# Variable Clustering via Distributionally Robust Nodewise Regression

Kaizheng Wang<sup>\*</sup> Xiao Xu<sup>†</sup> Xun Yu Zhou<sup>‡</sup>

December, 2022

#### Abstract

We study a multi-factor block model for variable clustering and connect it to the regularized subspace clustering by formulating a distributionally robust version of the nodewise regression. To solve the latter problem, we derive a convex relaxation, provide guidance on selecting the size of the robust region, and hence the regularization weighting parameter, based on the data, and propose an ADMM algorithm for implementation. We validate our method in an extensive simulation study. Finally, we propose and apply a variant of our method to stock return data, obtain interpretable clusters that facilitate portfolio selection and compare its out-of-sample performance with other clustering methods in an empirical study.

**Keywords.** Variable clustering, subspace clustering, nodewise regression, regularization, distributionally robust optimization, portfolio selection.

## 1 Introduction

The rapid development of technologies has created an enormous amount of data in many fields. Such high-dimensional data often have many similar variables, in the sense that they convey a similar message and hence are replaceable with one another for certain tasks. It would then be useful to identify groups of similar variables and reduce the data complexity. This problem is called variable clustering.

Generally speaking, variable clustering is the problem of grouping similar components of a *d*-dimensional random vector  $X = (X_1, \ldots, X_d)$ . The resulting groups are referred to as *clusters*.

Author names are sorted alphabetically.

<sup>\*</sup>Department of Industrial Engineering and Operations Research & The Data Science Institute, Columbia University, New York, NY 10027, USA. Email: kaizheng.wang@columbia.edu.

<sup>&</sup>lt;sup>†</sup>Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027, USA. Email: xiao.xu@columbia.edu.

<sup>&</sup>lt;sup>‡</sup>Department of Industrial Engineering and Operations Research & The Data Science Institute, Columbia University, New York, NY 10027, USA. Email: xz2574@columbia.edu.

In many applications, the problem of interest is to recover the clusters from a sample of n independent copies, or observations, of X. This is essentially clustering the d vectors, each having the n observations. Examples of the application of variable clustering include gene expression data (Jiang et al., 2004) and protein profile data (Bernardes et al., 2015). Recently, Tang et al. (2022) apply variable clustering to returns data of financial assets. Based on the clustering results, a small set of diversified stocks are selected from a large stock universe, and the resulting portfolios show outstanding performances.

Many solutions to the variable clustering problem, e.g., K-means, graph-based hierarchical clustering, spectral clustering, are usually based on heuristics. Model-based approaches are pursued for being more principled and more interpretable. A recent development in variable clustering is the G-block model proposed by Bunea et al. (2020), which offers clearly defined population-level clusters. Under the G-block model, the covariance matrix of X has a block structure, and the blocks correspond to the clusters in a partition G, hence the name "G-block". In the G-block model, each cluster has one latent factor. Each variable is comprised of the factor in its cluster and an idiosyncratic component. Consequently, all variables in the same cluster are noisy realizations of the same latent factor. Since their observations lie near a single point in  $\mathbb{R}^n$ , it is natural to use centroid-based clustering approaches such as k-means (Bunea et al., 2020).

A more flexible model is the multi-factor block model in Ando and Bai (2017), where each cluster may have several latent factors. Each variable is represented as a linear combination of its clusterspecific factors and an idiosyncratic component. Observations of variables in the same cluster lie near a low-dimensional subspace in  $\mathbb{R}^n$  spanned by the same set of factors. All of the *d* vectors in  $\mathbb{R}^n$ , each representing the observations of a variable, reside near a union of low-dimensional subspaces. The variable clustering problem can be solved by identifying these subspaces and their corresponding variables. This task is usually referred to as subspace clustering, for which many techniques have been developed and applied to various real-world problems ranging from computer vision to machine learning (Parsons et al., 2004; Vidal, 2011).

A majority of commonly-used approaches in subspace clustering exploit the subspace structure by nodewise regression, where each variable is regressed against all other variables. The hope is that the regressions will favor the other variables in the same cluster over variables in different clusters. This way, the regression coefficients create an association matrix that defines a weighted graph among variables. Then, the clusters can be recovered easily using, for example, spectral methods.

To deal with high dimensional data where the dimension exceeds the number of observations, regularization is added in the nodewise regression to make the regressions well-posed. To this end, sparse subspace clustering (SSC), which adds  $L_1$  regularization to the nodewise regression, is widely studied (Elhamifar and Vidal, 2013; Soltanolkotabi and Candés, 2012; Wang and Xu, 2016). However, a few drawbacks persist when using nodewise sparse regression for subspace clustering. First, tuning the parameter that controls the  $L_1$  regularization depends on the unknown variance of the idiosyncratic components and is difficult. In addition, pursuing sparsity in the regression coefficients might be unnatural since the true association matrix can be dense as long as the subspaces are not orthogonal to each other or many variables in the same subspace have non-negligible correlations. To address these drawbacks, we propose a method that naturally incorporates regularization in the nodewise regression from the perspective of distributionally robust optimization (DRO). For a review on DRO, see Rahimian and Mehrotra (2022).

The main contributions of this article are the following. First, we connect a multi-factor block model for variable clustering to the subspace clustering problem. Specifically, we show that under a multi-factor block model, the observations of the variables lie near a union of low-dimensional subspaces. Hence subspace clustering methods are directly applicable to such models. Second, we formulate a distributionally robust version of the commonly used nodewise regression method in subspace clustering. To our best knowledge, we are the first to apply DRO to nodewise regression and subsequently subspace clustering. This version of nodewise regression is motivated by the uncertainty in the data and leads to an interpretable regularization. Third, we obtain a convenient convex relaxation of distributionally robust nodewise regression, provide guidance on the choice of the size of the robust region, and propose an ADMM algorithm for efficient implementation. Instead of artificially introducing sparsity through the  $L_1$  regularization, the convex relaxation results in a natural spectral-norm regularization that does not necessarily pursue sparsity. The regularization parameter is related explicitly to the size of the robust region, which can be informed purely by the data and easily computed given a generic confidence level. In addition, our proposed ADMM algorithm significantly speeds up the calculation compared to off-the-shelf convex optimizers, which enables us to test it on dimensions higher than the hundreds. Our method is validated in an extensive simulation study. Lastly, we make an empirical contribution in obtaining interpretable clusters and utilizing them to facilitate portfolio selection. Backtesting results show that clustering information can offer a substantial advantage in obtaining diversified portfolios, and our method generally improves existing algorithms for portfolio selection.

The rest of the paper is organized as follows. In Section 2, we provide a more detailed overview of variable clustering, subspace clustering, and nodewise regression. In Section 3, we introduce our DRO nodewise regression method and present theoretical results. A simulation study and an empirical study for portfolio backtesting are reported respectively in Sections 4 and 5.

**Notation** We define a set of rules for notations that will be used throughout this paper. Boldfaced capital letters, e.g., B, represent matrices. Non-boldfaced capital letters, e.g., X, represent random variables. Non-boldfaced lower-case letters represent vectors or scalars, which will be made clear in the context. All vectors are column vectors unless otherwise stated. Lower-case letters with subscripts represent specific columns, rows, or elements of a matrix of the corresponding capital letters. For a matrix B,  $b_{\cdot i}$  or simply  $b_i$  denotes the *i*-th column vector,  $b_{j}$ . denotes the *j*-th row vector, and  $b_{ij}$  denotes the *j*-th element of the *i*-th row.

## 2 Multi-Factor Block Model and Nodewise Regression

#### 2.1 Problem setup

Given a *d*-dimensional random vector  $X = (X_1, \ldots, X_d)$  and a sample of *n* independent observations of X, we want to find similar components of this random vector.

First of all, consider the following single-factor block model, in which each random variable  $X_i$  belongs to one of the K clusters, indexed by  $z(i) \in \{1, \ldots, K\}$ . All random variables in the same

cluster k are associated with the same latent factor  $F_k$ . Formally,

$$X_i = F_{z(i)} + U_i,$$

where  $\operatorname{Cov}(F_{z(i)}, U_i) = 0$ ,  $\operatorname{Cov}(F) = \Sigma_F$ , and the idiosyncratic parts  $U_i$  are uncorrelated, i.e.,  $\operatorname{Cov}(U) = \Gamma$  which is diagonal.

This single-factor block model naturally leads to the G-block model (Bunea et al., 2020). In the G-block model, the covariance matrix of X has a block structure, with the blocks corresponding to groups of similar variables. Specifically, given a partition  $G := \{G_1, \ldots, G_K\}$  of the variable indices  $\{1, \ldots, d\}$  such that each  $G_k$  is a set of  $m_k$  indices, define the membership matrix  $\mathbf{A} \in \mathbb{R}^{d \times K}$  associated with G:  $a_{ik} = 1$  if  $i \in G_k$  and  $a_{ik} = 0$  otherwise. Suppose that G is the true underlying cluster partition of the random variables  $X_1, \ldots, X_d$ . The model assumes that the covariance matrix  $\Sigma$  of the random vector  $X \in \mathbb{R}^d$  follows a block decomposition, in which the blocks correspond to the groups in the partition G. This block structure means that variables in the same cluster have the same covariance with all other variables, and the covariance matrix  $\Sigma$  of X can be decomposed as:

$$\boldsymbol{\Sigma} = \boldsymbol{A}\boldsymbol{\Sigma}_{F}\boldsymbol{A}^{\top} + \boldsymbol{\Gamma},$$

where A is associated with the partition G,  $\Sigma_F$  is a symmetric  $K \times K$  matrix, and  $\Gamma$  is diagonal. When such a decomposition exists, we say that X follows an (exact) G-block model.

The single-factor block model justifies methods where a single variable is used to represent the entire cluster. For example, the K-means algorithm essentially approximates  $F_{z(i)}$  using the average of all  $X_i$ 's with the same z(i), and Tang et al. (2022) apply a variant of this model to clustering financial time series. The above model is arguably restrictive, as it assumes that each cluster is controlled by only one latent factor and all variables therein have the same loading. One may benefit from considering a more general model that allows the variables in the same cluster to be controlled by *a set of* factors. This motivates us to study the *multi-factor block model*, which is a natural extension of the single-factor block model.

Specifically, consider a d-dimensional random vector  $X = (X_1, \ldots, X_d)$ , and an underlying partition  $G := \{G_1, \ldots, G_K\}$  of the indices  $\{1, \ldots, d\}$ . Denote by  $m_k$  the size of cluster  $G_k$ . For each  $k = 1, \ldots, K$ , let  $F_G^k$  be a  $d_k$ -dimensional random vector that represents the factors controlling the k-th cluster, and without loss of generality, assume that  $\operatorname{Cov}(F_G^k) = I$ . We also assume that  $m_k > d_k$ , i.e., there are more variables than factors in each cluster. For each  $i = 1, \ldots, d$ , denote by  $z(i) \in 1, \ldots, K$  the index of the cluster that  $X_i$  belongs to.

**Definition 2.1** (Multi-factor block model). Under the multi-factor block model, for each i, the random variable  $X_i$  satisfies:

$$X_{i} = (F_{G}^{z(i)})^{\top} \beta_{i} + U_{i}, \qquad (2.1)$$

where  $\beta_i \in \mathbb{R}^{d_k}$  is the loadings of the *i*-th variable on the factors  $F_G^{z(i)}$  and  $U_i$  is a one-dimensional random variable that represents the idiosyncratic part satisfying  $\text{Cov}(U_i, U_j) = 0$  for  $i \neq j$ .

We note that, although the factors in the same cluster are assumed to be uncorrelated in this model, factors controlling different clusters are not necessarily uncorrelated.

The multi-factor block model studied in this paper is similar to the multi-factor asset pricing model studied in Ando and Bai (2017), which we briefly review below.

Let  $R = (R_1, \ldots, R_d)^{\top}$  represent the returns of d stocks and let  $R_i^* := (R_i - \mathbb{E}(R_i))/\sqrt{\operatorname{Var}(R_i)}$ be the standardized return. Let  $\{G_k\}_{k=1}^K$  form a partition of  $\{1, \ldots, d\}$ , i.e., for  $i \in \{1, \ldots, d\}$ , there exists a unique  $z(i) \in \{1, \ldots, K\}$  such that  $j \in G_{z(i)}$ . The returns of the stocks are characterized by:

$$R_i^* = F_C^{\top} \beta_C(i) + F_H^{\top} \beta_H(i) + (F_G^{z(i)})^{\top} \beta_G(i) + \varepsilon_i \quad \forall i = 1, \dots d.$$

$$(2.2)$$

where

- $F_C \in \mathbb{R}^p$  is the random vector of p observable common factors;
- $F_H \in \mathbb{R}^r$  is the random vector of r unobservable common factors;
- $F_G^{z(i)} \in \mathbb{R}^{d_{z(i)}}$  is the random vector of unobservable group-specific factors that control the returns of only group z(i);
- $\beta_C(j)$ ,  $\beta_H(j)$  and  $\beta_G(j)$  are the corresponding (unknown) factor loadings;
- The idiosyncratic error term  $U = (\varepsilon_1, \ldots, \varepsilon_d)$  is assumed to be uncorrelated with the factors.

The multi-factor asset pricing model in Ando and Bai (2017) generalizes the multi-factor block model by including observable global factors. Their model is specifically motivated by and applied to financial asset pricing, while our study focuses on identifying and analyzing the subspace structure for constructing clusters.

Following the multi-factor block model (2.1), we can see that the covariance matrix also displays a block structure and can be decomposed similarly to the *G*-block model.

**Fact 1** (Multi-factor block model, matrix form). Let  $F_G \in \mathbb{R}^D$  be the vector of all latent factors stacked together:  $F_G := (F_G^{1^{\top}}, \ldots, F_G^{K^{\top}})^{\top}$ , with  $D := d_1 + \cdots + d_K$  being the total number of latent factors. We can write

$$\Sigma = A \Sigma_F A^{ op} + \Gamma,$$

where the *i*-th row of  $\mathbf{A} \in \mathbb{R}^{d \times D}$  shows loadings of  $X_i$  on all the *D* factors: if z(i) = k, then  $(d_1 + \dots + d_{k-1}) \ 0$ 's

$$a_{i\cdot} = ($$
  $0, \ldots, 0$   $, \beta_i^{\top}, \underbrace{0, \ldots, 0}_{(d_{k+1} + \cdots + d_K) \ 0's}), \Sigma_F = \operatorname{Cov}(F_G), and \Gamma = \operatorname{Cov}(U).$ 

To better illustrate the model, we provide a simple example with 5 variables, 3 factors, and 2 clusters.

**Example 1.** Let  $X = (X_1, X_2, X_3, X_4, X_5)$  be a random vector in  $\mathbb{R}^5$ . Consider a partition  $G := \{G_1, G_2\} := \{\{1, 2, 3\}, \{4, 5\}\}$ , where the first three random variables are in the same cluster, and the last two in the same cluster. Let  $d_1 = 2$ ,  $d_2 = 1$ , i.e. the first cluster is controlled by two factors, and the second cluster by one factor. Denote the latent factors by  $F_G$ , and  $F_G = (F_G^{\uparrow\uparrow}, F_G^{2\uparrow})^{\uparrow\uparrow} = (F_1^{\uparrow\uparrow}, F_2^{\uparrow\uparrow}, F_3^{\uparrow\uparrow})^{\uparrow\uparrow}$ , where  $F_1, F_2, F_3$  represent the three latent factors, whose covariance matrix is

$$\boldsymbol{\Sigma}_F = \begin{bmatrix} 1 & 0.1 & 0.5 \\ 0.1 & 1 & 0.5 \\ 0.5 & 0.5 & 1 \end{bmatrix}.$$

Each random variable  $X_i$  only loads on the latent factors controlling the corresponding cluster. The loading matrix  $\mathbf{A} \in \mathbb{R}^{5 \times 3}$  is

$$\boldsymbol{A} = \begin{bmatrix} 0.4 & 0.6 & 0 \\ 0.7 & 0.3 & 0 \\ 0.4 & 0.6 & 0 \\ 0 & 0 & 0.8 \\ 0 & 0 & 0.7 \end{bmatrix}.$$

Let the covariance of the idiosyncratic components, denoted by  $\Gamma$ , be a diagonal matrix with diagonal entries (0.1, 0.1, 0.1, 0.1, 0.1), then the covariance of the random vector X can be calculated:

$$\boldsymbol{\Sigma} = \boldsymbol{A}\boldsymbol{\Sigma}_{F}\boldsymbol{A}^{\top} + \boldsymbol{\Gamma} = \begin{bmatrix} 0.722 \ 0.478 \ 0.586 \ 0.4 \ 0.35 \\ 0.478 \ 0.722 \ 0.514 \ 0.4 \ 0.35 \\ 0.586 \ 0.514 \ 0.668 \ 0.4 \ 0.35 \\ 0.4 \ 0.4 \ 0.4 \ 0.74 \ 0.56 \\ 0.35 \ 0.35 \ 0.35 \ 0.35 \ 0.56 \ 0.59 \end{bmatrix}.$$

We observe that the covariance matrix  $\Sigma$  displays a near-block structure. Figure 1b illustrates this observation with a heatmap. One can see four blocks and similar values within the blocks. The  $3 \times 3$  block on the top left and the  $2 \times 2$  block on the bottom right have slightly higher values than the off-diagonal blocks.

#### 2.2 Subspace clustering and nodewise regression

Let  $X \in \mathbb{R}^{n \times d}$  be the data matrix of n observations of X. Suppose that the (unobserved) realizations of the latent factors are  $F_G \in \mathbb{R}^{n \times D}$ , then X can be decomposed as X = Y + U, where  $Y = F_G A$  is the group-specific component controlled by the factors, and U is the idiosyncratic components. Both Y and U are unobservable. One can see that the factor part of the *i*-th variable  $y_i = F_G^{z(i)}\beta_i$ , which is an *n*-dimensional vector, lies in a  $d_{z(i)}$ -dimensional subspace, spanned by (unobserved) factor realizations  $F_G^{z(i)} \in \mathbb{R}^{n \times d_{z(i)}}$ . For the subspaces to be meaningful, we assume that the number of observations is strictly larger than the maximum dimension of the subspaces, i.e.,  $n \ge d_k + 1$ ,  $k = 1, \ldots, K$ . Let  $S_k$  be the linear subspace of  $\mathbb{R}^n$  spanned by the columns of  $F_G^k$ , then each column of Y lies in the union of the K subspaces:

$$\mathcal{S}_1 \cup \mathcal{S}_2 \cup \ldots \cup \mathcal{S}_K.$$

The task of variable clustering is now to identify these subspaces and thus uncover the clustering structure among the random variables  $X_i$ 's, using only the observations X.

Notice that under this subspace structure, each  $y_i$  can be written as a linear combination of all other  $y_j$ 's that lie in the same subspace. To exploit the subspace structure of the group-specific components, it is then natural to perform a nodewise regression on the observations of the variables in the data matrix X, where each column  $x_i$  is regressed against all other  $x_j$ 's. Specifically, for each  $i = 1, \ldots, d$ , we solve

$$\min_{b_i \in \mathbb{R}^d} \|x_i - Xb_i\|_2^2 \quad \text{s.t.} \quad b_{ii} = 0,$$
(2.3)

In matrix form, we can write equivalently

$$\min_{\boldsymbol{B}\in\mathbb{R}^{d\times d}} \|\boldsymbol{X}-\boldsymbol{X}\boldsymbol{B}\|_{F}^{2} \quad \text{s.t.} \quad \operatorname{diag}(\boldsymbol{B}) = 0$$
(2.4)

where  $\|\cdot\|_F$  is the Frobenius norm of a matrix. The hope is that the resulting regression coefficients **B** will mainly connect vectors that are in the same cluster, i.e.,  $|b_{ij}| \approx 0$  if  $z(i) \neq z(j)$ . Then, the clusters can be easily recovered by performing, for example, spectral clustering on the symmetrized matrix  $\mathbf{C} := \mathbf{B}_{abs}^{\top} + \mathbf{B}_{abs}$ , where  $(\mathbf{B}_{abs})_{ij} = |b_{ij}|$ .

To demonstrate the efficacy of nodewise regression in subspace clustering, let us recall Example 1. With the population covariance matrix, we can calculate in closed-form the true nodewise regression coefficients **B**. This is the optimal **B** that we would obtain from (2.4) as *n* approaches infinity. The symmetrized similarity matrix **C** calculated from the optimal **B**, i.e.,  $\mathbf{C} := \mathbf{B}_{abs}^{\top} + \mathbf{B}_{abs}$ is shown below.

$$\boldsymbol{C} = \left[ \begin{array}{cccc} 0 & 0.141 & 1.357 & 0.107 & 0.085 \\ 0.141 & 0 & 0.865 & 0.179 & 0.145 \\ 1.357 & 0.865 & 0 & 0.101 & 0.079 \\ 0.107 & 0.179 & 0.101 & 0 & 1.530 \\ 0.085 & 0.145 & 0.079 & 1.530 & 0 \end{array} \right]$$

Figure 1 visualizes and compares C and  $\Sigma$  in heatmaps. We can see that in C, the two blocks along the diagonal have much larger values than the off-diagonal blocks, whereas, in  $\Sigma$ , the same blocks are more difficult to distinguish.



Figure 1: The heatmap of C compared with that of the covariance matrix from population-level nodewise regression on variables in Example 1. Diagonals of C are filled with value 2.

In summary, the subspace structure can be exploited through the following two-step clustering scheme:

- 1. Compute a similarity matrix C between vectors, ideally connecting only vectors in the same subspace with non-zero edges.
- 2. Construct clusters by applying spectral clustering techniques to C.

#### 2.3 A review of nodewise regression

In this paper, we focus on step 1, specifically using nodewise regression to obtain the similarity matrix. In this section, we briefly review some existing variants of nodewise regression applied to subspace clustering and how nodewise regression connects to other areas of research.

In the existing literature, regularization of the regression coefficients is often added to the nodewise regression to overcome overfitting due to noise in the data and to deal with the issue of the regression (2.3) becoming ill-posed when d > n. One of the commonly used regularizers is the  $L_0$  regularizer, which penalizes the number of non-zero regression coefficients. This  $L_0$  seminorm is usually relaxed to the  $L_1$ -norm as its tightest convex relaxation. The regression thus becomes the Lasso, which promotes sparse solutions and can be solved efficiently. Such subspace clustering methods using nodewise sparse regression for subspace clustering are called "sparse subspace clustering" (SSC) (Elhamifar and Vidal, 2013; Soltanolkotabi et al., 2014). Others use nodewise regression with a nuclear-norm regularization, penalizing  $\|\boldsymbol{B}\|_*$  in (2.4), thus encouraging it to be low-rank. This type of method is called "low-rank representation" (LRR) and is used in subspace clustering, segmentation, and feature extraction (Chen and Yang, 2014; Favaro et al., 2011; Liu et al., 2013; Liu and Yan, 2011).

Much of the current subspace clustering algorithms using nodewise regression can be improved. Take the Lasso-type SSC algorithm (Elhamifar and Vidal, 2013; Soltanolkotabi et al., 2014) as an example. SSC solves the following optimization problem. For every  $j = 1, \ldots, d$ ,

$$\min_{b_j \in \mathbb{R}^d} \|x_j - \mathbf{X}b_j\|_2^2 + \lambda_j \|b_j\|_1 \quad \text{s.t.} \quad b_{jj} = 0,$$
(2.5)

for all  $j = 1, \ldots, d$ . Or in matrix form,

$$\min_{\boldsymbol{B}\in\mathbb{R}^{d\times d}} \|\boldsymbol{X}-\boldsymbol{X}\boldsymbol{B}\|_{F}^{2} + \|\boldsymbol{B}\boldsymbol{\Lambda}\|_{1} \quad \text{s.t.} \quad \operatorname{diag}(\boldsymbol{B}) = \boldsymbol{0},$$
(2.6)

where  $\Lambda$  is a  $d \times d$  diagonal matrix whose diagonals are the parameters controlling the regularization  $(\lambda_1, \ldots, \lambda_d)$ . However, this approach has a few drawbacks. First of all, even under the true model, the coefficients are not necessarily sparse but usually a dense combination. Second, strong correlations among variables make  $L_1$ -based sparse recovery difficult. In addition, the appropriate  $\lambda_i$  depends on the variance of the idiosyncratic components, and heterogeneous and unknown variances of idiosyncratic components make tuning these parameters hard. As such, pursuing sparsity in the nodewise regression is unnatural and difficult to implement in practice. In comparison, we propose a method that naturally derives regularization in the nodewise regression by reformulating (2.4) in the context of distributionally robust optimization (DRO). This formulation results in a spectral-norm regularization. Importantly, the DRO analysis leads to an endogenous choice of the regularization parameter that is data driven, easy to compute and interpretable.

As a widely used technique in structural learning, the application of nodewise regression is not limited to subspace clustering. Aside from being a natural approach to subspace clustering, nodewise regression is also closely related to the popular k-means clustering. The k-means algorithm for variable clustering (Bunea et al., 2020) amounts to the following problem:

$$\min_{z_1,\dots,z_d \in [K],\mu_1,\dots,\mu_k \in \mathbb{R}^n} \left\{ \sum_{i=1}^d \|x_i - \mu_{z_i}\|_2^2 \right\},\,$$

where  $x_i \in \mathbb{R}^n$  is the observations of the *i*-th variable,  $z_i$  is the index of the cluster that  $x_i$  is assigned to, and  $\mu_k$  is the mean of all variables in cluster k. According to the analysis of k-means in Peng and Wei (2007), this optimization problem can be reformulated as a nodewise regression problem with constraints:

$$\begin{split} \min_{\boldsymbol{B} \in \mathbb{R}^{d \times d}} & \|\boldsymbol{X} - \boldsymbol{X}\boldsymbol{B}\|_F^2, \\ \text{s.t.} & \boldsymbol{B} \in \{\boldsymbol{Z}(\boldsymbol{Z}^\top \boldsymbol{Z})^{-1} \boldsymbol{Z}^\top : \ \boldsymbol{Z} \in \{0,1\}^{d \times K}, \ \boldsymbol{Z} \boldsymbol{1}_K = \boldsymbol{1}_d, \ \boldsymbol{Z}^\top \boldsymbol{1}_d > 0\}, \end{split}$$

where  $1_K$  and  $1_d$  are vectors of all ones with lengths K and d, respectively;  $\mathbf{Z}^{\top} 1_d > 0$  is an entrywise positivity constraint.

Outside of clustering, the idea of nodewise regression is also used in graphical model selection. For example, Meinshausen and Bühlmann (2006) use an  $L_1$ -norm-regularized nodewise regression to recover neighbors by estimating a sparse inverse covariance matrix. This technique is recently applied to Markowitz-type portfolio selection by Callot et al. (2019).

Finally, nodewise regression also relates to network analysis. Mantegna (1999) analyzed the financial market by creating a minimum-spanning-tree out of correlation network of stock returns. Since then, network analysis of the financial market has included structural analysis (Aste et al., 2010; Tumminello et al., 2007), clustering analysis (Musmeci et al., 2015; Rosén, 2006), and portfolio construction (Korzeniewski, 2018; León et al., 2017; Pozzi et al., 2013). Such analysis requires some sort of filtering of the fully connected network to reveal its structure, and topological filters like the maximum spanning tree or the PMFG (Tumminello et al., 2005) have been used. Similar filtering techniques are also utilized to estimate the sparse inverse covariance matrix in graphical models (Barfuss et al., 2016) and to analyze the structure of partial correlation networks (Musmeci et al., 2017). As mentioned above, regularized nodewise regression, e.g., the  $L_1$ -norm-regularized sparse nodewise regression, can also be used in estimating a sparse inverse covariance matrix, which can serve as a natural filter for the partial correlation network.

### **3** Distributionally Robust Nodewise Regression

In this section, we put the nodewise regression (2.4) in a probabilistic context. Consider the *d*dimensional random vector X whose coordinates have zero mean and unit variance. Denote by  $\mathbb{P}^*$ the true probability measure underlying the distribution of X, and  $\mathbb{E}_{\mathbb{P}^*}$  the expectation under  $\mathbb{P}^*$ . The classical least-square nodewise regression problem (2.4) is to solve:

$$\min_{\boldsymbol{B}\in\mathbb{R}^{d\times d}}\mathbb{E}_{\mathbb{P}^{*}}\left[\left\|\boldsymbol{X}-\boldsymbol{B}^{\top}\boldsymbol{X}\right\|_{2}^{2}\right] \quad \text{s.t.} \quad \operatorname{diag}(\boldsymbol{B})=0$$
(3.1)

Suppose that we have a data matrix  $\mathbf{X} := (x_1, \ldots, x_n)^{\top}$ , where  $x_t \in \mathbb{R}^d$  is the *t*-th observation of the standardized random vector X. Denote by  $\mathbb{P}_n$  the empirical distribution:

$$\mathbb{P}_n(dx) := \frac{1}{n} \sum_{t=1}^n \delta_{x_t} dx.$$

Given a cost function  $c : \mathbb{R}^m \times \mathbb{R}^m \to [0, \infty]$  where  $c(u, w) := ||w - u||_2^2$  and two probability distributions  $\mathbb{P}$  and  $\mathbb{Q}$  supported on  $\mathbb{R}^m$ , we define the optimal transport cost or discrepancy between  $\mathbb{P}$  and  $\mathbb{Q}$ , denoted by  $\mathcal{D}_c(\mathbb{P}, \mathbb{Q})$  as

$$\mathcal{D}_{c}(\mathbb{P},\mathbb{Q}) = \inf\{\mathbb{E}_{\pi}\left[c(U,W)\right] : \pi \in \mathcal{P}(\mathbb{R}^{m} \times \mathbb{R}^{m}), \pi_{U} = \mathbb{P}, \pi_{W} = \mathbb{Q}\}\$$
$$= \inf\{\mathbb{E}_{\pi}\left[\|w - u\|_{2}^{2}\right] : \pi \in \mathcal{P}(\mathbb{R}^{m} \times \mathbb{R}^{m}), \pi_{U} = \mathbb{P}, \pi_{W} = \mathbb{Q}\}.$$

Notice that this discrepancy function is the squared Wasserstein distance of order two; it can be extended to any lower semi-continuous function c such that c(u, u) = 0 for every  $u \in \mathbb{R}^m$ . As long as  $c^{1/\rho}$  is a metric for some  $\rho > 1$ ,  $\mathcal{D}^{1/\rho}(\mathbb{P}, \mathbb{Q})$  is also a metric (Villani, 2009).

Recall that our original goal is to solve (3.1), which is the expected loss under the true distribution. The plug-in method, i.e., optimizing (3.1) under  $\mathbb{P}_n$  generally yields unfavorable results that are poor out-of-sample, or under the true distribution  $\mathbb{P}^*$ . However, we cannot observe  $\mathbb{P}^*$  but can only access the empirical distribution  $\mathbb{P}_n$  inferred from the observations. The DRO approach is to postulate that  $\mathbb{P}^*$  lies somewhere close to  $\mathbb{P}_n$ , e.g., within a region of radius  $\delta$  around  $\mathbb{P}_n$ , leading to the following problem:

$$\underset{\boldsymbol{B} \in \mathbb{R}^{d \times d}, \operatorname{diag}(\boldsymbol{B}) = 0}{\operatorname{minimize}} \sup_{\mathbb{P}: \mathcal{D}_{c}(\mathbb{P}, \mathbb{P}_{n}) \leq \delta} \mathbb{E}_{\mathbb{P}} \left[ \left\| \boldsymbol{X} - \boldsymbol{B}^{\top} \boldsymbol{X} \right\|_{2}^{2} \right].$$
(3.2)

By solving (3.2), we try to find coefficients B that optimize the *worst* regression error of (3.1) among all probability distributions within a region around  $\mathbb{P}_n$ . This region  $\mathcal{U}_{\delta}(\mathbb{P}_n) := \{\mathbb{P} : \mathcal{D}_c(\mathbb{P}, \mathbb{P}_n) \leq \delta\}$ is called the uncertainty region with radius  $\delta$  (Blanchet et al., 2019). If  $\mathbb{P}^*$  indeed lies in this region, we are guaranteed that the loss under  $\mathbb{P}^*$  will be no larger than what is achieved in (3.2).

At first glance, (3.2) appears very difficult to solve, as it involves the supremum over an infinitedimensional space of probability measures. However, as we will show, this DRO problem can be relaxed as a finite-dimensional convex optimization problem. We also provide an ADMM algorithm that efficiently solves this convex optimization problem. Finally, we provide a simple recipe for choosing the appropriate radius  $\delta$  of the uncertainty region.

#### 3.1 Transforming the DRO problem to convex optimization

Blanchet et al. (2019) have presented the equivalence between distributionally robust linear regression with Wasserstein discrepancy of order p and  $L_q$  regularization, where  $1 \le p \le \infty$  and 1/p + 1/q = 1. Based on that, one might want to separate (3.2) into d distributionally robust linear regressions and then solve their equivalent  $L_2$ -regularized formulations. However, this is not correct because the variables in those linear regressions are coupled. We will instead analyze the program (3.2) as a whole. The theorem below, whose proof is deferred to the appendix, relaxes the DRO problem (3.2) to a convex optimization problem.

**Theorem 3.1.** With cost function  $c(u, w) = ||w - u||_2^2$ , we have:

$$\sup_{\mathbb{P}:\mathcal{D}_{c}(\mathbb{P},\mathbb{P}_{n})\leq\delta}\mathbb{E}_{\mathbb{P}}\left[\left\|X-\boldsymbol{B}^{\mathsf{T}}X\right\|_{2}^{2}\right]\leq\left(\frac{1}{\sqrt{n}}\left\|\boldsymbol{X}-\boldsymbol{X}\boldsymbol{B}\right\|_{F}+\sqrt{\delta}\left\|\boldsymbol{I}-\boldsymbol{B}\right\|_{2}\right)^{2},\qquad\forall\boldsymbol{B}\in\mathbb{R}^{d\times d},$$

where  $\|\cdot\|_2$  represents the spectral norm of a matrix.

Thanks to Theorem 3.1, the DRO problem (3.2) can be relaxed as:

$$\underset{\boldsymbol{B} \in \mathbb{R}^{d \times d}, \text{diag}(\boldsymbol{B}) = 0}{\text{minimize}} \left( \frac{1}{\sqrt{n}} \left\| \boldsymbol{X} - \boldsymbol{X} \boldsymbol{B} \right\|_{F} + \sqrt{\delta} \left\| \boldsymbol{I} - \boldsymbol{B} \right\|_{2} \right)^{2},$$

which is equivalent to a convex program

$$\min_{\boldsymbol{B} \in \mathbb{R}^{d \times d}, \operatorname{diag}(\boldsymbol{B}) = 0} \left\{ \frac{1}{\sqrt{n}} \left\| \boldsymbol{X} - \boldsymbol{X} \boldsymbol{B} \right\|_{F} + \sqrt{\delta} \left\| \boldsymbol{I} - \boldsymbol{B} \right\|_{2} \right\}.$$
(3.3)

Note that the right-hand-side of (3.3) is convex in **B**. This characteristic offers our result more generalizability beyond just nodewise regression. A good example is distributionally robust *regularized* nodewise regression, such as the  $L_1$ -regularized nodewise regression (2.6) for sparse subspace clustering. Specifically, the direct DRO formulation of (2.6) is

$$\min_{\boldsymbol{B} \in \mathbb{R}^{d \times d}, \operatorname{diag}(\boldsymbol{B}) = 0} \sup_{\mathbb{P}: \mathcal{D}_{c}(\mathbb{P}, \mathbb{P}_{n}) \leq \delta} \left\{ \mathbb{E}_{\mathbb{P}} \left[ \left\| X - \boldsymbol{B}^{\mathsf{T}} X \right\|_{2}^{2} \right] + \left\| \boldsymbol{B} \boldsymbol{\Lambda} \right\|_{1} \right\}.$$

According to Theorem 3.1, a convex relaxation is

$$\min_{\boldsymbol{B} \in \mathbb{R}^{d \times d}, \operatorname{diag}(\boldsymbol{B}) = 0} \left\{ \left( \frac{1}{\sqrt{n}} \left\| \boldsymbol{X} - \boldsymbol{X} \boldsymbol{B} \right\|_{F} + \sqrt{\delta} \left\| \boldsymbol{I} - \boldsymbol{B} \right\|_{2} \right)^{2} + \left\| \boldsymbol{B} \boldsymbol{\Lambda} \right\|_{1} \right\}.$$

In fact, we may replace the  $L_1$  penalty  $\|B\Lambda\|_1$  with any convex function and still get a convex relaxation.

The regularization weight parameter  $\delta$  is nothing but the diameter of the uncertainty region. In simple words, the distributionally robust nodewise regression with squared loss (3.2) is approximately equivalent to a spectral-norm-regularized nodewise regression with square-root loss, where the strength of the regularization is controlled by the radius of the uncertainty region  $\delta$ . The spectral-norm (also called the operator norm) is widely used in the machine learning literature to describe the generalizability of a model by measuring its vulnerability against adversarial attacks (see, e.g., Szegedy et al., 2014). The same intuition applies to our problem. Let  $l(x_t, \mathbf{B}) = ||x_t - \mathbf{B}^{\top} x_t||_2^2$  be the regression loss for a given parameter  $\mathbf{B}$  associated with an observation  $x_t$ . If the observation is modified with some perturbation  $\xi$ , then we have

$$\frac{\left|l(x_t + \xi, \boldsymbol{B}) - l(x_t, \boldsymbol{B})\right|}{\|\xi\|_2^2} = \frac{\left\|(\boldsymbol{I} - \boldsymbol{B})^\top \xi\right\|_2^2}{\|\xi\|_2^2} \le \|\boldsymbol{I} - \boldsymbol{B}\|_2^2.$$

This means that the magnitude of relative changes in the loss compared with the magnitude of the perturbation can be bounded by the spectral norm of I - B. This interpretation is consistent with the intuition that DRO minimizes the worst-case loss inside a plausible region.

#### 3.2 An ADMM algorithm

Program (3.3) is a convex optimization problem and can be solved by off-the-shelf optimizers. However, we find it prohibitively expensive in practice as soon as the dimension d reaches the hundreds. We propose an efficient algorithm based on the alternating direction method of multipliers (ADMM). For an overview of ADMM, see Boyd et al. (2010).

First we rewrite (3.3) as:

$$\min_{\boldsymbol{B}_1, \boldsymbol{B}_2 \in \mathbb{R}^{d \times d}} \left\{ \frac{1}{\sqrt{n}} \left\| \boldsymbol{X} - \boldsymbol{X} \boldsymbol{B}_1 \right\|_F + \sqrt{\delta} \left\| \boldsymbol{B}_2 \right\|_2 \right\}$$
s.t.  $\boldsymbol{B}_1 + \boldsymbol{B}_2 = \boldsymbol{I}, \quad \text{diag}(\boldsymbol{B}_1) = 0.$ 

We now describe the ADMM algorithm. Given an arbitrarily initialized  $B_2^0$  and  $\Lambda^0$ , we repeat the following steps: At iteration t, update:

$$\boldsymbol{B}_{1}^{t+1} \leftarrow \operatorname*{argmin}_{\operatorname{diag}(\boldsymbol{B})=0} \left\{ \frac{1}{\sqrt{n}} \left\| \boldsymbol{X} - \boldsymbol{X} \boldsymbol{B} \right\|_{F} + \frac{\rho}{2} \left\| \boldsymbol{B} + \boldsymbol{B}_{2}^{t} - \boldsymbol{I} + \boldsymbol{\Lambda}^{t} \right\|_{F}^{2} \right\}$$
(3.4)

where  $\rho$  is an exogenous penalty parameter, and  $B_2^0$  and  $\Lambda^0$  are initialized as zeros. The above is repeated until the magnitudes of the updates are smaller than a predetermined threshold.

Each of the two sub-problems (3.4) and (3.5) are easily solved. Problem (3.4) is strongly convex due to the quadratic penalty and thus can be solved by first-order algorithms. Problem (3.5) has a partially closed-form solution based on the singular decomposition of  $I - B_1^{t+1} - \Lambda^t$ , as stated in the following lemma.

Lemma 3.1. Consider the optimization problem:

$$\underset{\boldsymbol{B}\in\mathbb{R}^{m\times n}}{\operatorname{minimize}} \quad \|\boldsymbol{B}-\boldsymbol{C}\|_{F}^{2}+\lambda\,\|\boldsymbol{B}\|_{2}\,,\tag{3.6}$$

where  $C \in \mathbb{R}^{m \times n}$ . Let  $C = U\Sigma V^{\top}$  be the singular value decomposition of C, where  $U \in \mathbb{R}^{m \times r}$ ,  $V \in \mathbb{R}^{n \times r}$ , and  $\Sigma$  is an  $r \times r$  diagonal matrix whose diagonals are the singular values  $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r \geq 0$  of C. Define  $\sigma_{r+1} = 0$ . Then the optimal solution  $\widehat{B}$  can be expressed as  $\widehat{B} = U\widehat{S}V^{\top}$ , for some diagonal  $\widehat{S} \in \mathbb{R}^{r \times r}$  whose diagonal s satisfies for some  $k \in \{1, \ldots, r\}$ ,  $s = \left(\overbrace{t, \cdots, t}^{k \text{ terms}}, \sigma_{k+1}, \sigma_{k+2}, \cdots, \sigma_r\right)$ , where  $u \coloneqq \operatorname{argmin}_{\sigma_{k+1} \leq u \leq \sigma_k} \left\{ \sum_{j=1}^k (\sigma_j - u)^2 + \lambda u \right\}$ .

We defer the proof of Lemma 3.1 to the appendix. By virtue of the above lemma, we can easily find the solution to (3.5) by computing the singular decomposition of  $I - B_1^{t+1} - \Lambda^t = U\Sigma V^{\top}$ , and then comparing the losses  $\sum_{j=1}^k (\sigma_j - u)^2 + 2t\sqrt{\delta}/\rho$  among all  $k \in \{1, \ldots, d\}$ .

#### **3.3** Choice of $\delta$

In machine learning, the strength of the regularization, controlled by  $\delta$ , is usually determined exogenously or by cross-validation. However, since  $\delta$  is the radius of the uncertainty region in our setting, the choice of  $\delta$  should be informed by the degree of uncertainty in the data. Specifically, we determine a distributional uncertainty region in a way that it is just large enough so that the correct set of regression coefficients, which we would obtain if the true distribution were known, becomes a plausible choice with a sufficiently high confidence level. A simple, actionable recipe for choosing  $\delta$  is provided at the end of this subsection.

Define the covariance of a random matrix  $\boldsymbol{M} \in \mathbb{R}^{d \times d}$ , denoted by  $\operatorname{Cov}(\boldsymbol{M})$ , as a  $(d \times d) \times (d \times d)$  tensor, with  $\operatorname{Cov}(\boldsymbol{M})_{ij,kl} := \operatorname{Cov}(M_{ij}, M_{kl}), i, j, k, l \in [d]$ . Before describing our method for choosing  $\delta$ , we introduce the following assumptions:

Assumption 3.1. The time series  $\{X(t) \in \mathbb{R}^d : t \ge 0\}$  underlying the observations is a stationary, ergodic process satisfying  $\mathbb{E}_{\mathbb{P}^*}\left(\|X(t)\|_2^4\right) < \infty$  for each  $t \ge 0$ . Moreover, for each measurable function  $g: \mathbb{R}^d \to \mathbb{R}^{d \times d}$  such that  $\sum_{i,j} |g(x)_{ij}| \le c(1 + ||x||_2^2)$  for some c > 0, the limit

$$\Upsilon_g := \lim_{n \to \infty} \operatorname{Cov}_{\mathbb{P}^*} \left( n^{-1/2} \sum_{t=1}^n g(X(t)) \right) \in \mathbb{R}^{(d \times d) \times (d \times d)}$$

exists, and the central limit theorem holds:

$$n^{1/2}\left[\mathbb{E}_{\mathbb{P}_n}\left(g(X)\right) - \mathbb{E}_{\mathbb{P}^*}\left(g(X)\right)\right] \Rightarrow N(0,\Upsilon_g),$$

where " $\Rightarrow$ " denotes weak convergence as  $n \to \infty$  with fixed d, and  $N(0, \Upsilon_g)$  represents a random matrix  $\mathbf{Z}$  whose entries follow a normal distribution with  $\mathbb{E}[Z_{ij}] = 0$  and  $\operatorname{Cov}(Z_{ij}, Z_{kl}) = (\Upsilon_g)_{ij,kl}$ .

Assumption 3.2. The classical optimization problem (3.1) has a unique solution  $B^*$ .

**Assumption 3.3.** X(t) has a density for each  $t \ge 0$ .

Assumption 3.1 is standard for most time series models. Assumption 3.2 holds when the true underlying covariance matrix is invertible, which is true when no random variable is exactly a linear combination of other random variables. This condition is easily satisfied when, for example, each random variable is generated with an idiosyncratic noise.

In order to choose an appropriate  $\delta$ , we follow the idea behind the robust Wasserstein profile inference (RWPI) approach introduced in Blanchet et al. (2019). Intuitively, the uncertainty region  $\mathcal{U}_{\delta}(\mathbb{P}_n) := \{\mathbb{P} : \mathcal{D}_c(\mathbb{P}, \mathbb{P}_n) \leq \delta\}$  contains all the probability measures that are plausible variations of  $\mathbb{P}_n$  implied by the data. Let  $\mathbf{H} := \mathbf{I} - \mathbf{B}$  for simpler notation. We denote by  $\mathcal{Q}(\mathbb{P})$  the classical regression problem with  $\mathbb{P}$  being the underlying probability distribution:

$$\underset{\boldsymbol{H} \in \mathbb{R}^{d \times d}}{\text{minimize}} \quad \mathbb{E}_{\mathbb{P}} \left[ X^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} X \right], \quad \text{s.t.} \quad \text{diag}(\boldsymbol{H}) = 1.$$

Also, denote by  $H_{\mathbb{P}}$  a solution to  $\mathcal{Q}(\mathbb{P})$  and by  $\mathcal{H}_{\mathbb{P}}$  the set of all such solutions. According to Assumption 3.2, we have  $\mathcal{H}_{\mathbb{P}^*} = \{H^*\}$  for some  $H^* := I - B^*$ . Therefore, there exist unique Lagrange multipliers  $\lambda_1^*, \lambda_2^*, \ldots, \lambda_d^*$  such that

$$\mathbb{E}_{\mathbb{P}^*}\left[XX^{\top}\right]\boldsymbol{H}^* - \boldsymbol{\Lambda}^* = \boldsymbol{0}, \quad \operatorname{diag}(\boldsymbol{H}^*) = 1,$$

where  $\Lambda^*$  is the diagonal matrix with entries  $\lambda_1^*, \lambda_2^*, \ldots, \lambda_d^*$ .

We choose  $\delta > 0$  such that  $\mathcal{U}_{\delta}(\mathbb{P}_n)$  contains all probability distributions that are plausible variations of  $\mathbb{P}_n$ , any hence  $\mathbf{H}_{\mathbb{P}}$  with  $\mathbb{P} \in \mathcal{U}_{\delta}(\mathbb{P}_n)$  is a plausible estimate of  $\mathbf{H}^*$ . Thus, if we collect all such plausible estimates as the set:

$$\Lambda_{\delta}(\mathbb{P}_n) = \bigcup_{\mathbb{P} \in \mathcal{U}_{\delta}(\mathbb{P}_n)} \mathcal{H}_{\mathbb{P}},$$

then  $\Lambda_{\delta}(\mathbb{P}_n)$  is a natural confidence region for  $H^*$ . Therefore,  $\delta$  should be chosen as the smallest number  $\delta_n^*$  such that  $H^*$  falls in this region with a given confidence level. Namely,

$$\delta_n^* = \min\left\{\delta : \mathbb{P}^*\left(\boldsymbol{H}^* \in \Lambda_{\delta}(\mathbb{P}_n)\right) \ge 1 - \alpha\right\},\$$

where  $1 - \alpha$  is a user-defined confidence level (typically 95%).

In order to be able to compute  $\delta_n^*$ , we provide a simpler representation using an auxiliary function called the Robust Wasserstein Profile (RWP) function. First observe that any  $\boldsymbol{H} \in \Lambda_{\delta}(\mathbb{P}_n)$  if and only if there exist  $\mathbb{P} \in \mathcal{U}_{\delta}(\mathbb{P}_n)$  along with  $\lambda_1, \lambda_2, \ldots, \lambda_d \in (-\infty, \infty)$  and their corresponding diagonal matrix  $\boldsymbol{\Lambda}$  such that

$$\mathbb{E}_{\mathbb{P}}\left[XX^{\top}\right]\boldsymbol{H}-\boldsymbol{\Lambda}=\boldsymbol{0}, \quad \mathrm{diag}(\boldsymbol{H})=1.$$

By plugging the second equation into the first, we have

$$\lambda_i = -\left(\mathbb{E}_{\mathbb{P}}\left[XX^{\top}\right] \boldsymbol{H}\right)_{ii} - \mathbb{E}_{\mathbb{P}}\left[X_i^2\right](1-h_{ii}), \text{ for } i \in [d],$$

where  $h_{ii}$  is the *i*-th element of the *i*-th row of H. Then the system of equations that H needs to satisfy becomes:

$$1 - h_{ii} = 0$$
 and  $\left( \mathbb{E}_{\mathbb{P}} \left[ X X^{\top} \right] \boldsymbol{H} \right)_{ij} = 0$ , for  $i, j \in [d]$  and  $i \neq j$ .

Now we define the following RWP function

$$\mathcal{R}_{n}(\boldsymbol{H}) := \inf \left\{ \mathcal{D}_{c}(\mathbb{P}, \mathbb{P}_{n}) : 1 - h_{ii} = 0, \left( \mathbb{E}_{\mathbb{P}} \left[ X X^{\top} \right] \boldsymbol{H} \right)_{ij} = 0, \text{ for } i, j \in [d] \text{ and } i \neq j \right\}$$

for  $\boldsymbol{H} \in \mathbb{R}^{d \times d}$  where  $\mathcal{S}^{d \times d}_+$ . Then, we can rewrite  $\delta^*_n$  as:

$$\delta_n^* = \inf\left\{\delta : \mathbb{P}^*\left(\mathcal{R}_n(\boldsymbol{H}^*) \le \delta\right) \ge 1 - \alpha\right\}.$$
(3.7)

In other words,  $\delta_n^*$  is now the  $1 - \alpha$  quantile of the distribution of  $\mathcal{R}_n(\mathbf{H}^*)$ . If we can asymptotically approximate the distribution of  $\mathcal{R}_n(\mathbf{H}^*)$ ,  $\delta_n^*$  can then easily determined.

Before presenting the asymptotic distribution of  $\mathcal{R}_n(\mathbf{H}^*)$ , we first introduce the notation for asymptotic stochastic upper bound:

$$n\mathcal{R}_n(\boldsymbol{H}^*) \lesssim_D \mathcal{R},$$

which means that, for every continuous and bounded non-decreasing function  $f(\cdot)$ , we have

$$\limsup_{n \to \infty} \mathbb{E}\left[f\left(n\mathcal{R}_n(\boldsymbol{H}^*)\right)\right] \leq \mathbb{E}\left[f(\bar{\mathcal{R}})\right].$$

Similarly, we write  $\gtrsim_D$  for an asymptotic stochastic lower bound, namely

$$\liminf_{n \to \infty} \mathbb{E}\left[f\left(n\mathcal{R}_n(\boldsymbol{H}^*)\right)\right] \geq \mathbb{E}\left[f(\bar{\mathcal{R}})\right].$$

If both the stochastic upper and lower bounds hold for the same  $\overline{\mathcal{R}}$ , then  $n\mathcal{R}_n(H^*) \Rightarrow \overline{\mathcal{R}}$ .

Now let us state an asymptotic stochastic upper bound for  $n\mathcal{R}_n(H^*)$ .

**Theorem 3.2.** Under Assumptions 3.1 and 3.3, write  $\Sigma_* := \mathbb{E}_{\mathbb{P}^*}[XX^{\top}]$  and  $g(X) := XX^{\top}$ . If  $\Sigma_*$  is invertible, then

$$n\mathcal{R}_n(\boldsymbol{H}^*) \lesssim_D \bar{\mathcal{R}} := \sum_{i=1}^d \frac{1}{4} Z_{\cdot i}^{\top} \boldsymbol{\Sigma}_*^{-1} Z_{\cdot i}$$

where  $Z_{i}$  is the *i*-th column of  $\mathbf{Z} \sim N(0, \Upsilon_{g})$ .

The result of Theorem 3.2 involves  $\Sigma_*^{-1}$ . The true covariance matrix  $\Sigma_*$  can be estimated using the sample second-moment matrix  $\Sigma_n = \mathbb{E}_{\mathbb{P}_n} [XX^{\top}] = \frac{1}{n-1} \sum_{t=1}^n g(x_t)$ . However, estimating  $\Sigma_*^{-1}$  with  $\Sigma_n^{-1}$  is not possible when n < d. Even when n is moderately large but of the same order as d, the sample covariance matrix has been shown to be unreliable (e.g., Johnstone, 2001). Here, we apply a commonly used remedy in machine learning, i.e., estimating  $\Sigma_*^{-1}$  by only keeping

the diagonals of  $\Sigma_*$  when calculating its inverse; see e.g., Bickel and Levina, 2004. After this, we can obtain  $\delta_n^*$  as the  $1 - \alpha$  quantile of  $\overline{\mathcal{R}}/n$ , as long as we know the distribution of Z. We can draw samples from the distribution of Z and then numerically estimate the quantile of  $\overline{\mathcal{R}}$ . Z follows a normal distribution with a covariance matrix  $\Upsilon_g$ , which can be estimated using the sample covariances of observations of  $g(x_t) \in \mathbb{R}^{d \times d}$ ,  $t = 1, \ldots, n$ . We note that since Z is a random symmetric matrix in  $\mathbb{R}^{d \times d}$ , the covariances of its entries  $\Upsilon_g$  is a  $(d \times d) \times (d \times d)$  tensor. Nonetheless,  $(\Upsilon_g)_{ij,kl}$  which represents the covariance between  $Z_{ij}$  and  $Z_{kl}$  can be approximated by the sample covariance  $\frac{1}{n-1}\sum_{t=1}^n \left(g(x_t)_{ij} - \bar{g}_{ij}\right) \left(g(x_t)_{kl} - \bar{g}_{kl}\right)$ , where  $\bar{g}_{ij} := \frac{1}{n}\sum_{t=1}^n g(x_t)_{ij}$ . One should be mindful that applying this method is not always realistic in practice. First of all,  $\Upsilon_g$  has size  $d^4$ and can be difficult to fit in the RAM of a consumer computer (e.g., when d = 500,  $\Upsilon_g$  is roughly 250 GB in float 32 format). Further, it would require  $n > d^2$  observations for the sample covariance matrix to be positive definite. In many applications, the number of observations of n is on the same order as d, so the  $\Upsilon_q$  estimated this way could be highly unstable. An alternative method is to simply disregard the covariances assuming entries in Z are independent, and only calculate the diagonals. Recall that  $\Upsilon_g = \lim_{n \to \infty} \operatorname{Cov}_{\mathbb{P}^*} \left( n^{-1/2} \sum_{t=1}^n g(X(t)) \right)$ . Because  $g(x) := xx^{\top}$ ,  $\sum_{t=1}^n g(X(t))$  follows the Wishart distribution with degree of freedom n if we further assume that X is normal, and its variance is  $n \left[ \sigma_{ii} \sigma_{jj} + (\sigma_{ij})^2 \right]$ . Then the diagonals of  $\Upsilon_g$  can be computed<sup>1</sup>:  $(\Upsilon_g)_{ij,ij} = \sigma_{ii}\sigma_{jj} + (\sigma_{ij})^2$ . The independence also greatly simplifies the sampling of Z.

We now provide a simple recipe for choosing  $\delta$ :

- 1. Collect standardized data  $\{x_t\}_{t=1}^n, x_t \in \mathbb{R}^d$ .
- 2. Calculate second moments  $\{g(x_t) = x_t x_t^{\mathsf{T}}\}_{t=1}^n$ .
- 3. Use the sample second-moment matrix  $\Sigma_n = \mathbb{E}_{\mathbb{P}_n} [XX^{\top}] = \frac{1}{n-1} \sum_{t=1}^n g(x_t)$  to approximate  $\Sigma_*$ . Then estimate  $\Sigma_*^{-1}$  by only keeping the diagonals of  $\Sigma_*$ .
- 4. Calculate  $\Upsilon_g$  using either of the following methods:
  - (a)  $(\Upsilon_g)_{ij,kl} = \frac{1}{n-1} \sum_{t=1}^n \left( g(x_t)_{ij} \bar{g}_{ij} \right) \left( g(x_t)_{kl} \bar{g}_{kl} \right).$
  - (b)  $(\Upsilon_g)_{ij,ij} = \sigma_{ii}\sigma_{jj} + (\sigma_{ij})^2$ ,  $(\Upsilon_g)_{ij,kl} = 0$  if  $(k,l) \neq (i,j)$ .
- 5. Draw *M* samples  $\{\mathbf{Z}_m\}_{m=1}^M$  from the distribution  $N(0, \Upsilon_g)$  to numerically estimate the  $1 \alpha$  quantile of  $\overline{\mathcal{R}}/n$ . Apply Theorem 3.2 and (3.7) to set  $\delta = \delta_n^*$  to this quantile.

### 4 Simulation Experiments

We demonstrate our results through simulation. For an easier illustration, we take a simplified version of (2.2) as the underlying model, where we disregard the observable common factors  $F_C$  and reduce the number of unobservable common factors to 1:

$$X_i = \beta_H(i)F_H + F_{z(i)}^{\top}\beta_G(i) + \varepsilon_i, \quad \text{for } i = 1, \dots, d.$$
(4.1)

With given parameters  $n, d, K, \beta_H(i), d_k$ , and  $\operatorname{Var}(\varepsilon_i)$ , the samples X are generated as follows. First, the sizes of the clusters,  $\{m_k\}_{k=1}^K$ , are determined following the multinomial distribution

<sup>&</sup>lt;sup>1</sup>The off-diagonals can also be computed: http://personal.psu.edu/drh20/asymp/fall2002/lectures/ln08.pdf

with equal probabilities d/K. For example, the first  $m_1$  variables are marked as Cluster 1, then the next  $m_2$  Cluster 2. Then, a pool of  $\min(n, d)$  candidate group-specific factors are generated as i.i.d. standard normal vectors. The direction of each factor is uniformly sampled from the unit sphere in  $\mathbb{R}^n$ . From this pool of candidate factors,  $d_k$  group-specific factors are then randomly chosen for each cluster k. We note that two clusters may share one or more group-specific factors, as all clusters randomly pick factors from the same pool. Even if they do not, the two corresponding subspaces might not be orthogonal to each other since different group-specific factors can be correlated. Next, the factor loadings are determined. Loadings of variable i, represented by the  $d_{z(i)}$ -dimensional vector  $\beta_G(i)$ , are determined by sampling from the standard normal distribution.  $\beta_G(i)$  is then normalized so that  $\|\beta_G(i)\|_2^2 + \beta_H(i)^2 = 1$ . Then, a hidden global factor  $F_H$  is sampled from an n-dimensional standard normal distribution, and similarly  $\varepsilon_i$ 's are drawn independently from a normal distribution with given variance  $\operatorname{Var}(\varepsilon_i)$ . In the end, the samples for each random variable are standardized to have zero mean and unit variance.

In the following experiments, we compare our DRO nodewise regression subspace clustering method ("DRO") with the Lasso nodewise regression subspace clustering method ("Lasso") as described in (2.5), the ACC algorithm proposed by Tang et al. (2022), and the k-medoids algorithm (or Partitioning Around Medoids, "PAM", Kaufman and Rousseeuw, 1990). For DRO, the parameter  $\delta$  is determined following the recipe described at the end of Section 3.3, where method (b) is used to calculate  $\Upsilon_g$ , and M = 1000 samples of Z are generated to determine the quantile, for which we use a confidence level  $1 - \alpha = 0.95$ . For Lasso, we use a uniform parameter  $\lambda$  for all regressions and determine the value for  $\lambda$  using cross-validation by minimizing the validation regression error. For the two methods above, the regression coefficients B are obtained then symmetrized by calculating  $C := B_{abs}^{\top} + B_{abs}$ , where  $(B_{abs})_{ij} = |b_{ij}|$ . Clusters are then calculated using spectral clustering (see, e.g., von Luxburg, 2007) with C being the similarity matrix. For ACC, we use a slightly modified version where the dissimilarity measure is defined as

$$\operatorname{CORD}(i,j) := \min\left(\max_{l \neq i,j} \left| \rho_{il} - \rho_{jl} \right|, \max_{l \neq i,j} \left| \rho_{il} + \rho_{jl} \right| \right)$$

in order to accommodate both positive and negative factor loadings. We also fix the number of desired clusters, instead of letting the algorithm decide. For k-medoids, we use the distance measure  $1 - \rho^2$  with  $\rho$  being the sample correlation between two variables.

For DRO only, we have also tested two additional values for  $\alpha$ :  $\alpha = 0.1$  and  $\alpha = 0.01$ , corresponding to a 90% confidence level and a 99% confidence level, respectively. We find that the clustering results are nearly identical to  $\alpha = 0.05$ , which means that the algorithm is not sensitive to the confidence level  $1 - \alpha$ . Below we will only report the results of  $\alpha = 0.05$ , which corresponds to a 95% confidence level.

We create a total of K = 25 clusters among d = 500 variables, fixing  $\beta_H(i) = 0$  for all *i*, which means no hidden factor. The number of factors controlling each cluster *k* is randomly chosen from 1 to  $m_k - 1$ , where  $m_k$  is the number of variables in cluster *k*. We keep the noise level fixed for all variables at  $\operatorname{Var}(\varepsilon) = 0.1$  and generate n = 250 i.i.d. samples. Note that although our theoretical results for the DRO are stated when *n* grows to infinity with *d* fixed, we choose *n* to be much smaller than *d* to test the robustness of the DRO result. Figure 2a shows the true clustering structure among the d = 50 variables, and Figure 2b shows a heatmap of the sample correlation matrix. The blocks along the diagonal of Figure 2b are much more difficult to distinguish, likely due to the presence of multiple group-specific factors. Figures 3a and 3b show the *C* matrices from DRO and Lasso, respectively. We observe that the blocks along the diagonal are more prominent visually with both methods. DRO maintains more connections than Lasso in the C matrix, reflected in the light grey background off the diagonals. In contrast, Lasso leaves more blanks on the off-diagonals. On the diagonals, the blocks are also darker in DRO than Lasso. This is consistent with our intuition that DRO does not artificially pursue sparsity and thus has the advantage of keeping more true connections while weakening, instead of eliminating, irrelevant connections.



Figure 2: The heatmap of the true similarity matrix (a) and the sample correlation matrix (b), with  $\beta_H = 0$ ,  $\sigma_e^2 = 0.1$ , and  $d_k$  randomly chosen. Variables in the same cluster have similarity 1 and otherwise 0.



Figure 3: The heatmap of the C matrices for DRO (a) and Lasso (b), with  $\beta_H = 0$ ,  $\sigma_e^2 = 0.1$ , and  $d_k$  randomly chosen.

To further validate the result, we repeat the experiment on 10 random trials and examine the quality of the clusters by calculating the adjusted mutual information ("AMI", Vinh et al., 2010) between the obtained clusters and the ground truth. The AMI is a measure of similarity between two partitions; an AMI of 1 represents identical partitions, while uniformly random cluster assignments will have an AMI close to 0. The higher the AMI is for a partition compared with the ground truth, the more accurate the clustering results are. Table 1 shows the average AMI of each clustering method over the 10 random trials. Between the two subspace clustering methods, DRO achieves an average AMI of 0.96, followed by Lasso with an average AMI of 0.94. ACC and k-medoids underperform in this experiment, with average AMIs of 0.34 and 0.55, respectively. The under-performance of ACC and k-medoids is expected, because they are not tailored to the subspace clustering problem: the model underlying ACC assumes variables from the same cluster are generated around the same single factor, and similarly, k-medoids only seeks points that are spatially close to each other.

Table 1: Average AMI of different clustering methods compared with ground truth, over 10 different random trials.

	DRO	Lasso	ACC	k-medoids
Average AMI	0.96	0.94	0.34	0.55

We also test different noise levels by setting  $\sigma_e^2 = 0.1, 0.2, \ldots, 2.0$  and repeating the experiment on 10 random trials for each value of  $\sigma_e^2$ . These values of  $\sigma_e^2$  represent signal-to-noise ratios from 10:1 to 1:2. The average AMI of each method is shown in Figure 4. Overall, the average AMI decays for all methods as the level of noise increases. The DRO subspace clustering methods perform similarly with Lasso, and both consistently outperform ACC and k-medoids in this experiment.



Figure 4: Comparison of average AMI between different clustering methods, with  $\beta_H = 0$  and  $d_k$  randomly chosen.

Finally, we examine the performance of the clustering algorithms in the presence of a global hidden factor and with heterogeneous factor loadings and noise levels. In these experiments, DRO again leads the cohort overall, and we include the detailed results and analysis in Appendix B.1.

### 5 Empirical Experiments on Financial Data

#### 5.1 Overview

In this section, we test subspace clustering algorithms by applying them to clustering financial time series data. Our task is to cluster the stocks in the S&P 500 universe using historical returns, and based on these clusters, construct stock portfolios. More specifically, we pick one representative stock from each cluster, and construct an optimized portfolio using these representative stocks. The underlying rationale is that by identifying stocks capable of representing others, one can create portfolios with a small number of stocks compared to the size of the full universe, yet still achieve a sufficient level of diversification. For more detailed discussion on clustering and diversification, please see Tang et al. (2022).

#### 5.2 Data preparation

We take the constituents of the S&P 500 as the universe. The data is obtained from Compustat through Wharton Research Data Services (WRDS), which consists of

- the daily closing prices of the constituents;
- the historical constituents data; and
- the daily closing S&P 500 total return index with dividends reinvested,

between January 1996 and January 2020.

We apply clustering, stock selection, portfolio optimization, and backtesting for the period between February 2001 and January 2020. Partitions and portfolios are calculated on the first trading day of each February, starting with February 2001, in the then S&P 500 constituent stock universe. Specifically, at the end of the first trading day of each February, we choose the stocks in the S&P 500 Index according to the historical constituent data. Of all the current constituents, we discard stocks with less than 5 years of history and those with more than 5% missing data in the past n = 500 days. If the same company has multiple classes of stocks in the S&P 500 Index (e.g. Alphabet Inc's GOOG and GOOGL), we only keep the class with the longest history. The numbers of eligible stocks that remained after the above filtering range between 468 and 487 over the backtesting period. For these eligible stocks, any missing prices are linearly interpolated using the previous and subsequent prices. Then, partitions are estimated based on the daily returns of the past n = 500 trading days. A smaller set of stocks are then selected, and portfolios are constructed with stock weights optimized. These steps are described in detail in the following subsections.

#### 5.3 Clustering

We first describe the clustering methods to be tested:

- 1. **DRO-ACC:** First create  $K_1$  clusters using the DRO subspace clustering algorithm, then further split each cluster into  $K_2$  sub-clusters using the ACC algorithm.
- 2. ACC Create  $K_1 \times K_2$  clusters using the ACC algorithm.
- 3. k-medoids Create  $K_1 \times K_2$  clusters using the k-medoids algorithm.

Method 1 is a combination of the DRO subspace clustering and the ACC clustering in a hierarchical fashion. We believe clusters generated by this approach is more suitable for stock selection, compared to, for instance, clusters generated directly by subspace clustering algorithms. This is because stocks in the same low-dimensional subspace may still be quite different from each other (vectors in the same subspace can point to rather different directions), and it may be difficult to use a single stock to represent a whole cluster. With the DRO clustering at the higher level followed by the ACC at the lower level, the former breaks down the universe into stocks driven by *groups* of factors, and the latter then easily finds stocks most closely associated with each single factor. We parsimoniously choose  $K_1 = K_2$  since we have no prior knowledge of how many subspace there should be vs. the dimensions of these subspaces. We compare this method with Method 2, which directly applies the ACC algorithm on the full universe. The ACC algorithm works very well in this task as demonstrated in Tang et al. (2022). We also include k-medoids as another benchmark.

At the end of the first trading day of each February, the above three clustering methods are applied to daily log returns in the backward 500-trading-day window for valid constituent stocks as described in Section 5.2. For each clustering method, we create a total of 19 sets of clusters, one for each year. The choice of parameters for the clustering algorithms is the same as described in Section 4.

#### 5.4 Portfolio construction

Under each clustering method described earlier, once clusters have been constructed, we select the stock with the lowest volatility from each cluster and then form a portfolio on the resulting smaller set of stocks. Here the volatility is measured by the sample variance of daily returns in the past n = 500 trading days, the same window used for clustering. From a practical and empirical perspective, the reason why we choose low volatility as the criterion is twofold. First, volatility as a criterion does not involve the estimation of the mean returns. All clustering algorithms tested avoid using the stocks' mean returns. It would then be inconsistent if we selected stocks from the clusters based on return-related criteria, e.g., mean return or Sharpe ratio. More importantly, the estimation of mean returns is well known to be often inaccurate (the "mean-blur" problem; see, e.g., Merton, 1980), rendering return-related criteria unreliable of indicating future performance. The second reason is that stocks with low volatility have been observed to outperform the benchmarks over time, which is contrary to CAPM and is documented as the "low-risk anomaly" (e.g., Zaremba and Shemer, 2017). We only choose one stock with the lowest volatility from each cluster, yielding the same number of stocks as clusters for each clustering method every time we update the portfolio.

Once a set of stocks is determined by the above procedure, we construct portfolios using the minimum variance allocation strategy to determine the weights of the stocks. The minimum variance allocation strategy is similar to Markowitz's mean-variance optimization but without the expected return constraint:

$$\begin{array}{ll} \min & \boldsymbol{w}^{\top}\boldsymbol{\Sigma}\boldsymbol{w} \\ \text{s.t.} & \boldsymbol{w}^{\top}\boldsymbol{1} = \boldsymbol{1}, \quad \boldsymbol{w} \geq \boldsymbol{0}. \end{array}$$

We choose the minimum variance allocation because it also does not involve the estimation of the mean return. Similar to why we use low volatility as a criterion to select stocks from the clusters, we aim to keep the experiment consistent by avoiding the estimation of the mean returns throughout the experiment.

The portfolios are updated annually. At each portfolio update, a new set of stocks are selected according to the clustering result. Their allocations are calculated using all daily returns in the past 500 trading days, starting with the first day when all stocks are available. The positions are then held until the first trading day of the following update. Any dividends are immediately reinvested in the same stock. We assume no transaction cost for simplicity. As a benchmark, we take the S&P 500 ETF (NYSE ticker: SPY), which is the largest ETF in the world and designed to track the S&P 500 Index.

#### 5.5 Results and analysis

Below we present the results when  $K_1 = K_2 = 6$ . At each update in February,  $6 \times 6 = 36$  clusters are recovered, each of which contributes one stock in the portfolio. While this is a reasonable number of stocks to have in a portfolio, we also present results of  $K_1 = K_2 = 3, 4, 5$  in Appendix B.2. In the results below, we update the portfolios once every year after the first trading day in February, when we re-do stock selection and re-compute allocation. Figure 5 shows the cumulative performance of these portfolios, and Table 2 reports the performance of the portfolios based on metrics commonly used in the wealth management industry. The DRO-ACC portfolio outperforms the others significantly in many important metrics, including Sharpe, Sortino, and Calmar ratios, annualized return, maximum drawdown, and recovery time, while it performs similarly to the best performers in other metrics.



Figure 5: Cumulative performance of portfolios; DRO-ACC creates 6x6 clusters, ACC and k-medoids create 36 clusters

To examine the compositions of the clusters, we compare the clusters with sectors defined by the Global Industry Classification Standard (GICS)<sup>2</sup>. Figure 6 shows the clustering results obtained by the DRO-ACC clustering method on Feb 1st, 2019, date of the last portfolio update in our experiment. The clusters are first ordered by the 6 major clusters from DRO, and then by size within each major cluster. In other words, Clusters 1 to 6 are from the first DRO major cluster, Clusters 7 to 12 are from the second, and so on.

One observation that immediately stands out is the similarity to GICS sectors in the 2nd through 6th DRO major clusters (noting the colors starting from Cluster 7). Each of the 2nd through 5th major clusters covers a different sector and often includes most companies in that sector. The last DRO cluster (Clusters 31-36) includes two sectors, namely, Consumer Discretionary and Consumer Staples, which are closely related to each other. In comparison, clusters within the first DRO cluster tend to be larger, especially Cluster 1. They also include companies from many different sectors, such as Communication Services, Consumers Discretionary, Industrials, and Information Technology. Intuitively, these sectors appear to be more closely related to the notion of the "day-to-day" economy, than some sectors represented by the other major clusters, like Real Estate, Energy, Utilities, and Financials. The reason why the sectors in the DRO clusters 2 through 6 (Clusters 7-36 in Figure 6) stand out is likely because they are the most distinguishable sectors from the rest

<sup>&</sup>lt;sup>2</sup>Available at https://www.msci.com/gics

Table 2: Performance metrics of the minimum variance portfolios; DRO-ACC creates 6x6 clusters, ACC and k-medoids create 36 clusters

	DRO-ACC	ACC	k-medoids	SPY
Ending VAMI	10694.3	5830.69	7701.07	3442.53
Max Drawdown	27.72%	36.24%	31.87%	55.25%
Peak-To-Valley	2008-09-08 -	2007-12-10 -	2007-12-10 -	2007-10-09 -
	2009-03-09	2009-03-09	2009-03-09	2009-03-09
Recovery	194 Days	446 Days	250 Days	774 Days
Sharpe Ratio	1.05	0.77	0.9	0.36
Sortino Ratio	1.73	1.25	1.46	0.57
Calmar Ratio	0.48	0.27	0.36	0.12
Ann. Volatility	12.73%	12.59%	12.69%	18.63%
Ann. Downside Volatility	7.71%	7.77%	7.81%	11.83%
Correlation	0.78	0.8	0.81	1.0
Beta	0.53	0.54	0.55	1.0
Annualized Return	13.32%	9.75%	11.38%	6.74%
Positive Periods	2608(54.63%)	2611(54.69%)	2595(54.36%)	2600(54.46%)
Negative Periods	2166~(45.37%)	2163(45.31%)	2179(45.64%)	2174(45.54%)

of the market. DRO being able to single them out in the first stage of clustering guarantees that a sufficient number of stocks are selected from each of these distinguishable sectors, which may facilitate diversification and lead to the good performance of the DRO-ACC portfolio.



Figure 6: DRO-ACC clustering results on 2019-02-01 compared with GICS sectors

### 6 Conclusion

In this paper, we propose a distributionally robust nodewise regression method and apply it to variable clustering. We derive a convenient convex relaxation of the problem. The uncertainty level in the distributionally robust regression can be chosen in a data-driven way. Compared with the popular sparse subspace clustering that uses nodewise Lasso, our method is tuning-free and has a naturally interpretable regularization. The only exogenous parameter of the algorithm is a confidence level  $1 - \alpha$ , which the algorithm is very insensitive to, as results are nearly identical with  $\alpha = 0.01, 0.05, 0.1$ . Simulation experiments show that our subspace clustering method outperforms the sparse subspace clustering (SSC), the newly developed ACC clustering algorithm, and the classic k-medoids clustering algorithm. We also apply the clustering method to financial time series data for asset selection and obtain promising results in portfolio backtesting.

Our work can be extended in the following directions. First, we have only worked with distributionally robust nodewise regression with quadratic cost  $c(u, w) = ||w - u||_2^2$ , which leads to Wasserstein distance of order 2 as the discrepancy measure. A similar analysis may be attempted for other cost functions and thus other discrepancy measures. For example, if one could find a cost function that connects distributionally robust nodewise regression to the LASSO nodewise regression, we would obtain an easy way to choose the parameter for the  $L_1$  regularization term. In addition, existing results for linear regression (Blanchet et al., 2019) are exact reformulations of the distributionally robust problems as non-robust optimization problems, while our result is a *relaxation*. Exact formulations are difficult with the Wasserstein distance of order 2 as it appears in our analysis but could be possible for other distance measures. Finally, the regularization parameter  $\delta$  is determined using a quantile in the asymptotic distribution of  $\mathcal{R}_n(\mathbf{H}^*)$ , which could be different from the actual distribution when the data is limited. It remains an interesting question whether it is possible to derive the actual distribution of  $\mathcal{R}_n(\mathbf{H}^*)$  with limited data.

## Acknowledgements

Wang gratefully acknowledges financial support through an NSF grant DMS-2210907 and a start-up grant at Columbia University. Zhou gratefully acknowledges financial support through a start-up grant and the Nie Center for Intelligent Asset Management at Columbia University.

## Appendix A Proof of Theorems

#### A.1 Proof of Lemma 3.1

We know that the optimal solution  $\hat{B}$  to (3.6) can be expressed as  $\hat{B} = U\hat{S}V^{\top}$ , for some  $\hat{S} \in \mathbb{R}^{r \times r}$ . Then we have:

$$\left\|\widehat{B} - C\right\|_{F}^{2} + \lambda \left\|\widehat{B}\right\|_{2} = \left\|\widehat{S} - \Sigma\right\|_{F}^{2} + \lambda \left\|\widehat{S}\right\|_{2}.$$

To proceed, we need to use the following lemma.

Lemma A.1 (Pinching Inequality (e.g., Bani-Domi and Kittaneh, 2008)). If a matrix A has a block form:

$$m{A} = egin{bmatrix} m{A}_{11} & m{A}_{12} & \dots \ m{A}_{21} & m{A}_{22} & \dots \ dots & dots & \ddots \ dots & dots & \ddots \ \end{bmatrix},$$

then for any weakly unitary invariant norm  $\|\cdot\|$ ,

$$\|\boldsymbol{A}\| \geq \left\| \begin{bmatrix} \boldsymbol{A}_{11} & \boldsymbol{0} & \dots \\ \boldsymbol{0} & \boldsymbol{A}_{22} & \dots \\ \vdots & \vdots & \ddots \end{bmatrix} \right\|.$$

Let  $\tilde{\mathbf{S}} := \text{diag}(\hat{s}_{11}, \hat{s}_{22}, \dots, \hat{s}_{rr})$ . Because  $\Sigma$  is diagonal and the Frobenius norm and the spectral norm are both weakly unitary invariant, by the pinching inequality, we have

$$\left\| \widetilde{\boldsymbol{S}} - \boldsymbol{\Sigma} \right\|_{F} \leq \left\| \widehat{\boldsymbol{S}} - \boldsymbol{\Sigma} \right\|_{F}, \quad \left\| \widetilde{\boldsymbol{S}} \right\|_{2} \leq \left\| \widehat{\boldsymbol{S}} \right\|_{2},$$

Because of the optimality of  $\hat{S}$ , we know that  $\tilde{S} = \hat{S}$  and thus  $\hat{S}$  is diagonal. Hence, (3.6) is equivalent to

$$\min_{\boldsymbol{S} \in \mathbb{R}^{r \times r}, \ \boldsymbol{S} \text{ is diagonal}} \left\{ \left\| \boldsymbol{S} - \boldsymbol{\Sigma} \right\|_{F}^{2} + \lambda \left\| \boldsymbol{S} \right\|_{2} \right\},\$$

which is just

$$\min_{S \in \mathbb{R}^r} \left\{ \sum_{j=1}^r (s_j - \sigma_j)^2 + \lambda \max_j |s_j| \right\},\$$

where  $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r \geq 0$  are singular values of C, and  $s_1, \ldots, s_r$  are diagonal entries of S. This can be further transformed to

$$\min_{S,t} \left\{ \sum_{j=1}^{r} (s_j - \sigma_j)^2 + \lambda t \right\}$$
  
s.t.  $0 \le s_1 \le s_2 \le \cdots, \le s_r \le t$ 

which is now easy to solve by noticing that for  $\sigma_{k+1} \leq t \leq \sigma_k$ ,  $k = 1, \ldots, r$ , the optimal s is

$$s = (\overbrace{t, \cdots, t}^{k \text{ terms}}, \sigma_{k+1}, \sigma_{k+2}, \cdots, \sigma_r),$$

and the loss for such t is  $\sum_{j=1}^{k} (\sigma_j - t)^2 + \lambda t$ , so the optimal t is  $(2\sum_{j=1}^{k} \sigma_j + \lambda)/k$ .

#### A.2 Proof of Theorem 3.1

We argue by strong duality using a lemma of Theorem 1 in Blanchet and Murthy (2019).

**Lemma A.2.** For  $\gamma \ge 0$  and loss functions l(x, B) that are upper semi-continuous in x for each B, define:

$$\phi_{\gamma}(x_t; \mathbf{B}) := \sup_{u \in \mathbb{R}^d} \{ l(u; \mathbf{B}) - \gamma c(u, x_t) \}.$$
(A.1)

Then,

$$\sup_{\mathbb{P}:\mathcal{D}_{c}(\mathbb{P},\mathbb{P}_{n})\leq\delta}\mathbb{E}_{\mathbb{P}}[l(X;\boldsymbol{B})] = \min_{\gamma\geq0}\left\{\gamma\delta + \frac{1}{n}\sum_{t=1}^{n}\phi_{\gamma}(x_{t};\boldsymbol{B})\right\}.$$
(A.2)

Recall that our loss function is the total squared error:  $l(X, \mathbf{B}) = ||X - \mathbf{B}^{\top}X||_2^2 = X^{\top}\mathbf{H}\mathbf{H}^{\top}X$ where  $\mathbf{H} := \mathbf{I} - \mathbf{B}$ , and the cost function is  $c(u, w) = ||w - u||_2^2$ . Using Lemma A.2, we can reduce the inner supremum of (3.2) to

$$\sup_{\mathbb{P}:\mathcal{D}_c(\mathbb{P},\mathbb{P}_n)\leq\delta} \mathbb{E}_{\mathbb{P}}[l(X;\boldsymbol{B})] = \min_{\gamma\geq0} \{\gamma\delta + \frac{1}{n}\sum_{t=1}^n \phi_{\gamma}(x_t;\boldsymbol{B})\},\tag{A.3}$$

where

$$\phi_{\gamma}(x_t; \boldsymbol{B}) := \sup_{u \in \mathbb{R}^d} \{ l(u; \boldsymbol{B}) - \gamma c(u, x_t) \}$$
$$= \sup_{u \in \mathbb{R}^d} \{ u^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} u - \gamma \| x_t - u \|_2^2 \}$$

Rewriting  $\Delta := u - x_t$ , we have:

$$\begin{split} \phi_{\gamma}(x_t; \boldsymbol{B}) &:= \sup_{\Delta \in \mathbb{R}^d} \{ (\Delta + x_t)^\top \boldsymbol{H} \boldsymbol{H}^\top (\Delta + x_t) - \gamma \, \|\Delta\|_2^2 \} \\ &= x_t^\top \boldsymbol{H} \boldsymbol{H}^\top x_t + \sup_{\Delta \in \mathbb{R}^d} \{ \Delta^\top \boldsymbol{H} \boldsymbol{H}^\top \Delta + 2x_t^\top \boldsymbol{H} \boldsymbol{H}^\top \Delta - \gamma \, \|\Delta\|_2^2 \} \\ &= x_t^\top \boldsymbol{H} \boldsymbol{H}^\top x_t + \sup_{\Delta \in \mathbb{R}^d} \{ -\Delta^\top (\gamma \boldsymbol{I} - \boldsymbol{H} \boldsymbol{H}^\top) \Delta + 2x_t^\top \boldsymbol{H} \boldsymbol{H}^\top \Delta \} \end{split}$$

Observe that inside the supremum is a quadratic function of  $\Delta$ , so the supreme is only finite if the quadratic function is concave. This means that  $\gamma \mathbf{I} - \mathbf{H}\mathbf{H}^{\top}$  needs to be positive definite, and thus invertible, which requires that  $\gamma > \lambda_1$  where  $\lambda_1$  is the largest eigenvalue of  $(\mathbf{H}\mathbf{H}^{\top})$ . Then, according to the first order condition, the supremum is achieved when  $(\gamma \mathbf{I} - \mathbf{H}\mathbf{H}^{\top})\Delta = \mathbf{H}\mathbf{H}^{\top}x_t$ , i.e.,  $\Delta = (\gamma \mathbf{I} - \mathbf{H}\mathbf{H}^{\top})^{-1}\mathbf{H}\mathbf{H}^{\top}x_t$ . Plugging in the value for  $\Delta$ , we have

$$\phi_{\gamma}(x_t; \boldsymbol{B}) = x_t^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} x_t + x_t^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} (\gamma \boldsymbol{I} - \boldsymbol{H} \boldsymbol{H}^{\top})^{-1} \boldsymbol{H} \boldsymbol{H}^{\top} x_t.$$

Through eigendecomposition, we can write

$$(\gamma \boldsymbol{I} - \boldsymbol{H} \boldsymbol{H}^{\mathsf{T}})^{-1} = \boldsymbol{Q} (\gamma \boldsymbol{I} - \boldsymbol{\Lambda})^{-1} \boldsymbol{Q}^{\mathsf{T}},$$

where  $\Lambda$  is the diagonal matrix with  $\Lambda_{ii} = \lambda_i$  being the *i*-th largest eigenvalue of  $HH^{\top}$ , and  $Q = [Q_1 \ Q_2 \ \dots \ Q_d]$  is the matrix whose *i*-th column is the eigenvector  $Q_i$  of  $HH^{\top}$  corresponding to the eigenvalue  $\lambda_i$ . Then

$$\phi_{\gamma}(x_t; \boldsymbol{B}) = x_t^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} x_t + x_t^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} \boldsymbol{Q} (\gamma \boldsymbol{I} - \boldsymbol{\Lambda})^{-1} \boldsymbol{Q}^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} x_t$$

$$= x_t^{\mathsf{T}} \boldsymbol{Q} \boldsymbol{\Lambda} \boldsymbol{Q}^{\mathsf{T}} x_t + x_t^{\mathsf{T}} \boldsymbol{Q} \boldsymbol{\Lambda} (\gamma \boldsymbol{I} - \boldsymbol{\Lambda})^{-1} \boldsymbol{\Lambda} \boldsymbol{Q}^{\mathsf{T}} x_t$$
$$= x_t^{\mathsf{T}} \boldsymbol{Q} \begin{bmatrix} \frac{\gamma \lambda_1}{\gamma - \lambda_1} & 0 & \dots & 0\\ 0 & \frac{\gamma \lambda_2}{\gamma - \lambda_2} & \dots & 0\\ \vdots & \vdots & \ddots & \vdots\\ 0 & 0 & \dots & \frac{\gamma \lambda_d}{\gamma - \lambda_d} \end{bmatrix} \boldsymbol{Q}^{\mathsf{T}} x_t$$
$$= \sum_{i=1}^d \frac{\gamma \lambda_i}{\gamma - \lambda_i} (Q_i^{\mathsf{T}} x_t)^2.$$

Now the minimum in (A.3) becomes:

$$\min_{\gamma > \lambda_1} \left\{ \gamma \delta + \frac{1}{n} \sum_{i=1}^d \frac{\gamma \lambda_i \sum_{t=1}^n (Q_i^{\mathsf{T}} x_t)^2}{\gamma - \lambda_i} \right\}.$$
 (A.4)

Inside the minimum of (A.4) is a convex function of  $\gamma$  on  $\gamma > \lambda_1$  that tends to infinity as  $\gamma \to \infty$  or  $\gamma \to \lambda_1$ . The optimal  $\gamma$  should follow the first-order condition

$$\delta - \frac{1}{n} \sum_{i=1}^{d} \left( \frac{\lambda_i^2}{(\gamma - \lambda_i)^2} \sum_{t=1}^{n} (Q_i^{\mathsf{T}} x_t)^2 \right) = 0.$$
 (A.5)

It is easy to see that this equation has a solution on  $(\lambda_1, \infty)$ , because the left hand side goes to  $-\infty$  as  $\gamma$  approaches  $\lambda_1$ , and goes to  $\delta > 0$  as  $\gamma$  approaches  $\infty$ . However, analytically solving this equation involves the 2*d*-th order product  $\prod_{i=1}^{d} (\gamma - \lambda_i)^2$  and is therefore difficult. So we introduce an approximation by replacing  $(\gamma - \lambda_i)^2$  with  $(\gamma - \lambda_1)^2$  in the denominator and replacing one of the  $\lambda_i$ 's in the numerator with  $\lambda_1^3$ . In other words, we try to find the  $\gamma$  that satisfies:

$$\delta - \frac{1}{n} \sum_{i=1}^{d} \left( \frac{\lambda_1 \lambda_i}{(\gamma - \lambda_1)^2} \sum_{t=1}^{n} (Q_i^{\mathsf{T}} x_t)^2 \right) = 0, \tag{A.6}$$

which yields

$$\gamma = \lambda_1 + \frac{1}{\sqrt{\delta}}\sqrt{\lambda_1}\sqrt{\frac{1}{n}\sum_{t=1}^n\sum_{i=1}^d\lambda_i(Q_i^{\mathsf{T}}x_t)^2} = \lambda_1 + \frac{1}{\sqrt{\delta}}\sqrt{\lambda_1}\sqrt{\frac{1}{n}\sum_{t=1}^nx_t^{\mathsf{T}}\boldsymbol{H}\boldsymbol{H}^{\mathsf{T}}x_t}$$

Using this value for  $\gamma$ , we obtain an upper bound of (A.4):

$$\min_{\gamma > \lambda_1} \left\{ \gamma \delta + \frac{1}{n} \sum_{i=1}^d \frac{\gamma \lambda_i \sum_{t=1}^n (Q_i^\top x_t)^2}{\gamma - \lambda_i} \right\}$$

<sup>&</sup>lt;sup>3</sup>The informed reader might have noticed that the rest of the proof still follows if, instead of replacing with  $\lambda_1$ , we replace  $\lambda_i$  with any number larger than or equal to  $\lambda_1$ . We choose  $\lambda_1$  because it offers the tightest approximation of this simple form.

$$\begin{split} &\leq \min_{\gamma>\lambda_{1}} \left\{ \gamma\delta + \frac{1}{n} \sum_{i=1}^{d} \frac{\gamma\lambda_{i} \sum_{t=1}^{n} (Q_{i}^{\top}x_{t})^{2}}{\gamma - \lambda_{1}} \right\} \\ &\leq \min_{\gamma>\lambda_{1}} \left\{ \gamma\delta + \frac{\gamma\frac{1}{n} \sum_{t=1}^{n} x_{t}^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} x_{t}}{\gamma - \lambda_{1}} \right\} \\ &= \lambda_{1}\delta + \sqrt{\delta}\sqrt{\lambda_{1}} \sqrt{\frac{1}{n} \sum_{t=1}^{n} x_{t}^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} x_{t}} \\ &+ \frac{\left(\lambda_{1} + \frac{\sqrt{\lambda_{1}}}{\sqrt{\delta}} \sqrt{\frac{1}{n} \sum_{t=1}^{n} x_{t}^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} x_{t}}\right) \left(\frac{1}{n} \sum_{t=1}^{n} x_{t}^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} x_{t}\right)}{\frac{\sqrt{\lambda_{1}}}{\sqrt{\delta}} \sqrt{\frac{1}{n} \sum_{t=1}^{n} x_{t}^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} x_{t}}} \\ &= \lambda_{1}\delta + 2\sqrt{\delta\lambda_{1}} \frac{1}{n} \sum_{t=1}^{n} x_{t}^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} x_{t} + \frac{1}{n} \sum_{t=1}^{n} x_{t}^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} x_{t}}{\left(\sqrt{\frac{1}{n} \sum_{t=1}^{n} x_{t}^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} x_{t}} + \sqrt{\delta}\sqrt{\lambda_{1}}\right)} \\ &= \left(\sqrt{\frac{1}{n} \sum_{t=1}^{n} x_{t}^{\top} \boldsymbol{H} \boldsymbol{H}^{\top} x_{t}} + \sqrt{\delta} \|\boldsymbol{H}\|_{2})\right)^{2} \end{split}$$

where  $\|\boldsymbol{H}\|_2$  is the spectral norm of  $\boldsymbol{H}$  and is equal to its largest singular value  $\sqrt{\lambda_1}$ .

#### 

## A.3 Proof of Theorem 3.2

By Proposition 3 in Blanchet et al. (2019), we have for any  $\boldsymbol{H} \in \mathbb{R}^{d \times d}$  and  $\operatorname{diag}(\boldsymbol{H}) = 1$ ,

$$\mathcal{R}_{n}(\boldsymbol{H}) = \sup_{\boldsymbol{\Lambda} \in \mathbb{R}^{d \times d}} \left\{ -\mathbb{E}_{\mathbb{P}_{n}} \left[ \sup_{u \in \mathbb{R}^{d}} \left\{ \sum_{i,j \in [d], i \neq j} \lambda_{ij} \left( uu^{\top} \boldsymbol{H} \right)_{ij} - \|u - X\|_{2}^{2} \right\} \right] \right\}$$
$$= \sup_{\boldsymbol{\Lambda} \in \mathbb{R}^{d \times d}: \operatorname{diag}(\boldsymbol{\Lambda}) = 0} \left\{ -\mathbb{E}_{\mathbb{P}_{n}} \left[ \sup_{u \in \mathbb{R}^{d}} \left\{ \operatorname{tr}(\boldsymbol{\Lambda}^{\top} uu^{\top} \boldsymbol{H}) - \|u - X\|_{2}^{2} \right\} \right] \right\}$$

Define  $h(X, \mathbf{H}) := X X^{\top} \mathbf{H}$ , then observe that the inner-most supremum

$$\sup_{u \in \mathbb{R}^d} \left\{ \operatorname{tr}(\mathbf{\Lambda}^\top u u^\top \mathbf{H}) - \|u - X\|_2^2 \right\}$$
  
= 
$$\sup_{\Delta \in \mathbb{R}^d} \left\{ \operatorname{tr}(\mathbf{\Lambda}^\top h(X + \Delta, \mathbf{H})) - \|\Delta\|_2^2 \right\}$$
  
= 
$$\sup_{\Delta \in \mathbb{R}^d} \left\{ \operatorname{tr}(\mathbf{\Lambda}^\top [h(X + \Delta, \mathbf{H}) - h(X, \mathbf{H})]) - \|\Delta\|_2^2 \right\} + \operatorname{tr}(\mathbf{\Lambda}^\top h(X, \mathbf{H})).$$

We can write

$$\operatorname{tr}(\boldsymbol{\Lambda}^{\top}[h(X+\Delta,\boldsymbol{H})-h(X,\boldsymbol{H})]) = \int_0^1 \frac{d}{dt} \operatorname{tr}(\boldsymbol{\Lambda}^{\top}h(X+t\Delta,\boldsymbol{H})) dt.$$

Calculating the derivative, we have

$$\frac{d}{dt}\operatorname{tr}(\boldsymbol{\Lambda}^{\top}h(X+t\Delta,\boldsymbol{H})) = 2\operatorname{tr}(\boldsymbol{H}\boldsymbol{\Lambda}^{\top}(X+t\Delta)\Delta^{\top})$$
$$= 2\operatorname{tr}(\boldsymbol{H}\boldsymbol{\Lambda}^{\top}X\Delta^{\top}) + 2t\Delta^{\top}\boldsymbol{H}\boldsymbol{\Lambda}^{\top}\Delta,$$

which is linear in t. So we deduce

$$\mathcal{R}_{n}(\boldsymbol{H}) = \sup_{\boldsymbol{\Lambda} \in \mathbb{R}^{d \times d} : \operatorname{diag}(\boldsymbol{\Lambda}) = 0} \left\{ -\mathbb{E}_{\mathbb{P}_{n}} \left[ \sup_{\Delta \in \mathbb{R}^{d}} \left\{ 2\operatorname{tr}(\boldsymbol{H}\boldsymbol{\Lambda}^{\mathsf{T}}\boldsymbol{X}\Delta^{\mathsf{T}}) + \Delta^{\mathsf{T}}\boldsymbol{H}\boldsymbol{\Lambda}^{\mathsf{T}}\Delta - \|\Delta\|_{2}^{2} \right\} + \operatorname{tr}(\boldsymbol{\Lambda}^{\mathsf{T}}\boldsymbol{X}\boldsymbol{X}^{\mathsf{T}}\boldsymbol{H}) \right] \right\}.$$

Introduce the scaling  $\Delta = \bar{\Delta}/n^{1/2}$  and  $\bar{\Lambda} = \Lambda n^{1/2}$ . Then we have

$$n\mathcal{R}_{n}(\boldsymbol{H}) = \sup_{\bar{\boldsymbol{\Lambda}}\in\mathbb{R}^{d\times d}:\operatorname{diag}(\bar{\boldsymbol{\Lambda}})=0} \left\{ -\mathbb{E}_{\mathbb{P}_{n}} \left[ \sup_{\bar{\boldsymbol{\Delta}}\in\mathbb{R}^{d}} \left\{ 2\operatorname{tr}(\boldsymbol{H}\bar{\boldsymbol{\Lambda}}^{\top}X\bar{\boldsymbol{\Delta}}^{\top}) + \bar{\boldsymbol{\Delta}}^{\top}\boldsymbol{H}\bar{\boldsymbol{\Lambda}}^{\top}\bar{\boldsymbol{\Delta}}/n^{1/2} - \left\|\bar{\boldsymbol{\Delta}}\right\|_{2}^{2} \right\} + n^{1/2}\operatorname{tr}(\bar{\boldsymbol{\Lambda}}^{\top}XX^{\top}\boldsymbol{H}) \right] \right\}.$$

Under Assumption 3.3, we have, for any matrix  $\Lambda \in \mathbb{R}^{d \times d}$  such that  $\operatorname{diag}(\Lambda) = 0$  and  $\Lambda \neq 0$ ,

$$\mathbb{P}^*\left(\sum_{i=1}^d \left(\operatorname{tr}(Xh_{i\cdot}^{*^{\top}}\boldsymbol{\Lambda}^{\top} + X\lambda_{i\cdot}^{\top}\boldsymbol{H}^{*^{\top}})\right)^2 > 0\right) > 0,$$

where  $h_{i.}^{*\top}$  represents the *i*-th row of  $H^*$  and  $\lambda_{i.}^{\top}$  the *i*-th row of  $\Lambda$ . Then, Assumptions A2) - A4) in Blanchet et al. (2019) are satisfied, and by Lemma 2 in Blanchet et al. (2019), for every  $\varepsilon > 0$ , there exists  $n_0 > 0$  and  $b \in (0, \infty)$  such that for all  $n \ge n_0$ ,

$$\mathbb{P}\left(\sup_{\|\bar{\mathbf{A}}\|_{F} \geq b} \left\{ -\mathbb{E}_{\mathbb{P}_{n}} \left[\sup_{\bar{\Delta} \in \mathbb{R}^{d}} \left\{ 2\operatorname{tr}(\boldsymbol{H}\bar{\mathbf{\Lambda}}^{\top}X\bar{\Delta}^{\top}) + \bar{\Delta}^{\top}\boldsymbol{H}\bar{\mathbf{\Lambda}}^{\top}\bar{\Delta}/n^{1/2} - \|\bar{\Delta}\|_{2}^{2} \right\} + n^{1/2}\operatorname{tr}(\bar{\mathbf{\Lambda}}^{\top}XX^{\top}\boldsymbol{H}) \right] \right\} > 0 \right) \leq \varepsilon.$$

This result means that if we want the value in the outer supremum to be larger than 0 with high probability as n approaches infinity, we need  $\|\bar{\Lambda}\|_F$  smaller than a finite b. In other words, the  $\bar{\Lambda}^*$  that attains the supremum will have  $\|\bar{\Lambda}^*\|_F$  smaller than a finite b. In this case, for any fixed H,  $\|H\|_F$  should be finite, then

$$\bar{\Delta}^{\top} \boldsymbol{H} \bar{\boldsymbol{\Lambda}}^{*\top} \bar{\Delta} / n^{1/2} \leq \left\| \bar{\Delta} \right\|_{2}^{2} \left\| \boldsymbol{H} \right\|_{F} \left\| \bar{\boldsymbol{\Lambda}}^{*} \right\|_{F} / n^{1/2} \leq b \left\| \bar{\Delta} \right\|_{2}^{2} \left\| \boldsymbol{H} \right\|_{F} / n^{1/2},$$

which is negligible is negligible compared with  $\|\bar{\Delta}\|_2^2$  as  $n \to \infty$ . The remaining terms in the inner supremum can be simplified:

$$\begin{split} \sup_{\bar{\Delta}\in\mathbb{R}^d} &\left\{ 2\operatorname{tr}(\boldsymbol{H}\bar{\boldsymbol{\Lambda}}^{\mathsf{T}}X\bar{\Delta}^{\mathsf{T}}) - \left\|\bar{\Delta}\right\|_2^2 \right\} \\ = \sup_{\bar{\Delta}\in\mathbb{R}^d} &\left\{ 2\left\|\boldsymbol{H}\bar{\boldsymbol{\Lambda}}^{\mathsf{T}}X\right\|_2 \left\|\bar{\Delta}\right\|_2 - \left\|\bar{\Delta}\right\|_2^2 \right\} \\ = &\left\|\boldsymbol{H}\bar{\boldsymbol{\Lambda}}^{\mathsf{T}}X\right\|_2^2. \end{split}$$

Also, we can write

$$\mathbb{E}_{\mathbb{P}_n}\left[n^{1/2}\operatorname{tr}(\bar{\boldsymbol{\Lambda}}^{\top}XX^{\top}\boldsymbol{H})\right] = \operatorname{tr}\left(n^{1/2}\left(\bar{\boldsymbol{\Lambda}}^{\top}\mathbb{E}_{\mathbb{P}_n}\left[XX^{\top}\right]\boldsymbol{H} - \bar{\boldsymbol{\Lambda}}^{\top}\mathbb{E}_{\mathbb{P}^*}\left[XX^{\top}\right]\boldsymbol{H}\right)\right)$$

because the diagonals of  $\bar{\mathbf{\Lambda}}^{\top}$  are zero, and by definition, the off-diagonals of  $\mathbb{E}_{\mathbb{P}^*}[XX^{\top}] \mathbf{H}$  are zero, thus the additional term  $\bar{\mathbf{\Lambda}}^{\top} \mathbb{E}_{\mathbb{P}^*}[XX^{\top}] \mathbf{H} = \mathbf{0}$ . Then by Assumption 3.1, as  $n \to \infty$ ,

$$\mathbb{E}_{\mathbb{P}_n}\left[n^{1/2}\operatorname{tr}(\bar{\boldsymbol{\Lambda}}^{\mathsf{T}}XX^{\mathsf{T}}\boldsymbol{H})\right] \Longrightarrow \operatorname{tr}\left(\bar{\boldsymbol{\Lambda}}^{\mathsf{T}}\boldsymbol{Z}\boldsymbol{H}\right)$$

where  $\boldsymbol{Z} \sim N(0, \Upsilon_g)$ , and  $g(X) \coloneqq XX^{\intercal}$ . Finally, as  $n \to \infty$ ,

$$\mathbb{E}_{\mathbb{P}_n}\left[\left\|\boldsymbol{H}\bar{\boldsymbol{\Lambda}}^{\mathsf{T}}\boldsymbol{X}\right\|_2^2\right] \Rightarrow \mathbb{E}_{\mathbb{P}^*}\left[\left\|\boldsymbol{H}\bar{\boldsymbol{\Lambda}}^{\mathsf{T}}\boldsymbol{X}\right\|_2^2\right].$$

Because the dimension of  $\overline{\Lambda}$  is fixed at d, we can safely take the limit inside the supremum. Therefore, we conclude that, as  $n \to \infty$ ,

$$n\mathcal{R}_n(\boldsymbol{H}) \Rightarrow \sup_{\bar{\boldsymbol{\Lambda}} \in \mathbb{R}^{d \times d} : \operatorname{diag}(\bar{\boldsymbol{\Lambda}}) = 0} \left\{ -\mathbb{E}_{\mathbb{P}^*} \left[ \left\| \boldsymbol{H} \bar{\boldsymbol{\Lambda}}^\top \boldsymbol{X} \right\|_2^2 \right] - \operatorname{tr} \left( \bar{\boldsymbol{\Lambda}}^\top \boldsymbol{Z} \boldsymbol{H} \right) \right\}.$$

This supremum can be bounded from above by substituting  $H\Lambda^{\top}$  with any  $G \in \mathbb{R}^{d \times d}$ :

$$\sup_{\bar{\boldsymbol{\Lambda}} \in \mathbb{R}^{d \times d} : \operatorname{diag}(\bar{\boldsymbol{\Lambda}}) = 0} \left\{ -\mathbb{E}_{\mathbb{P}^*}[\|\boldsymbol{H}\bar{\boldsymbol{\Lambda}}^\top X\|_2^2] - \operatorname{tr}(\bar{\boldsymbol{\Lambda}}^\top \boldsymbol{Z}\boldsymbol{H}) \right\}$$
$$\leq \sup_{\boldsymbol{G} \in \mathbb{R}^{d \times d}} \left\{ -\mathbb{E}_{\mathbb{P}^*}[\|\boldsymbol{G}X\|_2^2] - \operatorname{tr}(\boldsymbol{G}\boldsymbol{Z}) \right\}.$$

Breaking up G into rows, where the *i*-th row is  $G_{i}$ , and let  $Z_{i}$  be the *i*-th column of Z, we have

$$n\mathcal{R}_n(\boldsymbol{H}) \lesssim_D \sup_{\bar{\boldsymbol{\Lambda}} \in \mathbb{R}^{d \times d} : \operatorname{diag}(\bar{\boldsymbol{\Lambda}}) = 0} \left\{ \sum_{i=1}^d \left( -\mathbb{E}_{\mathbb{P}^*}[(G_{i}^\top X)^2] - G_{i}^\top Z_{\cdot i} \right) \right\}.$$

Taking the derivative with respect to  $G_{i}$ , we obtain

$$-2\mathbb{E}_{\mathbb{P}^*}[X^{\top}G_{i}X] - Z_{\cdot i} = 0.$$
(A.7)

Let  $\Sigma_* = \mathbb{E}_{\mathbb{P}^*}(XX^{\top})$ , which we assume to be invertible. Then (A.7) can be written as

$$-2\Sigma_*G_{i\cdot}-Z_{\cdot i}=0,$$

which has a unique solution:

$$G_{i\cdot} = -\frac{1}{2}\boldsymbol{\Sigma}_*^{-1} Z_{\cdot i},$$

where  $\Sigma_*^{-1}$  is the inverse of  $\Sigma_*$ . Therefore,

$$-\mathbb{E}_{\mathbb{P}^*}[(G_{i\cdot}^{\mathsf{T}}X)^2] - G_{i\cdot}^{\mathsf{T}}Z_{\cdot i} = -G_{i\cdot}^{\mathsf{T}}\boldsymbol{\Sigma}_*G_{i\cdot} - G_{i\cdot}^{\mathsf{T}}Z_{\cdot i}$$
$$= -\frac{1}{4}Z_{\cdot i}^{\mathsf{T}}\boldsymbol{\Sigma}_*^{-1}\boldsymbol{\Sigma}_*\boldsymbol{\Sigma}_*^{-1}Z_{\cdot i} + \frac{1}{2}Z_{\cdot i}^{\mathsf{T}}\boldsymbol{\Sigma}_*^{-1}Z_{\cdot i}$$
$$= \frac{1}{4}Z_{\cdot i}^{\mathsf{T}}\boldsymbol{\Sigma}_*^{-1}Z_{\cdot i},$$

and we conclude that

$$n\mathcal{R}_n(\boldsymbol{H}^*) \lesssim_D \sum_{i=1}^d rac{1}{4} Z_{\cdot i}^{ op} \boldsymbol{\Sigma}_*^{-1} Z_{\cdot i}$$

where  $Z_{i}$  is the *i*-th column of  $Z \sim N(0, \Upsilon_g)$ .

## Appendix B Additional Numerical Results

#### **B.1** Simulation analysis

We set a moderate noise level  $\sigma_e^2 = 1$  and test  $\beta_H^2(i) = 0.1, 0.2, \ldots, 0.9$  for all *i*, each value on 10 random trials. The higher  $\beta_H(i)$  is, the less group-specific information there is in the data, and when  $\beta_H(i) = 1$ , there is no group-specific information. The average AMI of each method is shown in Figure 7. The performances all decrease as the common factor becomes more dominant. DRO performs noticeably better than Lasso in this experiment, while both outperform ACC and *k*-medoids again.



Figure 7: Comparison of average AMI between different common factor loadings, with  $\sigma_e^2 = 1.0$ and  $d_k$  randomly chosen.

In the previous experiments, we assume fixed common factor loadings and noise levels across all variables, i.e.,  $\beta_H(i)$  is the same for all *i*. However, in a more realistic setting each variable may have a different factor loading and a different signal-to-noise ratio. To further test the performance of the algorithms in those situations, we let each  $\beta_H(i)^2$  and  $\operatorname{Var}(\varepsilon_i)$  be drawn independently and uniformly from [0,0.5]. We keep the dimension of the problem at d = 500 and the number of observations at n = 250. As done previously, the number of factors controlling each cluster k is randomly chosen from 1 to  $m_k - 1$ , where  $m_k$  is the number of variables in cluster k. We run the experiment on 10 different random trials and calculate the average AMI.

Before showing the clustering results, we first visualize the sample correlation matrix and compare it with the C matrix extracted by the subspace clustering methods. Figure 8a shows the true clustering structure, and Figure 8b shows a heatmap of the sample correlation matrix. The blocks along the diagonal of Figure 8b are nearly indistinguishable. Figures 9a and 9b show the C matrices from our regularized regressions. Similar to the previous experiments, both methods can extract the blocks by making them visually more prominent, with Lasso extracting a sparse Cmatrix while DRO keeps more entries in the matrix but at lower magnitudes.



Figure 8: The heatmap of the true similarity matrix (a) and the sample correlation matrix (b), with  $\beta_H(i)^2$  and  $\operatorname{Var}(\varepsilon_i)$  drawn independently and uniformly from [0, 0.5], and  $d_k$  randomly chosen. Variables in the same cluster have similarity 1 and otherwise 0.



Figure 9: The heatmap of the C matrices for DRO (a) and Lasso (b), with  $\beta_H(i)^2$  and  $\operatorname{Var}(\varepsilon_i)$  drawn independently and uniformly from [0,0.5], and  $d_k$  randomly chosen.

Table 3 shows the average AMI of each clustering method over the 10 random trials. DRO performs the best and much better compared to Lasso. Both methods still outperform ACC and k-medoids. As expected, neither ACC nor k-medoids can meaningfully recover the clusters in this experiment due to the additional complexity in the underlying model.

#### B.2 Additional portfolio backtesting results

Here we present Sharpe ratios of the portfolios in backtesting with different values for  $K_1 = K_2$ . As shown in Table 4, with annual portfolio updates, the DRO-ACC performs well with  $3 \times 3$  and Table 3: Average AMI of different clustering methods compared with ground truth, over 10 different random trials.

	DRO	Lasso	ACC	k-medoids
Average AMI	0.92	0.83	0.15	0.33

 $4 \times 4$  clusters. With  $5 \times 5$  clusters, the DRO-ACC portfolio underperforms the other portfolios in Sharpe ratio but still outperforms the benchmark SPY.

Number of clusters	$3 \times 3 = 9$	$4 \times 4 = 16$	$5 \times 5 = 25$	$6 \times 6 = 36$
DRO-ACC	0.87	0.87	0.73	1.05
CORD	0.9	0.85	0.79	0.77
k-medoids	0.82	0.85	0.84	0.9
SPY	0.36	0.36	0.36	0.36

Table 4: Sharpe Ratio of portfolios with different numbers of clusters, updated annually.

We also present the Sharpe ratios with monthly and quarterly portfolio updates. This means that the stocks are selected monthly/quarterly and weights re-calculated using the newest clustering results. As shown in Tables 5 and 6, DRO-ACC portfolios are also robust to the stock selection and allocation update frequency, as they tend to outperform with both monthly and quarterly updates.

Table 5: Sharpe Ratio of portfolios with different numbers of clusters, updated monthly.

Number of clusters	$3 \times 3 = 9$	$4 \times 4 = 16$	$5 \times 5 = 25$	$6 \times 6 = 36$
DRO-ACC	0.8	0.83	0.88	0.89
CORD	0.65	0.84	0.77	0.74
k-medoids	0.78	0.8	0.81	0.87
SPY	0.36	0.36	0.36	0.36

Table 6: Sharpe Ratio of portfolios with different numbers of clusters, updated quart
---

Number of clusters	$3 \times 3 = 9$	$4 \times 4 = 16$	$5 \times 5 = 25$	$6 \times 6 = 36$
DRO-ACC	0.91	0.83	0.78	0.87
CORD	0.76	0.82	0.88	0.84
k-medoids	0.72	0.78	0.78	0.82
SPY	0.36	0.36	0.36	0.36

## References

- Ando, T., & Bai, J. (2017). Clustering huge number of financial time series: A panel data approach with high-dimensional predictors and factor structures. Journal of the American Statistical Association, 112(519), 1182–1198.
- Aste, T., Shaw, W., & Di Matteo, T. (2010). Correlation structure and dynamics in volatile markets. New Journal of Physics, 12(8), 085009.

- Bani-Domi, W., & Kittaneh, F. (2008). Norm equalities and inequalities for operator matrices. Linear Algebra and Its Applications, 429(1), 57–67.
- Barfuss, W., Massara, G. P., Di Matteo, T., & Aste, T. (2016). Parsimonious modeling with information filtering networks. *Physical Review E*, 94(6), 1–17.
- Bernardes, J. S., Vieira, F. R., Costa, L. M., & Zaverucha, G. (2015). Evaluation and improvements of clustering algorithms for detecting remote homologous protein families. *BMC Bioinformatics*, 16(1), 1–14.
- Bickel, P. J., & Levina, E. (2004). Some theory for Fisher's linear discriminant function, 'naive Bayes', and some alternatives when there are many more variables than observations. Bernoulli, 10(6), 989–1010.
- Blanchet, J., Kang, Y., & Murthy, K. (2019). Robust Wasserstein profile inference and applications to machine learning. *Journal of Applied Probability*, 56(3), 830–857.
- Blanchet, J., & Murthy, K. (2019). Quantifying distributional model risk via optimal transport. Mathematics of Operations Research, 44 (2), 565–600.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., & Eckstein, J. (2010). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1), 1–122.
- Bunea, F., Giraud, C., Luo, X., Royer, M., & Verzelen, N. (2020). Model assisted variable clustering: Minimax-optimal recovery and algorithms. *The Annals of Statistics*, 48(1), 111–137.
- Callot, L., Caner, M., Önder, A. Ö., & Ulaşan, E. (2019). A nodewise regression approach to estimating large portfolios. *Journal of Business and Economic Statistics*, 2021(2), 520–531.
- Chen, J., & Yang, J. (2014). Robust subspace segmentation via low-rank representation. *IEEE Transactions on Cybernetics*, 44 (8), 1432–1445.
- Elhamifar, E., & Vidal, R. (2013). Sparse subspace clustering: Algorithm, theory, and applications. IEEE Transactions on Pattern Analysis and Machine Intelligence, 35(11), 2765–2781.
- Favaro, P., Vidal, R., & Ravichandran, A. (2011). A closed form solution to robust subspace estimation and clustering. Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 1801–1807.
- Jiang, D., Tang, C., & Zhang, A. (2004). Cluster analysis for gene expression data: A survey. IEEE Transactions on Knowledge and Data Engineering, 16(11), 1370–1386.
- Johnstone, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *The Annals of Statistics*, 29(2), 295–327.
- Kaufman, L., & Rousseeuw, P. J. (1990). Partitioning around medoids (program PAM). Finding groups in data: An introduction to cluster analysis (pp. 68–125). John Wiley & Sons, Ltd.
- Korzeniewski, J. (2018). Efficient stock portfolio construction by means of clustering. Acta Universitatis Lodziensis. Folia Oeconomica, 1(333), 85–92.
- León, D., Aragón, A., Sandoval, J., Hernández, G., Arévalo, A., & Niño, J. (2017). Clustering algorithms for risk-adjusted portfolio construction. *Proceedia Computer Science*, 108, 1334– 1343.
- Liu, G., Lin, Z., Yan, S., Sun, J., Yu, Y., & Ma, Y. (2013). Robust recovery of subspace structures by low-rank representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(1), 171–184.
- Liu, G., & Yan, S. (2011). Latent low-rank representation for subspace segmentation and feature extraction. Proceedings of the IEEE International Conference on Computer Vision, 1615– 1622.
- Mantegna, R. (1999). Hierarchical structure in financial markets. The European Physical Journal B, 11(1), 193–197.

- Meinshausen, N., & Bühlmann, P. (2006). High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, 34(3), 1436–1462.
- Merton, R. C. (1980). On estimating the expected return on the market: An exploratory investigation. Journal of Financial Economics, 8(4), 323–361.
- Musmeci, N., Aste, T., & Di Matteo, T. (2015). Relation between financial market structure and the real economy: Comparison between clustering methods. *PLOS ONE*, 10(3), e0116201.
- Musmeci, N., Nicosia, V., Aste, T., Di Matteo, T., & Latora, V. (2017). The multiplex dependency structure of financial markets. *Complexity*, 2017, 1–13.
- Parsons, L., Haque, E., & Liu, H. (2004). Subspace clustering for high dimensional data: A review. ACM SIGKDD Explorations Newsletter, 6(1), 90–105.
- Peng, J., & Wei, Y. (2007). Approximating k-means-type clustering via semidefinite programming. SIAM Journal on Optimization, 18(1), 186–205.
- Pozzi, F., Di Matteo, T., & Aste, T. (2013). Spread of risk across financial markets: Better to invest in the peripheries. *Scientific Reports*, *3*, 1665.
- Rahimian, H., & Mehrotra, S. (2022). Frameworks and results in distributionally robust optimization. Open Journal of Mathematical Optimization, 3, Article 4.
- Rosén, F. (2006). Correlation based clustering of the Stockholm stock exchange (Doctoral dissertation).
- Soltanolkotabi, M., & Candés, E. J. (2012). A geometric analysis of subspace clustering with outliers. The Annals of Statistics, 40(4), 2195–2238.
- Soltanolkotabi, M., Elhamifar, E., & Candès, E. J. (2014). Robust subspace clustering. The Annals of Statistics, 42(2), 669–699.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., & Fergus, R. (2014). Intriguing properties of neural networks. 2nd International Conference on Learning Representations, ICLR 2014 - Conference Track Proceedings.
- Tang, W., Xu, X., & Zhou, X. Y. (2022). Asset selection via correlation blockmodel clustering. Expert Systems with Applications, 195, 116558.
- Tumminello, M., Aste, T., Matteo, T. D., & Mantegna, R. N. (2005). A tool for filtering information in complex systems. Proceedings of the National Academy of Sciences, 102(30), 10421– 10426.
- Tumminello, M., Di Matteo, T., Aste, T., & Mantegna, R. N. (2007). Correlation based networks of equity returns sampled at different time horizons. *The European Physical Journal B*, 55(2), 209–217.
- Vidal, R. (2011). Subspace clustering. *IEEE Signal Processing Magazine*, 28(2), 52–68.
- Villani, C. (2009). Optimal transport: Old and new (Vol. 338). Springer Berlin Heidelberg.
- Vinh, N. X., Epps, J., & Bailey, J. (2010). Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *Journal of Machine Learning Research*, 11(95), 2837–2854.
- von Luxburg, U. (2007). A tutorial on spectral clustering. Statistics and Computing, 17, 395–416.
- Wang, Y. X., & Xu, H. (2016). Noisy sparse subspace clustering. Journal of Machine Learning Research, 17, 320–360.
- Zaremba, A., & Shemer, J. (2017). Is risk always rewarded? low-volatility anomalies. Country asset allocation: Quantitative country selection strategies in global factor investing (pp. 81–104). Palgrave Macmillan US.