimal solution<sup>1</sup>  $A^*$ , and the infimum of (2.1) is attained by the affine estimator  $\psi^*(y) = A^*y + b^*$ , where  $b^* = \widehat{\mu}_x - A^*(H\widehat{\mu}_x + \widehat{\mu}_w)$ .

*Proof of Theorem 2.7.* Throughout this proof we denote by  $\psi_{A,b} \in \mathcal{A}$  the affine estimator  $\psi_{A,b}(y) = Ay + b$  corresponding to the sensitivity matrix  $A \in \mathbb{R}^{n \times m}$  and the vector  $b \in \mathbb{R}^n$  of intercepts. In the following we fix some  $A \in \mathbb{R}^{n \times m}$  and define  $K = I_n - AH$  in order to simplify the notation. By the definitions of the average risk  $\mathcal{R}(\psi, \mathbb{Q})$  and the Gelbrich ambiguity set  $\mathbb{G}(\widehat{\mathbb{P}})$ , we then have

$$(2.3) \quad \inf_b \sup_{\mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}})} \mathcal{R}(\psi_{A,b}, \mathbb{Q}) = \begin{cases} \inf_b \sup_{\substack{\mu_x, \mu_w \\ \Sigma_x, \Sigma_w \geq 0}} \langle K^\top K, \Sigma_x + \mu_x \mu_x^\top \rangle + \langle A^\top A, \Sigma_w + \mu_w \mu_w^\top \rangle + b^\top b \\ \quad - 2\mu_x^\top K^\top A \mu_w - 2b^\top (K \mu_x - A \mu_w) \\ \text{s. t.} \quad \mathbb{G}((\mu_x, \Sigma_x), (\widehat{\mu}_x, \widehat{\Sigma}_x))^2 \leq \rho_x^2 \\ \quad \mathbb{G}((\mu_w, \Sigma_w), (\widehat{\mu}_w, \widehat{\Sigma}_w))^2 \leq \rho_w^2. \end{cases}$$

The outer minimization problem in (2.3) is convex because the objective function of the minimax problem is convex in  $b$  for any fixed  $(\mu_x, \mu_w, \Sigma_x, \Sigma_w)$  and because convexity is preserved under maximization. Moreover, the inner maximization problem in (2.3) is non-convex because its objective function is convex in  $(\mu_x, \mu_w)$ . This observation prompts us to maximize over  $(\mu_x, \mu_w)$  and  $(\Sigma_x, \Sigma_w)$  sequentially and to reformulate (2.3) as

$$(2.4) \quad \begin{aligned} \inf_b \quad & \sup_{\substack{\mu_x, \mu_w \\ \|\mu_x - \widehat{\mu}_x\| \leq \rho_x \\ \|\mu_w - \widehat{\mu}_w\| \leq \rho_w}} \sup_{\Sigma_x, \Sigma_w \geq 0} \langle K^\top K, \Sigma_x + \mu_x \mu_x^\top \rangle + \langle A^\top A, \Sigma_w + \mu_w \mu_w^\top \rangle + b^\top b \\ & - 2\mu_x^\top K^\top A \mu_w - 2b^\top (K \mu_x - A \mu_w) \\ \text{s. t.} \quad & \mathbb{G}((\mu_x, \Sigma_x), (\widehat{\mu}_x, \widehat{\Sigma}_x))^2 \leq \rho_x^2 \\ & \mathbb{G}((\mu_w, \Sigma_w), (\widehat{\mu}_w, \widehat{\Sigma}_w))^2 \leq \rho_w^2. \end{aligned}$$

As  $\|\mu_x - \widehat{\mu}_x\| \leq \mathbb{G}((\mu_x, \Sigma_x), (\widehat{\mu}_x, \widehat{\Sigma}_x))$  and as this inequality is tight for  $\Sigma_x = \widehat{\Sigma}_x$ , the extra constraint  $\|\mu_x - \widehat{\mu}_x\| \leq \rho_x$  is actually redundant and merely ensures that the maximization problem over  $\Sigma_x$  remains feasible for any admissible choice of  $\mu_x$ . An analogous statement holds for  $\mu_w$  and  $\Sigma_w$ . By the definition of the Gelbrich distance, the innermost maximization problem over  $(\Sigma_x, \Sigma_w)$  in (2.4) admits the Lagrangian dual

$$(2.5) \quad \begin{aligned} \inf_{\gamma_x, \gamma_w \geq 0} \sup_{\Sigma_x, \Sigma_w \geq 0} \quad & \langle K^\top K, \Sigma_x + \mu_x \mu_x^\top \rangle + \langle A^\top A, \Sigma_w + \mu_w \mu_w^\top \rangle + b^\top b - 2\mu_x^\top K^\top A \mu_w - 2b^\top (K \mu_x - A \mu_w) \\ & + \gamma_x(\rho_x^2 - \|\mu_x - \widehat{\mu}_x\|^2 - \text{Tr}[\Sigma_x + \widehat{\Sigma}_x - 2(\widehat{\Sigma}_x^{\frac{1}{2}} \Sigma_x \widehat{\Sigma}_x^{\frac{1}{2}})^{\frac{1}{2}}]) \\ & + \gamma_w(\rho_w^2 - \|\mu_w - \widehat{\mu}_w\|^2 - \text{Tr}[\Sigma_w + \widehat{\Sigma}_w - 2(\widehat{\Sigma}_w^{\frac{1}{2}} \Sigma_w \widehat{\Sigma}_w^{\frac{1}{2}})^{\frac{1}{2}}]). \end{aligned}$$

Strong duality holds by [7, Proposition 5.5.4], which applies because the primal problem has a non-empty compact feasible set. Next, we observe that the inner maximization problem in (2.5) can be solved analytically

<sup>1</sup>We say that  $A^*$  solves (2.2) if adding the constraint  $A = A^*$  does not change the infimum of (2.2). Note that the infimum of the resulting problem over  $(\gamma_x, \gamma_w)$  may not be attained, *i.e.*, the existence of a solution  $A^*$  does not imply that (2.2) is solvable.

by using Proposition A.3 in the appendix, and thus the dual problem (2.5) is equivalent to

$$(2.6) \quad \inf_{\substack{\gamma_x, \gamma_w \\ \gamma_x I_n \succ K^\top K \\ \gamma_w I_m \succ A^\top A}} \langle K^\top K, \mu_x \mu_x^\top \rangle + \langle A^\top A, \mu_w \mu_w^\top \rangle + b^\top b - 2\mu_x^\top K^\top A \mu_w - 2b^\top (K \mu_x - A \mu_w) \\ + \gamma_x (\rho_x^2 - \|\mu_x - \hat{\mu}_x\|^2 - \text{Tr} [\widehat{\Sigma}_x] + \gamma_x \langle (\gamma_x I_n - K^\top K)^{-1}, \widehat{\Sigma}_x \rangle) \\ + \gamma_w (\rho_w^2 - \|\mu_w - \hat{\mu}_w\|^2 - \text{Tr} [\widehat{\Sigma}_w] + \gamma_w \langle (\gamma_w I_m - A^\top A)^{-1}, \widehat{\Sigma}_w \rangle).$$

Substituting (2.6) back into (2.4) then allows us to reformulate the Gelbrich MMSE estimation problem (2.6) as

$$(2.7) \quad \inf_b \sup_{\substack{\mu_x, \mu_w \\ \|\mu_x - \hat{\mu}_x\| \leq \rho_x \\ \|\mu_w - \hat{\mu}_w\| \leq \rho_w}} \inf_{\substack{\gamma_x, \gamma_w \\ \gamma_x I_n \succ K^\top K \\ \gamma_w I_m \succ A^\top A}} \langle K^\top K, \mu_x \mu_x^\top \rangle + \langle A^\top A, \mu_w \mu_w^\top \rangle + b^\top b - 2\mu_x^\top K^\top A \mu_w - 2b^\top (K \mu_x - A \mu_w) \\ + \gamma_x (\rho_x^2 - \|\mu_x - \hat{\mu}_x\|^2 - \text{Tr} [\widehat{\Sigma}_x] + \gamma_x \langle (\gamma_x I_n - K^\top K)^{-1}, \widehat{\Sigma}_x \rangle) \\ + \gamma_w (\rho_w^2 - \|\mu_w - \hat{\mu}_w\|^2 - \text{Tr} [\widehat{\Sigma}_w] + \gamma_w \langle (\gamma_w I_m - A^\top A)^{-1}, \widehat{\Sigma}_w \rangle).$$

The infimum of the inner minimization problem over  $(\gamma_x, \gamma_w)$  in (2.7) is convex quadratic in  $b$ . Moreover, it is concave in  $(\mu_x, \mu_w)$  because  $K^\top K - \gamma_x I_n \prec 0$  and  $A^\top A - \gamma_w I_m \prec 0$  for any feasible choice of  $(\gamma_x, \gamma_w)$  and because concavity is preserved under minimization. Finally, the feasible set for  $(\mu_x, \mu_w)$  is convex and compact. By Sion's classical minimax theorem, we may therefore interchange the infimum over  $b$  with the supremum over  $(\mu_x, \mu_w)$ . The minimization problem over  $b$  thus reduces to an unconstrained (strictly) convex quadratic program that has the unique optimal solution  $b = K \mu_x - A \mu_w$ . Substituting this expression back into (2.7) then yields

$$(2.8) \quad \sup_{\mu_x, \mu_w} \inf_{\substack{\gamma_x, \gamma_w \\ \gamma_x I_n \succ K^\top K \\ \gamma_w I_m \succ A^\top A}} \gamma_x (\rho_x^2 - \|\mu_x - \hat{\mu}_x\|^2 - \text{Tr} [\widehat{\Sigma}_x]) + \gamma_x^2 \langle (\gamma_x I_n - K^\top K)^{-1}, \widehat{\Sigma}_x \rangle \\ + \gamma_w (\rho_w^2 - \|\mu_w - \hat{\mu}_w\|^2 - \text{Tr} [\widehat{\Sigma}_w]) + \gamma_w^2 \langle (\gamma_w I_m - A^\top A)^{-1}, \widehat{\Sigma}_w \rangle.$$

It is easy to verify that the resulting maximization problem over  $(\mu_x, \mu_w)$  is solved by  $\mu_x = \hat{\mu}_x$  and  $\mu_w = \hat{\mu}_w$ . Substituting the corresponding optimal value into (2.3) finally yields

$$\inf_b \sup_{\mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}})} \mathcal{R}(\psi_{A,b}, \mathbb{Q}) = \begin{cases} \inf_{\substack{\gamma_x, \gamma_w \\ \gamma_x I_n \succ K^\top K \\ \gamma_w I_m \succ A^\top A}} \gamma_x (\rho_x^2 - \text{Tr} [\widehat{\Sigma}_x]) + \gamma_x^2 \langle (\gamma_x I_n - K^\top K)^{-1}, \widehat{\Sigma}_x \rangle \\ + \gamma_w (\rho_w^2 - \text{Tr} [\widehat{\Sigma}_w]) + \gamma_w^2 \langle (\gamma_w I_m - A^\top A)^{-1}, \widehat{\Sigma}_w \rangle. \end{cases}$$

From the above equation and the definition of  $K$  it is evident that the Gelbrich MMSE estimation problem

$$(2.9) \quad \inf_{\psi \in \mathcal{A}} \sup_{\mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}})} \mathcal{R}(\psi, \mathbb{Q}) = \inf_{A, b} \sup_{\mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}})} \mathcal{R}(\psi_{A,b}, \mathbb{Q})$$

is indeed equivalent to the finite convex optimization problem (2.2).

Assume now that  $\rho_x > 0$  and  $\rho_w > 0$ . In this case we know from Proposition 2.6 (ii) that the Gelbrich MMSE estimation problem (2.9) admits an optimal affine estimator  $\psi^*(y) = A^* y + b^*$  for some  $A^* \in \mathbb{R}^{n \times m}$  and  $b^* \in \mathbb{R}^m$ . The reasoning in the first part of the proof then implies that  $A^*$  solves (2.2). Moreover, it implies that  $b^*$  is optimal in (2.3) when we fix  $A = A^*$ . As (2.3) is equivalent to (2.7) and as the unique optimal solution of (2.7) for  $A = A^*$  is given by  $b = \hat{\mu}_x - A^*(H \hat{\mu}_x + \hat{\mu}_w)$ , we may finally conclude that

$$b^* = \hat{\mu}_x - A^*(H \hat{\mu}_x + \hat{\mu}_w).$$

By reversing these arguments, one can further show that if  $A^*$  solves (2.2) and  $b^*$  is defined as above, then the affine estimator  $\psi^*(y) = A^* y + b^*$  is optimal in (2.9). This observation completes the proof.  $\square$

The strict semidefinite inequalities in (2.2) ensure that the inverse matrices in the objective function are well-defined. Using Schur complement arguments, the convex program (2.2) can be further simplified to a standard semidefinite program (SDP), which can be addressed with off-the-shelf solvers.

**Corollary 2.8** (SDP reformulation). *The Gelbrich MMSE estimation problem (2.1) is equivalent to the SDP*

$$(2.10) \quad \begin{aligned} & \inf \quad \gamma_x(\rho_x^2 - \text{Tr}[\widehat{\Sigma}_x]) + \text{Tr}[U_x] + \gamma_w(\rho_w^2 - \text{Tr}[\widehat{\Sigma}_w]) + \text{Tr}[U_w] \\ & \text{s. t.} \quad A \in \mathbb{R}^{n \times m}, \quad \gamma_x, \gamma_w \in \mathbb{R}_+ \\ & \quad U_x \in \mathbb{S}_+^n, \quad V_x \in \mathbb{S}_+^n, \quad U_w \in \mathbb{S}_+^m, \quad V_w \in \mathbb{S}_+^m \\ & \quad \begin{bmatrix} U_x & \gamma_x \widehat{\Sigma}_x^{\frac{1}{2}} \\ \gamma_x \widehat{\Sigma}_x^{\frac{1}{2}} & V_x \end{bmatrix} \succeq 0, \quad \begin{bmatrix} \gamma_x I_n - V_x & I_n - H^\top A^\top \\ I_n - AH & I_n \end{bmatrix} \succeq 0 \\ & \quad \begin{bmatrix} U_w & \gamma_w \widehat{\Sigma}_w^{\frac{1}{2}} \\ \gamma_w \widehat{\Sigma}_w^{\frac{1}{2}} & V_w \end{bmatrix} \succeq 0, \quad \begin{bmatrix} \gamma_w I_m - V_w & A^\top \\ A & I_n \end{bmatrix} \succeq 0. \end{aligned}$$

*Proof of Corollary 2.8.* Define the extended real-valued function  $h_w : \mathbb{R}^{n \times m} \times \mathbb{R}_+ \rightarrow (-\infty, \infty]$  through

$$h_w(A, \gamma_w) = \begin{cases} \gamma_w^2 \langle (\gamma_w I_m - A^\top A)^{-1}, \widehat{\Sigma}_w \rangle & \text{if } \gamma_w I_m - A^\top A \succ 0, \\ \infty & \text{otherwise.} \end{cases}$$

If  $\gamma_w I_m - A^\top A \succ 0$ , then, we have

$$(2.11) \quad \begin{aligned} h_w(A, \gamma_w) &= \inf_{U_w \succeq 0} \left\{ \text{Tr}[U_w] : U_w \succeq \gamma_w^2 \widehat{\Sigma}_w^{\frac{1}{2}} (\gamma_w I_m - A^\top A)^{-1} \widehat{\Sigma}_w^{\frac{1}{2}} \right\} \\ &= \inf_{U_w \succeq 0, V_w \succ 0} \left\{ \text{Tr}[U_w] : U_w \succeq \gamma_w^2 \widehat{\Sigma}_w^{\frac{1}{2}} V_w^{-1} \widehat{\Sigma}_w^{\frac{1}{2}}, \gamma_w I_m - A^\top A \succeq V_w \right\} \\ &= \inf_{U_w \succeq 0, V_w \succ 0} \left\{ \text{Tr}[U_w] : \begin{bmatrix} \gamma_w I_m - V_w & A^\top \\ A & I_n \end{bmatrix} \succeq 0, \begin{bmatrix} U_w & \gamma_w \widehat{\Sigma}_w^{\frac{1}{2}} \\ \gamma_w \widehat{\Sigma}_w^{\frac{1}{2}} & V_w \end{bmatrix} \succeq 0 \right\}, \end{aligned}$$

where the first equality holds due to the cyclicity of the trace operator and because  $U_w \succeq \bar{U}_w$  implies  $\text{Tr}[U_w] \geq \text{Tr}[\bar{U}_w]$  for all  $U_w, \bar{U}_w \succeq 0$ , the second equality holds because  $V_w \succeq \bar{V}_w$  is equivalent to  $V_w^{-1} \preceq \bar{V}_w^{-1}$  for all  $V_w, \bar{V}_w \succ 0$ , and the last equality follows from standard Schur complement arguments; see, e.g., [9, § A.5.5]. If  $\gamma_w I_m - A^\top A \not\succeq 0$ , on the other hand, then the first matrix inequality in (2.11) implies that  $V_w$  must have at least one non-positive eigenvalue, which contradicts the constraint  $V_w \succ 0$ . The SDP (2.11) is therefore infeasible, and its infimum evaluates to  $\infty$ . Thus,  $h_w(A, \gamma_w)$  coincides with the optimal value of the SDP (2.11) for all  $A \in \mathbb{R}^{n \times m}$  and  $\gamma_w \in \mathbb{R}_+$ .

A similar SDP reformulation can be derived for the function  $h_x : \mathbb{R}^{n \times m} \times \mathbb{R}_+ \rightarrow (-\infty, \infty]$  defined through

$$h_x(A, \gamma_x) = \begin{cases} \gamma_x^2 \langle [\gamma_x I_n - (I_n - AH)^\top (I_n - AH)]^{-1}, \widehat{\Sigma}_x \rangle & \text{if } \gamma_x I_n - (I_n - AH)^\top (I_n - AH) \succ 0, \\ \infty & \text{otherwise.} \end{cases}$$

The claim now follows by substituting the SDP reformulations for  $h_w(A, \gamma_w)$  and  $h_x(A, \gamma_x)$  into (2.2). In doing so, we may relax the strict semidefinite inequalities  $V_w \succ 0$  and  $V_x \succ 0$  to weak inequalities  $V_w \succeq 0$  and  $V_x \succeq 0$ , which amounts to taking the closure of the (non-empty) feasible set and does not change the infimum of problem (2.2). This observation completes the proof.  $\square$

**Remark 2.9** (Numerical stability). *The SDP (2.10) requires the square roots of the nominal covariance matrices as inputs. Unfortunately, iterative methods for computing matrix square roots often suffer from numerical instability in high dimensions. As a remedy, one may replace those matrix inequalities in (2.10) that involve  $\widehat{\Sigma}_x^{\frac{1}{2}}$  and  $\widehat{\Sigma}_w^{\frac{1}{2}}$  with*

$$\begin{bmatrix} U_x & \gamma_x \Lambda_x^\top \\ \gamma_x \Lambda_x & V_x \end{bmatrix} \succeq 0 \quad \text{and} \quad \begin{bmatrix} U_w & \gamma_w \Lambda_w^\top \\ \gamma_w \Lambda_w & V_w \end{bmatrix} \succeq 0,$$

where  $\Lambda_x$  and  $\Lambda_w$  represent the lower triangular Cholesky factors of  $\widehat{\Sigma}_x$  and  $\widehat{\Sigma}_w$ , respectively. Thus, we have  $\widehat{\Sigma}_x = \Lambda_x \Lambda_x^\top$  and  $\widehat{\Sigma}_w = \Lambda_w \Lambda_w^\top$ . We emphasize that  $\Lambda_x$  and  $\Lambda_w$  can be computed reliably in high dimensions.

## 3. THE DUAL WASSERSTEIN MMSE ESTIMATION PROBLEM OVER NORMAL PRIORS

We now examine the dual Wasserstein MMSE estimation problem

$$(3.1) \quad \text{maximize}_{\mathbb{Q} \in \mathbb{B}(\widehat{\mathbb{P}})} \inf_{\psi \in \mathcal{F}} \mathcal{R}(\psi, \mathbb{Q}),$$

which is obtained from (1.7) by interchanging the order of minimization and maximization. Any maximizer  $\mathbb{Q}^*$  of this dual estimation problem, if it exists, will henceforth be called a *least favorable prior*. Unfortunately, problem (3.1) is generically intractable. Below we will demonstrate, however, that (3.1) becomes tractable if the nominal distribution  $\widehat{\mathbb{P}}$  is normal.

**Definition 3.1** (Normal distributions). *We say that  $\mathbb{P}$  is a normal distribution on  $\mathbb{R}^d$  with mean  $\mu \in \mathbb{R}^d$  and covariance matrix  $\Sigma \in \mathbb{S}_+^d$ , that is,  $\mathbb{P} = \mathcal{N}(\mu, \Sigma)$ , if  $\mathbb{P}$  is supported on  $\text{supp}(\mathbb{P}) = \{\mu + Ev : v \in \mathbb{R}^k\}$ , and if the density function of  $\mathbb{P}$  with respect to the Lebesgue measure on  $\text{supp}(\mathbb{P})$  is given by*

$$q_{\mathbb{P}}(\xi) = \frac{1}{\sqrt{(2\pi)^k \det(D)}} e^{-(\xi - \mu)^\top E D^{-1} E^\top (\xi - \mu)},$$

where  $k = \text{rank}(\Sigma)$ ,  $D \in \mathbb{S}_{++}^k$  is the diagonal matrix of the positive eigenvalues of  $\Sigma$ , and  $E \in \mathbb{R}^{d \times k}$  is the matrix whose columns correspond to the orthonormal eigenvectors of the positive eigenvalues of  $\Sigma$ .

Definition 3.1 also accounts for degenerate normal distributions with singular covariance matrices. We now recall some basic properties of normal distributions that are crucial for the results of this paper.

**Proposition 3.2** (Affine transformations [27, Theorem 2.16]). *If  $\xi \in \mathbb{R}^d$  follows the normal distribution  $\mathcal{N}(\mu, \Sigma)$ , while  $A \in \mathbb{R}^{k \times d}$  and  $b \in \mathbb{R}^k$ , then  $A\xi + b \in \mathbb{R}^k$  follows the normal distribution  $\mathcal{N}(A\mu + b, A\Sigma A^\top)$ .*

**Proposition 3.3** (Affine conditional expectations [10, Corollary 5]). *Assume that  $\xi \in \mathbb{R}^d$  follows the normal distribution  $\mathbb{P} = \mathcal{N}(\mu, \Sigma)$  and that*

$$\xi = \begin{bmatrix} \xi_1 \\ \xi_2 \end{bmatrix}, \quad \mu = \begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix} \quad \text{and} \quad \Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix},$$

where  $\xi_1, \mu_1 \in \mathbb{R}^{d_1}$ ,  $\xi_2, \mu_2 \in \mathbb{R}^{d_2}$ ,  $\Sigma_{11} \in \mathbb{R}^{d_1 \times d_1}$ ,  $\Sigma_{22} \in \mathbb{R}^{d_2 \times d_2}$  and  $\Sigma_{12} = \Sigma_{21}^\top \in \mathbb{R}^{d_1 \times d_2}$  for some  $d_1, d_2 \in \mathbb{N}$  with  $d_1 + d_2 = d$ . Then, there exist  $A \in \mathbb{R}^{d_1 \times d_2}$  and  $b \in \mathbb{R}^{d_1}$  such that  $\mathbb{E}_{\mathbb{P}}[\xi_1 | \xi_2] = A\xi_2 + b$   $\mathbb{P}$ -almost surely.

Another useful but lesser known property of normal distributions is that their Wasserstein distances can be expressed analytically in terms of the distributions' first- and second-order moments.

**Proposition 3.4** (Wasserstein distance between normal distributions [32, Proposition 7]). *The Wasserstein distance between two normal distributions  $\mathbb{Q}_1 = \mathcal{N}(\mu_1, \Sigma_1)$  and  $\mathbb{Q}_2 = \mathcal{N}(\mu_2, \Sigma_2)$  equals the Gelbrich distance between their mean vectors and covariance matrices, that is,  $\mathbb{W}(\mathbb{Q}_1, \mathbb{Q}_2) = \mathbb{G}((\mu_1, \Sigma_1), (\mu_2, \Sigma_2))$ .*

Assume now that the nominal distributions of the parameter  $x \in \mathbb{R}^n$  and the noise  $w \in \mathbb{R}^m$  are normal, that is, assume that  $\widehat{\mathbb{P}}_x = \mathcal{N}(\widehat{\mu}_x, \widehat{\Sigma}_x)$  and  $\widehat{\mathbb{P}}_w = \mathcal{N}(\widehat{\mu}_w, \widehat{\Sigma}_w)$ . Thus, the joint nominal distribution  $\widehat{\mathbb{P}} = \widehat{\mathbb{P}}_x \times \widehat{\mathbb{P}}_w$  is also normal, that is,

$$(3.2) \quad \widehat{\mathbb{P}} = \mathcal{N}(\widehat{\mu}, \widehat{\Sigma}) \quad \text{where} \quad \widehat{\mu} = \begin{bmatrix} \widehat{\mu}_x \\ \widehat{\mu}_w \end{bmatrix} \quad \text{and} \quad \widehat{\Sigma} = \begin{bmatrix} \widehat{\Sigma}_x & 0 \\ 0 & \widehat{\Sigma}_w \end{bmatrix}.$$

We highlight that normal distributions are natural candidates for  $\widehat{\mathbb{P}}$ . One reason for this is that the normal distribution has maximum entropy among all distributions with prescribed first- and second-order

moments [14, § 12]. Therefore, it has appeal as the least prejudiced baseline model. Similarly, if the parameter  $x$  in (1.1) is normally distributed, then a normal distribution minimizes the mutual information between  $x$  and the observation  $y$  among all noise distributions with bounded variance [17, Lemma II.2]. In this sense, normally distributed noise renders the observations least informative. Conversely, if the noise in (1.1) is normally distributed, then a normal distribution maximizes the MMSE across all distributions of  $x$  with bounded variance [34, Proposition 15]. In this sense, normally distributed parameters are the hardest to estimate. Using normal nominal distributions thus amounts to adopting a worst-case perspective.

Armed with the fundamental results on normal distributions summarized above, we are now ready to address the dual Wasserstein MMSE estimation problem (3.1) with a normal nominal distribution. In analogy to Section 2, where we proposed the Gelbrich MMSE estimation problem as an easier conservative approximation for the original *primal* estimation problem (1.7), we will now construct an easier conservative approximation for the original *dual* estimation problem (3.1). To this end, we define the restricted ambiguity set

$$\mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}}) = \left\{ \mathbb{Q}_x \times \mathbb{Q}_w \in \mathcal{M}(\mathbb{R}^n) \times \mathcal{M}(\mathbb{R}^m) : \begin{array}{l} \exists \Sigma_x \in \mathbb{S}_+^n, \Sigma_w \in \mathbb{S}_{++}^m \text{ with} \\ \mathbb{Q}_x = \mathcal{N}(\widehat{\mu}_x, \Sigma_x), \mathbb{Q}_w = \mathcal{N}(\widehat{\mu}_w, \Sigma_w), \\ \mathbb{W}(\mathbb{Q}_x, \widehat{\mathbb{P}}_x) \leq \rho_x, \mathbb{W}(\mathbb{Q}_w, \widehat{\mathbb{P}}_w) \leq \rho_w \end{array} \right\}.$$

By construction,  $\mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})$  contains all *normal* distributions  $\mathbb{Q} = \mathbb{Q}_x \times \mathbb{Q}_w$  from within the original Wasserstein ambiguity set  $\mathbb{B}(\widehat{\mathbb{P}})$  that have the same mean vector  $(\widehat{\mu}_x, \widehat{\mu}_w)$  as the nominal distribution  $\widehat{\mathbb{P}} = \widehat{\mathbb{P}}_x \times \widehat{\mathbb{P}}_w$ , and where the covariance matrix of  $\mathbb{Q}_w$  is strictly positive definite. Thus, we have  $\mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}}) \subseteq \mathbb{B}(\widehat{\mathbb{P}})$ . Note also that  $\mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})$  is non-convex because mixtures of normal distributions usually fail to be normal.

By restricting the original Wasserstein ambiguity set  $\mathbb{B}(\widehat{\mathbb{P}})$  to its subset  $\mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})$ , we obtain the following conservative approximation for the dual Wasserstein MMSE estimation problem (3.1).

$$(3.3) \quad \begin{array}{l} \text{maximize} \inf \mathcal{R}(\psi, \mathbb{Q}) \\ \mathbb{Q} \in \mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}}) \quad \psi \in \mathcal{F} \end{array}$$

We will henceforth refer to (3.3) as the dual Wasserstein MMSE estimation problem *over normal priors*. The following main theorem shows that (3.3) is equivalent to a finite convex optimization problem.

**Theorem 3.5** (Dual Wasserstein MMSE estimation problem over normal priors). *Assume that the Wasserstein ambiguity set  $\mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})$  is centered at a normal distribution  $\widehat{\mathbb{P}}$  of the form (3.2). Then, the dual Wasserstein MMSE estimation problem over normal priors (3.3) is equivalent to the finite convex optimization problem*

$$(3.4) \quad \begin{array}{l} \sup \quad \text{Tr} [\Sigma_x - \Sigma_x H^\top (H \Sigma_x H^\top + \Sigma_w)^{-1} H \Sigma_x] \\ \text{s. t.} \quad \Sigma_x \in \mathbb{S}_+^n, \quad \Sigma_w \in \mathbb{S}_{++}^m \\ \text{Tr} [\Sigma_x + \widehat{\Sigma}_x - 2(\widehat{\Sigma}_x^{\frac{1}{2}} \Sigma_x \widehat{\Sigma}_x^{\frac{1}{2}})^{\frac{1}{2}}] \leq \rho_x^2 \\ \text{Tr} [\Sigma_w + \widehat{\Sigma}_w - 2(\widehat{\Sigma}_w^{\frac{1}{2}} \Sigma_w \widehat{\Sigma}_w^{\frac{1}{2}})^{\frac{1}{2}}] \leq \rho_w^2. \end{array}$$

If  $\widehat{\Sigma}_w \succ 0$ , then (3.4) is solvable, and the maximizer denoted by  $(\Sigma_x^*, \Sigma_w^*)$  satisfies  $\Sigma_x^* \succeq \lambda_{\min}(\widehat{\Sigma}_x) I_n$  and  $\Sigma_w^* \succeq \lambda_{\min}(\widehat{\Sigma}_w) I_m$ . Moreover, the supremum of (3.3) is attained by the normal distribution  $\mathbb{Q}^* = \mathbb{Q}_x^* \times \mathbb{Q}_w^*$  defined through  $\mathbb{Q}_x^* = \mathcal{N}(\widehat{\mu}_x, \Sigma_x^*)$  and  $\mathbb{Q}_w^* = \mathcal{N}(\widehat{\mu}_w, \Sigma_w^*)$ .

*Proof of Theorem 3.5.* If  $(x, w)$  is governed by a normal distribution  $\mathbb{Q} \in \mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})$ , then the linear transformation  $(x, y) = (x, Hx + w)$  is also normally distributed by virtue of Proposition 3.2, and the average risk  $\mathcal{R}(\psi, \mathbb{Q})$  is minimized by the Bayesian MMSE estimator  $\psi_{\mathbb{B}}^*(y) = \mathbb{E}_{\mathbb{P}_{x|y}}[x]$ , which is affine due to Proposition 3.3. Thus, in the dual Wasserstein MMSE estimation problem with normal priors, the set  $\mathcal{F}$  of *all* estimators may be

restricted to the set  $\mathcal{A}$  of all *affine* estimators without sacrificing optimality, that is,

$$(3.5) \quad \sup_{\mathbb{Q} \in \mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})} \inf_{\psi \in \mathcal{F}} \mathcal{R}(\psi, \mathbb{Q}) = \sup_{\mathbb{Q} \in \mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})} \inf_{\psi \in \mathcal{A}} \mathcal{R}(\psi, \mathbb{Q}).$$

As the average risk  $\mathcal{R}(\psi, \mathbb{Q})$  of an affine estimator  $\psi \in \mathcal{A}$  simply evaluates the expectation of a quadratic function in  $(x, w)$ , it depends on  $\mathbb{Q}$  only through its first and second moments. Moreover, as  $\mathbb{Q}$  and  $\widehat{\mathbb{P}}$  are normal distributions, their Wasserstein distance coincides with the Gelbrich distance between their mean vectors and covariance matrices; see Proposition 3.4. Thus, the maximization problem over  $\mathbb{Q}$  on the right hand side of (3.5) can be recast as an equivalent maximization problem over the first and second moments of  $\mathbb{Q}$ . Specifically, by the definitions of  $\mathcal{R}(\psi, \mathbb{Q})$  and  $\mathbb{B}_{\phi}(\widehat{\mathbb{P}})$  we find

$$\sup_{\mathbb{Q} \in \mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})} \inf_{\psi \in \mathcal{A}} \mathcal{R}(\psi, \mathbb{Q}) = \begin{cases} \sup_{\Sigma_x, \Sigma_w} \inf_{\substack{A, K \\ K = I_n - AH}} \inf_b \langle K^\top K, \Sigma_x + \widehat{\mu}_x \widehat{\mu}_x^\top \rangle + \langle A^\top A, \Sigma_w + \widehat{\mu}_w \widehat{\mu}_w^\top \rangle + b^\top b \\ \quad - 2\widehat{\mu}_x^\top K^\top A \widehat{\mu}_w - 2b^\top (K \widehat{\mu}_x - A \widehat{\mu}_w) \\ \text{s. t. } \mathbf{G}((\widehat{\mu}_x, \Sigma_x), (\widehat{\mu}_x, \widehat{\Sigma}_x)) \leq \rho_x, \mathbf{G}((\widehat{\mu}_w, \Sigma_w), (\widehat{\mu}_w, \widehat{\Sigma}_w)) \leq \rho_w \\ \Sigma_x \succeq 0, \Sigma_w \succ 0, \end{cases}$$

where the auxiliary decision variable  $K = I_n - AH$  has been introduced to simplify the objective function. The innermost minimization problem over  $b$  constitutes an unconstrained (strictly) convex quadratic program that has the unique optimal solution  $b = K \widehat{\mu}_x - A \widehat{\mu}_w$ . Substituting this minimizer back into the objective function of the above problem and recalling the definition of the Gelbrich distance then yields

$$(3.6) \quad \sup_{\mathbb{Q} \in \mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})} \inf_{\psi \in \mathcal{A}} \mathcal{R}(\psi, \mathbb{Q}) = \begin{cases} \sup_{\Sigma_x, \Sigma_w} \inf_{\substack{A, K \\ K = I_n - AH}} \langle K^\top K, \Sigma_x \rangle + \langle A^\top A, \Sigma_w \rangle \\ \text{s. t. } \text{Tr} [\Sigma_x + \widehat{\Sigma}_x - 2(\widehat{\Sigma}_x^{\frac{1}{2}} \Sigma_x \widehat{\Sigma}_x^{\frac{1}{2}})^{\frac{1}{2}}] \leq \rho_x^2 \\ \text{Tr} [\Sigma_w + \widehat{\Sigma}_w - 2(\widehat{\Sigma}_w^{\frac{1}{2}} \Sigma_w \widehat{\Sigma}_w^{\frac{1}{2}})^{\frac{1}{2}}] \leq \rho_w^2 \\ \Sigma_x \succeq 0, \Sigma_w \succ 0. \end{cases}$$

By using the equality  $K = I_n - AH$  to eliminate  $K$ , the inner minimization problem in (3.6) can be reformulated as an unconstrained quadratic program in  $A$ . As  $\Sigma_w \succ 0$ , this quadratic program is strictly convex, and an elementary calculation reveals that its unique optimal solution is given by

$$A^* = \Sigma_x H^\top (H \Sigma_x H^\top + \Sigma_w)^{-1}.$$

Substituting  $A^*$  as well as the corresponding auxiliary decision variable  $K^* = I_n - A^* H$  into the objective function of (3.6) finally yields the postulated convex program (3.4).

Assume now that  $\widehat{\Sigma}_w \succ 0$ , and define

$$\mathcal{S}_x = \left\{ \Sigma_x \in \mathbb{S}_+^n : \mathbf{G}((\widehat{\mu}_x, \Sigma_x), (\widehat{\mu}_x, \widehat{\Sigma}_x)) \leq \rho_x \right\} \quad \text{and} \quad \mathcal{S}_w = \left\{ \Sigma_w \in \mathbb{S}_+^m : \mathbf{G}((\widehat{\mu}_w, \Sigma_w), (\widehat{\mu}_w, \widehat{\Sigma}_w)) \leq \rho_w \right\}.$$

Equations (3.5) and (3.6) imply that

$$(3.7) \quad \begin{aligned} \sup_{\mathbb{Q} \in \mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})} \inf_{\psi \in \mathcal{F}} \mathcal{R}(\psi, \mathbb{Q}) &\leq \sup_{\Sigma_x \in \mathcal{S}_x} \sup_{\Sigma_w \in \mathcal{S}_w} \inf_{\substack{A, K \\ K = I_n - AH}} \langle K^\top K, \Sigma_x \rangle + \langle A^\top A, \Sigma_w \rangle \\ &= \sup_{\substack{\Sigma_x \in \mathcal{S}_x \\ \Sigma_x \geq \lambda_{\min}(\widehat{\Sigma}_x) I_n}} \sup_{\substack{\Sigma_w \in \mathcal{S}_w \\ \Sigma_w \geq \lambda_{\min}(\widehat{\Sigma}_w) I_m}} \inf_{\substack{A, K \\ K = I_n - AH}} \langle K^\top K, \Sigma_x \rangle + \langle A^\top A, \Sigma_w \rangle, \end{aligned}$$

where the inequality holds because we relax the requirement that  $\Sigma_w$  be strictly positive definite, and the equality follows from applying Lemma A.5 consecutively to each of the two maximization problems. If  $\widehat{\Sigma}_w \succ 0$ ,

then problem (3.7) constitutes a restriction of (3.6) and therefore provides also a lower bound on the dual Wasserstein MMSE estimation problem. In summary, we thus have

$$(3.8) \quad \sup_{\mathbb{Q} \in \mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})} \inf_{\psi \in \mathcal{F}} \mathcal{R}(\psi, \mathbb{Q}) = \begin{cases} \sup_{\Sigma_x, \Sigma_w} \inf_{\substack{A, K \\ K = I_n - AH}} \langle K^\top K, \Sigma_x \rangle + \langle A^\top A, \Sigma_w \rangle \\ \text{s. t.} \quad \text{Tr} [\Sigma_x + \widehat{\Sigma}_x - 2(\widehat{\Sigma}_x^{\frac{1}{2}} \Sigma_x \widehat{\Sigma}_x^{\frac{1}{2}})^{\frac{1}{2}}] \leq \rho_x^2 \\ \text{Tr} [\Sigma_w + \widehat{\Sigma}_w - 2(\widehat{\Sigma}_w^{\frac{1}{2}} \Sigma_w \widehat{\Sigma}_w^{\frac{1}{2}})^{\frac{1}{2}}] \leq \rho_w^2 \\ \Sigma_x \succeq \lambda_{\min}(\widehat{\Sigma}_x) I_n, \quad \Sigma_w \succeq \lambda_{\min}(\widehat{\Sigma}_w) I_m. \end{cases}$$

This reasoning implies that if  $\widehat{\Sigma}_w \succ 0$ , then the constraints  $\Sigma_x \succeq \lambda_{\min}(\widehat{\Sigma}_x) I_n$  and  $\Sigma_w \succeq \lambda_{\min}(\widehat{\Sigma}_w) I_m$  can be appended to problem (3.6) and, consequently, to problem (3.4) without altering their common optimal value. Problem (3.4) with the additional constraints  $\Sigma_x \succeq \lambda_{\min}(\widehat{\Sigma}_x) I_n$  and  $\Sigma_w \succeq \lambda_{\min}(\widehat{\Sigma}_w) I_m$  has a continuous objective function over a compact feasible set and is thus solvable. Any of its optimal solutions is also optimal in problem (3.4), which has no redundant constraints. Thus, problem (3.4) is solvable.

It remains to show that  $\mathbb{Q}^*$  as constructed in the theorem statement is optimal in (3.3). The feasibility of  $(\Sigma_x^*, \Sigma_w^*)$  in (3.4) implies that  $\mathbb{Q}^* \in \mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})$ , and thus  $\mathbb{Q}^*$  is feasible in (3.3). Moreover, we have

$$(3.9) \quad \sup_{\mathbb{Q} \in \mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})} \inf_{\psi \in \mathcal{F}} \mathcal{R}(\psi, \mathbb{Q}) \geq \inf_{\psi \in \mathcal{F}} \mathcal{R}(\psi, \mathbb{Q}^*) = \text{Tr} [\Sigma_x^* - \Sigma_x^* H^\top (H \Sigma_x^* H^\top + \Sigma_w^*)^{-1} H \Sigma_x^*],$$

where the equality follows from elementary algebra, recalling that the affine estimator  $\psi(y) = A^* y + b^*$  with

$$A^* = \Sigma_x^* H^\top (H \Sigma_x^* H^\top + \Sigma_w^*)^{-1} \quad \text{and} \quad b^* = \mu_x - A^* (H \widehat{\mu}_x + \widehat{\mu}_w)$$

is the Bayesian MMSE estimator for the normal distribution  $\mathbb{Q}^*$ . As the right hand side of (3.9) coincides with the maximum of (3.4) and as problem (3.4) is equivalent to the dual Wasserstein MMSE estimation problem (3.3) over normal priors, we may thus conclude that the inequality in (3.9) is tight. Thus, we find

$$\sup_{\mathbb{Q} \in \mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})} \inf_{\psi \in \mathcal{F}} \mathcal{R}(\psi, \mathbb{Q}) = \inf_{\psi \in \mathcal{F}} \mathcal{R}(\psi, \mathbb{Q}^*),$$

which in turn implies that  $\mathbb{Q}^*$  is optimal in (3.3). This observation completes the proof.  $\square$

**Remark 3.6** (Singular covariance matrices). *A nonlinear SDP akin to (3.4) has been derived in [64] under the stronger assumption that the covariance matrix of the nominal distribution  $\widehat{\mathbb{P}}$  is non-degenerate, which implies that  $\widehat{\Sigma}_x \succ 0$  and  $\widehat{\Sigma}_w \succ 0$ . However, the weaker condition  $\widehat{\Sigma}_w \succ 0$  is sufficient to ensure that the matrix inversion in the objective function of problem (3.4) is well-defined. Therefore, Theorem 3.5 remains valid if the nominal covariance matrix  $\widehat{\Sigma}_x$  is singular, which occurs in many applications. On the other hand, it is common to require that  $\widehat{\Sigma}_w = \sigma^2 I_m$  for some  $\sigma > 0$ , see, e.g., [11].*

Corollary 3.7 below asserts that the convex program (3.4) admits a canonical linear SDP reformulation. The proof is omitted as it relies on standard Schur complement arguments familiar from the proof of Corollary 2.8.

**Corollary 3.7** (SDP reformulation). *Assume that the Wasserstein ambiguity set  $\mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})$  is centered at a normal distribution  $\widehat{\mathbb{P}}$  of the form (3.2) with noise covariance matrix  $\widehat{\Sigma}_w \succ 0$ . Then, the dual Wasserstein*

MMSE estimation problem (3.3) over normal priors is equivalent to the SDP

$$\begin{aligned}
(3.10) \quad & \max \quad \text{Tr} [\Sigma_x] - \text{Tr} [U] \\
& \text{s. t.} \quad \Sigma_x \in \mathbb{S}_+^n, \Sigma_w \in \mathbb{S}_+^m, V_x \in \mathbb{S}_+^n, V_w \in \mathbb{S}_+^m, U \in \mathbb{S}_+^n \\
& \begin{bmatrix} \widehat{\Sigma}_x^{\frac{1}{2}} \Sigma_x \widehat{\Sigma}_x^{\frac{1}{2}} & V_x \\ V_x & I_n \end{bmatrix} \succeq 0, \quad \begin{bmatrix} \widehat{\Sigma}_w^{\frac{1}{2}} \Sigma_w \widehat{\Sigma}_w^{\frac{1}{2}} & V_w \\ V_w & I_m \end{bmatrix} \succeq 0 \\
& \text{Tr} [\Sigma_x + \widehat{\Sigma}_x - 2V_x] \leq \rho_x^2, \quad \text{Tr} [\Sigma_w + \widehat{\Sigma}_w - 2V_w] \leq \rho_w^2 \\
& \begin{bmatrix} U & \Sigma_x H^\top \\ H \Sigma_x & H \Sigma_x H^\top + \Sigma_w \end{bmatrix} \succeq 0, \quad \Sigma_x \succeq \lambda_{\min}(\widehat{\Sigma}_x) I_n, \quad \Sigma_w \succeq \lambda_{\min}(\widehat{\Sigma}_w) I_m.
\end{aligned}$$

We emphasize that the lower bounds on  $\Sigma_x$  and  $\Sigma_w$  are redundant but have been made explicit in (3.10).

#### 4. NASH EQUILIBRIUM AND OPTIMALITY OF AFFINE ESTIMATORS

If  $\widehat{\mathbb{P}}$  is a normal distribution of the form (3.2), then we have

$$(4.1) \quad \inf_{\psi \in \mathcal{A}} \sup_{\mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}})} \mathcal{R}(\psi, \mathbb{Q}) \geq \inf_{\psi \in \mathcal{F}} \sup_{\mathbb{Q} \in \mathbb{B}(\widehat{\mathbb{P}})} \mathcal{R}(\psi, \mathbb{Q}) \geq \sup_{\mathbb{Q} \in \mathbb{B}(\widehat{\mathbb{P}})} \inf_{\psi \in \mathcal{F}} \mathcal{R}(\psi, \mathbb{Q}) \geq \sup_{\mathbb{Q} \in \mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})} \inf_{\psi \in \mathcal{F}} \mathcal{R}(\psi, \mathbb{Q}),$$

where the first inequality follows from the inclusions  $\mathcal{A} \subseteq \mathcal{F}$  and  $\mathbb{B}(\widehat{\mathbb{P}}) \subseteq \mathbb{G}(\widehat{\mathbb{P}})$ , the second inequality exploits weak duality, and the last inequality holds due to the inclusion  $\mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}}) \subseteq \mathbb{B}(\widehat{\mathbb{P}})$ . Note that the left-most minimax problem is the Gelbrich MMSE estimation problem (2.1) studied in Section 2, and the right-most maximin problem is the dual Wasserstein MMSE estimation problem (3.3) over normal priors studied in Section 3. We also highlight that these restricted primal and dual estimation problems sandwich the original Wasserstein estimation problems (1.7) and (3.1), which coincide with the second and third problems in (4.1), respectively. The following theorem asserts that all inequalities in (4.1) actually collapse to equalities.

**Theorem 4.1** (Sandwich theorem). *If  $\widehat{\mathbb{P}}$  is a normal distribution of the form (3.2), then the optimal values of the restricted primal and dual estimation problems (2.1) and (3.3) coincide, i.e.,*

$$\inf_{\psi \in \mathcal{A}} \sup_{\mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}})} \mathcal{R}(\psi, \mathbb{Q}) = \sup_{\mathbb{Q} \in \mathbb{B}_{\mathcal{N}}(\widehat{\mathbb{P}})} \inf_{\psi \in \mathcal{F}} \mathcal{R}(\psi, \mathbb{Q}).$$

*Proof of Theorem 4.1.* By Theorem 2.7, the Gelbrich MMSE estimation problem (2.1) can be expressed as

$$\inf_{\psi \in \mathcal{A}} \sup_{\mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}})} \mathcal{R}(\psi, \mathbb{Q}) = \left\{ \begin{array}{ll} \inf_{\substack{A, K \\ K = I_n - AH}} & \inf_{\substack{\gamma_x, \gamma_w \\ \gamma_x I_n \succ K^\top K \\ \gamma_w I_m \succ A^\top A}} \gamma_x (\rho_x^2 - \text{Tr} [\widehat{\Sigma}_x]) + \gamma_x^2 \langle (\gamma_x I_n - K^\top K)^{-1}, \widehat{\Sigma}_x \rangle \\ & + \gamma_w (\rho_w^2 - \text{Tr} [\widehat{\Sigma}_w]) + \gamma_w^2 \langle (\gamma_w I_m - A^\top A)^{-1}, \widehat{\Sigma}_w \rangle, \end{array} \right.$$

where the auxiliary variable  $K = I_n - AH$  has been introduced to highlight the problem's symmetries. Next, we introduce the feasible sets

$$\mathcal{S}_x = \left\{ \Sigma_x \in \mathbb{S}_+^n : \text{Tr} [\Sigma_x + \widehat{\Sigma}_x - 2(\widehat{\Sigma}_x^{\frac{1}{2}} \Sigma_x \widehat{\Sigma}_x^{\frac{1}{2}})^{\frac{1}{2}}] \leq \rho_x^2 \right\}$$

and

$$\mathcal{S}_w = \left\{ \Sigma_w \in \mathbb{S}_+^m : \text{Tr} [\Sigma_w + \widehat{\Sigma}_w - 2(\widehat{\Sigma}_w^{\frac{1}{2}} \Sigma_w \widehat{\Sigma}_w^{\frac{1}{2}})^{\frac{1}{2}}] \leq \rho_w^2 \right\},$$



both of which are convex and compact by virtue of Lemma A.6. Using Proposition A.4 (i) to reformulate the inner minimization problem over  $\gamma_x$  and  $\gamma_w$ , we then obtain

$$\begin{aligned}
\inf_{\psi \in \mathcal{A}} \sup_{\mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}})} \mathcal{R}(\psi, \mathbb{Q}) &= \inf_{\substack{A, K \\ K=I_n - AH}} \sup_{\substack{\Sigma_x \in \mathcal{S}_x \\ \Sigma_w \in \mathcal{S}_w}} \langle K^\top K, \Sigma_x \rangle + \langle A^\top A, \Sigma_w \rangle \\
&= \sup_{\Sigma_x \in \mathcal{S}_x} \sup_{\Sigma_w \in \mathcal{S}_w} \inf_{\substack{A, K \\ K=I_n - AH}} \langle K^\top K, \Sigma_x \rangle + \langle A^\top A, \Sigma_w \rangle \\
&= \sup_{\substack{\Sigma_x \in \mathcal{S}_x \\ \Sigma_x \succeq \lambda_{\min}(\widehat{\Sigma}_x) I_n}} \sup_{\substack{\Sigma_w \in \mathcal{S}_w \\ \Sigma_w \succeq \lambda_{\min}(\widehat{\Sigma}_w) I_m}} \inf_{\substack{A, K \\ K=I_n - AH}} \langle K^\top K, \Sigma_x \rangle + \langle A^\top A, \Sigma_w \rangle \\
&= \sup_{\mathbb{Q} \in \mathbb{B}_N(\widehat{\mathbb{P}})} \inf_{\psi \in \mathcal{F}} \mathcal{R}(\psi, \mathbb{Q}),
\end{aligned}$$

where the second equality holds due to Sion's minimax theorem [65], and the third equality follows from Lemma A.5 applied twice separately to  $\Sigma_x$  and  $\Sigma_w$ . The last equality has already been derived in the proof of Theorem 3.5; see Equation (3.8). Thus, the claim follows.  $\square$

Theorem 4.1 suggests that solving any of the restricted estimation problems is tantamount to solving both original primal and dual estimation problems. This intuition is formalized in the following corollary.

**Corollary 4.2** (Nash equilibrium). *If  $\widehat{\mathbb{P}}$  is a normal distribution of the form (3.2) with  $\widehat{\Sigma}_w \succ 0$ , then the affine estimator  $\psi^*$  that solves (2.1) is optimal in the primal Wasserstein MMSE estimation problem (1.7), while the normal distribution  $\mathbb{Q}^*$  that solves (3.3) is optimal in the dual Wasserstein MMSE estimation problem (3.1). Moreover,  $\psi^*$  and  $\mathbb{Q}^*$  form a Nash equilibrium for the game between the statistician and nature, that is,*

$$(4.2) \quad \mathcal{R}(\psi^*, \mathbb{Q}) \leq \mathcal{R}(\psi^*, \mathbb{Q}^*) \leq \mathcal{R}(\psi, \mathbb{Q}^*) \quad \forall \psi \in \mathcal{F}, \mathbb{Q} \in \mathbb{B}(\widehat{\mathbb{P}}).$$

*Proof of Corollary 4.2.* As  $\widehat{\Sigma}_w \succ 0$ , the Gelbrich MMSE estimation problem (2.1) is solved by the affine estimator  $\psi^*$  defined in Theorem 2.7, and the dual Wasserstein MMSE estimation problem (3.1) over normal priors is solved by the normal distribution  $\mathbb{Q}^*$  defined in Theorem 3.5. Thus, we have

$$\mathcal{R}(\psi^*, \mathbb{Q}^*) \geq \inf_{\psi \in \mathcal{F}} \mathcal{R}(\psi, \mathbb{Q}^*) = \max_{\mathbb{Q} \in \mathbb{B}_N(\widehat{\mathbb{P}})} \inf_{\psi \in \mathcal{F}} \mathcal{R}(\psi, \mathbb{Q}) = \min_{\psi \in \mathcal{A}} \sup_{\mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}})} \mathcal{R}(\psi, \mathbb{Q}) = \sup_{\mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}})} \mathcal{R}(\psi^*, \mathbb{Q}) \geq \mathcal{R}(\psi^*, \mathbb{Q}^*),$$

where the three equalities follow from the definition of  $\mathbb{Q}^*$ , Theorem 4.1 and the definition of  $\psi^*$ , respectively. As the left and the right hand sides of the above expression coincide, we may then conclude that

$$\mathcal{R}(\psi^*, \mathbb{Q}^*) = \mathcal{R}(\psi^*, \mathbb{Q}) \leq \mathcal{R}(\psi^*, \mathbb{Q}^*) \leq \mathcal{R}(\psi, \mathbb{Q}^*) \quad \forall \psi \in \mathcal{F}, \mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}}).$$

Moreover, as  $\mathbb{B}(\widehat{\mathbb{P}}) \subseteq \mathbb{G}(\widehat{\mathbb{P}})$ , the above relation implies (4.2).

It remains to be shown that  $\psi^*$  and  $\mathbb{Q}^*$  solve the primal and dual Wasserstein MMSE estimation problems (1.7) and (3.1), respectively. As for  $\psi^*$ , we have

$$\sup_{\mathbb{Q} \in \mathbb{B}(\widehat{\mathbb{P}})} \mathcal{R}(\psi^*, \mathbb{Q}) \leq \sup_{\mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}})} \mathcal{R}(\psi^*, \mathbb{Q}) = \inf_{\psi \in \mathcal{A}} \sup_{\mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}})} \mathcal{R}(\psi, \mathbb{Q}) = \inf_{\psi \in \mathcal{F}} \sup_{\mathbb{Q} \in \mathbb{B}(\widehat{\mathbb{P}})} \mathcal{R}(\psi, \mathbb{Q}),$$

where the first equality follows from the definition of  $\psi^*$ , while the second equality exploits Theorem 4.1, which implies that all inequalities in (4.1) are in fact equalities. This reasoning shows that  $\psi^*$  is optimal in (1.7). The optimality of  $\mathbb{Q}^*$  in (3.1) can be proved similarly. Details are omitted for brevity.  $\square$

Corollary 4.2 implies that  $\psi^*$  can be viewed as a Bayesian estimator for the least favorable prior  $\mathbb{Q}^*$  and that  $\mathbb{Q}^*$  represents a worst-case distribution for the optimal estimator  $\psi^*$ . Next, we will argue that  $\psi^*$  can not only be constructed from the solution of the convex program (2.2), which is equivalent to the Gelbrich MMSE estimation problem (2.1), but also from the solution of the convex program (3.4), which is equivalent to the

dual MMSE estimation problem (3.3) over normal priors. This alternative construction is useful because problem (3.4) is amenable to highly efficient first-order methods to be derived in Section 6.

**Corollary 4.3** (Dual construction of the optimal estimator). *If  $\widehat{\mathbb{P}}$  is a normal distribution of the form (3.2) with  $\widehat{\Sigma}_w \succ 0$ , and  $(\Sigma_x^*, \Sigma_w^*)$  is a maximizer of (3.3), then the affine estimator  $\psi^*(y) = A^*y + b^*$  with*

$$(4.3) \quad A^* = \Sigma_x^* H^\top (H \Sigma_x^* H^\top + \Sigma_w^*)^{-1} \quad \text{and} \quad b^* = \widehat{\mu}_x - A^*(H \widehat{\mu}_x + \widehat{\mu}_w)$$

*solves the Wasserstein MMSE estimation problem (1.7).*

*Proof of Corollary 4.3.* Define  $\psi^*$  as the affine estimator that solves (2.1) and  $\mathbb{Q}^*$  as the normal distribution that solves (3.3). By Corollary 4.2, the second inequality in (4.2) holds for all admissible estimators  $\psi \in \mathcal{F}$ , which implies that  $\psi^* \in \arg \min_{\psi \in \mathcal{F}} \mathcal{R}(\psi, \mathbb{Q}^*)$ , that is,  $\psi^*$  solves the Bayesian MMSE estimation problem corresponding to  $\mathbb{Q}^*$ . As any Bayesian MMSE estimator satisfies  $\psi^*(y) = \mathbb{E}_{\mathbb{Q}_x^*|y}[x]$  for  $\mathbb{Q}^*$ -almost all  $y$  and as  $\Sigma_w^* \succ 0$ , we may use the known formulas for conditional normal distributions to conclude that the unique affine Bayesian MMSE estimator for  $\mathbb{Q}^*$  is of the form  $\psi^*(y) = A^*y + b^*$  with parameters defined as in (4.3).  $\square$

**Remark 4.4** (Non-normal nominal distributions). *The results of this section imply that the optimal values of the finite convex programs (2.2) and (3.4) typically coincide even if the nominal distribution fails to be normal. To see this, assume that  $\widehat{\mathbb{P}} = \widehat{\mathbb{P}}_x \times \widehat{\mathbb{P}}_w$ , where  $\widehat{\mathbb{P}}_x$  and  $\widehat{\mathbb{P}}_w$  are arbitrary signal and noise distributions with finite mean vectors  $\widehat{\mu}_x$  and  $\widehat{\mu}_w$  and covariance matrices  $\widehat{\Sigma}_x$  and  $\widehat{\Sigma}_w$ , respectively, and denote by  $\widehat{\mathbb{P}}' = \widehat{\mathbb{P}}'_x \times \widehat{\mathbb{P}}'_w$  the normal distribution with the same first and second moments as  $\widehat{\mathbb{P}}$ . If  $\widehat{\Sigma}_w \succ 0$ , then the optimal values of (2.2) and (3.4) are equal by virtue of the Theorems 2.7, 3.5 and 4.1 applied to  $\widehat{\mathbb{P}}'$ . As (2.2) and (3.4) depend only on the first and second moments of  $\widehat{\mathbb{P}}'$ , their optimal values do not change if  $\widehat{\mathbb{P}}'$  is replaced with  $\widehat{\mathbb{P}}$ . Therefore, the optimal values of (2.2) and (3.4) coincide for any nominal distribution  $\widehat{\mathbb{P}}$  with finite first and second moments provided that  $\widehat{\Sigma}_w \succ 0$ . In this case, however, the minimum of the Gelbrich MMSE estimation problem (2.1) may strictly exceed the maximum of the dual Wasserstein MMSE estimation problem (3.3) over normal priors. Note that in this case the ambiguity set  $\mathbb{B}_N(\widehat{\mathbb{P}})$  may even be empty. Moreover, while typically suboptimal for the original Wasserstein MMSE estimation problem (1.7), the affine estimator constructed in Corollary 4.3 remains optimal for the Gelbrich MMSE estimation problem (2.1) even if  $\widehat{\mathbb{P}}$  fails to be normal.*

## 5. ELLIPTICAL NOMINAL DISTRIBUTIONS

We will now show that the results of Sections 2–4 remain valid if  $\widehat{\mathbb{P}}$  is an arbitrary elliptical (but maybe non-normal) distribution. To this end, we first review some basic results on elliptical distributions.

**Definition 5.1** (Elliptical distributions). *The distribution  $\mathbb{P}$  of  $\xi \in \mathbb{R}^d$  is called elliptical if the characteristic function  $\Phi_{\mathbb{P}}(t) = \mathbb{E}_{\mathbb{P}}[\exp(it^\top \xi)]$  of  $\mathbb{P}$  is given by  $\Phi_{\mathbb{P}}(t) = \exp(it^\top \mu) \phi(t^\top S t)$  for some location parameter  $\mu \in \mathbb{R}^d$ , dispersion matrix  $S \in \mathbb{S}_+^d$  and characteristic generator  $\phi: \mathbb{R}_+ \rightarrow \mathbb{R}$ . In this case we write  $\mathbb{P} = \mathcal{E}_\phi^d(\mu, S)$ . The class of all  $d$ -dimensional elliptical distributions with characteristic generator  $\phi$  is denoted by  $\mathcal{E}_\phi^d$ .*

The class of elliptical distributions was introduced in [44] with the aim to generalize the family of normal distributions, which are obtained by setting the characteristic generator to  $\phi(u) = e^{-u/2}$ . We emphasize that, unlike the moment-generating function  $M_{\mathbb{P}}(t) = \mathbb{E}_{\mathbb{P}}[\exp(t^\top \xi)]$ , the characteristic function  $\Phi_{\mathbb{P}}(t)$  is always finite for all  $t \in \mathbb{R}^d$  even if some moments of  $\mathbb{P}$  do not exist. Thus, Definition 5.1 is general enough to cover also heavy-tailed distributions with non-zero tail dependence coefficients [37]. Examples of elliptical distributions include the Laplace, logistic and  $t$ -distribution etc. Useful theoretical properties of elliptical distributions are discussed in [10, 27]. We also highlight that elliptical distributions are central to a wide spectrum of diverse applications ranging from genomics [60] and medical imaging [61] to finance [39, § 6.2.1], to name a few.

If the dispersion matrix  $S \in \mathbb{S}_+^d$  has rank  $r$ , then there exists  $\Lambda \in \mathbb{R}^{d \times r}$  with  $S = \Lambda \Lambda^\top$ , and there exists a generalized inverse  $\Lambda^{-1} \in \mathbb{R}^{r \times d}$  with  $\Lambda^{-1} \Lambda = I_r = \Lambda^\top (\Lambda^{-1})^\top$ . One easily verifies that if  $\xi \in \mathbb{R}^d$  follows an elliptical distribution  $\mathbb{P} = \mathcal{E}_\phi^d(\mu, S)$ , then  $\tilde{\xi} = \Lambda^{-1}(\xi - \mu) \in \mathbb{R}^r$  follows the spherically symmetric distribution  $\tilde{\mathbb{P}} = \mathcal{E}_\phi^d(0, I_r)$  with characteristic function  $\Phi_{\tilde{\mathbb{P}}}(t) = \phi(\|t\|^2)$ . Thus, the choice of the characteristic generator  $\phi$  is constrained by the implicit condition that  $\phi(\|t\|^2)$  must be an admissible characteristic function. For instance, the normalization of probability distributions necessitates that  $\phi(0) = 1$ , while the dominated convergence theorem implies that  $\phi$  must be continuous etc. As any distribution is uniquely determined by its characteristic function, and as  $\phi(\|t\|^2)$  depends only on the norm of  $t$ , the spherical distribution  $\tilde{\mathbb{P}}$  is indeed invariant under rotations. This implies that  $\mathbb{E}_{\tilde{\mathbb{P}}}[\tilde{\xi}] = 0$  and, via the linearity of the expectation, that  $\mathbb{E}_{\mathbb{P}}[\xi] = \mu$  provided that  $\tilde{\xi}$  and  $\xi$  are integrable, respectively. Thus, the location parameter  $\mu$  of an elliptical distribution coincides with its mean vector whenever the mean exists. By the definition of the characteristic function, the covariance matrix of  $\tilde{\mathbb{P}}$ , if it exists, can be expressed as

$$\tilde{\Sigma} = -\nabla_t^2 \Phi_{\tilde{\mathbb{P}}}(t)|_{t=0} = -\nabla_t^2 \phi(\|t\|^2)|_{t=0} = -2\phi'(0)I_r,$$

where  $\phi'(0)$  denotes the right derivative of  $\phi(u)$  at  $u = 0$ . Hence,  $\tilde{\Sigma}$  exists if and only if  $\phi'(0)$  exists and is finite. Similarly, the covariance matrix of  $\mathbb{P}$  is given by  $\Sigma = -2\phi'(0)S$ , if it exists [10, Theorem 4]. Below we will focus on elliptical distributions with finite first- and second-order moments (*i.e.*, we will only consider characteristic generators with  $|\phi'(0)| < \infty$ ), and we will assume that  $\phi'(0) = -\frac{1}{2}$ , which ensures that the dispersion matrix  $S$  equals the covariance matrix  $\Sigma$ . The latter assumption does not restrict generality. In fact, changing the characteristic generator to  $\phi(\frac{-u}{2\phi'(0)})$  and the dispersion matrix to  $-2\phi'(0)S$  has no impact on the elliptical distribution  $\mathbb{P}$  but matches the dispersion matrix  $S$  with the covariance matrix  $\Sigma$ .

The elliptical distributions inherit many desirable properties from the normal distributions but are substantially more expressive as they include also heavy- and light-tailed distributions. For example, any class of elliptical distributions with a common characteristic generator is closed under affine transformations and affine conditional expectations. Moreover, the Wasserstein distance between two elliptical distributions with the same characteristic generator equals the Gelbrich distance between their mean vectors and covariance matrices [31, Theorem 2.4]. Thus, the Propositions 3.2, 3.3 and 3.4 readily extend from the class of normal distributions to *any* class of elliptical distributions that share the same characteristic generator.

The above discussion suggests that the results of Sections 2–4 carry over almost verbatim to MMSE estimation problems involving elliptical nominal distributions. In the following we will therefore assume that

$$(5.1) \quad \hat{\mathbb{P}} = \mathcal{E}_\phi^{n+m}(\hat{\mu}, \hat{\Sigma}) \quad \text{with} \quad \hat{\mu} = \begin{bmatrix} \hat{\mu}_x \\ \hat{\mu}_w \end{bmatrix} \quad \text{and} \quad \hat{\Sigma} = \begin{bmatrix} \hat{\Sigma}_x & 0 \\ 0 & \hat{\Sigma}_w \end{bmatrix},$$

where  $\phi$  denotes a prescribed characteristic generator. As the class of all elliptical distributions with characteristic generator  $\phi$  is closed under affine transformations, the marginal distributions  $\hat{\mathbb{P}}_x$  and  $\hat{\mathbb{P}}_w$  of  $x$  and  $w$  under  $\hat{\mathbb{P}}$  are also elliptical distributions with the same characteristic generator  $\phi$ .

Note that while the signal  $x$  and the noise  $w$  are uncorrelated under  $\hat{\mathbb{P}}$  irrespective of  $\phi$ , they fail to be independent unless  $\hat{\mathbb{P}}$  is a normal distribution. When working with generic elliptical nominal distributions, we must therefore abandon any independence assumptions. Otherwise, the ambiguity set would be empty for small radii  $\rho_x$  and  $\rho_w$ . This insight prompts us to redefine the Wasserstein ambiguity set as

$$(5.2) \quad \mathbb{B}(\hat{\mathbb{P}}) = \left\{ \mathbb{Q} \in \mathcal{M}(\mathbb{R}^{n+m}) : \mathbb{E}_{\mathbb{Q}}[xw^\top] = \mathbb{E}_{\mathbb{Q}}[x] \cdot \mathbb{E}_{\mathbb{Q}}[w]^\top, \mathbb{W}(\mathbb{Q}_x, \hat{\mathbb{P}}_x) \leq \rho_x, \mathbb{W}(\mathbb{Q}_w, \hat{\mathbb{P}}_w) \leq \rho_w \right\},$$

which relaxes the independence condition in (1.8) and merely requires  $x$  and  $w$  to be uncorrelated. When using the new ambiguity set (5.2) to model the distributional uncertainty, we can again compute a Nash equilibrium between the statistician and nature by solving a tractable convex optimization problem.

**Theorem 5.2** (Elliptical distributions). *Assume that  $\widehat{\mathbb{P}}$  is an elliptical distribution of the form (5.1) with characteristic generator  $\phi$  and noise covariance matrix  $\widehat{\Sigma}_w \succ 0$ , and define the ambiguity set  $\mathbb{B}(\widehat{\mathbb{P}})$  as in (5.2). If  $(\Sigma_x^*, \Sigma_w^*)$  solves the finite convex program (3.3), then the affine estimator  $\psi^*(y) = A^*y + b^*$  with*

$$A^* = \Sigma_x^* H^\top (H \Sigma_x^* H^\top + \Sigma_w^*)^{-1} \quad \text{and} \quad b^* = \widehat{\mu}_x - A^*(H \widehat{\mu}_x + \widehat{\mu}_w)$$

*solves the Wasserstein MMSE estimation problem (1.7), while the elliptical distribution*

$$\mathbb{Q}^* = \mathcal{E}_\phi^{n+m}(\widehat{\mu}, \Sigma^*) \quad \text{with} \quad \Sigma^* = \begin{bmatrix} \Sigma_x^* & 0 \\ 0 & \Sigma_w^* \end{bmatrix}$$

*solves the dual Wasserstein MMSE estimation problem (3.1). Moreover,  $\psi^*$  and  $\mathbb{Q}^*$  form a Nash equilibrium for the game between the statistician and nature, that is,*

$$\mathcal{R}(\psi^*, \mathbb{Q}) \leq \mathcal{R}(\psi^*, \mathbb{Q}^*) \leq \mathcal{R}(\psi, \mathbb{Q}^*) \quad \forall \psi \in \mathcal{F}, \mathbb{Q} \in \mathbb{B}(\widehat{\mathbb{P}}).$$

*Proof of Theorem 5.2.* The proof replicates the arguments used to establish Theorems 2.7, 3.5 and 4.1 as well as Corollary 4.2 with obvious minor modifications. Details are omitted for brevity.  $\square$

Theorem 5.2 asserts that the optimal estimator depends only on the first and second moments of the nominal distribution  $\widehat{\mathbb{P}}$  but *not* on its characteristic generator. Whether  $\widehat{\mathbb{P}}$  displays heavier or lighter tails than a normal distribution has therefore no impact on the prediction of the signal. Note, however, that the characteristic generator of  $\widehat{\mathbb{P}}$  determines the shape of the least favorable prior.

## 6. NUMERICAL SOLUTION OF WASSERSTEIN MMSE ESTIMATION PROBLEMS

By Corollaries 2.8 and 3.7, the primal and dual Wasserstein MMSE estimation problems (1.7) and (3.1) can be addressed with off-the-shelf SDP solvers. Unfortunately, however, general-purpose interior-point methods quickly run out of memory when the signal dimension  $n$  and the noise dimension  $m$  grow. It is therefore expedient to look for customized first-order algorithms that can handle larger problem instances.

In this section we develop a Frank-Wolfe method for the nonlinear SDP (3.4), which is equivalent to the dual Wasserstein MMSE estimation problem (3.1). This approach is meaningful because any solution to (3.4) allows us to construct both an optimal estimator as well as a least favorable prior that form a Nash equilibrium in the sense of Corollary 4.2; see also Corollary 4.3. Addressing the nonlinear SDP (3.4) directly with a Frank-Wolfe method has great promise because the subproblems that identify the local search directions can be shown to admit quasi-closed form solutions and can therefore be solved very quickly.

In Section 6.1 we first review three variants of the Frank-Wolfe algorithm corresponding to a static, an adaptive and a more flexible *fully* adaptive stepsize rule, and we prove that the fully adaptive rule offers a linear convergence guarantee under standard regularity conditions. In Section 6.2 we then show that the nonlinear SDP (3.4) is amenable to the fully adaptive Frank-Wolfe algorithm and can thus be solved efficiently.

### 6.1. Frank-Wolfe Algorithm for Generic Convex Optimization Problems

Consider a generic convex minimization problem of the form

$$(6.1) \quad f^* = \min_{s \in \mathcal{S}} f(s)$$

with a convex compact feasible set  $\mathcal{S} \subseteq \mathbb{R}^d$  and a convex differentiable objective function  $f : \mathcal{S} \rightarrow \mathbb{R}$ . We assume that for each precision  $\delta \in [0, 1]$  we have access to an inexact oracle  $F : \mathcal{S} \rightarrow \mathcal{S}$  that maps any  $s \in \mathcal{S}$

to a  $\delta$ -approximate solution of an auxiliary problem linearized around  $s$ . More precisely, we assume that

$$(6.2) \quad (F(s) - s)^\top \nabla f(s) \leq \delta \min_{z \in \mathcal{S}} (z - s)^\top \nabla f(s).$$

By the standard optimality condition for convex optimization problems, the minimum on the right hand side of (6.2) vanishes if and only if  $s$  solves the original problem (6.1). Otherwise, the minimum is strictly negative. If  $\delta = 1$ , then the oracle returns an exact minimizer of the linearized problem. If  $\delta = 0$ , on the other hand, then the oracle returns any solution that is weakly preferred to  $s$  in the linearized problem. Given an oracle satisfying (6.2), one can design a Frank-Wolfe algorithm whose iterates obey the recursion

$$(6.3) \quad s_{t+1} = s_t + \eta_t (F(s_t) - s_t) \quad \forall t \in \mathbb{N} \cup \{0\},$$

where  $s_0 \in \mathcal{S}$  is an arbitrary initial feasible solution,  $\delta$  is a prescribed precision, and  $\eta_t \in [0, 1]$  is a stepsize that may depend on the current iterate  $s_t$ . The Frank-Wolfe algorithm was originally developed for quadratic programs [28] and later extended to general convex programs with differentiable objective functions and compact convex feasible sets [46, 16, 21, 19, 20]. Convergence guarantees for the Frank-Wolfe algorithm typically rely on the assumption that the gradient of  $f$  is Lipschitz continuous [46, 19, 20, 30, 29], that  $f$  has a bounded curvature constant [13, 38], or that the gradient of  $f$  is Hölder continuous [53].

Throughout this section we will assume that the decision variable can be represented as  $s = (s^{[1]}, \dots, s^{[K]})$ , where  $s^{[k]} \in \mathbb{R}^{d_k}$  and  $\sum_{k=1}^K d_k = d$ . Moreover, we will assume that the feasible set  $\mathcal{S} = \times_{k=1}^K \mathcal{S}^{[k]}$  constitutes a  $K$ -fold Cartesian product, where the marginal feasible set  $\mathcal{S}^{[k]} \subseteq \mathbb{R}^{d_k}$  is convex and compact for each  $k = 1, \dots, K$ . This assumption is unrestrictive because we are free to set  $K = 1$  and  $\mathcal{S}^{[1]} = \mathcal{S}$ . For ease of notation, we use from now on  $\nabla_{[k]} f(s)$  to denote the partial gradient with respect to the subvector  $s^{[k]} \in \mathcal{S}^{[k]}$ ,  $k = 1, \dots, K$ .

The subsequent convergence analysis will rely on the following regularity conditions.

**Assumption 6.1** (Regularity conditions).

(i) *The objective function is  $\beta$ -smooth for some  $\beta > 0$ , i.e.,*

$$\|\nabla f(s) - \nabla f(\bar{s})\| \leq \beta \|s - \bar{s}\| \quad \forall s, \bar{s} \in \mathcal{S}.$$

(ii) *The marginal feasible sets are  $\alpha$ -strongly convex with respect to  $f$  for some  $\alpha > 0$ , i.e.,*

$$\theta s^{[k]} + (1 - \theta) \bar{s}^{[k]} - \theta(1 - \theta) \frac{\alpha}{2} \left\| s^{[k]} - \bar{s}^{[k]} \right\|^2 \frac{\nabla_{[k]} f(s)}{\|\nabla_{[k]} f(s)\|} \in \mathcal{S}^{[k]} \quad \forall s, \bar{s} \in \mathcal{S}, \theta \in [0, 1], k = 1, \dots, K.$$

(iii) *The objective function is  $\varepsilon$ -steep for some  $\varepsilon > 0$ , i.e.,*

$$\|\nabla_{[k]} f(s)\| \geq \varepsilon \quad \forall s \in \mathcal{S}, k = 1, \dots, K.$$

Assumption 6.1 (ii) relaxes the standard strong convexity condition prevailing in the literature, which is obtained by setting  $K = 1$  and requiring that the condition stated here remains valid when the normalized gradient  $\nabla_{[1]} f(s) / \|\nabla_{[1]} f(s)\|$  is replaced with any other vector in the Euclidean unit ball, see, e.g., [40, Equation (25)]. We emphasize that our weaker condition is sufficient for the standard convergence proofs of the Frank-Wolfe algorithm but is necessary for our purposes because the feasible set of problem (3.4) fails to be strongly convex in the traditional sense. Similarly, Assumption 6.1 (iii) generalizes the usual  $\varepsilon$ -steepness condition from the literature, which is recovered by setting  $K = 1$ , see, e.g., [40, Assumption 1].

In the following we will distinguish three variants of the Frank-Wolfe algorithm with different stepsize rules. The *vanilla Frank-Wolfe* algorithm employs the harmonically decaying static stepsize

$$\eta_t = \frac{2}{2+t},$$

which results in a sublinear  $\mathcal{O}(1/t)$  convergence whenever Assumption 6.1 (i) holds [28, 21]. The *adaptive Frank-Wolfe* algorithm uses the stepsize

$$(6.4) \quad \eta_t = \min \left\{ 1, \frac{(s_t - F(s_t))^\top \nabla f(s_t)}{\beta \|s_t - F(s_t)\|^2} \right\},$$

which adapts to the iterate  $s_t$ . If all of the Assumptions 6.1 (i)–(iii) hold, then the adaptive Frank-Wolfe algorithm enjoys a linear  $\mathcal{O}(c^t)$  convergence guarantee, where  $c \in (0, 1)$  is an explicit function of the oracle precision  $\delta$ , the smoothness parameter  $\beta$ , the strong convexity parameter  $\alpha$  and the steepness parameter  $\varepsilon$  [46, 30]. Note that the stepsize (6.4) is constructed as the unique solution of the univariate quadratic program

$$\min_{\eta \in [0, 1]} f(s_t) - \eta (s_t - F(s_t))^\top \nabla f(s_t) + \frac{1}{2} \beta \eta^2 \|s_t - F(s_t)\|^2,$$

which minimizes a quadratic majorant of the objective function  $f$  along the line segment from  $s_t$  to  $F(s_t)$ .

The adaptive stepsize rule (6.4) has undergone further scrutiny in [58], where it was discovered that one may improve the algorithm's convergence behavior by replacing the global smoothness parameter  $\beta$  in (6.4) with an adaptive smoothness parameter  $\beta_t$  that captures the smoothness of  $f$  along the line segment from  $s_t$  to  $F(s_t)$ . This extra flexibility is useful because  $\beta_t$  can be chosen smaller than the unnecessarily conservative global smoothness parameter  $\beta$  and because  $\beta_t$  is easier to estimate than  $\beta$ , which may not even be accessible.

Following [58], we will henceforth only require that  $\beta_t > 0$  satisfies the inequality

$$(6.5) \quad f\left(s_t - \eta_t(\beta_t)(s_t - F(s_t))\right) \leq f(s_t) - \eta_t(\beta_t)(s_t - F(s_t))^\top \nabla f(s_t) + \frac{1}{2} \beta_t \eta_t(\beta_t)^2 \|s_t - F(s_t)\|^2,$$

where  $\eta_t(\beta_t)$  is defined as the adaptive stepsize (6.4) with  $\beta$  replaced by  $\beta_t$ . As it adapts both to  $s_t$  and  $\beta_t$ , we will from now on refer to  $\eta_t = \eta_t(\beta_t)$  as the *fully adaptive stepsize*. The above discussion implies that (6.5) is always satisfiable if Assumption 6.1 (i) holds, in which case one may simply set  $\beta_t$  to the global smoothness parameter  $\beta$ . In practice, however, the inequality (6.5) is often satisfiable for much smaller values  $\beta_t \ll \beta$  that may not even be related to the smoothness properties of the objective function. A close upper bound on the smallest  $\beta_t > 0$  that satisfies (6.5) can be found efficiently via backtracking line search. Specifically, the *fully adaptive Frank-Wolfe* algorithm sets  $\beta_t$  to the smallest element of the discrete search space  $\frac{\beta_{t-1}}{\zeta} \cdot \{1, \tau, \tau^2, \tau^3, \dots\}$  that satisfies (6.5), where  $\tau > 1$  and  $\zeta > 1$  are prescribed line search parameters. A detailed description of the fully adaptive Frank-Wolfe algorithm in pseudocode is provided in Algorithm 1.

It has been shown in [58] that Algorithm 1 enjoys the same sublinear  $\mathcal{O}(1/t)$  convergence guarantee as the vanilla Frank-Wolfe algorithm when Assumption 6.1 (i) holds. Below we will leverage techniques from [46, 30] to show that Algorithm 1 offers indeed a linear convergence rate if all of the Assumptions 6.1 (i)–(iii) hold.

**Theorem 6.2** (Linear convergence of the fully adaptive Frank-Wolfe algorithm). *If Assumption 6.1 holds and  $\bar{\beta} = \max\{\tau\beta, \beta_{-1}\}$ , then Algorithm 1 enjoys the linear convergence guarantee*

$$f(s_t) - f^* \leq \max \left\{ 1 - \frac{\delta}{2}, 1 - \frac{(1 - \sqrt{1 - \delta})\alpha\varepsilon}{4\bar{\beta}} \right\}^t (f(s_0) - f^*) \quad \forall t \in \mathbb{N}.$$

The proof of Theorem 6.2 relies on the following preparatory lemma.

**Lemma 6.3** (Bounds on the surrogate duality gap). *The surrogate duality gap  $g_t = -d_t^\top \nabla f(s_t)$  corresponding to the search direction  $d_t = F(s_t) - s_t$  admits the following lower bounds.*

- (i) *If the objective function  $f$  is convex, then  $g_t \geq \delta(f(s_t) - f^*)$ .*

**Algorithm 1** Fully adaptive Frank-Wolfe algorithm

---

**Input:** initial feasible point  $s_0 \in \mathcal{S}$ , initial smoothness parameter  $\beta_{-1} > 0$   
line search parameters  $\tau > 1$ ,  $\zeta > 1$ , initial iteration counter  $t = 0$

**while** stopping criterion is not met **do**

solve the oracle subproblem to find  $\tilde{s}_t = F(s_t)$   
set  $d_t \leftarrow \tilde{s}_t - s_t$  and  $g_t \leftarrow -d_t^\top \nabla f(s_t)$   
set  $\beta_t \leftarrow \beta_{t-1}/\zeta$  and  $\eta \leftarrow \min\{1, g_t/(\beta_t \|d_t\|^2)\}$

**while**  $f(s_t + \eta d_t) > f(s_t) - \eta g_t + \frac{\eta^2 \beta_t}{2} \|d_t\|^2$  **do**

$\beta_t \leftarrow \tau \beta_t$  and  $\eta \leftarrow \min\{1, g_t/(\beta_t \|d_t\|^2)\}$

**end while**

set  $\eta_t \leftarrow \eta$  and  $s_{t+1} \leftarrow s_t + \eta_t d_t$   
set  $t \leftarrow t + 1$

**end while**

**Output:**  $s_t$

---

(ii) If the marginal feasible sets are  $\alpha$ -strongly convex with respect to  $f$  for some  $\alpha > 0$  in the sense of Assumption 6.1(ii), then

$$g_t \geq \min_{k \in \{1, \dots, K\}} \frac{(1 - \sqrt{1 - \delta})\alpha}{2\delta} \|d_t\|^2 \|\nabla_{[k]} f(s_t)\|.$$

*Proof of Lemma 6.3.* By the definition of  $g_t$  we have

$$(6.6) \quad g_t = (s_t - F(s_t))^\top \nabla f(s_t) \geq \delta (s_t - s)^\top \nabla f(s_t) \quad \forall s \in \mathcal{S},$$

where the inequality follows from the defining property (6.2) of the inexact oracle with precision  $\delta$ . Setting  $s$  in (6.6) to a global minimizer  $s^*$  of (6.1) then implies via the first-order convexity condition for  $f$  that

$$g_t \geq \delta (s_t - s^*)^\top \nabla f(s_t) \geq \delta (f(s_t) - f^*).$$

This observation establishes assertion (i). To prove assertion (ii), we first rewrite the estimate (6.6) as

$$(6.7) \quad g_t \geq \delta (s_t - s)^\top \nabla f(s) = \delta \sum_{t=1}^T (s_t^{[k]} - s^{[k]})^\top \nabla_{[k]} f(s_t) \quad \forall s \in \mathcal{S}.$$

In the following, we denote by  $F^{[k]} : \mathcal{S} \rightarrow \mathcal{S}^{[k]}$  the  $k$ -th suboracle for  $k = 1, \dots, K$ , which is defined through the identity  $F = (F^{[1]}, \dots, F^{[K]})$ . Similarly, for any  $\theta \in [0, 1]$ , we define  $s(\theta) = (s^{[1]}(\theta), \dots, s^{[K]}(\theta))$  through

$$s^{[k]}(\theta) = \theta F^{[k]}(s_t) + (1 - \theta) s_t^{[k]} - \frac{\alpha}{2} \theta(1 - \theta) \left\| F^{[k]}(s_t) - s_t^{[k]} \right\|^2 \frac{\nabla_{[k]} f(s_t)}{\|\nabla_{[k]} f(s_t)\|}.$$

By Assumption 6.1 (ii), we have  $s^{[k]}(\theta) \in \mathcal{S}^{[k]}$  for every  $k = 1, \dots, K$ . Thanks to the rectangularity of the feasible set this implies that  $s(\theta) \in \mathcal{S}$ . Setting  $s$  in (6.7) to  $s(\theta)$ , we thus find

$$\begin{aligned} g_t &\geq \delta \sum_{t=1}^T \left( \theta(s_t^{[k]} - F^{[k]}(s_t)) + \frac{\alpha}{2} \theta(1 - \theta) \left\| F^{[k]}(s_t) - s_t^{[k]} \right\|^2 \frac{\nabla_{[k]} f(s_t)}{\|\nabla_{[k]} f(s_t)\|} \right)^\top \nabla_{[k]} f(s_t) \\ &= \delta \left( \theta g_t + \frac{\alpha}{2} \theta(1 - \theta) \left[ \sum_{k=1}^K \left\| F^{[k]}(s_t) - s_t^{[k]} \right\|^2 \|\nabla_{[k]} f(s_t)\| \right] \right) \\ &\geq \delta \left( \theta g_t + \frac{\alpha}{2} \theta(1 - \theta) \|F(s_t) - s_t\|^2 \left( \min_{k \in \{1, \dots, K\}} \|\nabla_{[k]} f(s_t)\| \right) \right) \quad \forall \theta \in [0, 1], \end{aligned}$$

where the equality follows from the definition of  $g_t$ , and the last inequality exploits the Pythagorean theorem. Reordering the above inequality to bring  $g_t$  to the left hand side yields

$$(6.8) \quad g_t \geq \min_{k \in \{1, \dots, K\}} \frac{\alpha}{2} \|F(s_t) - s_t\|^2 \|\nabla_{[k]} f(s_t)\| \frac{\delta \theta(1 - \theta)}{1 - \delta \theta} \quad \forall \theta \in [0, 1].$$

A tedious but straightforward calculation shows that the lower bound on the right hand side of (6.8) is maximized by  $\theta^* = (1 - \sqrt{1 - \delta})/\delta$ . Assertion (ii) then follows by substituting  $\theta^*$  into (6.8).  $\square$

*Proof of Theorem 6.2.* By Assumption 6.1 (i) the function  $f$  is  $\beta$ -smooth, and thus one can show that

$$(6.9) \quad f(s_t + \eta d_t) \leq f(s_t) - \eta g_t + \frac{\eta^2 \beta}{2} \|d_t\|^2 \quad \forall \eta \in [0, 1],$$

where the surrogate duality gap  $g_t \geq 0$  and the search direction  $d_t \in \mathbb{R}^d$  are defined as in Lemma 6.3. We emphasize that (6.9) holds in fact for all  $\eta \in \mathbb{R}$ . However, the next iterate  $s_{t+1} = s_t + \eta d_t$  may be infeasible unless  $\eta \in [0, 1]$ . The inequality (6.9) implies that any  $\beta_t \geq \beta$  satisfies the condition of the inner while loop of Algorithm 1, and thus the loop must terminate at the latest after  $\lceil \log(\zeta \beta / \beta_{-1}) / \log(\tau) \rceil$  iterations, outputting a smoothness parameter  $\beta_t$  and a stepsize  $\eta_t$  that satisfy the inequality (6.5). We henceforth denote by  $h_t = f(s_t) - f^*$  the suboptimality of the  $k$ -th iterate and note that

$$(6.10) \quad h_{t+1} = f(s_t + \eta_t d_t) - f(s_t) + h_t \leq -g_t + \frac{1}{2} \beta_t \eta_t^2 \|d_t\|^2 + h_t,$$

where the inequality exploits (6.5) and the definitions of  $g_t$  and  $d_t$ . In order to show that  $h_t$  decays geometrically, we distinguish the cases (i)  $g_t / (\beta_t \|d_t\|^2) \geq 1$  and (ii)  $g_t / (\beta_t \|d_t\|^2) < 1$ . In case (i), the stepsize  $\eta_t$  defined in (6.5) satisfies  $\eta_t = \min\{1, g_t / (\beta_t \|d_t\|^2)\} = 1$ , and thus we have

$$(6.11) \quad h_{t+1} \leq \left( \frac{\beta_t \|d_t\|^2}{2g_t} - 1 \right) g_t + h_t \leq -\frac{g_t}{2} + h_t \leq \left( 1 - \frac{\delta}{2} \right) h_t,$$

where the first inequality follows from (6.10), while the third inequality holds due to Lemma 6.3 (i).

In case (ii), the stepsize satisfies  $\eta_t = g_t / \beta_t \|d_t\|^2 < 1$ , and thus we find

$$(6.12) \quad \begin{aligned} h_{t+1} &\leq -g_t + \frac{g_t^2}{2\beta_t \|d_t\|^2} + h_t \leq -\frac{g_t^2}{2\beta_t \|d_t\|^2} + h_t \leq \left( 1 - \frac{\delta g_t}{2\beta_t \|d_t\|^2} \right) h_t \\ &\leq \left( 1 - \min_{k \in \{1, \dots, K\}} \frac{(1 - \sqrt{1 - \delta}) \alpha}{4\beta_t} \|\nabla_{[k]} f(s_t)\| \right) h_t \leq \left( 1 - \frac{(1 - \sqrt{1 - \delta}) \alpha \varepsilon}{4\beta} \right) h_t, \end{aligned}$$

where the first and the second inequalities follow from (6.10) and from multiplying  $-g_t$  with  $\eta_t < 1$ , respectively, while the third and the fourth inequalities exploit Lemmas 6.3 (i) and 6.3 (ii), respectively. The



last inequality in (6.12) holds because of Assumption 6.1 (iii) and because  $\beta_t \leq \bar{\beta}$  for all  $t \in \mathbb{N}$ ; see [58, Proposition 2]. By the estimates (6.11) and (6.12), the suboptimality of the current iterate decays at least by

$$\max \left\{ 1 - \frac{\delta}{2}, 1 - \frac{(1 - \sqrt{1 - \delta})\alpha\varepsilon}{4\bar{\beta}} \right\} < 1$$

in each iteration of the algorithm. This observation completes the proof.  $\square$

## 6.2. Frank-Wolfe Algorithm for Wasserstein MMSE Estimation Problems

We now use the fully adaptive Frank-Wolfe algorithm of Section 6.1 to solve the nonlinear SDP (3.4), which is equivalent to the dual Wasserstein MMSE estimation problem over normal priors. Recall from Corollary 4.3 that any solution of (3.4) can be used to construct a least favorable prior and an optimal estimator that form a Nash equilibrium. Unlike the generic convex program (6.1), the nonlinear SDP (3.4) is a convex *maximization* problem. This prompts us to apply Algorithm 1 to the convex minimization problem obtained from problem (3.4) by turning the objective function upside down.

Throughout this section we assume that  $\hat{\Sigma}_x \succ 0$ ,  $\hat{\Sigma}_w \succ 0$ ,  $\rho_x > 0$  and  $\rho_w > 0$ , which implies via Theorem 3.5 that the nonlinear SDP (3.4) is solvable and can be reformulated more concisely as

$$(6.13) \quad \max_{\Sigma_x \in \mathcal{S}_x^+, \Sigma_w \in \mathcal{S}_w^+} f(\Sigma_x, \Sigma_w),$$

where the objective function  $f : \mathcal{S}_x^+ \times \mathcal{S}_w^+ \rightarrow \mathbb{R}$  is defined through

$$f(\Sigma_x, \Sigma_w) = \text{Tr} [\Sigma_x - \Sigma_x H^\top (H \Sigma_x H^\top + \Sigma_w)^{-1} H \Sigma_x],$$

and where the separate feasible sets for  $\Sigma_x$  and  $\Sigma_w$  are given by

$$\mathcal{S}_x^+ = \left\{ \Sigma_x \in \mathbb{S}_+^n : \text{Tr} [\Sigma_x + \hat{\Sigma}_x - 2(\hat{\Sigma}_x^{\frac{1}{2}} \Sigma_x \hat{\Sigma}_x^{\frac{1}{2}})^{\frac{1}{2}}] \leq \rho_x^2, \Sigma_x \succeq \lambda_{\min}(\hat{\Sigma}_x) I_n \right\}$$

and

$$\mathcal{S}_w^+ = \left\{ \Sigma_w \in \mathbb{S}_+^m : \text{Tr} [\Sigma_w + \hat{\Sigma}_w - 2(\hat{\Sigma}_w^{\frac{1}{2}} \Sigma_w \hat{\Sigma}_w^{\frac{1}{2}})^{\frac{1}{2}}] \leq \rho_w^2, \Sigma_w \succeq \lambda_{\min}(\hat{\Sigma}_w) I_m \right\},$$

respectively. One readily verifies that  $f$  is concave and differentiable. Moreover, in the terminology of Section 6.1, the feasible set of the nonlinear SDP (6.13) constitutes a Cartesian product of  $K = 2$  marginal feasible sets  $\mathcal{S}_x^+$  and  $\mathcal{S}_w^+$ , both of which are convex and compact thanks to Lemma A.6. Note that  $\mathcal{S}_x^+$  and  $\mathcal{S}_w^+$  constitute restrictions of the feasible sets  $\mathcal{S}_x$  and  $\mathcal{S}_w$ , respectively, which appeared in the proofs of Theorems 3.5 and 4.1. The oracle problem that linearizes the objective function of the nonlinear SDP (6.13) around a fixed feasible solution  $\Sigma_x \in \mathcal{S}_x^+$  and  $\Sigma_w \in \mathcal{S}_w^+$  can now be expressed concisely as

$$(6.14) \quad \max_{L_x \in \mathcal{S}_x^+, L_w \in \mathcal{S}_w^+} \langle L_x - \Sigma_x, D_x \rangle + \langle L_w - \Sigma_w, D_w \rangle,$$

where  $D_x = \nabla_{\Sigma_x} f(\Sigma_x, \Sigma_w)$  and  $D_w = \nabla_{\Sigma_w} f(\Sigma_x, \Sigma_w)$  represent the gradients of  $f$  with respect to  $\Sigma_x$  and  $\Sigma_w$ . Lemma A.7 offers analytical formulas for  $D_x$  and  $D_w$  and shows that they are both positive semidefinite.

The oracle problem (6.14) is manifestly separable in  $L_x$  and  $L_w$  and can therefore be decomposed into a sum of two structurally identical marginal subproblems. The Frank-Wolfe algorithm is an ideal method to address the nonlinear SDP (6.13) because these two marginal oracle subproblems admit quasi-closed form solutions. Specifically, Proposition A.4 (iii) in the appendix implies that problem (6.14) is uniquely solved by

$$L_x^* = (\gamma_x^*)^2 (\gamma_x^* I_n - D_x)^{-1} \hat{\Sigma}_x (\gamma_x^* I_n - D_x)^{-1} \quad \text{and} \quad L_w^* = (\gamma_w^*)^2 (\gamma_w^* I_m - D_w)^{-1} \hat{\Sigma}_w (\gamma_w^* I_m - D_w)^{-1},$$

where  $\gamma_x^* \in (\lambda_{\max}(D_x), \infty)$  and  $\gamma_w^* \in (\lambda_{\max}(D_w), \infty)$  are the unique solutions of the algebraic equations

$$(6.15) \quad \rho_x^2 - \langle \hat{\Sigma}_x, (I_n - \gamma_x^* (\gamma_x^* I_n - D_x)^{-1})^2 \rangle = 0 \quad \text{and} \quad \rho_w^2 - \langle \hat{\Sigma}_w, (I_m - \gamma_w^* (\gamma_w^* I_m - D_w)^{-1})^2 \rangle = 0,$$

respectively. In practice, these algebraic equations need to be solved numerically. However, the numerical errors in  $\gamma_x^*$  and  $\gamma_w^*$  must be contained to ensure that  $L_x^*$  and  $L_w^*$  give rise to a  $\delta$ -approximate solution for (6.14) in the sense of (6.2). In the following we will show that  $\delta$ -approximate solutions for each of the two oracle subproblems in (6.14) and for each  $\delta \in (0, 1)$  can be computed with an efficient bisection algorithm.

---

**Algorithm 2** Bisection algorithm for the oracle subproblem

---

**Input:** nominal covariance matrix  $\widehat{\Sigma} \in \mathbb{S}_{++}^d$ , radius  $\rho \in \mathbb{R}_{++}$ ,  
reference covariance matrix  $\Sigma \in \mathbb{S}_+^d$  feasible in (6.14),  
gradient matrix  $D \in \mathbb{S}_+^d$ ,  $D \neq 0$ , precision  $\delta \in (0, 1)$ ,  
dual objective function  $\varphi(\gamma)$  defined in Theorem 6.4

set  $\lambda_1 \leftarrow \lambda_{\max}(D)$ , and let  $v_1 \in \mathbb{R}^d$  be an eigenvector for  $\lambda_1$

set  $\underline{\gamma} \leftarrow \lambda_1(1 + (v_1^\top \widehat{\Sigma} v_1)^{\frac{1}{2}}/\rho)$  and  $\bar{\gamma} \leftarrow \lambda_1(1 + \text{Tr}[\widehat{\Sigma}]^{\frac{1}{2}}/\rho)$

**repeat**

Set  $\tilde{\gamma} \leftarrow (\bar{\gamma} + \underline{\gamma})/2$  and  $\tilde{L} \leftarrow (\tilde{\gamma})^2(\tilde{\gamma}I_d - D)^{-1}\widehat{\Sigma}(\tilde{\gamma}I_d - D)^{-1}$

**if**  $\frac{d\varphi}{d\gamma}(\tilde{\gamma}) < 0$  **then** Set  $\underline{\gamma} \leftarrow \tilde{\gamma}$  **else** Set  $\bar{\gamma} \leftarrow \tilde{\gamma}$  **endif**

**until**  $\frac{d\varphi}{d\gamma}(\tilde{\gamma}) > 0$  and  $\langle \tilde{L} - \Sigma, D \rangle \geq \delta \varphi(\tilde{\gamma})$

**Output:**  $\tilde{L}$

---

**Theorem 6.4** (Approximate oracle). *For any fixed  $\rho \in \mathbb{R}_{++}$ ,  $\widehat{\Sigma} \in \mathbb{S}_{++}^d$  and  $D \in \mathbb{S}_+^d$ ,  $D \neq 0$ , consider the generic oracle subproblem*

$$(6.16) \quad \begin{aligned} & \max_{L \in \mathbb{S}_+^d} \langle L - \Sigma, D \rangle \\ & \text{s. t.} \quad \text{Tr} \left[ L + \widehat{\Sigma} - 2(\widehat{\Sigma}^{\frac{1}{2}} L \widehat{\Sigma}^{\frac{1}{2}})^{\frac{1}{2}} \right] \leq \rho^2, \quad L \succeq \lambda_{\min}(\widehat{\Sigma})I_d, \end{aligned}$$

where  $\Sigma \in \mathbb{S}_+^d$  represents a feasible reference solution. Moreover, denote the feasible set of problem (6.16) by  $\mathcal{S}^+$ , let  $\delta \in (0, 1)$  be the desired oracle precision, and define  $\varphi(\gamma) = \gamma(\rho^2 + \langle \gamma(\gamma I_d - D)^{-1} - I_d, \widehat{\Sigma} \rangle) - \langle \Sigma, D \rangle$  for any  $\gamma > \lambda_{\max}(D)$ . Then, Algorithm 2 returns in finite time a matrix  $\tilde{L} \in \mathbb{S}_+^d$  with the following properties.

(i) *Feasibility:*  $\tilde{L} \in \mathcal{S}^+$

(ii)  *$\delta$ -Suboptimality:*  $\langle \tilde{L} - \Sigma, D \rangle \geq \delta \max_{L \in \mathcal{S}^+} \langle L - \Sigma, D \rangle$

*Proof of Theorem 6.4.* Proposition A.4 (iii) in the appendix guarantees that the lower bound on  $L$  in (6.16) is redundant and can be omitted without affecting the problem's optimal value. By Proposition A.4 (i), the oracle subproblem (6.16) thus admits the strong Lagrangian dual

$$\min_{\gamma > \lambda_{\max}(D)} \varphi(\gamma),$$

where the convex and differentiable function  $\varphi(\gamma)$  is defined as in the theorem statement. In the following we denote by  $\lambda_1 > 0$  the largest eigenvalue of  $D$  and let  $v_1 \in \mathbb{R}^d$  be a corresponding eigenvalue. By Proposition A.4 (iii), the dual oracle subproblem admits a minimizer  $\gamma^* \in [\underline{\gamma}, \bar{\gamma}]$  that is uniquely determined by the first-order optimality condition  $\frac{d\varphi}{d\gamma}(\gamma^*) = 0$ , where

$$\underline{\gamma} = \lambda_1 \left( 1 + \sqrt{v_1^\top \widehat{\Sigma} v_1 / \rho} \right) \quad \text{and} \quad \bar{\gamma} = \lambda_1 \left( 1 + \sqrt{\text{Tr}[\widehat{\Sigma}] / \rho} \right),$$

while the primal problem (6.16) admits a unique maximizer  $L^* = L(\gamma^*)$ , where

$$L(\gamma) = \gamma^2(\gamma I_d - D)^{-1}\widehat{\Sigma}(\gamma I_d - D)^{-1}.$$

From the proof of Proposition A.4 (iii) it is evident that  $L(\gamma) \succ \lambda_{\min}(\widehat{\Sigma})I_d$  for every  $\gamma > 0$ . A direct calculation further shows that

$$\frac{d\varphi}{d\gamma}(\gamma) = \rho^2 - \langle \widehat{\Sigma}, (I_d - \gamma(\gamma I_d - D)^{-1})^2 \rangle = \rho^2 - \text{Tr} [L(\gamma) + \widehat{\Sigma} - 2(\widehat{\Sigma}^{\frac{1}{2}}L(\gamma)\widehat{\Sigma}^{\frac{1}{2}})^{\frac{1}{2}}].$$

Recalling that  $\varphi(\gamma)$  is convex and that its derivative is non-negative for all  $\gamma \geq \gamma^*$ , the above reasoning implies that  $L(\gamma) \in \mathcal{S}^+$  for all  $\gamma \geq \gamma^*$ . Note also that the optimal value of the primal problem (6.16) is non-negative because  $\Sigma \in \mathcal{S}^+$ . The continuity of  $\langle L(\gamma) - \Sigma, D \rangle$  at  $\gamma = \gamma^*$  thus ensures that there exists  $\delta' > 0$  with

$$\langle L(\gamma) - \Sigma, D \rangle \leq \delta \langle L(\gamma^*) - \Sigma, D \rangle = \delta \max_{L \in \mathcal{S}^+} \langle L - \Sigma, D \rangle \quad \forall \gamma \in [\gamma^*, \gamma^* + \delta'].$$

In summary, computing a feasible and  $\delta$ -suboptimal matrix  $\tilde{L} \in \mathbb{S}_+^d$  is tantamount to finding  $\tilde{\gamma} \in [\gamma^*, \gamma^* + \delta']$ . Algorithm 2 uses bisection over the interval  $[\underline{\gamma}, \bar{\gamma}]$  to find a  $\tilde{\gamma}$  with these properties.  $\square$

Theorem 6.4 complements [64, Theorem 3.2], which constructs an approximate oracle for a nonlinear SDP similar to (6.13) that offers an *additive* error guarantee. The *multiplicative* error guarantee of the oracle constructed here is needed to ensure the linear convergence of the fully adaptive Frank-Wolfe algorithm. Next, we prove that the nonlinear SDP (6.13) satisfies all regularity conditions listed in Assumption 6.1.

**Proposition 6.5** (Regularity conditions of the nonlinear SDP (6.13)). *If  $\rho_x \in \mathbb{R}_{++}$ ,  $\rho_w \in \mathbb{R}_{++}$ ,  $\widehat{\Sigma}_x \in \mathbb{S}_+^n$  and  $\widehat{\Sigma}_w \in \mathbb{S}_+^m$ , then the nonlinear SDP (6.13) obeys the following regularity conditions.*

(i) *The objective function of problem (6.13) is  $\beta$ -smooth in the sense of Assumption 6.1 (i), where*

$$\beta = 2\lambda_{\min}^{-1}(\widehat{\Sigma}_w) (C + C \lambda_{\max}^2(H^\top H) + \lambda_{\max}(H^\top H)),$$

*which depends on the auxiliary constant  $C = \lambda_{\max}(H^\top H) \cdot \lambda_{\min}^{-2}(\widehat{\Sigma}_w) \cdot (\rho_x + \text{Tr} [\widehat{\Sigma}_x]^{\frac{1}{2}})^4$ .*

(ii) *The marginal feasible sets  $\mathcal{S}_x^+$  and  $\mathcal{S}_w^+$  of problem (6.13) are  $\alpha$ -strongly convex with respect to  $-f$  in the sense of Assumption 6.1 (ii), where  $\alpha = \min \{\alpha_x, \alpha_w\}$ , which depends on the auxiliary constants*

$$\alpha_x = \frac{\lambda_{\min}^{\frac{5}{4}}(\widehat{\Sigma}_x)}{2\rho_x(\rho_x + \text{Tr} [\widehat{\Sigma}_x]^{\frac{1}{2}})^{\frac{7}{2}}} \quad \text{and} \quad \alpha_w = \frac{\lambda_{\min}^{\frac{5}{4}}(\widehat{\Sigma}_w)}{2\rho_w(\rho_w + \text{Tr} [\widehat{\Sigma}_w]^{\frac{1}{2}})^{\frac{7}{2}}}.$$

(iii) *The objective function of problem (6.13) is  $\varepsilon$ -steep in the sense of Assumption 6.1 (iii), where  $\varepsilon = \min \{\varepsilon_x, \varepsilon_w\}$ , which depends on the auxiliary constants*

$$\varepsilon_x = \left( \frac{\lambda_{\min}(\widehat{\Sigma}_w)}{(\rho_x + \text{Tr} [\widehat{\Sigma}_x]^{\frac{1}{2}})^2 \lambda_{\max}(H^\top H) + (\rho_w + \text{Tr} [\widehat{\Sigma}_w]^{\frac{1}{2}})^2} \right)^2$$

and

$$\varepsilon_w = \lambda_{\max}(H^\top H) \left( \frac{\lambda_{\min}(\widehat{\Sigma}_x)}{(\rho_w + \text{Tr} [\widehat{\Sigma}_w]^{\frac{1}{2}})^2 + \lambda_{\min}(\widehat{\Sigma}_x) \lambda_{\max}(H^\top H)} \right)^2.$$

*Proof of Proposition 6.5.* The proof repeatedly uses the fact that, for any  $A \in \mathbb{R}^{d_1 \times d_2}$  and  $B \in \mathbb{S}_+^{d_2}$ , we have

$$(6.17) \quad \lambda_{\max}(ABA^\top) = \lambda_{\max}(A^\top AB) \leq \lambda_{\max}(A^\top A) \lambda_{\max}(B).$$

The equality in (6.17) holds because all eigenvalues of  $ABA^\top$  are non-negative and because the non-zero spectrum of  $ABA^\top$  is identical to that of  $A^\top AB$  due to [6, Proposition 4.4.10]. The inequality follows from the observation that  $\lambda_{\max}(A^\top A)$  and  $\lambda_{\max}(B)$  coincide with the operator norms of the positive semidefinite matrices  $A^\top A$  and  $B$ , respectively.

As for assertion (i), recall first that the objective function  $f$  of the nonlinear SDP (3.4) is concave. In order to show that  $f$  is  $\beta$ -smooth for some  $\beta > 0$ , it thus suffices to prove that the largest eigenvalue of the positive semidefinite Hessian matrix of  $-f$  admits an upper bound uniformly across  $\mathcal{S}_x^+ \times \mathcal{S}_w^+$ . By Lemma A.7, the partial gradients of  $f$  evaluated at  $\Sigma_x \succ 0$  and  $\Sigma_w \succ 0$  are given by

$$\begin{aligned} D_x &= \nabla_{\Sigma_x} f(\Sigma_x, \Sigma_w) = (I_n - \Sigma_x H^\top G^{-1} H)^\top (I_n - \Sigma_x H^\top G^{-1} H) \\ D_w &= \nabla_{\Sigma_w} f(\Sigma_x, \Sigma_w) = G^{-1} H \Sigma_x^2 H^\top G^{-1}, \end{aligned}$$

where  $G = H \Sigma_x H^\top + \Sigma_w$ . Moreover, the Hessian matrix

$$\mathcal{H} = \begin{bmatrix} \mathcal{H}_{xx} & \mathcal{H}_{xw} \\ \mathcal{H}_{xw}^\top & \mathcal{H}_{ww} \end{bmatrix} \succeq 0$$

of the convex function  $-f$  evaluated at  $\Sigma_x \succ 0$  and  $\Sigma_w \succ 0$  consists of the submatrices

$$\begin{aligned} \mathcal{H}_{xx} &= -\nabla_{xx}^2 f(\Sigma_x, \Sigma_w) = 2D_x \otimes H^\top G^{-1} H \\ \mathcal{H}_{xw} &= -\nabla_{xw}^2 f(\Sigma_x, \Sigma_w) = H^\top G^{-1} \otimes (H^\top D_w - \Sigma_x H^\top G^{-1}) + (H^\top D_w - \Sigma_x H^\top G^{-1}) \otimes H^\top G^{-1} \\ \mathcal{H}_{ww} &= -\nabla_{ww}^2 f(\Sigma_x, \Sigma_w) = 2D_w \otimes G^{-1}, \end{aligned}$$

where  $\nabla_x$  and  $\nabla_w$  are used as shorthands for the nabla operators with respect to  $\text{vec}(\Sigma_x)$  and  $\text{vec}(\Sigma_w)$ , respectively. To construct an upper bound on  $\lambda_{\max}(\mathcal{H})$  uniformly across  $\mathcal{S}_x^+ \times \mathcal{S}_w^+$ , we note first that

$$(6.18) \quad \lambda_{\max}(\mathcal{H}) \leq \lambda_{\max}(\mathcal{H}_{xx}) + \lambda_{\max}(\mathcal{H}_{ww}) = 2 \left( \lambda_{\max}(D_x) \lambda_{\max}(H^\top G^{-1} H) + \lambda_{\max}(D_w) \lambda_{\max}(G^{-1}) \right),$$

where the inequality follows from [6, Fact 5.12.20] and the subadditivity of the maximum eigenvalue, whereas the equality exploits the trace rule of the Kronecker product [6, Proposition 7.1.10]. In the remainder of the proof, we derive an upper bound for each term on the right hand side of the above expression.

By the definition of  $G$  and because  $\Sigma_w \in \mathcal{S}_w^+$ , we have  $G \succeq \lambda_{\min}(\Sigma_w) I_m \succeq \lambda_{\min}(\widehat{\Sigma}_w) I_m$ . As  $\widehat{\Sigma}_w \succ 0$  by assumption, we may thus conclude that  $\lambda_{\max}(G^{-1}) \leq \lambda_{\min}^{-1}(\widehat{\Sigma}_w)$ , which in turn implies via (6.17) that

$$\lambda_{\max}(H^\top G^{-1} H) \leq \lambda_{\max}(G^{-1}) \lambda_{\max}(H H^\top) \leq \lambda_{\min}^{-1}(\widehat{\Sigma}_w) \lambda_{\max}(H^\top H).$$

Similarly, by the definition of  $D_w$  we find

$$\begin{aligned} \lambda_{\max}(D_w) &= \lambda_{\max}(G^{-1} H \Sigma_x^2 H^\top G^{-1}) \\ &\leq \lambda_{\max}^2(\Sigma_x) \lambda_{\max}^2(G^{-1}) \lambda_{\max}(H^\top H) \leq (\rho_x + \text{Tr}[\widehat{\Sigma}_x]^{\frac{1}{2}})^4 \lambda_{\min}^{-2}(\widehat{\Sigma}_w) \lambda_{\max}(H^\top H), \end{aligned}$$

where the first inequality follows from applying the estimate (6.17) twice, while the last inequality reuses the bound on  $\lambda_{\max}(G^{-1})$  derived above and exploits Lemma A.6. Finally, by the definition of  $D_x$  we have

$$\begin{aligned} \lambda_{\max}(D_x) &\leq 1 + \lambda_{\max}(-H^\top G^{-1} H \Sigma_x) + \lambda_{\max}(-\Sigma_x H^\top G^{-1} H) + \lambda_{\max}(H^\top G^{-1} H \Sigma_x^2 H^\top G^{-1} H) \\ &\leq 1 + \lambda_{\max}(H^\top G^{-1} H \Sigma_x^2 H^\top G^{-1} H) \\ &\leq 1 + \lambda_{\max}^2(\Sigma_x) \lambda_{\max}^2(G^{-1}) \lambda_{\max}^2(H^\top H) \\ &\leq 1 + (\rho_x + \text{Tr}[\widehat{\Sigma}_x]^{\frac{1}{2}})^4 \lambda_{\min}^{-2}(\widehat{\Sigma}_w) \lambda_{\max}^2(H^\top H). \end{aligned}$$

where the first inequality holds due to the subadditivity of the maximum eigenvalue and [6, Proposition 4.4.10], which implies that the nonzero spectra of  $-\Sigma_x H^\top G^{-1} H$  and  $-H^\top G^{-1} H \Sigma_x$  are both real and coincide with the nonzero spectrum of the negative semidefinite matrix  $-\Sigma_x^{\frac{1}{2}} H^\top G^{-1} H \Sigma_x^{\frac{1}{2}}$ . The third inequality follows from applying the estimate (6.17) three times, and the fourth inequality reuses the bound on  $\lambda_{\max}(G^{-1})$  and exploits Lemma A.6. Substituting all the above bounds into (6.18) completes the proof of assertion (i).

As for assertion (ii), we first show that the feasible set  $\mathcal{S}_x^+$  is  $\alpha_x$ -strongly convex with respect to  $-f$  in the sense of Assumption 6.1 (ii). To see this, fix any  $\Sigma_x, \Sigma'_x \in \mathcal{S}_x^+$  and  $\theta \in [0, 1]$ , and set

$$(6.19) \quad \Sigma_\theta = \theta \Sigma_x + (1 - \theta) \Sigma'_x + \theta(1 - \theta) \frac{\alpha_x}{2} \frac{\|D_x\|^2}{\|D_x\|},$$

where  $\alpha_x > 0$  is defined as in the proposition statement, and  $D_x$  denotes again the partial gradient of  $f$  with respect to  $\Sigma_x$ . To prove strong convexity of  $\mathcal{S}_x^+$  with respect to  $-f$ , we will show that  $\Sigma_\theta \in \mathcal{S}_x^+$ . Note first that  $\Sigma_\theta \succeq \lambda_{\min}(\widehat{\Sigma}_x) I_n$  because  $\Sigma_x, \Sigma'_x \in \mathcal{S}_x^+$  and because  $D_x$  is positive semidefinite. Next, define

$$\overline{\mathcal{S}}_x = \left\{ \Sigma_x \in \mathbb{S}_+^n : \lambda_{\min}(\widehat{\Sigma}_x) I_n \preceq \Sigma_x \preceq (\rho_x + \text{Tr}[\widehat{\Sigma}_x]^{\frac{1}{2}}) I_n \right\},$$

and note that  $\mathcal{S}_x^+ \subseteq \overline{\mathcal{S}}_x$  by Lemma A.6. Moreover, [8, Theorem 1] implies that the function  $g_x : \mathbb{S}_+^n \rightarrow \mathbb{R}$  defined through  $g_x(\Sigma_x) = \text{Tr}[\Sigma_x + \widehat{\Sigma}_x - 2(\widehat{\Sigma}_x^{\frac{1}{2}} \Sigma_x \widehat{\Sigma}_x^{\frac{1}{2}})^{\frac{1}{2}}]$  is  $\kappa_1$ -strongly convex and  $\kappa_2$ -smooth over  $\overline{\mathcal{S}}_x$ , where

$$\kappa_1 = \frac{\lambda_{\min}^{\frac{1}{2}}(\widehat{\Sigma}_x)}{2(\rho_x + \text{Tr}[\widehat{\Sigma}_x]^{\frac{1}{2}})^3} \quad \text{and} \quad \kappa_2 = \frac{\rho_x + \text{Tr}[\widehat{\Sigma}_x]^{\frac{1}{2}}}{2\lambda_{\min}^{\frac{3}{2}}(\widehat{\Sigma}_x)}.$$

By [40, Theorem 12], the sublevel set  $\mathcal{S}_x^+ = \{\Sigma_x \in \overline{\mathcal{S}}_x : g_x(\Sigma_x) \leq \rho_x^2\}$  is thus strongly convex in the canonical sense—relative to  $\overline{\mathcal{S}}_x$ —with convexity parameter  $\alpha_x = \kappa_1 / (\sqrt{2\kappa_2\rho_x})$ . This insight implies that  $g_x(\overline{\Sigma}_x) \leq \rho_x^2$ , which in turn ensures that  $\overline{\Sigma}_x \in \mathcal{S}_x^+$ . As  $\theta \in [0, 1]$  was chosen arbitrarily, we may conclude that  $\mathcal{S}_x^+$  is  $\alpha_x$ -strongly convex with respect to  $f$ . Using an analogous argument, one can show that  $\mathcal{S}_w^+$  is  $\alpha_w$ -strongly convex with respect to  $f$ , where  $\alpha_w > 0$  is defined as in the proposition statement. In summary,  $\mathcal{S}_x^+ \times \mathcal{S}_w^+$  is therefore  $\alpha$ -strongly convex with respect to  $f$  in the sense of Assumption 6.1 (ii), where  $\alpha = \min\{\alpha_x, \alpha_w\}$ .

In order to prove assertion (iii), we will establish lower bounds on  $\lambda_{\max}(D_x)$  and  $\lambda_{\max}(D_w)$  uniformly across  $\mathcal{S}_x^+ \times \mathcal{S}_w^+$ . The claim then follows from the observation that  $\lambda_{\max}(D_x) \leq \|D_x\|$  and  $\lambda_{\max}(D_w) \leq \|D_w\|$ . We first derive a lower bound on  $\lambda_{\max}(D_x)$ . To this end, set  $T_1 = I_n - \Sigma_x H^\top G^{-1} H$  and  $T_2 = H \Sigma_x H^\top G^{-1}$ , and note that both  $T_1$  and  $T_2$  have real spectra thanks to [6, Proposition 4.4.10]. As  $D_x = T_1^\top T_1$ , we find

$$(6.20) \quad \lambda_{\max}(D_x) = \lambda_{\max}(T_1 T_1^\top) \geq \max_{\lambda \in \text{spec}(T_1)} |\lambda|^2 = \max_{\lambda \in \text{spec}(T_2)} |1 - \lambda|^2 = \lambda_{\max}^2(I - T_2) = \lambda_{\max}^2(\Sigma_w G^{-1}),$$

where  $\text{spec}(T)$  denotes the eigenvalue spectrum of any square matrix  $T$ . The inequality in (6.20) follows from Browne's theorem [6, Fact 5.11.21], the second equality holds because the non-zero spectrum of  $\Sigma_x H^\top G^{-1} H$  matches that of  $H \Sigma_x H^\top G^{-1}$  thanks to [6, Proposition 4.4.10], and the last equality follows from the identity  $T_2 = I_m - \Sigma_w G^{-1}$ . Notice that all eigenvalues of  $\Sigma_w G^{-1}$  are real because  $T_2$  has a real spectrum.

The estimate (6.20) implies that a uniform lower bound on the largest eigenvalue of  $D_x$  can be obtained by maximizing the largest eigenvalue of  $\Sigma_w G^{-1}$  over  $\mathcal{S}_x^+ \times \mathcal{S}_w^+$ . By the definition of  $G$ , we have

$$(6.21) \quad \begin{aligned} \inf_{\substack{\Sigma_x \in \mathcal{S}_x^+ \\ \Sigma_w \in \mathcal{S}_w^+}} \lambda_{\max}(\Sigma_w G^{-1}) &\geq \inf_{\substack{\Sigma_x \in \mathcal{S}_x^+ \\ \Sigma_w \in \mathcal{S}_w^+}} \lambda_{\min}(\Sigma_w (H \Sigma_x H^\top + \Sigma_w)^{-1}) \\ &\geq \inf_{\substack{\Sigma_x \in \mathcal{S}_x^+ \\ \Sigma_w \in \mathcal{S}_w^+}} \lambda_{\min}(\Sigma_w) \lambda_{\min}((H \Sigma_x H^\top + \Sigma_w)^{-1}) \\ &\geq \inf_{\substack{\Sigma_x \in \mathcal{S}_x^+ \\ \Sigma_w \in \mathcal{S}_w^+}} \frac{\lambda_{\min}(\Sigma_w)}{\lambda_{\max}(\Sigma_x) \lambda_{\max}(H^\top H) + \lambda_{\max}(\Sigma_w)} \\ &\geq \frac{\lambda_{\min}(\widehat{\Sigma}_w)}{(\rho_x + \text{Tr}[\widehat{\Sigma}_x]^{\frac{1}{2}})^2 \lambda_{\max}(H^\top H) + (\rho_w + \text{Tr}[\widehat{\Sigma}_w]^{\frac{1}{2}})^2}, \end{aligned}$$

where the second inequality holds because  $\lambda_{\min}(T) = \lambda_{\max}^{-1}(T^{-1})$  for any  $T \succ 0$  and because the maximum eigenvalue of a positive definite matrix coincides with its operator norm. The third inequality exploits (6.17) and the subadditivity of the maximum eigenvalue, and the last inequality follows from Lemma A.6. Combining (6.20) and (6.21) shows that  $\|D_x\| \geq \lambda_{\max}(D_x) \geq \varepsilon_x$ , where  $\varepsilon_x$  is defined as in the proposition statement.

Using similar arguments, we can also derive a uniform lower bound on  $\lambda_{\max}(D_w)$ . Specifically, we have

$$\begin{aligned}
(6.22) \quad \lambda_{\max}(D_w)\lambda_{\max}(H^\top H) &\geq \lambda_{\max}(H^\top D_w H) = \lambda_{\max}(H^\top G^{-1} H \Sigma_x (H^\top G^{-1} H \Sigma_x)^\top) \\
&\geq \lambda_{\max}^2(H \Sigma_x H^\top G^{-1}) = (1 - \lambda_{\min}(\Sigma_w G^{-1}))^2 \\
&= \left(1 - \frac{1}{1 + \lambda_{\max}(H \Sigma_x H^\top \Sigma_w^{-1})}\right)^2 \\
&= \left(1 - \frac{1}{1 + \lambda_{\max}\left(\Sigma_w^{-\frac{1}{2}} H \Sigma_x H^\top \Sigma_w^{-\frac{1}{2}}\right)}\right)^2,
\end{aligned}$$

where the first inequality follows from (6.17), and the first equality holds due to the definition of  $D_w$ . Moreover, the second inequality exploits Brown's theorem [6, Fact 5.11.21], and the second equality uses the definition of  $G$ . Finally, the third equality follows from the relation  $\lambda_{\min}(\Sigma_w G^{-1}) = \lambda_{\max}^{-1}(G \Sigma_w)$ , and the fourth equality holds due to [6, Proposition 4.4.10]. A uniform lower bound on  $D_w$  can thus be obtained from the estimate

$$\Sigma_w^{-\frac{1}{2}} H \Sigma_x H^\top \Sigma_w^{-\frac{1}{2}} \succeq \lambda_{\min}(\Sigma_x) \Sigma_w^{-\frac{1}{2}} H H^\top \Sigma_w^{-\frac{1}{2}} \succeq \frac{\lambda_{\min}(\Sigma_x)}{\lambda_{\max}(\Sigma_w)} H H^\top,$$

which implies via Lemma A.6 that

$$(6.23) \quad \inf_{\substack{\Sigma_x \in \mathcal{S}_x^+ \\ \Sigma_w \in \mathcal{S}_w^+}} \lambda_{\max}(\Sigma_w^{-\frac{1}{2}} H \Sigma_x H^\top \Sigma_w^{-\frac{1}{2}}) \geq \inf_{\substack{\Sigma_x \in \mathcal{S}_x^+ \\ \Sigma_w \in \mathcal{S}_w^+}} \frac{\lambda_{\min}(\Sigma_x)}{\lambda_{\max}(\Sigma_w)} \lambda_{\max}(H^\top H) \geq \frac{\lambda_{\min}(\widehat{\Sigma}_x)}{(\rho_w + \text{Tr}[\widehat{\Sigma}_w]^{\frac{1}{2}})^2} \lambda_{\max}(H^\top H).$$

Combining (6.22) and (6.23) shows that  $\|D_w\| \geq \lambda_{\max}(D_w) \geq \varepsilon_w$ , where  $\varepsilon_w$  is defined as in the proposition statement. We thus conclude that  $f$  is  $\varepsilon$ -step in the sense of Assumption 6.1 (iii) with  $\varepsilon = \min\{\varepsilon_x, \varepsilon_w\}$ .  $\square$

By Theorem 6.2, which is applicable because of Proposition 6.5, the fully adaptive Frank-Wolfe algorithm (see Algorithm 1) solves the (minimization problem equivalent to the) nonlinear SDP (6.13) at a linear convergence rate. Moreover, Theorem 6.4 ensures that the oracle problem (6.14), which needs to be solved in each iteration of Algorithm 1, can be solved highly efficiently via bisection (see Algorithm 2).

We emphasize that if  $\widehat{\Sigma}_x$  is singular, then the strong convexity parameter  $\alpha$  of Assumption 6.1 (ii) vanishes, and therefore Theorem 6.2 is no longer applicable. In that case, however, Algorithm 1 is still guaranteed to converge, albeit at a sublinear rate; see [58] for further details.

## 7. NUMERICAL EXPERIMENTS

All experiments are run on an Intel XEON CPU with 3.40GHz clock speed and 16GB of RAM. All (linear) SDPs are solved with MOSEK 8 using the YALMIP interface [49]. In order to ensure the reproducibility of our experiments, we make all source codes available at <https://github.com/sorooshafiee/WMMSE>.

### 7.1. Scalability of the Frank-Wolfe Algorithm

We first compare the convergence behavior of the Frank-Wolfe algorithm developed in Section 6 against that of MOSEK. Each experiment consists of 10 independent simulation runs, in all of which we fix the signal and noise dimensions to  $n = m = d$  and the Wasserstein radii to  $\rho_x = \rho_w = \sqrt{d}$  for some  $d \in \mathbb{N}$ . In each


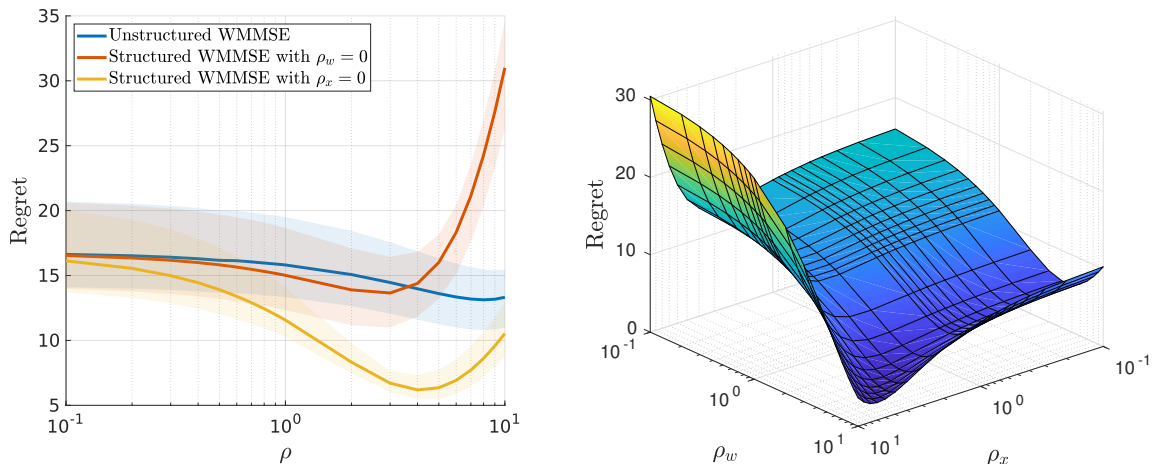


FIGURE 1. Scalability properties of different methods for computing the Wasserstein MMSE estimator (shown are the means (solid lines) and the ranges (shaded areas) of the respective performance measures across 10 simulation runs)

simulation run we randomly generate two nominal covariance matrices  $\widehat{\Sigma}_x$  and  $\widehat{\Sigma}_w$  as follows. First, we sample  $Q_x$  and  $Q_w$  from the standard normal distribution on  $\mathbb{R}^{d \times d}$ , and we denote by  $R_x$  and  $R_w$  the orthogonal matrices whose columns correspond to the orthonormal eigenvectors of  $Q_x + Q_x^\top$  and  $Q_w + Q_w^\top$ , respectively. Then, we set  $\widehat{\Sigma}_x = R_x \Lambda_x (R_x)^\top$  and  $\widehat{\Sigma}_w = R_w \Lambda_w (R_w)^\top$ , where  $\Lambda_x$  and  $\Lambda_w$  are diagonal matrices whose main diagonals are sampled uniformly from  $[1, 5]^d$  and  $[1, 2]^d$ , respectively. Finally, we set  $\widehat{\mu}_x = 0$  and  $\widehat{\mu}_w = 0$ . The Wasserstein MMSE estimator can then be computed either by solving the nonlinear SDP (3.4) with a Frank-Wolfe algorithm or by solving the linear SDP (3.10) with MOSEK. Figures 1(a) and 1(b) show the execution time and the number of iterations needed by the vanilla, adaptive and fully adaptive versions of the Frank-Wolfe (FW) algorithm as well as by MOSEK to drive the (surrogate) duality gap below  $10^{-3}$ . MOSEK runs out of memory for all dimensions  $d > 100$ . Figure 1(c) visualizes the empirical convergence behavior of the three different Frank-Wolfe algorithms. We observe that the fully adaptive Frank-Wolfe algorithm finds highly accurate solutions already after 20 iterations for problem instances of dimension  $d = 1,000$ .

## 7.2. The Value of Structural Information

In the second experiment we study the predictive power of different MMSE estimators. The experiment consists of 100 independent simulation runs. In each run, we use the same procedure as in Section 7.1 to generate two random covariance matrices  $\Sigma_x$  and  $\Sigma_w$  of dimensions  $n = m = 50$ , and we set the true signal and noise distributions to  $\mathbb{P}_x = \mathcal{N}(0, \Sigma_x)$  and  $\mathbb{P}_w = \mathcal{N}(0, \Sigma_w)$ , respectively. Next, we define  $\widehat{\Sigma}_x$  and  $\widehat{\Sigma}_w$  as the sample covariance matrices corresponding to 100 independent samples from  $\mathbb{P}_x$  and  $\mathbb{P}_w$ , respectively. Moreover, we set  $H = I_n$ . To assess the value of structural information, we compare the Wasserstein MMSE estimator proposed in this paper against the Bayesian MMSE estimator associated with the nominal signal and noise distributions and the unstructured Wasserstein MMSE estimator proposed in [64]. The latter uses a single Wasserstein ball to model the ambiguity of the joint distribution of  $x$  and  $y$ , thereby ignoring the structural information that  $w = Hy - x$  is independent of  $x$ . Both the structured and unstructured Wasserstein MMSE estimators collapse to the nominal Bayesian MMSE estimator when the underlying Wasserstein radii are set to zero. Recall also that the nominal Bayesian MMSE estimator is optimal in distributionally robust estimation problems whose ambiguity sets are defined via information divergences [47, 48, 71, 72]. This robustness property makes the nominal Bayesian MMSE estimator an interesting benchmark.



(a) Regret of different Wasserstein MMSE estimators involving a single hyperparameter  $\rho$  (b) Regret of the Wasserstein MMSE estimator involving two hyperparameters  $\rho_x$  and  $\rho_w$

FIGURE 2. Regret of different estimators averaged across 100 simulation runs

We quantify the performance of a given estimator by its *regret*, which is defined as the difference between the estimator’s average risk and the least possible average risk of *any* estimator under the unknown true distributions  $\mathbb{P}_x$  and  $\mathbb{P}_w$ . Note that the minimum average risk is attained by the (affine) Bayesian MMSE estimator corresponding to the (normal) distributions  $\mathbb{P}_x$  and  $\mathbb{P}_w$ . Figure 2(a) shows the regret of the Wasserstein MMSE estimator with  $\rho_x = \rho$  and  $\rho_w = 0$ , the Wasserstein MMSE estimator with  $\rho_x = 0$  and  $\rho_w = \rho$  and the unstructured Wasserstein MMSE estimator from [64] with Wasserstein radius  $\rho$  for all  $\rho \in [0.1, 10]$ . The solid lines represent the averages and the shaded areas the ranges of the regret across all 100 simulation runs. The regret of the structured Wasserstein MMSE estimator with Wasserstein radii  $\rho_x, \rho_w \in [0.1, 10]$  averaged across all 100 simulation runs is visualized by the surface plot in Figure 2(b).

We observe that the average regret of the nominal Bayesian MMSE estimator amounts to 16.7 (the leftmost value of all curves in Figure 2(a)), while the best unstructured Wasserstein MMSE estimator attains a significantly lower average regret of 13.1 (the minimum of the blue curve in Figure 2(a)). The best structured Wasserstein MMSE estimator without noise ambiguity ( $\rho_w = 0$ ) displays a similar performance, attaining an average regret of 13.2 (the minimum of the red curve in Figure 2(a)), while the one without signal ambiguity ( $\rho_x = 0$ ) further reduces the average regret by more than 50% to 6.0 (the minimum of the yellow curve in Figure 2(a)). Finally, the best among *all* structured Wasserstein MMSE estimators, which is obtained by tuning both  $\rho_x$  and  $\rho_w$ , attains an even lower average regret of 2.2 (the minimum of the surface plot in Figure 2(b)). This experiment confirms our hypothesis that structural (independence) information as well as information about distributional ambiguity can improve the predictive power of MMSE estimators.

**Acknowledgments.** We are grateful to Erick Delage for valuable comments on an earlier version of this paper. This research was supported by the Swiss National Science Foundation grant number BSCGIO\_157733.



## APPENDIX A. AUXILIARY RESULTS

We first prove Pythagoras' theorem for the Wasserstein distance.

*Proof of Lemma 1.2.* By the definition of the Wasserstein distance and by Pythagoras' theorem for the Euclidean distance we have

$$\begin{aligned} \mathbb{W}(\mathbb{Q}_x^1 \times \mathbb{Q}_w^1, \mathbb{Q}_x^2 \times \mathbb{Q}_w^2)^2 &= \inf_{\pi \in \Pi(\mathbb{Q}_x^1 \times \mathbb{Q}_w^1, \mathbb{Q}_x^2 \times \mathbb{Q}_w^2)} \int_{\mathbb{R}^{n+m} \times \mathbb{R}^{n+m}} \|x_1 - x_2\|^2 + \|w_1 - w_2\|^2 \pi(dx_1, dw_1, dx_2, dw_2) \\ &\leq \inf_{\pi_x \in \Pi(\mathbb{Q}_x^1, \mathbb{Q}_x^2)} \int_{\mathbb{R}^n \times \mathbb{R}^n} \|x_1 - x_2\|^2 \pi_x(dx_1, dx_2) + \inf_{\pi_w \in \Pi(\mathbb{Q}_w^1, \mathbb{Q}_w^2)} \int_{\mathbb{R}^m \times \mathbb{R}^m} \|w_1 - w_2\|^2 \pi_w(dw_1, dw_2) \\ &= \mathbb{W}(\mathbb{Q}_x^1, \mathbb{Q}_x^2)^2 + \mathbb{W}(\mathbb{Q}_w^1, \mathbb{Q}_w^2)^2, \end{aligned}$$

where the inequality follows from the restriction to factorizable transportation plans of the form  $\pi = \pi_x \times \pi_w$  for some  $\pi_x \in \Pi(\mathbb{Q}_x^1, \mathbb{Q}_x^2)$  and  $\pi_w \in \Pi(\mathbb{Q}_w^1, \mathbb{Q}_w^2)$ . To prove the converse inequality, we define  $\Pi_x(\mathbb{Q}_x^1, \mathbb{Q}_x^2)$  as the set of all joint distributions  $\pi \in \mathcal{M}(\mathbb{R}^{n+m} \times \mathbb{R}^{n+m})$  of  $(x_1, w_1) \in \mathbb{R}^{n+m}$  and  $(x_2, w_2) \in \mathbb{R}^{n+m}$  under which  $x_1$  and  $x_2$  have marginal distributions  $\mathbb{Q}_x^1$  and  $\mathbb{Q}_x^2$ , respectively. Similarly, we define  $\Pi_w(\mathbb{Q}_w^1, \mathbb{Q}_w^2)$  as the set of all joint distributions  $\pi \in \mathcal{M}(\mathbb{R}^{n+m} \times \mathbb{R}^{n+m})$  of  $(x_1, w_1) \in \mathbb{R}^{n+m}$  and  $(x_2, w_2) \in \mathbb{R}^{n+m}$  under which  $w_1$  and  $w_2$  have marginal distributions  $\mathbb{Q}_w^1$  and  $\mathbb{Q}_w^2$ , respectively. Using this notation, we find

$$\begin{aligned} \mathbb{W}(\mathbb{Q}_x^1 \times \mathbb{Q}_w^1, \mathbb{Q}_x^2 \times \mathbb{Q}_w^2)^2 &\geq \inf_{\pi \in \Pi(\mathbb{Q}_x^1 \times \mathbb{Q}_w^1, \mathbb{Q}_x^2 \times \mathbb{Q}_w^2)} \int_{\mathbb{R}^{n+m} \times \mathbb{R}^{n+m}} \|x_1 - x_2\|^2 \pi(dx_1, dw_1, dx_2, dw_2) \\ &\quad + \inf_{\pi \in \Pi(\mathbb{Q}_x^1 \times \mathbb{Q}_w^1, \mathbb{Q}_x^2 \times \mathbb{Q}_w^2)} \int_{\mathbb{R}^{n+m} \times \mathbb{R}^{n+m}} \|w_1 - w_2\|^2 \pi(dx_1, dw_1, dx_2, dw_2) \\ &\geq \inf_{\pi \in \Pi_x(\mathbb{Q}_x^1, \mathbb{Q}_x^2)} \int_{\mathbb{R}^{n+m} \times \mathbb{R}^{n+m}} \|x_1 - x_2\|^2 \pi(dx_1, dw_1, dx_2, dw_2) \\ &\quad + \inf_{\pi \in \Pi_w(\mathbb{Q}_w^1, \mathbb{Q}_w^2)} \int_{\mathbb{R}^{n+m} \times \mathbb{R}^{n+m}} \|w_1 - w_2\|^2 \pi(dx_1, dw_1, dx_2, dw_2) \\ &= \mathbb{W}(\mathbb{Q}_x^1, \mathbb{Q}_x^2)^2 + \mathbb{W}(\mathbb{Q}_w^1, \mathbb{Q}_w^2)^2, \end{aligned}$$

where the first inequality exploits the superadditivity of the infimum operator, while the second inequality holds because  $\Pi(\mathbb{Q}_x^1 \times \mathbb{Q}_w^1, \mathbb{Q}_x^2 \times \mathbb{Q}_w^2)$  contains both  $\Pi_x(\mathbb{Q}_x^1, \mathbb{Q}_x^2)$  and  $\Pi_w(\mathbb{Q}_w^1, \mathbb{Q}_w^2)$  as subsets. The equality in the last line follows from the observation that for any  $\pi \in \Pi_x(\mathbb{Q}_x^1, \mathbb{Q}_x^2)$  the marginal distribution  $\pi_x$  of  $(x_1, x_2)$  is an element of  $\Pi(\mathbb{Q}_x^1, \mathbb{Q}_x^2)$ , and for any  $\pi \in \Pi_w(\mathbb{Q}_w^1, \mathbb{Q}_w^2)$  the marginal distribution  $\pi_w$  of  $(w_1, w_2)$  is an element of  $\Pi(\mathbb{Q}_w^1, \mathbb{Q}_w^2)$ . Thus, the claim follows.  $\square$

In order to prove Proposition 2.6 in the main text, we establish first general conditions for the solvability and stability of parametric minimax problems and prove that the matrix square root is Hölder continuous.

**Lemma A.1** (Parametric minimax problems). *Consider the parametric minimax problem*

$$(A.1) \quad J(\theta) = \inf_{u \in \mathbb{U}} \sup_{v \in V(\theta)} f(u, v),$$

where  $\mathbb{U}$ ,  $\mathbb{V}$  and  $\Theta$  are non-empty convex subsets of Euclidean spaces equipped with the respective subspace topologies,  $f : \mathbb{U} \times \mathbb{V} \rightarrow \mathbb{R}$  and  $g : \mathbb{V} \times \Theta \rightarrow \mathbb{R}$  are continuous functions, while  $V : \Theta \rightrightarrows \mathbb{V}$  and  $V^\circ : \Theta \rightrightarrows \mathbb{V}$  are set-valued mappings defined through  $V(\theta) = \{v \in \mathbb{V} : g(v, \theta) \leq 0\}$  and  $V^\circ(\theta) = \{v \in \mathbb{V} : g(v, \theta) < 0\}$ , respectively. Assume that for each  $\theta' \in \Theta$  there exists a compact neighborhood  $\Theta' \subseteq \Theta$  of  $\theta'$  such that

- (i)  $V(\theta)$  is convex and bounded uniformly across all  $\theta \in \Theta'$ ;

- (ii)  $V^\circ(\theta)$  is non-empty and coincides with the interior of  $V(\theta)$  for all  $\theta \in \Theta'$ ;
- (iii) there exist  $v' \in \mathbb{V}$  and  $J' \in \mathbb{R}$  such that  $v' \in V(\theta)$  and  $J(\theta) \leq J'$  for all  $\theta \in \Theta'$  and such that the set  $\mathbb{U}' = \{u \in \mathbb{U} : f(u, v') \leq J'\}$  is non-empty and compact.

Then, the minimax problem (A.1) is solvable for all  $\theta \in \Theta$ , and  $J$  is continuous on  $\Theta$ .

*Proof of Lemma A.1.* Define  $F(u, \theta) = \sup_{v \in V(\theta)} f(u, v)$ , and note that  $F(u, \theta)$  is finite because it is the maximum of a continuous function over a compact feasible set. As  $V$  is locally bounded thanks to assumption (i) and as the graph of  $V$  is closed thanks to the continuity of  $g$ , the closed graph theorem [1, Proposition 1.4.8] ensures that the set-valued mapping  $V$  is upper semicontinuous in the sense of [5, Chapter VI]. By [5, Theorem 2, p. 116], the optimal value function  $F$  is thus upper semicontinuous on  $\mathbb{U} \times \Theta$ . Next, note that  $F(u, \theta) = \sup_{v \in V^\circ(\theta)} f(u, v)$  because  $f$  is continuous and because  $V^\circ(\theta)$  coincides with the non-empty interior of  $V(\theta)$  thanks to assumption (ii). As  $V^\circ$  is convex-valued thanks to assumption (i) and as the graph of  $V$  is open thanks to the continuity of  $g$ , the open graph theorem [70, Proposition 2] implies that the set-valued mapping  $V$  is lower semicontinuous in the sense of [5, Chapter VI]. By [5, Theorem 1, p. 115], the function  $F$  is thus lower semicontinuous on  $\mathbb{U} \times \Theta$ . In summary,  $F$  is therefore continuous on  $\mathbb{U} \times \Theta$ .

To prove that  $J$  is lower semicontinuous, we select an arbitrary  $\theta' \in \Theta$  and a compact neighborhood  $\Theta' \subseteq \Theta$  of  $\theta'$  that satisfies assumption (iii). For any  $\theta \in \Theta'$ , we may then restrict the feasible set of the outer minimization problem in (A.1) to  $\mathbb{U}'$  without affecting  $J(\theta)$ . Indeed, any  $u \notin \mathbb{U}'$  and  $\theta \in \Theta'$  satisfies

$$F(u, \theta) = \sup_{v \in V(\theta)} f(u, v) \geq f(u, v') > J' \geq J(\theta),$$

where the three inequalities hold because  $v' \in V(\theta)$  for all  $\theta \in \Theta'$ ,  $u \notin \mathbb{U}'$  and  $J(\theta) \leq J'$  for all  $\theta \in \Theta'$ , respectively. Thus,  $u \notin \mathbb{U}'$  is strictly suboptimal, and we have  $J(\theta) = \inf_{u \in \mathbb{U}'} F(u, \theta)$  for all  $\theta \in \Theta'$ . As  $\mathbb{U}'$  is compact and  $F(u, \theta)$  is continuous, the restricted minimax problem with  $\mathbb{U}'$  replacing  $\mathbb{U}$  is solvable, and any minimizer also solves (A.1). Moreover, the constant feasible set mapping that assigns each  $\theta \in \Theta'$  the same set  $\mathbb{U}'$  is continuous in the sense of [5, Chapter VI]. By Berge's maximum theorem [5, Theorem 1, p. 115], the function  $J$  is thus continuous on  $\Theta'$  and, as  $\theta' \in \Theta$  was chosen arbitrarily, on all of  $\Theta$ .  $\square$

**Lemma A.2** (Hölder continuity of the matrix square root). *The square root of a positive semidefinite matrix is Hölder-continuous with exponent 1/2. More precisely, we have*

$$\|\sqrt{A_1} - \sqrt{A_2}\| \leq 2\sqrt{d} \|A_1 - A_2\|^{1/2} \quad \forall A_1, A_2 \in \mathbb{S}_+^d.$$

*Proof of Lemma A.2.* The proof exploits the following two facts.

- (i) By [62, Lemma 2.2], we have

$$\|\sqrt{A_1} - \sqrt{A_2}\| \leq \frac{1}{\lambda_{\min}(A_1) + \lambda_{\min}(A_2)} \|A_1 - A_2\| \quad \forall A_1, A_2 \in \mathbb{S}_{++}^d.$$

- (ii) One can show that

$$\|\sqrt{A + \varepsilon I_d} - \sqrt{A}\| \leq d\sqrt{\varepsilon} \quad \forall A \in \mathbb{S}_+^d, \varepsilon \geq 0.$$

To prove assertion (ii), denote by  $\lambda_i \geq 0$ ,  $i = 1, \dots, d$ , the eigenvalues of  $A \in \mathbb{S}_+^d$  and by  $a_i \in \mathbb{R}^d$ ,  $i = 1, \dots, d$ , the corresponding orthonormal eigenvectors, which implies that  $A = \sum_{i=1}^d \lambda_i a_i a_i^\top$ . Thus, we find

$$\begin{aligned} \|\sqrt{A + \varepsilon I_d} - \sqrt{A}\| &= \left\| \sum_{i=1}^d (\sqrt{\lambda_i + \varepsilon} - \sqrt{\lambda_i}) a_i a_i^\top \right\| = \left\| \sum_{i=1}^d \varepsilon (\sqrt{\lambda_i + \varepsilon} + \sqrt{\lambda_i})^{-1} a_i a_i^\top \right\| \\ &\leq \sum_{i=1}^d \frac{\varepsilon}{\sqrt{\lambda_i + \varepsilon} + \sqrt{\lambda_i}} \|a_i a_i^\top\| \leq d\sqrt{\varepsilon}, \end{aligned}$$

where the first inequality follows from the triangle inequality, and the second inequality holds because the denominator of the  $i$ -th fraction grows with  $\lambda_i$  and is therefore minimal for  $\lambda_i = 0$ .

To prove the lemma, select now arbitrary  $A_1, A_2 \in \mathbb{S}_+^d$  and  $\varepsilon > 0$ . The triangle inequality implies that

$$\begin{aligned} \|\sqrt{A_1} - \sqrt{A_2}\| &\leq \|\sqrt{A_1} - \sqrt{A_1 + \varepsilon I_d}\| + \|\sqrt{A_1 + \varepsilon I_d} - \sqrt{A_2 + \varepsilon I_d}\| + \|\sqrt{A_2 + \varepsilon I_d} - \sqrt{A_2}\| \\ &\leq d\sqrt{\varepsilon} + \frac{1}{\lambda_{\min}(\sqrt{A_1 + \varepsilon I_d}) + \lambda_{\min}(\sqrt{A_2 + \varepsilon I_d})} \|A_1 - A_2\| + d\sqrt{\varepsilon} \leq 2d\sqrt{\varepsilon} + \frac{\|A_1 - A_2\|}{2\sqrt{\varepsilon}}, \end{aligned}$$

where the second exploits the assertions (i) and (ii). The claim now follows by setting  $\varepsilon = \|A_1 - A_2\|/4d$ , which actually minimizes the right hand side of the last expression.  $\square$

Armed with Lemmas A.1 and A.2, we are now ready to prove Proposition 2.6.

*Proof of Proposition 2.6.* The Gelbrich MMSE estimation problem (2.1) upper bounds the Wasserstein MMSE estimation problem (1.7) because  $\mathcal{A} \subseteq \mathcal{F}$  and  $\mathbb{G}(\widehat{\mathbb{P}}) \supseteq \mathbb{B}(\widehat{\mathbb{P}})$ ; see Corollary 2.5. Thus, assertion (i) follows. Recall now that any  $\psi \in \mathcal{A}$  can be represented as  $\psi(y) = Ay + b$  for some  $A \in \mathbb{R}^{n \times m}$  and  $b \in \mathbb{R}^n$ . Moreover, for any distribution  $\mathbb{Q} = \mathbb{Q}_x \times \mathbb{Q}_w \in \mathbb{G}(\widehat{\mathbb{P}})$ , denote by  $\mu_x$  and  $\mu_w$  the mean vectors and by  $\Sigma_x$  and  $\Sigma_w$  the covariance matrices of  $\mathbb{Q}_x$  and  $\mathbb{Q}_w$ , respectively. Hence, the objective function and the constraints of (2.1) depend on  $\psi$  and  $\mathbb{Q}$  only through  $u = (A, b)$ , which ranges over  $\mathbb{U} = \mathbb{R}^{n \times m} \times \mathbb{R}^n$ , and  $v = (\mu_x, \mu_w, \Sigma_x, \Sigma_w)$ , which ranges over  $\mathbb{V} = \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{S}_+^n \times \mathbb{S}_+^m$ . Indeed, the average risk of  $\psi$  under  $\mathbb{Q}$  satisfies

$$\begin{aligned} \mathcal{R}(\psi, \mathbb{Q}) &= \mathbb{E}_{\mathbb{Q}} [\|x - A(Hx + w) - b\|^2] = \mathbb{E}_{\mathbb{Q}} [\|(I_n - AH)x - Aw - b\|^2] \\ &= \langle (I_n - AH)^\top (I_n - AH), \Sigma_x + \mu_x \mu_x^\top \rangle + \langle A^\top A, \Sigma_w + \mu_w \mu_w^\top \rangle + b^\top b \\ &\quad - 2\mu_x^\top (I_n - AH)^\top A \mu_w - 2b^\top ((I_n - AH)\mu_x - A\mu_w) = f(u, v). \end{aligned}$$

Similarly, the constraint  $\mathbb{Q} \in \mathbb{G}(\widehat{\mathbb{P}})$  can be reformulated as  $g(v, \theta) \leq 0$ , where  $\theta = (\widehat{\mu}_x, \widehat{\mu}_w, \widehat{\Sigma}_x, \widehat{\Sigma}_w)$  is a shorthand for the problem's input parameters ranging over the set  $\Theta = \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{S}_+^n \times \mathbb{S}_+^m$ , and where

$$g(v, \theta) = \max \left\{ \mathbb{G}((\mu_x, \Sigma_x), (\widehat{\mu}_x, \widehat{\Sigma}_x))^2 - \rho_x^2, \mathbb{G}((\mu_w, \Sigma_w), (\widehat{\mu}_w, \widehat{\Sigma}_w))^2 - \rho_w^2 \right\}.$$

Thus, the Gelbrich MMSE estimation problem (2.1) can be re-expressed more concisely as

$$(A.2) \quad J(\theta) = \inf_{u \in \mathbb{U}} \sup_{v \in V(\theta)} f(u, v),$$

where  $V(\theta) = \{v \in \mathbb{V} : g(v, \theta) \leq 0\}$  for all  $\theta \in \Theta$ . Assume now that  $\rho_x > 0$  and  $\rho_w > 0$ . In the remainder we will prove that the optimal value  $J(\theta)$  of the minimax problem (A.2) is attained and continuous in  $\theta$  by showing that all assumptions of Lemma A.1 in the appendix are satisfied.

To this end, note first that  $f$  is continuous by construction and that  $g$  is convex and continuous thanks to Proposition 2.3. Next, select an arbitrary reference point  $\theta' = (\widehat{\mu}'_x, \widehat{\mu}'_w, \widehat{\Sigma}'_x, \widehat{\Sigma}'_w) \in \Theta$ , and define

$$\Theta' = \left\{ (\widehat{\mu}_x, \widehat{\mu}_w, \widehat{\Sigma}_x, \widehat{\Sigma}_w) \in \Theta : \begin{array}{l} \mathbb{G}((\widehat{\mu}_x, \widehat{\Sigma}_x), (\widehat{\mu}'_x, \widehat{\Sigma}'_x)) \leq \frac{\rho_x}{2} \\ \mathbb{G}((\widehat{\mu}_w, \widehat{\Sigma}_w), (\widehat{\mu}'_w, \widehat{\Sigma}'_w)) \leq \frac{\rho_w}{2} \end{array} \right\}.$$

Note that  $\Theta'$  is a neighborhood of  $\theta'$  because  $\rho_x$  and  $\rho_w$  are strictly positive and because Proposition 2.3 ensures that the (squared) Gelbrich distance is continuous. Moreover,  $\Theta'$  is compact due to Lemma A.6.

In order to verify assumption (i) of Lemma A.1, we note first that  $V(\theta)$  is a convex set for every  $\theta \in \Theta$  because  $g$  is a convex function. Moreover, we have

$$\begin{aligned} V(\theta) &= \left\{ (\mu_x, \mu_w, \Sigma_x, \Sigma_w) \in \mathbb{V} : \begin{array}{l} \mathbb{G}((\mu_x, \Sigma_x), (\hat{\mu}_x, \hat{\Sigma}_x)) \leq \rho_x \\ \mathbb{G}((\mu_w, \Sigma_w), (\hat{\mu}_w, \hat{\Sigma}_w)) \leq \rho_w \end{array} \right\} \\ &\subseteq \left\{ (\mu_x, \mu_w, \Sigma_x, \Sigma_w) \in \mathbb{V} : \begin{array}{l} \mathbb{G}((\mu_x, \Sigma_x), (\hat{\mu}'_x, \hat{\Sigma}'_x)) \leq \frac{3\rho_x}{2} \\ \mathbb{G}((\mu_w, \Sigma_w), (\hat{\mu}'_w, \hat{\Sigma}'_w)) \leq \frac{3\rho_w}{2} \end{array} \right\} = V' \quad \forall \theta \in \Theta', \end{aligned}$$

where the first equality holds due to the definition of  $V(\theta)$ , the inclusion holds due to the definition of  $\Theta'$  and because the Gelbrich distance satisfies the triangle inequality, and the last equality defines the set  $V'$ , which is compact thanks to Lemma A.6. This reasoning shows that  $V(\theta)$  is uniformly bounded on  $\Theta'$ .

In order to verify assumption (ii) of Lemma A.1, define

$$V^\circ(\theta) = \{v \in \mathbb{V} : g(v, \theta) < 0\} = \left\{ (\mu_x, \mu_w, \Sigma_x, \Sigma_w) \in \mathbb{V} : \begin{array}{l} \mathbb{G}((\mu_x, \Sigma_x), (\hat{\mu}_x, \hat{\Sigma}_x)) < \rho_x \\ \mathbb{G}((\mu_w, \Sigma_w), (\hat{\mu}_w, \hat{\Sigma}_w)) < \rho_w \end{array} \right\}$$

for any  $\theta \in \Theta$ . As the Gelbrich distance satisfies the identity of indiscernibles and as  $\rho_x$  and  $\rho_w$  are strictly positive, we have  $\theta \in V^\circ(\theta)$ , which implies that  $V^\circ(\theta)$  is non-empty. It remains to be shown that  $V^\circ(\theta)$  coincides with the interior of  $V(\theta)$ . As the Gelbrich distance is continuous by virtue of Proposition 2.3, it is clear that  $V^\circ(\theta)$  is an open subset of  $V(\theta)$  and thus contained in  $\text{int}(V(\theta))$ . To prove the converse inclusion, assume for the sake of argument that  $V^\circ(\theta)$  is a strict subset of  $\text{int}(V(\theta))$ . Thus, there must exist an open set  $O \subseteq \text{int}(V(\theta)) \setminus V^\circ(\theta)$ . Otherwise, each point in  $\text{int}(V(\theta)) \setminus V^\circ(\theta)$  would belong to the boundary of  $\text{int}(V(\theta))$ , which is impossible because  $\text{int}(V(\theta))$  is open. As  $O \subseteq V(\theta) \setminus V^\circ(\theta)$ , it is clear that at least one of the equalities  $\mathbb{G}((\mu_x, \Sigma_x), (\hat{\mu}_x, \hat{\Sigma}_x)) = \rho_x$  or  $\mathbb{G}((\mu_w, \Sigma_w), (\hat{\mu}_w, \hat{\Sigma}_w)) = \rho_w$  is satisfied at any point in  $O$ . In fact, as the Gelbrich distance is continuous, one of these equalities holds throughout an open set  $O' \subseteq O$ . Without loss of generality we may thus assume that  $\mathbb{G}((\mu_x, \Sigma_x), (\hat{\mu}_x, \hat{\Sigma}_x))^2 = \rho_x^2$  on an open set  $O' \subseteq O$ , which implies that any point in  $O'$  is a local minimizer of the squared Gelbrich distance. As the squared Gelbrich distance is convex due to Proposition 2.3, this means that all points in  $O'$  are in fact global minimizers. This is not possible, however, because the (squared) Gelbrich distance adopts its global minimum only at points where  $\mu_x = \hat{\mu}_x$  and  $\Sigma_x = \hat{\Sigma}_x$ . No such point belongs to  $O'$  because  $O' \cap V^\circ(\theta) = \emptyset$ . Thus, we have  $\text{int}(V(\theta)) = V^\circ(\theta)$ .

In order to verify assumption (iii) of Lemma A.1, select a point  $v' = (\mu'_x, \mu'_w, \Sigma'_x, \Sigma'_w) \in \Theta'$  with  $\Sigma'_w \succ 0$ . Such a point exists because  $\rho_w > 0$  and because the (squared) Gelbrich distance is continuous by virtue of Proposition 2.3, which implies that

$$(\mu'_x, \mu'_w, \Sigma'_x, \Sigma'_w) = (\hat{\mu}_x, \hat{\mu}_w, \hat{\Sigma}_x, \hat{\Sigma}_w + \lambda \cdot I_m) \in \Theta'$$

for all sufficiently small  $\lambda > 0$ . The triangle inequality for the Gelbrich distance then ensures that  $v' \in V(\theta)$  for all  $\theta \in \Theta'$ . Next, set  $J' = \sup_{v \in V'} f(0, v)$ , which is finite because  $V'$  is compact and  $f$  is continuous, and note that

$$J(\theta) = \inf_{u \in \mathbb{U}} \sup_{v \in V(\theta)} f(u, v) \leq \sup_{v \in V'} f(0, v) = J' \quad \forall \theta \in \Theta',$$

where the inequality holds because  $V(\theta) \subseteq V'$  whenever  $\theta \in \Theta'$ . Finally, define  $\mathbb{U}' = \{u \in \mathbb{U} : f(u, v') \leq J'\}$ , which is non-empty. Indeed,  $\mathbb{U}'$  contains at least the point  $u = 0$  because  $v' \in \Theta' \subseteq V'$ . As  $\Sigma'_w \succ 0$ , it is easy to verify that  $f(u, v')$  is strictly convex and quadratic in  $u$ , which implies that  $\mathbb{U}'$  is a compact ellipsoid.

In summary, problem (A.2), which is equivalent to the Gelbrich MMSE estimation problem (2.1), satisfies all assumptions of Lemma A.1. Therefore, its optimal value  $J(\theta)$  is attained and continuous in  $\theta$ .  $\square$

In order to derive a tractable reformulation for the Gelbrich MMSE estimation problem studied in Section 2, we need to be able to solve nonlinear SDPs of the form

$$(A.3) \quad J^* = \sup_{\Sigma \succeq 0} \langle D, \Sigma \rangle - \gamma \operatorname{Tr} \left[ \Sigma - 2 \left( \widehat{\Sigma}^{\frac{1}{2}} \Sigma \widehat{\Sigma}^{\frac{1}{2}} \right)^{\frac{1}{2}} \right]$$

parameterized by  $D \in \mathbb{S}^d$ ,  $\widehat{\Sigma} \in \mathbb{S}_+^d$  and  $\gamma \in \mathbb{R}_+$ . It is known that, under certain regularity conditions, problem (A.3) admits a unique optimal solution that is available in closed form [54]. In the following we review the construction of this optimal solution under slightly more general conditions.

**Proposition A.3** (Closed form solution of (A.3)). *For any  $D \in \mathbb{S}^d$ ,  $\widehat{\Sigma} \in \mathbb{S}_+^d$  and  $\gamma \in \mathbb{R}_+$  the optimal value of the nonlinear SDP (A.3) is given by*

$$J^* = \begin{cases} \gamma^2 \langle (\gamma I_d - D)^{-1}, \widehat{\Sigma} \rangle & \text{if } \gamma > \lambda_{\max}(D), \\ \liminf_{\bar{\gamma} \downarrow \gamma} \bar{\gamma}^2 \langle (\bar{\gamma} I_d - D)^{-1}, \widehat{\Sigma} \rangle & \text{if } \gamma = \lambda_{\max}(D), \\ +\infty & \text{if } \gamma < \lambda_{\max}(D). \end{cases}$$

Moreover, problem (A.3) is solved by  $\Sigma^* = \gamma^2 (\gamma I_d - D)^{-1} \widehat{\Sigma} (\gamma I_d - D)^{-1}$  whenever  $\gamma > \lambda_{\max}(D)$ . This solution is unique if  $\widehat{\Sigma} \succ 0$ .

*Proof of Proposition A.3.* Assume first that  $\gamma > \lambda_{\max}(D)$ . Moreover, in order to simplify the proof, assume temporarily that  $\widehat{\Sigma} \succ 0$ . By applying the nonlinear variable transformation  $B \leftarrow \left( \widehat{\Sigma}^{\frac{1}{2}} \Sigma \widehat{\Sigma}^{\frac{1}{2}} \right)^{\frac{1}{2}}$ , which implies that  $\Sigma = \widehat{\Sigma}^{-\frac{1}{2}} B^2 \widehat{\Sigma}^{-\frac{1}{2}}$ , we can reformulate problem (A.3) as

$$\begin{aligned} J^* &= \sup_{B \succeq 0} \langle D, \widehat{\Sigma}^{-\frac{1}{2}} B^2 \widehat{\Sigma}^{-\frac{1}{2}} \rangle - \gamma \operatorname{Tr} \left[ \widehat{\Sigma}^{-\frac{1}{2}} B^2 \widehat{\Sigma}^{-\frac{1}{2}} - 2B \right] \\ &= \sup_{B \succeq 0} \langle B^2, \widehat{\Sigma}^{-\frac{1}{2}} (D - \gamma I_d) \widehat{\Sigma}^{-\frac{1}{2}} \rangle + 2\gamma \langle B, I_d \rangle, \end{aligned}$$

where the second equality exploits the cyclicity of the trace operator. Introducing the auxiliary parameter  $\Delta = \widehat{\Sigma}^{-\frac{1}{2}} (D - \gamma I_d) \widehat{\Sigma}^{-\frac{1}{2}}$ , we can then rewrite the last maximization problem over  $B$  more concisely as

$$(A.4) \quad J^* = \sup_{B \succeq 0} \langle B^2, \Delta \rangle + 2\gamma \langle B, I_d \rangle.$$

Note that (A.4) represents a convex maximization problem because  $\gamma > \lambda_{\max}(D)$  and  $\widehat{\Sigma} \succ 0$ , which imply that  $\Delta \prec 0$ . Ignoring the positive semidefiniteness constraint on  $B$ , the objective function of (A.4) is uniquely minimized by the solution  $B^* = -\gamma \Delta^{-1}$  of the first-order optimality condition  $B\Delta + \Delta B + 2\gamma I_d = 0$ . Uniqueness follows from [36, Theorem 12.5]. As it is strictly positive definite,  $B^*$  thus uniquely solves (A.4), which in turn implies that  $\Sigma^* = \widehat{\Sigma}^{-\frac{1}{2}} (B^*)^2 \widehat{\Sigma}^{-\frac{1}{2}} = \gamma^2 (\gamma I_d - D)^{-1} \widehat{\Sigma} (\gamma I_d - D)^{-1}$  uniquely solves (A.3). Substituting  $\Sigma^*$  back into the objective function of (A.3) further shows that  $J^* = \gamma^2 \langle (\gamma I_d - D)^{-1}, \widehat{\Sigma} \rangle$ .

Next, we will argue that the analytical formula for  $J^*$  in the regime  $\gamma > \lambda_{\max}(D)$  remains valid even when  $\widehat{\Sigma}$  is rank-deficient. To see this, define

$$J^*(\widehat{\Sigma}) = \gamma^2 \langle (\gamma I_d - D)^{-1}, \widehat{\Sigma} \rangle \quad \text{and} \quad \Sigma^*(\widehat{\Sigma}) = \gamma^2 (\gamma I_d - D)^{-1} \widehat{\Sigma} (\gamma I_d - D)^{-1}$$

as explicit continuous functions of the parameter  $\widehat{\Sigma} \in \mathbb{S}_+^d$ . Similarly, define the function

$$F(\Sigma, \widehat{\Sigma}) = \langle D, \Sigma \rangle - \gamma \operatorname{Tr} \left[ \Sigma - 2 \left( \widehat{\Sigma}^{\frac{1}{2}} \Sigma \widehat{\Sigma}^{\frac{1}{2}} \right)^{\frac{1}{2}} \right],$$

which is jointly continuous in  $\Sigma \in \mathbb{S}_+^d$  and  $\widehat{\Sigma} \in \mathbb{S}_+^d$ . We then have

$$J^*(\widehat{\Sigma}) = \liminf_{\varepsilon \downarrow 0} J^*(\widehat{\Sigma} + \varepsilon I_d) = \liminf_{\varepsilon \downarrow 0} \sup_{\Sigma \succeq 0} F(\Sigma, \widehat{\Sigma} + \varepsilon I_d) \geq \sup_{\Sigma \succeq 0} F(\Sigma, \widehat{\Sigma}) \geq F(\Sigma^*(\widehat{\Sigma}), \widehat{\Sigma}) = J^*(\widehat{\Sigma}),$$

where the first equality follows from the continuity of  $J^*(\widehat{\Sigma})$ , while the second equality holds because  $\widehat{\Sigma} + \varepsilon I_d \succ 0$  for every  $\varepsilon > 0$  and because  $J^*(\widehat{\Sigma}') = \sup_{\Sigma \succ 0} F(\Sigma, \widehat{\Sigma}')$  for every  $\widehat{\Sigma}' \succ 0$ , which was established in the first part of the proof. The first inequality exploits the fact that a pointwise supremum of continuous functions is lower semicontinuous, and the second inequality holds because  $\Sigma^*(\widehat{\Sigma} + \varepsilon I_d) \succ 0$  for every  $\varepsilon > 0$ . Finally, the last equality follows from elementary algebra. The above arguments imply that  $J^*(\widehat{\Sigma})$  and  $\Sigma^*(\widehat{\Sigma})$  represent the optimal value and an optimal solution of (A.3), respectively, even if  $\widehat{\Sigma} \in \mathbb{S}_+^d$  is rank-deficient.

Assume next that  $\gamma < \lambda_{\max}(D)$ , and denote by  $\bar{v} \in \mathbb{R}^d$  a normalized eigenvector of  $D$  corresponding to the eigenvalue  $\lambda_{\max}(D)$ . By optimizing only over matrices of the form  $\Sigma = t \bar{v} \bar{v}^\top$  for some  $t \geq 0$ , we find

$$\begin{aligned} J^* &\geq \sup_{t \geq 0} t \langle D - \gamma I_d, \bar{v} \bar{v}^\top \rangle + 2\sqrt{t} \gamma \operatorname{Tr} \left[ (\widehat{\Sigma}^{\frac{1}{2}} \bar{v} \bar{v}^\top \widehat{\Sigma}^{\frac{1}{2}})^{\frac{1}{2}} \right] \\ &= \sup_{t \geq 0} t (\lambda_{\max}(D) - \gamma) + 2\sqrt{t} \gamma \operatorname{Tr} \left[ (\widehat{\Sigma}^{\frac{1}{2}} \bar{v} \bar{v}^\top \widehat{\Sigma}^{\frac{1}{2}})^{\frac{1}{2}} \right] = \infty. \end{aligned}$$

Assume finally that  $\gamma = \lambda_{\max}(D)$ . To investigate this limiting case, note that the objective function of (A.3) is linear in  $\gamma$ , which implies that the optimal value of (A.3) is convex and lower semicontinuous in  $\gamma$ . Given the results for  $\gamma \neq \lambda_{\max}(D)$ , it is thus clear that for  $\gamma = \lambda_{\max}(D)$  the optimal value of (A.3) must be given by  $J^* = \liminf_{\bar{\gamma} \downarrow \gamma} \bar{\gamma}^2 \langle (\bar{\gamma} I_d - D)^{-1}, \widehat{\Sigma} \rangle$ . This observation completes the proof.  $\square$

In order to derive search directions for the Frank-Wolfe algorithm developed in Section 6, we need to be able to solve constrained nonlinear SDPs of the form

$$(A.5) \quad \begin{aligned} &\sup_{\Sigma \succeq 0} \langle D, \Sigma \rangle \\ &\text{s. t.} \quad \operatorname{Tr} \left[ \Sigma + \widehat{\Sigma} - 2(\widehat{\Sigma}^{\frac{1}{2}} \Sigma \widehat{\Sigma}^{\frac{1}{2}})^{\frac{1}{2}} \right] \leq \rho^2 \end{aligned}$$

parameterized by  $D \in \mathbb{S}^d$ ,  $\widehat{\Sigma} \in \mathbb{S}_+^d$  and  $\rho \in \mathbb{R}_+$ . It is known that problem (A.5) admits a unique optimal solution that is available in quasi-closed form [54]. Below we review the construction of this optimal solution under more general conditions and uncover several previously unknown properties of this solution.

**Proposition A.4** (Quasi-closed form solution of (A.5)). *The following statements hold.*

- (i) *If  $D \in \mathbb{S}^d$ ,  $\widehat{\Sigma} \in \mathbb{S}_+^d$  and  $\rho \in \mathbb{R}_+$ , then problem (A.5) is solvable, and its maximum matches that of the univariate convex minimization problem*

$$(A.6) \quad \inf_{\substack{\gamma \geq 0 \\ \gamma > \lambda_{\max}(D)}} \gamma \left( \rho^2 + \langle \gamma I_d - D \rangle^{-1} - I_d, \widehat{\Sigma} \right).$$

- (ii) *If  $D \neq 0$ ,  $\widehat{\Sigma} \succ 0$  and  $\rho > 0$ , then problem (A.6) has a unique minimizer  $\gamma^* \in (\lambda_{\max}(D), \infty)$ , and problem (A.5) is solved by  $\Sigma^* = \gamma^{*2} (\gamma^* I_d - D)^{-1} \widehat{\Sigma} (\gamma^* I_d - D)^{-1}$ .*

- (iii) *If  $D \succeq 0$ ,  $D \neq 0$ ,  $\widehat{\Sigma} \succ 0$  and  $\rho > 0$ , then  $\gamma^*$  is the unique solution of the algebraic equation*

$$(A.7) \quad \rho^2 - \langle \widehat{\Sigma}, (I_d - \gamma^* (\gamma^* I_d - D)^{-1})^2 \rangle = 0$$

*in the interval  $(\lambda_{\max}(D), \infty)$ , the matrix  $\Sigma^*$  defined in assertion (iii) is the unique maximizer of (A.5), the Gelbrich distance constraint in (A.5) is binding at  $\Sigma^*$ , and  $\Sigma^* \succ \lambda_{\min}(\widehat{\Sigma}) I_d$ . Moreover, we have*

$$\underline{\gamma} = \lambda_1 \left( 1 + \sqrt{v_1^\top \widehat{\Sigma} v_1 / \rho} \right) \leq \gamma^* \leq \lambda_1 \left( 1 + \sqrt{\operatorname{Tr}[\widehat{\Sigma}] / \rho} \right) = \bar{\gamma},$$

*where  $\lambda_1$  is the largest eigenvalue of  $D$ , and  $v_1$  an eigenvector corresponding to  $\lambda_1$ .*

*Proof of Proposition A.4.* As for assertion (i), note that the Lagrangian dual of (A.5) can be represented as

$$(A.8) \quad \inf_{\gamma \geq 0} \sup_{\Sigma \succeq 0} \langle D, \Sigma \rangle - \gamma \operatorname{Tr} [\Sigma + \widehat{\Sigma} - 2(\widehat{\Sigma}^{\frac{1}{2}} \Sigma \widehat{\Sigma}^{\frac{1}{2}})^{\frac{1}{2}}] + \gamma \rho^2.$$

Strong duality as well as primal solvability follow from [7, Proposition 5.5.4], which applies because the primal problem (A.5) has a continuous objective function and—by virtue of Lemma A.6 below—a non-empty, compact and convex feasible set. The postulated reformulation (A.6) then follows immediately from replacing the supremum of the inner maximization problem in (A.8) with the analytical formula derived in Proposition A.3. We emphasize that for  $\gamma = \lambda_{\max}(D)$ , depending on the problem data, the inner supremum in (A.8) may evaluate to any non-negative real number or to  $+\infty$ . In order to avoid cumbersome case distinctions, we thus exclude the point  $\gamma = \lambda_{\max}(D)$  from the feasible set of (A.6) without affecting the problem's infimum.

As for assertion (ii), note that  $\Sigma = \widehat{\Sigma}$  represents a Slater point for the primal problem (A.5) because  $\rho > 0$ . Thus, the dual problem (A.8) is solvable by [7, Proposition 5.3.1]. To prove that (A.6) is also solvable, it remains to be shown that (A.8) does not attain its maximum at the boundary point  $\gamma = \lambda_{\max}(D)$ , which has been excluded from the feasible set of (A.6). This is the case, however, because of the assumption that  $\widehat{\Sigma} \succ 0$  and  $D \neq 0$ , which ensures that the objective function value of  $\gamma = \lambda_{\max}(D)$  in (A.8) amounts to  $+\infty$ . We may thus conclude that (A.6) admits a minimizer  $\gamma^* \in (\lambda_{\max}(D), \infty)$ . This minimizer is unique because the objective function of (A.6) is strictly convex when  $\widehat{\Sigma} \succ 0$ . Finally, the Karush-Kuhn-Tucker optimality conditions [7, Proposition 5.3.2] imply that any solution of the primal problem (A.5) also solves the inner maximization problem in (A.8) at  $\gamma = \gamma^*$ . The formula for  $\Sigma^*$  thus follows from Proposition A.3.

To prove assertion (iii), note that the assumptions  $D \succeq 0$  and  $D \neq 0$  imply that  $\gamma^* > \lambda_{\max}(D) > 0$ . Therefore, none of the constraints in (A.6) are binding at optimality. As the objective function of (A.6) is smooth and strictly convex,  $\gamma^*$  is thus uniquely determined by the first-order optimality condition (A.7), which forces the gradient of the objective function to vanish. The uniqueness of  $\Sigma^*$  follows from the uniqueness of  $\gamma^*$  and the uniqueness of the inner maximizer in (A.8); see Proposition A.3. Moreover, as  $\gamma^* > 0$ , the Gelbrich distance constraint in (A.6) is binding at  $\Sigma_\rho^*$  due to complementary slackness. Next, we have

$$\begin{aligned} \frac{1}{\lambda_{\min}(\Sigma^*)} &= \lambda_{\max}((\Sigma^*)^{-1}) = \lambda_{\max}\left(\left(\gamma^{*2}(\gamma^* I_n - D)^{-1} \widehat{\Sigma} (\gamma^* I_d - D)^{-1}\right)^{-1}\right) \\ &= \lambda_{\max}\left(\left(I_d - D/\gamma^*\right) \widehat{\Sigma}^{-1} \left(I_d - D/\gamma^*\right)\right) \\ &\leq \lambda_{\max}(I_d - D/\gamma^*)^2 \lambda_{\max}(\widehat{\Sigma}^{-1}) < \lambda_{\max}(\widehat{\Sigma}^{-1}) = \frac{1}{\lambda_{\min}(\widehat{\Sigma})}, \end{aligned}$$

where the strict inequality holds because  $\gamma^* > \lambda_{\max}(D) > 0$ . Thus, we conclude that  $\lambda_{\min}(\Sigma^*) > \lambda_{\min}(\widehat{\Sigma})$ .

To in order derive upper and lower bounds on  $\gamma^*$ , we let  $D = \sum_{i=1}^d \lambda_i v_i v_i^\top$  be the eigendecomposition of  $D$ , where  $\lambda_1, \dots, \lambda_d$  denote the eigenvalues of  $D$  indexed in descending order, while  $v_1, \dots, v_d$  represent the corresponding normalized eigenvectors. The left hand side of (A.7) can thus be re-expressed as

$$\rho^2 - \sum_{i=1}^d \left(\frac{\lambda_i}{\gamma - \lambda_i}\right)^2 v_i^\top \widehat{\Sigma} v_i$$

This expression is manifestly non-decreasing in  $\gamma \in (\lambda_1, \infty)$ . Moreover, the sum admits the simple bounds

$$\left(\frac{\lambda_1}{\gamma - \lambda_1}\right)^2 v_1^\top \widehat{\Sigma} v_1 \leq \sum_{i=1}^d \left(\frac{\lambda_i}{\gamma - \lambda_i}\right)^2 v_i^\top \widehat{\Sigma} v_i \leq \left(\frac{\lambda_1}{\gamma - \lambda_1}\right)^2 \operatorname{Tr}[\widehat{\Sigma}].$$

Equating these lower and upper bounds to  $\rho^2$  and solving the resulting equation for  $\gamma$  yields  $\underline{\gamma}$  and  $\bar{\gamma}$ , respectively. This observation concludes the proof.  $\square$

In Sections 3 and 4 we repeatedly encounter nonlinear SDPs of the form

$$(A.9) \quad \begin{aligned} & \sup_{\Sigma \succeq 0} \inf_{L \in \mathcal{C}} \langle L^\top L, \Sigma \rangle + f(L) \\ & \text{s. t. } \quad \text{Tr} [\Sigma + \widehat{\Sigma} - 2(\widehat{\Sigma}^{\frac{1}{2}} \Sigma \widehat{\Sigma}^{\frac{1}{2}})^{\frac{1}{2}}] \leq \rho^2 \end{aligned}$$

parameterized by  $\widehat{\Sigma} \in \mathbb{S}_+^d$  and  $\rho \in \mathbb{R}_+$ , where  $\mathcal{C} \subseteq \mathbb{R}^{\ell \times d}$  is a convex set and  $f : \mathcal{C} \rightarrow \mathbb{R}$  a convex continuous function. Problem (A.9) is reminiscent of (A.5) but accommodates a nonlinear convex objective function. We do not attempt to characterize the maximizers of (A.9) for arbitrary choices of  $\mathcal{C}$  and  $f$ , but we can prove that there is at least one well-behaved maximizer that is bounded away from 0.

**Lemma A.5** (Structural properties of the maximizers of (A.9)). *Assume that  $\widehat{\Sigma} \in \mathbb{S}_+^d$  and  $\rho \in \mathbb{R}_+$ . If  $\mathcal{C} \subseteq \mathbb{R}^{\ell \times d}$  is a non-empty convex set and  $f : \mathcal{C} \rightarrow \mathbb{R}$  is a convex continuous function, then the nonlinear SDP (A.9) admits a maximizer  $\Sigma^* \succeq \lambda_{\min}(\widehat{\Sigma})I_d$ .*

*Proof of Lemma A.5.* Note that if  $\rho = 0$  or  $\lambda_{\min}(\widehat{\Sigma}) = 0$ , then the claim holds trivially. Thus, we may henceforth assume without loss of generality that  $\rho > 0$  and  $\widehat{\Sigma} \succ 0$ . Denoting the feasible set of (A.9) by

$$\mathcal{S} = \left\{ \Sigma \in \mathbb{S}_+^d : \text{Tr} [\Sigma + \widehat{\Sigma} - 2(\widehat{\Sigma}^{\frac{1}{2}} \Sigma \widehat{\Sigma}^{\frac{1}{2}})^{\frac{1}{2}}] \leq \rho^2 \right\},$$

we then find

$$\begin{aligned} \sup_{\Sigma \in \mathcal{S}} \inf_{L \in \mathcal{C}} \langle L^\top L, \Sigma \rangle + f(L) &= \inf_{L \in \mathcal{C}} \sup_{\Sigma \in \mathcal{S}} \langle L^\top L, \Sigma \rangle + f(L) \\ &= \inf_{L \in \mathcal{C}} \sup_{\substack{\Sigma \in \mathcal{S} \\ \Sigma \succeq \lambda_{\min}(\widehat{\Sigma})I_d}} \langle L^\top L, \Sigma \rangle + f(L) = \sup_{\substack{\Sigma \in \mathcal{S} \\ \Sigma \succeq \lambda_{\min}(\widehat{\Sigma})I_d}} \inf_{L \in \mathcal{C}} \langle L^\top L, \Sigma \rangle + f(L), \end{aligned}$$

where the first and the third equality follow from Sion's minimax theorem [65], which applies because  $\langle L^\top L, \Sigma \rangle$  is convex and continuous in  $L$  for every fixed  $\Sigma \succeq 0$  and because  $\mathcal{S}$  is convex and compact by virtue of Lemma A.6. The second equality follows readily from Proposition A.4 (iii). The last maximization problem in the above expression has a solution  $\Sigma^* \succeq \lambda_{\min}(\widehat{\Sigma})I_d$  because its feasible set is compact and its objective function is upper semicontinuous. Clearly,  $\Sigma^*$  also solves (A.9), and thus the claim follows.  $\square$

The proofs of Proposition A.4 and Lemma A.5 rely on the following auxiliary lemma, which extends [64, Lemma A.6] to situations where  $\widehat{\Sigma}$  may be an arbitrary positive semidefinite covariance matrix.

**Lemma A.6** (Compactness of the feasible set). *For any  $\widehat{\Sigma} \in \mathbb{S}_+^d$  and  $\rho \in \mathbb{R}_+$ , the set*

$$\mathcal{S} = \left\{ \Sigma \in \mathbb{S}_+^d : \text{Tr} [\Sigma + \widehat{\Sigma} - 2(\widehat{\Sigma}^{\frac{1}{2}} \Sigma \widehat{\Sigma}^{\frac{1}{2}})^{\frac{1}{2}}] \leq \rho^2 \right\}$$

*is convex and compact. Moreover, for any  $\Sigma \in \mathcal{S}$  we have  $\text{Tr} [\Sigma] \leq (\rho + \text{Tr}[\widehat{\Sigma}]^{\frac{1}{2}})^2$ .*

*Proof of Lemma A.6.* The convexity of the feasible set  $\mathcal{S}$  follows from the convexity of the squared Gelbrich distance proved in Proposition 2.3. To prove that  $\mathcal{S}$  is compact, we recall from [51, Proposition 2] that

$$\begin{aligned} \text{Tr} [(\widehat{\Sigma}^{\frac{1}{2}} \Sigma \widehat{\Sigma}^{\frac{1}{2}})^{\frac{1}{2}}] &= \max_{C \in \mathbb{R}^{d \times d}} \left\{ \text{Tr} [C] : \begin{bmatrix} \Sigma & C \\ C^\top & \widehat{\Sigma} \end{bmatrix} \succeq 0 \right\} \\ &\leq \max_{C \in \mathbb{R}^{d \times d}} \left\{ \text{Tr} [C] : C_{ij}^2 \leq \Sigma_{ii} \widehat{\Sigma}_{jj} \forall i, j = 1, \dots, d \right\} = \sum_{i=1}^d \sqrt{\Sigma_{ii} \widehat{\Sigma}_{ii}} \leq \sqrt{\text{Tr} [\Sigma] \text{Tr} [\widehat{\Sigma}]}, \end{aligned}$$



where the first inequality holds because all second principal minors of a positive semidefinite matrix are non-negative, and the second inequality follows from the Cauchy-Schwarz inequality. Thus, any  $\Sigma \in \mathcal{S}$  satisfies

$$\rho^2 \geq \text{Tr} [\Sigma + \widehat{\Sigma} - 2(\widehat{\Sigma}^{\frac{1}{2}} \Sigma \widehat{\Sigma}^{\frac{1}{2}})^{\frac{1}{2}}] \geq \left( \sqrt{\text{Tr} [\Sigma]} - \sqrt{\text{Tr} [\widehat{\Sigma}]} \right)^2,$$

which implies that  $\text{Tr} [\Sigma] \leq (\rho + \text{Tr} [\widehat{\Sigma}]^{\frac{1}{2}})^2$ . This allows us to conclude that  $\mathcal{S}$  is bounded. Moreover,  $\mathcal{S}$  is closed due to the continuity of the matrix square root established in Lemma A.2.  $\square$

Finally, we derive the second-order Taylor expansion of the objective function

$$f(\Sigma_x, \Sigma_w) = \text{Tr} [\Sigma_x - \Sigma_x H^\top (H \Sigma_x H^\top + \Sigma_w)^{-1} H \Sigma_x]$$

of the nonlinear SDP (6.13), which is needed for the proof of Proposition 6.5 in the main text.

**Lemma A.7** (Gradient and Hessian of  $f$ ). *If  $(\Sigma_x, \Sigma_w) \in \mathbb{S}_{++}^n \times \mathbb{S}_{++}^m$  and  $G = H \Sigma_x H^\top + \Sigma_w \in \mathbb{S}_{++}^m$ , then*

$$\begin{aligned} D_x &= \nabla_{\Sigma_x} f(\Sigma_x, \Sigma_w) = (I_n - \Sigma_x H^\top G^{-1} H)^\top (I_n - \Sigma_x H^\top G^{-1} H) \\ D_w &= \nabla_{\Sigma_w} f(\Sigma_x, \Sigma_w) = G^{-1} H \Sigma_x^2 H^\top G^{-1} \\ \mathcal{H}_{xx} &= -\nabla_{xx}^2 f(\Sigma_x, \Sigma_w) = 2D_x \otimes H^\top G^{-1} H \\ \mathcal{H}_{xw} &= -\nabla_{xw}^2 f(\Sigma_x, \Sigma_w) = H^\top G^{-1} \otimes (H^\top D_w - \Sigma_x H^\top G^{-1}) + (H^\top D_w - \Sigma_x H^\top G^{-1}) \otimes H^\top G^{-1} \\ \mathcal{H}_{ww} &= -\nabla_{ww}^2 f(\Sigma_x, \Sigma_w) = 2D_w \otimes G^{-1}, \end{aligned}$$

where  $\nabla_x$  and  $\nabla_w$  stand for the nabla operators with respect to  $\text{vec}(\Sigma_x)$  and  $\text{vec}(\Sigma_w)$ , respectively.

*Proof of Lemma A.7.* We first derive the second-order Taylor expansion of  $G^{-1}$ . Specifically, if  $\Delta_x \in \mathbb{S}^n$  and  $\Delta_w \in \mathbb{S}^m$  represent symmetric perturbation directions of  $\Sigma_x$  and  $\Sigma_w$ , respectively, then we have

$$\begin{aligned} & (H [\Sigma_x + t \Delta_x] H^\top + [\Sigma_w + t \Delta_w])^{-1} \\ &= (G + t [H \Delta_x H^\top + \Delta_w])^{-1} \\ &= G^{-\frac{1}{2}} (I_m + t G^{-\frac{1}{2}} [H \Delta_x H^\top + \Delta_w] G^{-\frac{1}{2}})^{-1} G^{-\frac{1}{2}} \\ &= G^{-\frac{1}{2}} \left( I_m - t G^{-\frac{1}{2}} [H \Delta_x H^\top + \Delta_w] G^{-\frac{1}{2}} + t^2 \left( G^{-\frac{1}{2}} [H \Delta_x H^\top + \Delta_w] G^{-\frac{1}{2}} \right)^2 + \mathcal{O}(|t|^3) \right) G^{-\frac{1}{2}} \\ &= G^{-1} - t G^{-1} [H \Delta_x H^\top + \Delta_w] G^{-1} + t^2 G^{-1} [H \Delta_x H^\top + \Delta_w] G^{-1} [H \Delta_x H^\top + \Delta_w] G^{-1} + \mathcal{O}(|t|^3), \end{aligned}$$

where the third equality follows from a Neumann series expansion [6, Proposition 9.4.13]. Thanks to [6, Fact 7.4.9], the second-order Taylor expansion of  $f$  can thus be represented as

$$\begin{aligned} & f(\Sigma_x + t \Delta_x, \Sigma_w + t \Delta_w) \\ &= \text{Tr} [ [\Sigma_x + t \Delta_x] - [\Sigma_x + t \Delta_x] H^\top (H [\Sigma_x + t \Delta_x] H^\top + [\Sigma_w + t \Delta_w])^{-1} H [\Sigma_x + t \Delta_x] ] \\ &= f(\Sigma_x, \Sigma_w) + t \langle D_x, \Delta_x \rangle + t \langle D_w, \Delta_w \rangle - \frac{t^2}{2} \begin{pmatrix} \text{vec}(\Delta_x) \\ \text{vec}(\Delta_w) \end{pmatrix}^\top \begin{pmatrix} \mathcal{H}_{xx} & \mathcal{H}_{xw} \\ \mathcal{H}_{xw}^\top & \mathcal{H}_{ww} \end{pmatrix} \begin{pmatrix} \text{vec}(\Delta_x) \\ \text{vec}(\Delta_w) \end{pmatrix} + \mathcal{O}(|t|^3), \end{aligned}$$

where the matrices  $D_x$ ,  $D_w$ ,  $\mathcal{H}_{xx}$ ,  $\mathcal{H}_{xw}$  and  $\mathcal{H}_{ww}$  are defined as in the statement of the lemma.  $\square$

## REFERENCES

- [1] J.-P. AUBIN AND H. FRANKOWSKA, *Set-Valued Analysis*, Birkhäuser, 1990.
- [2] A. BECK, A. BEN-TAL, AND Y. C. ELДАР, *Robust mean-squared error estimation of multiple signals in linear systems affected by model and noise uncertainties*, *Mathematical Programming*, 107 (2006), pp. 155–187.

- [3] A. BECK AND Y. C. ELДАР, *Regularization in regression with bounded noise: A Chebyshev center approach*, SIAM Journal on Matrix Analysis and Applications, 29 (2007), pp. 606–625.
- [4] A. BECK, Y. C. ELДАР, AND A. BEN-TAL, *Mean-squared error estimation for linear systems with block circulant uncertainty*, SIAM Journal on Matrix Analysis and Applications, 29 (2007), pp. 712–730.
- [5] C. BERGE, *Topological Spaces: Including a Treatment of Multi-Valued Functions, Vector Spaces, and Convexity*, Courier Corporation, 1963.
- [6] D. S. BERNSTEIN, *Matrix Mathematics: Theory, Facts, and Formulas*, Princeton University Press, 2009.
- [7] D. BERTSEKAS, *Convex Optimization Theory*, Athena Scientific, 2009.
- [8] R. BHATIA, T. JAIN, AND Y. LIM, *Strong convexity of sandwiched entropies and related optimization problems*, Reviews in Mathematical Physics, (2018).
- [9] S. BOYD AND L. VANDENBERGHE, *Convex Optimization*, Cambridge University Press, 2004.
- [10] S. CAMBANIS, S. HUANG, AND G. SIMONS, *On the theory of elliptically contoured distributions*, Journal of Multivariate Analysis, 11 (1981), pp. 368–385.
- [11] S. G. CHANG, B. YU, AND M. VETTERLI, *Adaptive wavelet thresholding for image denoising and compression*, IEEE Transactions on Image Processing, 9 (2000), pp. 1532–1546.
- [12] C. CHATFIELD, *The Analysis of Time Series: An Introduction*, CRC Press, 2016.
- [13] K. L. CLARKSON, *Coresets, sparse greedy approximation, and the frank-wolfe algorithm*, ACM Transactions on Algorithms, 6 (2010), p. 63.
- [14] T. M. COVER AND J. A. THOMAS, *Elements of Information Theory*, Wiley-Interscience, 2006.
- [15] M. CUTURI, *Sinkhorn distances: Lightspeed computation of optimal transport*, in Advances in Neural Information Processing Systems 26, 2013, pp. 2292–2300.
- [16] V. F. DEMYANOV AND A. M. RUBINOV, *Approximate Methods in Optimization Problems*, American Elsevier Publishing, 1970.
- [17] S. N. DIGGAVI AND T. M. COVER, *The worst additive noise under a covariance constraint*, IEEE Transactions on Information Theory, 47 (2001), pp. 3072–3081.
- [18] D. C. DOWSON AND B. V. LANDAU, *The Fréchet distance between multivariate normal distributions*, Journal of Multivariate Analysis, 12 (1982), pp. 450–455.
- [19] J. C. DUNN, *Rates of convergence for conditional gradient algorithms near singular and nonsingular extremals*, SIAM Journal on Control and Optimization, 17 (1979), pp. 187–211.
- [20] ———, *Convergence rates for conditional gradient sequences generated by implicit step length rules*, SIAM Journal on Control and Optimization, 18 (1980), pp. 473–487.
- [21] J. C. DUNN AND S. HARSHBARGER, *Conditional gradient algorithms with open loop step size rules*, Journal of Mathematical Analysis and Applications, 62 (1978), pp. 432–444.
- [22] Y. C. ELДАР, *Robust competitive estimation with signal and noise covariance uncertainties*, IEEE Transactions on Information Theory, 52 (2006), pp. 4532–4547.
- [23] Y. C. ELДАР, A. BECK, AND M. TEBoulLE, *A minimax Chebyshev estimator for bounded error estimation*, IEEE Transactions on Signal Processing, 56 (2008), pp. 1388–1397.
- [24] Y. C. ELДАР, A. BEN-TAL, AND A. NEMIROVSKI, *Linear minimax regret estimation of deterministic parameters with bounded data uncertainties*, IEEE Transactions on Signal Processing, 52 (2004), pp. 2177–2188.
- [25] ———, *Robust mean-squared error estimation in the presence of model uncertainties*, IEEE Transactions on Signal Processing, 53 (2004), pp. 168–181.
- [26] Y. C. ELДАР AND N. MERHAV, *A competitive minimax approach to robust estimation of random parameters*, IEEE Transactions on Signal Processing, 52 (2004), pp. 1931–1946.

- [27] K.-T. FANG, S. KOTZ, AND K. W. NG, *Symmetric Multivariate and Related Distributions*, Chapman & Hall, 1990.
- [28] M. FRANK AND P. WOLFE, *An algorithm for quadratic programming*, Naval Research Logistics, 3 (1956), pp. 95–110.
- [29] R. M. FREUND AND P. GRIGAS, *New analysis and results for the Frank–Wolfe method*, Mathematical Programming, 155 (2016), pp. 199–230.
- [30] D. GARBER AND E. HAZAN, *Faster rates for the Frank–Wolfe method over strongly-convex sets*, in International Conference on International Conference on Machine Learning, 2015, pp. 541–549.
- [31] M. GELBRICH, *On a formula for the  $L^2$  Wasserstein metric between measures on Euclidean and Hilbert spaces*, Mathematische Nachrichten, 147 (1990), pp. 185–203.
- [32] C. R. GIVENS AND R. M. SHORTT, *A class of Wasserstein metrics for probability distributions*, The Michigan Mathematical Journal, 31 (1984), pp. 231–240.
- [33] F. GOLNARAGHI AND B. KUO, *Automatic Control Systems*, McGraw-Hill Education, 2017.
- [34] D. GUO, Y. WU, S. S. SHITZ, AND S. VERDU, *Estimation in Gaussian noise: Properties of the minimum mean-square error*, IEEE Transactions on Information Theory, 57 (2011), pp. 2371–2385.
- [35] J. HAMILTON, *Time Series Analysis*, Princeton University Press, 1994.
- [36] J. P. HESPANHA, *Linear Systems Theory*, Princeton University Press, 2009.
- [37] H. HULT AND F. LINDSKOG, *Multivariate extremes, aggregation and dependence in elliptical distributions*, Advances in Applied Probability, 34 (2002), pp. 587–608.
- [38] M. JAGGI, *Revisiting Frank–Wolfe: Projection-free sparse convex optimization*, in Proceedings of the 30th International Conference on Machine Learning, 2013, pp. 427–435.
- [39] E. JONDEAU, S. POON, AND M. ROCKINGER, *Financial Modeling under Non-Gaussian Distributions*, Springer, 2007.
- [40] M. JOURNÉE, Y. NESTEROV, P. RICHTÁRIK, AND R. SEPULCHRE, *Generalized power method for sparse principal component analysis*, Journal of Machine Learning Research, 11 (2010), pp. 517–553.
- [41] A. JUDITSKY AND A. NEMIROVSKI, *Lectures on statistical inferences via convex optimization*, 2018.
- [42] A. JUDITSKY, A. NEMIROVSKI, ET AL., *Near-optimality of linear recovery in Gaussian observation scheme under  $\|\cdot\|_2$ -loss*, The Annals of Statistics, 46 (2018), pp. 1603–1629.
- [43] S. M. KAY, *Fundamentals of Statistical Signal Processing: Estimation Theory*, Prentice Hall, 1993.
- [44] D. KELKER, *Distribution theory of spherical distributions and a location-scale parameter generalization*, Sankhyā: The Indian Journal of Statistics, Series A, (1970), pp. 419–430.
- [45] M. KNOTT AND C. S. SMITH, *On the optimal mapping of distributions*, Journal of Optimization Theory and Applications, 43 (1984), pp. 39–49.
- [46] E. S. LEVITIN AND B. T. POLYAK, *Constrained minimization methods*, USSR Computational Mathematics and Mathematical Physics, 6 (1966), pp. 1–50.
- [47] B. C. LEVY AND R. NIKOUKHAH, *Robust least-squares estimation with a relative entropy constraint*, IEEE Transactions on Information Theory, 50 (2004), pp. 89–104.
- [48] ———, *Robust state space filtering under incremental model perturbations subject to a relative entropy tolerance*, IEEE Transactions on Automatic Control, 58 (2012), pp. 682–695.
- [49] J. LÖFBERG, *YALMIP: A toolbox for modeling and optimization in MATLAB*, in 2004 IEEE International Conference on Robotics and Automation, 2004, pp. 284–289.
- [50] D. MACKAY, *Information Theory, Inference and Learning Algorithms*, Cambridge University Press, 2003.
- [51] L. MALAGÒ, L. MONTRUCCHIO, AND G. PISTONE, *Wasserstein Riemannian geometry of Gaussian densities*, Information Geometry, 1 (2018), pp. 137–179.

- [52] K. MURPHY, *Machine Learning: A Probabilistic Perspective*, MIT Press, 2012.
- [53] Y. NESTEROV, *Complexity bounds for primal-dual methods minimizing the model of objective function*, *Mathematical Programming*, 171 (2018), pp. 311–330.
- [54] V. A. NGUYEN, D. KUHN, AND P. MOHAJERIN ESFAHANI, *Distributionally robust inverse covariance estimation: The Wasserstein shrinkage estimator*, arXiv preprint arXiv:1805.07194, (2018).
- [55] K. OGATA, *Modern Control Engineering*, Pearson, 2009.
- [56] I. OLKIN AND F. PUKELSHEIM, *The distance between two random vectors with given dispersion matrices*, *Linear Algebra and its Applications*, 48 (1982), pp. 257–263.
- [57] A. V. OPPENHEIM AND G. C. VERGHESE, *Signals, Systems and Inference*, Pearson, 2015.
- [58] F. PEDREGOSA, A. ASKARI, G. NEGIAR, AND M. JAGGI, *Step-size adaptivity in projection-free optimization*, arXiv preprint arXiv:1806.05123, (2018).
- [59] G. PEYRÉ AND M. CUTURI, *Computational optimal transport*, *Foundations and Trends® in Machine Learning*, 11 (2019), pp. 355–607.
- [60] A. POSEKANY, K. FELSENSTEIN, AND P. SYKACEK, *Biological assessment of robust noise models in microarray data analysis*, *Bioinformatics*, 27 (2011), pp. 807–814.
- [61] U. E. RUTTIMANN, M. UNSER, R. R. RAWLINGS, D. RIO, N. F. RAMSEY, V. S. MATTAY, D. W. HOMMER, J. A. FRANK, AND D. R. WEINBERGER, *Statistical analysis of functional MRI data in the wavelet domain*, *IEEE Transactions on Medical Imaging*, 17 (1998), pp. 142–154.
- [62] B. A. SCHMITT, *Perturbation bounds for matrix square roots and Pythagorean sums*, *Linear Algebra and its Applications*, 174 (1992), pp. 215–227.
- [63] S. SHAFIEEZADEH-ABADEH, P. MOHAJERIN ESFAHANI, AND D. KUHN, *Regularization via mass transportation*, *Journal of Machine Learning Research*, 20 (2019), pp. 1–68.
- [64] S. SHAFIEEZADEH-ABADEH, V. NGUYEN, D. KUHN, AND P. MOHAJERIN ESFAHANI, *Wasserstein distributionally robust Kalman filtering*, in *Advances in Neural Information Processing Systems* 31, 2018, pp. 8483–8492.
- [65] M. SION, *On general minimax theorems*, *Pacific Journal of Mathematics*, 8 (1958), pp. 171–176.
- [66] J. SOLOMON, F. D. GOES, G. PEYRÉ, M. CUTURI, A. BUTSCHER, A. NGUYEN, T. DU, AND L. GUIBAS, *Convolutional Wasserstein distances: Efficient optimal transportation on geometric domains*, *ACM Transactions on Graphics*, 34 (2015), p. 66.
- [67] J. STOCK AND M. WATSON, *Introduction to Econometrics*, Prentice Hall, 2015.
- [68] B. TAŞKESEN, S. SHAFIEEZADEH-ABADEH, AND D. KUHN, *A unified approach for solving optimal transport problem*, Working paper, (2019).
- [69] J. WOOLDRIDGE, *Econometric Analysis of Cross Section and Panel Data*, MIT Press, 2010.
- [70] J. ZHOU, *On the existence of equilibrium for abstract economies*, *Journal of Mathematical Analysis and Applications*, 193 (1995), pp. 839–858.
- [71] M. ZORZI, *Robust Kalman filtering under model perturbations*, *IEEE Transactions on Automatic Control*, 62 (2016), pp. 2902–2907.
- [72] ———, *On the robustness of the Bayes and Wiener estimators under model uncertainty*, *Automatica*, 83 (2017), pp. 133–140.