

Confidence Intervals for Maximin Effects in Inhomogeneous Large-Scale Data

Dominik Rothenhäusler, Nicolai Meinshausen, and Peter Bühlmann

Abstract One challenge of large-scale data analysis is that the assumption of an identical distribution for all samples is often not realistic. An optimal linear regression might, for example, be markedly different for distinct groups of the data. Maximin effects have been proposed as a computationally attractive way to estimate effects that are common across all data without fitting a mixture distribution explicitly. So far just point estimators of the common maximin effects have been proposed in Meinshausen and Bühlmann (Ann Stat 43(4):1801–1830, 2015). Here we propose asymptotically valid confidence regions for these effects.

1 Introduction

Large-scale regression analysis often has to deal with inhomogeneous data in the sense that samples are not drawn independently from the same distribution. The optimal regression coefficient might for example be markedly different in distinct groups of the data or vary slowly over a chronological ordering of the samples. One option is then to either model the exact variation of the regression vector with a varying-coefficient model in the latter case [5, 7] or to fit a mixture distribution in the former [1, 6, 8]. For large-scale analysis with many groups of data samples or many predictor variables this approach might be too expensive computationally and also yield more information than necessary in settings where one is just interested in effects that are present in all sub-groups of data. A maximin effect was defined in [9] as the effect that is common to all sub-groups of data and a simple estimator based on subsampling of the data was proposed in [3]. However, the estimators for maximin effects proposed so far just yield point estimators but we are interested here in confidence intervals. While we are mostly dealing with low-dimensional data where the sample size exceeds the number of samples, the results could potentially be extended to high-dimensional regression using similar ideas as proposed for

D. Rothenhäusler (✉) • N. Meinshausen • P. Bühlmann
Seminar für Statistik, ETH Zürich, Rämistrasse 101, 8092 Zürich, Switzerland
e-mail: rothenhaeusler@stat.math.ethz.ch; meinshausen@stat.math.ethz.ch;
buehlmann@stat.math.ethz.ch

© Springer International Publishing Switzerland 2016
A. Frigessi et al. (eds.), *Statistical Analysis for High-Dimensional Data*,
Abel Symposia 11, DOI 10.1007/978-3-319-27099-9_12

255

example in [12] or [11] for the estimation of optimal linear regression effects for high-dimensional data.

1.1 Model and Notation

We first present a model for inhomogeneous data as considered in [9]. Specifically, we look at a special case where the data are split into several known groups $g = 1, \dots, G$. In each group g , we assume a linear model of the form

$$Y_g = \mathbf{X}_g b_g^0 + \varepsilon_g, \quad (1)$$

where Y_g is a n -dimensional response vector of interest, b_g^0 a deterministic p -dimensional regression parameter vectors and \mathbf{X}_g a $n \times p$ -dimensional design matrix containing in the columns the n observations of p predictor variables. The noise contributions ε_g are assumed to be independent with distribution $\mathcal{N}_n(0, \sigma^2 \text{Id}_n)$. We assume the sample size n to be identical in each group. Generalizations to varying-coefficient models [5, 7] are clearly possible but notationally more cumbersome. Inhomogeneity is caused by the different parameter vectors in the group. We define \mathbf{X} as the row-wise concatenation of the design matrices $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_G$ and assume that the groups are known, that is we know which observations belong to the groups $g = 1, \dots, G$, respectively. For the distribution of \mathbf{X}_g , $g = 1, \dots, G$ we consider different scenarios.

1.1.1 Scenario 1

Random design. The observations of the predictor variables are independent samples of an unknown multivariate distribution F with finite fourth moments. We assume this distribution to be common across all groups $g = 1, \dots, G$.

1.1.2 Scenario 2

Random design in each group. The observation in each group are independent samples of an unknown distribution F_g with finite fourth moments. Observations in different groups are independent. The distribution F_g may be different in different groups.

In the following if not mentioned otherwise we assume Scenario 1. The generalization to Scenario 2 is to a large extent only notational.

1.2 Aggregation

The question arises how the inhomogeneity of the optimal regression across groups is taken into account when trying to estimate the relationship between the predictor variables and the outcome of interest. Several known alternatives such as mixed effects models [10], mixture models [8] and clusterwise regression models [4] are possibilities and are useful especially in cases where the group structure is unknown. They are at the same time computationally quite demanding.

A computationally attractive alternative (especially for the discussed case of known groups but also more generally) is to estimate the optimal regression coefficient separately in each group, which are either known (as assumed in the following) or sampled in some appropriate form [9]. As estimates for the b_g^0 we use in the following standard least squares estimators

$$\hat{b}_g = \arg \min_{b \in \mathbb{R}^p} \|Y_g - \mathbf{X}_g b\|_2^2.$$

The restriction to this estimator is only for the purpose of simplicity. Regularization can be added if necessary but the essential issues are already visible for least-squares estimation.

Now a least-squares estimator is obtained in each group of data and the question is how these different estimators can be aggregated. The simplest and perhaps most widely-used aggregation scheme is bagging (bootstrap aggregation), as proposed by Breiman [2], where the aggregated estimator is given by

$$\text{Bagging : } \hat{b} := \sum_g w_g \hat{b}_g, \quad \text{where } w_g = \frac{1}{G} \quad \forall g = 1, \dots, G. \quad (2)$$

If the data from different groups originate from an independent sampling mechanism, the bagging is a useful aggregation scheme. In particular, computing the bagged estimator is computationally more attractive than computing a single least-squares estimator as it allows the data to be split up into distinct subsets and processed independently before the aggregation step. For inhomogeneous data, the variability of the estimates \hat{b}_g for $g = 1, \dots, G$ allows to gain some insight into the nature of the inhomogeneity. However, as argued in [3], averaging is the wrong aggregation mechanism for inhomogeneous data.

1.3 Maximin Effect and Merging

For inhomogeneous data, instead of looking for an estimator that works best on average, Meinshausen and Bühlmann [9] proposed to aim to maximize the minimum explained variance across several settings $g = 1, \dots, G$. To be more precise, in our

setting,

$$b_{\text{maximin}} := \arg \max_{b \in \mathbb{R}^p} \min_{g=1, \dots, G} V(b, b_g^0),$$

where $V(b, b_g^0)$ is the explained variance in group g (with true regression vector b_g^0) when using a regression vector b . That is

$$\begin{aligned} V(b, b_g^0) &:= \mathbb{E} \|Y_g\|_2^2 - \mathbb{E} \|Y_g - X_g b\|_2^2 \\ &= 2b^t \Sigma^0 b_g^0 - b^t \Sigma^0 b, \end{aligned}$$

where $\Sigma^0 := \mathbb{E} \hat{\Sigma}$ with $\hat{\Sigma} := (nG)^{-1} \mathbf{X}' \mathbf{X}$ is the sample covariance matrix. In words, the maximin effect is defined as the estimator that maximises the explained variance in the most adversarial scenario (“group”). In this sense, the maximin effect is the effect that is common among all groups in the data and ignores the effects that are present in some groups but not in others. It was shown in [9] that the definition above is equivalent to

$$b_{\text{maximin}} = \arg \min_{b \in \text{CVX}(B^0)} b^t \Sigma^0 b,$$

where $B^0 = (b_1^0, \dots, b_G^0) \in \mathbb{R}^{p \times G}$ the matrix of the regression parameter vectors and $\text{CVX}(B^0)$ denotes the closed convex hull of the G vectors in B^0 . The latter definition motivates **maximin aggregating**, or **magging** [3], which is the convex combination that minimizes the ℓ_2 -norm of the fitted values:

$$\begin{aligned} \text{Magging: } \hat{b} &:= \sum_{g=1}^G \alpha_g \hat{b}_g, \quad \text{where } \alpha := \arg \min_{\alpha \in C_G} \left\| \sum_{g=1}^G \alpha_g \mathbf{X} \hat{b}_g \right\|_2 \text{ and} \\ C_G &:= \left\{ \alpha \in \mathbb{R}^G : \min_g \alpha_g \geq 0 \text{ and } \sum_g \alpha_g = 1 \right\} \end{aligned}$$

The magging regression vector is unique if $\mathbf{X}' \mathbf{X}$ is positive definite. Otherwise, we can only identify the prediction effect $\mathbf{X} b_{\text{maximin}}$ and the solution above is meant to be any member of the feasible set of solutions. To compute the estimator, the dataset is split into several smaller datasets and we assume here that the split separates the data into already known groups. After computing estimators on all of these groups separately, possibly in parallel, magging can be used to find common effects of all datasets. This is in particular interesting if there is inhomogeneity in the data. For known groups, as in our setting, magging can be interpreted as the plug-in estimate of the maximin effect.

In the following we need additional notation. For $B := (b_1, \dots, b_G) \in \mathbb{R}^{p \times G}$ and for $\Sigma \in \mathbb{R}^{p \times p}$ positive definite define

$$M_\Sigma(B) := \arg \min_{b \in \text{CVX}(B)} b' \Sigma b$$

We obtain the original definition of the magging estimator for $M_{\hat{\Sigma}}(\hat{B})$ with $\hat{B} = (\hat{b}_1, \dots, \hat{b}_G)$ and the maximin effect with $M_{\Sigma^0}(B^0)$.

1.4 Novel Contribution and Organization of the Paper

So far only point estimators of maximin effects have been proposed in the literature. In Sect. 2 we discuss an asymptotic approach to construct confidence regions for the maximin effect. Specifically, we calculate the asymptotic distribution of $\sqrt{n}(M_{\hat{\Sigma}}(\hat{B}) - M_{\Sigma^0}(B^0))$ and derive corresponding asymptotically valid confidence regions. This gives us (asymptotically) tight confidence regions and will shed more light on the (asymptotic) nature of the fluctuations of the magging estimator. We evaluate the actual coverage of this approximation on simulated datasets in Sect. 3. The proofs of the corresponding theorems and an alternative non-asymptotic approach can be found in the Appendix. The advantages and disadvantages of the approaches are discussed in Sect. 4.

2 Confidence Intervals for Maximin Effects

In Scenario 1, the random design of the predictor variables is identical across all groups of data. For fixed G and $n \rightarrow \infty$, we can then use the delta method to derive the asymptotic distribution of the scaled difference between the true and estimated magging effects

$$\sqrt{n}(M_{\hat{\Sigma}}(\hat{B}) - M_{\Sigma^0}(B^0)).$$

This in turn allows to construct confidence intervals for the true maximin effects. Let $W(\hat{B}, \hat{\Sigma})$ be a consistent estimator of the (positive definite) variance of the Gaussian distribution

$$\lim_{n \rightarrow \infty} \sqrt{n}(M_{\hat{\Sigma}}(\hat{B}) - M_{\Sigma^0}(B^0)).$$

Let $\alpha > 0$. Choose τ as the $(1 - \alpha)$ -quantile of the χ_p^2 -distribution. Define then a confidence region as

$$\mathbf{C}(\hat{\Sigma}, \hat{B}) := \{M \in \mathbb{R}^p : (M_{\hat{\Sigma}}(\hat{B}) - M)' W(\hat{B}, \hat{\Sigma})^{-1} (M_{\hat{\Sigma}}(\hat{B}) - M) \leq \frac{\tau}{n}\} \quad (3)$$

The definition of $W(\hat{B}, \hat{\Sigma})$ is deferred to the Appendix, section “Definitions and Formulas”. We will show in the following that we obtain asymptotically valid confidence intervals with this approach. For simplicity, we work with Scenario 1 here and assume that the noise contributions ε_g in Eq. (1) are independent with distribution $\mathcal{N}_n(0, \sigma^2 \text{Id}_n)$. Furthermore, each $\mathbf{X}_g \in \mathbb{R}^{n \times p}$ is assumed to have full rank, requiring $p \leq n$. Though the framework for the result is a Gaussian linear model, it can be easily extended to more general settings.

The following theorem describes the coverage properties of the confidence interval (3). In the following, for $x, y \in \mathbb{R}^p$ and $\Sigma \in \mathbb{R}^{p \times p}$ positive definite define $\langle x, y \rangle_\Sigma := x^t \Sigma y$.

Theorem 1 *Let Σ^0 be positive definite. Let $M_{\Sigma^0}(B^0) = \sum_{g=1}^G \alpha_g b_g^0$ with $\alpha_g \geq 0$, $\sum_{g=1}^G \alpha_g = 1$ and let this representation be unique. Let $|\{g : \alpha_g \neq 0\}| > 1$. Suppose that the hyperplane orthonormal to the maximin effect contains only “active” b_g^0 , i.e. $\{b_g^0 : g = 1, \dots, G\} \cap \{M \in \mathbb{R}^p : \langle M - M_{\Sigma^0}(B^0), M_{\Sigma^0}(B^0) \rangle_{\Sigma^0} = 0\} \subset \{b_g^0 : \alpha_g \neq 0\}$. Then*

$$\lim_{n \rightarrow \infty} \mathbb{P}[M_{\Sigma^0}(B^0) \in \mathbf{C}(\hat{\Sigma}, \hat{B})] = 1 - \alpha.$$

In other words, the set defined in (3) is an asymptotically valid confidence region for $M_{\Sigma^0}(B^0)$ under the made assumptions. If the true coefficients b_g^0 in each group are drawn from a multivariate density, then the assumptions are fulfilled with probability one.

The special case $|\{g : \alpha_g \neq 0\}| = 1$ is excluded, as the magging estimator is identical to a solution in one individual group in this case, which is equivalent to $M_{\hat{\Sigma}}(\hat{B}) = \hat{b}_g$ for a $g \in \{1, \dots, G\}$, up to an asymptotically negligible set. This case is mainly excluded for notational reasons. The assumptions of Theorem 1 guarantee that the derivative of magging $M_\Sigma(B)$ exists and is continuous at B^0 and Σ^0 . If the latter condition is violated, it is still possible to obtain asymptotic bounds in the more general setting, as $\lim_{n \rightarrow \infty} \sqrt{n}(M_{\hat{\Sigma}}(\hat{B}) - M_{\Sigma^0}(B^0))$ is still subgaussian. We explore the violation of these assumptions with simulation studies in the next section. The proof of Theorem 1 is an application of Slutsky’s Theorem, combined with the following result about the asymptotic variance of the magging estimator.

Theorem 2 *Let the assumptions of Theorem 1 be true. Then, for $n \rightarrow \infty$,*

$$\begin{aligned} & \sqrt{n} \left(M_{\hat{\Sigma}}(\hat{B}) - M_{\Sigma^0}(B^0) \right) \\ & \rightarrow \mathcal{N} \left(0, \sigma^2 \sum_{g \in A(B^0, \Sigma^0)} D_g^t M_{\Sigma^0}(B^0) \Sigma^{-1} D_g M_{\Sigma^0}(B^0) + V(B_{A(B^0, \Sigma^0)}^0, \Sigma^0) \right). \end{aligned} \quad (4)$$

Here, D_g denotes the differential in direction b_g . This derivative is calculated in the Appendix, see section “Definitions and Formulas”. The set $A(B, \Sigma) \subset \{1, \dots, G\}$ denotes indices g for which b_g has nonvanishing coefficient α_g in one of

the convex combinations $M_{\Sigma}(B) = \sum_{g=1,\dots,G} \alpha_g b_g$ with $\alpha_g \geq 0$, $\sum_{g=1,\dots,G} \alpha_g = 1$. Note that by the assumptions of Theorem 1 this convex combination is unique for $M_{\Sigma^0}(B^0)$. The definition of $V(B_{A(B,\Sigma)}, \Sigma)$ is somewhat lengthy and can be found in the Appendix, section “Definitions and Formulas”.

The first summand in the variance in formula (4) is due to fluctuations of the estimator of B^0 , the second summand is due to fluctuations of the estimator of Σ^0 . If Σ^0 is known in advance, we can use $\hat{\Sigma} := \Sigma^0$ and in the theorem above $V = 0$. Figure 1 is an illustration of Theorem 2.

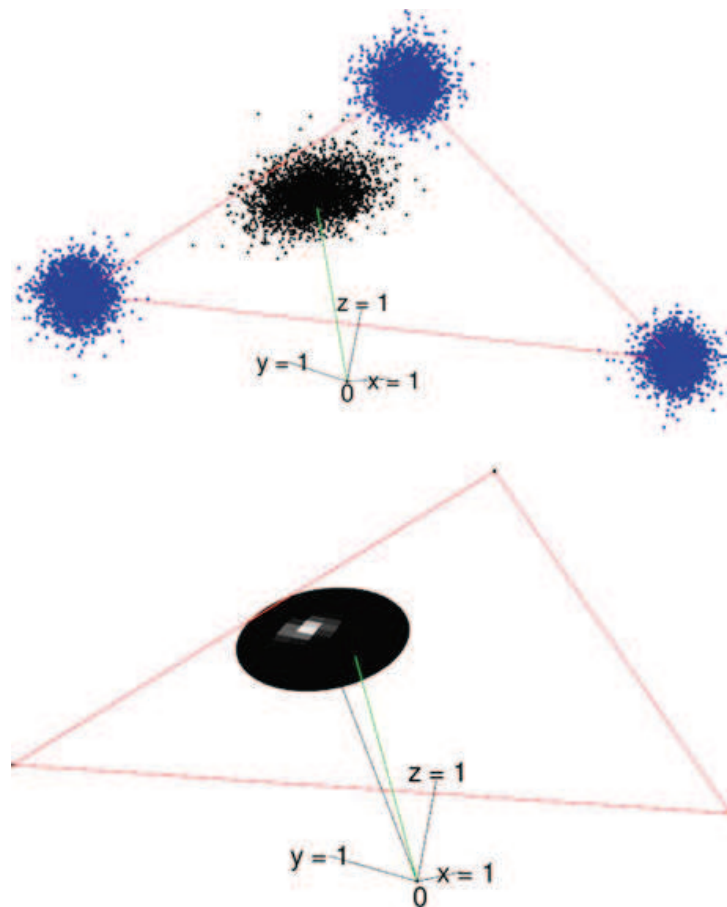


Fig. 1 An illustration of Theorems 1 and 2. On the *upper plot* the *blue dots* represent 3000 realizations of \hat{b}_g , $g = 1, 2, 3$ with dimension $p = 3$. The *black dots* are the corresponding mapping estimates $M_{\hat{\Sigma}}(\hat{B})$. The *green line* indicates the true maximin effect $M_{\Sigma^0}(B^0)$. On the *lower plot*, the *black line* indicates one of the $M_{\hat{\Sigma}}(\hat{B})$ with the corresponding approximate 95 %-confidence region calculated with the terms of Eq. (3)

3 Numerical Examples

The aim of this section is to evaluate the actual coverage of the approximate confidence regions as defined above. We study several examples. They have in common that the entries in \mathbf{X} are i.i.d. $\mathcal{N}(0, 1)$. Furthermore the ε_g are i.i.d. $\mathcal{N}(0, \text{Id}_n)$ and independent of \mathbf{X} . The tables show the coverage of the true maximin effect $M_{\Sigma^0}(B^0)$ by the proposed 95% confidence regions. We calculate the confidence intervals only for $p < n$ scenarios as long as least squares estimators are used (Tables 1, 2, and 3), while the case of $p \geq n$ is covered in Tables 4 and 5 by the use of a ridge penalty. All simulations were run 1000 times.

In the setting of Table 1 all assumptions of Theorem 1 are satisfied. As expected, for large p the convergence of the actual coverage seems to be slower. Note that for validity of Theorem 1 it is not necessary that $p = G$, as we have asymptotically tight coverage for all $1 < G \leq p$.

In Tables 2 and 3 we explore the violation of one of the assumptions in Theorem 1. The maximin effect is $M_{\Sigma^0}(B^0) = (1, 0, 0, \dots)$, and the convex combination $M_{\Sigma^0}(B^0) = \sum_{g=1}^G \alpha_g b_g^0$ with $\alpha_g \geq 0$, $\sum \alpha_g = 1$ is not unique. In both cases, this seems to lead to too conservative confidence regions. Generally, in these

Table 1 $b_g^0 = e_g$, $g = 1, \dots, G = p$, where the e_g denote the vectors of the standard basis, 1000 iterations

| | $n = 5$ | 10 | 15 | 100 | 200 | 500 | 1000 | 2000 | 4000 |
|---------|---------|------|------|------|------|------|------|------|------|
| $p = 3$ | 0.70 | 0.78 | 0.82 | 0.92 | 0.94 | 0.95 | 0.94 | 0.94 | 0.95 |
| 5 | | 0.69 | 0.76 | 0.90 | 0.93 | 0.95 | 0.94 | 0.95 | 0.95 |
| 10 | | | 0.62 | 0.84 | 0.88 | 0.94 | 0.95 | 0.96 | 0.94 |
| 15 | | | | 0.78 | 0.85 | 0.93 | 0.92 | 0.95 | 0.95 |
| 20 | | | | 0.72 | 0.83 | 0.90 | 0.91 | 0.95 | 0.94 |
| 40 | | | | 0.54 | 0.63 | 0.79 | 0.88 | 0.91 | 0.94 |
| 80 | | | | 0.57 | 0.38 | 0.50 | 0.74 | 0.85 | 0.92 |

The coverage can be seen to be approximately correct if n is sufficiently large

Table 2 $b_g^0 = e_1 + z_g e_2$, $g = 1, \dots, G = p$, $z_g \sim \mathcal{N}(0, 1)$ independent

| | $n = 5$ | 10 | 15 | 100 | 200 | 500 | 1000 | 2000 | 4000 |
|---------|---------|------|------|------|------|------|------|------|------|
| $p = 3$ | 0.64 | 0.84 | 0.91 | 0.97 | 0.96 | 0.82 | 0.98 | 0.96 | 0.97 |
| 5 | | 0.61 | 0.79 | 0.99 | 0.97 | 0.88 | 0.82 | 0.91 | 1.00 |
| 10 | | | 0.23 | 0.99 | 0.99 | 1.00 | 0.99 | 0.93 | 0.98 |
| 15 | | | | 0.99 | 0.99 | 1.00 | 1.00 | 0.99 | 0.99 |
| 20 | | | | 0.99 | 1.00 | 0.99 | 1.00 | 1.00 | 0.99 |
| 40 | | | | 0.94 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| 80 | | | | 0.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |

The assumptions are violated, yielding too conservative confidence intervals. The 0.00 at $n = 100$, $p = 80$ is due to a large bias of $M_{\hat{\Sigma}}(\hat{B})$ towards 0. For larger n , however, this bias quickly vanishes and we get the desired coverage (starting at approximately $n = 120$)

Table 3 $b_g^0 = e_1, g = 1, \dots, G = [0.8p]$

| | $n = 5$ | 10 | 15 | 100 | 200 | 500 | 1000 | 2000 | 4000 |
|---------|---------|------|------|------|------|------|------|------|------|
| $p = 3$ | 0.76 | 0.87 | 0.90 | 0.99 | 0.99 | 0.99 | 1.00 | 1.00 | 1.00 |
| 5 | | 0.65 | 0.78 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| 10 | | | 0.33 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| 15 | | | | 0.99 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| 20 | | | | 0.99 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| 40 | | | | 0.93 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| 80 | | | | 0.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |

The assumptions are again violated and coverage is too high. At $p = 80$ and $n = 100$ we observe the same effect as in Table 2. In this scenario the estimated confidence regions can become arbitrarily large. This stems from the fact that if some of the \hat{b}_g corresponding to $A(\hat{B}, \hat{\Sigma})$ are very close, the estimated variance of mapping may become large. In this setting a different approach, for example as discussed in section “Relaxation-Based Approach” in the Appendix makes more sense

Table 4 $b_g^0 = e_g, g = 1, \dots, G = p$

| | $n = 5$ | 10 | 15 | 100 | 200 | 500 | 1000 | 2000 | 4000 |
|---------|---------|------|------|------|------|------|------|------|------|
| $p = 3$ | 0.71 | 0.77 | 0.84 | 0.92 | 0.94 | 0.96 | 0.95 | 0.94 | 0.93 |
| 5 | 0.74 | 0.69 | 0.76 | 0.90 | 0.94 | 0.94 | 0.95 | 0.95 | 0.95 |
| 10 | 0.55 | 0.70 | 0.60 | 0.86 | 0.88 | 0.93 | 0.94 | 0.94 | 0.95 |
| 15 | 0.52 | 0.53 | 0.70 | 0.77 | 0.86 | 0.91 | 0.94 | 0.95 | 0.95 |
| 20 | 0.53 | 0.48 | 0.52 | 0.73 | 0.81 | 0.89 | 0.93 | 0.92 | 0.94 |
| 40 | 0.40 | 0.47 | 0.37 | 0.52 | 0.62 | 0.81 | 0.87 | 0.90 | 0.94 |
| 80 | 0.20 | 0.40 | 0.37 | 0.56 | 0.38 | 0.52 | 0.72 | 0.84 | 0.90 |

The diagonal elements of $\hat{\Sigma}$ and $\hat{\Sigma}_g$ where increased by a value 10^{-4} in order to make them invertible and not too ill-conditioned for $n \leq p$. Again, coverage is approximately correct for n sufficiently large

Table 5 This table shows the average maximum eigenvalues of the estimated covariance matrix of $\sqrt{n}(M_{\Sigma^0}(B^0) - M_{\hat{\Sigma}}(\hat{B}))$, analogous to Table 4

| | $n = 5$ | 10 | 15 | 100 | 200 | 500 | 1000 | 2000 | 4000 |
|---------|---------|---------|---------|--------|------|------|------|------|------|
| $p = 3$ | 41.70 | 2.97 | 1.59 | 0.59 | 0.53 | 0.49 | 0.47 | 0.47 | 0.46 |
| 5 | 831.50 | 13.52 | 4.83 | 0.42 | 0.34 | 0.30 | 0.28 | 0.26 | 0.26 |
| 10 | 6.56 | 1935.77 | 27.78 | 0.29 | 0.20 | 0.16 | 0.14 | 0.13 | 0.12 |
| 15 | 0.29 | 19.83 | 3844.87 | 0.26 | 0.16 | 0.12 | 0.10 | 0.09 | 0.08 |
| 20 | 0.08 | 4.25 | 41.04 | 0.29 | 0.15 | 0.09 | 0.08 | 0.07 | 0.06 |
| 40 | 0.01 | 0.04 | 4.61 | 2.71 | 0.16 | 0.07 | 0.05 | 0.04 | 0.03 |
| 80 | 0.00 | 0.00 | 0.01 | 205.85 | 1.09 | 0.06 | 0.03 | 0.02 | 0.02 |

settings the difficulty arises from the fact that the derivative of $M_{\Sigma}(B)$ does not exist at $M_{\Sigma^0}(B^0)$. As a result, the fluctuations of $\lim_n \sqrt{n}(M_{\Sigma^0}(B^0) - M_{\hat{\Sigma}}(\hat{B}))$ —provided that this limit exists—are not necessarily Gaussian anymore.

In the last simulation, depicted in Table 4 the \hat{b}_g , $g = 1, \dots, G$ were not calculated by ordinary least squares but ridge regression. The diagonal elements of $\hat{\Sigma}$ and $\hat{\Sigma}_g$ were increased by a value 10^{-4} in order to make them invertible and not too ill-conditioned for $n \leq p$. Apart from that we used the same setting as in Table 1. As in Table 1, for large n the coverage seems to be (approximately) correct but severe undercoverage can still occur for $n \ll p$. In these high-dimensional settings, the tuning ridge parameter would need to be better adjusted for a useful balance between bias and variance and the bias of the ridge penalty would have to be adjusted for, something which is beyond the current scope. In Table 5 the corresponding maximum eigenvalues of the estimated variance of $\sqrt{n}(M_{\Sigma^0}(B^0) - M_{\hat{\Sigma}}(\hat{B}))$ were plotted, each entry being the average over all 1000 runs. We observe a spike for $p = n$. This peaking is similar to a related effect in ridge and lasso regression. Specifically, for fixed p and varying n , the norm of the regression estimate is growing as n is increased, reaching its peak at approximately $p = n$ while then decreasing again as the solution converges towards the true parameter as n grows very large.

4 Discussion

We derived the asymptotic distribution of the magging estimator and proposed asymptotically tight and valid confidence regions for the maximin effect. The corresponding theorems requires a rather weak assumption on the true regression coefficients b_1^0, \dots, b_G^0 . However, if this assumption is not satisfied, as studied in simulations, the resulting confidence regions seem to become too conservative. Especially when all of the “active” vectors $\{\hat{b}_g : g \in A(\hat{\Sigma}, \hat{B})\}$ are very close to each other, the proposed confidence regions tend to become large. Furthermore, in this scenario the magging estimator may suffer from a large bias. Then it may make more sense to use an approach based on relaxation. Such an approach is outlined in the Appendix in section “Relaxation-Based Approach” and it would also allow for non-asymptotic confidence intervals at the price of coverage probabilities well above the specified level. The proposed asymptotic confidence interval on the other hand is arguably more intuitive and yields in most scenarios tight bounds for large sample sizes.

Appendix

The structure is as follows: the first part is devoted to the most important definitions and explicit formulas which were omitted in the main section of the paper. The second part contains the proof of Theorem 2 and several lemmata. The third part contains the proof of Theorem 1. Finally, the last part contains a relaxation-based idea to construct confidence intervals for maximin effects.

Definitions and Formulas

Definition 1 ($A(B, \Sigma)$) The set $A(B, \Sigma) \subset \{1, \dots, G\}$ denotes indices g for which b_g has nonvanishing coefficient α_g in one of the convex combinations $M_\Sigma(B) = \sum_{g=1, \dots, G} \alpha_g b_g$ with $\alpha_g \geq 0$, $\sum_{g=1, \dots, G} \alpha_g = 1$. Note that by the assumptions of Theorem 1 or Theorem 2 the α_g are unique for $M_{\Sigma^0}(B^0)$.

Definition 2 ($W(B, \Sigma)$) $W(\hat{B}, \hat{\Sigma})$ is a consistent estimator of the variance of $\lim_n \sqrt{n}(M_{\hat{\Sigma}}(\hat{B}) - M_{\Sigma^0}(B^0))$, see proof of Theorem 2.

$$W(B, \Sigma) = \sigma^2 \sum_{g \in A(B, \Sigma)} D_g^t M_\Sigma(B_{A(B, \Sigma)}) \Sigma^{-1} D_g M_\Sigma(B_{A(B, \Sigma)}) + V(B_{A(B, \Sigma)}, \Sigma)$$

Definitions and explicit formulas of these terms can be found below. We estimate Σ^0 by $\hat{\Sigma} = \frac{1}{nG} \mathbf{X}' \mathbf{X}$. $D_g^t M_\Sigma(B)$ denotes the derivative of $M_\Sigma(B)$ with respect to b_g .

Explicit Formula for $V(\hat{B}_{A(\hat{B}, \hat{\Sigma})}, \hat{\Sigma})$

(Compare with Lemma 5)

Consistent estimator of the additional variance of $\lim_n \sqrt{n}(M_{\hat{\Sigma}}(\hat{B}) - M_{\Sigma^0}(B^0))$ “caused” by not knowing Σ^0 , see proof of Theorem 2 and Lemma 5.

$$V(\hat{B}_{A(\hat{B}, \hat{\Sigma})}, \hat{\Sigma}) = \hat{D}(\hat{D}^t \hat{\Sigma} \hat{D})^{-1} \hat{D}^t \hat{C} \hat{D}(\hat{D}^t \hat{\Sigma} \hat{D})^{-1} \hat{D}^t,$$

where \hat{C} is the empirical covariance matrix of the p -dimensional vectors $\frac{1}{\sqrt{G}} \mathbf{X}_k^t \mathbf{X}_k \cdot M_{\hat{\Sigma}}(\hat{B})$, $k = 1, \dots, (nG)$. Furthermore, with $\tilde{B} = \hat{B}_{A(\hat{B}, \hat{\Sigma})}$, $G' = |A(\hat{B}, \hat{\Sigma})|$:

$$\hat{D} := (\tilde{b}_2, \dots, \tilde{b}_{G'}) - (\tilde{b}_1, \dots, \tilde{b}_1).$$

Explicit Formula for $D_g M_{\hat{\Sigma}}(\hat{B}_{A(\hat{\Sigma}, \hat{B})})$

(Compare with Lemma 1)

Let us again write $\tilde{B} = \hat{B}_{A(\hat{B}, \hat{\Sigma})}$, $G' = |A(\hat{B}, \hat{\Sigma})|$,

$$\begin{aligned} D_g M_{\hat{\Sigma}}(\hat{B}_{A(\hat{\Sigma}, \hat{B})}) &= - \frac{\|M_{\hat{\Sigma}}(\tilde{B})\|_{\hat{\Sigma}}}{\|(\text{Id} - \widehat{\text{PA}}^{(g)}) \tilde{b}_g\|_{\hat{\Sigma}}} \frac{(\text{Id} - \widehat{\text{PA}}^{(g)}) \tilde{b}_g}{\|(\text{Id} - \widehat{\text{PA}}^{(g)}) \tilde{b}_g\|_{\hat{\Sigma}}} \frac{M_{\hat{\Sigma}}(\tilde{B})^t}{\|M_{\hat{\Sigma}}(\tilde{B})\|_{\hat{\Sigma}}} \hat{\Sigma} \\ &+ \frac{\|(\text{Id} - \widehat{\text{PA}}^{(g)}) M_{\hat{\Sigma}}(\tilde{B})\|_{\hat{\Sigma}}}{\|(\text{Id} - \widehat{\text{PA}}^{(g)}) \tilde{b}_g\|_{\hat{\Sigma}}} \hat{\Pi}_{\tilde{B}}. \end{aligned}$$

Here, $\widehat{\text{PA}}^{(g)}$ denotes the affine projection on the smallest affine space containing $\tilde{b}_1, \dots, \tilde{b}_{g-1}, \tilde{b}_{g+1}, \dots, \tilde{b}_{G'}$. Let $\Pi_{\tilde{B}} \in \mathbb{R}^{p \times p}$ denote the projection on $\langle \tilde{b}_2 - \tilde{b}_1, \dots, \tilde{b}_{G'} - \tilde{b}_1 \rangle^\perp$. These geometric definitions are meant with respect to the scalar product $\langle x, y \rangle_{\hat{\Sigma}} = x^t \hat{\Sigma} y$.

Proof of Theorem 2

Proof The proof is based on the delta method. As $\hat{B} \rightarrow B^0$ and $\hat{\Sigma} \rightarrow \Sigma^0$, by Lemma 2, $A(B^0, \Sigma^0) = A(\hat{B}, \hat{\Sigma})$ up to an asymptotically negligible set. Hence $M_{\Sigma^0}(B^0) = M_{\Sigma^0}(B_{A(B^0, \Sigma^0)}^0)$ and $M_{\hat{\Sigma}}(\hat{B}) = M_{\hat{\Sigma}}(\hat{B}_{A(B^0, \Sigma^0)})$ up to an asymptotically negligible set. So without loss of generality let us assume (without changing the definition of $\hat{\Sigma}$) that $A(B^0, \Sigma^0) = A(\hat{B}, \hat{\Sigma}) = \{1, \dots, G\}$, and hence $B^0 = B_{A(B^0, \Sigma^0)}^0$, $\hat{B} = \hat{B}_{A(\hat{B}, \hat{\Sigma})}$. By Lemmas 1 and 3, $M_{\Sigma}(B)$ is continuously differentiable in a neighborhood of B^0 and Σ^0 . Using Taylor in a neighborhood of B^0 and Σ^0 we can write

$$\begin{aligned} \sqrt{n} \left(M_{\hat{\Sigma}}(\hat{B}) - M_{\Sigma^0}(B^0) \right) &= D_B M_{\mathcal{E}}(\xi) \sqrt{n}(\hat{B} - B^0) \\ &\quad + D_{\Sigma} M_{\mathcal{E}}(\xi) \sqrt{n}(\hat{\Sigma} - \Sigma^0) + o_{\mathbb{P}}(1) \\ &= (D_B M_{\mathcal{E}}(\xi) - D_B M_{\Sigma^0}(B^0)) \sqrt{n}(\hat{B} - B^0) \\ &\quad + (D_{\Sigma} M_{\mathcal{E}}(\xi) - D_{\Sigma} M_{\Sigma^0}(B^0)) \sqrt{n}(\hat{\Sigma} - \Sigma^0) \\ &\quad + D_B M_{\Sigma^0}(B^0) \sqrt{n}(\hat{B} - B^0) \\ &\quad + D_{\Sigma} M_{\Sigma^0}(B^0) \sqrt{n}(\hat{\Sigma} - \Sigma^0) + o_{\mathbb{P}}(1), \end{aligned}$$

with $\xi = \gamma B^0 + (1 - \gamma)\hat{B}$ and $\mathcal{E} = \gamma \Sigma^0 + (1 - \gamma)\hat{\Sigma}$ for some random variable $\gamma \in [0, 1]$. We now want to show that the first and second term are negligible, and calculate the asymptotic Gaussian distributions of the last two terms. Furthermore we want to show that the last two terms are asymptotically independent. This guarantees that the variance of $\lim_n \sqrt{n} \left(M_{\hat{\Sigma}}(\hat{B}) - M_{\Sigma^0}(B^0) \right)$ is the sum of the variances of the two asymptotic Gaussian distributions.

Hence, to prove (4) it suffices to show:

- (1) $D_B M_{\mathcal{E}}(\xi) - D_B M_{\Sigma^0}(B^0) = o_{\mathbb{P}}(1)$
- (2) $D_{\Sigma} M_{\mathcal{E}}(\xi) - D_{\Sigma} M_{\Sigma^0}(B^0) = o_{\mathbb{P}}(1)$
- (3) $\sqrt{n}(\hat{b}_g - b_g^0) \rightarrow \mathcal{N}(0, \sigma^2(\Sigma^0)^{-1})$ for $g = 1, \dots, G$.
- (4) $D_B M_{\Sigma^0}(B^0) \sqrt{n}(\hat{B} - B^0) \rightarrow \mathcal{N} \left(0, \sigma^2 \sum_{g \in A(B^0, \Sigma^0)} D_g^t M_{\Sigma^0}(B^0) (\Sigma^0)^{-1} D_g M_{\Sigma^0}(B^0) \right)$

- (5) $D_{\Sigma} M_{\Sigma^0}(B^0) \sqrt{n}(\hat{\Sigma} - \Sigma^0) \rightarrow \mathcal{N}(0, V(B^0, \Sigma^0))$
 (6) For $\delta_n := \sqrt{n}(\hat{B} - B^0)$ and $\Delta_n := \sqrt{n}(\hat{\Sigma} - \Sigma^0)$ we have $(\delta_n, \Delta_n) \rightarrow (\delta, \Delta)$ with $\delta_g, g = 1, \dots, G$ and Δ independent.

Part (1) and (2): By Lemmas 1 and 3 the derivatives are continuous at B^0 and Σ^0 and $\hat{\Sigma} \rightarrow \Sigma^0, \hat{B} \rightarrow B^0$ in probability (which implies $\xi \rightarrow B^0$ and $\mathcal{E} \rightarrow \Sigma^0$ in probability).

Part (3): This is immediate, as under the chosen model, conditioned on \mathbf{X} ,

$$\hat{b}_g \sim \mathcal{N}(b_g, \sigma^2(\mathbf{X}_g^t \mathbf{X}_g)^{-1})$$

and $\frac{1}{n} \mathbf{X}_g^t \mathbf{X}_g \rightarrow \Sigma$ in probability.

Part (4): Part (3) and a linear transformation.

Part (5): We defer this part to Lemma 5.

Part (6): We saw the convergence of δ_n in part (3). The convergence of Δ_n is deferred to Lemma 4. In the following we use the notation $\delta = (\delta_1, \dots, \delta_G)$ and $\delta_n = (\delta_{n,1}, \dots, \delta_{n,G})$. For the asymptotic independence of part (6). we have to show that for any bounded continuous function g ,

$$\mathbb{E}g(\delta_n, \Delta_n) \rightarrow \int \int g(\delta, \Delta) \frac{(\det \Sigma^0)^{G/2}}{(2\pi\sigma^2)^{G/2}} \prod_{g=1}^G \exp\left(-\delta_g^t \frac{\Sigma^0}{2\sigma^2} \delta_g\right) d\delta_1 \cdots d\delta_G \mathbb{P}[d\Delta].$$

In the following equation the inner integral is bounded by 2, and for $n \rightarrow \infty$, $\frac{1}{n} \mathbf{X}_g^t \mathbf{X}_g \rightarrow \Sigma^0$ in probability. Hence, by dominated convergence on the inner and outer integral,

$$\begin{aligned} & \int \int \left| \prod_{g=1}^G \frac{\sqrt{\det \frac{1}{n} \mathbf{X}_g^t \mathbf{X}_g}}{(2\pi\sigma^2)^{1/2}} \exp\left(-\delta_{n,g}^t \frac{\mathbf{X}_g^t \mathbf{X}_g}{2n\sigma^2} \delta_{n,g}\right) \right. \\ & \left. - \prod_{g=1}^G \frac{\sqrt{\det \Sigma^0}}{(2\pi\sigma^2)^{1/2}} \exp\left(-\delta_{n,g}^t \frac{\Sigma^0}{2\sigma^2} \delta_{n,g}\right) \right| d\delta_{n,1} \cdots d\delta_{n,G} \mathbb{P}[d\Delta_n] \rightarrow 0. \end{aligned}$$

Using this,

$$\limsup_{n \rightarrow \infty} |\mathbb{E}g(\delta_n, \Delta_n) - g(\delta, \Delta_n)| = 0,$$

where δ is independent of Δ_n , $\delta_g \sim \mathcal{N}(0, \sigma^2(\Sigma^0)^{-1})$ i.i.d.. Finally, with Δ independent of δ , $\Delta \sim \lim_n \sqrt{n}(\hat{\Sigma} - \Sigma^0)$,

$$\limsup_{n \rightarrow \infty} |\mathbb{E}g(\delta_n, \Delta_n) - \mathbb{E}g(\delta, \Delta)|$$

$$= \limsup_{n \rightarrow \infty} |\mathbb{E}g(\delta, \Delta_n) - \mathbb{E}g(\delta, \Delta)|$$

$$\begin{aligned}
&= \limsup_{n \rightarrow \infty} \left| \int \mathbb{E}[(g(\delta, \Delta_n) - g(\delta, \Delta))|\delta] \frac{(\det \Sigma^0)^{G/2}}{(2\pi\sigma^2)^{G/2}} \prod_{g=1}^G \exp\left(-\delta_g^t \frac{\Sigma^0}{2\sigma^2} \delta_g\right) d\delta_1 \cdots d\delta_G \right| \\
&= 0.
\end{aligned}$$

In the second line we used Eq. (4), in the last line we used dominated convergence and $\Delta_n \rightharpoonup \Delta$. This concludes the proof.

Let $\Sigma \in \mathbb{R}^{p \times p}$ be symmetric positive definite. In the following, we work in the Hilbert space $(\mathbb{R}^p, \langle \cdot, \cdot \rangle_\Sigma)$, where for $x, y \in \mathbb{R}^p$,

$$\langle x, y \rangle_\Sigma := x^t \Sigma y,$$

and induced norm

$$\|x\|_\Sigma = \sqrt{x^t \Sigma x}.$$

This means that projections and orthogonality etc. are always meant with respect to this space. Let PA denote the affine projection on the smallest affine space containing b_1, \dots, b_G . Let $PA^{(g)}$ denote the affine projection on the smallest affine space containing $b_1, \dots, b_{g-1}, b_{g+1}, \dots, b_G$. Note that for $g = 1$ this space can be expressed as $b_2 + \langle b_3 - b_2, \dots, b_G - b_2 \rangle$. Let $\Pi_B \in \mathbb{R}^{p \times p}$ denote the projection on $\langle b_2 - b_1, \dots, b_G - b_1 \rangle^\perp$.

Lemma 1 *If $M_\Sigma(B) = \alpha_1 b_1 + \dots + \alpha_G b_G$ with $0 < \alpha_g < 1$ for $g = 1, \dots, G > 1$ and this representation is unique (i.e. $B = (b_1, \dots, b_G)$ has full rank), then M_Σ is continuously differentiable in a neighborhood of B with*

$$\begin{aligned}
D_{g,v} M_\Sigma(B) &= - \frac{\|M_\Sigma(B)\|_\Sigma}{\|(Id - PA^{(g)})b_g\|_\Sigma} \left\langle \frac{M_\Sigma(B)}{\|M_\Sigma(B)\|_\Sigma}, v \right\rangle_\Sigma \frac{(Id - PA^{(g)})b_g}{\|(Id - PA^{(g)})b_g\|_\Sigma} \\
&\quad + \frac{\|(Id - PA^{(g)})M_\Sigma(B)\|_\Sigma}{\|(Id - PA^{(g)})b_g\|_\Sigma} \Pi_B v.
\end{aligned} \tag{5}$$

Here, $D_{g,v} M_\Sigma(B)$ denotes the differential with respect to the variable b_g in direction v .

Remark 1 In the proof of Theorem 2, we could assume that without loss of generality $\{1, \dots, G\} = A(B, \Sigma)$, i.e. $B = B_{A(B, \Sigma)}$. We saw that in a neighborhood of B and Σ , magging depends only on $B_{A(B, \Sigma)}$. Hence, for using the formula of $D_g M_\Sigma(B)$ in the context of Theorems 1 and 2, replace in the definition B by $B_{A(B, \Sigma)}$. The derivatives with respect to $b_g, g \in \{1, \dots, G\} - A(B, \Sigma)$ are zero.

Proof Without loss of generality, let us assume that $g = 1$. We will show that the partial derivatives exist and are continuous.

Let $\Delta_1 \in \langle b_2 - b_1, \dots, b_G - b_1 \rangle^\perp$ and $\Delta_2 \in \langle b_2 - b_1, \dots, b_G - b_1 \rangle$ and define $\tilde{B} := (b_1 + \Delta_1 + \Delta_2, b_2, \dots, b_G)$. First, we want to show that, if $\|\Delta_1 + \Delta_2\|_\Sigma$ small,

$$M_\Sigma(\tilde{B}) = \text{PA}^{(1)}M_\Sigma(B) - \frac{\langle \text{PA}^{(1)}M_\Sigma(B), (\text{Id} - \text{PA}^{(1)})\tilde{b}_1 \rangle_\Sigma}{\|(\text{Id} - \text{PA}^{(1)})\tilde{b}_1\|_\Sigma^2} (\text{Id} - \text{PA}^{(1)})\tilde{b}_1. \quad (6)$$

Let us denote the r.h.s. by $\xi(\tilde{B})$. We have to show:

1. $\xi(\tilde{B}) \perp (\text{Id} - \text{PA}^{(1)})\tilde{b}_1$
2. $\xi(\tilde{B}) \perp \langle b_3 - b_2, \dots, b_G - b_2 \rangle$
3. $\xi(\tilde{B}) \in \text{CVX}(\tilde{B})$, the convex hull generated by the columns of \tilde{B} .

Note that 1. and 2. guarantee that the r.h.s. in (6) is perpendicular to the linear space generated by the columns of \tilde{B} .

1. is trivial. 2. By definition, $(\text{Id} - \text{PA}^{(1)})\tilde{b}_1 \perp \langle b_3 - b_2, \dots, b_G - b_2 \rangle$. $\text{PA}^{(1)}M_\Sigma(B) \perp \langle b_3 - b_2, \dots, b_G - b_2 \rangle$ as we can decompose into $\text{PA}^{(1)}M_\Sigma(B) = M_\Sigma(B) - (\text{Id} - \text{PA}^{(1)})M_\Sigma(B)$, which are both, by definition, perpendicular to $\langle b_3 - b_2, \dots, b_G - b_2 \rangle$.

Now let us show 3.: $M_\Sigma(B) = \sum_{g=1}^G \alpha_g b_g$ for some $0 < \alpha_g$ and $\sum_{g=1}^G \alpha_g = 1$, i.e. $(B^t B)^{-1} B^t M_\Sigma(B) = (B \upharpoonright_{\langle b_1, \dots, b_G \rangle})^{-1} M_\Sigma(B) = \alpha$. Similarly, as $\xi(\tilde{B})$ lies on the affine space generated by $\tilde{b}_1, \dots, \tilde{b}_G$, we have $\xi(\tilde{B}) = \sum_{g=1}^G \tilde{\alpha}_g \tilde{b}_g$ with $\sum_{g=1}^G \tilde{\alpha}_g = 1$. For small $\|\Delta_1 + \Delta_2\|_\Sigma$, \tilde{B} has full rank and as $\xi(\tilde{B}) \rightarrow M_\Sigma(B)$,

$$\lim_{\Delta \rightarrow 0} (\tilde{B} \upharpoonright_{\langle \tilde{b}_1, \dots, \tilde{b}_G \rangle})^{-1} \xi(\tilde{B}) = \lim_{\Delta \rightarrow 0} (\tilde{B}^t \tilde{B})^{-1} \tilde{B}^t \xi(\tilde{B}) = \alpha.$$

Hence, for small $\|\Delta_1 + \Delta_2\|_\Sigma$, $\tilde{\alpha}_g > 0$ and $\sum_{g=1}^G \tilde{\alpha}_g = 1$, hence $\xi(\tilde{B}) \in \text{CVX}(\tilde{B})$ and thus $M_\Sigma(\tilde{B}) = \xi(\tilde{B})$. This concludes the proof of (6).

Note that, as $\Delta_1 \perp \langle b_2 - b_1, \dots, b_G - b_1 \rangle = \langle b_1 - b_2, b_3 - b_2, \dots, b_G - b_2 \rangle$,

$$\begin{aligned} (\text{Id} - \text{PA}^{(1)})\tilde{b}_1 &= \tilde{b}_1 - \arg \min_{\gamma \in b_2 + \langle b_3 - b_2, \dots, b_G - b_2 \rangle} \|\gamma - b_1 - \Delta_1 - \Delta_2\|_\Sigma^2 \\ &= \tilde{b}_1 - \arg \min_{\gamma \in b_2 + \langle b_3 - b_2, \dots, b_G - b_2 \rangle} \|\gamma - b_1 - \Delta_2\|_\Sigma^2 + \|\Delta_1\|_\Sigma^2 \\ &= \Delta_1 + (\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2). \end{aligned} \quad (7)$$

$(\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2)$ and $(\text{Id} - \text{PA}^{(1)})b_1$ are linearly dependent. To see this, observe that both lie in the one-dimensional space $\langle b_2 - b_1, \dots, b_G - b_1 \rangle \cap \langle b_3 - b_2, \dots, b_G - b_2 \rangle^\perp$. This implies that

$$\begin{aligned} &\frac{\langle \text{PA}^{(1)}M_\Sigma(B), (\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2) \rangle_\Sigma}{\|(\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2)\|_\Sigma^2} (\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2) \\ &= \frac{\langle \text{PA}^{(1)}M_\Sigma(B), (\text{Id} - \text{PA}^{(1)})b_1 \rangle_\Sigma}{\|(\text{Id} - \text{PA}^{(1)})b_1\|_\Sigma^2} (\text{Id} - \text{PA}^{(1)})b_1 \end{aligned} \quad (8)$$

Now we can put these pieces together: In the first step we use (6) and (7), in the second we use $\Delta_1 \in \langle b_2 - b_1, \dots, b_G - b_1 \rangle^\perp$.

$$\begin{aligned} M_\Sigma(\tilde{B}) &= \text{PA}^{(1)}M_\Sigma(B) - \frac{\langle \text{PA}^{(1)}M_\Sigma(B), \Delta_1 + (\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2) \rangle_\Sigma}{\|\Delta_1 + (\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2)\|_\Sigma^2} \\ &\quad (\Delta_1 + (\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2)) \\ &= \text{PA}^{(1)}M_\Sigma(B) - \frac{\langle \text{PA}^{(1)}M_\Sigma(B), \Delta_1 + (\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2) \rangle_\Sigma}{\|\Delta_1\|^2 + \|(\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2)\|_\Sigma^2} \\ &\quad (\Delta_1 + (\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2)). \end{aligned}$$

In the first step we do an expansion of the equation above and in the second, we use (8) and $(\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2) = (\text{Id} - \text{PA}^{(1)})b_1 + \mathcal{O}(\|\Delta_2\|_\Sigma)$:

$$\begin{aligned} &M_\Sigma(\tilde{B}) \\ &= \text{PA}^{(1)}M_\Sigma(B) - \frac{\langle \text{PA}^{(1)}M_\Sigma(B), (\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2) \rangle_\Sigma}{\|(\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2)\|_\Sigma^2} (\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2) \\ &\quad - \frac{\langle \text{PA}^{(1)}M_\Sigma(B), \Delta_1 \rangle_\Sigma}{\|(\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2)\|_\Sigma^2} (\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2) \\ &\quad - \frac{\langle \text{PA}^{(1)}M_\Sigma(B), (\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2) \rangle_\Sigma}{\|(\text{Id} - \text{PA}^{(1)})(b_1 + \Delta_2)\|_\Sigma^2} \Delta_1 + \mathcal{O}(\|\Delta_1\|_\Sigma^2 + \|\Delta_2\|_\Sigma^2) \\ &= \text{PA}^{(1)}M_\Sigma(B) - \frac{\langle \text{PA}^{(1)}M_\Sigma(B), (\text{Id} - \text{PA}^{(1)})b_1 \rangle_\Sigma}{\|(\text{Id} - \text{PA}^{(1)})b_1\|_\Sigma^2} (\text{Id} - \text{PA}^{(1)})b_1 \\ &\quad - \frac{\langle \text{PA}^{(1)}M_\Sigma(B), \Delta_1 \rangle_\Sigma}{\|(\text{Id} - \text{PA}^{(1)})b_1\|_\Sigma^2} (\text{Id} - \text{PA}^{(1)})b_1 \\ &\quad - \frac{\langle \text{PA}^{(1)}M_\Sigma(B), (\text{Id} - \text{PA}^{(1)})b_1 \rangle_\Sigma}{\|(\text{Id} - \text{PA}^{(1)})b_1\|_\Sigma^2} \Delta_1 + \mathcal{O}(\|\Delta_1\|_\Sigma^2 + \|\Delta_2\|_\Sigma^2). \end{aligned}$$

From this and (6) we obtain

$$\begin{aligned} &M_\Sigma(\tilde{B}) - M_\Sigma(B) \\ &= - \frac{\langle \text{PA}^{(1)}M_\Sigma(B), \Delta_1 \rangle_\Sigma}{\|(\text{Id} - \text{PA}^{(1)})b_1\|_\Sigma^2} (\text{Id} - \text{PA}^{(1)})b_1 \\ &\quad - \frac{\langle \text{PA}^{(1)}M_\Sigma(B), (\text{Id} - \text{PA}^{(1)})b_1 \rangle_\Sigma}{\|(\text{Id} - \text{PA}^{(1)})b_1\|_\Sigma^2} \Delta_1 + \mathcal{O}(\|\Delta_1\|_\Sigma^2 + \|\Delta_2\|_\Sigma^2). \end{aligned}$$

Now let us write $\Delta_1 + \Delta_2 = \gamma v$, $\Delta_1 = \gamma(M_\Sigma(B)/\|M_\Sigma(B)\|_\Sigma + \mu v_\perp)$ with $v_\perp \perp M_\Sigma(B)$ and $v_\perp \perp \langle b_2 - b_1, \dots, b_G - b_1 \rangle$. By noting that

$$\begin{aligned} \langle \text{PA}^{(1)}M_\Sigma(B), \Delta_1 \rangle_\Sigma &= \langle M_\Sigma(B) + (\text{PA}^{(1)} - \text{Id})M_\Sigma(B), \gamma \frac{M_\Sigma(B)}{\|M_\Sigma(B)\|_\Sigma} + \mu v_\perp \rangle_\Sigma \\ &= \gamma \|M_\Sigma(B)\|_\Sigma \\ &= \gamma \left\langle \frac{M_\Sigma(B)}{\|M_\Sigma(B)\|_\Sigma}, v \right\rangle_\Sigma \|M_\Sigma(B)\|_\Sigma, \end{aligned}$$

and, as $(\text{Id} - \text{PA}^{(1)})M_\Sigma(B)$ and $(\text{Id} - \text{PA}^{(1)})b_1$ are linearly dependent (both lie in the one-dimensional space $\langle b_2 - b_1, \dots, b_G - b_1 \rangle \cap \langle b_3 - b_2, \dots, b_G - b_2 \rangle^\perp$),

$$\begin{aligned} -\langle \text{PA}^{(1)}M_\Sigma(B), (\text{Id} - \text{PA}^{(1)})b_1 \rangle_\Sigma &= \langle (\text{Id} - \text{PA}^{(1)})M_\Sigma(B), (\text{Id} - \text{PA}^{(1)})b_1 \rangle_\Sigma \\ &= \|(\text{Id} - \text{PA}^{(1)})M_\Sigma(B)\|_\Sigma \|(\text{Id} - \text{PA}^{(1)})b_1\|_\Sigma. \end{aligned}$$

We obtain:

$$\begin{aligned} &M_\Sigma(\tilde{B}) - M_\Sigma(B) \\ &= -\gamma \frac{\|M_\Sigma(B)\|_\Sigma}{\|(\text{Id} - \text{PA}^{(1)})b_1\|_\Sigma^2} \left\langle \frac{M_\Sigma(B)}{\|M_\Sigma(B)\|_\Sigma}, v \right\rangle_\Sigma (\text{Id} - \text{PA}^{(1)})b_1 \\ &\quad - \gamma \frac{\|(\text{Id} - \text{PA}^{(1)})M_\Sigma(B)\|_\Sigma \|(\text{Id} - \text{PA}^{(1)})b_1\|_\Sigma}{\|(\text{Id} - \text{PA}^{(1)})b_1\|_\Sigma^2} \Pi_B v + \mathcal{O}(\|\Delta_1\|_\Sigma^2 + \|\Delta_2\|_\Sigma^2). \end{aligned}$$

Hence the directional derivative exists and is equal to (5). The assertion follows by existence and continuity of the directional derivatives in a neighborhood of B .

Lemma 2 *Let Σ^0 be positive definite. $M_\Sigma(B)$ is continuous in B and Σ in a neighborhood of Σ^0 . Furthermore, under the assumptions of Theorem 1 (or Theorem 2), in a neighborhood of B^0 and Σ^0 , $A(B, \Sigma)$ is constant.*

Proof First, let us prove that magging is continuous. Proof by contradiction: Assume there exist sequences $B_k \rightarrow B$, $\Sigma_k \rightarrow \Sigma$ positive definite such that $M_{\Sigma_k}(B_k) \not\rightarrow M_\Sigma(B)$. Without loss of generality, as Σ is invertible, $M_{\Sigma_k}(B_k)$ converges, too. By definition of $M_{\Sigma_k}(B_k)$ we have

$$\|M_{\Sigma_k}(B_k)\|_{\Sigma_k} \leq \|\Pi_{B_k} M_\Sigma(B)\|_{\Sigma_k},$$

where Π_{B_k} denotes the projection (in $\langle \cdot, \cdot \rangle$) on the convex set $\text{CVX}(B_k)$. By continuity,

$$\|\lim_k M_{\Sigma_k}(B_k)\|_\Sigma \leq \|M_\Sigma(B)\|_\Sigma.$$

We have $M_{\Sigma_k}(B_k) \in \text{CVX}(B_k)$ and hence by continuity $\lim_k M_{\Sigma_k}(B_k) \in \text{CVX}(B)$. As magging is unique (Σ is positive definite), this yields a contradiction.

Consider b_g^0 with $g \in A(B^0, \Sigma^0)$. By the assumptions of Theorem 1, $M_{\Sigma^0}(B^0) = \sum_{i \in A(B^0, \Sigma^0)} \alpha_i b_i^0$ with $0 < \alpha_i < 1$. Hence for small $\gamma \in \mathbb{R}$, $(1 - \gamma)M_{\Sigma^0}(B^0) + \gamma b_g^0 \in \text{CVX}(B^0)$ and by definition of magging

$$\|M_{\Sigma^0}(B^0)\|_{\Sigma^0} \leq \|(1 - \gamma)M_{\Sigma^0}(B^0) + \gamma b_g^0\|_{\Sigma^0} \quad (9)$$

Using this inequality for small $\gamma > 0$ and small $\gamma < 0$ we obtain $\langle M_{\Sigma^0}(B^0), b_g^0 - M_{\Sigma^0}(B^0) \rangle = 0$. Hence, for all $g \in A(B^0, \Sigma^0)$, $M_{\Sigma^0}(B^0)$ is perpendicular (with respect to $\langle \cdot, \cdot \rangle_{\Sigma^0}$) to $b_g^0 - M_{\Sigma^0}(B^0)$. Hence $A(B^0, \Sigma^0) \subset M_{\Sigma^0}(B^0) + M_{\Sigma^0}(B^0)^\perp$.

Furthermore, by assumptions of Theorem 1, if $g \notin A(B^0, \Sigma^0)$ we have $b_g^0 \notin M_{\Sigma^0}(B^0) + M_{\Sigma^0}(B^0)^\perp$. By continuity, for $B = (b_1, \dots, b_G)$ close to B^0 and Σ close to Σ^0 (in $\|\cdot\|_2$) we have $b_g \notin M_\Sigma(B) + M_\Sigma(B)^\perp$. By an analogous argument as in Eq. (9), $g \notin A(B, \Sigma)$. This proves $A(B^0, \Sigma^0) \subset A(B, \Sigma)$.

It remains to show $A(B, \Sigma) \subset A(B^0, \Sigma^0)$: For notational simplicity let us assume $A(B^0, \Sigma^0) = \{1, \dots, G\}$. For B close to B^0 and Σ close to Σ^0 , $M_\Sigma(B) = B\tilde{\alpha}$ with $\sum_{i=1}^G \tilde{\alpha}_i = 1$, $0 \leq \tilde{\alpha}_i \leq 1$. We want to show that for B close to B^0 and Σ close to Σ^0 (in $\|\cdot\|_2$), $0 < \tilde{\alpha}_i < 1$.

To this end, note that by the assumptions of Theorem 1 we have that $B_{A(B^0, \Sigma^0)}^0$ (here without loss of generality: B^0) has full rank, hence for B close to B^0 and Σ close to Σ^0 , $(B^t B)^{-1} B^t M_\Sigma(B) = \tilde{\alpha}$ with $\tilde{\alpha}_i \geq 0$, $\sum_i \tilde{\alpha}_i = 1$. Furthermore,

$$\lim_{B \rightarrow B^0, \Sigma \rightarrow \Sigma^0} (B^t B)^{-1} B^t M_\Sigma(B) = ((B^0)^t B^0)^{-1} (B^0)^t M_{\Sigma^0}(B^0) = \alpha.$$

Hence for B close to B^0 and Σ close to Σ^0 (in $\|\cdot\|_2$), $0 < \tilde{\alpha}_i < 1$. This concludes the proof.

Lemma 3 *Let $G > 2$. Let $M_\Sigma(B) = \alpha_1 b_1 + \dots + \alpha_G b_G$ with unique $0 < \alpha_g < 1$ satisfying $\sum_{g=1}^G \alpha_g = 1$. Then the mapping*

$$\begin{aligned} \{ \text{positive definite matrices in } \mathbb{R}^{p \times p} \} &\rightarrow \mathbb{R}^p \\ \Sigma &\mapsto M_\Sigma(B) \end{aligned}$$

is continuously differentiable at B, Σ . Let Δ be a symmetric matrix. The differential in direction Δ is

$$D_\Sigma M_\Sigma(B) \Delta = -D(D^t \Sigma D)^{-1} D^t \Delta M_\Sigma(B),$$

where

$$D := (b_2, \dots, b_G) - (b_1, \dots, b_1).$$

Proof By elementary analysis, it suffices to show that the directional derivatives exist in a neighborhood and that they are continuous.

For a small symmetric perturbation $\lambda\Delta$, by continuity of magging (Lemma 2), $M_{\Sigma+\lambda\Delta}(B)$ has to satisfy

$$M_{\Sigma+\lambda\Delta}(B) = M_{\Sigma}(B) + D\gamma$$

for some (small) vector $\gamma \in \mathbb{R}^{G-1}$. By definition of magging, and as $0 < \alpha_g < 1$ we have $\|M_{\Sigma+\lambda\Delta}(B)\|_{\Sigma+\lambda\Delta} \leq \|M_{\Sigma+\lambda\Delta}(B) + D\gamma'\|_{\Sigma+\lambda\Delta}$ for all small vectors $\gamma' \in \mathbb{R}^{G-1}$. Hence,

$$M_{\Sigma+\lambda\Delta}(B)'(\Sigma + \lambda\Delta)D = 0. \quad (10)$$

Putting these two conditions together, we get

$$(M_{\Sigma}(B) + D\gamma)'(\Sigma + \lambda\Delta)D = 0.$$

Furthermore, analogously as in Eq. (10) we obtain

$$M_{\Sigma}(B)'\Sigma D = 0.$$

By combining the last two equations,

$$\gamma'D'(\Sigma + \lambda\Delta)D = -M_{\Sigma}(B)'\lambda\Delta D.$$

As $D'(\Sigma + \lambda\Delta)D$ is invertible (D has full rank as B has full rank. B has full rank as the α_g are unique),

$$\begin{aligned} \gamma' &= -M_{\Sigma}(B)'\lambda\Delta D(D'(\Sigma + \lambda\Delta)D)^{-1}, \\ D\gamma &= -D(D'(\Sigma + \lambda\Delta)D)^{-1}D'\lambda\Delta M_{\Sigma}(B). \end{aligned}$$

Dividing by λ and letting $\lambda \rightarrow 0$ gives the desired result.

Lemma 4 Let $\mathbf{X}_{k\cdot} \sim F$, $k = 1, \dots, nG$ denote the i.i.d. rows of \mathbf{X} . Let $\mathbb{E}[\|\mathbf{X}_{1\cdot}^t \mathbf{X}_{1\cdot}\|_2^2] < \infty$ and $\Sigma^0 = \mathbb{E}[\mathbf{X}_{1\cdot}^t \mathbf{X}_{1\cdot}]$ positive definite. Then, for $n \rightarrow \infty$,

$$\frac{1}{G\sqrt{n}} \sum_{k=1}^{nG} (\mathbf{X}_{k\cdot}^t \mathbf{X}_{k\cdot} - \Sigma^0) \rightarrow \Delta$$

where the symmetric matrix Δ has centered multivariate normal distributed entries under and on the diagonal with covariance

$$c_{ijkl} := \text{Covar}(\Delta_{ij}, \Delta_{kl}) = \frac{1}{G} \mathbb{E}[(\mathbf{X}_{1i} \mathbf{X}_{1j} - \mathbb{E}[\mathbf{X}_{1i} \mathbf{X}_{1j}])(\mathbf{X}_{1k} \mathbf{X}_{1l} - \mathbb{E}[\mathbf{X}_{1k} \mathbf{X}_{1l}])].$$

Proof Apply the CLT.

In the following Lemma, we want to calculate the distribution of

$$-D(D^t \Sigma D)^{-1} D^t \Delta M_\Sigma(B).$$

Lemma 5 *Let us use setting of Lemma 3 and 4.*

$$D_\Sigma M_\Sigma(B) \sqrt{n}(\hat{\Sigma} - \Sigma) \rightarrow \mathcal{N}(0, V(B, \Sigma))$$

with

$$V(B, \Sigma) = D(D^t \Sigma D)^{-1} D^t C D(D^t \Sigma D)^{-1} D^t,$$

where

$$C_{ij} = \sum_{k,l=1}^p M_\Sigma(B)_k M_\Sigma(B)_l C_{iklj},$$

is the covariance matrix of $\Delta M_\Sigma(B)$ and

$$D := (b_2, \dots, b_G) - (b_1, \dots, b_1).$$

Remark 2 In the proof of Theorem 2, we could assume that without loss of generality $\{1, \dots, G\} = A(B, \Sigma)$, i.e. $B = B_{A(B, \Sigma)}$. For using the definition of V in the context of Theorems 1 and 2, replace in the definition B by $B_{A(B, \Sigma)}$. The G in the definition of C stays the same, i.e. it is still the total number of groups.

Proof With Lemmas 3 and 4 it suffices to calculate the distribution of

$$-D(D^t \Sigma D)^{-1} D^t \Delta M_\Sigma(B),$$

i.e. the nontrivial part is to calculate the distribution of $\Delta M_\Sigma(B)$. We know it is Gaussian and centered, hence it suffices to determine the covariance matrix:

$$\begin{aligned} \mathbb{E}(\Delta M_\Sigma(B) M_\Sigma(B)^t \Delta)_{ij} &= \mathbb{E} \sum_{k,l=1}^p \Delta_{ik} (M_\Sigma(B) M_\Sigma(B)^t)_{kl} \Delta_{lj} \\ &= \sum_{k,l=1}^p M_\Sigma(B)_k M_\Sigma(B)_l \mathbb{E} \Delta_{ik} \Delta_{lj} \\ &= \sum_{k,l=1}^p M_\Sigma(B)_k M_\Sigma(B)_l C_{iklj}. \end{aligned}$$

In the last line we used Lemma 4. This concludes the proof.

Proof of Theorem 1

Proof First, note that by Lemma 1, $W(\Sigma^0, B^0)$ is invertible. Using Lemma 2, in a neighborhood of B^0 and Σ^0 the set-valued function $A(B, \Sigma)$ is constant. Hence, by Lemma 1 and Lemma 3, the derivatives of $M_\Sigma(B) = M_\Sigma(B_{A(B, \Sigma)})$ are continuous at B^0 and Σ^0 . Furthermore, $V(B_{A(B, \Sigma)}, \Sigma)$ is continuous in C and in B and Σ at B^0 and Σ^0 . All together, $W(\Sigma, B)$ is continuous at B^0 and Σ^0 in all its variables. By the definition of C in Lemma 5 and the definition of \hat{C} in section “Definitions and Formulas” in this Appendix, $\hat{C} \rightarrow C$.

Hence, $W(\hat{\Sigma}, \hat{B}) \rightarrow W(\Sigma^0, B^0)$ in probability and we obtain that $W(\hat{B}, \hat{\Sigma})^{-1} \rightarrow W(B^0, \Sigma^0)^{-1}$ in probability. By Theorem 2 and Slutsky’s Theorem we obtain

$$\sqrt{n}(M_{\hat{\Sigma}}(\hat{B}) - M_{\Sigma^0}(B^0))^t W(\hat{B}, \hat{\Sigma})^{-1} \sqrt{n}(M_{\hat{\Sigma}}(\hat{B}) - M_{\Sigma^0}(B^0)) \rightarrow \chi^2(p)$$

for $n \rightarrow \infty$. Hence

$$\begin{aligned} & \mathbb{P}[M_{\Sigma^0}(B^0) \in \mathbf{C}(\hat{\Sigma}, \hat{B})] \\ &= \mathbb{P}[(M_{\hat{\Sigma}}(\hat{B}) - M_{\Sigma^0}(B^0))^t W(\hat{B}, \hat{\Sigma})^{-1} (M_{\hat{\Sigma}}(\hat{B}) - M_{\Sigma^0}(B^0)) \leq \frac{\tau}{n}] \\ & \rightarrow 1 - \alpha \end{aligned}$$

for $n \rightarrow \infty$. This concludes the proof.

Relaxation-Based Approach

A simple approach is as follows: For given $\alpha > 0$, take random sets $\mathcal{R}_B, \mathcal{R}_\Sigma$ such that

$$\mathbb{P}[\Sigma^0 \in \mathcal{R}_\Sigma, B^0 \in \mathcal{R}_B] \geq 1 - \alpha,$$

where $B^0 = (b_1^0, \dots, b_G^0)$ is the matrix of regression coefficients in all G groups. A generic approach is to choose a confidence region for Σ^0 on the confidence level $1 - \alpha/2$ and confidence regions for b_g^0 on the confidence level $1 - \alpha/(2G)$. However, this approach can easily be improved by taking larger regions around \hat{b}_g that are far away from zero (thus have negligible influence on $M_{\hat{\Sigma}}(\hat{B})$) and smaller regions around \hat{b}_g that are close to zero. Then calculate

$$\mathcal{R} = \{M_{\hat{\Sigma}}(\hat{B}) : \hat{\Sigma} \in \mathcal{R}_\Sigma, \hat{B} \in \mathcal{R}_B\} \subset \mathbb{R}^p,$$

which is a $1 - \alpha$ confidence region for the maximin effect. However, direct computation of this confidence region is computationally cumbersome.

For known Σ^0 the idea can be relaxed to the following scheme:

For $m \in \mathbb{R}^p$ and $\Sigma \in \mathbb{R}^{p \times p}$ positive definite let us define $\|m\|_{\Sigma} := \sqrt{m^T \Sigma m}$. Note that this defines a norm on \mathbb{R}^p . Now,

$$\begin{aligned} \|M_{\Sigma^0}(B')\|_{\Sigma^0} &= \min_{\gamma \geq 0, \sum_{g=1}^G \gamma_g = 1} \|B' \gamma\|_{\Sigma^0} \\ &= \min_{\gamma \geq 0, \sum_{g=1}^G \gamma_g = 1} \|B' \gamma\|_{\Sigma^0} - \|B \gamma\|_{\Sigma^0} + \|B \gamma\|_{\Sigma^0} \\ &\leq \sup_{\gamma \geq 0, \sum_{g=1}^G \gamma_g = 1} |\|B' \gamma\|_{\Sigma^0} - \|B \gamma\|_{\Sigma^0}| + \min_{\gamma \geq 0, \sum_{g=1}^G \gamma_g = 1} \|B \gamma\|_{\Sigma^0} \\ &\leq \sup_{\gamma \geq 0, \sum_{g=1}^G \gamma_g = 1} \|(B' - B)\gamma\|_{\Sigma^0} + \min_{\gamma \geq 0, \sum_{g=1}^G \gamma_g = 1} \|B \gamma\|_{\Sigma^0} \end{aligned}$$

and hence

$$\begin{aligned} \|M_{\Sigma^0}(B')\|_{\Sigma^0} &\leq \sup_{\gamma \geq 0, \sum_{g=1}^G \gamma_g = 1} \sum_{g=1}^G \gamma_g \|b'_g - b_g\|_{\Sigma^0} + \min_{\gamma \geq 0, \sum_{g=1}^G \gamma_g = 1} \|B \gamma\|_{\Sigma^0} \\ &= \max_{g=1, \dots, G} \|b'_g - b_g\|_{\Sigma^0} + \min_{\gamma \geq 0, \sum_{g=1}^G \gamma_g = 1} \|B \gamma\|_{\Sigma^0} \\ &= \max_{g=1, \dots, G} \|b'_g - b_g\|_{\Sigma^0} + \|M_{\Sigma}(B)\|_{\Sigma^0} \end{aligned}$$

By symmetry,

$$|\|M_{\Sigma^0}(B')\|_{\Sigma^0} - \|M_{\Sigma^0}(B)\|_{\Sigma^0}| \leq \max_{g=1, \dots, G} \|b'_g - b_g\|_{\Sigma^0}. \quad (11)$$

We can now choose a covering of the confidence region \mathcal{R}_B with $B^{(k)} \in \mathcal{R}_B, k = 1, \dots, K$ such that balls $\mathcal{B}_{\epsilon_k}(B^{(k)})$ with radius ϵ_k around $B^{(k)}$ cover \mathcal{R}_B with respect to the maximum norm $\|B\|_{\max} := \max_g \|b_g\|_{\Sigma^0}$.

A confidence region of the maximin effect can then be constructed as

$$\tilde{\mathcal{R}} = \bigcup_{k=1, \dots, K} \{M : |\|M\|_{\Sigma^0} - \|M_{\Sigma^0}(B^{(k)})\|_{\Sigma^0}| \leq \epsilon_k\} \cap \text{CVX}(\mathcal{B}_{\epsilon_k}(B^{(k)})).$$

This confidence region is valid: For all $M_{\Sigma^0}(B') \in \mathcal{R}_B$ there exists $k \in \{1, \dots, K\}$ such that $\|B' - B^{(k)}\|_{\max} \leq \epsilon_k$. By Eq. (11), $|\|M_{\Sigma^0}(B')\|_{\Sigma^0} - \|M_{\Sigma^0}(B^{(k)})\|_{\Sigma^0}| \leq \epsilon_k$, hence $M_{\Sigma^0}(B') \in \tilde{\mathcal{R}}$. This implies $\mathcal{R}_B \subset \tilde{\mathcal{R}}$;

$$\mathbb{P}[M_{\Sigma^0}(B^0) \in \tilde{\mathcal{R}}] \geq \mathbb{P}[M_{\Sigma^0}(B^0) \in \mathcal{R}] \geq \mathbb{P}[B^0 \in \mathcal{R}_B] \geq 1 - \alpha.$$

If Σ^0 is unknown, using the approach above we need to estimate lower and upper bounds for $\|\cdot\|_{\Sigma^0}$.

References

1. Aitkin, M., Rubin, D.: Estimation and hypothesis testing in finite mixture models. *J. R. Stat. Soc. Ser. B* **47**, 67–75 (1985)
2. Breiman, L.: Bagging predictors. *Mach. Learn.* **24**, 123–140 (1996)
3. Bühlmann, P., Meinshausen, N.: Magging: maximin aggregation for inhomogeneous large-scale data. *IEEE Proc.* (2014, preprint). arxiv:1409.2638
4. DeSarbo, W., Cron, W.: A maximum likelihood methodology for clusterwise linear regression. *J. Classif.* **5**, 249–282 (1988)
5. Fan, J., Zhang, W.: Statistical estimation in varying coefficient models. *Ann. Stat.* **27**, 1491–1518 (1999)
6. Figueiredo, M., Jain, A.: Unsupervised learning of finite mixture models. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**, 381–396 (2002)
7. Hastie, T., Tibshirani, R.: Varying-coefficient models. *J. R. Stat. Soc. Ser. B* **55**, 757–796 (1993)
8. McLachlan, G., Peel, D.: *Finite Mixture Models*. Wiley, New York (2004)
9. Meinshausen, N., Bühlmann, P.: Maximin effects in inhomogeneous large-scale data. *Ann. Stat.* **43**(4), 1801–1830 (2015)
10. Pinheiro, J., Bates, D.: *Mixed-Effects Models in S and S-PLUS*. Springer, New York (2000)
11. Van de Geer, S., Bühlmann, P., Ritov, Y., Dezeure, R.: On asymptotically optimal confidence regions and tests for high-dimensional models. *Ann. Stat.* **42**(3), 1166–1202 (2014)
12. Zhang, C.H., Zhang, S.S.: Confidence intervals for low dimensional parameters in high dimensional linear models. *J. R. Stat. Soc. Ser. B (Stat Methodol.)* **76**(1), 217–242 (2014)