

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Distributionally Robust Formulation and Model Selection for the Graphical Lasso

Permalink

<https://escholarship.org/uc/item/2d41n78t>

Author

Cisneros, Pedro

Publication Date

2019

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Distributionally Robust Formulation and Model Selection for the Graphical Lasso

A thesis submitted in partial satisfaction
of the requirements for the degree

Master of Arts
in
Statistics

by

Pedro A. Cisneros

Committee in charge:

Professor Alexander Petersen, Co-Chair
Professor Sang-Yun Oh, Co-Chair
Professor Wendy Meiring

March 2020

The Thesis of Pedro A. Cisneros is approved.

Professor Wendy Meiring

Professor Sang-Yun Oh, Committee Co-Chair

Professor Alexander Petersen, Committee Co-Chair

December 2019

Distributionally Robust Formulation and Model Selection for the Graphical Lasso

Copyright © 2019

by

Pedro A. Cisneros

Para mi familia (Pedro, Lourdes, Sergio y abuelos); y para una
amiga muy especial de Santa Paula.

Acknowledgements

I am deeply grateful to Prof. Alexander Petersen and Prof. Sang-Yun Oh for all the help and good advising. Without them, this accomplishment could not be possible. I am also grateful to Prof. Francesco Bullo for supporting me in pursuing this masters degree, and all the professors from the different classes I took in the PSTAT department. Finally, I am grateful to God for all the blessings throughout my studies.

Curriculum Vitæ

Pedro A. Cisneros

Education

- 2021 Ph.D. in Electrical and Computer Engineering (Expected), University of California, Santa Barbara.
- 2020 M.A. in Mathematical Statistics (Expected), University of California, Santa Barbara.
- 2018 M.S. in Electrical and Computer Engineering, University of California, Santa Barbara.

Publications

- P. Cisneros-Velarde, K. S. Chan, and F. Bullo, "Polarization and Fluctuations in Signed Social Networks," Submitted, 2019.
- P. Cisneros-Velarde, N. E. Friedkin, A. V. Proskurnikov, F. Bullo, "Structural Balance via Gradient Flows over Signed Graphs," Submitted, 2019.
- P. Cisneros-Velarde, and F. Bullo, "Signed Network Formation Games and Clustering Balance," Dynamic Games and Applications, 2020.
- W. Mei, P. Cisneros-Velarde, G. Chen, N. E. Friedkin and F. Bullo, "Dynamic social balance and convergent appraisals via homophily and influence mechanisms," Automatica, Volume 110, 2019.
- P. Cisneros-Velarde, D. Oliveira, and K. Chan, "Spread and Control of Misinformation with Heterogeneous Agents," International Workshop on Complex Networks, Springer International Publishing, 2019.
- P. Cisneros-Velarde et al., "Automatic pneumonia detection based on ultrasound video analysis," 2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), 2016.
- D. Arroyo, C. Lucho, P. Cisneros and F. Cuellar, "Prime: Towards the Design of a Small Interactive Office Robot," 6th International Conference on Social Robotics (ICSR 2014), Springer International Publishing, 2014.
- P. Cisneros and P. Rodriguez, "Practical hand tracking solution by alternating the use of a priori information," 2014 IEEE 5th Latin American Symposium on Circuits and Systems, Santiago, 2014.

Abstract

Distributionally Robust Formulation and Model Selection for the Graphical Lasso

by

Pedro A. Cisneros

Building on a recent framework for distributionally robust optimization, we consider estimation of the inverse covariance matrix for multivariate data. We provide a novel notion of a Wasserstein ambiguity set specifically tailored to this estimation problem, leading to a tractable class of regularized estimators. Special cases include penalized likelihood estimators for Gaussian data, specifically the graphical lasso estimator. As a consequence of this formulation, the radius of the Wasserstein ambiguity set is directly related to the regularization parameter in the estimation problem. Using this relationship, the level of robustness of the estimation procedure can be shown to correspond to the level of confidence with which the ambiguity set contains a distribution with the population covariance. Furthermore, a unique feature of our formulation is that the radius can be expressed in closed-form as a function of the ordinary sample covariance matrix. Taking advantage of this finding, we develop a simple algorithm to determine a regularization parameter for graphical lasso, using only the bootstrapped sample covariance matrices, meaning that computationally expensive repeated evaluation of the graphical lasso algorithm is not necessary. Alternatively, the distributionally robust formulation can also quantify the robustness of the corresponding estimator if one uses an off-the-shelf method such as cross-validation. Finally, we numerically study the obtained regularization criterion and analyze the robustness of other automated tuning procedures used in practice.

Contents

Curriculum Vitae	vi
Abstract	vii
1 Introduction	1
2 Distributionally Robust Optimization, Graphical lasso and Model Selection	6
2.1 A Distributionally Robust Formulation of the Graphical lasso	6
2.2 Selection of the regularization parameter	13
3 Numerical Evaluation	22
A Appendix Title	27
A.1 Using a different ambiguity set	27
A.2 Applicability of the dual representations of the RWP function and the DRO formulation	29
A.3 Additional figures, tables and additional analysis for the numerical results	31
Bibliography	36

Chapter 1

Introduction

In statistics and machine learning, the covariance matrix Σ of a random vector $X \in \mathbb{R}^d$ is a fundamental quantity for characterizing marginal pairwise dependencies between variables. Furthermore, the inverse covariance matrix $\Omega = \Sigma^{-1}$ provides information about the *conditional* linear dependency structure between the variables. For example, in the case that X is Gaussian, $\Omega_{jk} = 0$ if and only if the j th and k th variables of X are conditionally independent given the rest. Such relationships are of interest in many applications such as environmental science, biology, and neuroscience [1, 2, 3], and have given rise to various statistical and machine learning methods for inverse covariance estimation.

Given an independent sample $X_i \sim X$, $i = 1, \dots, n$, the sample covariance $A_n = \frac{1}{n} \sum_{i=1}^n X_i X_i^\top$, which is the maximum likelihood estimator if X is Gaussian, can be a poor estimator of Σ unless d/n is very small. Driven by a so-called high dimensional setting where $n \ll d$, where A_n is not invertible, regularized estimation of the precision matrix has gained significant interest [4, 5, 6, 7, 8, 9]. Such regularization procedures are useful even when A_n is a stable estimate (i.e., positive definite with small condition number), since the inverse covariance estimate A_n^{-1} is dense and will not reflect the

sparsity of corresponding nonzero elements in Ω .

Distributionally Robust Optimization: Let \mathbb{S}_d be the set of $d \times d$ symmetric matrices, and $\mathbb{S}_d^{++} \subset \mathbb{S}_d$ be the subset of positive definite symmetric matrices. Given a loss function $l(X; K)$ for $K \in \mathbb{S}_d^{++}$ and $X \in \mathbb{R}^d$, a classical approach would be to estimate Ω by minimizing the empirical loss $\sum_{i=1}^n l(X_i; K)$ over $K \in \mathbb{S}_d^{++}$, perhaps including a regularization or penalty term. The error in this estimate arises from the discrepancy between the true data-generating distribution and the observed training samples, and can be assessed by various tools such as concentration bounds or rates of convergence. In contrast, *distributionally robust optimization* (DRO) is a technique that explicitly incorporates uncertainty about the distribution of the X_i into the estimation procedure. For an introduction on the general topic of DRO, we refer to the works [10, 11, 12] and the references therein. In the context of inverse covariance estimation, a distributionally robust estimate of Ω is obtained by solving

$$\inf_{K \in \mathbb{S}_d^{++}} \sup_{P \in \mathcal{S}} E_P[l(K; X)], \quad (1.1)$$

where \mathcal{S} , known as an *ambiguity set*, is a collection of probability measures on \mathbb{R}^d . As pointed out in [11], a natural choice for \mathcal{S} is the neighborhood $\{P \mid D(P, \mu) \leq \delta\}$, with μ being a chosen baseline model, δ being some tolerance level which defines the uncertainty size of the ambiguity set, and D being some discrepancy metric between two probability measures. In a practical setting, we have access to some samples (or data points) from the unknown distribution, and thus, a good candidate for the baseline model is the empirical measure.

Very recent work by [13], also analyzed by [14], used the DRO framework to construct a new regularized (dense) inverse covariance estimator. Working under the assumption that X is Gaussian, the authors construct an ambiguity set \mathcal{S} of Gaussian distributions

that, up to a certain tolerance level, are consistent with the observed data. Recent work on DRO in other machine learning problems has revealed explicit connections to well-known regularized estimators, specifically regularized logistic regression [10] and the square-root lasso for linear models [15]; however, such a connection to regularized *sparse* inverse covariance estimators that are used in practice has yet to be made.

The Graphical Lasso: One of the most common methods to recover the sparsity pattern in Ω is to add an l_1 -regularization term to the Gaussian likelihood function, motivated by the consideration of the *Gaussian Graphical Model* (GGM). A sparse estimate of $\Omega = \Sigma^{-1}$ is produced by minimizing

$$\mathcal{L}_\lambda(K) = \frac{1}{n} \sum_{i=1}^n X_i^\top K X_i - \log |K| + \lambda \sum_{i=1}^d \sum_{j=1}^d |k_{ij}|, \quad (1.2)$$

where k_{ij} is the (i, j) entry of K and $\lambda > 0$ is a user-specified regularization parameter [16, 6, 9]. Although several algorithms exist to solve this objective function [6, 17, 18], the minimizer of (1.2) is often referred to as *graphical lasso* estimator [6]. The first two terms of (1.2) are related to Stein's loss [19] when evaluated at the empirical measure, and also correspond to the negative log-likelihood up to an additive constant if X is Gaussian. The performance of the graphical lasso estimator in high-dimensional settings has been investigated [20, 21], as well as modifications and extensions that implement some notion of *robustness*, i.e., for making it robust to outliers or relaxing the normality assumptions in the data [7, 22, 23, 24, 25].

Besides its theoretical relevance, the graphical lasso and its extensions also enjoy many practical advantages. For example, it has been used as a network inference tool. In these applications, the precision matrix can indicate which nodes in a network are conditionally independent given information from remaining nodes, thus giving an indication of network functionality. This has been important in neuroscience applications when studying

the inference of brain connectivity [26, 27, 2]. Applications in gene regulatory networks and metabolomics have also been reported [28, 29, 3].

The performance of the graphical lasso estimator hinges critically on the choice of λ . While there have been studies on how to properly tune λ to obtain a consistent estimator or to establish correct detection of nonzero elements in the precision matrix [20, 16, 30], in practice, this selection is often made through automated methods like cross-validation.

Contributions: In this paper, we propose a distributionally robust reformulation of the graphical lasso estimator in (1.2). Following [10, 15, 31], we utilize the Wasserstein metric to quantify distributional uncertainty for the construction of the ambiguity set. The following points summarize our main contributions.

- We formulate a class of DRO problems for inverse covariance estimation, leading to a tractable class of ℓ_p -norm regularized estimators. As the graphical lasso estimator (1.2) is a special case, this provides us with a new interpretation of this popular technique. This DRO formulation is made possible by a novel type of ambiguity set, now defined as a collection of measures on matrices. This nontrivial adaptation is necessary due to the fact that a direct generalization of other DRO approaches using vector-valued data (e.g., [12]) does not result in a closed-form regularization problem, and thus does not provide the desired connection to the graphical lasso.
- We use this formulation to suggest a criterion for the selection of the regularization parameter in the estimation problem in the classical regime $n > d$. This criterion follows the *Robust Wasserstein Profile* (RWP) inference recently introduced by [15], which makes no assumption on the normality of the data, and which we tailor to our specific problem. The proposed criterion expresses the regularization parameter as an explicit function of the sample covariance A_n , unlike other instances where RWP has been implemented which rely on stochastic dominance arguments.

- We formulate a novel *robust selection* (RobSel) algorithm for regularization parameter choice. Focusing on the graphical lasso, we provide numerical results that compare the performance of cross-validation and our proposed algorithm for the selection of the regularization term.

The paper is organized as follows. In Section 2 we describe our main theoretical result: the distributionally robust formulation of the regularized inverse covariance (log-likelihood) estimation, from which graphical lasso is a particular instance. In Section 3 we propose a criterion for choosing the regularization parameter inspired by this formulation and outline the bootstrap-based RobSel algorithm for its computation. In Section 4 we present some numerical results comparing the proposed criterion of Section 3 with cross-validation. Finally, we state some concluding remarks and future research directions in Section 5. All proofs of theoretical results can be found in the supplementary material.

Chapter 2

Distributionally Robust Optimization, Graphical lasso and Model Selection

2.1 A Distributionally Robust Formulation of the Graphical lasso

First, we provide preliminary details on notation. Given a matrix $A \in \mathbb{R}^{d \times d}$, a_{jk} denotes its (j, k) entry and $\mathbf{vec}(A) \in \mathbb{R}^{d^2}$ denotes its vectorized form, which we assume to be in a row major fashion. For matrices denoted by Greek letters, its entries are simply denoted by appropriate subscripts, i.e. Σ_{jk} . The operator $|\cdot|$, when applied to a matrix, denotes its determinant; when applied to a scalar or a vector, it denotes the absolute value or entry-wise absolute value, respectively. The ℓ_p -norm of a vector is denoted by $\|\cdot\|_p$. We use the symbol \Rightarrow to denote convergence in distribution.

Recall that $X \in \mathbb{R}^d$ is a zero-mean random vector with covariance matrix $\Sigma \in \mathbb{S}_d^{++}$.

Let \mathbb{Q}_0 be the probability law for X and $\Omega = \Sigma^{-1}$ be the precision matrix. Define the *graphical loss function* as

$$\begin{aligned} l(X; K) &= X^\top K X - \log |K| \\ &= \text{trace}(K X X^\top) - \log |K|. \end{aligned} \tag{2.1}$$

Then $E_{\mathbb{Q}_0}[l(K; X)] = \text{trace}(K \Sigma) - \log |K|$ is a convex function of K over the convex cone \mathbb{S}_d^{++} . Using the first-order optimality criterion, we observe that $K = \Omega$ sets the gradient $\frac{\partial}{\partial K} E_{\mathbb{Q}_0}[l(X, K)] = \Sigma - K^{-1}$ equal to the zero matrix (see [32, Appendix A] for details on this differentiation). Hence,

$$\arg \min_{K \in \mathbb{S}_d^{++}} E_{\mathbb{Q}_0}[l(X; K)] = \Omega.$$

so that (2.1) is a consistent loss function.

Now, if we consider an iid random sample $X_1, \dots, X_n \sim X$, $n > d$, with empirical measure \mathbb{Q}_n , then

$$\begin{aligned} \arg \min_{K \in \mathbb{S}_d^{++}} E_{\mathbb{Q}_n}[l(X; K)] &= \arg \min_{K \in \mathbb{S}_d^{++}} \frac{1}{n} \sum_{i=1}^n l(X_i; K) \\ &= A_n^{-1} \end{aligned}$$

with $A_n = \frac{1}{n} \sum_{i=1}^n X_i X_i^\top$. Thus, as described in the Introduction, a natural approach would be to implement the DRO procedure outlined in [31] by building an ambiguity set based on perturbations of \mathbb{Q}_n , leading to the DRO estimate given by (1.1). However, this approach does not convert (1.1) into a regularized estimation problem as desired, since the inner supremum cannot be explicitly given in closed-form. For more details, see section A in the supplementary material.

As an alternative, let \mathbb{P}_0 represent the measure of the random matrix $W = XX^\top$ on \mathbb{S}_d induced by \mathbb{Q}_0 and, similarly let \mathbb{P}_n be empirical measure of the sample $W_i = X_i X_i^\top$, $i = 1, \dots, n$. Redefining the graphical loss function $l : \mathbb{S}_d \times \mathbb{S}_d^{++}$ as

$$l(W; K) = \text{trace}(KW) - \log |K|, \quad (2.2)$$

then

$$\begin{aligned} \Omega &= \arg \min_{K \in \mathbb{S}_d^{++}} E_{\mathbb{P}_0}[l(W; K)] \\ A_n^{-1} &= \arg \min_{K \in \mathbb{S}_d^{++}} E_{\mathbb{P}_n}[l(W; K)]. \end{aligned}$$

This observation leads to a tractable DRO formulation by constructing ambiguity sets built around the empirical measure \mathbb{P}_n . The DRO formulation for inverse covariance estimation becomes

$$\min_{K \in \mathbb{S}_d^{++}} \sup_{P: \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta} E_P[l(W; K)]. \quad (2.3)$$

The ambiguity set in this formulation is specified by the collection of measures

$$\{P \mid \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta\},$$

which we now describe. Given two probability distributions P_1 and P_2 on \mathbb{S}_d and some transportation cost function $c : \mathbb{S}_d \times \mathbb{S}_d \rightarrow [0, \infty)$ (which we will specify below), we define the *optimal transport cost* between P_1 and P_2 as

$$\mathcal{D}_c(P_1, P_2) = \inf\{E_\pi[c(U, V)] \mid \pi \in \mathcal{P}(\mathbb{S}_d \times \mathbb{S}_d), \pi_U = P_1, \pi_V = P_2\} \quad (2.4)$$

where $\mathcal{P}(\mathbb{S}_d \times \mathbb{S}_d)$ is the set of joint probability distributions π of (U, V) supported on

$\mathbb{S}_d \times \mathbb{S}_d$, and π_U and π_V denote the marginals of U and V under π , respectively. In this paper, we are interested in cost functions

$$c(U, V) = \|\mathbf{vec}(U) - \mathbf{vec}(V)\|_q^\rho, \quad (2.5)$$

with $U, V \in \mathbb{S}_d$, $\rho \geq 1$, $q \in [1, \infty]$. As pointed out by [15], the resulting optimal transport cost $\mathcal{D}_c^{1/\rho}$ is the Wasserstein distance of order ρ . Our first theoretical result demonstrates that the optimization in (2.3) corresponds to a class of regularized estimators under the graphical loss function (2.2).

Theorem 2.1.1 (DRO formulation of regularized inverse covariance estimation)

Consider the cost function in (2.5) for a fixed $\rho \geq 1$. Then,

$$\min_{K \in \mathbb{S}_d^{++}} \sup_{P: \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta} E_P[l(W; K)] = \min_{K \in \mathbb{S}_d^{++}} \left\{ \text{trace}(K A_n) - \log |K| + \delta^{1/\rho} \|\mathbf{vec}(K)\|_p \right\}, \quad (2.6)$$

where $\frac{1}{p} + \frac{1}{q} = 1$.

Proof: Consider $K \in \mathbb{S}_d^{++}$. Observe that the function $l(\cdot; K) : \mathbb{S}_d \rightarrow \mathbb{R}$ is Borel measurable since it is a continuous function. Then, we use the duality result for the DRO formulation from Proposition A.2.2 from the appendix of this paper and obtain

$$\sup_{P: \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta} E_P[l(W; K)] = \inf_{\gamma \geq 0} \left\{ \gamma \delta + \frac{1}{n} \sum_{i=1}^n \left(\sup_{W \in \mathbb{S}_d} \{l(W; K) - \gamma c(W, W_i)\} \right) \right\}. \quad (2.7)$$

Let $\Delta := W - W_i$. Then,

$$\begin{aligned}
 & \sup_{W \in \mathbb{S}_d} \{l(W; K) - \gamma c(W, W_i)\} \\
 &= \sup_{W \in \mathbb{S}_d} \{\text{trace}(KW) - \log |K| - \gamma \|\mathbf{vec}(W) - \mathbf{vec}(W_i)\|_q^\rho\} \\
 &= \sup_{\Delta \in \mathbb{S}_d} \{\text{trace}(K(\Delta + W_i)) - \gamma \|\mathbf{vec}(\Delta)\|_q^\rho\} - \log |K| \tag{2.8} \\
 &= \sup_{\Delta \in \mathbb{S}_d} \{\text{trace}(K\Delta) - \gamma \|\mathbf{vec}(\Delta)\|_q^\rho\} + \text{trace}(KW_i) - \log |K| \\
 &= \sup_{\Delta \in \mathcal{M}(K)} \{\|\mathbf{vec}(\Delta)\|_q \|\mathbf{vec}(K)\|_p - \gamma \|\mathbf{vec}(\Delta)\|_q^\rho\} + \text{trace}(KW_i) - \log |K|
 \end{aligned}$$

with $\mathcal{M}(K) = \{\Delta \in \mathbb{S}_d \mid \text{trace}(K\Delta) > 0, |\Delta_{ij}|^q = \theta |k_{ij}|^p \text{ for some } \theta > 0\}$ so that the fourth line follows from selecting a $\Delta \in \mathbb{S}_d$ (since $K \in \mathbb{S}_d^{++}$) such that Holder's inequality holds tightly (with $\frac{1}{p} + \frac{1}{q} = 1$). In fact, Holder's inequality holds tightly if and only if $\Delta \in \mathcal{M}(K)$ [33, Chapter 9], even for the limiting case $q = \infty, p = 1$. Observe that there exist multiple $\Delta \in \mathbb{S}_d$ that can satisfy Holder's inequality tightly. As a consequence, we are still free to choose the magnitude of the q -norm of such $\mathbf{vec}(\Delta)$ (and this is what we will use next).

Now, the argument inside the supremum in the last line of (A.5) is a polynomial function on $\|\mathbf{vec}(\Delta)\|_q$. We have to analyze two cases.

Case 1: $\rho = 1$. In this case we observe that, by setting

$$\epsilon(\gamma, K) = \sup_{\Delta \in \mathcal{M}(K)} \{\|\mathbf{vec}(\Delta)\|_q (\|\mathbf{vec}(K)\|_p - \gamma)\}:$$

- if $\gamma \geq \|\mathbf{vec}(K)\|_p$, then $\epsilon(\gamma, K) = 0$ (in particular, if $\gamma = \|\mathbf{vec}(K)\|_p$, the optimizer is $\Delta = \mathbb{0}_{d \times d}$);
- if $\gamma < \|\mathbf{vec}(K)\|_p$, then $\epsilon(\gamma, K) = \infty$;

so that, recalling (2.7), due to the outside infimum to be taken over $\gamma \geq 0$; and so we

must have that

$$\sup_{P: \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta} E_P[l(W; K)] = \inf_{\gamma \geq \|\mathbf{vec}(K)\|_p} \left\{ \gamma \delta + \frac{1}{n} \sum_{i=1}^n (\text{trace}(KW_i) - \log |K|) \right\}$$

from which we immediately obtain (2.6).

Case 2: $\rho > 1$. By differentiation and basic calculus (e.g., using the first and second derivative test) we obtain that the maximizer

$$\Delta^* = \arg \sup_{\Delta \in \mathcal{M}(K)} \{ \|\mathbf{vec}(\Delta)\|_q \|\mathbf{vec}(K)\|_p - \gamma \|\mathbf{vec}(\Delta)\|_q^\rho \}$$

is such that $\|\mathbf{vec}(\Delta^*)\|_q = \left(\frac{\|\mathbf{vec}(K)\|_p}{\gamma \rho} \right)^{\frac{1}{\rho-1}}$. Then,

$$\begin{aligned} \sup_{W \in \mathbb{S}_d} \{l(W; K) - \gamma c(W, W_i)\} &= \frac{\|\mathbf{vec}(K)\|_p^{\frac{\rho}{\rho-1}}}{(\gamma \rho)^{\frac{1}{\rho-1}}} - \gamma \left(\frac{\|\mathbf{vec}(K)\|_p}{\gamma \rho} \right)^{\frac{\rho}{\rho-1}} + \text{trace}(KW_i) \\ &\quad - \log |K| \\ &= \|\mathbf{vec}(K)\|_p^{\frac{\rho}{\rho-1}} \frac{\rho-1}{\rho^{\frac{\rho}{\rho-1}} \gamma^{\frac{1}{\rho-1}}} + \text{trace}(KW_i) - \log |K|. \end{aligned} \tag{2.9}$$

Replacing this back in (2.7),

$$\begin{aligned}
& \sup_{P: \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta} E_P[l(W; K)] \\
&= \inf_{\gamma \geq 0} \left\{ \gamma \delta + \frac{1}{n} \sum_{i=1}^n \left(\|\mathbf{vec}(K)\|_p^{\frac{\rho}{\rho-1}} \frac{\rho-1}{\rho^{\frac{\rho}{\rho-1}} \gamma^{\frac{1}{\rho-1}}} + \text{trace}(KW_i) - \log |K| \right) \right\} \\
&= \inf_{\gamma \geq 0} \left\{ \gamma \delta + \|\mathbf{vec}(K)\|_p^{\frac{\rho}{\rho-1}} \frac{\rho-1}{\rho^{\frac{\rho}{\rho-1}} \gamma^{\frac{1}{\rho-1}}} + \text{trace} \left(K \frac{1}{n} \sum_{i=1}^n W_i \right) - \log |K| \right\} \\
&= \inf_{\gamma \geq 0} \left\{ \gamma \delta + \|\mathbf{vec}(K)\|_p^{\frac{\rho}{\rho-1}} \frac{\rho-1}{\rho^{\frac{\rho}{\rho-1}} \gamma^{\frac{1}{\rho-1}}} \right\} + \text{trace}(KA_n) - \log |K|.
\end{aligned} \tag{2.10}$$

Now, we observe that the argument inside the infimum in the last line of (2.10) is a function that grows to infinity when $\gamma \rightarrow 0$ or $\gamma \rightarrow \infty$, so that the minimum is attained for some optimal γ . By using the first and second derivative tests, we obtain that the minimizer is $\gamma^* = \frac{\|\mathbf{vec}(K)\|_p}{\rho \delta^{\frac{\rho-1}{\rho}}}$. Then, replacing this back in (2.10) and then this in (2.7), we finally obtain (2.6) after some algebraic simplification. \blacksquare

Theorem 2.1.1 is a remarkable theoretical result that provides a mapping between the regularization parameter and the uncertainty size δ of the ambiguity set in the DRO formulation. Then, the regularization problem reduces to determining a good criterion for choosing δ , which we explore in Section 2.2. Moreover, we obtain the *graphical lasso* formulation (1.2) by setting $q = \infty$ in (2.5). From (2.6), a smaller ambiguity set implies less robustness being introduced in the estimation problem by reducing the importance of the regularization term. Conversely, a larger regularization term increases the number of nuisance distributions inside the ambiguity set, and thus the robustness.

Remark 2.1.2 *The ambiguity set used in (2.6) makes no assumptions on the normality of the distribution of the samples $\{X_1, \dots, X_n\}$. Then, (2.6) tells us that adding a penalization to the precision matrix gives a robustness in terms of the distributions that the*

samples may have, which do not necessarily have to be Gaussian for this formulation to hold. Furthermore, it holds independent of the relationship between n and d .

2.2 Selection of the regularization parameter

This section follows closely the line of thought recently introduced by [15] in the analysis of regularized estimators under the DRO formulation. Specifically, we will demonstrate that the ambiguity set $\{P : \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta\}$ represents a confidence region for $\Omega = \Sigma^{-1}$, and use the techniques of [15] to explicitly connect the ambiguity size δ with a confidence level. As previously stated, $l(W; K)$ is a differentiable function on $K \in \mathbb{S}_d^{++}$ with $\frac{\partial}{\partial K} l(W; K) = W - K^{-1}$, so that

$$E_{\mathbb{P}_0} \left[\frac{\partial}{\partial K} l(W; K) \Big|_{K=\Omega} \right] = \mathbb{0}_{d \times d}. \quad (2.11)$$

Hence, even though the loss function $l(W; K)$ has been inspired from the log-likelihood estimation of the covariance matrix Σ for samples of Gaussian random vectors, equation (2.11) is transparent to any underlying distribution of the data. For any $K \in \mathbb{S}_d^{++}$, define the set

$$\mathcal{O}(K) := \left\{ P \in \mathcal{P}(\mathbb{S}_d) \mid E_P \left[\frac{\partial}{\partial K'} l(W; K') \Big|_{K'=K} \right] = \mathbb{0}_{d \times d} \right\}, \quad (2.12)$$

corresponding to all probability measures with covariance K^{-1} , i.e. for which K is an optimal loss minimization parameter; here, $\mathcal{P}(\mathbb{S}_d)$ denotes the set of all probability distributions supported on \mathbb{S}_d . Thus, $\mathcal{O}(\Omega)$ contains all probability measures with covariance matrix agreeing with that of X .

Implicitly, the Wasserstein ambiguity set $\{P : \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta\}$ is linked to the

collection of covariance matrices

$$\begin{aligned} \mathcal{C}_n(\delta) &:= \{K \in \mathbb{S}_d^{++} \mid \text{there exists } P \in \mathcal{O}(K) \cap \{P \mid \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta\}\} \\ &= \bigcup_{P: \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta} \arg \min_{K \in \mathbb{S}_d^{++}} E_P[l(W; K)]. \end{aligned} \quad (2.13)$$

We refer to $\mathcal{C}_n(\delta)$ as the set of *plausible* selections for Ω .

Lemma 2.2.1 (Interchangeability in the DRO formulation) *Consider the setting of Theorem 2.1.1. Then, for $n > d$, the following holds with probability one:*

$$\inf_{K \in \mathbb{S}_d^{++}} \sup_{P: \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta} E_P[l(W; K)] = \sup_{P: \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta} \inf_{K \in \mathbb{S}_d^{++}} E_P[l(W; K)]. \quad (2.14)$$

Proof: Consider $K \in \mathbb{S}_d^{++}$ and define $g(K) = \sup_{P: \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta} E_P[l(W; K)]$ for a fixed δ . We prove (2.14), by a direct application of Proposition 8 of [15, Appendix C], observing that we satisfy the three conditions for its application:

1. g is convex on \mathbb{S}_d^{++} and finite,
2. there exists $b \in \mathbb{R}$ such that the sublevel set $\kappa_b = \{K \mid g(K) \leq b\}$ is compact and non-empty,
3. $E_P[l(W; K)]$ is lower semi-continuous and convex as a function of K throughout κ_b for any $P \in \{P \mid \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta\}$.

First, we observe that

$$E_{\mathbb{P}_0}[l(W, K)] = E_{\mathbb{P}_0}[\text{trace}(KW) - \log |K|] \leq \text{trace}(KE_{\mathbb{P}_0}[W]) < \infty,$$

since $E_{\mathbb{Q}_0}(\|X\|_2^2) < \infty$. Then, using Theorem 2.1.1, the function

$$g(K) = \sup_{P: \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta} E_P[l(W; K)] = \text{trace}(KA_n) - \log |K| + \delta^{1/\rho} \|\mathbf{vec}(K)\|_p$$

is finite. Moreover, we also claim it is convex for all $K \in \mathbb{S}_d^{++}$. This follows from the fact that $\text{trace}(KA_n) - \log |K|$ and $\|\mathbf{vec}(K)\|_p$, $p \in [1, \infty]$ are two convex functions on $K \in \mathbb{S}_d^{++}$, and from the fact that the nonnegative weighted sum of two convex functions is another convex function [32]. This proves 1.

Now, we claim that $g(K)$ is radially unbounded, i.e., $g(K) \rightarrow \infty$ as $\|\mathbf{vec}(K)\|_p \rightarrow \infty$. To see this, recall that $\text{trace}(KA_n) - \log |K|$ is a differentiable convex function in K that is minimized whenever $K^{-1} = A_n$, since A_n is invertible for $n > d$ almost surely. Then,

$$g(K) = \text{trace}(KA_n) - \log |K| + \delta^{1/\rho} \|\mathbf{vec}(K)\|_p \geq d - \log |A_n^{-1}| + \delta^{1/\rho} \|\mathbf{vec}(K)\|_p,$$

from which it immediately follows that $g(K) \rightarrow \infty$ as $\|\mathbf{vec}(K)\|_p \rightarrow \infty$.

Now, again, since $g(K)$ is also convex and continuous in \mathbb{S}_d^{++} , we conclude that the level sets $\kappa_b = \{K \mid g(K) \leq b\}$ are compact and nonempty as long as $b > l(W; K) + \delta^{1/\rho} \|\mathbf{vec}(K)\|_p$. This proves 2. Moreover, since $l(W; K)$ is convex and continuous on any $K \in \mathbb{S}_d^{++}$, it follows that $E_P[l(W; K)]$ for any $P \in \{P \mid \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta\}$ is also continuous and convex on any $K \in \mathbb{S}_d^{++}$, thus 3 follows immediately. ■

Lemma 2.2.1 states that any estimator obtained by minimizing the left-hand side of (2.14) must be in $\mathcal{C}_n(\delta)$, otherwise the right-hand side of (2.14) would be strictly greater than the left. Thus, in line with the goal of providing a robust estimator, the idea is to choose δ so that $\mathcal{C}_n(\delta)$ also contains the true inverse covariance matrix Ω with high confidence.

As \mathbb{P}_0 is the weak limit of \mathbb{P}_n , we will eventually have that $\Omega \in \mathcal{C}_n(\delta)$ with high probability for any δ , so that $\mathcal{C}_n(\delta)$ is a confidence region for Ω . From this observation, we can choose the uncertainty size δ optimally by the criterion

$$\delta = \inf \{ \delta > 0 \mid P(\Omega \in \mathcal{C}_n(\delta)) \geq 1 - \alpha \}, \quad (2.15)$$

i.e., for a specified confidence level $1 - \alpha$, we choose δ so that $\mathcal{C}_n(\delta)$ is a $(1 - \alpha)$ -confidence region for Ω .

To continue our analysis, we make use of the so-called *Robust Wasserstein Profile (RWP) function* R_n introduced by [15],

$$\begin{aligned} R_n(K) &= \inf \{ \mathcal{D}_c(P, \mathbb{P}_n) \mid P \in \mathcal{O}(K) \} \\ &= \inf \left\{ \mathcal{D}_c(P, \mathbb{P}_n) \Big|_{E_P} \left[\frac{\partial}{\partial K'} l(W; K') \Big|_{K'=K} \right] = \mathbb{0}_{d \times d} \right\}, \end{aligned} \quad (2.16)$$

for $K \in \mathbb{S}_d^{++}$, which has the geometric interpretation of being the minimum distance between the empirical distribution and any distribution that satisfies the optimality condition for the precision matrix K . Then, using the equivalence of events $\{\Omega \in \mathcal{C}_n(\delta)\} = \{\mathcal{O}(\Omega) \cap \{P \mid \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta\} \neq \emptyset\} = \{R_n(\Omega) \leq \delta\}$, (2.15) becomes equivalent to

$$\delta = \arg \inf \{ \delta > 0 \mid P(R_n(\Omega) \leq \delta) \geq 1 - \alpha \}, \quad (2.17)$$

i.e., the optimal selection of δ is the $1 - \alpha$ quantile of $R_n(\Omega)$. Indeed, the set $\{P \mid \mathcal{D}_c(P, \mathbb{P}_n) \leq R_n(\Omega)\}$ is the smallest ambiguity set around the empirical measure \mathbb{P}_n such that there exists a distribution for which Ω is an optimal loss minimization parameter. In contrast to previously reported applications of the RWP function on linear regression and logistic regression [15], our problem allows for a (finite sample) closed form expression of this function. This is due to the fact that we have recast the

covariance Σ as the mean of the random matrix XX^\top , so that the following result gives a nontrivial generalization of [15, Example 3].

Theorem 2.2.2 (RWP function) *Consider the cost function in (2.5) for a fixed $\rho \geq 1$. For $K \in \mathbb{S}_d^{++}$, consider $R_n(K)$ as in (2.16). Then,*

$$R_n(K) = \|\mathbf{vec}(A_n - K^{-1})\|_q^\rho. \quad (2.18)$$

Proof:

Consider $K \in \mathbb{S}_d^{++}$. Setting $h(U; K) = U - K^{-1}$, it is clear that we satisfy the conditions for applying Proposition A.2.1 in the Appendix, and so we obtain

$$R_n(K) = \sup_{\Lambda \in \mathbb{S}_d} \left\{ -\frac{1}{n} \sum_{i=1}^n \sup_{U \in \mathbb{S}_d} \{ \text{trace}(\Lambda^\top (U - K^{-1})) - \|\mathbf{vec}(U) - \mathbf{vec}(W_i)\|_q^\rho \} \right\} \quad (2.19)$$

Now, letting $\Delta := U - W_i^\top$

$$\begin{aligned} & \sup_{U \in \mathbb{S}_d} \{ \text{trace}(\Lambda^\top (U - K^{-1})) - \|\mathbf{vec}(U) - \mathbf{vec}(W_i)\|_q^\rho \} \\ &= \sup_{\Delta \in \mathbb{S}_d} \{ \text{trace}(\Lambda^\top (\Delta + W_i - K^{-1})) - \|\mathbf{vec}(\Delta)\|_q^\rho \} \\ &= \sup_{\Delta \in \mathbb{S}_d} \{ \text{trace}(\Lambda^\top \Delta) - \|\mathbf{vec}(\Delta)\|_q^\rho \} + \text{trace}(\Lambda^\top (W_i - K^{-1})) \\ &= \sup_{\Delta \in \mathcal{M}(\Lambda)} \{ \|\mathbf{vec}(\Lambda)\|_p \|\mathbf{vec}(\Delta)\|_q - \|\mathbf{vec}(\Delta)\|_q^\rho \} \\ &+ \text{trace}(\Lambda^\top (W_i - K^{-1})) \end{aligned} \quad (2.20)$$

with $\mathcal{M}(\Lambda)$ as in the proof of Theorem 2.1.1, so that the third line follows from selecting a $\Delta \in \mathbb{S}_d$ such that Holder's inequality holds tightly (with $\frac{1}{p} + \frac{1}{q} = 1$), whose existence has been explained in the proof of Theorem 2.1.1. Thus, we are still free to choose the

magnitude of the q -norm of such $\mathbf{vec}(\Delta)$ (and this is what we will use next).

Now, the argument inside the supremum in the last line of (2.20) is a polynomial function on $\|\mathbf{vec}(\Delta)\|_q$. We have to analyze two cases.

Case 1: $\rho = 1$. In this case we observe that, by setting

$$\epsilon(\Lambda) = \sup_{\Delta \in \mathbb{S}_d} \{ \|\mathbf{vec}(\Delta)\|_q (\|\mathbf{vec}(\Lambda)\|_p - 1) \}:$$

- if $\|\mathbf{vec}(\Lambda)\|_p \leq 1$, then $\epsilon(\Lambda) = 0$ (in particular, if $\|\mathbf{vec}(\Lambda)\|_p < 1$, the optimizer is $\Delta = \mathbf{0}_{d \times d}$);
- if $\|\mathbf{vec}(\Lambda)\|_p > 1$, then $\epsilon(\Lambda) = \infty$;

so that, recalling (2.19), we see that if $\|\mathbf{vec}(\Lambda)\|_p > 1$, then $R_n(K) = -\infty$. Then, we obtain that

$$\begin{aligned} R_n(K) &= \sup_{\Lambda \in \mathbb{S}_d: \|\mathbf{vec}(\Lambda)\|_p \leq 1} \left\{ -\frac{1}{n} \sum_{i=1}^n \text{trace}(\Lambda^\top (W_i - K^{-1})) \right\} \\ &= \sup_{\Lambda \in \mathbb{S}_d: \|\mathbf{vec}(\Lambda)\|_p \leq 1} \left\{ -\text{trace}(\Lambda^\top (A_n - K^{-1})) \right\} \\ &= \sup_{\Lambda \in \mathbb{S}_d: \|\mathbf{vec}(\Lambda)\|_p \leq 1} \left\{ \mathbf{vec}(\Lambda)^\top \mathbf{vec}(A_n - K^{-1}) \right\} \\ &= \|\mathbf{vec}(A_n - K^{-1})\|_q \end{aligned}$$

where the third line results from the fact that Λ is a free variable so we can flip its sign, and the last line follows from the the analysis of conjugate norms and the fact that $\Lambda, A_n - K^{-1} \in \mathbb{S}_d$. We thus obtained (2.18).

Case 2: $\rho > 1$. By differentiation and basic calculus (e.g., using the first and second derivative test) we obtain that the maximizer

$$\Delta^* = \arg \sup_{\Delta \in \mathbb{S}_d} \{ \|\mathbf{vec}(\Delta)\|_q \|\mathbf{vec}(\Lambda)\|_p - \gamma \|\mathbf{vec}(\Delta)\|_q^\rho \}$$

is such that $\|\mathbf{vec}(\Delta^*)\|_q = \left(\frac{\|\mathbf{vec}(K)\|_p}{\rho}\right)^{\frac{1}{\rho-1}}$. Then, replacing this back in (2.19),

$$\begin{aligned}
R_n(K) &= \sup_{\Lambda \in \mathbb{S}_d} \left\{ -\frac{1}{n} \sum_{i=1}^n \left(\|\mathbf{vec}(\Lambda)\|_p^{\frac{\rho}{\rho-1}} \frac{\rho-1}{\rho^{\frac{\rho}{\rho-1}}} + \text{trace}(\Lambda^\top (W_i - K^{-1})) \right) \right\} \\
&= \sup_{\Lambda \in \mathbb{S}_d} \left\{ -\|\mathbf{vec}(\Lambda)\|_p^{\frac{\rho}{\rho-1}} \frac{\rho-1}{\rho^{\frac{\rho}{\rho-1}}} - \text{trace}(\Lambda^\top (A_n - K^{-1})) \right\} \\
&= \sup_{\Lambda \in \mathbb{S}_d} \left\{ \text{trace}(\Lambda^\top (A_n - K^{-1})) - \|\mathbf{vec}(\Lambda)\|_p^{\frac{\rho}{\rho-1}} \frac{\rho-1}{\rho^{\frac{\rho}{\rho-1}}} \right\} \\
&= \sup_{\Lambda \in \mathbb{S}_d} \left\{ \|\mathbf{vec}(\Lambda)\|_p \|\mathbf{vec}(A_n - K^{-1})\|_q - \|\mathbf{vec}(\Lambda)\|_p^{\frac{\rho}{\rho-1}} \frac{\rho-1}{\rho^{\frac{\rho}{\rho-1}}} \right\}.
\end{aligned}$$

Again, by differentiation and basic calculus, we obtain that the maximizer Λ^* is such that $\|\mathbf{vec}(\Lambda^*)\|_p = \rho \|\mathbf{vec}(A_n - K^{-1})\|_q^{\rho-1}$. Replacing this value back in our previous expression, we get that $R_n(K) = \|\mathbf{vec}(A_n - K^{-1})\|_q^\rho$, and thus we showed (2.18). \blacksquare

We now establish important convergence guarantees on the RWP function in the following corollary.

Corollary 2.2.3 (Asymptotic behavior of the RWP function) *Suppose that the conditions of Theorem 2.2.2 hold, and that $E_{\mathbb{Q}_0}(\|X\|_2^4) < \infty$. Let $H \in \mathbb{S}_d$ be a matrix of jointly Gaussian random variables with zero mean and such that $\text{Cov}(h_{ij}, h_{kl}) = E[w_{ij}w_{kl}] - \Sigma_{ij}\Sigma_{kl} = E[x_i x_j x_k x_\ell] - \Sigma_{ij}\Sigma_{kl}$. Then,*

$$n^{\rho/2} R_n(\Omega) \Rightarrow \|\mathbf{vec}(H)\|_q^\rho. \quad (2.21)$$

Proof: By the central limit theorem, we observe that $\sqrt{n}(A_n - \Sigma) \Rightarrow H$, and by the continuous mapping theorem, we get that

$$n^{\rho/2} R_n(\Omega) = \|\sqrt{n} \mathbf{vec}(A_n - \Sigma)\|_q^\rho \Rightarrow \|\mathbf{vec}(H)\|_q^\rho.$$

\blacksquare

Remark 2.2.4 Turning our attention back to Theorem 2.1.1, a robust selection for the ambiguity size or regularization parameter $\lambda = \delta^{1/\rho}$, as obtained from Theorem 2.2.2, is

$$\delta^{1/\rho} = \inf \left\{ \delta > 0 \mid P(\|\mathbf{vec}(A_n - \Sigma)\|_q \leq \delta) \geq 1 - \alpha \right\} \quad (2.22)$$

As a result, this robust selection for λ results in a class of estimators, given by minimizers of the right-hand side of (2.6), that are invariant to the choice of ρ in (2.5). Thus, for simplicity, we will set $\rho = 1$ in the remainder of the paper.

Remark 2.2.5 Let $r_{1-\alpha}$ be the $(1-\alpha)$ quantile from the distribution of the right-hand side of (2.21). Then, for any fixed α , the robust selection δ in (2.22) satisfies $n^{1/2}\delta \rightarrow r_{1-\alpha}$, so that the optimal decay rate of $n^{-1/2}$ for λ is automatically chosen by the RWP function.

As solving (2.22) requires knowledge of Σ , we now outline the robust selection (RobSel) algorithm for data-adaptive choice of the regularization parameter δ for our inverse covariance estimation with an ℓ_p penalization parameter. The special case $p = 1$ corresponds to the graphical lasso in (1.2), in which case we will also use the notation $\delta = \lambda$. The asymptotic result in Corollary 2.2.3 invokes a central limit theorem, and thus motivates the approximation of the RWP function through bootstrapping, which we further explain and evaluate its numerical performance in the next section. Let $\alpha \in (0, 1)$ be a prespecified confidence level and B a large integer such that $(B + 1)(1 - \alpha)$ is also an integer.

Algorithm RobSel algorithm for estimation of the regularization parameter λ

- 1: For $b = 1, \dots, B$, obtain a bootstrap sample $X_{1b}^*, \dots, X_{nb}^*$ by sampling uniformly and with replacement from the data, and compute the bootstrap RWP function $R_{n,b}^* = \|A_{n,b}^* - A_n\|_q$, with the empirical covariance $A_{n,b}^*$ computed from the bootstrap sample.
 - 2: Set λ to be the bootstrap order statistic $R_{n,((B+1)(1-\alpha))}^*$.
-

RobSel can potentially provide considerable computational savings over cross-validation in practice. Computing sample covariance matrices for each of B bootstrap samples has cost $O(Bnd^2)$. On the other hand, it is known that each iteration of graphical lasso can cost $O(d^3)$ in the worst case [30]; therefore, performing an F -fold cross-validation to search over L -grid of regularization parameters, each taking T -iterations of graphical lasso, would cost $O(FLTd^3)$.

Chapter 3

Numerical Evaluation

The true precision matrix $\Omega \in \mathbb{S}_d^{++}$ used to generate simulated data has been constructed as follows. First, generate an adjacency matrix of an undirected Erdős-Renyi graph with equal edge probability and without self-loops. Then, the weight of each edge is sampled uniformly between $[0.5, 1]$, and the sign of each non-zero weight is positive or negative with equal probability of 0.5. Finally, the diagonal entries of this weighted adjacency matrix are set to 1 and the matrix is made diagonally dominant by following a procedure described in [34], which ensures that the resulting matrix Ω is positive definite. Throughout this numerical study section, a randomly generated sparse matrix Ω (edge probability 0.1 and $d = 100$) is fixed. Using this Ω , a total of $N = 200$ datasets (of varying size n) were generated as independent observations from a multivariate zero-mean Gaussian distribution, i.e., $\mathcal{N}(\mathbf{0}_d, \Omega^{-1})$.

Consider the problem of choosing the regularization parameter λ (equivalently, the ambiguity size parameter δ) to obtain graphical lasso estimates \hat{K}_λ of Ω using the simulated datasets. An R software package, `glasso`, from CRAN was used throughout our numerical experiments. Below, we compare two different criteria for choosing λ . The first criterion is *Robust Selection (RS)*, which follows our proposed RobSel algorithm

with $B = 200$ sets of bootstrap samples. We present here results mainly for $n > d$, but additional results in the high-dimensional regime $n < d$ can be found in the supplementary material. The second criterion is a 5-fold *cross-validation (CV)* procedure. The performance on the validation set is the evaluation of the graphical loss function under the empirical measure of the samples on the training set.

Recall the elements in the confusion matrix to be true positives (TP), true negatives (TN), false positives (FP) and false negatives (FN). We compare model selection performance of λ chosen by the two different approaches: λ_{RS} and λ_{CV} . The following comparison metrics are used:

- *True positive rate (TPR) and false detection rate (FDR)*: $TPR = \frac{TP}{TP+FN}$ is the proportion of nonzero entries of Ω that are correctly identified in \hat{K}_λ , and $FDR = \frac{FN}{FN+TP}$ is the proportion of zero entries of Ω that are incorrectly identified as nonzeros in \hat{K}_λ .
- *Matthew's Correlation Coefficient (MCC)*: MCC summarizes all counts in the confusion matrix in contrast to other measures like TPR and FDR. More details about MCC is given in supplemental subsection D.1.

In the remainder of this section, we compare the model selection performance (FDR, TPR, MCC) from our simulation results. As mentioned in Remark 2.2.5, supplemental subsection D.2 shows that λ_{RS} decreases as n increases as is also observed to be true with λ_{CV} . Furthermore, across the tested range of α , the regularization λ_{RS} are all larger than λ_{CV} for any n . Then, our distributionally robust representation (2.6) allows us to observe that even for small values of n , CV always chooses a λ that corresponds to smaller ambiguity sets than RS.

To assess the accuracy of RobSel in estimating $\lambda = \delta$ for a given α , we approximated the right-hand side of (2.22) using the $N = 200$ data sets and the true covariance Σ ,

giving the “true” value λ_{RWP} . Figures in supplemental subsection D.3 show that the performance obtained by λ_{RWP} is similar to the one obtained by λ_{RS} for all comparison metrics. This finite sample behavior of RS indicates that the RobSel bootstrap algorithm reliably approximates the desired robustness level corresponding to the choice of α . These plots also indicate that the RS criterion is more conservative than CV in achieving a lower FDR across different sample sizes, due to providing larger values for λ (see supplemental subsection D.2). More specifically, Fig. 3.1 shows that RS gives a better performance than CV in terms of FDR even for smaller values of n and this performance improves even more as n increases.

Moreover, the trade-off between the preference for robustness and the preference for a higher density estimation of nonzero entries in the precision matrix can be observed in terms of the Matthews correlation coefficient (MCC), as shown in Fig. 3.2. Higher values in the curve of MCC implies values of λ that describe a better classification of the entries of Ω as either zero or nonzero (see supplemental subsection D.1 for more details). We observe that CV is placed at the left of the optimal value of the MCC and its under-performance is due to the overestimation of nonzero entries in Ω . On the other hand, RS is to the right of the optimal value and its under-performance is due to its conservative overestimation of zero entries in Ω induced by its robust nature. Then, it is up to the experimenter to know which method to use depending on whether she desires to control for the overestimation of zero or nonzero entries. Remarkably, for large numbers of n , RS seems to be much closer to the optimal performance than CV according to the MCC, and it does this by maintaining a lower FDR than CV while increasing its TPR.

Our results from the MCC analysis and supplemental subsection D.3 also indicate that we should aim for higher values of α if we want a performance closer to CV in terms of TPR when using the RS criterion, with the advantage of still maintaining a better performance than CV in terms of the FDR. In contrast, if we want more conservative

results, we should aim for lower values of α . This is a good property of RS: it allows the use of a single parameter $\alpha \in (0, 1)$ to adjust the importance of the regularization term. As mentioned before section 4, a practical importance of RS is that it provides a candidate for λ with potentially considerable computational savings over CV.




Figure 3.1: False Detection Rate (FDR) for seven different sample sizes, with both axes in logarithmic scale. For each sample size, the average FDR is plotted for both criteria, cross-validation (CV) and RobSel (RS). For RS, a point is plotted with a different symbol for each different value of the parameter α (some points may not be plotted for lower values of α , since those values gave no true positive detected, and so FDR was not well-defined).




Figure 3.2: Matthews correlation coefficient (MCC) for three different sample sizes. The curves in the upper plot are the average MCC obtained over $N = 200$ datasets. The lower plot are boxplots for the CV and RS (with parameter $\alpha = 0.9$) methods for the different choices of n .

Appendix A

Appendix Title

A.1 Using a different ambiguity set

Recall that $X \in \mathbb{R}^d$ is a zero-mean random vector with covariance matrix $\Sigma \in \mathbb{S}_d^{++}$ and measure \mathbb{Q}_0 , and that we consider an iid random sample $X_1, \dots, X_n \sim X$, $n > d$, with empirical measure \mathbb{Q}_n . Since we use the graphical loss function as in equation (2.1), we are interested in finding a tractable or closed-form expression for the optimization problem

$$\sup_{Q: \mathcal{D}'_{c'}(Q, \mathbb{Q}_n) \leq \delta} E_Q[l(X; K)] \tag{A.1}$$

with $K \in \mathbb{S}_d^{++}$. Ideally, the solution should be connected to the graphical lasso estimator, since it is one of the most commonly-used sparse inverse covariance estimators in practice. The ambiguity set in this formulation is specified by the collection of measures $\{Q \mid \mathcal{D}'_{c'}(Q, \mathbb{Q}_n) \leq \delta\}$, which we now describe. Given two probability distributions Q_1 and Q_2 on \mathbb{R}^d and some transportation cost function $c' : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, \infty)$ (which we

will specify below), we define the *optimal transport cost* between Q_1 and Q_2 as

$$\begin{aligned} \mathcal{D}'_{c'}(Q_1, Q_2) &= \inf\{E_\pi[c'(u, v)] \mid \pi \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d), \\ &\quad \pi_u = Q_1, \pi_v = Q_2\} \end{aligned} \quad (\text{A.2})$$

where $\mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$ is the set of joint probability distributions π of (u, v) supported on $\mathbb{R}^d \times \mathbb{R}^d$, and π_u and π_v denote the marginals of u and v under π , respectively. In this paper, we are interested in cost functions

$$c'(u, v) = \|u - v\|_q^\rho, \quad (\text{A.3})$$

with $u, v \in \mathbb{R}^d$, $\rho \geq 1$, $q \in [1, \infty]$.

Now, observe that the function $l(\cdot; K) : \mathbb{R}^d \rightarrow \mathbb{R}$ is Borel measurable since it is a continuous function. Then, we use the duality result from Proposition 4 of [15, version 2] and obtain

$$\sup_{P: \mathcal{D}'_{c'}(Q, Q_n) \leq \delta} E_Q[l(X; K)] = \inf_{\gamma \geq 0} \left\{ \gamma \delta + \frac{1}{n} \sum_{i=1}^n \left(\sup_{u \in \mathbb{R}^d} \{l(u; K) - \gamma c'(u, X_i)\} \right) \right\}. \quad (\text{A.4})$$

Let $\Delta := u - X_i$. Then

$$\begin{aligned} &\sup_{u \in \mathbb{R}^d} \{l(u; K) - \gamma c(u, X_i)\} \\ &= \sup_{u \in \mathbb{R}^d} \{u^T K u - \log |K| - \gamma \|u - X_i\|_q^\rho\} \\ &= \sup_{\Delta \in \mathbb{R}^d} \{(\Delta + X_i)^T K (\Delta + X_i) - \gamma \|\Delta\|_q^\rho\} - \log |K|. \end{aligned} \quad (\text{A.5})$$

Replacing this expression back in (A.4), it may be difficult, if not impossible, to obtain a closed form optimization problem over K . Even if such a simplification is possible, it will not provide the desired connection to the graphical lasso estimator. That is why, in

this paper, as outlined in Section 2.1, we redefine the ambiguity set to obtain a desired closed form as expressed in Theorem 2.1.1 in a more transparent way.

A.2 Applicability of the dual representations of the RWP function and the DRO formulation

The dual representations of the RWP function and the DRO formulation for the case in which the space of probability measures is $\mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$ is studied in the paper [15]. In this paper, we are interested in the case $\mathcal{P}(\mathbb{S}_d \times \mathbb{S}_d)$. In other words, we consider the samples to be in \mathbb{S}_d instead of \mathbb{R}^d . We want to emphasize that the derivations of these dual representations rely on the dual formulation of the so called “problem of moments” or a specific class of “Chebyshev-type inequalities” referenced in the work by [35]. The derivation by Isii is actually more general in the sense that is applied to more general probability spaces than the ones used in this paper and in [15] (in fact, it is stated for general spaces of non-negative measures).

Throughout this section, we consider an integrable function $h : \mathbb{S}_d \times \mathbb{S}_d \rightarrow \mathbb{S}_d$, and a lower semi-continuous function $c : \mathbb{S}_d \times \mathbb{S}_d \rightarrow [0, \infty)$ such that $c(U, U) = 0$ for any $U \in \mathbb{S}_d$ and such that the set

$$\Omega := \{(U', W') \in \mathbb{S}_d \times \mathbb{S}_d \mid c(U', W') < \infty\}$$

is Borel measurable and non-empty. Also consider an iid random sample $W_1, \dots, W_n \sim W$ with W coming from a distribution on $\mathcal{P}(\mathbb{S}_d)$.

Now, let us focus first on the RWP function in the following proposition which parallels [15, Proposition 3].

Proposition A.2.1 Consider $K \in \mathbb{S}_d^{++}$. Let $h(\cdot, K)$ be Borel measurable. Also, suppose that $\mathbb{0}_{d \times d}$ lies in the interior of the convex hull of $\{h(U', C) \mid U' \in \mathbb{S}_d\}$. Then,

$$R_n(K) = \sup_{\Lambda \in \mathbb{S}_d} \left\{ -\frac{1}{n} \sum_{i=1}^n \sup_{U \in \mathbb{S}_d} \{ \text{trace}(\Lambda^\top h(U; K)) - c(U, W_i) \} \right\}.$$

Proof: Consider the proof of [15, Proposition 3]. If we:

- set the estimating equation by $E[h(W; K)] = \mathbb{0}_{d \times d}$,
- set

$$R_n(K) = \inf \{ E_\pi[c(U, W)] \mid E_\pi[h(U; K)] = \mathbb{0}_{n \times n}, \pi_W = \mathbb{P}_n, \pi \in \mathcal{P}(\mathbb{S}_d \times \mathbb{S}_d) \},$$

with π_W denoting the marginal distribution of W ,

- consider the previously defined Ω ,

then we obtain that, following the rest of this proof (and using [35, Theorem 1] with its special case):

$$\begin{aligned} R_n(K) &= \sup_{\Lambda \in \mathbb{S}_d} \left\{ -\frac{1}{n} \sum_{i=1}^n \sup_{U \in \mathbb{S}_d} \{ \mathbf{vec}(\Lambda)^\top \mathbf{vec}(h(U; K)) - c(U, W_i) \} \right\} \\ &= \sup_{\Lambda \in \mathbb{S}_d} \left\{ -\frac{1}{n} \sum_{i=1}^n \sup_{U \in \mathbb{S}_d} \{ \text{trace}(\Lambda^\top h(U; K)) - c(U, W_i) \} \right\}, \end{aligned}$$

thus obtaining the dual representation of the RWP function. ■

The following proposition for the dual representation of the DRO formulation parallels [15, Proposition 1].

Proposition A.2.2 For $\gamma \geq 0$ and loss functions $l(U'; K)$ that are upper semi-continuous

in $U' \in \mathbb{S}_d$ for each $K \in \mathbb{S}_d^{++}$, let

$$\phi_\gamma(W_i; K) = \sup_{U \in \mathbb{S}_d} \{l(U; K) - \gamma c(U, W_i)\}. \quad (\text{A.6})$$

Then

$$\sup_{P: \mathcal{D}_c(P, \mathbb{P}_n) \leq \delta} E_P[l(W; K)] = \min_{\gamma \geq 0} \left\{ \gamma \delta + \frac{1}{n} \sum_{i=1}^n \phi_\gamma(W_i; K) \right\}.$$

Proof: The proof for the dual representation of the DRO for our domain of symmetric matrices is also very similar to the one described in Proposition 4 of [15, version 2], just by following appropriate similar changes as we did for the proof of A.2.1. ■

A.3 Additional figures, tables and additional analysis for the numerical results

A.3.1 Matthews correlation coefficient analysis

Let true positives (TP) be the number of nonzero off-diagonal entries of Ω that are correctly identified, false negatives (FN) be the number of its nonzero off-diagonal entries that are incorrectly identified as zeros, false positives (FP) be the number of its zero off-diagonal entries that are incorrectly identified as nonzeros, and true negatives (TN) be the number of its zero off-diagonal entries that are correctly identified. Given the estimated precision matrix \hat{K} , the Matthews correlation coefficient (MCC) [36] is defined as:

$$MCC = \frac{TP \cdot TN - FP \cdot FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}, \quad (\text{A.7})$$

and, whenever the denominator is zero, it can be arbitrarily set to one. It can be shown that $MCC \in [-1, +1]$, and $+1$ is interpreted as a perfect prediction (of both zero and

nonzero values), 0 is interpreted as prediction no better than a random one, and -1 is interpreted as indicating total disagreement between prediction and observation.

MCC has been argued to be one of the most informative coefficients for assessing the performance of binary classification (in this case, classifying if an entry of the precision matrix is zero or nonzero) since it summarizes all information from the TP, TN, FP and FN quantities [37, 36], in contrast to other measures like TPR and FPR.

A.3.2 Regularization parameters

All the plots related to the RS criterion for choosing λ have in their x-axis the values $\alpha \in \{0.10, 0.21, 0.33, 0.44, 0.56, 0.67, 0.79, 0.90\}$. We study the cases for sample sizes $n \in \{75, 200, 1000\}$.






Figure A.1: $n = 75$

Figure A.2: $n = 200$ Figure A.3: $n = 1000$

A.3.3 Performance figures for different choices of the regularization parameter

All the plots related to the RS and RWP criteria for choosing λ have in their x-axis the values

$\alpha \in \{0.10, 0.21, 0.33, 0.44, 0.56, 0.67, 0.79, 0.90\}$. We study the cases for sample sizes $n \in \{75, 200, 1000\}$.

$n = 75$




Figure A.4: True positive rate (%)




Figure A.5: False detection rate (%)

$n = 200$




Figure A.6: True positive rate (%)

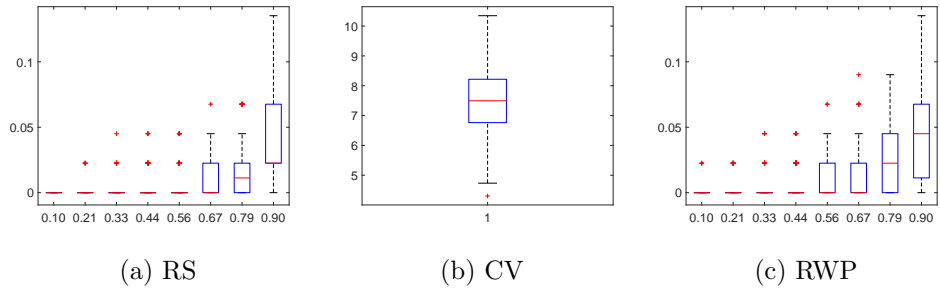


Figure A.7: False detection rate (%)

$n = 1000$

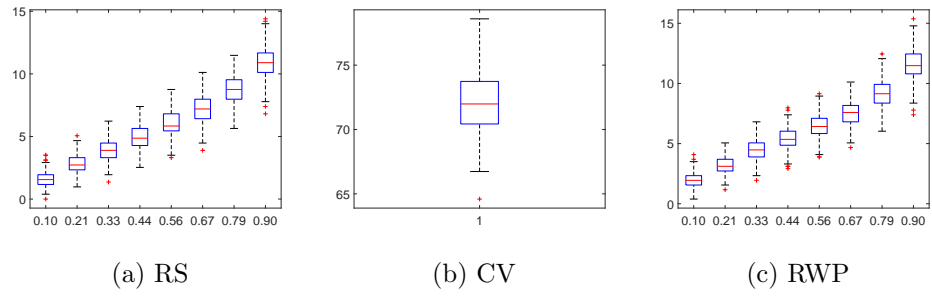


Figure A.8: True positive rate (%)

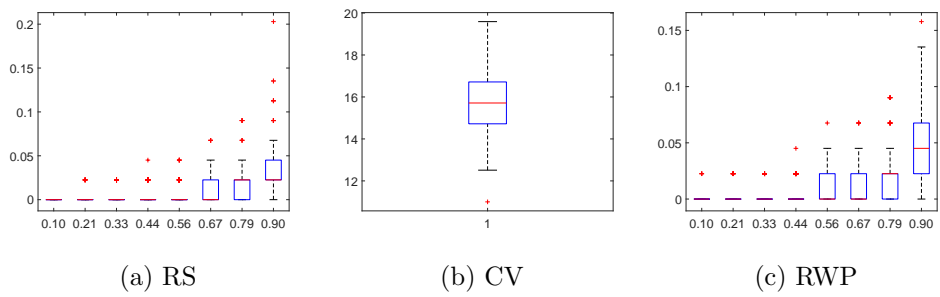


Figure A.9: False detection rate (%)

Bibliography

- [1] D. Guillot, B. Rajaratnam, and J. Emile-Geay, *Statistical paleoclimate reconstructions via markov random fields*, *Ann. Appl. Stat.* **9** (03, 2015) 324–352.
- [2] S. Huang, J. Li, L. Sun, J. Ye, A. Fleisher, T. Wu, K. Chen, and E. Reiman, *Learning brain connectivity of Alzheimer’s disease by sparse inverse covariance estimation*, *NeuroImage* **50** (2010), no. 3 935 – 949.
- [3] J. Krumsiek, K. Suhre, T. Illig, J. Adamski, and F. J. Theis, *Gaussian graphical modeling reconstructs pathway reactions from high-throughput metabolomics data*, *BMC Systems Biology* **5** (Jan, 2011) 21.
- [4] T. Cai, W. Liu, and X. Luo, *A constrained ℓ_1 minimization approach to sparse precision matrix estimation*, *Journal of the American Statistical Association* **106** (jun, 2011) 594–607.
- [5] D. Edwards, *Introduction to Graphical Modelling*. Springer-Verlag New York, 2010.
- [6] J. Friedman, T. Hastie, and R. Tibshirani, *Sparse inverse covariance estimation with the graphical lasso*, *Biostatistics* **9** (12, 2007) 432–441.
- [7] K. Khare, S.-Y. Oh, and B. Rajaratnam, *A convex pseudolikelihood framework for high dimensional partial correlation estimation with convergence guarantees*, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **77** (sep, 2015) 803–825.
- [8] J.-H. Won, J. Lim, S.-J. Kim, and B. Rajaratnam, *Condition-number-regularized covariance estimation*, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **75** (dec, 2012) 427–450.
- [9] M. Yuan and Y. Lin, *Model selection and estimation in the Gaussian graphical model*, *Biometrika* **94** (2007), no. 1 19–35.
- [10] S. Shafieezadeh-Abadeh, P. M. Esfahani, and D. Kuhn, *Distributionally robust logistic regression*, in *Advances in Neural Information Processing Systems 28*, pp. 1576–1584. 2015.

- [11] J. Blanchet and K. Murthy, *Quantifying distributional model risk via optimal transport*, .
- [12] S. Shafieezadeh-Abadeh, D. Kuhn, and P. M. Esfahani, *Regularization via mass transportation*, *Journal of Machine Learning Research* **20** (2019).
- [13] V. A. Nguyen, D. Kuhn, and P. M. Esfahani, *Distributionally robust inverse covariance estimation: The Wasserstein shrinkage estimator*, .
- [14] J. Blanchet and N. Si, *Optimal uncertainty size in distributionally robust inverse covariance estimation*, .
- [15] J. Blanchet, Y. Kang, and K. Murthy, *Robust Wasserstein profile inference and applications to machine learning*, .
- [16] O. Banerjee, L. E. Ghaoui, and A. d’Aspremont, *Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data*, *J. Mach. Learn. Res.* **9** (2008) 485–516.
- [17] B. Rolfs, B. Rajaratnam, D. Guillot, I. Wong, and A. Maleki, *Iterative thresholding algorithm for sparse inverse covariance estimation*, in *Advances in Neural Information Processing Systems 25* (F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, eds.), pp. 1574–1582. Curran Associates, Inc., 2012.
- [18] C.-J. Hsieh, M. A. Sustik, I. S. Dhillon, and P. Ravikumar, *Quic: Quadratic approximation for sparse inverse covariance estimation*, *Journal of Machine Learning Research* **15** (2014) 2911–2947.
- [19] W. James and C. Stein, *Estimation with quadratic loss*, in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, (Berkeley, Calif.), pp. 361–379, University of California Press, 1961.
- [20] A. J. Rothman, P. J. Bickel, E. Levina, and J. Zhu, *Sparse permutation invariant covariance estimation*, *Electron. J. Statist.* **2** (2008) 494–515.
- [21] J. Jankova and S. van de Geer, *Inference in high-dimensional graphical models*, .
- [22] C. Lam and J. Fan, *Sparsity and rates of convergence in large covariance matrix estimation*, *The Annals of Statistics* **37** (2009), no. 6B 4254–4278.
- [23] P.-L. Loh and X. L. Tan, *High-dimensional robust precision matrix estimation: Cellwise corruption under ϵ -contamination*, *Electron. J. Statist.* **12** (2018), no. 1 1429–1467.
- [24] L. Xue and H. Zou, *Regularized rank-based estimation of high-dimensional paranormal graphical models*, *The Annals of Statistics* **40** (2012), no. 5 2541–2571.

- [25] E. Yang and A. C. Lozano, *Robust Gaussian graphical modeling with the trimmed graphical lasso*, in *Advances in Neural Information Processing Systems 28*, pp. 2602–2610. 2015.
- [26] S. Yang, Q. Sun, S. Ji, P. Wonka, I. Davidson, and J. Ye, *Structural graphical lasso for learning mouse brain connectivity*, in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, 2015.
- [27] S. M. Smith, K. L. Miller, G. Salimi-Khorshidi, M. Webster, C. F. Beckmann, T. E. Nichols, J. D. Ramsey, and M. W. Woolrich, *Network modelling methods for fMRI*, *NeuroImage* **54** (2011), no. 2 875 – 891.
- [28] P. Menndez, Y. Kourmpetis, C. ter Braak, and F. van Eeuwijk, *Gene regulatory networks from multifactorial perturbations using graphical lasso: Application to the dream4 challenge*, *PLOS ONE* **5** (12, 2010) 1–8.
- [29] N. Sulaimanov, S. Kumar, H. Koepl, F. Burdet, M. Pagni, and M. Ibberson, *Inferring gene expression networks with hubs using a degree weighted Lasso approach*, *Bioinformatics* **35** (08, 2018) 987–994.
- [30] R. Mazumder and T. Hastie, *Exact covariance thresholding into connected components for large-scale graphical lasso*, *J. Mach. Learn. Res.* **13** (2012), no. 1 781–794.
- [31] P. M. Esfahani and D. Kuhn, *Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations*, *Mathematical Programming* **171** (2018), no. 1-2 115–166.
- [32] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [33] J. M. Steele, *The Cauchy-Schwarz Master Class: An Introduction to the Art of Mathematical Inequalities*. Cambridge University Press, 2004.
- [34] J. Peng, P. Wang, N. Zhou, and J. Zhu, *Partial correlation estimation by joint sparse regression models*, *Journal of the American Statistical Association* **104** (2009), no. 486 735–746.
- [35] K. Isii, *On sharpness of tchebycheff-type inequalities*, *Annals of the Institute of Statistical Mathematics* **14** (1962), no. 1 185–197.
- [36] D. M. Powers, *Evaluation: From precision, recall and f-factor to roc, informedness, markedness & correlation*, *Journal of Machine Learning Technologies* **2** (2011) 37–63.
- [37] D. Chicco, *Ten quick tips for machine learning in computational biology*, *BioData Mining* **10** (Dec, 2017) 35.